CHAPTER 4

Motherboards and Buses

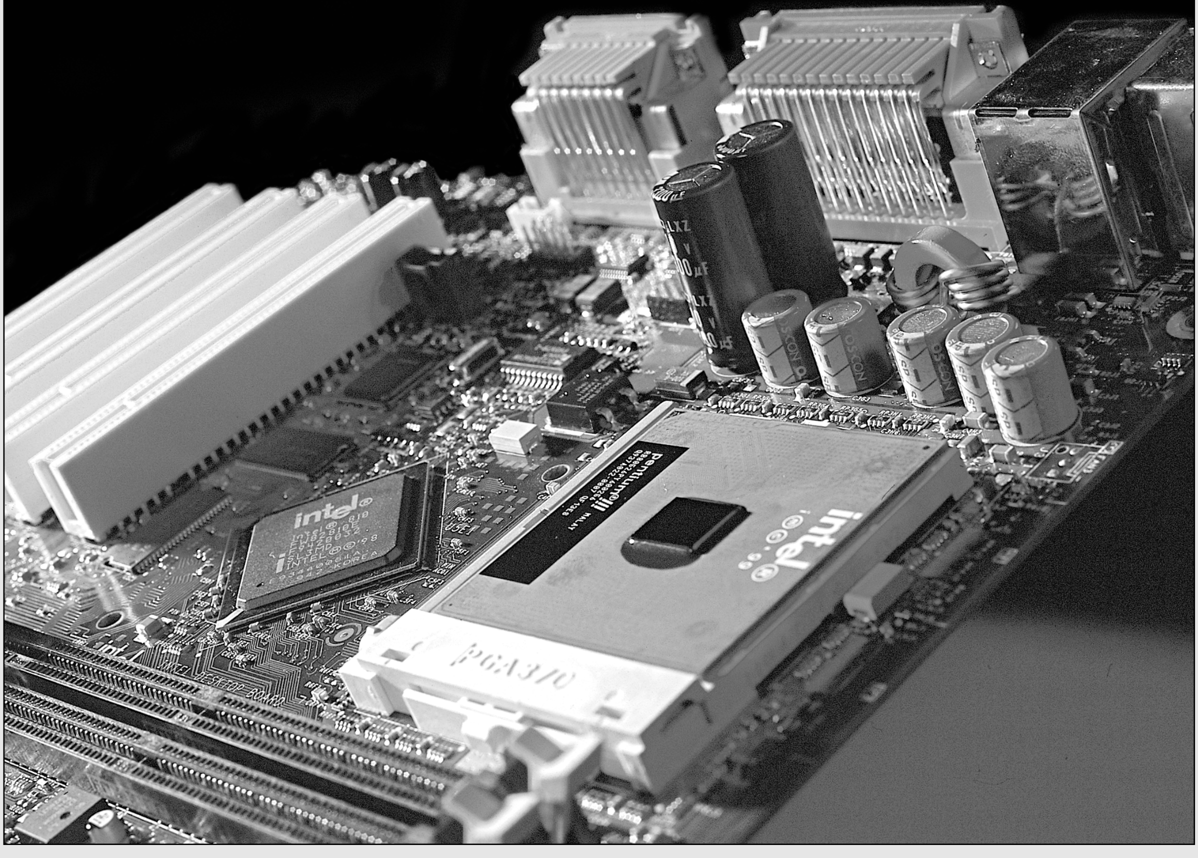

Motherboard Form Factors

Without a doubt, the most important component in a PC system is the main board or motherboard. Some companies refer to the motherboard as a system board or planar. The terms *motherboard*, *main board*, *system board*, and *planar* are interchangeable, although I prefer the *motherboard* designation. This chapter examines the various types of motherboards available and those components typically contained on the motherboard and motherboard interface connectors.

Several common form factors are used for PC motherboards. The *form factor* refers to the physical dimensions (size and shape) as well as certain connector, screw hole, and other positions that dictate into which type of case the board will fit. Some are true standards (meaning that all boards with that form factor are interchangeable), whereas others are not standardized enough to allow for interchangeability. Unfortunately, these nonstandard form factors preclude any easy upgrade or inexpensive replacement, which generally means they should be avoided. The more commonly known PC motherboard form factors include the following:

Obsolete Form Factors

- Baby-AT
- Full-size AT
- LPX (semiproprietary)
- WTX (no longer in production)
- ITX (flex-ATX variation, never produced)

Modern Form Factors

- ATX
- micro-ATX
- Flex-ATX
- Mini-ITX (flex-ATX variation)
- NLX

All Others

- Fully proprietary designs (certain Compaq, Packard Bell, Hewlett-Packard, notebook/portable systems, and so on)

Motherboards have evolved over the years from the original Baby-AT form factor boards used in the original IBM PC and XT to the current ATX and NLX boards used in most full-size desktop and tower systems. ATX has a number of variants, including micro-ATX (which is a smaller version of the ATX form factor used in the smaller systems) and flex-ATX (an even smaller version for the lowest-cost home PCs). A new form factor called mini-ITX is also available; it's really just a minimum-size version of flex-ATX designed for very small systems. NLX is designed for corporate desktop–type systems; WTX was designed for workstations and medium-duty servers, but never became popular. Table 4.1 shows the modern industry-standard form factors and their recommended uses.

Table 4.1 Current Industry-Standard Motherboard Form Factors

Form Factor	Use
ATX	Standard desktop, mini-tower, and full-tower systems; most common form factor today; most flexible design for power users, enthusiasts, low-end servers/workstations, and higher-end home systems; ATX boards support up to seven expansion slots.
Mini-ATX	A slightly smaller version of ATX that fits into the same case as ATX. Many so-called ATX motherboards are actually mini-ATX motherboards; mini-ATX boards support up to six expansion slots.
micro-ATX	A smaller version of ATX, used in Mid-range desktop or mini-tower systems. Fits micro-ATX or ATX chassis.
Flex-ATX	Smallest version of ATX, used in expensive or low-end small desktop or mini-tower systems; entertainment or appliance systems. Fits in flex-ATX, micro-ATX, or ATX chassis.
Mini-ITX	Minimum-size flex-ATX version, used in set-top boxes and compact/small form factor computers; highly integrated with one PCI expansion slot. Fits in mini-ITX, flex-ATX, micro-ATX, or ATX chassis.
NLX	Corporate desktop or mini-tower systems; fast and easy serviceability.

Although the Baby-AT, Full-size AT, and LPX boards were once popular, they have all but been replaced by more modern and interchangeable form factors. The modern form factors are true standards, which guarantees improved interchangeability within each type. This means that ATX boards can interchange with other ATX boards, NLX with other NLX, and so on. The additional features found on these boards as compared to the obsolete form factors, combined with true interchangeability, has made the migration to these newer form factors quick and easy. Today I recommend purchasing only systems with one of the modern industry-standard form factors. Each of these form factors, however, is discussed in more detail in the following sections.

Anything that does not fit into one of the industry-standard form factors is considered proprietary. Unless there are special circumstances, I do not recommend purchasing systems with proprietary board designs. They will be virtually impossible to upgrade and very expensive to repair later because the motherboard, case, and often power supply will not be interchangeable with other models. I call proprietary form factor systems "disposable" PCs because that's what you must normally do with them when they are too slow or need repair out of warranty.

Caution

"Disposable" PCs might be more common than ever. Some estimate that as much as 60% of all PCs sold today are disposable models, not so much because of the motherboards used, but because of the tiny power supplies and cramped micro-tower cases that are favored on most retail-market PCs today. Although low-cost PCs using small chassis and power supplies are theoretically more upgradeable than past disposable type systems, you'll still hit the wall over time if you need more than three expansion slots or want to use more than two or three internal drives. Because mini-tower systems are so cramped and limited, I consider them to be almost as disposable as the LPX systems they have largely replaced.

You also need to watch out for systems that only appear to meet industry standards, such as certain Dell computer models built from 1996 to the present—especially the XPS line of systems. These computers often use rewired versions of the ATX power supply (or even some that are completely nonstandard in size and shape) and modified motherboard power connectors, which makes both components completely incompatible with standard motherboards and power supplies. In some of the systems, the power supply has a completely proprietary shape as well and the motherboards are not fully standard ATX either. If you want to upgrade the power supply, you must use a special Dell-compatible power supply. And if you want to upgrade the motherboard (assuming you can find one that fits), you must buy a standard power supply to match. The best alternative is to replace the motherboard, power supply, and possibly the case with industry-standard components simultaneously. For more details about how to determine whether your Dell computer uses nonstandard power connectors, see Chapter 21, "Power Supply and Chassis/Case."

If you want to have a truly upgradeable system, insist on systems that use ATX motherboards in a mid-tower or larger case with at least five drive bays.

PC and XT

The first popular PC motherboard was, of course, the original IBM PC released in August 1981. Figure 4.1 shows how this board looked. IBM followed the PC with the XT motherboard in March 1983, which had the same size and shape as the PC board but had eight slots instead of five. Both the IBM PC and XT motherboards were 9"×13" in size. Also, the slots were spaced 0.8" apart in the XT instead of 1" apart as in the PC (see Figure 4.2). The XT also eliminated the little used cassette port in the back, which was supposed to be used to save BASIC programs on cassette tape instead of the much more expensive (at the time) floppy drive.

Note

The Technical Reference section of the DVD accompanying this book contains detailed information on the PC (5150) and XT (5160). All the information there is printable.

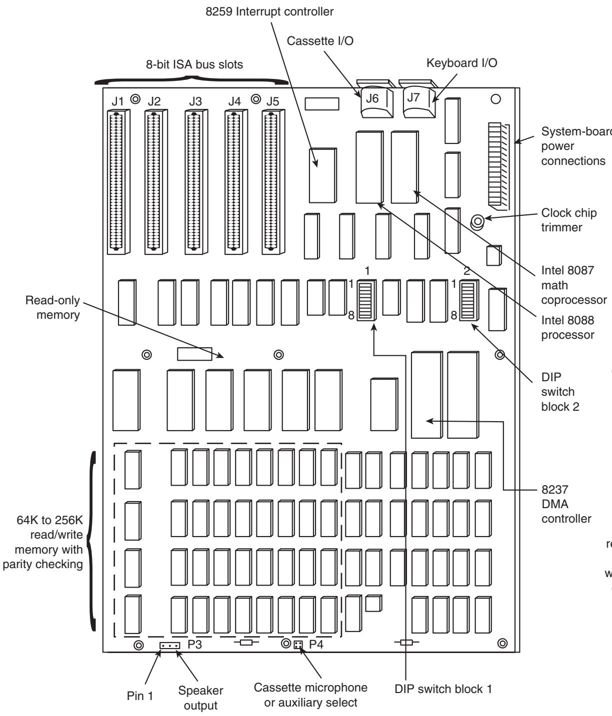

Figure 4.1 IBM PC motherboard (circa 1981).

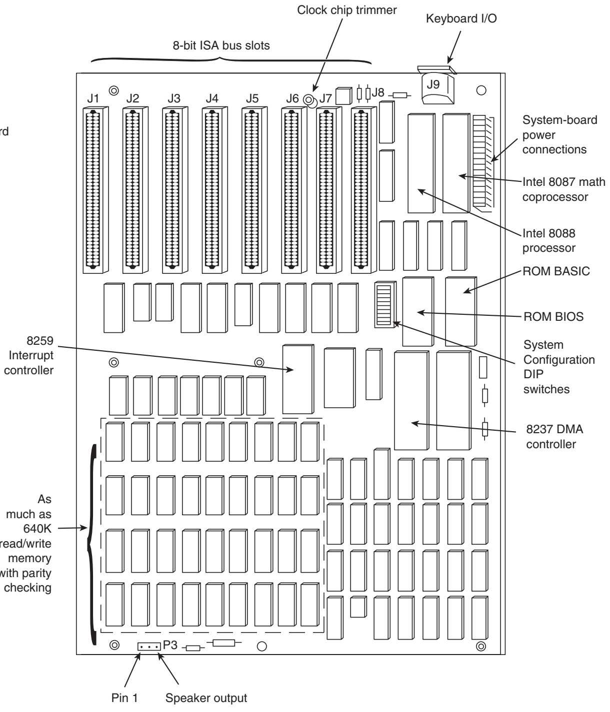

Figure 4.2 IBM PC-XT motherboard (circa 1983).

The minor differences in the slot positions and the deleted cassette connector on the back required a minor redesign of the case. In essence, the XT was a mildly enhanced PC, with a motherboard that was the same overall size and shape, used the same processor, and came in a case that was identical except for slot bracketry and the lack of a hole for the cassette port. Eventually, the XT motherboard design became very popular, and many other PC motherboard manufacturers of the day copied IBM's XT design and produced similar boards.

Full-Size AT

The full-size AT motherboard form factor matches the original IBM AT motherboard design. This allows for a very large board of up to 12" wide by 13.8" deep. The full-size AT board first debuted in August 1984, when IBM introduced the Personal Computer AT (advanced technology). To accommodate the 16-bit 286 processor and all the necessary support components at the time, IBM needed more room than the original PC/XT-sized boards could provide. So for the AT, IBM increased the size of the motherboard but retained the same screw hole and connector positions of the XT design. To accomplish this, IBM essentially started with a PC/XT-sized board and extended it in two directions (see Figure 4.3).

Note

The Technical Reference section of the DVD enclosed with this book contains detailed coverage of the AT and the XT Model 286.

A little more than a year after being introduced, the appearance of chipsets and other circuit consolidation allowed the same motherboard functionality to be built using fewer chips, so the board was redesigned to make it slightly smaller. Then, it was redesigned again as IBM shrank the board down to XT-size in a system it called the XT-286 (introduced in September 1986). The XT-286 board was virtually identical in size and shape to the original XT, a form factor which would later be known as Baby-AT.

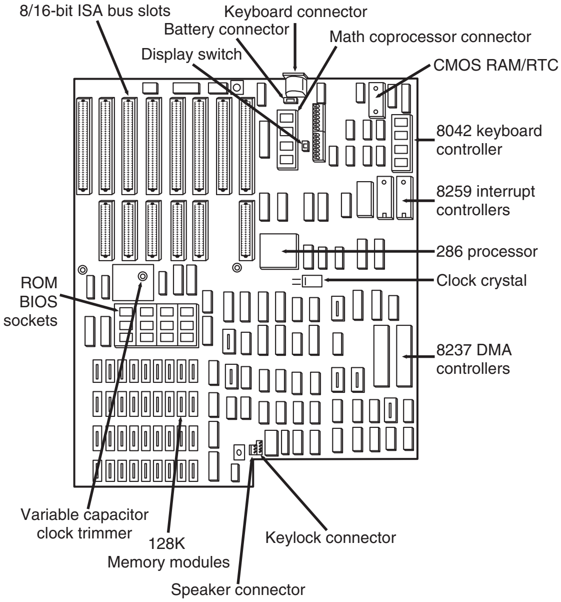

Figure 4.3 IBM AT motherboard (circa 1984).

The keyboard connector and slot connectors in the full-size AT boards still conformed to the same specific placement requirements to fit the holes in the XT cases already in use, but a larger case was still required to fit the larger board. Because of the larger size of the board, a full-size AT motherboard only fits into full-size AT desktop or tower cases. Because these motherboards do not fit into the smaller Baby-AT or mini-tower cases, and because of advances in component miniaturization, they are no longer being produced by most motherboard manufacturers—except in some cases for dual processor server applications.

The important thing to note about the full-size AT systems is that you can always replace a full-size AT motherboard with a Baby-AT (or XT-size) board, but the opposite is not true unless the case is large enough to accommodate the full-size AT design.

Baby-AT

After IBM released the AT in August 1984, component consolidation allowed subsequent systems to be designed using far fewer chips and requiring much less in the way of motherboard real estate. Therefore, all the additional circuits on the 16-bit AT motherboard could fit into boards using the smaller XT form factor.

IBM was one of the first to use the smaller boards when it introduced a system called the XT-286 in September 1986. Unfortunately the “XT” designation in the name of that system caused a lot of confusion, and many people did not want to buy a system they thought used older and slower technology. Sales of the XT-286 were dismal. By this time, other companies had also developed XT-size AT class systems. However, they decided that rather than calling these boards XT-size, which seemed to make people think they were 8-bit designs, they would refer to them as “Baby-AT” designs. The intention was to make people understand that these new boards had AT technology in a smaller form factor and were not souped-up versions of older technology as was seemingly implied by IBM’s XT-286 moniker.

Thus, the Baby-AT form factor is essentially the same form factor as the original IBM XT motherboard. The only difference is a slight modification in one of the screw hole positions to fit into an AT-style case. These motherboards also have specific placement of the keyboard and slot connectors to match the holes in the case. Note that virtually all full-size AT and Baby-AT motherboards use the standard 5-

pin DIN type connector for the keyboard. Baby-AT motherboards can be used to replace full-size AT motherboards and will fit into several case designs. Because of its flexibility, from 1983 into early 1996, the Baby-AT form factor was the most popular motherboard type. Starting in mid-1996, Baby-AT was replaced by the superior ATX motherboard design, which is not directly interchangeable. Most systems sold since 1996 have used the improved ATX, micro-ATX, or NLX design, and Baby-AT is getting harder and harder to come by. Figure 4.3 shows the onboard features and layout of a late-model Baby-AT motherboard. Older Baby-AT motherboards have the same general layout but lack advanced features, such as USB connectors, DIMM memory sockets, and the AGP slot.

Any case that accepts a full-size AT motherboard will also accept a Baby-AT design. PC motherboards using the Baby-AT design have been manufactured to use virtually any processor from the original 8088 to the Pentium III or Athlon, although the pickings are slim where the newer processors are concerned. As such, systems with Baby-AT motherboards were the original upgradeable systems. Because any Baby-AT motherboard can be replaced with any other Baby-AT motherboard, this is an interchangeable design. Even though the Baby-AT design (shown in Figure 4.4) is now obsolete, ATX carries on its philosophy of interchangeability. Figure 4.5 shows a more modern Baby-AT motherboard, which includes USB compatibility, SIMM and DIMM sockets, and even a supplemental ATX power supply connection.

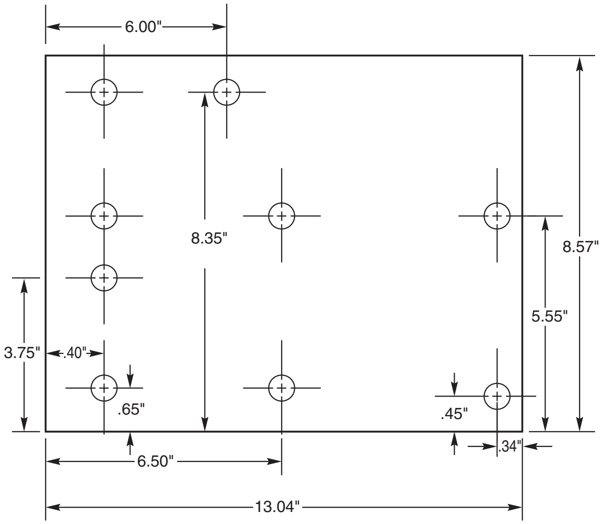

Figure 4.4 Baby-AT motherboard form factor dimensions.

The easiest way to identify a Baby-AT form factor system without opening it is to look at the rear of the case. In a Baby-AT motherboard, the cards plug directly into the board at a 90° angle; in other words, the slots in the case for the cards are perpendicular to the motherboard. Also, the Baby-AT motherboard has only one visible connector directly attached to the board, which is the keyboard connector. Typically, this connector is the full-size 5-pin DIN type connector, although some Baby-AT systems use the smaller 6-pin mini-DIN connector (sometimes called a PS/2 type connector) and might even have a mouse connector. All other connectors are mounted on the case or on card edge brackets and are attached to the motherboard via cables. The keyboard connector is visible through an appropriately placed hole in the case.

▶▶ See "Keyboard/Mouse Interface Connectors," p. 995.

Baby-AT boards all conform to specific widths and screw hole, slot, and keyboard connector locations, but one thing that can vary is the length of the board. Versions have been built that are smaller than the full 9"×13" size; these are often called mini-AT, micro-AT, or even things such as 2/3-Baby or

1/2-Baby. Even though they might not be the full size, they still bolt directly into the same case as a standard Baby-AT board and can be used as a direct replacement for one.

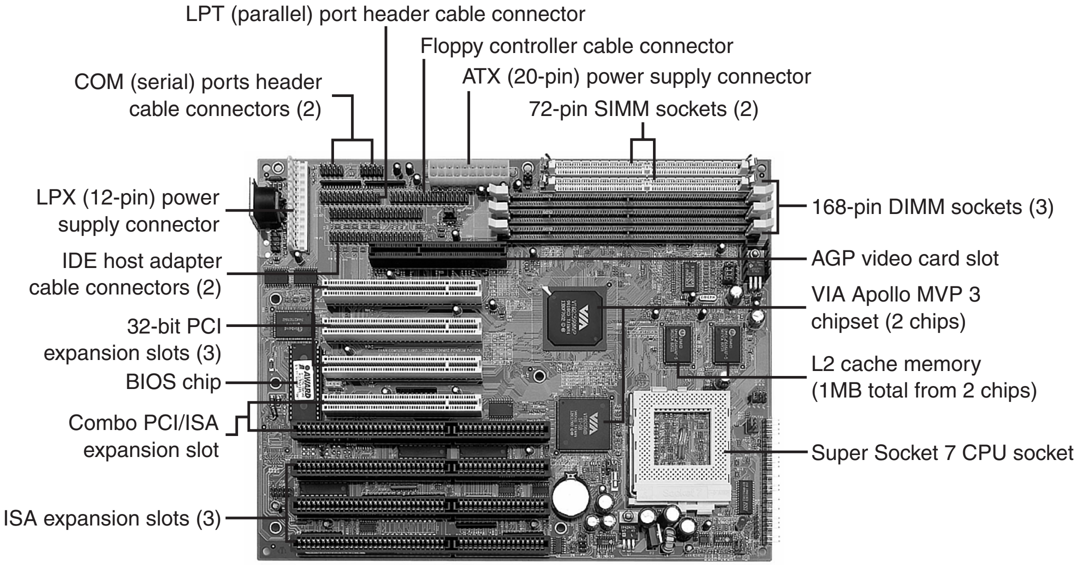

Figure 4.5 A late-model Baby-AT motherboard, the Tyan Trinity 100AT (S1590). *Photo courtesy of Tyan Computer Corporation.*

LPX

The LPX and mini-LPX form factor boards were a semiproprietary design that Western Digital originally developed in 1987 for some of its motherboards. The *LP* in LPX stands for Low Profile, which is so named because these boards incorporate slots that are parallel to the main board, enabling the expansion cards to install sideways. This allows for a slim or low-profile case design and overall a smaller system than the Baby-AT.

Although Western Digital no longer produces PC motherboards, the form factor lives on, and many other motherboard manufacturers have duplicated the design. Unfortunately, because the specifications were never laid out in exact detail—especially with regard to the bus riser card portion of the design—these boards are termed semiproprietary and are not interchangeable between manufacturers. Some vendors, such as IBM and HP, for example, have built LPX systems that use a T-shaped riser card that allows expansion cards to be mounted at the normal 90° angle to the motherboard but still above the motherboard. This lack of standardization means that if you have a system with an LPX board, in most cases you can’t replace the motherboard with a different LPX board later. You essentially have a system you can’t upgrade or repair by replacing the motherboard with something better. In other words, you have what I call a disposable PC, something I would not normally recommend that anybody purchase.

Most people were not aware of the semiproprietary nature of the design of these boards, and they were extremely popular in what I call “retail store” PCs from the late 1980s through the late 1990s. This would include primarily Compaq and Packard Bell systems, as well as many others who used this form factor in their lower-cost systems. These boards were most often used in low-profile or Slimline case systems but were found in tower cases, too. These were often lower-cost systems such as those sold at retail electronics superstores. Although scarce even in retail chains today, because of their proprietary nature, I recommend staying away from any system that uses an LPX motherboard.

Purchasing LPX Motherboards

Normally, I would never recommend upgrading an LPX system—they simply aren't worth the expense. However, a few vendors do sell LPX motherboards, so if it's absolutely necessary, an upgrade might be possible. The problem is the riser card, which is typically sold separately from the motherboard itself. It is up to you to figure out which riser card will work in your existing case, and often, if you choose the wrong riser card, you're stuck with it.

If you must locate LPX-compatible products, try these Web sites and vendors:

- Unicorn Computers (Taiwan): `www.unicorn-computer.com.tw`
- Hong Faith America (Taiwan): `america.hongfaith.com`
- FriendTech's LPX Zone: `www.friendtech.com/LPX_Zone/LPX_Zone.htm`

Note that the FriendTech Web site has links to other vendors and is developing a comprehensive knowledge base of LPX-related information and possible upgrades.

LPX boards are characterized by several distinctive features (see Figure 4.6). The most noticeable is that the expansion slots are mounted on a bus riser card that plugs into the motherboard. In most designs, expansion cards plug sideways into the riser card. This sideways placement allows for the low-profile case design. Slots are located on one or both sides of the riser card depending on the system and case design. Vendors who use LPX-type motherboards in tower cases sometimes use a T-shaped riser card instead, which puts the expansion slots at the normal right angle to the motherboard but on a raised shelf above the motherboard itself.

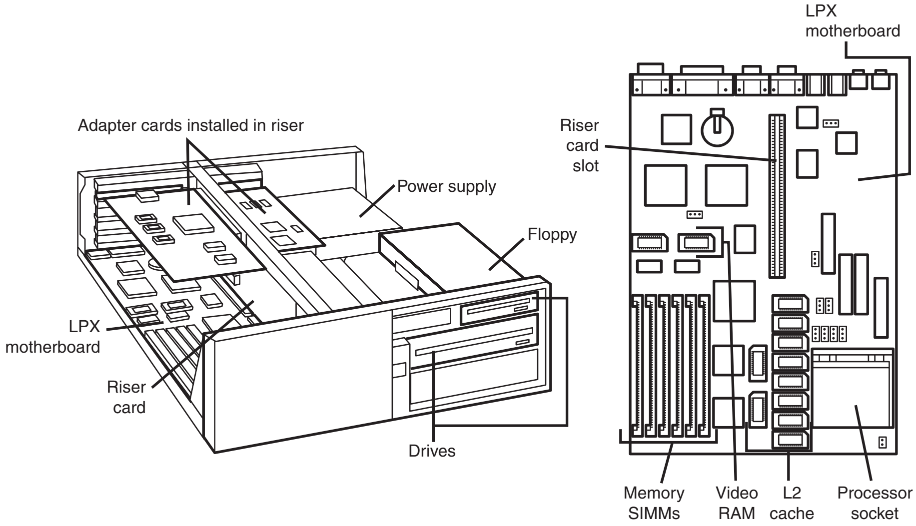

Figure 4.6 Typical LPX system chassis and motherboard.

Another distinguishing feature of the LPX design is the standard placement of connectors on the back of the board. An LPX board has a row of connectors for video (VGA 15-pin), parallel (25-pin), two serial ports (9-pin each), and mini-DIN PS/2 style mouse and keyboard connectors. All these connectors are mounted across the rear of the motherboard and protrude through a slot in the case. Some LPX motherboards might have additional connectors for other internal ports, such as network or SCSI adapters. Because LPX systems use a high degree of motherboard port integration, many vendors of LPX motherboards, cases, and systems often refer to LPX products as having an "all-in-one" design.

The standard form factor used for LPX and mini-LPX motherboards in many typical low-cost systems is shown in Figure 4.7.

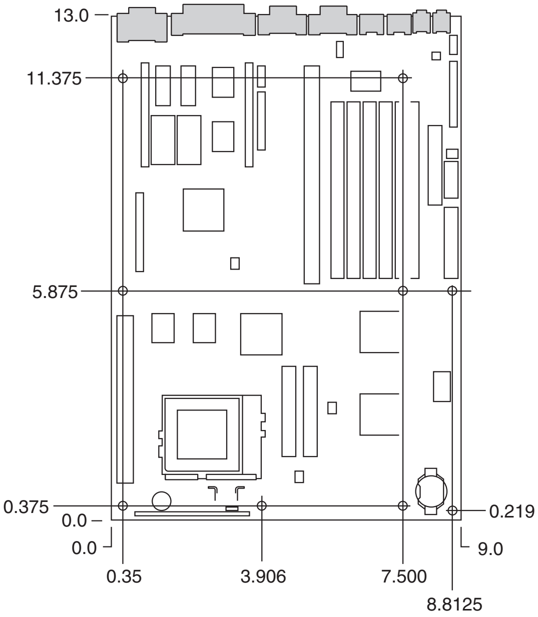

Figure 4.7 LPX motherboard dimensions.

I am often asked, “How can I tell whether a system has an LPX board without opening the cover?” Because of the many variations in riser card design, and because newer motherboards such as NLX also use riser cards, the most reliable way to distinguish an LPX motherboard from other systems is to look at the connector signature (the layout and pattern of connectors on the back of the board). As you can see in Figure 4.8, all LPX motherboards—regardless of variations in riser card shape, size, or location—place all external ports along the rear of the motherboard. By contrast, Baby-AT motherboards use case-mounted or expansion slot–mounted connectors for serial, parallel, PS/2 mouse, and USB ports, whereas ATX-family motherboards group all external ports together to the left side of the expansion slots.

On an LPX board, the riser is placed in the middle of the motherboard, whereas NLX boards have the riser to the side (the motherboard actually plugs into the riser in NLX).

Figure 4.8 shows two typical examples of the connectors on the back of LPX boards. Note that not all LPX boards have the built-in audio, so those connectors might be missing. Other ports (such as USB) might be missing from what is shown in these diagrams, depending on exactly which options are included on a specific board; however, the general layout will be the same.

The connectors along the rear of the board would interfere with locating bus slots directly on the motherboard, which accounts for why riser cards are used for adding expansion boards.

Although the built-in connectors on the LPX boards were a good idea, unfortunately the LPX design was semiproprietary (not a fully interchangeable standard) and thus, not a good choice. Newer motherboard form factors such as ATX, micro-ATX, and NLX have both built-in connectors and use a standard board design. The riser card design of LPX allowed system designers to create a low-profile desktop system, a feature now carried by the much more standardized NLX form factor. In fact, NLX was developed as the modern replacement for LPX.

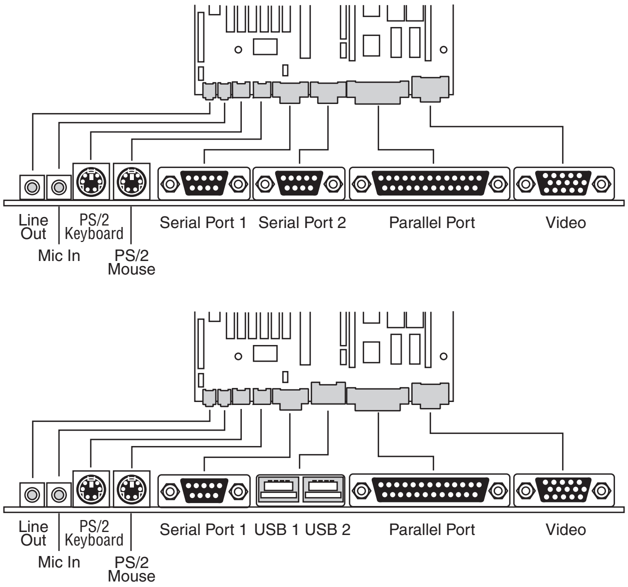

Figure 4.8 LPX motherboard back panel connectors.

ATX

The ATX form factor was the first of a dramatic evolution in motherboard form factors. ATX is a combination of the best features of the Baby-AT and LPX motherboard designs, with many new enhancements and features thrown in. The ATX form factor is essentially a Baby-AT motherboard turned sideways in the chassis, along with a modified power supply location and connector. The most important thing to know initially about the ATX form factor is that it is physically incompatible with either the previous Baby-AT or LPX design. In other words, a different case and power supply are required to match the ATX motherboard. These case and power supply designs have become common and are found in most new systems.

Intel initially released the official ATX specification in July 1995. It was written as an open specification for the industry. ATX boards didn't hit the market in force until mid-1996, when they rapidly began replacing Baby-AT boards in new systems. The ATX specification was updated to version 2.01 in February 1997, and has had several minor revisions since. The latest revision is ATX version 2.03 (with Engineering Change Revision P1), released in May 2000. Intel has published detailed specifications so other manufacturers can use the ATX design in their systems. The current specifications for ATX and other current motherboard types are available online from the Desktop Form Factors site: `www.formfactors.org`. Currently, ATX is the most popular motherboard form factor for new systems, and it is the one I recommend most people get in their systems today. An ATX system will be upgradeable for many years to come, exactly like Baby-AT was in the past.

ATX improved on the Baby-AT and LPX motherboard designs in several major areas:

- *Built-in double high external I/O connector panel.* The rear portion of the motherboard includes a stacked I/O connector area that is 6 1/4" wide by 1 3/4" tall. This enables external connectors to be located directly on the board and negates the need for cables running from internal connectors to the back of the case as with Baby-AT designs.
- *Single main keyed internal power supply connector.* This is a boon for the average end user who always had to worry about interchanging the Baby-AT power supply connectors and subsequently blowing the motherboard. The ATX specification includes a keyed and shrouded main power connector that

is easy to plug in and can't be installed incorrectly. This connector also features pins for supplying 3.3V to the motherboard, so ATX motherboards do not require built-in voltage regulators that are susceptible to failure. The ATX specification was extended to include two additional optional keyed power connectors called the Auxiliary Power connector (3.3V and 5V) and the ATX12V connector for systems that require more power than the original specification would allow.

▶▶ See "Motherboard Power Connectors," p. 1138.

- *Relocated CPU and memory.* The CPU and memory modules are relocated so they can't interfere with any bus expansion cards and can easily be accessed for upgrade without removing any of the installed bus adapters. The CPU and memory are relocated next to the power supply, which is where the primary system fan is located. The improved airflow concentrated over the processor, in the case of some older processors, eliminates the need for extra-cost CPU cooling fans (lower power configurations only). There is room for a CPU and a heatsink and fan combination of up to 2.8" in height, as well as more than adequate side clearance provided in that area.

Note

Most systems require cooling in addition to the fan in the power supply—from a secondary case-mounted fan or an active heatsink on the processor with an integral fan. Intel and AMD supply processors with attached high-quality (ball bearing) fans for CPUs sold to smaller vendors. These are so-called "boxed" processors because they are sold in single-unit box quantities instead of cases of 100 or more like the raw CPUs sold to the larger vendors. The included fan heatsink is an excellent form of thermal insurance because most smaller vendors and system self-assemblers lack the engineering knowledge necessary to perform thermal analysis, temperature measurements, and the testing required to select the properly sized passive heatsinks. The only thermal requirement spelled out for the boxed processors is that the temperature of the air entering the active heatsink (usually the same as the system interior ambient temperature) is kept to 45°C (113°F) or less in Pentium III or earlier models, or 40°C (104°F) or less in the Pentium 4 or later processors. By putting a high-quality fan on these "boxed" processors, Intel and AMD can put a warranty on the boxed processors that is independent of the system warranty. Larger vendors have the engineering talent to select the proper passive heatsink, thus reducing the cost of the system as well as increasing reliability. With an OEM non-boxed processor, the warranty is with the system vendor and not the processor manufacturer directly. Heatsink mounting instructions usually are included with a motherboard if non-boxed processors are used.

- *Relocated internal I/O connectors.* The internal I/O connectors for the floppy and hard disk drives are relocated to be near the drive bays and out from under the expansion board slot and drive bay areas. Therefore, internal cables to the drives can be much shorter, and accessing the connectors does not require card or drive removal.
- *Improved cooling.* The CPU and main memory are designed and positioned to improve overall system cooling. This can decrease—but not necessarily eliminate—the need for separate case or CPU cooling fans. Most higher-speed systems still need additional cooling fans for the CPU and chassis. Note that the ATX specification originally specified that the ATX power supply fan blows into the system chassis instead of outward. This reverse flow, or positive pressure design, pressurizes the case and minimizes dust and dirt intrusion. More recently, the ATX specification was revised to allow the more normal standard flow, which negatively pressurizes the case by having the fan blow outward. Because the specification technically allows either type of airflow, and because some overall cooling efficiency is lost with the reverse flow design, most power supply manufacturers provide ATX power supplies with fans that exhaust air from the system, otherwise called a negative pressure design. See Chapter 21 for more detailed information.
- *Lower cost to manufacture.* The ATX specification eliminates the need for the rat's nest of cables to external port connectors found on Baby-AT motherboards, additional CPU or chassis cooling fans, or onboard 3.3V voltage regulators. Instead, ATX allows for shorter internal drive cables and no cables for standard external serial or parallel ports. These all conspire to greatly reduce the cost of the motherboard and the cost of a complete system—including the case and power supply.

Figure 4.9 shows the ATX system layout and chassis features, as you would see them looking in with the lid off on a desktop, or sideways in a tower with the side panel removed. Notice how virtually the entire motherboard is clear of the drive bays and how the devices such as CPU, memory, and internal drive connectors are easy to access and do not interfere with the bus slots. Also notice how the processor is positioned near the power supply.

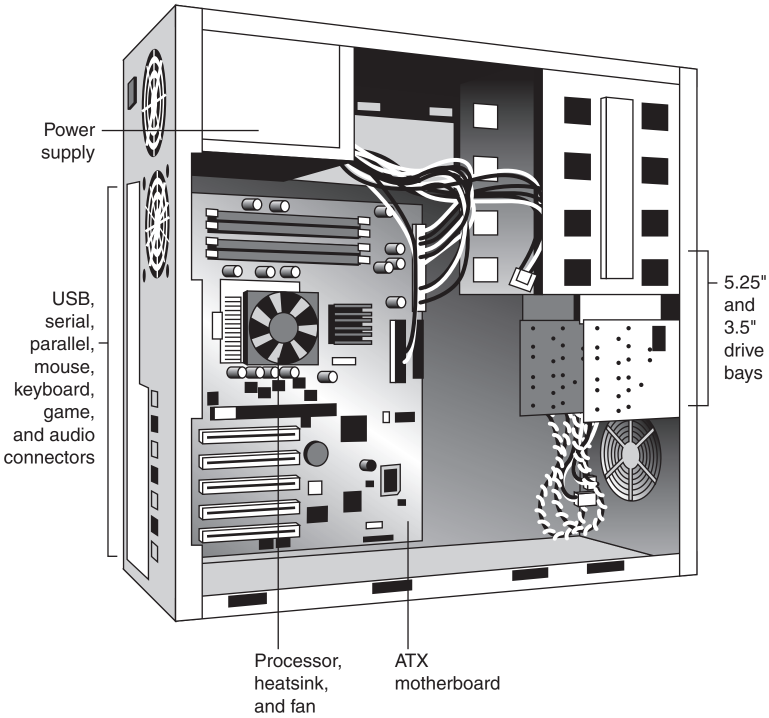

Figure 4.9 When mounted inside the case, the ATX motherboard is oriented so that the CPU socket is near the power supply fan and case fan (if your case includes one).

The ATX motherboard shape is basically a Baby-AT design rotated sideways 90°. The expansion slots are now parallel to the shorter side dimension and do not interfere with the CPU, memory, or I/O connector sockets (see Figure 4.10). There are actually two basic sizes of standard ATX boards. In addition to a full-size ATX layout, Intel also specified a mini-ATX design, which is a fully compatible subset of ATX that fits into the same case:

- A full-size ATX board is 12" wide × 9.6" deep (305mm×244mm).
- The mini-ATX board is 11.2" × 8.2" (284mm×208mm).

Mini-ATX is not an official standard; instead it is simply referenced as a slightly smaller version of ATX. In fact, all references to mini-ATX were removed from the ATX 2.1 and later specifications. Two smaller official versions of ATX exist, called micro-ATX and flex-ATX. They are discussed in the following sections.

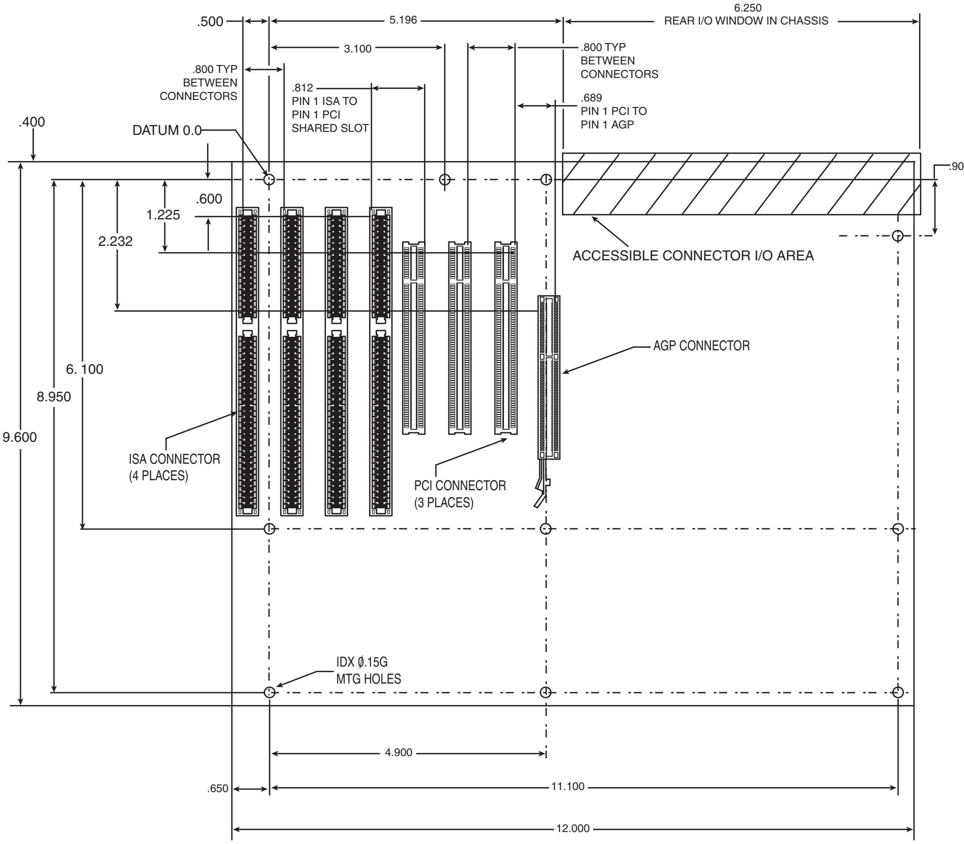

Figure 4.10 ATX specification 2.1 motherboard dimensions.

Although the case holes are similar to the Baby-AT case, cases for Baby-AT and ATX are generally incompatible. The ATX power supply design is identical in physical size to the standard Slimline power supply used with Baby-AT systems; however, they also use different connectors and supply different voltages.

The ATX form factor's design advantages have swept Baby-AT and LPX motherboards off the market. Although other form factors are now available, I have been recommending only ATX (or compatible variations such as micro-ATX or flex-ATX) systems for new system purchases since late 1996 and will probably continue to do so for the next several years.

The best way to tell whether your system has an ATX-family motherboard design without removing the lid is to look at the back of the system. Two distinguishing features identify ATX. One is that the expansion boards plug directly into the motherboard. There is usually no riser card as with LPX or NLX, so the slots are perpendicular to the plane of the motherboard. Also, ATX boards have a unique double-high connector area for all the built-in connectors on the motherboard (see Figure 4.11 and Table 4.2). This is found just to the side of the bus slot area and can be used to easily identify an ATX board.

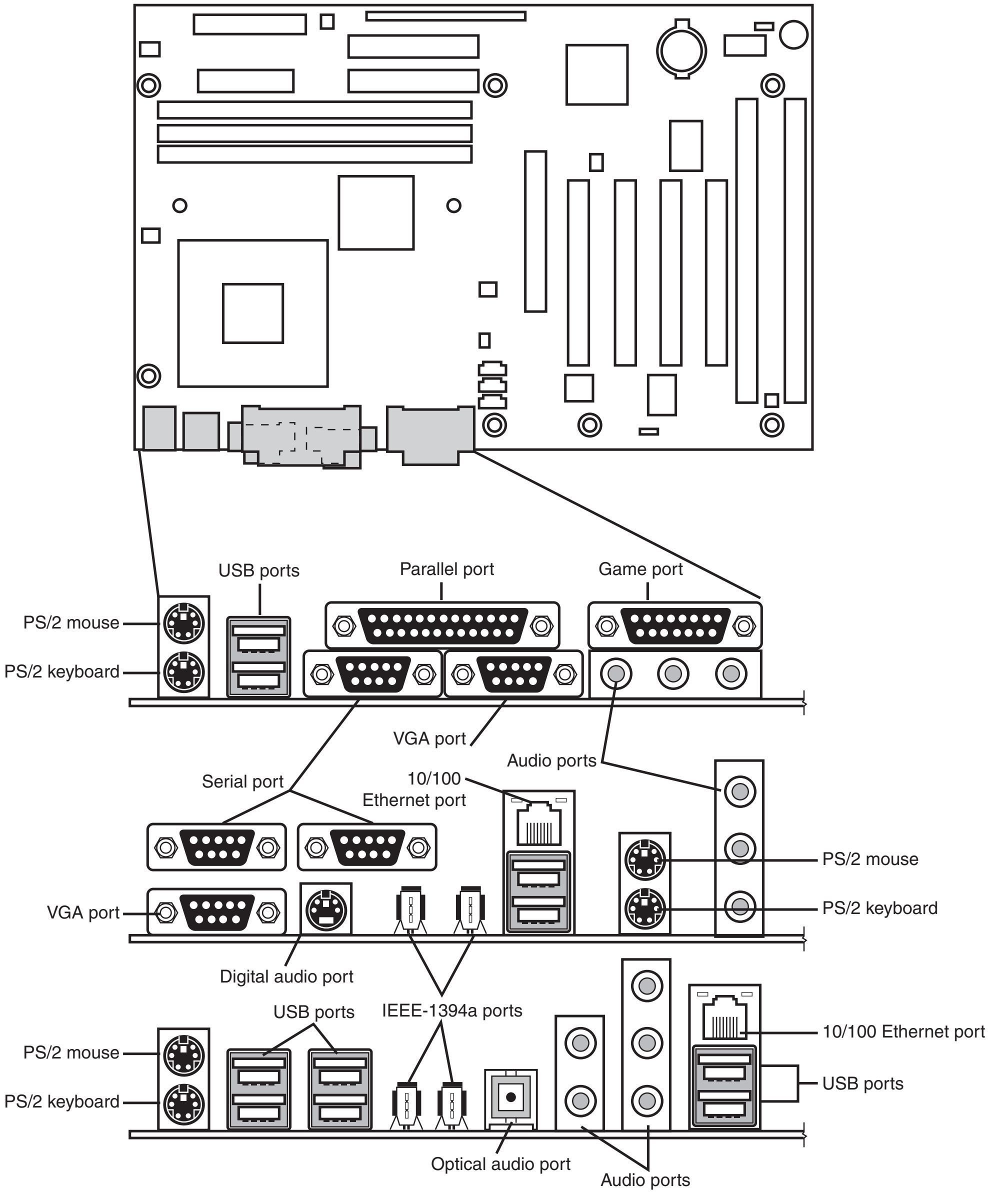

Figure 4.11 ATX motherboard and rear panel connections from systems with onboard sound and video (top and middle), networking and IEEE-1394/FireWire (middle and bottom), and a "legacy-free" system (bottom).

Table 4.2 Built-in Ports Usually Found on ATX Motherboards

Port Description	Connector Type	Connector Color
PS/2 mouse port	6-pin Mini-DIN	Green
PS/2 keyboard port	6-pin Mini-DIN	Purple
USB ports	Dual Stack USB	Black
Parallel port	25-pin D-Submini	Burgundy
Serial port	9-pin D-Submini	Teal
VGA analog video port	15-pin HD D-Submini	Dark blue
MIDI/Game port	15-pin D-Submini	Gold
Audio ports: L/R in, front L/R out, rear L/R out, center/LFE out, Microphone L/R in	1/8'' (3.5mm) Mini-Phone	Light blue, lime green, black, black, pink
S-Video TV out	4-pin Mini-DIN	Black
IEEE-1394/FireWire port	6-pin IEEE-1394	Gray
10/100/1000 Ethernet LAN	8-pin RJ-45	Black
Optical S/PDIF audio out	TOSLINK	Black
DVI digital video out (not shown)	DDWG-DVI	White
Digital S/PDIF audio out (not shown)	RCA Jack	Orange
SCSI (not shown)	50/68-pin HD SCSI	Black
Modem (not shown)	4-pin RJ-11	Black
Composite Video out (not shown)	RCA Jack	Yellow

DIN = Deutsches Institut für Normung e.V.
USB = Universal serial bus
VGA = Video graphics array
HD = High density
MIDI = Musical Instrument Digital Interface
L/R = Left and right channel
LFE = Low frequency effects (subwoofer)
S-Video = Super Video
IEEE = Institute of Electrical and Electronics Engineers
LAN = Local area network
RJ = Registered jack
S/PDIF = Sony/Philips Digital Interface
TOSLINK = Toshiba optical link
DVI = Digital visual interface
DDWG = Digital Display Working Group
RCA = Radio Corporation of America
SCSI = Small computer system interface

Note

Most ATX motherboards feature connectors with industry-standardized color codes (shown in the previous table). This makes plugging in devices much easier and more foolproof: You merely match up the colors. For example, most keyboards have a cable with a purple plug, whereas most mice have a cable with a green plug. Even though the keyboard and mouse connectors on the motherboard appear the same (both are 6-pin Mini-DIN types), their color-coding matches the plugs on the respective devices. Thus, to plug them in properly, you merely insert the purple plug into the purple connector and the green plug into the green connector. This saves you from having to bend down to try to decipher small labels on the connectors to ensure you get them right.

The specification and related information about the ATX, mini-ATX, micro-ATX, flex-ATX, or NLX form factor specifications are available from the Form Factors Web site at `www.formfactors.org`. The Form Factors site provides form factor specifications and design guides, as well as design considerations for new technologies, information on initiative supporters, vendor products, and a form factor discussion forum.

Note

Some motherboards, especially those used in server systems, come in nonstandard ATX variations collectively called *extended ATX*. This is a term applied to boards that are compatible with ATX but that are deeper. Standard ATX is 12"×9.6" (305mm×244mm), whereas extended ATX boards are up to 12"×13" (305mm×330mm). Because technically no official "extended ATX" standard exists, compatibility problems can exist with boards and chassis claiming to support extended ATX. When purchasing an extended ATX board, be sure it will fit in the chassis you intend to use. Dual Xeon processors fit in a standard ATX-size board, so choose a standard ATX-size board for maximum compatibility with the existing ATX chassis.

ATX Riser

In December 1999, Intel introduced a riser card design modification for ATX motherboards. The design includes the addition of a 22-pin (2×11) connector to one of the PCI slots on the motherboard, along with a two- or three-slot riser card that plugs in. The riser enables two or three PCI cards to be installed, but it does not support AGP.

ATX motherboards typically are found in vertically oriented tower-type cases, but often a horizontal desktop system is desired for a particular application. When ATX boards are installed in desktop cases, PCI cards can be as tall as 4.2", thus requiring a case that is at least 6"–7" tall. For Slimline desktop systems, most manufacturers now use the NLX format, but the more complex design and lower popularity of NLX makes that a more expensive alternative. A low-cost way to use an industry-standard ATX form factor board in a Slimline desktop case is therefore needed. The best long-term solution to this problem is the eventual adoption of a lower-profile PCI card design that is shorter than the current 4.2". The PCI Low-Profile specification was released for engineering review by the Peripheral Component Interconnect Special Interest Group (PCI SIG) on February 14, 2000, and some PCI card products have been produced in this shorter (2.5") form factor. Until Low-Profile PCI becomes widespread, Intel has suggested a riser card approach to enable standard-height PCI cards to be used in Slimline and rack-mount systems.

By adding a small 22-pin extension connector to one of the PCI slots on a motherboard, the necessary additional signals for riser card support could be implemented. The current design enables the use of a two- or three-slot riser that is either 2" or 2.8" tall, respectively. To this riser, you can attach full-length cards sideways in the system, and the motherboard can be used with or without the riser. The only caveat is that, if a riser card is installed, the remaining PCI slots on the motherboard can't be used. You can have expansion cards plugged in to only the riser or the motherboard, but not both. Also, the riser card supports only PCI cards—not AGP or ISA cards. A sample ATX board with riser installed is shown in Figure 4.12.

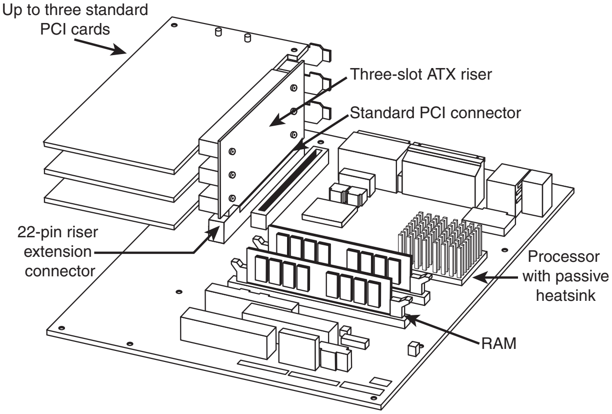

Figure 4.12 A three-slot ATX riser implementation on a micro-ATX motherboard.

The 22-pin extension connector usually is installed in line with PCI slot 6, which is the second one from the right; the slots are usually numbered from right to left (facing the board) starting with 7 as the one closest to the processor. Some boards number the slots from right to left starting with 1; in that case, the extension connector is on PCI slot 2. The pinout of the ATX 22-pin riser extension connector is shown in Figure 4.13.

Signal	Pin	Pin	Signal
Ground	B1	A1	PCI_GNT1#
PCI_ CLK1	B2	A2	Ground
Ground	B3	A3	PCI_GNT2#
PCI_REQ1#	A4	B4	Ground
Ground	A5	B5	PCI_CLK3
PCI_CLK2	A6	B6	RISER_ID1
Ground	A7	B7	Reserved
PCI_REQ2#	A8	B8	RISER_ID2
Ground	A9	B9	NOGO
PC/PCI_DREQ#	A10	B10	+12V
PC/PCI_DGNT#	A11	B11	SER_IRQ

Figure 4.13 An ATX 22-pin riser extension connector pinout.

The PCI connector that is in line with the riser extension connector is just a standard PCI slot; none of the signals are changed.

Systems that use the riser generally are low-profile designs. Therefore, they don't fit normal PCI or AGP cards in the remaining (nonriser-bound) slots. Although the ATX riser standard originally was developed for use with low-end boards—which have integrated video, sound, and network support—many rack-mounted servers are also using the ATX riser because these boards also have most of their required components already integrated. In fact, the ATX riser appears to be more popular for rack-mounted servers than for the originally intended target market of Slimline desktop systems.

ATX riser cards, compatible cases, and compatible motherboards are available from a variety of vendors, allowing you to build your own Slimline ATX system.

Micro-ATX

Micro-ATX is a motherboard form factor Intel originally introduced in December 1997, as an evolution of the ATX form factor for smaller and lower-cost systems. The reduced size as compared to standard ATX allows for a smaller chassis, motherboard, and power supply, thereby reducing the cost of the entire system. The micro-ATX form factor is also backward-compatible with the ATX form factor and can be used in full-size ATX cases. Of course, a micro-ATX case doesn't take a full-size ATX board. This form factor has become popular in the low-cost, sub-$1,000 PC market. Currently, mini-tower chassis systems dominate the low-cost PC market, although their small sizes and cramped interiors severely limit future upgradeability.

The main differences between micro-ATX and standard or mini-ATX are as follows:

- Reduced width motherboard (9.6" [244mm] instead of 12" [305mm] or 11.2" [284mm])
- Fewer I/O bus expansion slots (four maximum, although most boards feature only three)
- Smaller power supply optional (SFX/TFX form factors)

The micro-ATX motherboard maximum size is only 9.6"×9.6" (244mm×244mm) as compared to the full-size ATX size of 12"×9.6" (305mm×244mm) or the mini-ATX size of 11.2"×8.2" (284mm×208mm). Even smaller boards can be designed as long as they conform to the location of the mounting holes, connector positions, and so on, as defined by the standard. Fewer slots aren't a problem for typical home or small-business PC users because more components such as sound and video are usually integrated on the motherboard and therefore don't require separate slots. This higher integration reduces motherboard and system costs. External buses, such as USB, 10/100 Ethernet, and optionally SCSI or 1394 (FireWire), can provide additional expansion out of the box. The specifications for micro-ATX motherboard dimensions are shown in Figure 4.14.

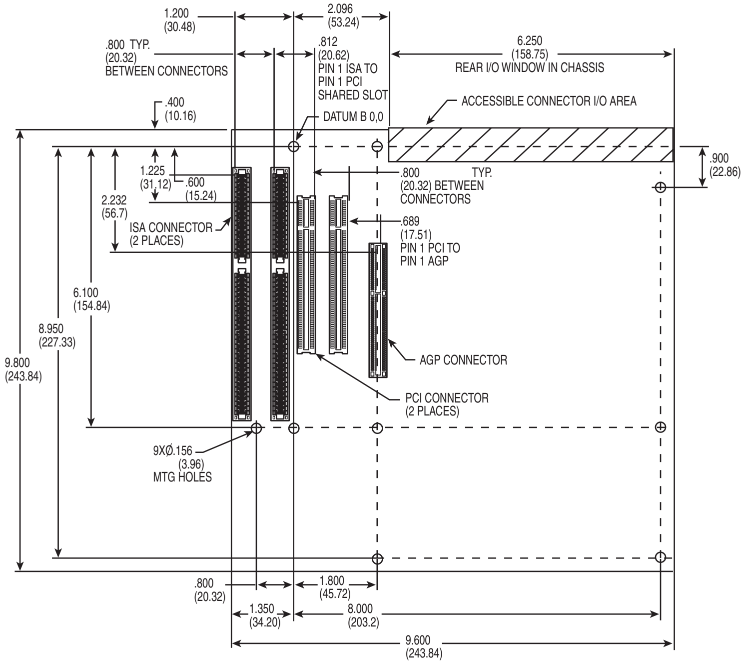

Figure 4.14 Micro-ATX specification 1.1 motherboard dimensions.

Smaller form factor (called SFX or TFX) power supplies have been defined for optional use with micro-ATX systems, although the standard ATX supply also works fine because the connectors are the same. The smaller size SFX/TFX power supplies encourage flexibility in choosing mounting locations within the chassis and allows for smaller systems that consume less power overall. Although the smaller supplies can be used, they may lack sufficient power output for faster or more fully configured systems. Because of the high power demands of most modern systems, most third-party micro-ATX chassis are designed to accept standard ATX power supplies, although micro-ATX systems sold by vendors such as Compaq, HP, and eMachines typically use some type of SFX or TFX power supply to reduce costs.

▶▶ See "Power Supply Form Factors," p. 1129.

The micro-ATX form factor is similar to ATX for compatibility. The similarities include the following:

- Standard ATX 20-pin power connector
- Standard ATX I/O panel
- Mounting holes and dimensions are a subset of ATX

These similarities ensure that a micro-ATX motherboard can easily work in a standard ATX chassis with a standard ATX power supply, as well as the smaller micro-ATX chassis and SFX/TFX power supply.

The overall system size for a micro-ATX is very small. A typical case is only 12"–14" tall, about 7" wide, and 12" deep. This results in a kind of micro-tower or desktop size. A typical micro-ATX motherboard is shown in Figure 4.15.

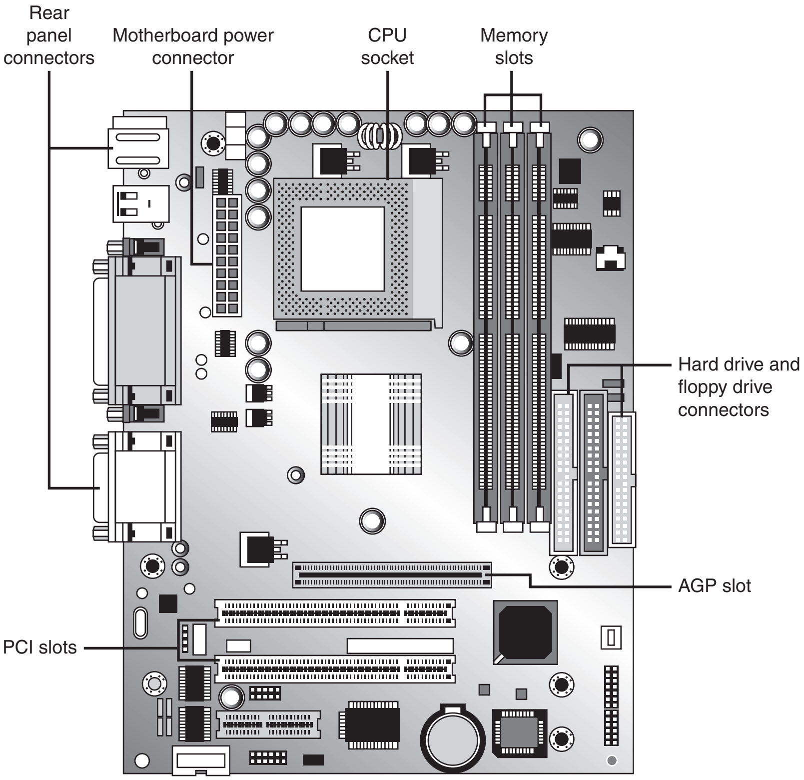

Figure 4.15 A typical micro-ATX motherboard's dimensions are 9.6"×9.6".

As with ATX, Intel released micro-ATX to the public domain to facilitate adoption as a de facto standard. The specification and related information on micro-ATX are available through the Desktop Form Factors site (`www.formfactors.org`).

Flex-ATX

In March 1999, Intel released the flex-ATX addendum to the micro-ATX specification. This added a new and even smaller variation of the ATX form factor to the motherboard scene. Flex-ATX's smaller design is intended to allow a variety of new PC designs, especially extremely inexpensive, smaller, consumer-oriented, appliance-type systems. Some of these designs might not even have expansion slots, allowing expansion only through USB or IEEE-1394/FireWire ports.

Flex-ATX defines a board that is up to 9"×7.5" (229mm×191mm) in size, which is the smallest of the ATX family boards. In all other ways, flex-ATX is the same as ATX and micro-ATX, making flex-ATX fully backward compatible with ATX or micro-ATX by using a subset of the mounting holes and the same I/O and power supply connector specifications (see Figure 4.16).

Most flex-ATX systems likely use SFX/TFX (small or thin form factor) type power supplies (introduced in the micro-ATX specification), although if the chassis allows it, a standard ATX power supply can also be used.

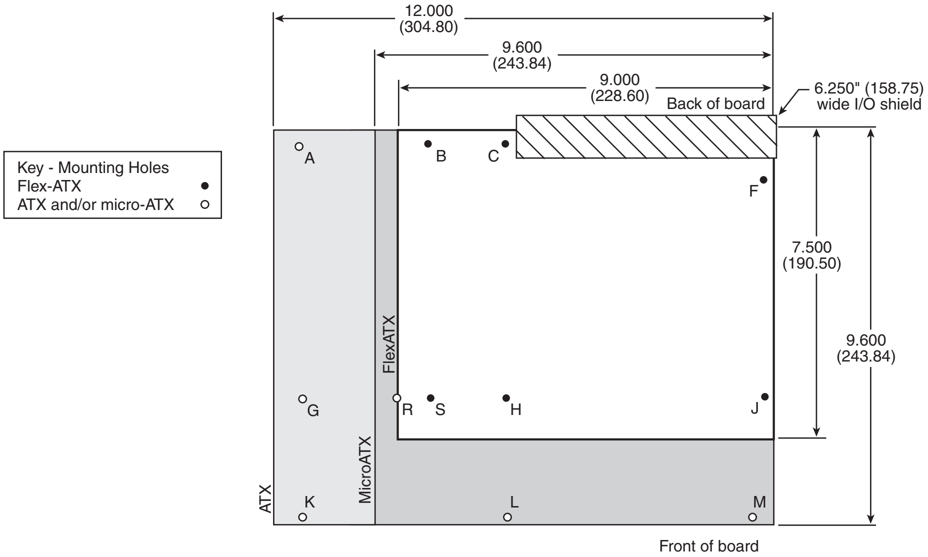

Form Factor	Mounting Hole Locations	Notes
• Flex-ATX	B, C, F, H, J, S	
• Micro-ATX	B, C, F, H, J, L, M, R, S	Holes R and S were added for micro-ATX form factor. Hole B was defined in full-AT format.
• ATX	A, C, F, G, H, J, K, L, M	Hole F must be implemented in all ATX 2.03–compliant chassis assemblies. The hole was optional in the ATX 1.1 specification.

Figure 4.16 Size and mounting hole comparison between ATX, micro-ATX, and flex-ATX motherboards.

With the addition of flex-ATX, the family of ATX boards has now grown to include four definitions of size (three are the official standards), as shown in Table 4.3.

Table 4.3 ATX Motherboard Form Factors

Form Factor	Max. Width	Max. Depth	Max. Area	Size Comparison
ATX	12.0'' (305mm)	9.6'' (244mm)	115 sq. in. (743 sq. cm)	—
Mini-ATX	11.2'' (284mm)	8.2'' (208mm)	92 sq. in. (593 sq. cm)	20% smaller
Micro-ATX	9.6'' (244mm)	9.6'' (244mm)	92 sq. in. (595 sq. cm)	20% smaller
Flex-ATX	9.0'' (229mm)	7.5'' (191mm)	68 sq. in. (435 sq. cm)	41% smaller

Note that these dimensions are the maximums allowed. Making a board smaller in any given dimension is always possible as long as it conforms to the mounting hole and connector placement requirements detailed in the respective specifications. Each board has the same basic screw hole and connector placement requirements, so if you have a case that fits a full-size ATX board, you could also mount a micro-ATX, or flex-ATX board in that same case. Obviously, if you have a smaller case designed for micro-ATX or flex-ATX, you won't be able to put the larger mini-ATX or full-size ATX boards in that case.

ITX and Mini-ITX

Flex-ATX is the smallest industry-standard form factor specification, and it defines a board that is *up to* 9"×7.5" in size. Note the *up to* part of the dimensions, which means that, even though those dimensions are the maximums, less is also allowed. Therefore, a flex-ATX board can be smaller than that, but how much smaller? By analyzing the flex-ATX specification—and, in particular, studying the required mounting screw locations—you can see that a flex-ATX board could be made small enough to use only four mounting holes (C, F, H, and J). Refer to Figure 4.16 for the respective hole locations.

According to the flex-ATX standard, the distance between holes H and J is 6.2", and the distance between hole J and the right edge of the board is 0.25". By leaving the same margin from hole H to the left edge, you could make a board with a minimum width of 6.7" (0.25" + 6.2" + 0.25") that would conform to the flex-ATX specification. Similarly, the distance between holes C and H is 6.1", and the distance between hole C and the back edge of the board is 0.4". By leaving a minimum 0.2" margin from hole H to the front edge, you could make a board with a minimum depth of 6.7" (0.4" + 6.1" + 0.2") that would conform to the flex-ATX specification. By combining the minimum width and depth, you can see that the minimum board size that would conform to the flex-ATX specification is 6.7"×6.7" (170mm×170mm).

VIA Technologies Platform Solutions Division wanted to create a motherboard as small as possible, yet not define a completely new and incompatible form factor. To accomplish this, in March 2001 VIA created a board that was slightly narrower in width (8.5" instead of 9") but still the same depth as flex-ATX, resulting in a board that was 6% smaller and yet still conformed to the flex-ATX specification. VIA called this ITX but then realized that the size savings were simply too small to justify developing it further, so it was discontinued before any products were released.

In April 2002, VIA created an even smaller board that featured the absolute minimum width and depth dimensions allowed by flex-ATX. It called it mini-ITX. In essence, all mini-ATX boards are simply flex-ATX boards that are limited to the minimum allowable dimensions. All other aspects, including the I/O aperture size and location, screw hole locations, and power supply connections, are pure flex-ATX. A mini-ITX board fits in any chassis that accepts a flex-ATX board; however, larger boards will not fit into a mini-ITX chassis.

The mini-ITX form factor was designed by VIA especially to support VIA's low-power embedded Eden and C3 E-Series processors. Only a very small number of motherboards is available in this form factor, and only from VIA and one or two other manufacturers. Because the processors used on these boards are substantially less powerful than even the Intel Celeron 4 or AMD Duron entry-level processors, the mini-ITX form factor is intended for use mainly in nontraditional settings such as set-top boxes and computing appliances. The size of the ITX and mini-ITX boards relate to flex-ATX as shown in Table 4.4.

Table 4.4 Comparing the Flex-ATX, ITX, and Mini-ITX Form Factors

Form Factor	Max. Width	Max. Depth	Max. Area	Size Comparison
Flex-ATX	9.0" (229mm)	7.5" (191mm)	68 sq. in. (435 sq. cm)	—
ITX	8.5" (215mm)	7.5" (191mm)	64 sq. in. (411 sq. cm)	6% smaller
Mini-ITX	6.7" (170mm)	6.7" (170mm)	45 sq. in. (290 sq. cm)	34% smaller

Again, I must point out that technically any ITX or mini-ITX board conforms to the flex-ATX specification. In particular, the mini-ITX is the smallest board that can conform. Although the still-born ITX format was virtually the same as flex-ATX in size (which is probably why it was discontinued before any were sold), mini-ITX motherboards are 170mm×170mm (6.7"×6.7"), which is 34% smaller than the maximum allowed by flex-ATX.

To take advantage of the smaller mini-ITX format, several chassis makers are producing very small chassis to fit these boards. Most are the shape of a small cube, with one floppy and one optical drive bay visible from the front. The layout of a typical mini-ITX motherboard, the VIA EPIA-V, is shown in Figure 4.17.

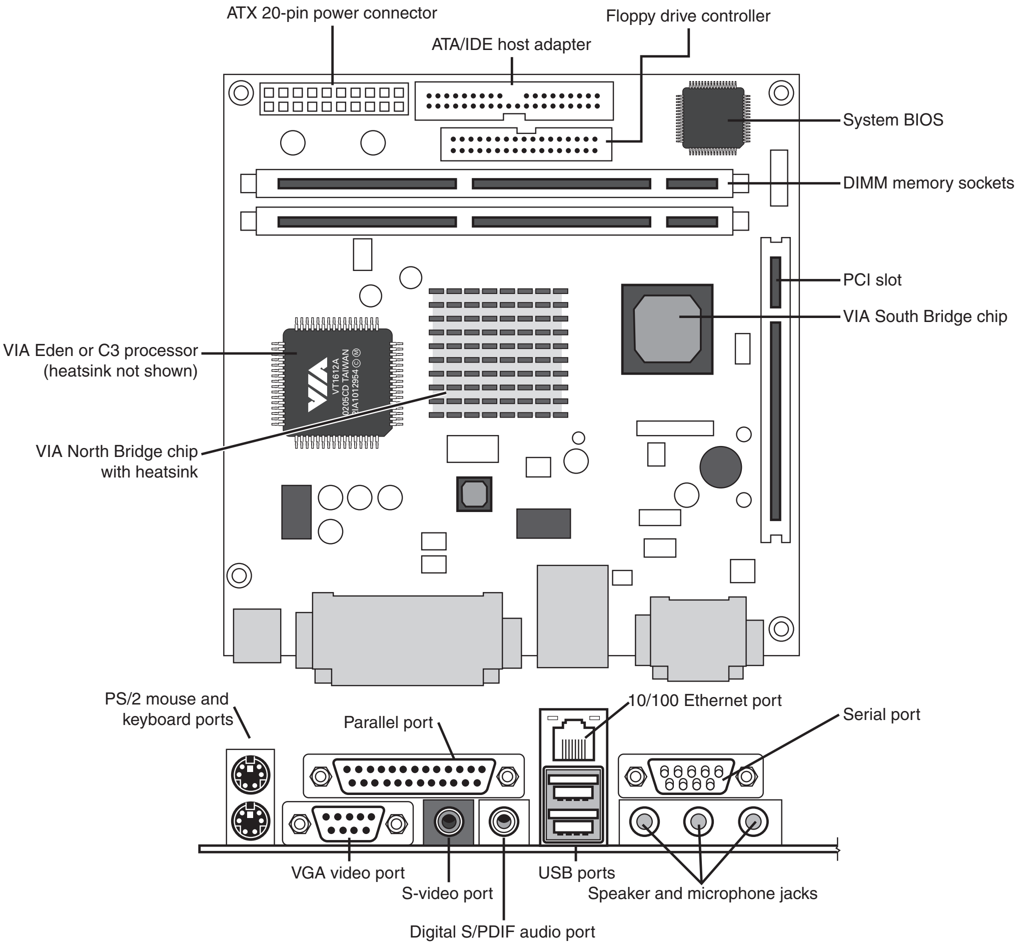

Figure 4.17 Top and rear views of the VIA EPIA-V motherboard, a typical mini-ITX motherboard. *Photo courtesy VIA Technologies, Inc.*

As Figure 4.17 makes clear, mini-ITX motherboards can offer a full range of input-output ports. However, several differences exist between the mini-ITX motherboards and other ATX designs:

- The processor on a mini-ITX motherboard is usually permanently soldered to the board, making future processor upgrades or replacements impossible.
- Most mini-ITX chassis use TFX power supplies, for which there are currently only a few suppliers. Consequently, replacements for them are more expensive and more difficult to find.
- The available TFX power supplies are rated for less output than larger supplies, typically up to 240 watts maximum.
- There is no provision for replacing onboard video with an AGP video card.

Because mini-ITX boards and chassis are made by only a small number of suppliers, future upgrades or parts replacements are limited. However, mini-ITX boards are actually flex-ATX boards, so they can be installed in any standard flex-ATX, micro-ATX, or full-size ATX chassis and use the corresponding power supplies. The only caveats are that the smaller mini-ITX chassis will not accept larger flex-ATX, micro-ATX, or full-size ATX boards and most mini-ITX chassis accept only TFX power supplies. When you select a mini-ITX system, you must be sure to select the appropriate processor type and speed necessary for the task you need it to perform because processor replacements or upgrades almost always require changing the entire motherboard.

Both the VIA C3 E-series and the VIA Eden are x86-compatible processors, so they run the same operating systems and applications as typical AMD and Intel processors, including Windows and Linux. However, the C3 and Eden processors are significantly slower than the Celeron 4, Pentium 4, and Athlon XP processors found in typical desktop and notebook computers. Table 4.5 provides technical information about these processors.

Table 4.5 VIA C3 E-Series and Eden Processors

Processor Model	Clock Speed	FSB Speed	Voltage	L1 Cache	L2 Cache
Eden ESP 4000	400MHz	100MHz	1.05V	128KB	64KB
Eden ESP 5000	533MHz	133MHz	1.20V	128KB	64KB
Eden ESP 6000	600MHz	133MHz	1.20V	128KB	64KB
C3 E-series	733MHz	133MHz	1.35V	64KB	128KB
C3 E-series	800MHz	133MHz	1.35V	64KB	128KB
C3 E-series	866MHz	133MHz	1.35V	64KB	128KB
C3 E-series	933MHz	133MHz	1.35V	64KB	128KB
C3 E-series	1GHz	133MHz	1.35V	64KB	128KB

The Eden ESP series of processors uses a simplified design optimized for the most common operations. Eden ESP processors are an excellent choice for set-top boxes and Internet clients, but they are not as powerful in features or clock speed as the C3 E-series. VIA recommends Eden processors for embedded systems because the processor can be operated without a fan by using a passive heatsink for cooling. For more traditional computer tasks in a small footprint, VIA recommends the C3 E-series because it has performance similar to Intel Celeron processors running at similar speeds. C3 E-series processors can also be run with a passive heatsink, but they run hotter than Eden-series processors because they use higher voltages. VIA recommends a cooling fan for C3 E-series processors unless the case is specially designed to provide adequate cooling for the processor.

VIA uses varying combinations of the following North Bridge and South Bridge chips in its mini-ITX motherboards. The North Bridge uses either the PLE133 or CLE266 chips, whereas the South Bridge uses either the VT8231 or VT8235 chips.

The PLE133 North Bridge chip features built-in Trident AGP 4x video and support for PC100 and PC133 SDRAM memory. The CLE266, on the other hand, features built-in S3 Savage 4 4x AGP video, a built-in MPEG2 decoder for excellent DVD playback, and support for DDR266 SDRAM memory.

The VT8231 South Bridge chip features AC’97 audio, MC’97 modem, an ATA-100 host adapter, and four USB 1.1 ports. Additional capabilities can be added through optional chips. The VT8235 South Bridge chip features six-channel audio, an ATA-133 host adapter, USB 2.0 ports, 10/100 Ethernet, PCI controller, and MC’97 modem. It also supports the new 8X V-Link interface with the North Bridge.

VIA offers a handful of mini-ITX motherboards through its VIA Platform Solutions Division (VPSD), including the following:

- EPIA
- EPIA V (refer to Figure 4.17)
- EPIA M
- EPIA M10000

Table 4.6 summarizes the differences in these motherboards.

Table 4.6 VPSD Mini-ITX Motherboards

Model	North Bridge; South Bridge	Memory Type (Number of Modules)	ATA Type (Number of Ports)	USB Type (Number of Ports)	IEEE-1394a (FireWire)	Floppy
EPIA	PLE133; VT8231	PC100 or PC133 (2)	100 (1)	USB 1.1 (4)	None	No
EPIA V	PLE133; VT8231	PC133 (2)	100 (2)	USB 1.1 (2)	None	Yes
EPIA M	CLE266; VT8235	DDR266 (1)	133 (2)	USB 2.0 (2)	2	Yes

The EPIA M motherboards are the most sophisticated of the three and are the best choice for digital multimedia because of their faster memory subsystem, USB 2.0 and IEEE-1394a ports, and optimized DVD playback.

Obviously, with top performance only on par with sub-1GHz Celeron systems, mini-ITX motherboards are not intended for power user applications. However, if you need a compact system for specialized uses such as home entertainment centers or for small-footprint computers for office suites and Internet access and don't mind investing in a small form factor that might make future upgrades or repairs extremely difficult, these tiny systems can be useful.

Note

The official site for ITX information is `www.via.com.tw/en/VInternet/mini_itx.jsp`. The site `www.mini-itx.com` is often mistaken for an official site, but it is actually a vendor that specializes in ITX systems and component sales.

NLX

NLX is a low-profile form factor designed to replace the nonstandard LPX design used in previous low-profile systems. First introduced in November 1996 by Intel, NLX was a popular form factor in the late 1990s for Slimline corporate desktop systems from vendors such as Compaq, HP, Toshiba, and others. Since 2000, many Slimline systems have used variations on the flex-ATX motherboard instead.

NLX is similar in initial appearance to LPX, but with numerous improvements designed to enable full integration of the latest technologies. NLX is basically an improved version of the proprietary LPX design, but, unlike LPX, NLX is fully standardized, which means you should be able to replace one NLX board with another from a different manufacturer—something that was not possible with LPX.

Another limitation of LPX boards is the difficulty in handling the larger physical size of the newer processors and their larger heatsinks, as well as newer bus structures such as AGP for video. The NLX form factor has been designed specifically to address these problems (see Figure 4.18). In fact, NLX provides enough room for some vendors to support dual Slot 1 Pentium III processors in this form factor.

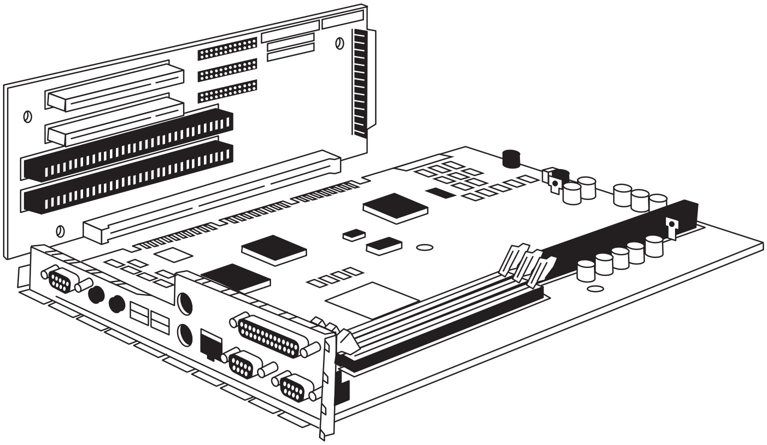

Figure 4.18 NLX motherboard and riser combination.

The main characteristic of an NLX system is that the motherboard plugs into the riser, unlike LPX where the riser plugs into the motherboard. Therefore, the motherboard can be removed from the system without disturbing the riser or any of the expansion cards plugged into it. In addition, the motherboard in a typical NLX system literally has no internal cables or connectors attached to it! All devices that normally plug into the motherboard—such as drive cables, the power supply, the front panel light, switch connectors, and so on—plug into the riser instead (see Figure 4.18). By using the riser card as a connector concentration point, you can remove the lid on an NLX system and literally slide the motherboard out the left side of the system without unplugging a single cable or connector on the inside. This allows for unbelievably quick motherboard changes; in fact, I have swapped motherboards in less than 30 seconds on NLX systems!

As Figure 4.19 shows, by using different sizes and types of riser cards, a system designer can customize the features of a given NLX system.

Such a design is a boon for the corporate market, where ease and swiftness of servicing is a major feature. Not only can components be replaced with lightning speed, but because of the industry-standard design, motherboards, power supplies, and other components can be interchanged even among different systems.

Specific advantages of the NLX form factor include

- *Support for all desktop system processor technologies.* This is especially important because, since the NLX form factor was developed, both AMD and Intel adopted and then abandoned bulkier slot-based processors and returned to more compact socketed processors. NLX can handle both types of processors.
- *Flexibility in the face of rapidly changing processor technologies.* Backplane-like flexibility has been built in to the form by allowing a new motherboard to be easily and quickly installed without tearing your entire system to pieces. But unlike traditional backplane systems, many industry leaders, such as Compaq, Toshiba, and HP, have sold NLX-based systems.
- *Support for newer technologies.* This includes Accelerated Graphics Port (AGP) high-performance graphic solutions, Universal Serial Bus (USB), and memory modules in DIMM or RIMM form.
- *Ease and speed of servicing and repair.* Compared to other industry-standard interchangeable form factors, NLX systems are by far the easiest to work on and allow component swaps or other servicing in the shortest amount of time.

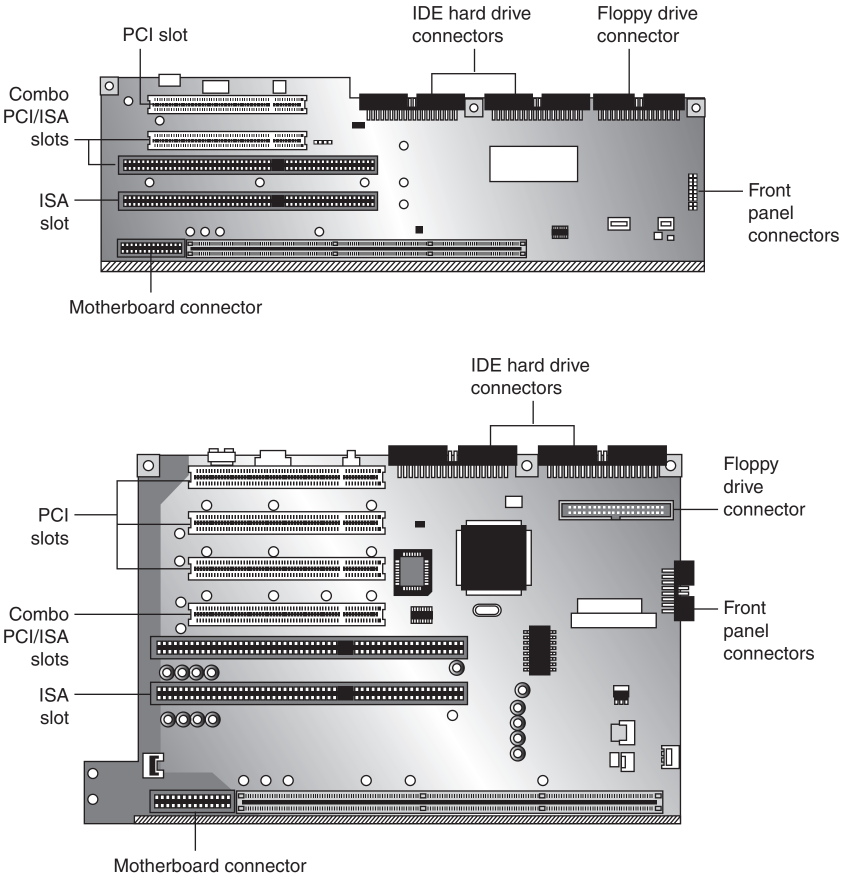

Figure 4.19 Typical NLX riser cards. Although most NLX systems use a low-profile riser card similar to the top riser card, others use a taller riser card to provide more slots for add-on cards.

Furthermore, with the importance of multimedia applications, connectivity support for such things as video playback, enhanced graphics, and extended audio has been built in to the motherboard. This should represent a good cost savings over expensive daughterboard arrangements that have been necessary for many advanced multimedia uses in the past. Although ATX also has this support, LPX and Baby-AT don't have the room for these additional connectors.

Figure 4.20 shows the basic NLX system layout. Notice that, similar to ATX, the motherboard is clear of the drive bays and other chassis-mounted components. Also, the motherboard and I/O cards (which, like the LPX form factor, are mounted parallel to the motherboard) can easily be slid in to and out of the side of the chassis, leaving the riser card and other cards in place. The processor can be easily accessed and enjoys greater cooling than in a more closed-in layout.

Note the position of the optional AGP slot shown previously in Figure 4.13. It is mounted on the motherboard itself, not on the riser card as with PCI or ISA slots. This location was necessary because AGP was developed well after the NLX form factor was introduced. Most NLX motherboards use chipset-integrated or motherboard-based video instead of a separate AGP card, but you must remove an AGP card installed in an NLX system before you can remove the motherboard for servicing. Also, the AGP card used in an NLX system must have a different form factor to enable it to clear the rear connector shield at the back of the NLX motherboard (see Figure 4.21).

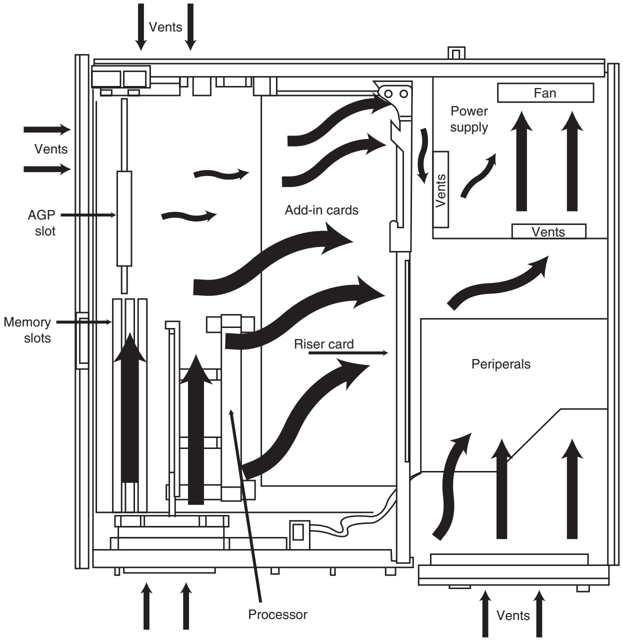

Figure 4.20 NLX system chassis layout and cooling airflow.

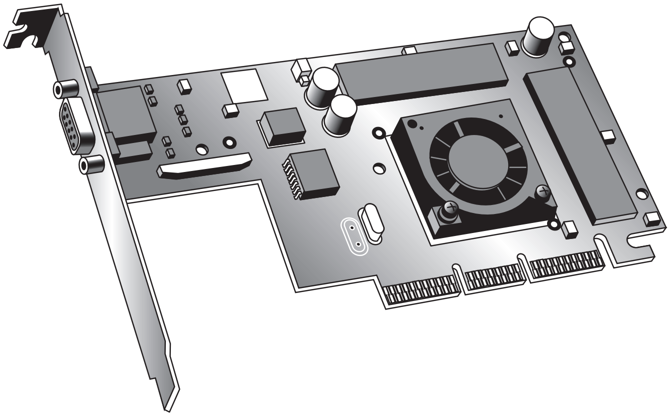

Figure 4.21 An AGP card that can be installed in either a standard ATX/Baby-AT system or an NLX system. This is because of the shape, which leaves room for the NLX motherboard's rear connector shield. *Photo courtesy Elsa AG.*

The NLX motherboard is specified in three lengths, front to back: 13.6", 11.2", or 10" total (see Figure 4.22). With proper bracketry, the shorter boards can go into a case designed for a longer board.

As with most of the form factors, you can identify NLX via the unique I/O shield or connector area at the back of the board (see Figure 4.23). You only need a quick look at the rear of any given system to determine which type of board is contained within. Figure 4.23 shows the unique stepped design of the NLX I/O connector area. This allows for a row of connectors all along the bottom and has room for double-stacked connectors on one side.

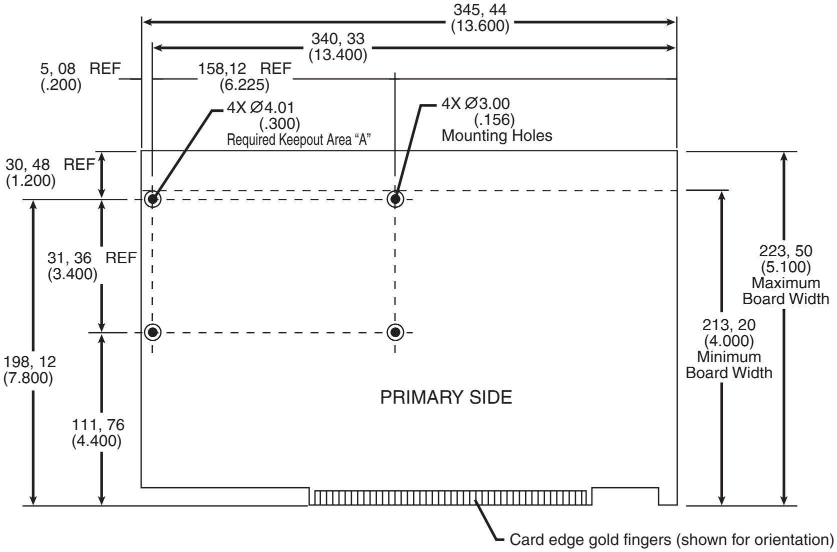

Figure 4.22 NLX form factor. This shows a 13.6" long NLX board. The NLX specification also allows shorter 11.2" and 10" versions.

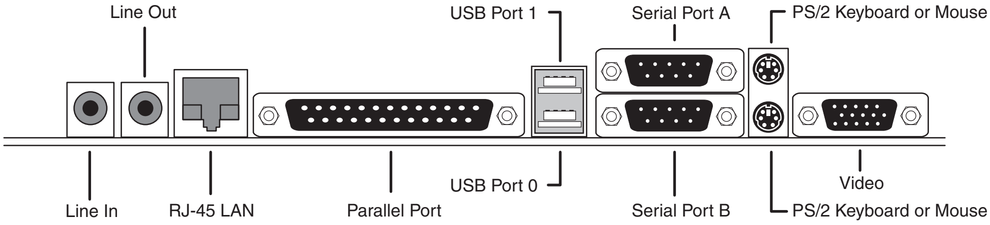

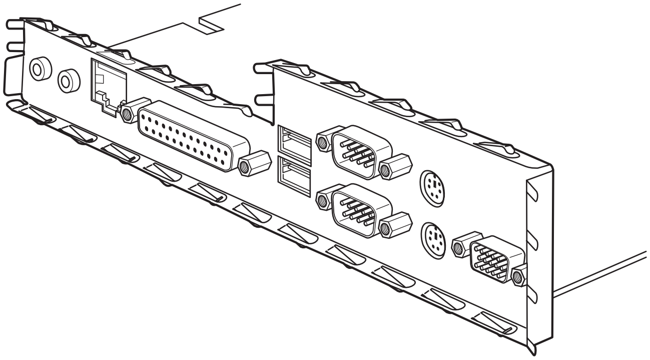

Figure 4.23 Typical NLX motherboard rear connector layout.

As you can see, the NLX form factor has been designed for maximum flexibility and space efficiency. Even extremely long I/O cards will fit easily without getting in the way of other system components—a problem with Baby-AT form factor systems.

The specification and related information about the NLX form factor are available at the Desktop Form Factor site located at `www.formfactors.org`. Although NLX is a standard form factor—just as the ATX family is—most NLX products have been sold as part of complete systems aimed at the corporate market. Very few aftermarket motherboards have been developed in this form factor. Thus, although NLX makes swapping motherboards easy, you are more likely to encounter it in a corporate environment than in a home or small-business computer. The micro-ATX and flex-ATX form factors have largely superseded NLX in the markets formerly dominated by LPX. Overall, one of the ATX variants is still the best choice for most new systems where expandability, upgradeability, low cost, and ease of service are of prime importance.

WTX

WTX was a board and system form factor developed for the mid-range workstation market; however, it didn't seem to catch on. WTX went beyond ATX and defined the size and shape of the board and the interface between the board and chassis, as well as required chassis features.

WTX was first released in September 1998 (1.0) and updated in February 1999 (1.1). The specification and other information on WTX used to be available at `www.wtx.org`; however, WTX has been discontinued and there will be no further updates. It is not recommended to build new systems using the WTX form factor.

The very few WTX form factor systems that were introduced were designed as servers. Figure 4.24 shows a typical WTX system with the cover removed. Note that easy access is provided to internal components via pull-out drawers and swinging side panels.

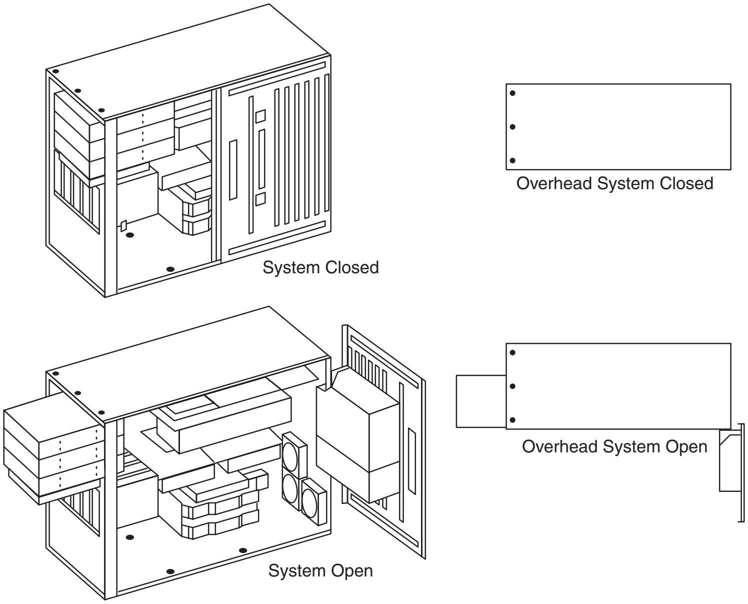

Figure 4.24 Typical WTX system chassis showing internal layout and ease of access.

WTX motherboards have a maximum width of 14" (356mm) and a maximum length of 16.75" (425mm), which is significantly larger than ATX. There are no minimum dimensions, so board designers are free to design smaller boards as long as they meet the mounting criteria.

The WTX specification offers flexibility by leaving motherboard mounting features and locations undefined. Instead of defining exact screw hole positions, WTX motherboards must mount to a standard mounting adapter plate, which must be supplied with the board. The WTX chassis is designed to accept the mounting plate with attached motherboard and not just a bare board alone.

Proprietary Designs

Motherboards that are *not* one of the industry standard form factors, such as AT/Baby-AT, NLX, or any of the ATX formats, are deemed *proprietary* or *semi-proprietary*. LPX, ITX, and mini-ITX systems fall into the semi-proprietary category for example, while other companies have fully proprietary systems that only they manufacture. Most people purchasing PCs should avoid proprietary designs because they do not allow for a future motherboard, power supply, or case upgrade, which limits future use and serviceability of the system. To me, proprietary systems are disposable PCs because you can neither upgrade them nor easily repair them. The problem is that the proprietary parts can come only from the original system manufacturer, and they usually cost much more than nonproprietary parts. Therefore, after your proprietary system goes out of warranty, not only is it not upgradeable, but it is also essentially no longer worth repairing. If the motherboard or any component on it goes bad, you will be better off purchasing a completely new standard system than paying five times the normal price for a new proprietary motherboard. In addition, a new motherboard in a standard form factor system would be one or more generations newer and faster than the one you would be replacing. In a proprietary system, the replacement board would not only cost way too much, but it would be the same as the one that failed.

Note that you might be able to perform limited upgrades to older systems with proprietary motherboards, in the form of custom (non-OEM) processor replacements with attached voltage regulators, usually called “overdrive” or “turbo” chips. Unfortunately, these often don’t perform up to the standards of a less expensive new processor and motherboard combination. Of course, I usually recommend upgrading the motherboard and processor together—but that is something that can’t be done with a proprietary system.

Until the late 1990s, the LPX motherboard design was at the heart of most proprietary systems. These systems were sold primarily in the retail store channel by Compaq, IBM’s Aptiva line, HP’s Vectra line, and Packard Bell (no longer in business in North America). As such, virtually all their systems have problems inherent with their proprietary designs.

If the motherboard in your current ATX form factor system dies, you can find any number of replacement boards that will bolt directly in—with your choice of processors and clock speeds—at great prices. You can also find replacements for Baby-AT motherboards, but this form factor doesn’t support the newest technologies and has not been used in new system designs for several years. However, if the motherboard dies in a proprietary form factor system, you’ll pay for a replacement available only from the original manufacturer, and you have little or no opportunity to select a board with a faster or better processor than the one that failed. In other words, upgrading or repairing one of these systems via a motherboard replacement is difficult and usually not cost-effective.

Systems sold by the leading mail-order suppliers, such as Gateway, Micron, Dell, and others, are available in industry-standard form factors such as ATX, micro-ATX, flex-ATX, and NLX. This allows for easy upgrading and system expansion in the future. These standard factors allow you to replace your own motherboards, power supplies, and other components easily and select components from any number of suppliers other than where you originally bought the system.

Backplane Systems

One type of design that has been used in some systems over the years is the *backplane system*. These systems do not have a motherboard in the true sense of the word. In a backplane system, the components typically found on a motherboard are located instead on an expansion adapter card plugged into a slot.

In these systems, the board with the slots is called a backplane, rather than a motherboard. Systems using this type of construction are called backplane systems.

Backplane systems come in two main types—passive and active. A *passive* backplane means the main backplane board does not contain any circuitry at all except for the bus connectors and maybe some buffer and driver circuits. All the circuitry found on a conventional motherboard is contained on one or more expansion cards installed in slots on the backplane. Some backplane systems use a passive design that incorporates the entire system circuitry into a single mothercard. The mothercard is essentially a complete motherboard designed to plug into a slot in the passive backplane. The passive backplane/mothercard concept enables the entire system to be easily upgraded by changing one or more cards. Because of the expense of the high-function mothercard, this type of system design is rarely found in standard PC systems today, although it was once favored by a few early 286/386 vendors such as Zenith Data Systems. The passive backplane design does enjoy popularity in industrial systems, which are often rack-mounted. Some high-end fileservers also feature this design. Figure 4.25 shows a typical Pentium III single-board computer used in passive backplane systems. Figure 4.26 shows a rack-mount chassis with a passive backplane.

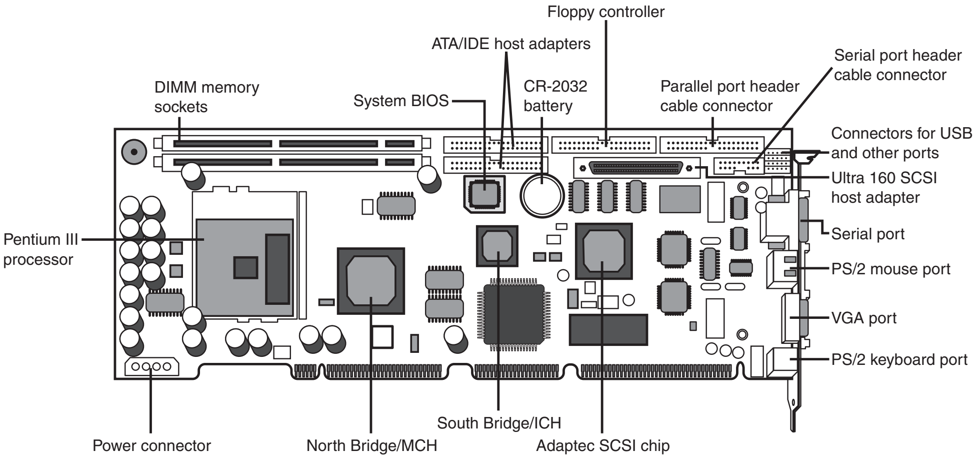

Figure 4.25 A typical Pentium III/Celeron III PICMG single-board computer. This single card provides PCI and ISA interfacing, integrated AGP video with support for DVI and CRT displays, three 10/100 Ethernet network interfaces, and a 68-pin Wide SCSI interface, as well as normal parallel, serial, IDE, and floppy interfaces.

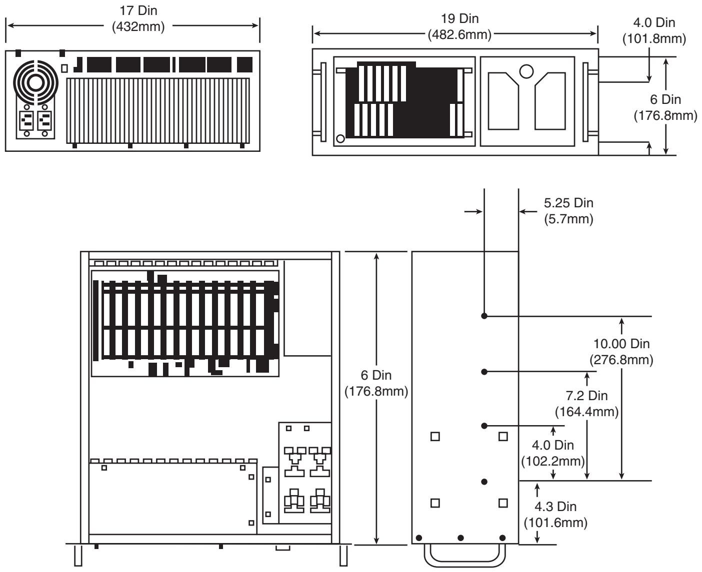

Figure 4.26 A rack-mount chassis with passive backplane.

Passive backplane systems with mothercards (often called single-board computers or SBCs) are by far the most popular backplane design. They are used in industrial or laboratory-type systems and are rack-mountable. They usually have a large number of slots and extremely heavy-duty power supplies; they also feature high-capacity, reverse flow cooling designed to pressurize the chassis with cool, filtered air. Many passive backplane systems, such as the one pictured in Figure 4.27, adhere to the PCI/ISA passive backplane and CompactPCI form factor standards set forth by the PCI Industrial Computer Manufacturers Group (PICMG). You can get more information about these standards from PICMG's Web site at `www.picmg.org`.

An *active* backplane means the main backplane board contains bus control and usually other circuitry as well. Most active backplane systems contain all the circuitry found on a typical motherboard except for what is then called the *processor complex*. The processor complex is the name of the circuit board that contains the main system processor and any other circuitry directly related to it, such as clock control, cache, and so forth. The processor's complex design enables the user to easily upgrade the system later to a new processor type by changing one card. In effect, it amounts to a modular motherboard with a replaceable processor section.

Many large PC manufacturers have built systems with an active backplane/processor complex. Both IBM and Compaq, for example, have used this type of design in some of their high-end (server class) systems. ALR once made a series of desktop and server PCs that also featured this design. This allows an easier and generally more affordable upgrade than the passive backplane/mothercard design because the processor complex board is usually much cheaper than a mothercard. Unfortunately, because no standards exist for the processor complex interface to the system, these boards are proprietary and can be purchased only from the system manufacturer. This limited market and availability causes the prices of these boards to be higher than most complete motherboards from other manufacturers.

The motherboard system design and the backplane system design have advantages and disadvantages. Most original PCs were designed as backplanes in the late 1970s. Apple and IBM shifted the market to the now traditional motherboard with a slot-type design because this kind of system generally is cheaper to mass-produce than one with the backplane design. In the late 1980s, Zenith Data manufactured a line of backplane-based 8088, 286, and 386-based systems but later abandoned this for a standard motherboard design similar to other vendors. The theoretical advantage of a backplane system, however, is that you can easily upgrade it to a new processor and level of performance by changing a single card. For example, you can upgrade a system's processor just by changing the card. In a motherboard-design system, you often must change the motherboard, a seemingly more formidable task. Unfortunately, the reality of the situation is that a backplane design is frequently much more expensive to upgrade. For example, because the bus remains fixed on the backplane, the backplane design precludes more comprehensive upgrades that involve adding local bus slots.

Another nail in the coffin of backplane designs is the upgradeable processor. Starting with the 486, Intel and AMD began standardizing the sockets or slots in which processors were to be installed, allowing a single motherboard to support a wider variety of processors and system speeds. Because board designs could be made more flexible, changing only the processor chip for a faster standard OEM type (not one of the kludgy "overdrive" chips) is the easiest and most cost-effective way to upgrade without changing the entire motherboard.

Because of the limited availability of the processor-complex boards or mothercards, they usually end up being more expensive than a complete new motherboard that uses an industry-standard form factor. The bottom line is that unless you have a requirement for a large-capacity industrial or laboratory-type system, especially one that would be rack-mounted, you are better off sticking with standard ATX form factor PCs. They will certainly be far less expensive.

Note

Some companies offer plug-in processor cards that essentially turn your existing motherboard into an active backplane, shutting down the main CPU and memory and having the card's processor and memory essentially take over. These are, unfortunately, much more expensive than a new motherboard and processor alone, use the more expensive SO-DIMM memory, don't provide AGP video, and are generally not recommended.

Motherboard Components

A modern motherboard has several components built in, including various sockets, slots, connectors, chips, and so on. This section examines the components found on a typical motherboard.

Most modern motherboards have at least the following major components on them:

- Processor socket/slot
- Chipset (North/South Bridge or memory and I/O controller hubs)
- Super I/O chip
- ROM BIOS (Flash ROM/firmware hub)
- SIMM/DIMM/RIMM (RAM memory) sockets
- ISA/PCI/AGP bus slots
- CPU voltage regulator
- Battery

Some motherboards also include integrated video, audio, networking, SCSI, Audio Modem Riser (AMR), Communications and Networking Riser (CNR) connectors, or other optional interfaces, depending on the individual board.

These standard components are discussed in the following sections.

Processor Sockets/Slots

The CPU is installed in either a socket or a slot, depending on the type of chip.

Starting with the 486 processors, Intel designed the processor to be a user-installable and replaceable part and developed standards for CPU sockets and slots that would allow different models of the same basic processor to plug in. These specifications were given a designation that is usually imprinted or embossed on the connector or board. If you know the type of socket or slot on your motherboard, you essentially know which type of processors are designed to plug in.

◀◀ See "Processor Socket and Slot Types," p. 85.

Sockets for processors prior to the 486 were not designated, and interchangeability was limited. Table 4.7 shows the designations for the various processor sockets/slots and lists the chips designed to plug into them.

Originally, all processors were mounted in sockets (or soldered directly to the motherboard). With the advent of the Pentium II and original Athlon processors, both Intel and AMD shifted to a slot-based approach for their processors because the processors now incorporated built-in L2 cache, purchased as separate chips from third-party Static RAM (SRAM) memory chip manufacturers. Therefore, the processor then consisted not of one but of several chips, all mounted on a daughterboard that was then plugged into a slot in the motherboard. This worked well, but there were additional expenses in the extra cache chips, the daughterboard itself, the slot, optional casings or packaging, and the support mechanisms and physical stands and latches for the processor and heatsink. All in all, slot-based processors were expensive to produce compared to the previous socketed versions.

Table 4.7 CPU Socket Specifications

Socket Number	Pins	Pin Layout	Voltage	Supported Processors
Socket 1	169	17×17 PGA	5V	486 SX/SX2, DX/DX2[1], DX4 OD
Socket 2	238	19×19 PGA	5V	486 SX/SX2, DX/DX2[1], DX4 OD, 486 Pentium OD
Socket 3	237	19×19 PGA	5V/3.3V	486 SX/SX2, DX/DX2, DX4, 486 Pentium OD, AMD 5x86
Socket 4	273	21×21 PGA	5V	Pentium 60/66, OD
Socket 5	320	37×37 SPGA	3.3V/3.5V	Pentium 75-133, OD
Socket 6[2]	235	1×19 PGA	3.3V	486 DX4, 486 Pentium OD
Socket 7	321	37×37 SPGA	VRM	Pentium 75-233+, MMX, OD, AMD K5/K6, Cyrix M1/II
Socket 8	387	Dual-pattern SPGA	Auto VRM	Pentium Pro, OD
Socket 370	370	37×37 SPGA	Auto VRM	Celeron/Pentium III PPGA/FC-PGA
Socket PAC418	418	38×22 split SPGA	Auto VRM	Itanium, Itanium 2
Socket 423	423	39×39 SPGA	Auto VRM	Pentium 4 FC-PGA
Socket A (462)	462	37×37 SPGA	Auto VRM	AMD Athlon/Athlon XP/Duron PGA
Socket 478	478	26×26 mPGA	Auto VRM	Pentium 4 FC-PGA2, Celeron FC-PGA2
Socket 603	603	31×25 mPGA	Auto VRM	Xeon (P4)
Socket 754	754	29×29 mPGA	Auto VRM	AMD Athlon 64
Slot A	242	Slot	Auto VRM	AMD Athlon SECC
Slot 1 (SC242)	242	Slot	Auto VRM	Pentium II/III, Celeron SECC
Slot 2 (SC330)	330	Slot	Auto VRM	Pentium II/III Xeon SECC

1. Non-overdrive DX4 or AMD 5x86 also can be supported with the addition of an aftermarket 3.3V voltage regulator adapter.

2. Socket 6 was never actually implemented in any systems.

Auto VRM = Voltage regulator module with automatic voltage selection determined by processor VID pins

FC-PGA = Flip-chip pin grid array

OD = OverDrive (retail upgrade processors)

PAC = Pin array cartridge

PGA = Pin grid array

PPGA = Plastic pin grid array

SC242 = Slot connector, 242 pins

SC330 = Slot connector, 330 pins

SECC = Single edge contact cartridge

SPGA = Staggered pin grid array

mPGA = Micro pin grid array

VRM = Voltage regulator module with variable voltage output determined by module type or manual jumpers

With the advent of the second-generation Celeron, Intel integrated the L2 cache directly into the processor die, meaning within the main CPU chip circuits with no extra chips required. The second-generation (code named Coppermine) Pentium III also received on-die L2 cache, as did the K6-3, Duron (code named Spitfire), and second-generation Athlon (code named Thunderbird) processors from AMD (some early Thunderbird Athlon CPUs were also made in the Slot A configuration). With on-die L2, the processor was back to being a single chip again, which also meant that mounting it on a separate board plugged into a slot was expensive and unnecessary. Because of on-die integrated L2 cache, the trend for processor packaging has shifted back to sockets and will likely continue that way for the foreseeable future. All modern processors now use the socket form. Besides allowing a return to socketed packaging, the on-die L2 cache runs at full processor speed, instead of the one-half or one-third speed of the previous integrated (but not on-die) L2 cache.

The 64-bit Itanium processor from Intel, however, features a cartridge design that includes L3 cache and yet plugs into a socket rather than a slot.

Chipsets

We can't talk about modern motherboards without discussing chipsets. The chipset *is* the motherboard; therefore, any two boards with the same chipsets are functionally identical. The chipset contains the processor bus interface (called front-side bus, or FSB), memory controllers, bus controllers, I/O controllers, and more. All the circuits of the motherboard are contained within the chipset. If the processor in your PC is like the engine in your car, the chipset represents the chassis. It is the framework in which the engine rests and is its connection to the outside world. The chipset is the frame, suspension, steering, wheels and tires, transmission, driveshaft, differential, and brakes. The chassis in your car is what gets the power to the ground, allowing the vehicle to start, stop, and corner. In the PC, the chipset represents the connection between the processor and everything else. The processor can't talk to the memory, adapter boards, devices, and so on without going through the chipset. The chipset is the main hub and central nervous system of the PC. If you think of the processor as the brain, the chipset is the spine and central nervous system.

Because the chipset controls the interface or connections between the processor and everything else, the chipset ends up dictating which type of processor you have; how fast it will run; how fast the buses will run; the speed, type, and amount of memory you can use; and more. In fact, the chipset might be the single most important component in your system, possibly even more important than the processor. I've seen systems with faster processors be outperformed by systems with slower processor but a better chipset, much like how a car with less power might win a race through better cornering and braking. When deciding on a system, I start by choosing the chipset first because the chipset decision then dictates the processor, memory, I/O, and expansion capabilities.

Chipset Evolution

When IBM created the first PC motherboards, it used several discrete (separate) chips to complete the design. Besides the processor and optional math coprocessor, many other components were required to complete the system. These other components included items such as the clock generator, bus controller, system timer, interrupt and DMA controllers, CMOS RAM and clock, and keyboard controller. Additionally, many other simple logic chips were used to complete the entire motherboard circuit, plus, of course, things such as the actual processor, math coprocessor (floating-point unit), memory, and other parts. Table 4.8 lists all the primary chip components used on the original PC/XT and AT motherboards.

Table 4.8 Primary Chip Components on PC/XT and AT Motherboards

Chip Function	PC/XT Version	AT Version
Processor	8088	80286
Math Coprocessor (Floating-Point Unit)	8087	80287
Clock Generator	8284	82284
Bus Controller	8288	82288
System Timer	8253	8254
Low-order Interrupt Controller	8259	8259
High-order Interrupt Controller	—	8259
Low-order DMA Controller	8237	8237
High-order DMA Controller	—	8237
CMOS RAM/Real-Time Clock	—	MC146818
Keyboard Controller	8255	8042

In addition to the processor/coprocessor, a six-chip set was used to implement the primary motherboard circuit in the original PC and XT systems. IBM later upgraded this to a nine-chip design in the

AT and later systems, mainly by adding more interrupt and DMA controller chips and the nonvolatile CMOS RAM/Real-time Clock chip. All these motherboard chip components came from Intel or an Intel-licensed manufacturer, except the CMOS/Clock chip, which came from Motorola. To build a clone or copy of one of these IBM systems required all these chips plus many smaller discrete logic chips to glue the design together, totaling 100 or more individual chips. This kept the price of a motherboard high and left little room on the board to integrate other functions.

In 1986, a company called Chips and Technologies introduced a revolutionary component called the 82C206—the main part of the first PC motherboard chipset. This was a single chip that integrated into it all the functions of the main motherboard chips in an AT-compatible system. This chip included the functions of the 82284 Clock Generator, 82288 Bus Controller, 8254 System Timer, dual 8259 Interrupt Controllers, dual 8237 DMA Controllers, and even the MC146818 CMOS/Clock chip. Besides the processor, virtually all the major chip components on a PC motherboard could now be replaced by a single chip. Four other chips augmented the 82C206 acting as buffers and memory controllers, thus completing virtually the entire motherboard circuit with five total chips. This first chipset was called the CS8220 chipset by Chips and Technologies. Needless to say, this was a revolutionary concept in PC motherboard manufacturing. Not only did it greatly reduce the cost of building a PC motherboard, but it also made designing a motherboard much easier. The reduced component count meant the boards had more room for integrating other items formerly found on expansion cards. Later, the four chips augmenting the 82C206 were replaced by a new set of only three chips, and the entire set was called the New Enhanced AT (NEAT) CS8221 chipset. This was later followed by the 82C836 Single Chip AT (SCAT) chipset, which finally condensed all the chips in the set down to a single chip.

The chipset idea was rapidly copied by other chip manufacturers. Companies such as Acer, Erso, Opti, Suntac, Symphony, UMC, and VLSI each gained an important share of this market. Unfortunately for many of them, the chipset market has been a volatile one, and many of them have long since gone out of business. In 1993, VLSI had become the dominant force in the chipset market and had the vast majority of the market share; by the next year, VLSI (which later was merged into Philips Semiconductors), along with virtually everybody else in the chipset market, was fighting to stay alive. This is because a new chipset manufacturer had come on the scene, and within a year or so of getting serious, it was totally dominating the chipset market. That company was Intel, and after 1994, it had a virtual lock on the chipset market. If you have a motherboard built since 1994 that uses or accepts an Intel processor, chances are good that it has an Intel chipset on it as well.

More recently, Intel has struggled somewhat with chipsets because of its reliance on RDRAM memory. Intel originally signed a contract with Rambus back in 1996 declaring it would support this memory as its primary focus for desktop PC chipsets through 2001. I suspect this has turned out to be something Intel regrets (the contract has since expired). RDRAM memory has had a significantly higher price than SDRAM memory—although the prices have recently come down a bit—and it does have some performance advantages when used in a dual-channel mode. There is a lot of momentum in the market for supporting double data rate (DDR) SDRAM. Consequently, Intel introduced the 845 chipset (code named Brookdale), which supports DDR-SDRAM with the Pentium 4. Intel's latest chipsets for Pentium 4, the i865 (code named Springdale) and i875P (code named Canterwood), continue the trend of support for faster DDR SDRAM memory and FSB bus speeds. Intel is not alone in the Pentium 4 chipset business: Silicon Integrated Systems (SiS), ATI, and ALi Corporation all make licensed chipsets for the Pentium 4. VIA Technologies also makes Pentium 4–compatible chipsets, but without a license, which has greatly limited their popularity with motherboard makers.

Although AMD has developed its own chipsets for the K6 and Athlon family of processors, it now emphasizes encouraging third-party chipset developers to support its products. Today, VIA Technologies is the leading developer of AMD Athlon/Athlon XP/Duron chipsets for both discrete and integrated uses. The popularity of AMD processors has also encouraged SiS, ATI, NVIDIA, and ALi Corporation to develop chipsets for both Intel- and AMD-based systems.

It is interesting to note that the original PC chipset maker, Chips and Technologies, survived by changing course to design and manufacture video chips and found a niche in that market specifically for laptop and notebook video chipsets. Chips and Technologies was subsequently bought out by Intel in 1998 as a part of Intel's video strategy.

Intel Chipsets

You can't talk about chipsets today without discussing Intel because it currently owns the vast majority of the chipset market. It is interesting to note that we probably have Compaq to thank for forcing Intel into the chipset business in the first place!

The thing that really started it all was the introduction of the EISA bus designed by Compaq in 1989. At that time, it had shared the bus with other manufacturers in an attempt to make it a market standard. However, Compaq refused to share its EISA bus chipset—a set of custom chips necessary to implement this bus on a motherboard.

Enter Intel, who decided to fill the chipset void for the rest of the PC manufacturers wanting to build EISA bus motherboards. As is well known today, the EISA bus failed to become a market success except for a short-term niche server business, but Intel now had a taste of the chipset business and this it apparently wouldn't forget. With the introduction of the 286 and 386 processors, Intel became impatient with how long it took the other chipset companies to create chipsets around its new processor designs; this delayed the introduction of motherboards that supported the new processors. For example, it took more than two years after the 286 processor was introduced for the first 286 motherboards to appear and just over a year for the first 386 motherboards to appear after the 386 had been introduced. Intel couldn't sell its processors in volume until other manufacturers made motherboards that would support them, so it thought that by developing motherboard chipsets for a new processor in parallel with the new processor, it could jumpstart the motherboard business by providing ready-made chipsets for the motherboard manufacturers to use.

Intel tested this by introducing the 420 series chipsets along with its 486 processor in April 1989. This enabled the motherboard companies to get busy right away, and in only a few months the first 486 motherboards appeared. Of course, the other chipset manufacturers weren't happy; now they had Intel as a competitor, and Intel would always have chipsets for new processors on the market first!

Intel then realized that it made both processors *and* chipsets, which were 90% of the components on a typical motherboard. What better way to ensure that motherboards were available for its Pentium processor when it was introduced than by making its own motherboards as well and having these boards ready on the new processor's introduction date. When the first Pentium processor debuted in 1993, Intel also debuted the 430LX chipset as well as a fully finished motherboard. Now, besides the chipset companies being upset, the motherboard companies weren't too happy, either. Intel was not only the major supplier of parts needed to build finished boards (processors and chipsets), but was now building and selling the finished boards as well. By 1994, Intel dominated the processor and chipset markets and had cornered the motherboard market as well.

Now as Intel develops new processors, it develops chipsets and motherboards simultaneously, which means they can be announced and shipped in unison. This eliminates the delay between introducing new processors and waiting for motherboards and systems capable of using them, which was common in the industry's early days. For the consumer, this means no waiting for new systems. Since the original Pentium processor in 1993, we have been able to purchase ready-made systems on the same day a new processor is released.

In my seminars, I ask how many people in the class have Intel-brand PCs. Of course, Intel does not sell or market a PC under its own name, so nobody thinks they have an "Intel-brand" PC. But, if your motherboard was made by Intel, for all intents and purposes you sure seem to have an Intel-brand PC, at least as far as the components are concerned. Does it really matter whether Dell, Gateway, or Micron put that same Intel motherboard into a slightly different looking case with their name on it?

If you look under the covers, you'll find that many, if not most, of the systems from the major manufacturers are really the same because they basically use the same parts. Although more and more major manufacturers are offering AMD Athlon- and Duron-based systems as alternatives to Intel's, no manufacturer dominates AMD motherboard sales the way Intel has dominated OEM sales to major system manufacturers.

To hold down pricing, many low-cost retail systems based on micro-ATX motherboards use non-Intel motherboards (albeit with Intel chipsets in most cases). But, even though many companies make Intel-compatible motherboards for aftermarket upgrades or local computer assemblers, Intel still dominates the major vendor OEM market for midrange and high-end systems.

Intel Chipset Model Numbers

Starting with the 486 in 1989, Intel began a pattern of numbering its chipsets as follows:

Chipset Number	Processor Family
420xx	P4 (486)
430xx	P5 (Pentium)
440xx	P6 (Pentium Pro/PII/PIII)
8xx	P6/P7 (PII/PIII/P4) with hub architecture
450xx	P6 server (Pentium Pro/PII/PIII Xeon)
E72xx	Xeon workstation with hub architecture
E75xx	Xeon server with hub architecture
460xx	Itanium processor
E88xx	Itanium 2 processor with hub architecture

The chipset numbers listed here are abbreviations of the actual chipset numbers stamped on the individual chips. For example, one of the popular Pentium II/III chipsets was the Intel 440BX chipset, which consisted of two components: the 82443BX North Bridge and the 82371EX South Bridge. Likewise, the 845 chipset supports the Pentium 4 and consists of two main parts, including the 82845 Memory Controller Hub (MCH; replaces the North Bridge) and an 82801BA I/O Controller Hub (ICH2; replaces the South Bridge). By reading the logo (Intel or others) as well as the part number and letter combinations on the larger chips on your motherboard, you can quickly identify the chipset your motherboard uses.

Intel has used two distinct chipset architectures: a North/South Bridge architecture and a newer hub architecture. All its more recent 800 series chipsets use the hub architecture.

AMD Athlon/Duron Chipsets

AMD took a gamble with its Athlon family of processors (Athlon, Athlon XP, Athlon MP, and the now-discontinued Duron). With these processors, AMD decided for the first time to create a chip that was Intel compatible with regards to software but not directly hardware or pin compatible. Whereas the K6 series would plug into the same Socket 7 that Intel designed for the Pentium processor line, the AMD Athlon and Duron would not be pin compatible with the Pentium II/III and Celeron chips. This also meant that AMD could not take advantage of the previously existing chipsets and motherboards when the Athlon and Duron were introduced; instead, AMD would have to either create its own chipsets and motherboards or find other companies who would.

The gamble seems to have paid off. AMD bootstrapped the market by introducing its own chipset, referred to as the AMD-750 chipset (code named Irongate). The AMD 750 chipset consists of the 751 System Controller (North Bridge) and the 756 Peripheral Bus Controller (South Bridge). More recently, AMD introduced the AMD-760 chipset for the Athlon/Duron processors, which is the first major chipset on the market supporting DDR SDRAM for memory. It consists of two chips—the AMD-761 System Bus Controller (North Bridge) and the AMD-766 Peripheral Bus Controller (South Bridge). Although AMD no longer puts much emphasis on chipset sales, its pioneering efforts have inspired other companies, such as VIA Technologies, NVIDIA, and SiS, to develop chipsets specifically designed for the Slot A and current Socket A and Socket 754 processors from AMD. This has enabled the motherboard companies to make a variety of boards supporting these chips and the Athlon processors to take a fair amount of market share away from Intel in the process.

North/South Bridge Architecture

Most of Intel's earlier chipsets (and virtually all non-Intel chipsets) are broken into a multi-tiered architecture incorporating what are referred to as North and South Bridge components, as well as a Super I/O chip:

- *The North Bridge.* So named because it is the connection between the high-speed processor bus (400/266/200/133/100/66MHz) and the slower AGP (533/266/133/66MHz) and PCI (33MHz) buses. The North Bridge is what the chipset is named after, meaning that, for example, what we call the 440BX chipset is derived from the fact that the actual North Bridge chip part number for that set is 82443BX.
- *The South Bridge.* So named because it is the bridge between the PCI bus (66/33MHz) and the even slower ISA bus (8MHz).
- *The Super I/O chip.* It's a separate chip attached to the ISA bus that is not really considered part of the chipset and often comes from a third party, such as National Semiconductor or Standard MicroSystems Corp. (SMSC). The Super I/O chip contains commonly used peripheral items all combined into a single chip. Note that most recent South Bridge chips now include Super I/O functions, so that most recent motherboards no longer include a separate Super I/O chip.

▶▶ See "Super I/O Chips," p. 304.

Chipsets have evolved over the years to support various processors, bus speeds, peripheral connections, and features.

Figure 4.27 shows a typical AMD Socket A motherboard using North/South Bridge architecture with the locations of all chips and components.

The North Bridge is sometimes referred to as the PAC (PCI/AGP Controller). It is essentially the main component of the motherboard and is the only motherboard circuit besides the processor that normally runs at full motherboard (processor bus) speed. Most modern chipsets use a single-chip North Bridge; however, some of the older ones actually consisted of up to three individual chips to make up the complete North Bridge circuit.

The South Bridge is the lower-speed component in the chipset and has always been a single individual chip. The South Bridge is a somewhat interchangeable component in that different chipsets (North Bridge chips) often are designed to use the same South Bridge component. This modular design of the chipset allows for lower cost and greater flexibility for motherboard manufacturers. The South Bridge connects to the 33MHz PCI bus and contains the interface or bridge to the 8MHz ISA bus. It also typically contains dual IDE hard disk controller interfaces, one or two USB interfaces, and in later designs even the CMOS RAM and real-time clock functions. The South Bridge contains all the components that make up the ISA bus, including the interrupt and DMA controllers.

The third motherboard component, the Super I/O chip, is connected to the 8MHz ISA bus and contains all the standard peripherals that are built in to a motherboard. For example, most Super I/O chips contain the serial ports, parallel port, floppy controller, and keyboard/mouse interface. Optionally, they might contain the CMOS RAM/Clock, IDE controllers, and game port interface as well. Systems that integrate IEEE-1394 and SCSI ports use separate chips for these port types.

Most recent motherboards that use North/South Bridge chipset designs incorporate a Super-South Bridge, which incorporates the South Bridge and Super I/O functions into a single chip.

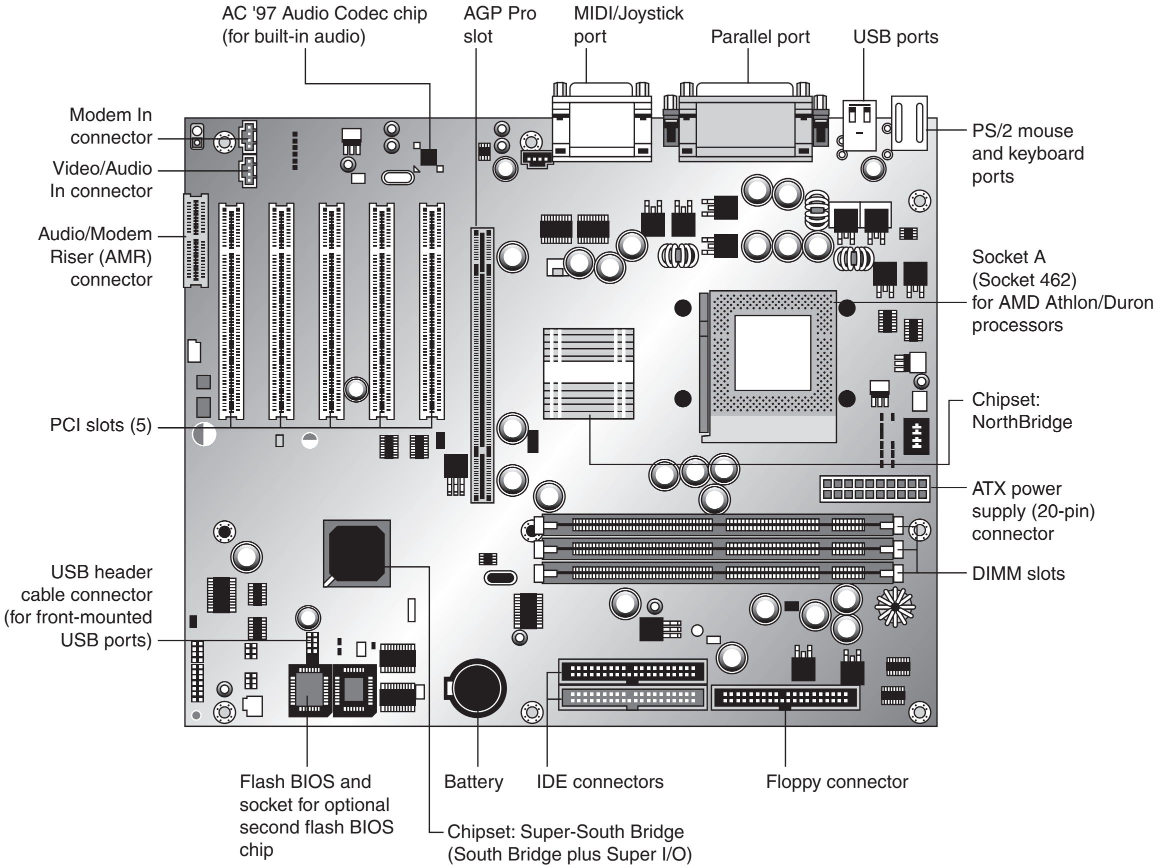

Figure 4.27 A typical Socket A (AMD Athlon/Duron) motherboard showing component locations.

Hub Architecture

The newer 800 series chips from Intel use a hub architecture in which the former North Bridge chip is now called a Memory Controller Hub (MCH) and the former South Bridge is called an I/O Controller Hub (ICH). Rather than connect them through the PCI bus as in a standard North/South Bridge design, they are connected via a dedicated hub interface that is twice as fast as PCI. The hub design offers several advantages over the conventional North/South Bridge design:

- *It's faster.* The hub interface is a 4X (quad-clocked) 66MHz 8-bit (4×66MHz×1 byte = 266MBps) interface, which has twice the throughput of PCI (33MHz×32 bits = 133MBps).
- *Reduced PCI loading.* The hub interface is independent of PCI and doesn't share or steal PCI bus bandwidth for chipset or Super I/O traffic. This improves performance of all other PCI bus connected devices because the PCI bus is not involved in these transactions.
- *Reduced board wiring.* Although twice as fast as PCI, the hub interface is only 8 bits wide and requires only 15 signals to be routed on the motherboard. By comparison, PCI requires no less than 64 signals be routed on the board, causing increased electromagnetic interference (EMI) generation, greater susceptibility to signal degradation and noise, and increased board manufacturing costs.

This hub interface design allows for a much greater throughput for PCI devices because there is no South Bridge chip (also carrying traffic from the Super I/O chip) hogging the PCI bus. Due to bypassing PCI, hub architecture also enables greater throughput for devices directly connected to the I/O Controller Hub (formerly the South Bridge), such as the new higher-speed ATA-100 and USB 2.0 interfaces.

The hub interface design is also very economical, being only 8 bits wide. Although this seems too narrow to be useful, there is a reason for the design. By making the interface only 8 bits wide, it uses only 15 signals, compared to the 64 signals required by the 32-bit-wide PCI bus interface used by North/South Bridge chip designs. The lower pin count means less circuit routing exists on the board, less signal noise and jitter occur, and the chips themselves have many fewer pins, making them smaller and more economical to produce.

Although it transfers only 8 bits at a time, the hub interface executes four transfers per cycle and cycles at 66MHz. This gives it an effective throughput of 4×66MHz×1 byte = 266MB per second (MBps). This is twice the bandwidth of PCI, which is 32 bits wide but runs only one transfer per 33MHz cycles for a total bandwidth of 133MBps. So, by virtue of a very narrow—but very fast—design, the hub interface achieves high performance with less cost and more signal integrity than with the previous North/South Bridge design.

The MCH interfaces between the high-speed processor bus (533/400/133/100/66MHz) and the hub interface (66MHz) and AGP bus (533/266/133/66MHz), whereas the ICH interfaces between the hub interface (66MHz) and the ATA (IDE) ports (66/100MHz) and PCI bus (33MHz).

The ICH also includes a new low-pin-count (LPC) bus, consisting basically of a stripped 4-bit wide version of PCI designed primarily to support the motherboard ROM BIOS and Super I/O chips. By using the same 4 signals for data, address, and command functions, only nine other signals are necessary to implement the bus, for a total of only 13 signals. This dramatically reduces the number of traces connecting the ROM BIOS chip and Super I/O chips in a system as compared to the 96 ISA bus signals necessary for older North/South Bridge chipsets that used ISA as the interface to those devices. The LPC bus has a maximum bandwidth of 6.67MBps, which is close to ISA and more than enough to support devices such as ROM BIOS and Super I/O chips.

Note

Although other chipset makers typically use the North Bridge/South Bridge nomenclature for their chipsets, several have developed high-speed connections similar to Intel's hub architecture. For example, most of VIA's recent chipsets use the V-Link Hub Architecture, which provides a dedicated 266MHz bus between the North and South Bridge chips. The high-speed HyperTransport bus between the North and South Bridges originally developed by AMD has been licensed by chipset vendors such as NVIDIA, VIA, and ALi Corporation, and SiS's MuTIOL Connect is used by recent SiS chipsets.

Figure 4.28 shows a typical Intel motherboard that uses bus architecture—the Intel D845PEBT2, which supports the Intel Pentium 4 processor. Unlike some of Intel's less-expensive hub-based motherboards, the 845PEBT2's Intel 845 chipset doesn't incorporate video.

Let's examine the popular chipsets, starting with those used in 486 motherboards and working all the way through to the latest Pentium III/Celeron, Pentium 4, Athlon XP, and Athlon 64 chipsets.

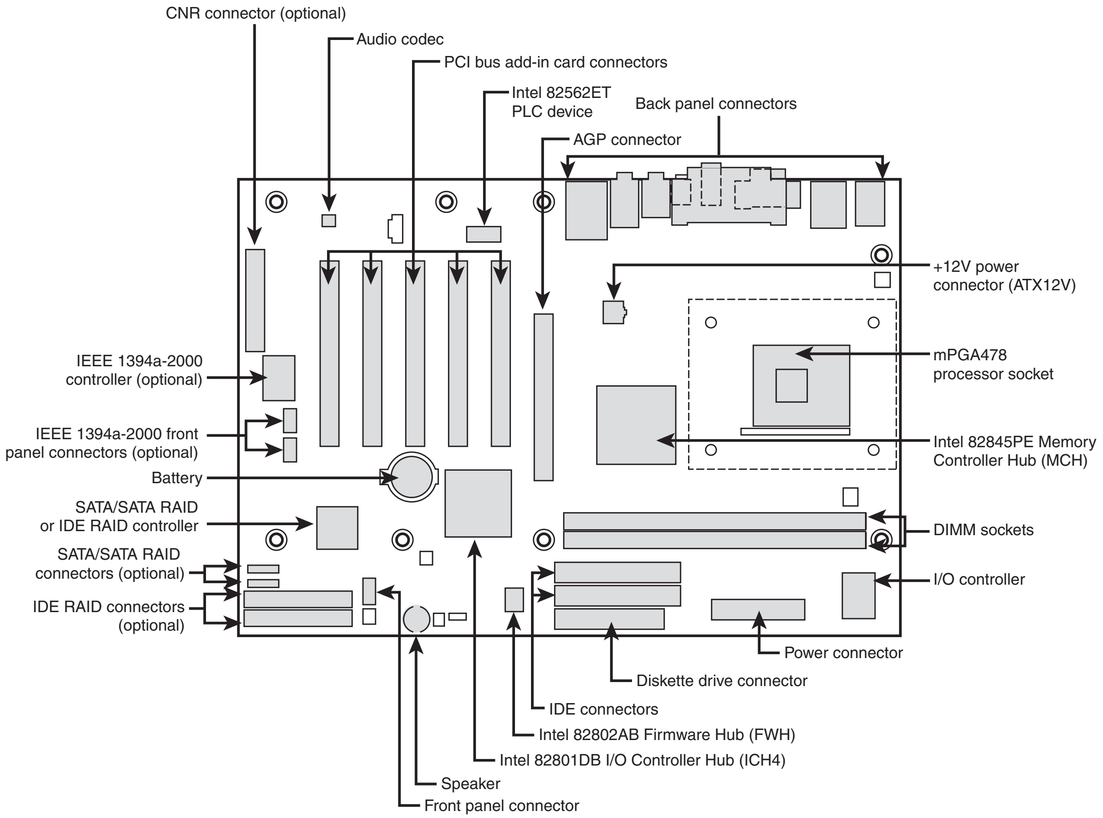

Figure 4.28 Intel D845PEBT2 motherboard showing component locations. *Illustration used by permission of Intel Corporation.*

Intel's Early 386/486 Chipsets

Intel's first real PC motherboard chipset was the 82350 chipset for the 386DX and 486 processors. This chipset was not very successful, mainly because the EISA bus was not very popular and many other manufacturers were making standard 386 and 486 motherboard chipsets at the time. The market changed very quickly, and Intel dropped the EISA bus support and introduced follow-up 486 chipsets that were much more successful.

Table 4.9 shows the Intel 486 chipsets.

Table 4.9 Intel 486 Motherboard Chipsets

Chipset	420TX	420EX	420ZX
Code name	Saturn	Aries	Saturn II
Date introduced	Nov. 1992	March 1994	March 1994
Processor	5V 486	5V/3.3V 486	5V/3.3V 486
Bus speed	Up to 33MHz	Up to 50MHz	Up to 33MHz
SMP (dual CPUs)	No	No	No
Memory types	FPM	FPM	FPM

Table 4.9 Continued

Chipset	420TX	420EX	420ZX
Parity/ECC	Parity	Parity	Parity
Max. memory	128MB	128MB	160MB
L2 cache type	Async	Async	Async
PCI support	2.0	2.0	2.1
AGP support	No	No	No

AGP = Accelerated graphics port
FPM = Fast page mode
PCI = Peripheral component interconnect
SMP = Symmetric multiprocessing (dual processors)
Note: PCI 2.1 supports concurrent PCI operations.

The 420 series chipsets were the first to introduce the North/South Bridge design that is still used in many chipsets today.

Fifth-Generation (P5 Pentium Class) Chipsets

With the advent of the Pentium processor in March 1993, Intel also introduced its first Pentium chipset: the 430LX chipset (code named Mercury). This was the first Pentium chipset on the market and set the stage as Intel took this lead and ran with it. Other manufacturers took months to a year or more to get their Pentium chipsets out the door. Since the debut of its Pentium chipsets, Intel has dominated the chipset market. Table 4.10 shows the Intel Pentium motherboard chipsets. Note that none of these chipsets support AGP; Intel first added support for AGP in its chipsets for the Pentium II/Celeron processors.

Table 4.10 Intel Pentium Motherboard Chipsets (North Bridge)

Chipset	430LX	430NX	430FX	430MX	430HX	430VX	430TX
Code name	Mercury	Neptune	Triton	Mobile Triton	Triton II	Triton III	n/a
Date introduced	March 1993	March 1994	Jan. 1995	Oct. 1995	Feb. 1996	Feb. 1996	Feb. 1997
CPU bus speed	66MHz	66MHz	66MHz	66MHz	66MHz	66MHz	66MHz
CPUs supported	P60/66	P75+	P75+	P75+	P75+	P75+	P75+
SMP (dual CPUs)	No	Yes	No	No	Yes	No	No
Memory types	FPM	FPM	FPM/EDO	FPM/EDO	FPM/EDO	FPM/EDO/SDRAM	FPM/EDO/SDRAM
Parity/ECC	Parity	Parity	Neither	Neither	Both	Neither	Neither
Max. memory	192MB	512MB	128MB	128MB	512MB	128MB	256MB
Max. cacheable	192MB	512MB	64MB	64MB	512MB	64MB	64MB
L2 cache type	Async	Async	Async/Pburst	Async/Pburst	Async/Pburst	Async/Pburst	Async/Pburst
PCI support	2.0	2.0	2.0	2.0	2.1	2.1	2.1
AGP support	No	No	No	No	No	No	No
South Bridge	SIO	SIO	PIIX	MPIIX	PIIX3	PIIX3	PIIX4

EDO = Extended data out
FPM = Fast page mode
PIIX = PCI ISA IDE Xcelerator
SDRAM = Synchronous dynamic RAM
SIO = System I/O
SMP = Symmetric multiprocessing (dual processors)

Note

PCI 2.1 supports concurrent PCI operations, enabling multiple PCI cards to perform transactions at the same time for greater speed.

Table 4.11 shows the Intel South Bridge chips used with Intel chipsets for Pentium processors. South Bridge chips are the second part of the modern Intel motherboard chipsets.

Table 4.11 Intel South Bridge Chips

Chip Name	SIO	PIIX	PIIX3	PIIX4	PIIX4E	ICH0	ICH
Part number	82378IB/ZB	82371FB	82371SB	82371AB	82371EB	82801AB	82801AA
IDE support	None	BMIDE	BMIDE	UDMA-33	UDMA-33	UDMA-33	UDMA-66
USB support	None	None	Yes	Yes	Yes	Yes	Yes
CMOS/clock	No	No	No	Yes	Yes	Yes	Yes
Power management	SMM	SMM	SMM	SMM	SMM/ACPI	SMM/ACPI	SMM/ACPI

ACPI = Advanced configuration and power interface
BMIDE = Bus master IDE (ATA)
ICH = I/O Controller Hub
IDE = Integrated Drive Electronics (AT attachment)
PIIX = PCI ISA IDE Xcelerator
SIO = System I/O
SMM = System management mode
UDMA = Ultra-DMA IDE (ATA)
USB = Universal serial bus

The Pentium chipsets listed in Tables 4.10 and 4.11 have been out of production for several years, and most computers that use these chipsets have been retired.

Intel 430LX (Mercury)

The 430LX was introduced in March 1993, concurrent with the introduction of the first Pentium processors. This chipset was used only with the original Pentiums, which came in 60MHz and 66MHz versions. These were 5V chips and were used on motherboards with Socket 4 processor sockets.

◀◀ See "Processor Socket and Slot Types," p. 85.

◀◀ See "P1 (086) First-Generation Processors," p. 114.

The 430LX chipset consisted of three total chips for the North Bridge portion. The main chip was the 82434LX system controller. This chip contained the processor-to-memory interface, cache controller, and PCI bus controller. There was also a pair of PCI bus interface accelerator chips, which were identical 82433LX chips.

The 430LX chipset was noted for the following:

- Single processor
- Support for up to 512KB of L2 cache
- Support for up to 192MB of standard DRAM

This chipset died off along with the 5V 60/66MHz Pentium processors.

Intel 430NX (Neptune)

Introduced in March 1994, the 430NX was the first chipset designed to run the new 3.3V second-generation Pentium processor. These were noted by having Socket 5 processor sockets and an onboard 3.3V/3.5V voltage regulator for both the processor and chipset. This chipset was designed primarily for Pentiums with speeds from 75MHz to 133MHz, although it was used mostly with 75MHz–100MHz systems. Along with the lower voltage processor, this chipset ran faster, cooler, and more reliably than the first-generation Pentium processor and the corresponding 5V chipsets.

◀◀ See "CPU Operating Voltages," p. 99.

◀◀ See "P2 (286) Second-Generation Processors," p. 115.

The 430NX chipset consisted of three chips for the North Bridge component. The primary chip was the 82434NX, which included the cache and main memory (DRAM) controller and the control interface to the PCI bus. The actual PCI data was managed by a pair of 82433NX chips called local bus accelerators. Together, these two chips, plus the main 82434NX chip, constituted the North Bridge.

The South Bridge used with the 430NX chipset was the 82378ZB System I/O (SIO) chip. This component connected to the PCI bus and generated the lower-speed ISA bus.

The 430NX chipset introduced the following improvements over the Mercury (430LX) chipset:

- Dual processor support
- Support for 512MB of system memory (up from 192MB for the LX Mercury chipset)

This chipset rapidly became the most popular chipset for the early 75MHz–100MHz systems, overshadowing the older 60MHz and 66MHz systems that used the 430LX chipset.

Intel 430FX (Triton)

The 430FX (Triton) chipset rapidly became popular after it was introduced in January 1995. This chipset is noted for being the first to support extended data out (EDO) memory, which subsequently became very popular. EDO was about 21% faster than the standard fast page mode (FPM) memory that had been used up until that time but cost no more than the slower FPM. Unfortunately, although it was known for faster memory support, the Triton chipset was also known as the first Pentium chipset without support for parity checking for memory. This was somewhat of a blow to PC reliability and fault tolerance, even though many did not know it at the time.

▶▶ See "Extended Data Out RAM," p. 432.

▶▶ See "Parity and ECC," p. 464.

The Triton chipset lacked parity support from the previous 430NX chipset, but it also supported only a single CPU. The 430FX was designed as a low-end chipset for home or non-mission-critical systems. As such, it did not replace the 430NX, which carried on in higher-end network fileservers and other more mission-critical systems.

The 430FX consisted of a three-chip North Bridge. The main chip was the 82437FX system controller that included the memory and cache controllers, CPU interface, and PCI bus controller, along with dual 82438FX data path chips for the PCI bus. The South Bridge was the first PIIX (PCI ISA IDE Xcelerator) chip that was a 82371FB. This chip not only acted as the bridge between the 33MHz PCI bus and the slower 8MHz ISA bus, but also incorporated for the first time a dual-channel IDE interface. By moving the IDE interface off the ISA bus and into the PIIX chip, it was now effectively connected to the PCI bus, enabling much faster Bus Master IDE transfers. This was key in supporting the ATA-2 or Enhanced IDE interface for better hard disk performance.

The major points on the 430FX are

- Support for EDO memory
- Support for higher speed—pipelined burst L2 cache
- PIIX South Bridge with high-speed Bus Master IDE
- Lack of support for parity-checked memory
- Only single CPU support
- Supported only 128MB of RAM, of which only 64MB could be cached

That last issue is one that many people are not aware of. The 430FX chipset can cache only up to 64MB of main memory. So, if you install more than 64MB of RAM in your system, performance suffers greatly. At the time, many didn't think this would be that much of a problem—after all, they didn't usually run enough software to load past the first 64MB anyway. That is another misunderstanding because Windows 9x and NT/2000 (as well as all other protected-mode operating systems including Linux and so on) load from the top down. So, for example, if you install 96MB of RAM (one 64MB and one 32MB bank), virtually all your software, including the main operating system, loads into the noncached region above 64MB. Needless to say, performance would suffer greatly. Try disabling the L2 cache via your CMOS Setup to see how slowly your system runs without it. That is the performance you can expect if you install more than 64MB of RAM in a 430FX-based system. Some thought this was a Windows limitation, but it is instead caused by the chipset design.

Intel 430HX (Triton II)

Intel created the Triton II 430HX chipset as a true replacement for the powerful 430NX chip. It added some of the high-speed memory features from the low-end 430FX, such as support for EDO memory and pipeline burst L2 cache. It also retained dual-processor support. In addition to supporting parity checking to detect memory errors, it added support for error correcting code (ECC) memory to detect and correct single bit errors on-the-fly. And the great thing was that this was implemented using plain parity memory.

The HX chipset's primary advantages over the FX are

- Symmetric multiprocessor (dual processor) support.
- Support for ECC and parity memory.
- 512MB maximum RAM support (versus 128MB).
- L2 cache functions over 512MB RAM versus 64MB (providing optional cache tag RAM is installed).
- Memory transfers in fewer cycles overall.
- PCI Level 2.1 compliance that allows concurrent PCI operations.
- PIIX3 supports different IDE/ATA transfer speed settings on a single channel.
- PIIX3 South Bridge component supports USB.

The memory problems with caching in the 430FX were corrected in the 430HX. This chipset allowed for the caching of the full 512MB of possible RAM as long as the correct amount of cache *tag* was installed. Tag is a small cache memory chip used to store the index to the data in the cache. Most 430HX systems shipped with a tag chip that could manage only 64MB of cached main memory, although you could optionally upgrade it to a larger capacity tag chip that would enable caching the full 512MB of RAM.

The 430HX chipset was a true one-chip North Bridge. It was also one of the first chips out in a ball-grid array (BGA) package, in which the chip leads are configured as balls on the bottom of the chip. This enabled a smaller chip package than the previous plastic quad flat pack (PQFP) packaging used on the older chips, and, because only one chip existed for the North Bridge, a very compact motherboard was possible. The South Bridge was the PIIX3 (82371SB) chip, which enabled independent timing of the dual IDE channels. Therefore, you could install two different speed devices on the same channel as master/slave and configure their transfer speeds independently. Previous PIIX chips allowed both devices on a single cable to work at the lowest common denominator speed supported by both. The PIIX3 also incorporated the USB for the first time on a PC motherboard. Unfortunately at the time, no devices were available to attach to USB, nor was there any operating systems or driver support for the bus. USB ports were a curiosity at the time, and nobody had a use for them.

▶▶ See "Universal Serial Bus," p. 947.

The 430HX supports the newer PCI 2.1 standard, which allowed for concurrent PCI operations and greater performance. Combined with the support for EDO and pipelined burst cache, this was perhaps the best Pentium chipset for the power user's system. It offered excellent performance, and with ECC memory it offered a truly reliable and stable system design.

The 430HX was the only modern Intel Pentium-class chipset to offer parity and error-corrected memory support. This made it the recommended Intel chipset at the time for mission-critical applications, such as fileservers, database servers, business systems, and so on.

Intel 430VX (Triton III)

The 430VX was designed to be a replacement for the low-end 430FX chipset; it was not a replacement for the higher-powered 430HX chipset. As such, the VX has only one significant technical advantage over the HX, but in almost all other respects it is more like the 430FX than the HX.

The VX has the following features:

- Supports 66MHz SDRAM
- No parity or ECC memory support
- Single processor only
- Supports only 128MB RAM
- Supports caching for only 64MB RAM

Most notable was the support for SDRAM, which was about 27% faster than the more popular EDO memory used at the time. Although the support for SDRAM was a nice bonus, the actual improvement in system speed derived from such memory was somewhat limited. This was because with a normal L1/L2 cache combination, the processor read from the caches 99% of the time. A combined miss (both L1 and L2 missing) happened only about 1% of the time while reading/writing memory. Thus with SDRAM, the system would be up to 27% faster, but only about 1% of the time. Therefore, the cache performance was actually far more important than main memory performance. See the section "How Cache Works" in Chapter 3, "Microprocessor Types and Specifications," for more information.

As with the 430FX, the VX has the limitation of being capable of caching only 64MB of main memory. Therefore, installing more than 64MB of memory actually slows down the system dramatically because none of the memory past that point can be cached. Because Windows loads from the top of memory down, installing any amount of memory greater than 64MB in a system using this chipset dramatically decreases performance.

The 430VX chipset was rapidly made obsolete in the market by the 430TX chipset that followed.

Intel 430TX

The 430TX was Intel's final Pentium chipset. It was designed not only to be used in desktop systems, but to replace the 430MX mobile Pentium chipset for laptop and notebook systems.

The 430TX had some refinements over the 430VX, but, unfortunately, it still lacked support for parity or ECC memory and retained the 64MB cacheable RAM limitation of the older FX and VX chipsets. The 430TX was not designed to replace the more powerful 430HX chipset, which still remained the chipset of choice for Pentium class mission-critical systems.

The TX chipset features include the following:

- 66MHz SDRAM support
- Cacheable memory still limited to 64MB
- Support for Ultra-ATA, or Ultra-DMA 33 (UDMA) IDE transfers
- Lower power consumption for mobile use
- No parity or ECC memory support
- Single processor only

▶▶ See "ATA/ATAPI-4 (AT Attachment with Packet Interface Extension-4)," p. 505.

Third-Party (Non-Intel) P5 Pentium Class Chipsets

The development of non-Intel Pentium-class chipsets was spurred by AMD's development of its own equivalents to the Pentium processor—the K5 and K6 processor families. Although the K5 was not a successful processor, the K6 family was very successful in the low-cost (under $1,000) market and as an upgrade for Pentium systems. AMD's own chipsets aren't used as often as other third-party chipsets, but AMD's capability to support its own processors with timely chipset deliveries has helped make the K6 and its successors—the Athlon, Athlon XP, and Duron—into credible rivals for Intel's Pentium MMX and Pentium II/III/4/Celeron families and has spurred other vendors, such as VIA, Acer Laboratories, and SiS, to support AMD's processors. Major third-party chipsets for Pentium-class processors include

- AMD 640
- VIA Apollo VP1, VP2, VPX, VP3, MVP3, and MVP4
- ALi Aladdin 4, Aladdin 5, and Aladdin 7
- SiS SiS540, SiS530/5595, SiS5598, SiS5581, SiS5582, SiS5571, SiS5591, and SiS5592

Most computers that use these chipsets have been retired. For more detailed information about these chipsets, see *Upgrading and Repairing PCs, 14th Edition*, which is included on the DVD packaged with this book.

Sixth-Generation (P6 Pentium Pro/II/III Class) Chipsets

Just as Intel clearly dominated the Pentium chipset world, it is also the leading vendor for chipsets supporting its P6 processor families. As discussed earlier, the biggest reason for this is that, since the Pentium first came out in 1993, Intel has been introducing new chipsets (and even complete ready-to-go motherboards) simultaneously with its new processors. This makes it hard for anybody else to catch up. Another problem for other chipset manufacturers is that they are required to license the CPU bus interface design before they can produce a matching chipset.

Note that because the Pentium Pro, Celeron, and Pentium II/III are essentially the same processor with different cache designs and minor internal revisions, the same chipset can be used for Socket 8 (Pentium Pro), Socket 370 (Celeron/Pentium III), and Slot 1 (Celeron/Pentium II/III) designs. Of course, the newer P6-class chipsets are optimized for the Socket 370 architecture and nobody is making any new designs for Socket 8 or Slot 1.

Table 4.12 shows the chipsets used on Pentium Pro motherboards.

Table 4.12 Pentium Pro Motherboard Chipsets (North Bridge)

Chipset	450KX	450GX	440FX
Code name	Orion	Orion Server	Natoma
Workstation date introduced	Nov. 1995	Nov. 1995	May 1996
Bus speed	66MHz	66MHz	66MHz
SMP (dual CPUs)	Yes	Yes (four CPUs)	Yes
Memory types	FPM	FPM	FPM/EDO/BEDO
Parity/ECC	Both	Both	Both
Maximum memory	8GB	1GB	1GB
L2 cache type	In CPU	In CPU	In CPU
Maximum cacheable	1GB	1GB	1GB
PCI support	2.0	2.0	2.1
AGP support	No	No	No
AGP speed	n/a	n/a	n/a
South Bridge	Various	Various	PIIX3

AGP = Accelerated graphics port
BEDO = Burst EDO
EDO = Extended data out
FPM = Fast page mode
Pburst = Pipeline burst (synchronous)
PCI = Peripheral component interconnect
PIIX = PCI ISA IDE Xcelerator
SDRAM = Synchronous dynamic RAM
SIO = System I/O
SMP = Symmetric multiprocessing (dual processors)

Note

PCI 2.1 supports concurrent PCI operations.

For the Celeron and Pentium II/III motherboards, Intel offers the chipsets in Table 4.13. 4xx series chipsets incorporate a North/South Bridge architecture, whereas 8xx series chipsets support the newer and faster hub architecture. P6/P7 (Pentium III/Celeron, Pentium 4, and Xeon) processor chipsets using hub architecture are shown in Table 4.14.

Table 4.13 P6 Processor Chipsets Using North/South Bridge Architecture

Chipset	440FX	440LX	440EX
Code name	Natoma	None	None
Date introduced	May 1996	Aug. 1997	April 1998
Part numbers	82441FX, 82442FX	82443LX	82443EX
Bus speed	66MHz	66MHz	66MHz
Supported processors	Pentium II	Pentium II	Celeron
SMP (dual CPUs)	Yes	Yes	No
Memory types	FPM/EDO/BEDO	FPM/EDO/SDRAM	EDO/SDRAM
Parity/ECC	Both	Both	Neither
Maximum memory	1GB	1GB EDO/512MB SDRAM	256MB
Memory banks	4	4	2
PCI support	2.1	2.1	2.1
AGP support	No	AGP 2x	AGP 2x
South Bridge	82371SB (PIIX3)	82371AB (PIIX4)	82371EB (PIIX4E)

Table 4.14 P6 (Pentium III/Celeron) Processor Chipsets Using Hub Architecture

Chipset	810	810E	815	815E	815EP
Code name	Whitney	Whitney	Solano	Solano	Solano
Date introduced	April 1999	Sept. 1999	June 2000	June 2000	Nov. 2000
Part number	82810	82810E	82815	82815	82815EP
Bus speed	66/100MHz	66/100/ 133MHz	66/100/ 133MHz	66/100/ 133MHz	66/100/ 133MHz
Supported processors	Celeron, Pentium II/III	Celeron, Pentium II/III	Celeron, Pentium II/III	Celeron, Pentium II/III	Celeron, Pentium II/III
SMP (dual CPUs)	No	No	No	No	No
Memory types	SDRAM (PC100), EDO	SDRAM (PC100)	SDRAM (PC133)	SDRAM (PC133)	SDRAM (PC133)
Parity/ECC	Neither	Neither	Neither	Neither	Neither
Maximum memory	512MB	512MB	512MB	512MB	512MB
Memory banks	2	2	3	3	3
PCI support	2.2	2.2	2.2	2.2	2.2
PCI speed/ width	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit
AGP slot	No	No	AGP 4x	AGP 4x	AGP 4x
Integrated video	AGP 2x[2]	AGP 2x[2]	AGP 2x[3]	AGP 2x[3]	No
South Bridge (Hub)	82801AA/ AB (ICH/ICH0)	82801AA (ICH)	82801AA (ICH)	82801BA (ICH2)	82801BA (ICH2)

1. The 440ZX is available in a cheaper 440ZX-66 version that runs only at 66MHz.

2. The 810 chipsets have integral AGP 2x 3D video that is NOT upgradable via an external AGP adapter.

3. The 815/815E chipsets have integral AGP 2x 3D video that IS upgradable via an APG 4x slot.

4. The only difference between the 815 and 815E is in which I/O controller hub (South Bridge) is used.

AGP = Accelerated graphics port

440BX	440GX	450NX	440ZX
None	None	None	None
April 1998	June 1998	June 1998	Nov. 1998
82443BX	82443GX	82451NX, 82452NX, 82453NX, 82454NX	82443ZX
66/100MHz	100MHz	100MHz	66/100MHz[1]
Pentium II/III, Celeron	Pentium II/III, Xeon	Pentium II/III, Xeon	Celeron, Pentium II/III
Yes	Yes	Yes, up to four	No
SDRAM	SDRAM	FPM/EDO	SDRAM
Both	Both	Both	Neither
1GB	2GB	8GB	256MB
4	4	4	2
2.1	2.1	2.1	2.1
AGP 2x	AGP 2x	No	AGP 2x
82371EB (PIIX4E)	82371EB (PIIX4E)	82371EB (PIIX4E)	82371EB (PIIX4E)

820	820E	840	815P	815EG	815G
Camino	Camino	Carmel	Solano	Solano	Solano
Nov. 1999	June 2000	Oct. 1999	March 2001	Sept. 2001	Sept. 2001
82820	82820	82840	82815EP	82815G	82815G
66/100/133MHz	66/100/133MHz	66/100/133MHz	66/100/133MHz	66/100/133MHz	66/100/133MHz
Pentium II/III, Celeron	Pentium II/III, Celeron	Pentium III, Xeon	Celeron, Pentium III	Celeron, Pentium III	Celeron, Pentium III
Yes	Yes	Yes	No	No	No
RDRAM (PC800)	RDRAM (PC800)	RDRAM (PC800) Dual-Channel	SDRAM (PC100/133)	SDRAM (PC66/100/133)	SDRAM (PC66/100/133)
Both	Both	Both	Neither	Neither	Neither
1GB	1GB	4GB	512MB	512MB	512MB
2	2	3x2	3	3	3
2.2	2.2	2.2	2.2	2.2	2.2
33MHz/32-bit	33MHz/32-bit	66MHz/64-bit	33MHz/32-bit	33MHz/32-bit	33MHz/32-bit
AGP 4x	AGP 4x	AGP 4x	AGP 4x	No	No
No	No	No	No	AGP 2x[3]	AGP 2x[3]
82801AA (ICH)	82801BA (ICH2)	82801AA (ICH)	82801AA/AB (ICH/ICH0)	82801BA (ICH2)	82801AA/AB (ICH/ICH0)

BEDO = Burst EDO
EDO = Extended data out
FPM = Fast page mode
ICH = I/O controller hub
Pburst = Pipeline burst (synchronous)
PCI = Peripheral component interconnect
PIIX = PCI ISA IDE Xcelerator
SDRAM = Synchronous dynamic RAM
SIO = System I/O
SMP = Symmetric multiprocessing (dual processors)

Note

Pentium Pro, Celeron, and Pentium II/III CPUs have their secondary caches integrated into the CPU package. Therefore, cache characteristics for these machines are not dependent on the chipset but are quite dependent on the processor instead.

Most Intel chipsets are designed as a two-part system, using a North Bridge and a South Bridge component. Often the same South Bridge component can be used with several different North Bridge chipsets. Table 4.15 shows a list of all the Intel South Bridge components used with P6-class processors and their capabilities. The ICH2 is also used as part of some of the first seventh-generation (Pentium 4/Celeron 4) Intel chipsets.

Table 4.15 Intel South Bridge—I/O Controller Hub Chips for P6

Chip Name	SIO	PIIX	PIIX3	PIIX4	PIIX4E	ICH0	ICH	ICH2
Part number	82378IB/ZB	82371FB	82371SB	82371AB	82371EB	82801AB	82801AA	82801BA
IDE support	None	BMIDE	BMIDE	UDMA-33	UDMA-33	UDMA-33	UDMA-66	UDMA-100
USB support	None	None	1C/2P	1C/2P	1C/2P	1C/2P	1C/2P	2C/4P
CMOS/clock	No	No	No	Yes	Yes	Yes	Yes	Yes
ISA support	Yes	Yes	Yes	Yes	Yes	No	No	No
LPC support	No	No	No	No	No	Yes	Yes	Yes
Power management	SMM	SMM	SMM	SMM	SMM/ACPI	SMM/ACPI	SMM/ACPI	SMM/ACPI

SIO = System I/O
PIIX = PCI ISA IDE (ATA) Xcelerator
ICH = I/O controller hub
USB = Universal serial bus
1C/2P = 1 controller, 2 ports
2C/4P = 2 controllers, 4 ports
IDE = Integrated Drive Electronics (ATA = AT attachment)
BMIDE = Bus master IDE (ATA)
UDMA = Ultra-DMA IDE (ATA)
ISA = Industry standard architecture bus
LPC = Low pin count bus
SMM = System management mode
ACPI = Advanced configuration and power interface

The following sections examine the chipsets for P6 processors up through the Celeron and Pentium III.

Intel 450KX/GX (Orion Workstation/Server)

The first chipsets to support the Pentium Pro were the 450KX and GX, both code named Orion. The 450KX was designed for networked or standalone workstations; the more powerful 450GX was designed for fileservers. The GX server chipset was particularly suited to the server role because it supports up to four Pentium Pro processors for symmetric multiprocessing (SMP) servers, up to 8GB of four-way interleaved memory with ECC or parity, and two bridged PCI buses. The 450KX is the workstation or standalone user version of Orion and as such it supports fewer processors (one or two) and less memory (1GB) than the GX. The 450GX and 450KX both have full support for ECC memory—a requirement for server and workstation use.

The 450GX and 450KX North Bridge comprises four individual chip components—an 82454KX/GX PCI bridge, an 82452KX/GX data path (DP), an 82453KX/GX data controller (DC), and an 82451KX/GX memory interface controller (MIC). Options for QFP or BGA packaging were available on the PCI Bridge and the DP. BGA uses less space on a board.

The 450's high reliability is obtained through ECC from the Pentium Pro processor data bus to memory. Reliability is also enhanced by parity protection on the processor bus, control bus, and all PCI signals. In addition, single-bit error correction is provided, thereby avoiding server downtime because of spurious memory errors caused by cosmic rays.

Until the introduction of the following 440FX chipset, these were used almost exclusively in fileservers. After the debut of the 440FX, the expensive Orion chips all but disappeared due to their complexity and high cost.

Intel 440FX (Natoma)

The first popular mainstream P6 (Pentium Pro or Pentium II) motherboard chipset was the 440FX, which was code named Natoma. Intel designed the 440FX to be a lower-cost and somewhat higher-performance replacement for the 450KX workstation chipset. It offered better memory performance through support of EDO memory, which the prior 450KX lacked.

The 440FX uses half the number of components that the previous Intel chipset used. It offers additional features, such as support for the PCI 2.1 (concurrent PCI) standard, support for USB, and reliability through ECC.

The concurrent PCI processing architecture maximizes system performance with simultaneous activity on the CPU, PCI, and ISA buses. Concurrent PCI provides increased bandwidth to better support 2D/3D graphics, video and audio, and processing for host-based applications. ECC memory support delivers improved reliability to business system users.

The main features of this chipset include

- Support for up to 1GB of EDO memory
- Full 1GB cacheability (based on the processor because the L2 cache and tag are in the CPU)
- Support for USB
- Support for Bus master IDE
- Support for full parity/ECC

The 440FX consists of a two-chip North Bridge. The main component is the 82441FX PCI Bridge and Memory controller, along with the 82442FX Data Bus accelerator for the PCI bus. This chipset uses the PIIX3 82371SB South Bridge chip that supports high-speed Bus Master DMA IDE interfaces and USB, and it acts as the bridge between the PCI and ISA buses.

Note that this was the first P6 chipset to support EDO memory, but it lacked support for the faster SDRAM. Also, the PIIX3 used with this chipset does not support the faster Ultra DMA IDE hard drives.

The 440FX was the chipset used on the first Pentium II motherboards, which have the same basic architecture as the Pentium Pro. The Pentium II was released several months before the chipset that was supposedly designed for it was ready, so early PII motherboards used the older 440FX chipset. This chipset was never designed with the Pentium II in mind, whereas the newer 440LX was optimized specifically to take advantage of the Pentium II architecture. For that reason, I normally recommended that people stay away from the original 440FX-based PII motherboards and wait for Pentium II systems that used the forthcoming 440LX chipset. When the new chipset was introduced, the 440FX was quickly superseded by the improved 440LX design.

Intel 440LX

The 440LX quickly took over in the marketplace after it debuted in August of 1997. This was the first chipset to really take full advantage of the Pentium II processor. Compared to the 440FX, the 440LX chipset offers several improvements:

- Support for the new AGP video card bus
- Support for 66MHz SDRAM memory
- Support for the Ultra DMA IDE interface
- Support for USB

The 440LX rapidly became the most popular chip for all new Pentium II systems from the end of 1997 through the beginning of 1998.

Intel 440EX

The 440EX was designed to be a low-cost, lower-performance alternative to the 440LX chipset. It was introduced in April 1998, along with the Intel Celeron processor. The 440EX lacks several features in the more powerful 440LX, including dual processor and ECC or parity memory support. This chipset is basically designed for low-end 66MHz bus-based systems that use the Celeron processor. Note that boards with the 440EX also fully support a Pentium II but lack some of the features of the more powerful 440LX or 440BX chipsets.

The main things to note about the 440EX are listed here:

- Designed with a feature set tuned for the low-end PC market
- Primarily for the Intel Celeron processor
- Supports AGP
- Does not support ECC or parity memory
- Single processor support only

The 440EX consists of an 82443EX PCI AGP Controller (PAC) North Bridge component and the new 82371EB (PIIX4E) South Bridge chip.

Note

The original 266MHz and 300MHz Celeron processors used with the 440EX chipset provided very low performance because these processors lacked any onboard Level 2 cache memory. Starting with the 300MHz Celeron 300A, Celeron added 128KB of Level 2 cache to its SEP packaging; all Socket 370 Celerons also include Level 2 cache. You should consider upgrading to a faster Celeron CPU with Level 2 cache if your 440EX-based system uses one of the original Celeron processors.

Intel 440BX

The Intel 440BX chipset was introduced in April 1998 and was the first chipset to run the processor host bus (often called the front-side bus, or FSB) at 100MHz. The 440BX was designed specifically to support the faster Pentium II/III processors at 350MHz and higher. A mobile version of this chipset is the first Pentium II/III chipset for notebook or laptop systems.

The main change from the previous 440LX to the BX is that the 440BX chipset improves performance by increasing the bandwidth of the system bus from 66MHz to 100MHz. Because the chipset can run at either 66MHz or 100MHz, it allows one basic motherboard design to support all Pentium II/III processor speeds based on either the 66MHz or 100MHz processor bus.

Here are the Intel 440BX highlights:

- Support for 100MHz SDRAM (PC100); the now-common PC133 RAM can also be installed, but it will still run at just 100MHz
- Support for both 100MHz and 66MHz system and memory bus designs
- Support for up to 1GB of memory in up to four banks (four DIMMs)
- Support for ECC memory
- Support for ACPI
- The first chipset to support the Mobile Intel Pentium II processor

The Intel 440BX consists of a single North Bridge chip called the 82443BX Host Bridge/Controller, which is paired with a new 82371EB PCI-ISA/IDE Xcelerator (PIIX4E) South Bridge chip. The new South Bridge adds support for the ACPI specification version 1.0. Figure 4.29 shows a typical system block diagram using the 440BX.

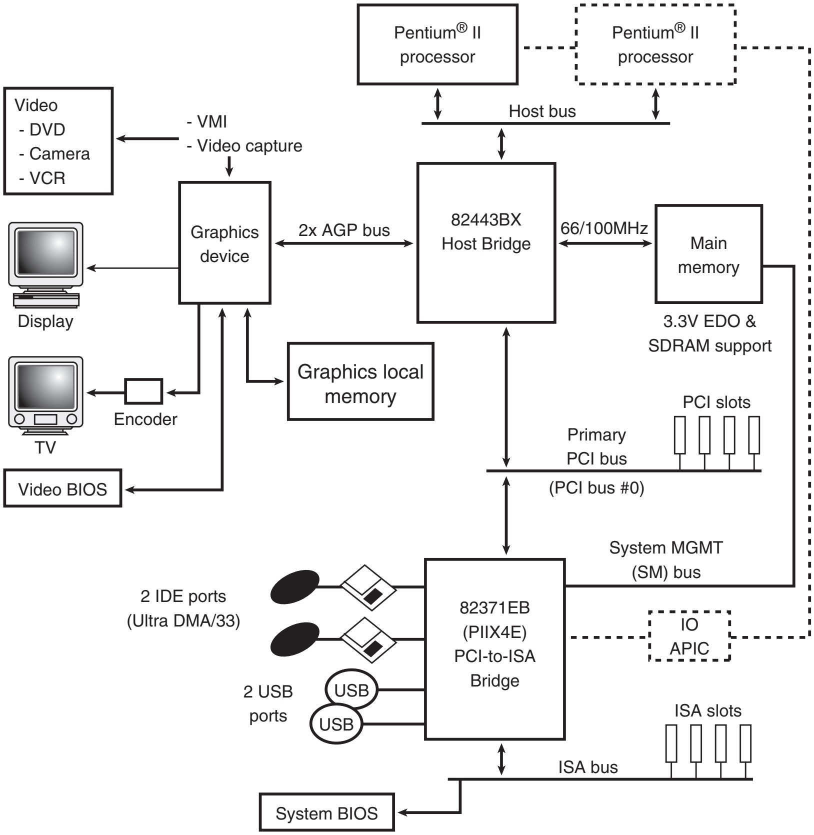

Figure 4.29 System block diagram using the Intel 440BX chipset.

The 440BX was a popular chipset during 1998 and into 1999. It offered superior performance and high reliability through the use of ECC, SDRAM, and DIMMs.

Intel 440ZX and 440ZX-66

The 440ZX was designed to be a low-cost version of the 440BX. The 440ZX brings 66MHz or 100MHz performance to entry-level Celerons (with or without Level 2 cache) and low-end Pentium II/III systems. The 440ZX is pin compatible with the more expensive 440BX, meaning existing 440BX motherboards can be easily redesigned to use this lower-cost chipset.

Note that two versions of the 440ZX are available: The standard one runs at 100MHz or 66MHz, and the 440ZX-66 runs only at the slower 66MHz.

The features of the 440ZX include the following:

- Support for Celeron and Pentium II/III processors at up to 100MHz bus speeds
- These main differences from the 440BX:
  - No parity or ECC memory support
  - Only two banks of memory (two DIMMs) supported
  - Maximum memory only 256MB

The 440ZX is not a replacement for the 440BX; instead, it was designed to be used in less expensive systems (such as those based on the micro-ATX form factor), in which the greater memory capabilities, performance, and data integrity functions (ECC memory) of the 440BX are unnecessary.

Intel 440GX

The Intel 440GX AGP set is the first chipset optimized for high-volume midrange workstations and lower-cost servers. The 440GX is essentially a version of the 440BX that has been upgraded to support the Slot 2 (also called SC330) processor slot for the Pentium II/III Xeon processor. The 440GX can still be used in Slot 1 designs, as well. It also supports up to 2GB of memory, twice that of the 440BX. Other than these items, the 440GX is essentially the same as the 440BX. Because the 440GX is core compatible with the 440BX, motherboard manufacturers could quickly and easily modify their existing Slot 1 440BX board designs into Slot 1 or 2 440GX designs.

The main features of the 440GX include the following:

- Support for Slot 1 and Slot 2
- Support for 100MHz system bus
- Support for up to 2GB of SDRAM memory

This chipset allows for lower-cost, high-performance workstations and servers using the Slot 2–based Xeon processors.

Intel 450NX

The 450NX chipset is designed for multiprocessor systems and standard high-volume servers based on the Pentium II/III Xeon processor. The Intel 450NX chipset consists of four components: the 82454NX PCI Expander Bridge (PXB), 82451NX Memory and I/O Bridge Controller (MIOC), 82452NX RAS/CAS Generator (RCG), and 82453NX Data Path Multiplexor (MUX).

The 450NX supports up to four Pentium II/III Xeon processors at 100MHz. Two dedicated PCI Expander Bridges can be connected via the Expander Bus. Each PXB provides two independent 32-bit, 33MHz PCI buses, with an option to link the two buses into a single 64-bit, 33MHz bus.

Figure 4.30 shows a typical high-end server block diagram using the 450NX chipset.

The 450NX supports one or two memory cards. Each card incorporates an RCG chip and two MUX chips, in addition to the memory DIMMs. Up to 8GB of memory is supported in total.

The primary features of the 450NX include the following:

- Slot 2 (SC330) processor bus interface at 100MHz
- Support for up to four-way processing
- Support for two dedicated PCI Expander Bridges
- Up to four 32-bit PCI buses or two 64-bit PCI buses

The 450NX chipset does not support AGP because high-end video is not an issue in network fileservers.

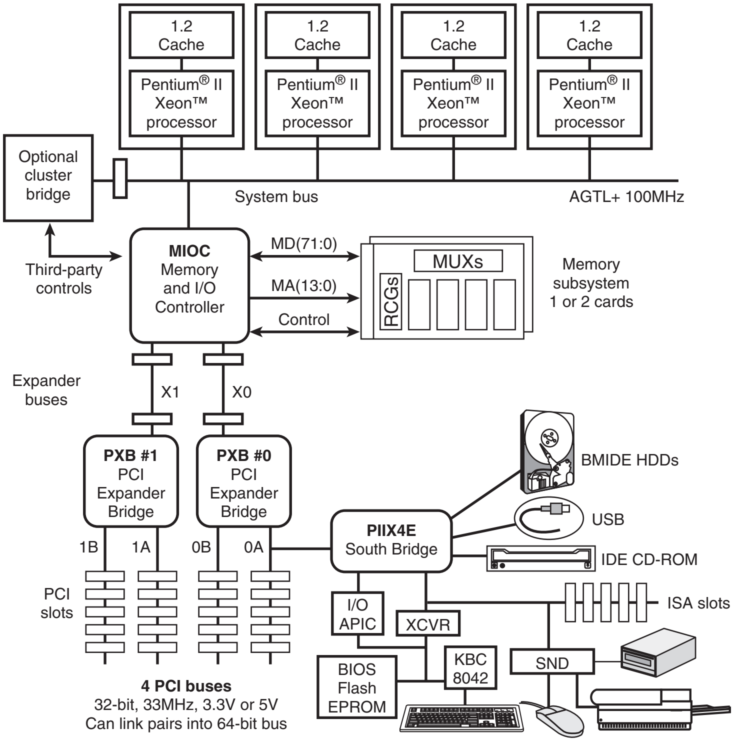

Figure 4.30 High-end server block diagram using the Intel 440NX chipset.

Intel 810, 810E, and 810E2

Introduced in April 1999, the Intel 810 chipset (code named Whitney) represents a major change in chipset design from the standard North and South Bridges that have been used since the 486 days. The 810 chipset allows for improvements in system performance, all for less cost and system complexity. The 810 (which supports 66MHz and 100MHz processor buses) was later revised as the 810E with support for the 133MHz processor bus.

Note

The 810E2 uses the same 82810E GMCH as the 810E but pairs it with the 82801BA I/O Controller Hub (ICH2) used by the Intel 815E. For information about the 82801BA ICH2 chip, see the section "Intel 815 Family," later in this chapter.

The major features of the 810E chipset include

- 66/100/133MHz system bus
- Integrated AGP 2x Intel 3D graphics
- Efficient use of system memory for graphics performance
- Optional 4MB of dedicated display cache video memory
- Digital Video Out port compatible with DVI specification for flat-panel displays
- Software MPEG-2 DVD playback with hardware motion compensation
- 266MBps hub interface

- Support for ATA-66
- Integrated Audio-Codec 97 (AC'97) controller
- Support for low-power sleep modes
- Random number generator (RNG)
- Integrated USB controller
- LPC bus for Super I/O and Firmware Hub (ROM BIOS) connection
- Elimination of ISA bus

The 810E chipset consists of three major components (see Figure 4.31):

- *82810E Graphics Memory Controller Hub (GMCH).* 421 BGA package (the original 810 chipset used the 82810 GMCH).
- *82801 Integrated Controller Hub (ICH).* 241 BGA package.
- *82802 Firmware Hub (FWH).* In either 32-pin plastic leaded chip carrier (PLCC) or 40-pin thin small outline package (TSOP) packages. Although a functional part of the chipset, this component is actually sold separately by Intel to motherboard developers.

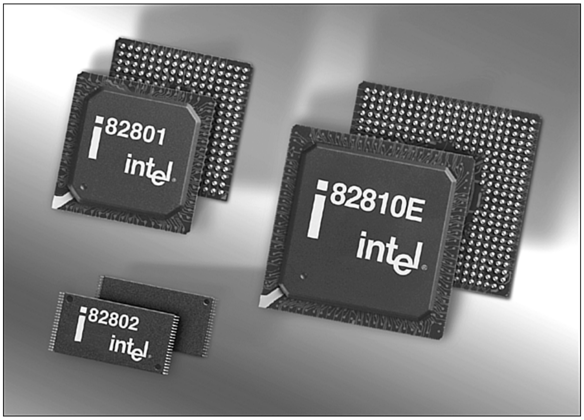

Figure 4.31 Intel 810E chipset showing the 82810E (GMCH), 82801 (ICH), and 82802 (FWH) chips. *Photograph used by permission of Intel Corporation.*

Compared to the previous North/South Bridge designs, there are some fairly significant changes in the 810 chipset. The previous system designs had the North Bridge acting as the memory controller, talking to the South Bridge chip via the PCI bus. This new design has the GMCH taking the place of the North Bridge, which talks to the ICH via a 66MHz dedicated interface called the accelerated hub architecture (AHA) bus instead of the previously used PCI bus. In particular, implementing a direct connection between the North and South Bridges in this manner was key in implementing the new UDMA-66 high-speed IDE interface for hard disks, DVD drives, and other IDE devices.

Figure 4.32 shows a system block diagram for the 810E chipset. With the 810 chipset family, ISA is finally dead.

The 82810E GMCH uses an internal Direct AGP (integrated AGP) interface to create 2D and 3D effects and images. The video capability integrated into the 82810E chip features hardware motion compensation to improve software DVD video quality; it also features both analog and direct digital video out ports, which enable connections to either traditional TVs (via an external converter module) or a direct digital flat panel display. The GMCH chip also incorporates the System Manageability Bus, which enables networking equipment to monitor the 810 chipset platform. Using ACPI specifications, the system manageability function enables low-power sleep mode and conserves energy when the system is idle.

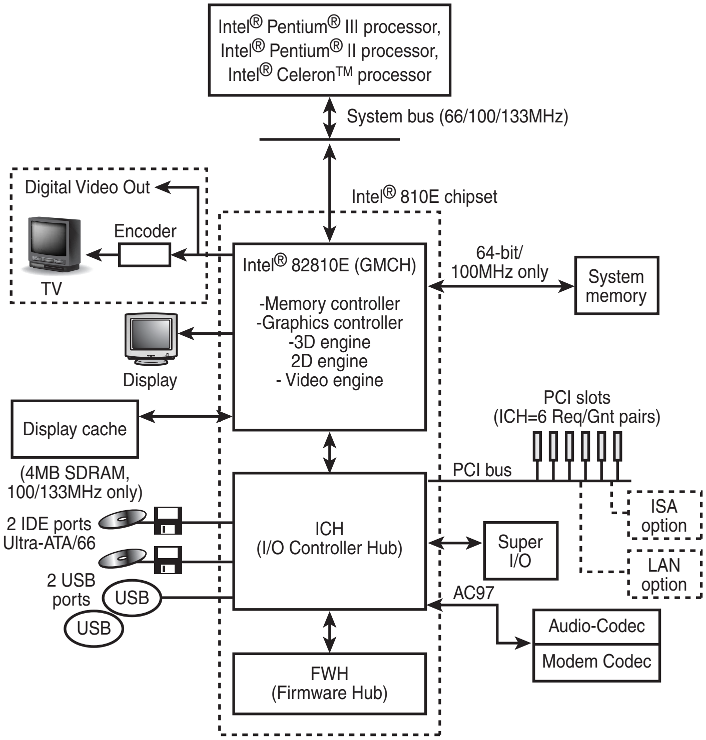

Figure 4.32 Intel 810E chipset system block diagram.

The 82801 I/O Controller Hub employs AHA for a direct connection from the GMCH chip. This is twice as fast (266MBps) as the previous North/South Bridge connections that used the PCI bus, and it uses far fewer pins for reduced electrical noise. Plus, the AHA bus is dedicated, meaning that no other devices will be on it. The AHA bus also incorporates optimized arbitration rules allowing more functions to run concurrently, enabling better video and audio performance.

The ICH also integrates dual IDE controllers, which run up to either 33MBps (UDMA-33 or Ultra-ATA/33) or 66MBps (UDMA-66 or Ultra-ATA/66). Note that two versions of the ICH chip exist. The 82801AA (ICH) incorporates the 66MBps-capable ATA/IDE and supports up to six PCI slots, whereas the 82801AB (ICH0) supports only 33MBps ATA/IDE maximum and supports up to four PCI slots.

The ICH also integrates an interface to an Audio-Codec 97 (AC97) controller, dual USB ports, and the PCI bus with up to four or six slots. The Integrated Audio-Codec 97 controller enables software audio and modem by using the processor to run sound and modem software via very simple digital-to-analog conversion circuits. Reusing existing system resources lowers the system cost by eliminating components.

The 82802 Firmware Hub (FWH) incorporates the system BIOS and video BIOS, eliminating a redundant nonvolatile memory component. The BIOS within the FWH is flash-type memory, so it can be field-updated at any time. In addition, the 82802 contains a hardware RNG. The RNG provides truly random numbers to enable fundamental security building blocks supporting stronger encryption, digital signing, and security protocols. Two versions of the FWH are available, called the 82802AB and 82802ACy. The AB version incorporates 512KB (4Mb) of Flash BIOS memory, and the AC version incorporates a full 1MB (8Mb) of BIOS ROM.

With the Intel 810 and 810E chipsets, Intel did something that many in the industry were afraid of: It integrated the video and graphics controller directly into the motherboard chipset with no means of upgrade. This means systems using the 810 chipset don't have an AGP slot and aren't capable of using conventional AGP video cards. For the low-end market for which this chipset is designed, lacking an AGP slot shouldn't be too much of a drawback. Higher-end systems, on the other hand, use the 815 or other chipsets that do support AGP slots. Intel calls the integrated interface Direct AGP, and it describes the direct connection between the memory and processor controllers with the video controller all within the same chip.

This means the video card as we know it, will be reserved only for midrange and higher-end systems, as well as gaming-oriented systems. With the 810 as well as subsequent chipsets with integrated video, Intel has let it be known in a big way that it has entered the PC video business.

In fact, the theme with the 810 chipset is one of integration. The integrated video means no video cards are required; the integrated AC97 interface means that conventional modems and sound cards are not required. Plus, there is an integrated CMOS/Clock chip (in the ICH), and even the BIOS is integrated in the FWH chip. All in all, the 810 should be taken as a sign for things to come in the PC industry, which means more integration, better overall performance for low-end and mainstream systems, and less overall cost.

Intel Random Number Generator

The 8xx chipset series features the Intel Random Number Generator (RNG). The RNG is built in to the 82802 FWH, which is the ROM BIOS component used on 8xx-based motherboards. The RNG provides software with true nondeterministic random numbers.

Most security routines, especially those providing authentication or encryption services, require random numbers for purposes such as key code generation. One method of cracking these types of codes is to predict the random numbers being used to generate the keys. Current methods that use system and user input as a seed to a conventional pseudorandom number generator have proven vulnerable to this type of attack. The Intel RNG uses thermal noise across a resistor contained in the FWH (that is, ROM BIOS in 8xx-based boards) to generate true nondeterministic, unpredictable random numbers. Therefore, "random" numbers generated by 8xx-series chipsets really are random.

Intel 815 Family

Introduced in June 2000, the 815 and 815E chipsets are mainstream PC chipsets with integral video that is also upgradable via an AGP 4x slot. An 815EP version introduced a few months later lacks the integrated video for lower cost. In March 2001, the 815P chipset, an improved version of the 815EP, was introduced. In September 2001, the last members of the family—the 815G and 815EG—were introduced. Note that the *G* indicates that these chipsets also include integrated video.

The 815 chipsets are designed for Slot-1 or Socket-370 processors, such as the Celeron or Pentium III. These are the first chipsets from Intel designed to directly support PC133 SDRAM memory, allowing for a more affordable solution than other chipsets using RDRAM memory. Similar to the other 8xx series chipsets from Intel, the 815 uses hub architecture that provides a 266MBps connection between the main chipset components and does not share the PCI bus like the prior North/South Bridge designs.

Although six variations on the 815 chipset are available, only five different parts are used to create the various members of the family: one memory controller hub (82815EP MCH: North Bridge replacement without integrated graphics), two graphics memory controller hubs (82815 or 82815G GMCH: North Bridge replacement with integrated graphics), and two I/O controller hubs (ICH and ICH2). Table 4.16 shows how these parts are combined to create the various members of the family.

Table 4.16 815 Chipset Family Components

Chipset Name	82815 GMCH	82815G GMCH	82815EP MCH	82801AA ICH	82801BA ICH2
815	*			*	
815E	*				*
815EP			*		*
815P			*	*	
815G		*		*	
815EG		*			*

Figure 4.33 illustrates one member of this chipset family, the 815E.

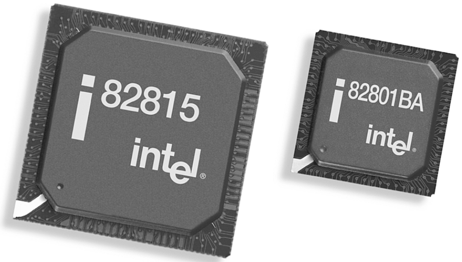

Figure 4.33 Intel 815E chipset showing the 82815 (GMCH) and 82801BA (ICH2) chips. *Photograph used by permission of Intel Corporation.*

All 815 chipsets support the following features:

- 66/100/133MHz system bus
- 266MBps hub interface
- ATA-100 (815E/EP/EG) or ATA-66 (815/P/G)
- PC100 or PC133 CL-2 SDRAM (also PC66 with 815G/GE)
- Up to 512MB RAM
- Integrated Audio-Codec 97 (AC97) controller
- Low-power sleep modes
- RNG for stronger security products
- One (815/P/G) or two (815E/EP/EG) integrated USB controllers with either two or four ports, respectively
- LPC bus for Super I/O and Firmware Hub (ROM BIOS) connection
- Elimination of ISA Bus

The 815/E/G/EG also support the following:

- Integrated Intel AGP 2x 3D graphics
- Efficient use of system memory for graphics performance
- Optional 4MB of dedicated display cache video memory
- Digital Video Out port compatible with DVI specification for flat panel displays
- Software MPEG-2 DVD playback with hardware motion compensation

The 815E/EP/EG uses the ICH2, which is most notable for providing ATA-100 support, allowing 100MBps drive performance. Of course, few drives can really take advantage of this much throughput, but in any case, there won't be a bottleneck there. The other notable feature is having two USB 1.1 controllers and four ports on board. This allows double the USB performance by splitting up devices over the two ports and can allow up to four connections before a hub is required.

Integrated Ethernet

Another important feature of the 815 series is the integration of a fast Ethernet controller directly into the chipset. The integrated LAN controller works with one of three new physical layer components from Intel and enables three distinct solutions for computer manufacturers. These include

- Enhanced 10/100Mbps Ethernet with Alert on LAN technology
- Basic 10/100Mbps Ethernet
- 1Mbps home networking

These physical layer components can be placed directly on the PC motherboard (additional chips) or installed via an adapter that plugs into the CNR slot. The CNR slot and cards enable PC assemblers to build network-ready systems for several markets.

AGP Inline Memory Module

Although the 815/815E feature is essentially the same built-in AGP 2x 3D video that comes with the 810 chipset, the difference is upgradeability. The video can easily be upgraded by adding a graphics performance accelerator (GPA) card (see Figure 4.34) or an AGP 4x card for maximum 3D graphics and video performance. The GPA card (also called the AGP Inline Memory Module, or AIMM) is essentially a high-performance video memory card that works in the AGP 4x slot and improves the performance of the integrated video by up to 30%. Unfortunately, these are not commonly sold and are somewhat expensive. For even more performance, you can install a full 4x AGP card in the AGP 4x slot, which disables the integrated video. By having the video integrated, very low-cost systems with reasonable video performance can be assembled. By later installing either the GPA or a full 4x AGP card, you can improve video performance up to 100% or more.

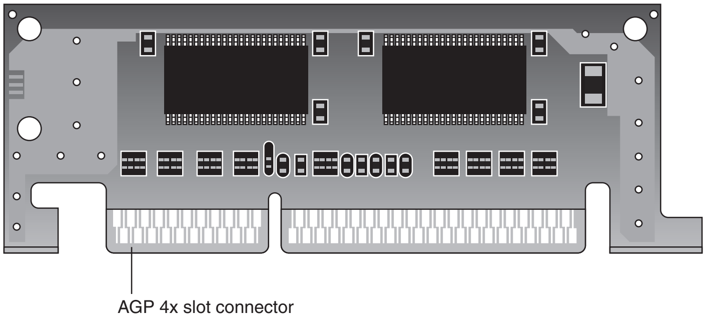

Figure 4.34 A typical 4MB GPA/AIMM module, which attaches to the AGP slot of a motherboard using the 815 or 815E chipset.

PC133 Memory Support

Another important feature of the 815 chipset is the support of PC133 memory. The 815 family also uses PC100 memory. With PC133 support, Intel has also officially set a standard for PC133 memory that was higher than some of the PC133 memory on the market at the time of introduction. To meet the Intel PC133 specification, the memory must support what is called 2-2-2 timing, sometimes also known as CAS-2 (column address strobe) or CL-2 timing. The numbers refer to the number of clock cycles for the following functions to complete:

- *Precharge command to Active command.* Charges the memory's storage capacitors to prepare them for data
- *Active command to Read command.* Selects rows and columns in memory array for reading
- *Read command to Data Out.* Reads data from selected rows and columns for transmission

Some of the PC133 memory on the market takes three cycles for each of these functions and would therefore be termed PC133 3-3-3, CAS-3, or CL-3 memory. Note that the faster PC133 CL-2 can be used in place of the slower CL-3 variety, but not the other way around.

As a result of the tighter cycling timing, PC133 CL-2 offers a lead-off latency of only 30ns, instead of the 45ns required by PC133 CL-3. This results in a 34% improvement in initial access due to the decreased latency.

The 815 chipset was a popular chipset for the mainstream PC market that didn't want to pay the higher prices for RDRAM memory. The 815 was essentially designed to replace the venerable 440BX chipset.

Intel 820 and 820E

The Intel 820 chipsets use the hub-based architecture like all the 800 series chips and are designed to support slot 1 or socket 370 processors, such as the Pentium III and Celeron (see Figure 4.35). The 820 chipset supports RDRAM memory technology, 133MHz system bus, and 4x AGP.

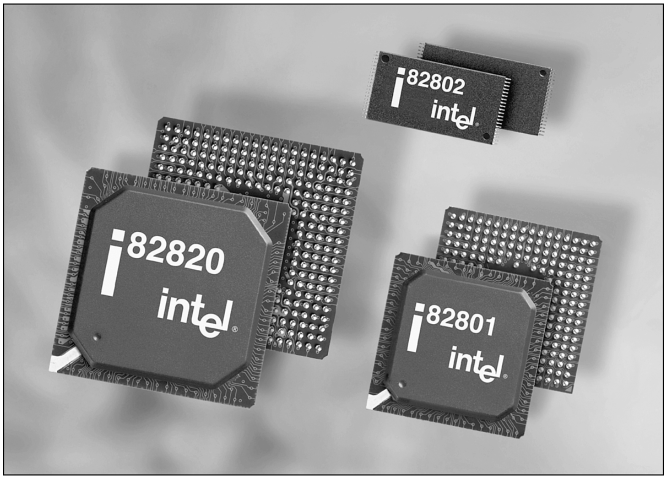

Figure 4.35 Intel 820 chipset showing the 82820 (MCH), 82801 (ICH), and 82802 (FWH) chips. *Photograph used by permission of Intel Corporation.*

The 82820 MCH provides the processor, memory, and AGP interfaces. Two versions are available: One supports a single processor (82820), whereas the other supports two processors (82820DP). Either is designed to work with the same 82801 ICH as used with the other 800 series chipsets, such as the 810 and 840. The 820 chipset also uses the 82802 FWH for BIOS storage and for the Intel RNG.

The connection between the MCH and ICH uses what is called the Intel Hub Architecture bus instead of the PCI bus as with prior North/South Bridge chipsets. The hub architecture bus provides twice the bandwidth of PCI at 266MB per second, enabling twice as much data to flow between them. The hub architecture bus also has optimized arbitration rules, allowing more functions to run concurrently, as well as far fewer signal pins, reducing the likelihood of encountering or generating noise and signal errors.

The 820 chipset is designed to use RDRAM memory, which has a maximum throughput of up to 1.6GBps. The 820 supports PC600, PC700, and PC800 RDRAM, delivering up to 1.6GBps of theoretical memory bandwidth in the PC800 version. PC800 RDRAM is a 400MHz bus running double-clocked and transferring 16 bits (2 bytes) at a time (2×400MHz×2 bytes = 1.6GBps). Two RIMM sockets are available to support up to 1GB of total system memory.

The AGP interface in the 820 enables graphics controllers to access main memory at AGP 4x speed, which is about 1GB per second—twice that of previous AGP 2x platforms. Figure 4.36 shows the 820 chipset architecture. Because the 820 was designed for mid-range to higher-end systems, it does not include integrated graphics, relying instead on the AGP 4x slot to contain a graphics card.

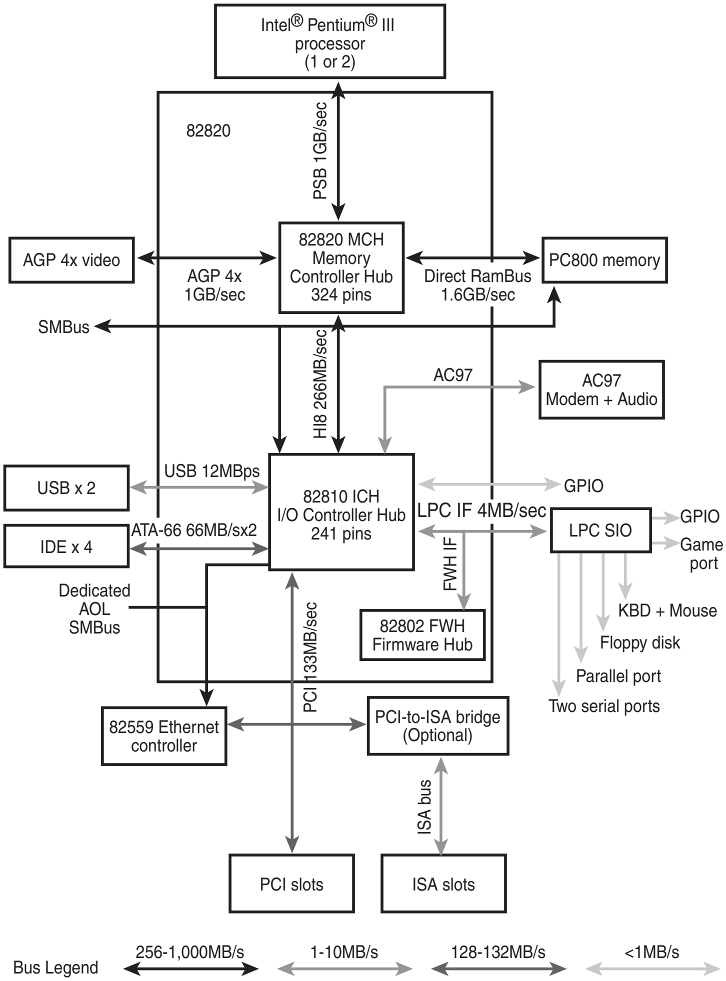

Figure 4.36 Intel 820 chipset architecture.

820 Chipset features include

- 100/133MHz processor bus
- Intel 266MBps hub interface
- PC800 RDRAM RIMM memory support
- AGP 4x support
- ATA-100 (820E) or ATA-66 interface
- Intel RNG
- LPC interface
- AC97 controller
- One (820) or two (820E) USB buses with either two or four ports, respectively

The 820 chipset consists of three main components with a few optional extras. The main component is the 82820 (single-processor) or 82820DP (dual-processor) MCH, which is a 324 BGA chip. That is paired with an 82801 ICH, which is a 241 BGA chip, and finally it has the 82802 FWH, which is really just a fancy Flash ROM BIOS chip. Optionally, there can be an 82380AB PCI-ISA bridge that is used only if the board is equipped with ISA slots.

The newer 820E version uses an updated 82801BA ICH2, which supports ATA-100 and incorporates dual USB controllers with two ports each, for a total of four USB ports.

820 Chipset MTH Bug

The 820 chipset is designed to support RDRAM memory directly. However, because the market still demanded lower-cost SDRAM, Intel created an RDRAM-to-SDRAM translator chip called the Memory Translator Hub (MTH). This enabled them to produce 820 chipset motherboards that supported SDRAM instead of the more expensive RDRAM.

Because the design of the MTH was proven defective, the chip (and any board using it) was simply discontinued. On May 10, 2000, Intel officially announced that it would replace any motherboards using the MTH with a new board lacking the component. The MTH translates signals from SDRAM memory to the Intel 820 chipset and is used only with motherboards utilizing SDRAM and the Intel 820 chipset; boards using RDRAM don't have an MTH and were not affected. Intel found electrical noise issues with the MTH that can cause some systems to intermittently reset, reboot, or hang. In addition, the noise issue can, under extreme conditions, potentially cause data corruption.

The MTH bug forced Intel to recall and replace more than a million motherboards in mid-2000, with new versions lacking the MTH and thus supporting only RDRAM memory. The final bill for this recall was reported at about $253 million, making it perhaps the most costly recall of computer components since the infamous Pentium math bug in 1994. I found it interesting that, due to the fact that Intel did more than $24.4 billion in sales the previous year, at least one article classified the cost of this recall as "chump change" to the chip giant!

Intel has an MTH I.D. Utility at `www.intel.com/support/mth` that will tell you whether you have that component and whether your board is eligible for replacement, including a 128MB RDRAM RIMM. Again, note that the 820 chipset was really designed to support RDRAM as the native type of memory, and RDRAM-based systems are not affected because they don't use the memory translator hub component.

Intel 840

The Intel 840 is a high-end chipset designed for use in high-performance multiprocessor systems using slot 1, slot 2 (Xeon processor), or Socket 370 processors. The 840 chipset uses the same hub architecture and modular design as the rest of the 800 family chipsets, with some additional components enabling more performance. See Figure 4.37 for a photo of the Intel 840 chipset.

As with the other 800 series chipsets, the 840 has three main components:

- *82840 Memory Controller Hub.* Provides graphics support for AGP 2x/4x, dual RDRAM memory channels, and multiple PCI bus segments for high-performance I/O.

- *82801 I/O Controller Hub.* Equivalent to the South Bridge in older chipset designs, except it connects directly to the MCH component via the high-speed Intel Hub Architecture bus. The ICH supports 32-bit PCI, IDE controllers, and dual USB ports.
- *82802 Firmware Hub.* Basically an enhanced Flash ROM chip that stores system BIOS and video BIOS, as well as an Intel RNG. The RNG provides truly random numbers to enable stronger encryption, digital signing, and security protocols.

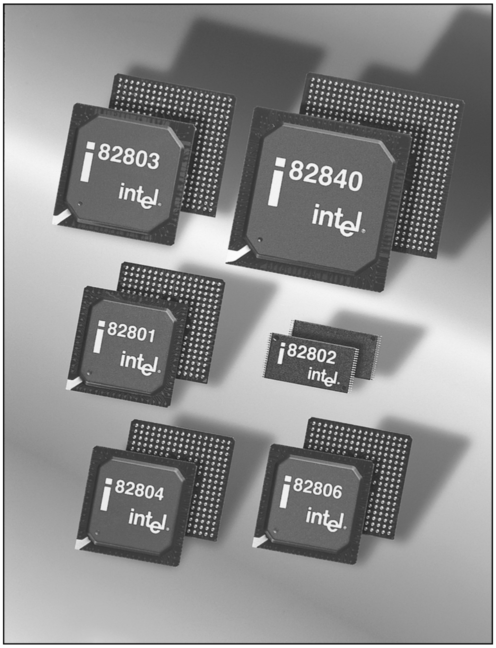

Figure 4.37 Intel 840 chipset showing the 82840 (MCH), 82801 (ICH), 82802 (FWH), 82803 (MRH-R), 82804 (MRH-S), and 82806 (P64H) chips. *Photograph used by permission of Intel Corporation.*

In addition to the core components, parts are available for scaling up to a more powerful design. Three additional components can be added:

- *82806 64-bit PCI Controller Hub (P64H).* Supports 64-bit PCI slots at speeds of either 33MHz or 66MHz. The P64H connects directly to the MCH using Intel Hub Architecture, providing a dedicated path for high-performance I/O. This is the first implementation of the 66MHz 66-bit PCI on a PC motherboard chipset, allowing for a PCI bus four times faster than the standard 32-bit 33MHz version.
- *82803 RDRAM-based Memory Repeater Hub (MRH-R).* Converts each memory channel into two memory channels for expanded memory capacity.
- *82804 SDRAM-based Memory Repeater Hub (MRH-S).* Translates the RDRAM protocol into SDRAM-based signals for system memory flexibility. This would be used only in 840 systems that supported SDRAM.

Figure 4.38 shows the 840 chipset architecture.

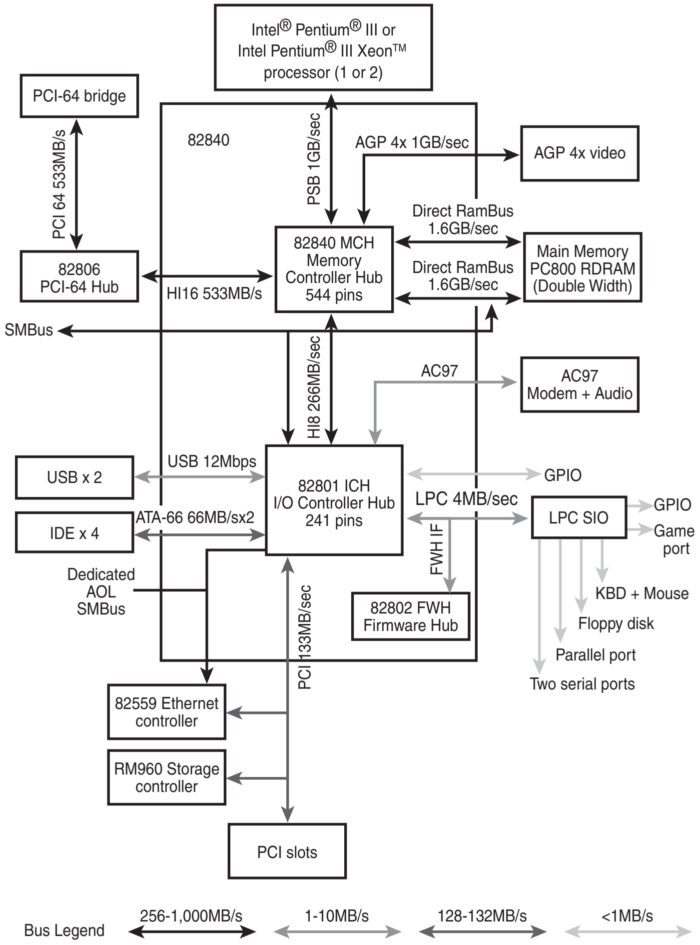

Figure 4.38 Intel 840 chipset architecture.

840 chipset features include

- 100/133MHz processor bus
- Dual RDRAM memory channels, operating simultaneously and providing up to 3.2GBps memory bandwidth
- 16-bit wide implementation of Intel Hub Architecture (HI16), which enables high-performance concurrent PCI I/O with the optional P64H component
- AGP 4x
- Prefetch cache, unique to the 840 chipset, which enables highly efficient data flow and helps maximize system concurrency
- Intel RNG (see the section "Intel Random Number Generator," earlier in this chapter)
- USB support

Optionally, network interface and RAID controller interface chips can be added as well.

Third-Party (Non-Intel) P6-Class Chipsets

Several companies produce chipsets designed to support P6-class processors, including ALi Corporation (formerly known as Acer Laboratories), VIA Technologies, and SiS. The following sections discuss the offerings from these companies.

ALi Chipsets for P6-Class Processors

ALi has a variety of chipsets for the P6-class processors. Table 4.17 provides an overview of these chipsets.

Table 4.17 ALi Chipsets for Pentium Pro-II-III-Celeron

Chipset	Aladdin Pro II	Aladdin Pro 4	Aladdin TNT2	Aladdin Pro 5
Date introduced	1999	2000	1999	2000, 2001(T)
Part number	M1621	M1641/M1641B	M1631	M1651, M1651T
Bus speed	60, 66, 100MHz	100, 133, 200, 266MHZ (B)	66, 100, 133MHz	66, 100, 133, 200, 266MHz
Supported processors	Pentium II, Pentium Pro	Pentium II, III, Celeron	Pentium II, III, Celeron	Pentium II, III, Celeron (T version supports Tualatin)
Form factor	Slot 1, Socket 370	Slot 1, Socket 370	Slot 1, Socket 370	Slot 1, Socket 370
SMP (dual CPUs)	Yes	No	No	No
Memory types	FPM, EDO, PC100	PC100, PC133, DDR200, DDR266 (B)	PC66, 100, 133, EDO	PC66, PC100, PC133, DDR200, DDR266
Parity/ECC	ECC	ECC	ECC	Neither
Maximum memory	1GB (SDRAM), 2GB (EDO)	1.5GB	1.5GB	3GB
PCI support	2.2	2.2	2.2	2.2
PCI speed/ width	33MHz/32-bit	33MHz/32-bit	33MHz/32-bit	33MHz/32-bit
AGP slot	1x/2x	1x/2x/4x	No	1x/2x/4x
Integrated video	No	No	Yes—TnT2	No
South Bridge	M1533 or M1543	M1535D	M1543C	M1535D

Table 4.18 provides an overview of the features of the South Bridge chips used in these chipsets.

Table 4.18 ALi South Bridge Chips Used with P6-Class Chipsets

South Bridge Chip	Number of USB Ports	ATA Support	Integrated Sound	Integrated Super I/O
M1533	2	ATA-33	No	No
M1543	2	ATA-33	No	Yes
M1535D	4	ATA-66	Yes[1]	Yes
M1535D+	6[2]	ATA-100	Yes[3]	Yes
M1543C	3	ATA-66	No	Yes

1. SoundBlaster 16 compatible with wavetable

2. Supports Legacy USB (mouse/keyboard)

3. 3D PCI audio with Direct3D (DirectX) support, MIDI, SPDIF, SoundBlaster compatibility

For more information about these chipsets, see *Upgrading and Repairing PCs, 14th Edition*, found in electronic form on the DVD packaged with this book.

VIA Technologies Chipsets for P6-Class Processors

VIA Technologies has a variety of chipsets for the P6 processors. They are discussed in the following sections and Table 4.19.

Apollo Pro

The Apollo Pro is a high-performance chipset for Slot 1 mobile and desktop PC systems; it can also support the Socket 8 Pentium Pro processor. The Apollo Pro includes support for advanced system power management capability for both desktop and mobile PC applications, PC100 SDRAM, AGP 2x mode, and multiple CPU/DRAM timing configurations. The Apollo Pro chipset is comparable in features to the 440BX and PIIX4e chipsets from Intel and represents one of the first non-Intel chipsets to support the Socket 1 architecture.

The VIA Apollo Pro consists of two devices—the VT82C691 North Bridge chip and the VT82C596, a BGA-packaged South Bridge with a full set of mobile, power-management features. For cost-effective desktop designs, the VT82C691 can also be configured with the VT82C586B South Bridge.

The VT82C691 Apollo Pro North Bridge supports all Slot 1 (Intel Pentium II) and Socket 8 (Intel Pentium Pro) processors. The Apollo Pro also supports the 66MHz external bus speed and the newer 100MHz CPU external bus speed required by the 350MHz and faster Pentium II processors. AGP v1.0 and PCI 2.1 are also supported, as are FPM, EDO, and SDRAM. Various DRAM types can be used in mixed combinations in up to eight banks and up to 1GB of DRAM. EDO memory timing is 5-2-2-2-2-2-2-2 for back-to-back accesses, and SDRAM timing is 6-1-1-1-2-1-1-1 for back-to-back accesses.

The VT82C596 South Bridge supports both ACPI and APM and includes an integrated USB controller and dual Ultra DMA-66 EIDE ports.

Apollo Pro Plus

The VIA Apollo Pro Plus is a high-performance chipset for Slot 1/Socket 370 mobile and desktop PC systems.

Features include the following:

- 66/100MHz CPU bus
- AGP 1x
- PCI 2.1 compliant
- Support for eight banks up to 1GB PC-100 SDRAM
- Support for both ACPI and legacy (APM) power management
- Integrated USB controller
- Ultra DMA-33 controller

The Apollo Pro consists of two chips: a VT82C693 North Bridge paired with either the VT82C596A mobile South Bridge or the VT82C686A Super South Bridge.

Table 4.19 VIA Technologies Chipsets for Pentium Pro-II-III-Celeron

Chipset	Apollo Pro	Apollo Pro Plus	Apollo KLE133 (PM601)	ProSavage PM133	Apollo Pro133	Apollo Pro133A	Apollo Pro PL133T	Apollo Pro 266/266T
Part number	VT82C691	VT82C693	VT8601	VT8605	VT82C693A	VT82694X	VT8605	VT8653
Bus speed	66, 100MHz	66, 100MHz	66, 100, 133MHz	66, 100, 133MHz	66, 100, 133MHz	66, 100, 133MHz	66, 100, 133MHz	66, 100, 133MHz
Supported processors	Pentium Pro, Pentium II, Celeron	Pentium II, Celeron	Pentium II, III, Celeron, VIA C3	Pentium II, III, Celeron, VIA C3	Pentium II, III, Celeron, VIA C3	Pentium II, III, Celeron, VIA C3	Pentium II, III Celeron (Tualatin), VIA C3	Pentium III, Celeron (Tualatin), VIA C3
Form factor	Socket 8, Slot 1	Slot 1, Socket 370	Slot 1, Socket 370	Slot 1, Socket 370	Slot 1, Socket 370	Slot 1, Socket 370	Slot 1, Socket 370	Socket 370
SMP (dual CPUs)	No	No	No	No	No	Yes	No	No
Memory types	FP, EDO, PC66, 100 SDRAM	FP, EDO, PC66, 100 SDRAM	PC66, 100, 133 SDRAM	PC66, 100, 133 SDRAM	PC66, 100, 133 SDRAM	PC66, 100, 133 SDRAM, EDO	PC100, 133 SDRAM	PC100, 133 SDRAM, DDR200, 266
Parity/ECC	No	No	No	No	No	Yes	No	No
Maximum memory	1GB	1GB	1GB	1.5GB	1.5GB	4GB	1.5GB	4GB
PCI support	2.1	2.1	2.1	2.2	2.1	2.2	2.2	2.2
PCI speed/width	33MHz/32-bit	33MHz/32-bit	33MHz/32-bit	33MHz/32-bit	33MHz/32-bit	33MHz/32-bit	33MHz/32-bit	33MHz/32-bit
AGP slot	1x, 2x	1x, 2x	1x, 2x	2x, 4x	1x, 2x	2x, 4x	2x, 4x	2x, 4x
Integrated video	No	No	Yes[1]	Yes	No	No	Yes[2]	No
South Bridge	VT82C596 or VT82C586B	VT82C596A	VT82C686A	VT8231	VT82C596B or VT82C586A	VT82C686A	VT8231	VT8233C[3]

1. Trident Blade3D

2. S3 Savage 4 (3D) integrating Savage 2000 (2D)

3. Supports VIA 4x V-Link 266MHz high-speed interconnect between North Bridge and South Bridge

Table 4.20 provides an overview of the features of the South Bridge chips used in these chipsets.

Table 4.20 VIA South Bridge Chips Used with P6-Class Chipsets

South Bridge Chip	Number of USB Ports	ATA Support	Integrated Sound	Integrated Super I/O	Integrated 10/100 Ethernet	Supports V-Link
VT82C596	2	ATA-33	No	No	No	No
VT82C596A	2	ATA-33	No	No	No	No
VT82C686A	4	ATA-66	AC'97	Yes	No	No
VT82C586B	2	ATA-33	No	No	No	No
VT8231	4	ATA-100	AC'97	Yes	No	No
VT82C596B	4	ATA-66	AC'97	Yes	No	No
VT82C586A	No	ATA-33	No	No	No	No
VT8233(C)	6	ATA-100	AC'97	Yes	Yes*	Yes

**3Com 10/100 Ethernet on C version only*

Apollo Pro133

The VIA Apollo Pro133 was the first chipset on the market designed to support PC-133 SDRAM memory. This chipset supports Slot 1 and Socket 370 processors, such as the Intel Pentium III, Intel Celeron, and VIA Cyrix III processors.

Key features include the following:

- Support for AGP 2x graphics bus
- Support for 133/100/66MHz processor bus
- PC-133 SDRAM memory interface
- Support for 1.5GB of RAM
- ATA-66 interface
- Support for four USB ports
- AC97 link for audio and modem
- Hardware monitoring
- Power management

The VIA Apollo Pro133 is a two-chip set consisting of the VT82C693A North Bridge controller and a choice of VT82C596B or VT82C686A South Bridge controllers.

Apollo KLE133

The VIA Apollo KLE133 (previously known as the PM601) chipset is a highly integrated chipset platform designed for the value PC and Internet appliance market. As such, this chipset has a built-in Trident Blade3D graphics engine and 10/100 Ethernet. The KLE133 is designed for Slot 1 and Socket 370 processors, such as the Pentium III, Celeron, and VIA C3.

Key features include the following:

- Support for AGP 2x
- Integrated Trident Blade3D AGP graphics engine
- 66/100/133MHz processor bus
- Support for PC-133 SDRAM memory
- Integrated 10/100 Ethernet
- Support for AC97 audio, MC'97 modem, Super I/O, and hardware monitoring
- Four USB ports
- Support for ATA-66
- Advanced power management

The VIA Apollo KLE133 is a two-chip set consisting of the VT8601 North Bridge controller and the VT82C686A South Bridge controller.

Apollo Pro133A

The VIA Apollo Pro133A chipset is a North/South Bridge chipset designed to support Slot 1 and Socket 370 processors such as the Intel Pentium III, Intel Celeron, and VIA Cyrix III. The Apollo Pro133A is based on the previous Pro133 with even more high-end features added.

Features of the Apollo Pro133A include

- AGP 4x graphics bus support
- Support for one or two processors
- 133/100/66MHz processor bus support
- PC-133 SDRAM memory interface
- ATA-66 interface
- Support for four USB ports
- AC97 link for audio and modem
- Hardware monitoring
- Power management

The VIA Apollo Pro133A chipset is a two-chip set consisting of the VT82C694X North Bridge controller and a choice of a VT82C596B or VT82C686A South Bridge controller.

ProSavage PM133

The VIA ProSavage PM133 integrates S3 Graphics' S3 Savage 4 and S3 Savage 2000 3D and 2D graphics engines with the Apollo Pro 133A chipset. The major features of the chipset are the same as for the Apollo Pro 133A, with the following additions:

- 2MB–32MB shared memory architecture integrated with Savage 4 3D and Savage 2000 2D video
- Z-buffering, 32-bit true-color rendering, massive 2K-by-2K textures, single-pass multiple textures, sprite antialiasing, and other 3D features
- Support for DVD playback, DVI LCD displays, and TV-out
- PCI 2.2 compliance

An optional AGP 4x interface allows the integrated AGP 4x video to be upgraded with an add-on card if desired. This two-chip set consists of the VT8605 North Bridge and VT8231 South Bridge.

The VT8231 South Bridge integrates the Super I/O and supports the LPC interface.

Apollo Pro266

The VIA Apollo Pro266 is a high-performance North/South Bridge chipset designed to support Socket 370 processors, including the Pentium III, Celeron, and VIA's own C3. The Apollo Pro 266 is the first chipset from VIA to replace the traditional PCI (133MBps) connection between North and South Bridge chips with VIA's 4x V-Link interconnect, which runs at 266MBps.

Features of the Apollo Pro266 include

- AGP 2x/4x graphics bus support
- 133/100/66MHz processor bus support
- PC-100/133 SDRAM and PC200/266 DDR SDRAM memory interface
- ATA-100 IDE interface
- Support for six USB ports

- Integrated AC97 six-channel audio
- Integrated MC'97 modem
- Integrated 10/100BASE-T Ethernet and 1/10MHz Home PNA networking
- Hardware monitoring
- ACPI/On Now! Power management
- VIA 4x V-Link North/South Bridge interconnect'

The VIA Apollo Pro266 chipset is a two-chip set consisting of the 552-pin BGA VT8633 North Bridge controller and the 376-pin BGA VT8233 South Bridge controller. Figure 4.39 shows the architecture of the Apollo Pro266 chipset.

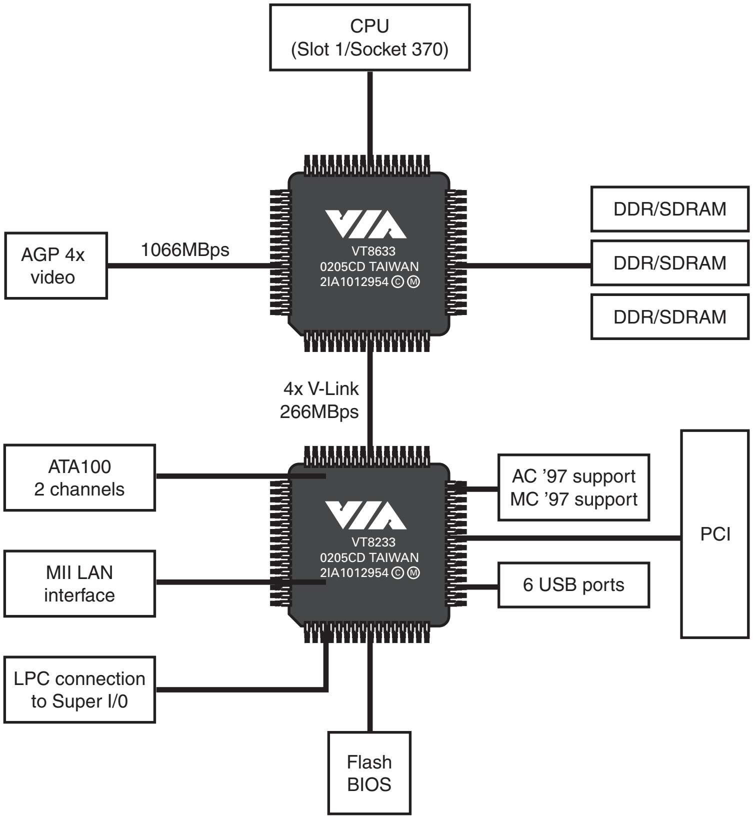

Figure 4.39 Apollo Pro266 chipset architecture.

Because of the V-Link high-speed interconnect between the North Bridge and South Bridge, PCI is managed by the South Bridge. This is similar to the way in which Intel hub architecture works, and this basic architecture has been followed by all subsequent VIA chipsets that use V-Link architecture.

Silicon Integrated Systems Chipsets for P6-Class Processors

Silicon Integrated Systems has a variety of chipsets for the P6-class processors. They are discussed in the following sections, and Table 4.21 provides a summary of them.

Table 4.21 Current SiS Chipsets for Pentium II/III/Celeron

Chipset	SiS620	SiS630	SiS630E	SiS630ET	SiS630S	SiS630ST
Bus speed	66, 100MHz	66, 100, 133MHz	66, 100, 133MHz	66, 100, 133MHz	66, 100, 133MHz	66, 100, 133MHz
Supported processors	Pentium II	Celeron, Pentium III	Celeron, Pentium III	Celeron, Pentium III, PIII Tualatin	Celeron, Pentium III	Celeron, Pentium III, PIII Tualatin
Form factor	Slot 1	Socket 370	Socket 370	Socket 370	Socket 370	Socket 370
SMP (dual CPUs)	No	No	No	No	No	No
Memory types	SDRAM PC66/100	SDRAM PC100/133	SDRAM PC100/133	SDRAM PC100/133	SDRAM PC100/133	SDRAM PC100/133
Parity/ECC	Neither	Neither	Neither	Neither	Neither	Neither
Maximum memory	1.5GB	3GB	3GB	3GB	3GB	3GB
PCI support	PCI 2.2	PCI 2.2	PCI 2.2	PCI 2.2	PCI 2.2	PCI 2.2
PCI speed/ width	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit
AGP slot	None	None	None	None	Yes	Yes
Integrated video	AGP 2.0	AGP 2.0	AGP 2.0	AGP 2.0	AGP 2.0	AGP 2.0
ATA support	ATA-33/66	ATA-33/66	ATA-33/66	ATA-33/ 66/100	ATA-33/ 66/100	ATA-33/ 66/100
USB support/ ports	USB 1.1/2 ports	USB 1.1/5 ports	USB 1.1/5 ports	USB 1.1/5 ports	USB 1.1/6 ports	USB 1.1/6 ports
10/100 Ethernet	No	Yes	Yes	Yes	Yes	Yes
Hardware audio	No	Yes	Yes	Yes	Yes	Yes
South Bridge chip	SiS 5595	No	No	No	No	No
SiS video bridge support	No	Yes	No	No	Yes	Yes

SiS630 Family

The SiS630 is a high-performance, low-cost, single-chip set with integrated 2D/3D graphics and support for processors such as the Pentium III, Celeron, and Cyrix III/VIA C3.

The integrated video is based on a 128-bit graphic display interface with AGP 4x performance. In addition to providing a standard analog interface for CRT monitors, the SiS630 also provides the digital flat panel port (DFP) for a digital flat panel monitor. An optional SiS301 video bridge supports NTSC/PAL TV output.

The SiS630 also includes integrated 10/100Mb Fast Ethernet, as well as an AC97-compliant interface that comprises a digital audio engine with 3D-hardware accelerator, on-chip sample rate converter, and professional wavetable along with a separate modem DMA controller. SiS630 also incorporates the LPC interface for attaching newer Super I/O chips and a dual USB host controller with four USB ports.

Features of the SiS630 include the following:

- Support for Intel/AMD/Cyrix/IDT Pentium CPU processor bus at 66/83/90/95/100/133MHz
- Integrated 2MB Level 2 cache controller
- PC133 SDRAM support
- Meets PC99 requirements
- PCI 2.2 compliant
- Ultra DMA66/33 support
- Integrated AGP 2x 2D/3D video/graphics accelerator
- Support for digital flat panel
- Hardware DVD decoding
- Built-in secondary CRT controller for independent secondary CRT, LCD, or TV digital output
- Low pin count interface
- Advanced PCI H/W audio and S/W modem
- Meets ACPI 1.0 requirements
- PCI Bus Power Management Interface Spec. 1.0
- Integrated keyboard/mouse controller
- Dual USB controller with five USB ports
- Integrated 10/100Mbps Ethernet controller

The SiS630E chip is almost identical in its feature set but doesn't support the secondary CRT controller (also called the video bridge). The SiS630S chip adds support for an AGP slot and the new Advanced Communication Riser to the basic SiS630 features. The SiS630ET and SiS630ST chips add support for the Tualatin 0.13 micron Pentium III and Celeron III processors.

What's really amazing about the SiS630 family is that all these features are combined into a single chip that's not much larger than the North Bridge used by other chipsets.

SiS633/635 Family

The SiS633/635 family is a series of high-performance single-chip chipsets that support the Pentium III and Celeron Socket 370 processors. T-series chipsets also support the newer .13-micron Pentium III Tualatin processors released in May 2001.

The SiS633 and SiS633T chipsets support PC133 SDRAM, whereas the SiS635 and SiS635T chipsets support PC133 SDRAM, DDR266 DDR SDRAM, or a mixture of PC133 and DDR266 SDRAM. All chipsets in the family support the following features:

- Integrated 4x AGP
- Up to six PCI masters
- UDMA/100 IDE host adapters
- 1.5GB RAM
- Six USB ports
- AC97 audio and AMR support

SiS635/635T also support ACR or CNR riser technology and integrate 10/100BASE-T Ethernet and 1/10MHz Home PNA networking.

These chips use a 677-pin BGA package. The 633/635 family has been discontinued, and it appears that very few motherboards were built using these chipsets.

SiS620/5595

The SiS620/5595 is a North/South Bridge chipset with integrated video that supports Slot 1 or Socket 370 processors.

Features include the following:

- 66/100MHz processor bus
- Support for PC-100 SDRAM memory
- Support for Ultra DMA 33/66
- PCI 2.2 compliant
- Integrated AGP 3D graphics accelerator
- Support for digital flat panel port for LCD panel

SiS600/5595 and 5600/5595

The SiS600/SiS5595 is a slot-1 North/South Bridge chipset designed for lower-cost systems.

Its features include

- 66/100MHz processor bus
- PC-100 SDRAM with ECC
- AGP 2x
- Ultra DMA/33
- Two USB ports
- Advanced Configuration and Power Interface revision 1.0

This chipset has been discontinued.

Table 4.22 Pentium 4 Chipsets from Intel Introduced 2000–2002

Chipset	850	850E	845	845E
Code name	Tehama	Tehama-E	Brookdale	Brookdale-E
Date introduced	Nov. 2000	May 2002	Sept. 2001 (SDRAM); Jan. 2002 (DDR)	May 2002
Part number	82850	82850E	82845	82845E
Bus speeds	400MHz	400/533MHz	400MHz	400/533MHz
Supported processors	Pentium 4, Celeron[1]	Pentium 4, Celeron[2]	Pentium 4, Celeron[2]	Pentium 4, Celeron[2,4]
SMP (dual CPUs)	No	No	No	No
Memory types	RDRAM (PC800) dual-channel	RDRAM (PC800, 1066 dual-channel)	PC133 SDRAM, DDR 200/ 266 SDRAM	DDR 200/ 266 SDRAM
Parity/ECC	Both	Both	ECC	ECC
Maximum memory	2GB	2GB (PC800); 1.5GB (PC1066)	2GB (PC2100 DDR); 3GB (PC133 SDRAM)	2GB
Memory banks	2	2	2 (PC2100); 3 (PC133)	2
PCI support	2.2	2.2	2.2	2.2
PCI speed/width	33MHz/32-bit	33MHz/32-bit	33MHz/32-bit	33MHz/32-bit
AGP slot	AGP 4x (1.5V)	AGP 4x (1.5V)	AGP 4x (1.5V)	AGP 4x (1.5V)
Integrated video	No	No	No	No
South Bridge (hub)	ICH2	ICH2	ICH2	ICH4

1. Supports Socket 423 and Socket 478 processors.

2. Supports Socket 478 processors only.

3. Stepping B-1 supports HT Technology (hyper-threading).

4. Supports HT Technology (hyper-threading).

Seventh-Generation (Pentium 4) Chipsets

Because of its long-time integration of processor and chipset design, it's not surprising that Intel dominates the Pentium 4 chipset market as it has the markets for Pentium (P5) and Pentium II/III/Celeron (P6) markets in the past. Although Intel has licensed Socket 423 (used by early Pentium 4 processors) and the current Socket 478 to rival chipset vendors such as SiS and ALi, Intel is still the leading developer of Pentium 4 chipsets. Third-party chipsets for Pentium 4 and Celeron 4 processors are discussed later in this chapter.

Because the Pentium 4 and Celeron processors using Socket 423 and those made for Socket 478 are essentially the same processors with different cache designs and minor internal revisions, the same chipset can be used for both processors.

Tables 4.22 and 4.23 show the chipsets made by Intel for Pentium 4 and Celeron 4 processors.

845GL	**845G**	**845GE**	**845GV**	**845PE**
Brookdale-GL	Brookdale-G	Brookdale-GE	Brookdale-GV	Brookdale-PE
July 2002	July 2002	Oct. 2002	Oct. 2002	Oct. 2002
82845GL	82845G	82845GE	82845GV	82845PE
400MHz	400/533MHz	400/533MHz	400/533MHz	400/533MHz
Pentium 4, Celeron[2]	Pentium 4, Celeron[2,3]	Pentium 4, Celeron[2,3]	Pentium 4, Celeron[2,4]	Pentium 4, Celeron[2,4]
No	No	No	No	No
PC133 SDRAM, DDR 200/ 266 SDRAM	PC133 SDRAM, DDR 200/ 266 SDRAM	DDR 333/ 266 SDRAM	DDR 200/ 266 SDRAM	DDR 333/ 266 SDRAM
Neither	ECC	Neither	Neither	Neither
2GB	2GB	2GB	2GB	2GB
2	2	2	2	2
2.2	2.2	2.2	2.2	2.2
33MHz/32-bit	33MHz/32-bit	33MHz/32-bit	33MHz/32-bit	33MHz/32-bit
None	AGP 4x (1.5V)	AGP 4x (1.5V)	None	AGP 4x (1.5V)
Intel Extreme Graphics 200MHz	Intel Extreme Graphics 200MHz	Intel Extreme Graphics 266MHz	Intel Extreme Graphics 200MHz	No
ICH4	ICH4	ICH4	ICH4	ICH4

Table 4.23 Intel Chipset Introduced in 2003 for Pentium 4

Chipset	865P	865PE	865G	875P
Code name	Springdale-P	Springdale-PE	Springdale-G	Canterwood
Date introduced	May '03	May '03	May '03	April '03
Part number	82865P	82865PE	82865G	82875P
Bus speeds	533/400MHz	800/533MHz	800/533MHz	800/533MHz
Supported processors	Pentium 4, Celeron[1]	Pentium 4, Celeron[1]	Pentium 4, Celeron[1]	Pentium 4, Celeron[1,2]
SMP (dual CPUs)	No	No	No	No
Memory types	DDR266/333 dual-channel[2]	DDR333/400 dual-channel	DDR333/400 dual-channel	DDR333/400 dual-channel
Parity/ECC	Neither	Neither	Neither	ECC
Maximum memory	4GB	4GB	4GB	4GB
Memory banks	2	2	2	2
PCI support	2.2	2.2	2.2	2.2
PCI speed/width	33MHz/32-bit	33MHz/32-bit	33MHz/32-bit	33MHz/32-bit
AGP slot	AGP 8x	AGP 8x	AGP 8x	AGP 8x
Integrated video	No	No	Yes	No
South Bridge (hub)	ICH5/ICH5R	ICH5/ICH5R	ICH5/ICH5R	ICH5/ICH5R

1. Supports current Socket 478 processors as well as Prescott core introduced 2003

2. Two modules per bank

Table 4.24 lists the ICH chips used by Pentium 4/Celeron 4 chipsets made by Intel.

Table 4.24 I/O Controller Hub Chips for Pentium 4 Chipsets

Chip Name	ICH0	ICH	ICH2	ICH4	ICH5	ICH5R
Part number	82801AB	82801AA	82801BA	828201DB	828201EB	828201ER
ATA support	UDMA-33	UDMA-66	UDMA-100	UDMA-100	UDMA-100	UDMA-100
SATA support	No	No	No	No	SATA-150	SATA-150
SATA Raid	No	No	No	No	No	RAID 0
USB support	1C/2P	1C/2P	2C/4P	3C/6P	4C/8P	4C/8P
USB 2.0	No	No	No	No	Yes	Yes
CMOS/clock	Yes	Yes	Yes	Yes	Yes	Yes
PCI Support	2.2	2.2	2.2	2.2	2.3	2.3
ISA support	No	No	No	Mo	No	No
LPC support	Yes	Yes	Yes	Yes	Yes	Yes
Power management	SMM/ ACPI 1.0	SMM/ ACPI 1.0	SMM/ ACPI 1.0	SMM/ ACPI 2.0	SMM/ ACPI 2.0	SMM/ ACPI 2.0
10/100 Ethernet	No	No	No	Yes	Yes	Yes

ICH = I/O controller hub
USB = Universal serial bus
xC/xP = # of controller / # of ports
ATA = AT attachment (IDE)
UDMA = Ultra-DMA ATA
ISA = Industry-standard architecture bus
LPC = Low pin count bus
SMM = System management mode
ACPI = Advanced configuration and power interface

Intel 850 Family

The Intel 850 family contains two members, the original 850 and an enhanced version called the 850E. The 850 is the first chipset for the Intel Pentium 4 processor and thus is also the first chipset to support the NetBurst microarchitecture. The 850 is designed for high-performance desktop computers and workstations and uses the same hub architecture and modular design as the rest of Intel's 8xx family of chipsets. See Figure 4.40 for a photo of the Intel 850 chipset.

The 850 has two main components, down from three in earlier 800-series chipsets:

- *82850 Memory Controller Hub.* Provides support for dual 400MHz RDRAM memory channels with a 3.2GBps bandwidth and a 100MHz system bus. The 82850 MCH also supports 1.5V AGP 4x video cards at a bandwidth exceeding 1GBps.
- *82801BA I/O Controller Hub 2.* The ICH2 (an enhanced version of the 82801 used by other 800-series chipsets) supports 32-bit PCI rev. 2.2, dual UDMA 33/66/100 IDE host adapters, four USB ports, an integrated LAN controller, six-channel AC-97 audio/modem codec, FWH interface support, SMBus support, and Alert on LAN and Alert on LAN 2 support.

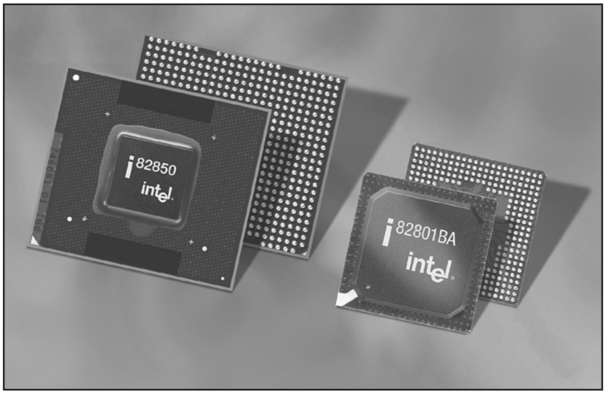

Figure 4.40 The Intel 850 chipset. *Photo used by permission of Intel Corporation.*

Optionally, the Intel 82562ET/82562EM Platform LAN communication chips can be added to the 850 chipset to provide support for 10BASE-T and Fast Ethernet networking, building on the LAN features in the 82801BA ICH2 chip.

The 850 chipset, similar to most recent Intel and non-Intel chipsets, also supports the CNR card for integrated audio, modem, and network capabilities. See Figure 4.41 for a diagram of the 850's chipset architecture.

The 850E is an enhanced version of the 850. Its 82850E MCH adds support for dual 533MHz Rambus RDRAM memory channels and support for PC1066 RIMM modules to the 850's standard features. It also uses the same ICH2 hub as the original 850.

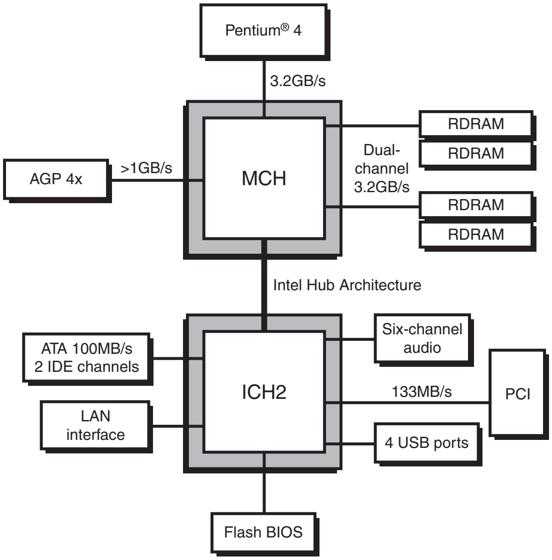

Figure 4.41 Architecture of the Intel 850 chipset.

Intel 845 Family

Unlike the 850 and 850E chipsets, the 845 family of chipsets is widely used by both Intel and third-party motherboard makers. If you purchased a Pentium 4 system from late 2001 through mid-2003, it probably uses some version of the 845 chipset. The 845, code named Brookdale during its development, was the first Pentium 4 chipset from Intel to support low-cost SDRAM instead of expensive RDRAM. Subsequent variations support DDR SDRAM at speeds up to DDR333, ATA/100, and USB 2.0.

The 845-series chipsets include the following models:

- 845
- 845GL
- 845GV
- 845G
- 845GE
- 845E
- 845PE

All members of the 845 family use the same hub-based architecture developed for the 845 family, but they also have onboard audio and support the communications and networking riser (CNR) card for integrated modem and 10/100 Ethernet networking. However, they differ in their support for different types and amounts of memory, integrated graphics, external AGP support, and which ICH chip they use.

Although the original version of the 845 supported only PC133 SDRAM memory, the so-called 845D model (a designation used by review sites but not by Intel) also supports 200/266MHz DDR SDDRAM. The Intel 845's 82845 MCH supports Socket 478–based Celeron or Pentium 4 processors and can support up to two DDR SDRAM modules or three standard SDRAM modules (depending on the motherboard). When DDR SDRAM is used, the 845 supports either 200MHz (PC2100) or 266MHz (PC2700) memory speeds, with an FSB speed of 400MHz. The 845 also supports ECC error correction when parity-checked memory modules are used and offers an AGP 4x video slot, but it has no onboard video.

The 845 uses the same ICH2 I/O controller hub chip (82801-BA) used by the Intel 850 and 850E chipsets in Rambus-based systems and the 815EP in low-cost SDRAM-based systems. The ICH2 supports ATA/100 hard disk interfacing, basic AC'97 sound, and four USB 1.1 ports.

All G-series 845 models feature Intel Extreme Graphics integrated video, which has faster core speeds and adds 3D performance to the bare-bones integrated video used by the 810 and 815 chipset families. Two chipsets—the 845G and 845GE—also offer support for AGP 4x video cards.

The 845E is an updated version of the current 845 model with ECC error correction and support for 533MHz FSB, whereas the 845PE supports the 533MHz FSB, DDR 266, and 333MHz memory, but it doesn't support ECC error correction. All models except the 845 (845D) use the enhanced ICH4 I/O Controller Hub 82801DB, which offers six USB 2.0 ports as well as integrated networking. Additionally, all models except the 845 and 845GL offer enhanced 20-bit audio.

Figure 4.42 compares the system block diagrams of the 845 and 845GE models.

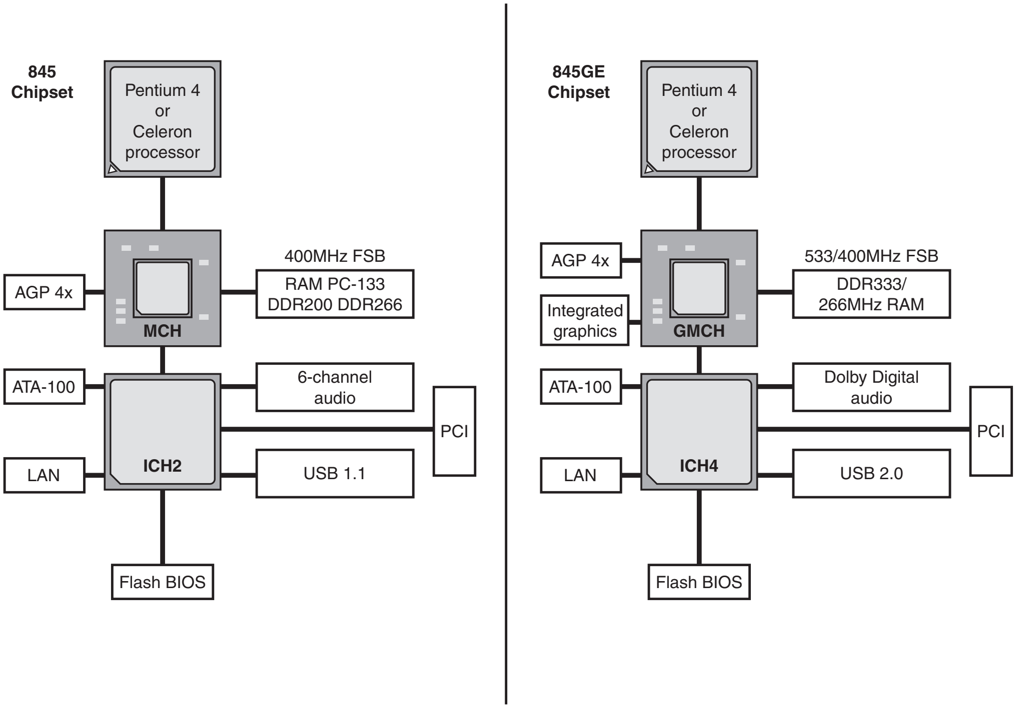

Figure 4.42 The 845GE (right) adds support for faster FSB speeds, memory, integrated graphics, and USB 2.0 to the basic 845 chipset architecture (left).

Intel Extreme Graphics Architecture

845-series chipsets with integrated video (845G-series models) support Intel's new Extreme Graphics Architecture, which supports 3D graphics and features the following four technologies to improve 3D rendering speed and quality:

- *Rapid Pixel and Texel Rendering Engine.* Uses pipelines to overlap 2D and 3D operations, provides 8x data compression to improve the use of memory bandwidth, and features a multitier cache for 3D operations
- *Zone Rendering.* Reduces memory bandwidth requirements by dividing the frame buffer in to rectangular zones, sorting the triangles into memory by zone, and processing each zone to memory
- *Dynamic Video Memory Technology.* Manages memory sharing between the display, applications, and operating system depending on the memory requirements of the programs running
- *Intelligent Memory Management.* Improves memory addressing, display buffer implementation, and memory efficiency

Extreme Graphics Architecture improves 3D rendering compared to Intel's earlier integrated video chipsets (the 810 and 815-series chipsets, which have no 3D functions at all), but its performance and features still lag behind even current mid-range video chipsets from NVIDIA and ATI. Extreme Graphics Architecture lacks hardware Transform and Lighting (T&L) features, a feature required by some of today's most popular games, and offers frame rates which, at best, are at about the level of the low-end NVIDIA GeForce 2 MX 200. Thus, even though you can play an occasional game with the G-series systems with integrated video, if you want serious game play, you'll need one of the 845 models that supports AGP 4x graphics cards and DDR333 memory, such as the 845GE or 845PE.

Intel 865 Family

The Intel 865 chipset family, code named Springdale, was released in May 2003. As the name might suggest, the 865 series is designed to replace the 845 series with chipsets that feature dual-channel memory support, the new communications streaming architecture (CSA) that provides a dedicated connection for the integrated network controller, faster performance, and support for the latest technologies (including Gigabit Ethernet and Serial ATA). The features of the 865 and 875 families are summarized in Table 4.23.

The 865 family includes the 865P, 865PE, and 865G chipsets. The 865PE and 865G support single-channel or dual-channel DDR266, as well as dual-channel DDR333 and DDR400 SDRAM with FSB speeds up to 800MHz. Dual-channel memory provides a wider memory bandwidth for faster performance. The 845P model supports DDR266/333 and FSB speeds up to 533MHz. All models support AGP 8x, and the G model includes Intel Extreme Graphics 2—a faster version of the integrated graphics technology found in G-series versions of the 845 chipset family.

All members of the 865 family use the new ICH5 I/O controller hub. A faster hub architecture called Hub Link 1.5 with 266MBps bandwidth connects the MCH/GMCH and ICH together.

ICH5 and ICH5R

ICH5 and ICH5R (Raid) are the latest generation of Intel's I/O controller hub, which is the equivalent of the South Bridge in Intel's hub-based architecture introduced with the 800 series of chipsets.

ICH5 and ICH5R feature four USB 2.0 controllers with eight external ports, two ATA/100 ports, and two Serial ATA/150 ports. ICH5R models add support only for RAID 0 (striping) on the SATA ports. ICH5/ICH5R also support the PCI 2.3 bus and include an integrated 10/100 Ethernet LAN controller.

Intel 875P

The Intel 875P chipset, code named Canterwood during its development, was introduced in April 2003. The 875P chipset supports Intel's HT Technology (hyper-threading), so it fully supports 3.06GHz and faster Pentium 4s, including the newer Prescott (90nm) core versions.

For faster memory access, the 875P supports four standard or ECC memory modules (up to 4GB total) using DDR333 or DDR400 memory in a dual-channel mode, and it offers a new Turbo mode that uses a faster path between DDR400 memory and the MCH to boost enhanced performance. Because multiple memory modules aren't always the same size or type, the 875P also features a new dynamic mode that optimizes system memory when different types or sizes of memory are used at the same time. The 875P also includes both Serial ATA and RAID support and uses the same ICH5 I/O controller hub used by the 865 series.

Third-Party Pentium 4 Chipsets

SiS, the Ali Corporation, ATI, and VIA all produce chipsets for the Intel Pentium 4 and Celeron 4 processors.

Although Intel's chipsets for the Pentium 4 have dominated the market up to this point, many of these chipsets offer unique features that are worth considering. The following sections discuss these chipsets by vendor.

High-Speed North-South Bridge Connections

As you learned earlier in this chapter, Intel has developed a replacement for the traditional North Bridge/South Bridge architecture known as *hub architecture*. This 266MBps interface provides a faster connection between the memory controller hub/graphics memory controller hub (North Bridge replacements) and the I/O controller hub (South Bridge replacements) in Intel 8xx-series chipsets for the Pentium III and Pentium 4 processor families.

Intel is not alone in replacing the slow PCI bus connection between North and South Bridge–type chips with a faster architecture that bypasses the PCI bus. Other companies introducing high-speed chipset interconnects include

- *VIA*. It created the V-Link architecture to connect its North and South Bridge chips at speeds matching or exceeding Intel hub architecture. V-Link uses a dedicated 8-bit data bus and is currently implemented in two versions: 4x V-Link and 8x V-Link. 4x V-Link transfers data at 266MBps (4×66MHz), which is twice the speed of PCI and matches the speed of Intel's hub architecture. 8x V-Link transfers data at 533MBps (4×133MHz), which is twice the speed of Intel's hub architecture. All VIA South Bridge chips in the VT82xx series support V-Link. The first chipsets to use V-Link were VIA's 266-series chipsets for the Pentium III, Pentium 4, and Athlon processor families. VIA's 333- and 400-series chipsets also use V-Link.
- *SiS*. Its MuTIOL architecture provides performance comparable to VIA's 4x V-Link. Chipsets that support MuTIOL use separate address, DMA, input data, and output data buses for each I/O bus master. MuTIOL buffers and manages multiple upstream and downstream data transfers over a bidirectional 16-bit data bus. South Bridge chips in the SiS961 and 962 series support MuTIOL at a data transfer rate of 533MBps (133MHz×4), whereas the SiS963 series supports the faster MuTIOL 1G, which supports data transfer rates exceeding 1GBps. The North Bridge chips that support MuTIOL are described in the sections listing SiS chipsets for Pentium 4 and Athlon-series processors.
- *ATI*. It uses a high-speed interconnect called A-Link in some of its IGP integrated chipsets. A-Link runs at 266MBps, matching Intel's hub architecture and first-generation V-Link and MuTIOL designs. Its RS- and RX-series chipsets use the HyperTransport bus. HyperTransport, which is now developed and managed by the nonprofit HyperTransport Technology Consortium (`www.hypertransport.org`), uses a packetized point-to-point interconnect IP bus that uses low voltage differential signaling. The 8x8 version of HyperTransport used by some ATI chipsets supports 800MHz clock speeds and a data transfer rate of 1.6GBps.
- *NVIDIA*. Its nForce, nForce2, and nForce3 chipsets use the HyperTransport bus originally developed by AMD.

Although the terms *North Bridge* and *South Bridge* continue to be used for chipsets using V-Link, MuTIOL, A-Link, or HyperTransport interconnects between these chipset components, these chipsets really use a hub-based architecture similar to Intel 8xx-series chipsets and receive corresponding boosts in performance as a consequence.

SiS Chipsets for Pentium 4

SiS offers several chipsets for the Pentium 4, including integrated chipsets, chipsets for use with discrete video accelerator cards, and one that supports Rambus RDRAM. Details of SiS's chipsets for the Pentium 4 are available in Tables 4.25 and 4.26. Unlike most of the chipsets SiS has created for the Pentium II/III/Celeron, the SiS chipsets for the Pentium 4 use one of several high-speed South Bridge equivalents (SiS 96x series Media I/O chips) instead of integrating North and South Bridge functions into a single chip. SiS North and South Bridge chips for the Pentium 4 use a high-speed 16-bit connection known as MuTIOL (Multi-Threaded I/O Link) instead of the slow PCI bus as with older chipsets.

Table 4.25 SiS North Bridge Chips for Pentium 4 Socket 478

Chipset	SiS650	SiS651	SiS645	SiS645DX	SiS648	SiS655	SiS R658
Bus speed	400MHz	400/ 533MHz	400MHz	400/ 533MHz	400/ 533MHz	400/ 533MHz	400/ 533MHZ
Supports hyper-threading	No	Yes*	No	Yes*	Yes*	Yes*	Yes*
SMP (dual CPUs)	No	No	No	No	No	No	No
Memory types	PC133, DDR266	PC100/133, DDR200 /266/333	PC133, DDR200/ 266	PC133, DDR266/ 333	DDR200/ 266/333	DDR266/ 333, dual-channel	1066/ 800MHz RDRAM
Parity/ECC	Neither	Neither	Neither	Neither	Neither	Neither	Neither
Maximum memory	3GB	3GB	3GB	3GB	3GB	4GB	4GB
PCI support	2.2	2.2	2.2	2.2	2.2	2.2	2.2
PCI speed/ width	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit
AGP slot	4x	4x	4x	4x	8x	8x	8x
Integrated video	Yes	Yes	No	No	No	No	No
South Bridge	SiS961	SiS962	SiS961	SiS961	SiS963	SiS963	SiS963
MuTIOL speed	533MBps	533MBps	533MBps	533MBps	1GBps	1GBps	1GBps

**B-stepping only*

Table 4.26 SiS Media I/O (South Bridge) Chips for Pentium 4

South Bridge Chip	USB Support	Number of USB Ports	ATA Support	Audio	10/100 Ethernet	HomePNA 1.0/2.0	IEEE-1394	MuTIOL Clock Speed
SiS961	1.1	6	33/66/ 100	AC'97, 5.1 channel	Yes	Yes	No	266MHz
SiS961B	1.1	6	33/66/ 100/133	AC'97, 5.1 channel	Yes	Yes	No	266MHz
SiS962	1.1, 2.0	6	33/66/ 100/133	AC'97, 5.1 channel	Yes	Yes	Yes	266MHz
SiS962L	1.1, 2.0	6	33/66/ 100/133	AC'97, 5.1 channel	Yes	Yes	No	266MHz
SiS963	1.1, 2.0	6	33/66/ 100/133	AC'97, 5.1 channel	Yes	Yes	Yes	533MHz
SiS963L	1.1, 2.0	6	33/66/ 100/133	AC'97, 5.1 channel	Yes	Yes	No	533MHz

SiS650/651 Chipsets

The SiS650 and 651 chipsets enable Pentium 4 system builders to create low-cost systems with onboard video that can be enhanced with AGP 4x video cards at a later date. The integrated video features support for high-quality DVD playback and the optional SiS301B video bridge for TV-out and DVI LCD panels.

Both chipsets also feature SiS's own MuTIOL technology for connecting the North Bridge and South Bridge chips with a three-layer high-speed (266MHz/533MBps bandwidth) data highway.

The 650 and 651 both support SDRAM and DDR SDRAM, and the 651 adds support for DDR333 memory, the 533MHz system bus of the latest Pentium 4 processors, and hyper-threading in its B-stepping version.

The 650's SiS961 South Bridge provides USB 1.1, ATA-100 (133 in its 961B version), AC'97 six-channel audio, and integrated Ethernet/HomePNA networking. The 651 also uses the newer SiS962 South Bridge, which provides ATA133 and USB 2.0 support.

SiS645/645DX

The 645 family of SiS chipsets does not include the integrated graphics of the 65x family, but they are otherwise similar. They support SDRAM and DDR SDRAM, AGP 4x, and the high-speed MuTIOL North Bridge/South Bridge interface. The 645DX supports DDR333 memory, the 533MHz system bus, and the HT technologies found in the most recent Pentium 4 processors.

Both the 645 and 645DX use the SiS961 South Bridge.

SiS648/655

The SiS648 chipset is a development of the 645DX chipset, with the following differences:

- Supports DDR memory only (up to DDR333)
- 8X AGP slot
- SiS963 South Bridge (USB 2.0, IEEE-1394a support)

The 655 chipset is essentially a dual-channel version of the 648, supporting up to 4GB of memory with DDR266/333 memory only.

SiS R658

The SiS R658 is the first SiS chipset ever to support Rambus RDRAM. Other features include

- Support for Pentium 4 processors with 533MHz and HT (B-stepping only) features
- Dual-channel support for PC1066/PC800 RDRAM (requires identical pairs of memory)
- 4GB maximum memory size
- AGP 8x interface
- MuTIOL 1G (533MHz clock speed providing more than 1GBps throughput) interface to the SiS963 South Bridge

Essentially, the R658 is an RDRAM version of the 655 chipset, and, like the 655, it uses the SiS963 South Bridge.

ALi Corporation Chipsets for Pentium 4

ALi Corporation (formerly known as Acer Laboratories) has produced several chipsets for the Pentium 4 and Celeron 4 processors. Tables 4.27 and 4.28 provide an overview of these chipsets, which are discussed in the following sections.

Table 4.27 ALi Chipsets for Pentium 4

Chipset	ALADDiN-P4	M1681
North Bridge chip	M1671	M1681
Bus speed	400MHz	400/533MHz*
SMP (dual CPUs)	No	No
Memory types	PC100/133, DDR200/266/333	PC100/133, DDR200/266/333/400
Parity/ECC	Neither	Neither
Maximum memory	3GB	?
PCI support	2.3	2.2
PCI speed/width	33MHz/32-bit	33MHz/32-bit
AGP slot	4x	8x
Integrated video	No	No
South Bridge	M1535 series	M1563
HyperTransport Link	N/A	400MBps

**Also supports hyper-threaded processors*

Table 4.28 ALi South Bridge Chips for Pentium 4 and Athlon XP

South Bridge Chip	USB Support	Number of USB Ports	ATA Support	Audio	Soft Modem	10/100 Ethernet	Super I/O	HyperTransport Throughput
M1535D	1.1	4	33/66	Yes[1]	Yes	No	Yes[4]	N/A
M1535D+	1.1	6	33/66/100/133	Yes[1,3]	Yes	No	Yes[4]	N/A
M1563[5]	2.0	6	66/100/133	Yes[2,3]	Yes	Yes	Yes	400MBps

1. Integrated DirectX 3D with wavetable and SoundBlaster Pro/16 compatibility; Also supports AC'97 audio codec

2. Supports 6-channel audio

3. Supports SPDIF digital audio interface

4. Includes fast IR interface, serial, parallel, PS/2 mouse, and keyboard ports

5. Also integrates Memory Stick and Secure Digital (SD) interfaces

Aladdin P4

The ALi Aladdin P4 was ALi's first Pentium 4-compatible chipset. Because it uses the same M1535-series South Bridge chips used by its earlier Pentium and Pentium II/III chipsets, the P4 is a traditional North Bridge/South Bridge solution. Thus, it relies on the slow (133MBps) PCI interface to carry data between the bridge chips.

The P4's major features include the following:

- 400MHz system bus
- Support for PC100/133 and DDR200/266/333 memory
- ATA-133 support (when used with the M1535D+ South Bridge)
- AGP 4x interface
- USB 1.1 ports
- ACPI power management

The P4 is also available in a version for notebook computers: the ALADDiN-P4M, which uses the D1535+ South Bridge.

M1581/M1563

ALi's M1581/M1563 chipset for the Pentium 4 processor brings ALi's offerings in line with the latest from other chipset vendors. In a break with ALi tradition, it uses the HyperTransport high-speed direct connection between North and South Bridge chips instead of relying on the PCI bus, as with previous designs.

Its major features include

- Support for hyper-threading and 533MHz system bus
- Support for DDR memory up to DDR400 and PC100/133 SDRAM
- ATA-133
- USB 2.0
- AGP 8x interface
- Memory Stick and SD (Secure Digital) flash memory device interfaces
- ACPI power management
- HyperTransport high-speed link between North and South Bridge chips, running at >400MBps bandwidth in each direction (800MBps total throughput)

ATI Chipsets for Pentium 4

ATI's line of chipsets for the Pentium 4 integrate Radeon VE-level 3D graphics, DVD playback, and dual-display features with high-performance North Bridge and South Bridge designs. ATI uses its high-speed A-Link bus to connect its North and South Bridge chips, but it also supports connections to third-party South Bridge chips via the PCI bus. This enables system designers to create an all-ATI or a mix-and-match solution. Many of the first Radeon IGP-based systems on the market actually used ALi or VIA South Bridge chips.

The Radeon IGP North Bridge chips for Pentium 4 include

- Radeon IGP 330 (uses IXP-series or third-party South Bridge)
- Radeon IGP 340 (uses IXP-series or third-party South Bridge)
- RS250 (uses SB-series)
- RS300 (uses SB-series)
- RS300VE (uses SB-series)

ATI's South Bridge chips include

- IXP 200 (for IGP series)
- IXP 250 (for IGP series)
- SB300C (for RS series)
- SB380 (for RS series)
- SB210 (for RS series)

Table 4.29 summarizes the major features of the North Bridge chips, and Table 4.30 summarizes the major features of the South Bridge chips.

Table 4.29 Radeon IGP (North Bridge) Chips for Pentium 4

North Bridge Chip	Radeon IGP 330	Radeon IGP 340	RS250	RS300	RS300VE
Bus speed	400MHz	400/ 533MHz	400/ 533MHz	400/533/ 800MHz[2]	400/533/ 800MHz[2]
SMP (dual CPUs)	No	No	No	No	No
Memory types	DDR200/ 266	DDR200/ 266/333	DDR200/ 266/333	DDR333/ 400[3]	DDR333/ 400
Parity/ECC	Neither	Neither	Neither	Neither	Neither
Maximum memory	1GB	1GB	1GB		
PCI support	2.2	2.2	2.3	2.3	2.3
PCI speed/ width	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit
AGP slot	4x	4x	4x	8x	8x
Integrated video	Radeon VE[1]	Radeon VE[1]	Radeon VE[1]	Radeon 9000	Radeon 9000
NS/SB interconnect speed	266MBps	266MBps	266MBps	Up to 12.8GBps	Up to 12.8GBps
NS/SB interconnect type	A-Link	A-Link	A-Link	HyperTransport 8×8	HyperTransport 8×8

1. Same core as ATI Radeon 7000, with support for dual displays

2. Also supports hyper-threaded processors

3. Dual-channel memory support

Table 4.30 ATI South Bridge Chips for Pentium 4 and Athlon

South Bridge Chip	USB Support	Number of USB Ports	ATA Support	Audio	Soft Modem	10/100 Ethernet	Super I/O	High-Speed Interconnect
IXP 200/ 250[1]	2.0	6	ATA33/ 66/100	AC'97, SDPIF	No	3Com	Yes	A-Link
SB300C	2.0	6	ATA33/66/ 100/133 SATA	AC'97 6-channel	No	3Com	Yes	A-Link
SB380	2.0	6	ATA33/66/ 100/133 SATA	AC'97 6-channel	No	3Com	Yes	HyperTransport 8×8
SB210[2]	2.0	6	ATA33/66/100	AC'97, SDPIF	No	3Com	Yes	A-Link

1. IXP250 identical features to IXP200, plus it supports Wake On LAN (WOL), Desktop Management Interface (DMI), manage boot agent (MBA), and the Alert Standards Forum (ASF) mechanism

2. Pin-compatible with SB300C and SB380, but has the same features as IXP200

VIA Chipsets for Pentium 4

Although VIA Technologies produces a line of chipsets for the Pentium 4, it lacks a license from Intel for the Socket 478 interface. VIA claims rights to Socket 478 through a patent exchange agreement with Intel, but Intel disputes this. As a result of the legal battles over the VIA chipsets for Pentium 4, most of the motherboards that use VIA's chipsets are manufactured by VIA's Platform Solutions Division (VPSD), although they might be sold under a variety of brand names.

Tables 4.31 and 4.32 provide an overview of VIA's chipsets for the Pentium 4, including ProSavage chipsets with integrated graphics.

Table 4.31 VIA Chipsets for Pentium 4

Chipset	P4X266	P4X266A	P4X266E	P4M266	P4X400	P4X400A
North Bridge chip	VT8753	VT8753A	VT8753E	VT8751	VT8754	VT8754CE
Bus speeds	400MHz	400MHz	400/533MHz	400MHz	400/533MHz	400/533/800MHz
SMP (dual CPUs)	No	No	No	No	No	No
Memory types	PC100/133, DDR200/266	PC100/133, DDR200/266	DDR200/266	PC100/133, DDR200/266	DDR200/266/333/400	DDR200/266/333/400
Parity/ECC	Neither	Neither	Neither	Neither	ECC	ECC
Maximum memory	4GB	4GB	4GB	4GB	32GB	32GB
PCI support	2.2	2.2	2.2	2.2	2.2	2.2
PCI speed/width	33MHz/32-bit*	33MHz/32-bit*	33MHz/32-bit*	33MHz/32-bit	33MHz/32-bit*	33MHz/32-bit*
AGP slot	4x	4x	4x	4x	8x	8x
Integrated video	No	No	No	S3 Graphics ProSavage8 3D	No	No
South Bridge	VT8233, VT8233C, VT8233A	VT8233, VT8233C, VT8233A	VT8233, VT8233C, VT8233A, VT8235	VT8233, VT8233C, VT8233A	VT8235	VT8235
V-Link speed	266MBps	266MBps	266MBps	266MBps	533MBps	533MBps

**Supports 66MHz/64-bit PCI when optional VPX-64 (VT8101) chip is used*

Table 4.32 lists the major features of the VIA South Bridge chips used in VIA's chipsets for the Pentium 4. Note that these same chips are also used by VIA chipsets for the AMD Athlon family of processors. All chipsets that use these South Bridge chips use VIA's high-speed V-Link interface between North and South Bridge chips. These chipsets connect to the VT1211 LPC (low pin count) or equivalent Super I/O chip for support of legacy devices such as serial, IR, and parallel ports and the floppy drive.

Table 4.32 VIA South Bridge Chips for Pentium 4

South Bridge Chip	USB Support	Number of USB Ports	ATA Support	Audio	10/100 Ethernet	HomePNA	V-Link Throughput
VT8233	1.1	6	33/66/100	AC'97, 6-channel[1]	Yes	Yes	266MBps
VT8233A	1.1	6	33/66/100/133	AC'97, 6-channel[1]	Yes	No	266MBps
VT8233C	1.1	6	33/66/100	AC'97, 6-channel[1]	Yes[2]	No	266MBps
VT8235	2.0	6	33/66/100/133	AC'97, 5.1 channel[1]	Yes	No	533MBps

[1] *Integrated audio requires separate audio codec chip on motherboard; it also supports MC'97 soft modem.*

[2] *3Com 10/100 Ethernet.*

VIA Modular Architecture Platforms (V-MAP) for Pentium 4

VIA's North and South Bridge chips for the Pentium 4 support VIA's Modular Architecture Platforms (V-MAP) designs, which enable motherboard designers to convert quickly to more advanced versions of a chipset because of a common pinout. The North Bridge chips used in the P4X266, 266A, 266E, P4M266, and P4X400 chipsets are all pin-compatible with each other, as are the 8233/8235-series South Bridge chips. Thus, motherboards using these chipsets can be built in a variety of configurations. All these chipsets also support VIA's V-Link high-speed connection between the North and South Bridge chips.

VIA Apollo P4X266 Family

The VIA Apollo P4X266 is its first chipset for the Pentium 4 and Celeron 4 processors, supporting AGP 4x, 4GB of RAM, and the 400MHz system bus used by early Pentium 4/Celeron 4 processors. The P4X266A improves the memory interface and queues more instructions (up to 12) in the processor bus interface to reduce latency and improve performance. The P4X266E adds support for the 533MHz bus used in the 2.53GHz (and faster) Pentium 4 processors. It also supports both the VT8233 and newer VT8235 series of South Bridge chips.

ProSavage P4M266

The VIA ProSavage P4M266 integrates the S3 Graphics ProSavage8 2D/3D graphics accelerators with the features of the P4X266 chipset. Unlike some other chipsets with integrated graphics, the P4M266 retains an AGP 4x slot, so users can upgrade to faster AGP 4x graphics in the future.

The ProSavage8 core uses 32MB of system RAM for its frame buffer, supports AGP 8x bandwidth internally with 128-bit data paths, and features DVD DXVA Motion Compensation to improve the quality of DVD playback. In addition, it supports all members of the 8233/8235 family of South Bridge chips.

Apollo P4X400 and P4X400A

The VIA Apollo P4X400 chipset is an improved version of the short-lived P4X333. It's suitable for both server and workstation/desktop computer use, thanks to its support for up to 32GB of RAM and ECC memory. It also supports 400MHz and 533MHz system bus speeds and DDR memory up to 333MHz. It uses the VT8235 South Bridge, so it also supports the latest I/O standards (USB 2.0 and ATA-133).

The P4X400A chipset features improved timings and support for DDR400 memory and uses the VT8235 South Bridge.

Athlon/Duron/Athlon XP Chipsets

The original AMD Athlon was a Slot A processor chip, but subsequent versions have used Socket A, as have the now-discontinued Duron and current Athlon XP. Although similar in some ways to the Pentium III and Celeron, the AMD chips use a different interface and require different chipsets. AMD was originally the only supplier for Athlon chipsets, but VIA Technology, ALi Corporation, SiS, NVIDIA, and ATI now provide a large number of chipsets with a wide range of features. These chipsets are covered in the following sections.

AMD Chipsets for Athlon/Duron Processors

AMD makes four chipsets for Athlon and Duron processors: the AMD-750 and AMD-760/MP/MPX. The major features of these chipsets are compared in Table 4.33 and described in greater detail in the following sections.

Table 4.33 AMD Athlon/Duron Processor Chipsets Using North/South Bridge Architecture

Chipset	AMD-750	AMD-760
Code name	Irongate	None
Date introduced	Aug. 1999	Oct. 2000
Part number	AMD-751	AMD-761
Bus speed	200MHz	200/266MHz
Supported processors	Athlon/Duron	Athlon/Duron
SMP (dual CPUs)	No	Yes
Memory type	SDRAM	DDR SDRAM
Memory speed	PC100	PC1600/PC2100
Parity/ECC	Both	Both
Maximum memory	768MB	2GB buffered, 4GB registered
PCI support	2.2	2.2
AGP support	AGP 2x	AGP 4x
South Bridge	AMD-756	AMD-766
ATA/IDE support	ATA-66	ATA-100
USB support	1C/4P	1C/4P
CMOS/clock	Yes	Yes
ISA support	Yes	No
LPC support	No	Yes
Power management	SMM/ACPI	SMM/ACPI

AGP = Accelerated graphics port
ATA = AT attachment (IDE) interface
DDR-SDRAM = Double data rate SDRAM
ECC = Error correcting code
ISA = Industry Standard Architecture
LPC = Low pin count bus
PCI = Peripheral component interconnect
SDRAM = Synchronous dynamic RAM
SMP = Symmetric multiprocessing (dual processors)
USB = Universal serial bus

AMD-750

AMD's first chipset for its own Slot A and Socket A processors, called the AMD-750, is a traditional North/South Bridge design specifically for the Athlon and Duron processors. The AMD-750 chipset consists of the AMD-751 North Bridge and the AMD-756 South Bridge.

The AMD-751 system controller connects between the AMD Athlon processor bus to the processor and features the memory controller, AGP 2x controller, and PCI bus controller. The AMD-756 South Bridge includes a PCI-to-ISA bridge, USB controller interface, and ATA 33/66 controller.

The AMD-750 chipset includes the following features:

- AMD Athlon 200MHz processor bus
- PCI 2.2 bus with up to six masters
- AGP 2x
- PC-100 SDRAM with ECC
- Up to 768MB of memory
- ACPI power management
- ATA-33/66 support
- USB controller
- ISA bus support
- Integrated 256-byte CMOS RAM with clock
- Integrated keyboard/mouse controller

AMD-760 Family

The AMD-760 chipset was introduced in October 2000 and is notable as the first chipset supporting DDR SDRAM memory. The AMD-760 chipset consists of the AMD-761 system controller (North Bridge) in a 569-pin plastic ball-grid array (PBGA) package and the AMD-766 peripheral bus controller (South Bridge) in a 272-pin PBGA package. See Figure 4.43 for details of the 760’s block diagram.

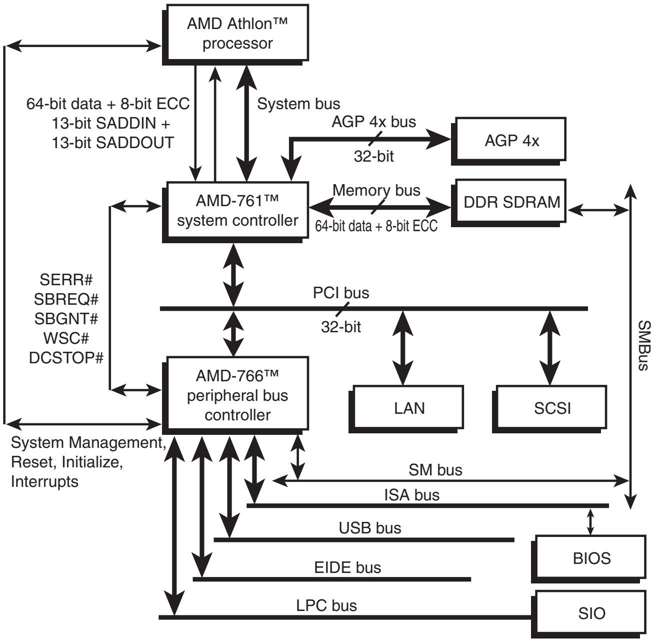

Figure 4.43 AMD-760 chipset block diagram.

The AMD-761 North Bridge features the AMD Athlon system bus, DDR-SDRAM system memory controller with support for either PC1600 or PC2100 memory, AGP 4x controller, and PCI bus controller. The 761 allows for 200MHz or 266MHz processor bus operation and supports the newer Athlon chips that use the 266MHz processor (also called front-side) bus.

The AMD-766 South Bridge includes a USB controller, dual UDMA/100 ATA/IDE interfaces, and the LPC bus for interfacing newer Super I/O and ROM BIOS components.

The AMD-760 chipset includes the following features:

- AMD Athlon 200/266MHz processor bus
- Dual processor support
- PCI 2.2 bus with up to six masters
- AGP 2.0 interface that supports 4x mode
- PC1600 or PC2100 DDR SDRAM with ECC
- Support for a maximum of 2GB buffered or 4GB registered DDR SDRAM
- ACPI power management
- ATA-100 support
- USB controller
- LPC bus for Super I/O support

The AMD-760MP chipset, which uses the AMD-762 North Bridge chip, is a development of the basic AMD-760 design that supports dual-processor Athlon MP systems. It differs from the standard 760 chipset in the following ways:

- Supports dual AMD Athlon MP processors with 200/266MHz processor bus speeds
- Up to 4GB PC2100 DDR (registered modules)
- Supports 33MHz PCI slots in 32-bit and 64-bit widths

The AMD-760MPX chipset uses the same AMD-762 North Bridge chip as the AMD-760MP to support multiple Athlon MP processors, but it uses the AMD-768 peripheral bus controller (South Bridge) chip. It differs from the 760MP chipset in the following ways:

- The AMD-762 North Bridge chip is used to support two 66MHz 32/64-bit PCI slots.
- The AMD-768 South Bridge chip is used to support 33MHz/32-bit PCI slots.

The 760MPX chipset is a better choice for a server because of its support for 66MHz and 64-bit PCI slots, whereas the 760MP is a suitable choice for a workstation.

None of these chipsets support USB 2.0, ATA-133, or DDR333 or faster memory. If you buy an Athlon, a Duron, or an Athlon XP desktop system, it's far more likely that your system will contain a third-party chipset than an AMD chipset; however, the 760MP and 760MPX chipsets are popular choices for AMD-based workstations and servers. The following sections cover the third-party chipsets made for the Athlon, Duron, and Athlon XP processors.

VIA Chipsets for Athlon, Duron, and Athlon XP

VIA Technologies, Inc., is the largest chipset and processor supplier outside of Intel and AMD. Originally founded in 1987, VIA is based in Taipei, Taiwan, and is the largest integrated circuit design firm on the island. VIA is a fabless company, which means it farms out the manufacturing to other companies with chip foundry capability. Although it is best known for its chipsets, in 1999 VIA purchased the Cyrix processor division from National Semiconductor and the Centaur processor division from IDT, respectively, thereby becoming a supplier of processors in addition to chipsets. VIA has also formed a joint venture with SONICblue (formerly S3) as a means of integrating graphics capabilities into various chipset products. This joint venture is known as S3 Graphics, Inc.

VIA makes chipsets for Intel, AMD, and Cyrix (VIA) processors. Table 4.34 provides an overview of VIA's Athlon/Duron chipsets, which use the traditional North/South Bridge architecture.

More recently, VIA has converted to a new architecture called V-Link, which uses a fast dedicated connection between the North Bridge and South Bridge. V-Link is similar to Intel's hub architecture, as well as HyperTransport (used by ALi, NVIDIA, and ATI) and A-Link (used by ATI). V-Link is also used by VIA's Pentium 4 chipsets. Tables 4.35 and 4.36 provide an overview of V-Link chipsets, which also support the VIA Modular Architecture Platform (V-MAP) design. As with VIA's chipsets for the Pentium 4, V-MAP uses identical pinouts for the V-Link North and South Bridge chips so vendors can reuse a motherboard design with more advanced chipsets as they are developed.

Both types of chipsets are discussed in the following sections.

Table 4.34 VIA Athlon/Duron/Athlon XP Processor Chipsets Using North/South Bridge Architecture

Chipset	Apollo KX133	Apollo KT133	Apollo KT133A	Apollo KLE133
Introduced	Aug. 1999	June 2000	Dec. 2000	March 2001
North Bridge	VT8371	VT8363	VT8363A	VT 8361
Processor support	Athlon	Athlon/Duron	Athlon/Duron	Duron
CPU interface	Slot-A	Socket-A (462)	Socket-A (462)	Socket-A (462)
CPU FSB	200MHz	200MHz	200/266MHz	200/266MHz
AGP	4x	4x	4x	Integrated AGP 2x
PCI	2.2	2.2	2.2	2.2
Memory type	SDRAM	SDRAM	SDRAM	SDRAM
Memory speed	PC133	PC133	PC100/133	PC100/133
Maximum memory	2GB	2GB	2GB	2GB
South Bridge	VT82C686A	VT82C686A	VT82C686B	VT82C686B
ATA/IDE	ATA-66	ATA-66	ATA-100	ATA-100
USB ports	1C4P	1C4P	1C4P	1C4P
Power management	SMM/ACPI	SMM/ACPI	SMM/ACPI	SMM/ACPI
Super I/O	Yes	Yes	Yes	Yes
CMOS/Clock	Yes	Yes	Yes	Yes

Table 4.35 VIA Athlon XP/Duron Processor Chipsets Using V-Link Architecture

Chipset	Apollo KT266	Apollo KT266A	Apollo KT333	ProSavage KM266	Apollo KT400	Apollo KT400A
North Bridge chip	VT8366	VT8633A	VT8753E	VT8375	VT8377	VT8377A
Bus speed	200, 266MHz	200, 266MHz	200, 266, 333MHz	200, 266MHz	200, 266, 333MHz	200, 266, 333MHz
SMP (dual CPUs)	No	No	No	No	No	No
Memory types	PC100/133, DDR200/266	PC100/133, DDR200/266	DDR200/266/333	PC100/133, DDR200/266	DDR200/266/333	DDR200/266/333/400
Parity/ECC	Neither	Neither	Neither	Neither	Neither	Neither
Maximum memory	4GB	4GB	4GB	4GB	4GB	4GB
PCI support	2.2	2.2	2.2	2.2	2.2	2.2

Table 4.35 Continued

Chipset	Apollo KT266	Apollo KT266A	Apollo KT333	ProSavage KM266	Apollo KT400	Apollo KT400A
PCI speed/ width	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit
AGP slot	4x	4x	4x	4x	8x	8x
Integrated video	No	No	No	S3 Graphics ProSavage8 3D	No	No
South Bridge	VT8233, VT8233C, VT8233A	VT8233, VT8233C, VT8233A	VT8233, VT8233C, VT8233A, VT8235	VT8233, VT8233C, VT8233A	VT8235	VT8235CE, VT8237
V-Link speed	266MBps	266MBps	266MBps	266MBps	533MBps	533MBps

Table 4.36 VIA South Bridge Chips for Athlon/Duron/Athlon XP

South Bridge Chip	USB Support	Number of USB Ports	ATA Support	Audio	10/100 Ethernet	HomePNA	V-Link Throughput
VT8233	1.1	6	33/66/100	AC'97, 6-channel[1]	Yes	Yes	266MBps
VT8233A	1.1	6	33/66/100/133	AC'97, 6-channel[1]	Yes	No	266MBps
VT8233C	1.1	6	33/66/100	AC'97, 6-channel[1]	Yes[2]	No	266MBps
VT8235	2.0	6	33/66/100/133	AC'97, 6-channel[1]	Yes	No	533MBps
VT8237	2.0	6	33/66/100/133; Serial ATA; Optional SATA RAID	AC'97, 6-channel[1]	Yes	No	533MBps

1. Integrated audio requires separate audio codec chip on motherboard; it also supports MC'97 soft modem.

2. 3Com 10/100 Ethernet.

VIA Technologies Apollo KX133

The VIA Apollo KX133 chipset brings AGP 4x, PC133, a 200MHz FSB, and ATA-66 technologies to the AMD Athlon processor platform, exceeding the performance of AMD's own 750 chipset. It was the first chipset to support AGP 4x.

Key features include the following:

- 200MHz processor bus
- AGP 4x graphics bus
- PC133 SDRAM memory
- 2GB maximum RAM
- ATA-66 support
- Four USB ports
- AC97 link for audio and modem
- Hardware monitoring
- Power management

The VIA Apollo KX133 is a two-chip set consisting of the VT8371 North Bridge controller and the VT82C686A South Bridge controller.

VIA Technologies Apollo KT133 and KT133A

The VIA Apollo KT133/A chipsets are designed to support the AMD Athlon and Duron processors in Socket-A (462) form. Based on the prior KX133 chipset, the KT133/A differ mainly in their support for Socket-A (462) over the previous Slot-A processor interface.

The VIA Apollo KT133 and KT133A are both two-chip sets consisting of the VT8363 North Bridge and the VT82C686A South Bridge (KT133) or the VT8363A North Bridge and the VT82C686B South Bridge (KT133A).

Both the KT133 and KT133A support the following standard features:

- Athlon/Duron Socket-A (462) processors
- 200MHz CPU (front-side) bus
- AGP 4x
- 2GB RAM maximum
- PC100/PC133MHz SDRAM
- PCI 2.2
- ATA-66
- USB support
- AC-97 audio
- Integrated Super I/O
- Integrated hardware monitoring
- ACPI power management

The KT133A (VT8363A North Bridge with VT82C686B South Bridge) adds the following features:

- 266MHz CPU (front-side) bus
- ATA-100

ProSavage PM133

The VIA ProSavage PM133 integrates S3 Graphics' S3 Savage 4 and S3 Savage 2000 3D and 2D graphics engines with the Apollo Pro KT133 chipset. The major features of the chipset are the same as for the Apollo Pro KT133, with the following additions:

- 2MB–32MB shared memory architecture integrated with Savage 4 3D and Savage 2000 2D video
- Z-buffering, 32-bit true-color rendering, massive 2K-by-2K textures, single-pass multiple textures, sprite antialiasing, and other 3D features
- Support for DVD playback, DVI LCD displays, and TV-out
- PCI 2.2 compliance

An optional AGP 4x interface enables the integrated AGP 4x video to be upgraded with an add-on card if desired. This two-chip chipset consists of the VT8365 North Bridge and VT8231 South Bridge.

The VT8231 South Bridge integrates the Super I/O and supports the LPC interface.

Apollo KT266 and KT266A

The Apollo KT266 is the first VIA chipset for Athlon-based systems to support VIA's high-speed V-Link system architecture. V-Link connects the 552-pin VT8366 North Bridge to the 376-pin VT8233 series South Bridge with a 266MBps data pathway, which is twice as fast as traditional PCI-based connections.

Major features of the KT266 include system bus speeds of 200/266MHz, AGP 2x/4x interface, and up to 4GB of DDR200/266 DDR SDRAM or PC100/133 SDRAM. Other features vary with the South Bridge (VT8233, VT8233A, or VT8233C) chip used with the VT8366.

The KT266A is a pin-compatible upgrade to the original KT266's North Bridge. The KT266A's VT8366A includes VIA's Performance Driven Design, which is not a technical term but a marketing term for the A-series' chips improved memory timing and deeper command queues to improve chipset performance. Basic features of the KT266A are otherwise similar to the KT266.

ProSavage KM266

The ProSavage KM266 combines the features of the KT266 with the graphics core of the S3 Graphics ProSavage 8 2D/3D accelerator. Unlike some other chipsets with integrated graphics, the P4M266 retains an AGP 4x slot, so users can upgrade to faster AGP 4x graphics in the future.

The ProSavage8 core uses 32MB of system RAM for its frame buffer, supports AGP 8x bandwidth internally with 128-bit data paths, and features DVD DXVA Motion Compensation to improve the quality of DVD playback. It supports all members of the VT8233 family of South Bridge chips and has a 266MBps 4x V-Link connection between North and South Bridge chips.

Apollo KT333

The Apollo KT333 is a pin-compatible development of the KT266A, adding support for a 333MHz system bus, 333MHz memory bus, and DDR333 memory. Unlike the KT266A, the KT333 no longer supports PC100/133 memory, but it uses the same KT8233 family of South Bridge chips.

Apollo KT400

The Apollo KT400 is the first VIA chipset for the Athlon XP processor to offer AGP 8x and a second-generation 533MB/s V-Link connection to the South Bridge. It uses the new VT8235 South Bridge, which is VIA's first South Bridge chip to support USB 2.0 as well as ATA-133.

The combination of faster video, fast memory and system bus speeds, and faster V-Link connections make the KT400 among the fastest Athlon XP chipsets.

Apollo KT400A

In previous A-series chipsets from VIA, the North Bridge component was replaced with an improved chip and the South Bridge component was retained. However, the KT400A features new designs for both its North Bridge (VT8377A) and South Bridge (VT8237) chips. Its major features include

- System bus support up to 333MHz
- DDR memory support up to DDR400
- Up to 4GB of memory
- An expanded array of prefetch buffers to reduce memory latency and improve throughput (FastStream64)
- AGP 8x interface
- Integrated six-channel Surround Sound AC'97 audio
- Eight USB 2.0 ports
- Integrated MC'97 modem
- Integrated 10/100 Ethernet
- Serial ATA
- ATA 33/66/100/133
- ACPI/OnNow power management

KT400A's FastStream64 enables the system to reach 3.2GBps memory transfer speeds without the need for more expensive dual-channel memory support.

Figure 4.44 shows the architecture of the KT400A chipset.

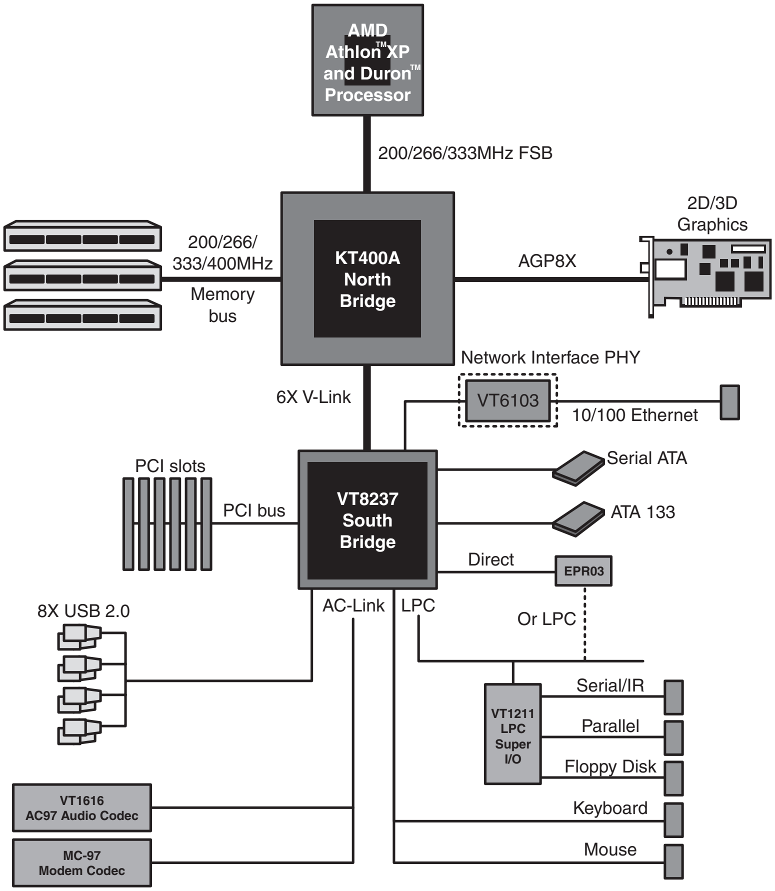

Figure 4.44 VIA VT400A block diagram.

Silicon Integrated Systems Chipsets for AMD Athlon/Duron Processors

SiS has a variety of chipsets for the Athlon, Duron, and Athlon XP processors. Tables 4.37 and 4.38 provide an overview of these chipsets, some of which use SiS's unique single-chip design and others of which use a high-speed two-chip design similar to other vendors' chipsets. These chipsets are discussed in the following sections.

Table 4.37 SiS Chipsets for Athlon/Duron/Athlon XP Processors

Chipset	SiS730S	SiS740	SiS733	SiS735	SiS745	SiS746	SiS746FX	SiS748
Bus speed	200/ 266MHz	266MHz	200/ 266MHz	200/ 266MHz	266MHz	266MHz	266MHz	Up to 400MHz
SMP (dual CPUs)	No	No	No	No	No	No	No	No
Memory type	PC133/ SDRAM	DDR266/ PC133	PC133	PC133/ DDDR266	DDR266/ 333	DDR266/ 333	DDR266/ 333/400	DDR266/ 333/400

Table 4.37 Continued

Chipset	SiS730S	SiS740	SiS733	SiS735	SiS745	SiS746	SiS746FX	SiS748
Parity/ECC	Neither	Neither	Neither	Neither	Neither	Neither	Neither	Neither
Maximum memory	1.5GB	1.5GB	1.5GB	1.5GB	3GB	3GB	3GB	3GB
PCI support	2.2	2.2	2.2	2.2	2.2	2.2	2.2	2.2
PCI speed/ width	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit
AGP slot	4x	None	4x	4x	4x	8x	8x	8x
Integrated video	Yes[1]	Yes[2]	No	No	No	No	No	No
South Bridge	N/A[3]	SiS96x series	N/A[3]	N/A[3]	N/A[3]	SiS963 series	SiS963 series	SiS963 series
MuTIOL speed	N/A	533MBps	N/A	N/A	N/A	1GBps[5]	1GBps[5]	1GBps[5]
ATA support	ATA-100	Varies[4]	ATA-100	ATA-100	ATA-100	Varies[4]	Varies[4]	Varies[4]
USB support	1.1/4 ports	Varies[4]	1.1/6 ports	1.1/6 ports	1.1/6 ports	Varies[4]	Varies[4]	Varies[4]
Audio support	Wavetable audio	Varies[4]	AC'97, SPDIF out	AC'97, SPDIF out	AC'97	Varies[4]	Varies[4]	Varies[4]
10/100 Ethernet	Yes	Varies[4]	No	Yes	No	Varies[4]	Varies[4]	Varies[4]
IEEE-1394a	No	Varies[4]	No	No	Yes	Varies[4]	Varies[4]	Varies[4]

1. 2D/3D accelerator with hardware DVD playback and optional SiS301 Video Bridge to TV and secondary monitor.

2. DirectX7-compliant 3D features, including two pixel rendering pipelines and four texture units.

3. Single-chip (combined N/S Bridge) design.

4. Varies with MuTIOL South Bridge chip used.

5. Chipset uses HyperStreaming technology, an improved version of MuTIOL.

Table 4.38 SiS MuTIOL South Bridge Chips for Athlon XP

South Bridge Chip	USB Support	Number of USB Ports	ATA Support	Audio	10/100 Ethernet	HomePNA 1.0/2.0	IEEE-1394	MuTIOL Clock Speed
SiS961	1.1	6	33/66/100	AC'97, 5.1 channel	Yes	Yes	No	266MHz
SiS961B	1.1	6	33/66/100/133	AC'97, 5.1 channel	Yes	Yes	No	266MHz
SiS962	1.1, 2.0	6	33/66/100/133	AC'97, 5.1 channel	Yes	Yes	Yes	266MHz
SiS962L	1.1, 2.0	6	33/66/100/133	AC'97, 5.1 channel	Yes	Yes	No	266MHz
SiS963	1.1, 2.0	6	33/66/100/133	AC'97, 5.1 channel	Yes	Yes	Yes	533MHz
SiS963L	1.1, 2.0	6	33/66/100/133	AC'97, 5.1 channel	Yes	Yes	No	533MHz

SiS MuTIOL High-Speed North/South Bridge Connection

The SiS96x-series South Bridge chips use a high-speed bus called MuTIOL to connect with compatible North Bridge chips. The original version of MuTIOL (supported by the SiS961- and 962-series chips) is a 16-bit wide 266MHz connection that provides 533MBps bandwidth, twice the speed of the Intel hub architecture used by Intel's 800-series chipsets.

The SiS963-series and matching North Bridge chips use a second generation of MuTIOL called MuTIOL 1G, which supports a 16-bit-wide 533MHz connection to achieve bandwidths exceeding 1GBps.

When connected to the SiS748 North Bridge, the SiS963-series chips use a further development of MuTIOL called HyperStreaming, which integrates the following four technologies to further improve the speed of data transfer:

- *Single Stream with Low Latency Technology.* Improves performance by 5%–43% depending on the activity.
- *Multiple Stream with Pipelining and Concurrent Execution Technology.* Uses concurrent parallel data pipelines and simultaneous processing of nonsequential data. In file-copy operations, for example, performance increases as data file size increases.
- *Specific Stream with Prioritized Channel Technology.* Improves playback of Internet music, video, and applications such as IP telephony and videoconferencing.
- *Smart Stream Flow Control Technology.* Analyzes characteristics of different interfaces and improves performance.

SiS730S

The SiS730S is a high-performance, low-cost, single-chip chipset with integrated 2D/3D graphics and support for Socket A versions of the AMD Athlon and Duron.

The integrated video is based on a 128-bit graphic display interface with AGP 4x performance. In addition to providing a standard analog interface for CRT monitors, the SiS730S also provides the DFP for a digital flat panel monitor. An optional SiS301 video bridge supports NTSC/PAL TV output. The SiS730S also supports an AGP 4x slot, enabling users to upgrade to a separate AGP card in the future.

The SiS730S also includes integrated 10/100Mb Fast Ethernet as well as an AC-97-compliant interface that comprises a digital audio engine with 3D-hardware accelerator, on-chip sample rate converter, and professional wavetable along with separate modem DMA controller. SiS730S also incorporates the LPC interface for attaching newer Super I/O chips and a dual USB host controller with six USB ports. The SiS730S can also be used with ISA slots if an optional LPC/ISA bridge chip is used.

Features of the SiS730S include

- Support for AMD Athlon/Duron processors with 200MHz system bus
- Support for PC133 SDRAM
- Meets PC99 requirements
- PCI 2.2 compliant
- Four PCI masters
- Support for Ultra DMA100
- Integrated AGP 2x 2D/3D video/graphics accelerator
- Support for digital flat panel
- Hardware DVD decoding
- Built-in secondary CRT controller for independent secondary CRT, LCD, or TV digital output
- LPC interface

- Advanced PCI H/W audio (Sound Blaster 16 and DirectSound 3D compliant) and modem
- Meets ACPI 1.0, APM 1.2 requirements
- PCI Bus Power Management Interface Spec. 1.0
- Integrated keyboard/mouse controller
- Dual USB controller with six USB ports
- Integrated 10/100Mbps Ethernet controller

SiS733 and SiS735

The SiS733 and SiS735 are high-performance single-chip sets that support the AMD Athlon and Duron Socket A processors. Similar to other SiS single-chip sets, the SiS733 and SiS735 incorporate the features of a traditional North Bridge, South Bridge, and Super I/O chip into a single chip.

The SiS733 supports PC133 SDRAM and uses a 682-pin BGA package. The SiS735 supports either PC133 or DDR266 SDRAM and integrates 10/100 Fast Ethernet and HomePNA 1Mbps/10Mbps Home Network interfaces. The SiS735 also uses a 682-pin BGA package.

The SiS733 and SiS735 share the following features:

- Support for 4x AGP
- Up to six PCI masters
- Dual UDMA/100 IDE host adapters
- 1.5GB RAM maximum
- Six USB ports
- AC'97 audio and AMR support
- Integrated RTC
- LPC interface for support of MIDI, joystick, and legacy BIOS devices
- PC2001 compliant

Figure 4.45 illustrates the SiS733's block diagram.

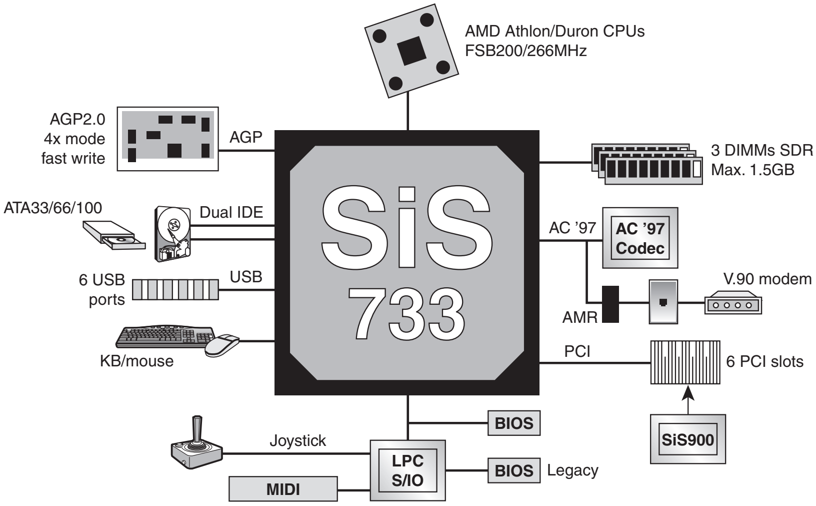

Figure 4.45 The SiS733 provides a single control interface for all major motherboard functions.

SiS740

The SiS740 is a dual-chip design that provides a high-speed integrated video solution for Athlon-class processors. Its North Bridge and South Bridge chips use the high-speed MuTIOL connection to transfer data. Its major features include

- Integrated Real256 2D/3D graphics core with full DirectX 7 compatibility
- Up to 128MB of shared memory
- Hardware DVD playback
- DDR266 memory support

It is designed to use the SiS961- or 962-series South Bridge chips.

SiS745

The SiS745 is the first single-chip solution to integrate IEEE-1394a (FireWire 400) as part of its I/O. It is designed to provide a high-performance legacy-free solution. Its key features include the following:

- DDR266/333 memory support up to 3GB
- Support for Athlon XP processor as well as earlier models
- Six USB 1.1 ports
- Three IEEE-1394a ports
- Legacy keyboard, mouse, floppy, MIDI, and joystick interfaces
- ATA-100
- AC’97 audio and AMR (audio modem riser) support for a V.90 soft modem

SiS746 and SiS746FX

The SiS746 is the first Athlon/Duron/Athlon XP–compatible chipset on the market to feature an AGP 8x interface. It is a two-piece chipset designed to connect with the SiS963-series South Bridge chipsets.

Key features include 266MHz system bus, support for DD266/333 memory, and a second-generation MuTIOL connection between the North and South Bridge chips that provides a 1BGps bandwidth.

The companion SiS963L South Bridge chip adds the following features: ATA-133 support, six USB 2.0 ports, six-channel AC’97 audio, and MII interface for HomePNA or 10/100 Ethernet networking.

The SiS963 South Bridge chip adds IEEE-1394a (FireWire 400) support.

The SiS746FX North Bridge is an enhanced version of the SiS746, adding support for the 333MHz system bus and approved DDR400 memory. It also uses the SiS963 series of South Bridge chips.

SiS748

Like the SiS746 series, the SiS748 also uses the SiS963 series of South Bridge chips, but it uses the latest HyperStreaming technology for a faster and more intelligent connection between the chips. Its other major features include

- Up to 400MHz system bus
- DDR266/333/400 memory support
- AGP 8x interface

Figure 4.46 shows the system architecture of the SiS748 chipset when using the SiS963L South Bridge chip. If the SiS963 South Bridge chip is used in place of the SiS963L pictured, three IEEE-1394a ports are also available.

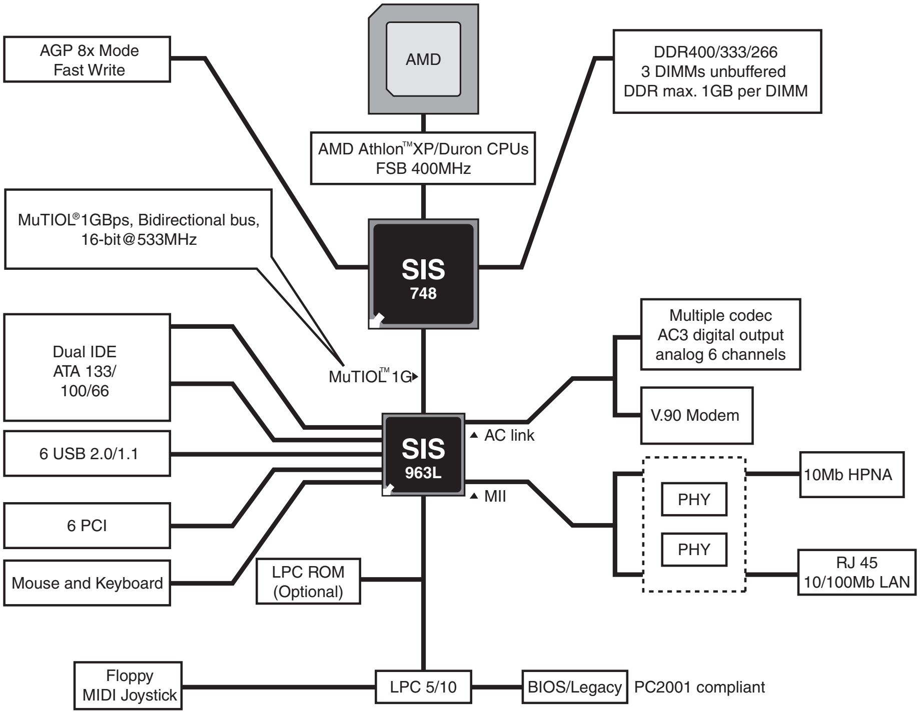

Figure 4.46 Block diagram for SiS748 chipset with SiS963L South Bridge.

ALiMagik1 for AMD Athlon/Duron Systems

ALi Corporation makes only one chipset for AMD Athlon/Duron processors: the ALiMagik1.

ALiMagik1 is a two-chip chipset that uses the M1647 Super North Bridge and the M1535D+ South Bridge (which is also used by its Pentium III/Celeron chipsets). The M1647 Super North Bridge is a 528-pin BGA chip.

The M1647 Super North Bridge supports SDRAM as well as DDR SDRAM at speeds of 200MHz or 266MHz. It supports up to 3GB of RAM but does not support ECC. Memory timing is x-1-1-1-1-1-1-1 in back-to-back SDRAM reads. Because the M1647 supports both conventional and DDR SDRAM, system manufacturers can use the same chipset for both memory types.

The M1647 supports AGP 4x video, PCI 2.2, up to six PCI masters beyond the North Bridge and PCI bridge, ACPI and Legacy green power management, PCI Mobile CLKRUN#, and AGP Mobile BUSY#/STOP#.

When combined with an M1535+ South Bridge chip, the chipset is called the MobileMagik1 and can be used on Athlon- or Duron-based portable systems.

Because this chipset, unlike others that use DDR memory, still uses the traditional 133MBps PCI bus connection between North and South Bridge chips, its performance is among the lowest of any Athlon chipset.

NVIDIA nForce Chipsets for Athlon/Duron/Athlon XP

NVIDIA, although best known for its popular GeForce line of graphics chipsets, is also becoming a popular vendor of chipsets for the AMD Athlon/Duron/Athlon XP processor family with its nForce and nForce2 product families.

nForce's advanced features include

- 400MBps HyperTransport link between chipset components. nForce is the first PC chipset to use HyperTransport.
- Dual-channel crossbar memory controller to provide high-speed memory access when identical pairs of memory are used. It uses independent 64-bit memory controllers.
- nView multidisplay hardware (420 chipset with integrated video) supports dual displays with integrated video.
- AGP 4x interface.
- Dynamic adaptive speculative preprocessor (DASP) to reduce latency and improve prefetching of data.
- StreamThru architecture to improve isochronous (time-dependent) data transfers for network and broadband through the chipset's integrated 10/100 Ethernet port.
- Integrated GeForce2 MX video (420 chipset with integrated video) with support for DVI LCD flat panels.
- True hardware-based audio processing with support for Dolby Digital (AC-3) 5.1 channel audio (SoundStorm chipsets).

nForce2 improvements over nForce include the following:

- DualDDR improved dual-channel memory controller that supports up to DDR400 memory and allows dual-channel operation with two or three DIMMs
- Optional IEEE-1394a support
- Optional GeForce4 MX integrated video
- AGP 8x interface

Table 4.39 provides an overview of the North Bridge chips in the nForce and nForce2 families, and Table 4.40 provides an overview of the nForce/nForce2 South Bridge chips. Unlike most other chip makers, NVIDIA does not make Pentium 4–compatible chipsets; the nForce is a descendant of the custom chipset NVIDIA created for the Microsoft Xbox console game system.

nForce North Bridge chips with integrated graphics are known as *integrated graphics processors (IGPs)*, whereas those that require separate AGP video are known as *system platform processors (SPPs)*. All South Bridge chips are known as media and communications processors (MCP). IGP/SPP and MCP chips communicate over an 800MBps HyperTransport connection.

Table 4.39 nForce/nForce2 IGP/SPP (North Bridge) Chips

North Bridge Chip	nForce 420	nForce 415	nForce2 IGP	nForce2 SPP
Bus speed	200/ 266MHz	200/ 266MHz	200/266/ 333MHz	200/ 266MHz
SMP (dual CPUs)	No	No	No	No
Memory type	DDR200/ 266MHz PC100/133	DDR200/ 266MHz PC100/133	DDR200/ 266/333/ 400[1]	DDR200/ 266/333/ 400[1]

Table 4.39 Continued

North Bridge Chip	nForce 420	nForce 415	nForce2 IGP	nForce2 SPP
Parity/ECC	Neither	Neither	Neither	Neither
Maximum memory	4GB	4GB	3GB	3GB
Dual-channel mode	Yes[2]	Yes[2]	Yes	Yes
PCI support	2.2	2.2	2.2	2.2
PCI speed/width	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit	33MHz/ 32-bit
AGP slot	4x	4x	8x	8x
Integrated video	GeForce2 MX	No	GeForce4 MX	No
HyperTransport speed	400MBps	400MBps	400MBps	400MBps
South Bridge chip	nForce MCP, MCP-D	nForce MCP, MCP-D	nForce2 MCP, MCP-T	nForce2 MCP, MCP-T

1. Requires external AGP card to support DDR400 memory.

2. Use only two identical memory modules to enable this mode.

Table 4.40 nForce/nForce2 MCP (South Bridge) Chips

South Bridge Chip	USB Support	Number of USB Ports	ATA Support	Audio	10/100 Ethernet	IEEE-1394	Used with
nForce MCP	1.1	6	33/66/100	AC'97, 5.1-channel	Yes	No	nForce IGP, SPP
nForce MCP-D[1]	1.1	6	33/66/100	NVIDIA Audio Processing Unit, Dolby Digital 5.1 channel, DirectX 8 3D audio, SPDIF	Yes	No	nForce IGP, SPP
nForce2 MCP	1.1, 2.0	6	33/66/ 100/133	AC'97, 2-channel, 3D audio	Yes	No	nForce2 IGP, SPP
nForce2 MCP-T[1]	1.1, 2.0	6	33/66/ 100/133	NVIDIA Audio Processing Unit, Dolby Digital 5.1 channel, DirectX 8 3D audio, SPDIF	Yes (dual NVIDIA[2] and 3Com)	Yes	nForce2 IGP, SPP

1. Also known as NVIDIA SoundStorm

2. Also supports HomePNA networking

The combination of advanced memory controllers, prefetch design, HyperTransport high-speed connection, and hardware audio processing in MCP-D and MCP-T chips makes the second-generation nForce2 chipsets among the fastest chipsets available for Athlon XP processors.

Figure 4.47 shows the architecture of the nForce2 IGP and MCP-T combination, which provides the greatest versatility. If the SPP North Bridge is used instead of the IGP, integrated video is not present. If the MCP South Bridge is used instead of the MCP-T, IEEE-1394a, hardware 5.1 Dolby Digital audio, and dual network ports are not available.

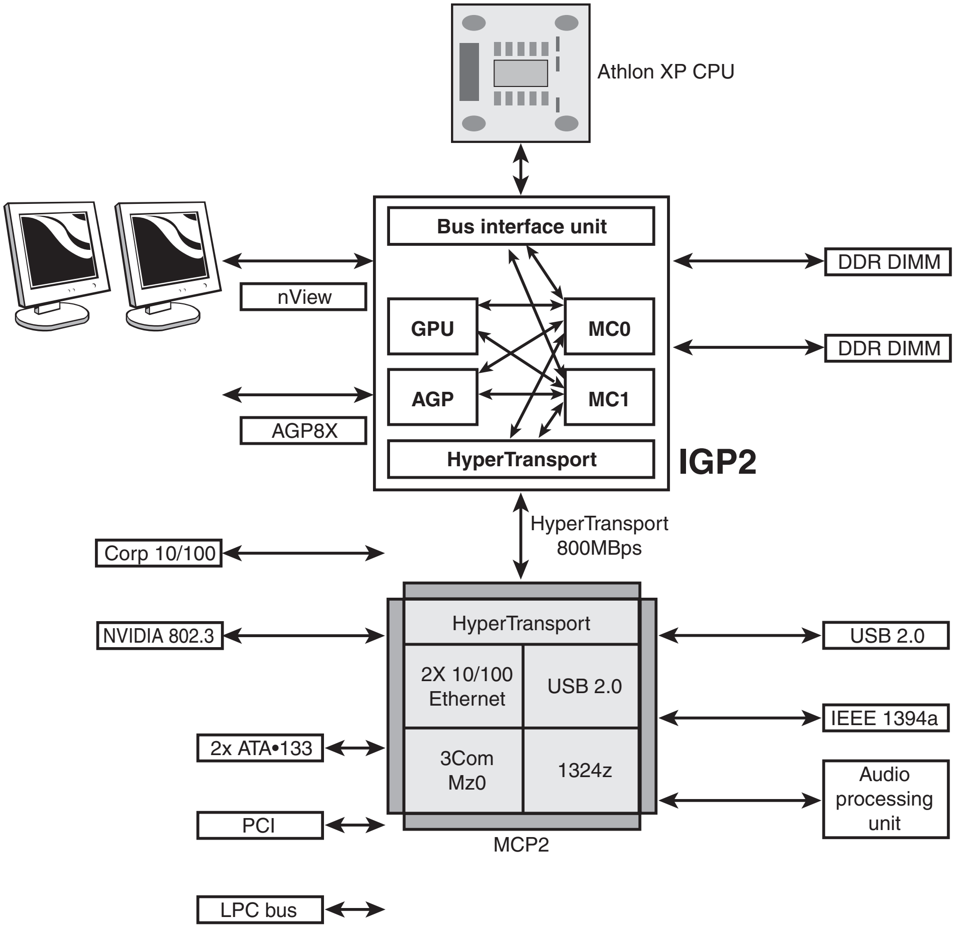

Figure 4.47 NVIDIA nForce2 IGP/MCP2 chipset architecture.

ATI Radeon IGP Chipsets for Athlon/Duron/Athlon XP

ATI's line of chipsets for the Athlon series of processors integrates Radeon VE-level 3D graphics, DVD playback, and dual-display features with high-performance North Bridge and South Bridge designs. ATI uses its high-speed A-Link bus to connect its North and South Bridge chips, but it also supports connections to third-party South Bridge chips via the PCI bus. This enables system designers to create an all-ATI or a mix-and-match solution. Many of the first Radeon IGP-based systems on the market used ALi or VIA South Bridge chips.

The Radeon IGP North Bridge chips for Athlon-series processors include the Radeon IGP 320. To create an all-ATI solution, the IGP 320 can be matched with either of ATI's South Bridge chips: the IXP 200 or IXP 250. Both of these chips support six USB 2.0 ports and ATA33/66/100.

Table 4.41 summarizes the major features of the IGP 320, and Table 4.42 summarizes the major features of the IXP 200 and 250.

Table 4.41 Radeon IGP (North Bridge) Chip for Athlon

North Bridge Chip	Radeon IGP 320	North Bridge Chip	Radeon IGP 320
Bus speed	200/266MHz	PCI support	2.2
SMP (dual CPUs)	No	PCI speed/width	33MHz/32-bit
Memory type	DDR200/266	AGP slot	4x
Parity/ECC	Neither	Integrated video	Radeon VE*
Maximum memory	1GB	A-Link speed	266MBps

**Same core as ATI Radeon 7000, with support for dual displays*

Table 4.42 ATI South Bridge Chips for Athlon

South Bridge Chip	IXP 200/250*	South Bridge Chip	IXP 200/250*
USB Support	6 USB 2.0 ports	Ethernet LAN	3Com 10/100
ATA Support	ATA100	Super I/O	Yes
Audio Support	AC'97, S/PDIF	High Speed Interconnect	A-Link

**The IXP250 has identical features to IXP200; plus it supports Wake On LAN (WOL), Desktop Management Interface (DMI), manage boot agent (MBA), and the Alert Standards Forum (ASF) mechanism.*

Intel Workstation Chipsets for Pentium 4 and Xeon

Intel has developed several chipsets for workstations based on the Pentium 4 and Xeon processors. The following sections describe these chipsets in detail; Table 4.43 provides a quick reference to their features.

Table 4.43 Workstation Chipsets

Chipset	860	E7205	E7505
Code name	Colusa	Granite Bay	Placer
Date introduced	May 2001	Dec. 2002	Dec. 2002
Part number	82860	E7205	E7505
Bus speeds	400MHz	533/400MHz	533/400MHz
Supported processors	Xeon	Pentium 4[5]	Xeon 533MHz FSB and 512KB L2 cache[5]
SMP (dual CPUs)	Yes (2)	No	Yes (2)
Memory type	4 RDRAM PC800[1]	DDR200/266 SDRAM (unbuffered)	DDR200/266 (in pairs)
Parity/ECC	Both	Both	Both
Maximum memory	4GB (with two MRHR chips)	4GB	16GB
Memory banks	Up to 4[3]	Up to 4	Up to 6 (registered memory) or 4 (unbuffered memory)
PCI support	2.2	2.2	2.2
PCI speed/width	32-bit/33MHz[2]	32-bit/33MHz	32-bit/33MHz[4]
AGP slot	AGP 4x/2x	AGP8x-1x	AGP8x-1x
Integrated video	None	None	None
South Bridge (Hub)	ICH2	ICH4	ICH4

1. Up to eight on motherboards with MRHR chips

2. 64-bit 33/66MHz on motherboards with P64H chips

3. Two banks when MRHR chip is not present

4. 64-bit 33/66MHz and PCI-X on motherboards with P64H2 chips

5. Supports HT Technology

Intel 860

The Intel 860 is a high-performance chipset designed for the Socket 602 (Pentium 4–based) Xeon processors for DP workstations. The 860 uses the same ICH2 as the Intel 850 but uses a different MCH—the 82860, which supports one or two Socket 602 ("Foster") Xeon processors. The other major features of the 82860 are similar to those of the 82850, including support for dual 400MHz RDRAM memory channels with a 3.2GBps bandwidth and a 100MHz system bus. The 82860 MCH also supports 1.5V AGP 4x video cards at a bandwidth exceeding 1GBps.

The 860 chipset uses a modular design, in which its two core chips can be supplemented by the 82860AA (P64H) 66MHz PCI Controller Hub and the 82803AA MRHR. The 82860AA supports 64-bit PCI slots at either 33MHz or 66MHz, and the 82803AA converts each RDRAM memory channel into two, which doubles memory capacity. Thus, whether a particular 860-based motherboard offers 64-bit or 66MHz PCI slots or dual-channel RDRAM memory depends on whether these supplemental chips are used in its design.

Intel E7205

The Intel E7205 chipset, known as Granite Bay during its development, is designed to support both workstation and high-performance PC applications. It supports DDR200/266 SDRAM modules with a system bus speed up to 533MHz and uses the ICH4 I/O controller hub, just as some versions of the 845 chipset do. However, the E7205 supports ECC and parity-checked memory for better system reliability and supports all standard-voltage speeds of AGP from 1x to 8x with an AGP Pro slot (nonstandard 3.5V versions of AGP once sold by some vendors such as 3dfx will not work). It supports hyper-threading for use with the 3.06GHz and faster Pentium 4 processors.

Intel E7505

The Intel E7505 chipset, known as Placer during its development, is in some ways an updated version of the 860 chipset, adding support for faster processors and more advanced hardware than the 860 offers.

The E7505 supports a system bus of up to 533MHz, matching the single or dual Xeon 533MHz FSB and 512KB L2 cache processors it supports; it also supports the HT Technology included in these processors. The E7505 supports pairs of DDR200/266 memory up to 16GB total, four times as much as the E7205 and the 860. It can use up to six registered or four unbuffered memory modules and supports ECC. Its Intel x4 single-device data correction (SDDC) can correct up to four errors per memory module for better system reliability.

Its AGP Pro slot supports all speeds of AGP from 1x to 8x (except for the nonstandard 3.5V versions of AGP once sold by some vendors), and it uses the ICH4 I/O controller hub. To achieve 66MHz/64-bit PCI and 133MHz PCI-X support, the E7505 can be used with up to three optional P64H2 (82870P2) chips, an improved version of the P64H chip that is an optional part of the 860 chipset.

Chipsets for Athlon 64

The Athlon 64 processor requires a new generation of chipsets, both to support its 64-bit processor architecture and to allow for integration of the memory controller into the processor (the memory controller has traditionally been located in the North Bridge chip or equivalent). AMD, VIA Technologies, NVIDIA, ATI, and ALi Corporation have developed chipsets for the Athlon 64.

Table 4.44 lists the major features of the first chipsets developed for the Athlon 64.

Table 4.44 Chipsets for Athlon 64

Vendor	Chipset Model	Opteron Support	Memory Support	Video Support	ATA/Serial ATA Support
AMD	8151	Yes	DDR200/266/333	AGP 8x	ATA-133, SATA
ALi	M1687/M1563	Yes	DDR266/333	AGP 8x	ATA-133
ATI	IGP 380	No	DDR200/266/333	AGP 8x, Radeon 9000	ATA-133
VIA	K8T400	Yes	DDR200/266/333	AGP 8x	ATA-133, SATA
VIA	K8M400	Yes	DDR200/266/333	SavageXP integrated graphics	ATA-133, SATA
VIA	K8T400M	Yes	DDR200/266/333	AGP 8x	ATA-133
NVIDIA	Crush K8	No	DDR200/266/333	AGP 8x	ATA-133
NVIDIA	Crush K8G	No	DDR200/266/333	AGP 8x, GeForce4 MX	ATA-133

Many of these chipsets introduce enhanced versions of features seen in previous chipsets. For example, the ALi M1687/M1653 is the first PC chipset to support the 16-bit/800MHz version of the HyperTransport interconnect. HyperTransport 16×8 has a maximum bandwidth of 3.2GBps to help prevent data-transfer bottlenecks.

Some chipsets add new features. For example, AMD's 8151 Graphics Tunnel (North Bridge) also uses HyperTransport at speeds equal or superior to ALi's to connect to the AMD 8111 HyperTransport I/O Hub, and it supports an optional AMD 8131 PCI-X Tunnel chip to permit system and motherboard designers to add high-speed PCI-X slots to an Athlon 64 system.

Figure 4.48 shows the architecture of the AMD 8151 chipset for Athlon 64.

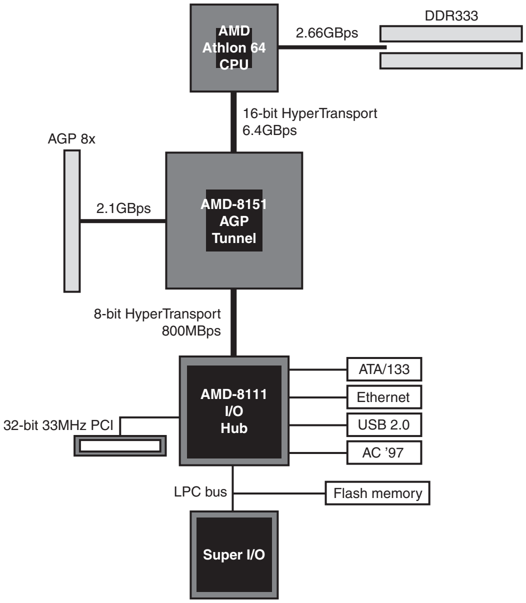

Figure 4.48 Block diagram of the AMD 8151 chipset for Athlon 64.

Super I/O Chips

The third major chip seen on many PC motherboards is called the Super I/O chip. This is a chip that integrates devices formerly found on separate expansion cards in older systems.

Most Super I/O chips contain, at a minimum, the following components:

- Floppy controller
- Dual serial port controllers
- Parallel port controller

The floppy controllers on most Super I/O chips handle two drives, but some of them can handle only one. Older systems often required a separate floppy controller card.

The dual serial port is another item that was formerly on one or more cards. Most of the better Super I/O chips implement a buffered serial port design known as a universal asynchronous receiver transmitter (UART), one for each port. Most mimic the standalone NS16550A high-speed UART, which was created by National Semiconductor. By putting the functions of two of these chips into the Super I/O chip, these ports are essentially built in to the motherboard.

Virtually all Super I/O chips also include a high-speed multimode parallel port. The better ones allow three modes: standard (bidirectional), Enhanced Parallel Port (EPP), and the Enhanced Capabilities Port (ECP) modes. The ECP mode is the fastest and most powerful, but selecting it also causes your port to use an ISA bus 8-bit DMA channel—usually DMA channel 3. As long as you account for this and don't set anything else to that channel (such as a sound card and so on), the EPC mode parallel port should work fine. Some of the newer printers and scanners that connect to the system via the parallel port use ECP mode, which was invented by Hewlett-Packard.

The Super I/O chip can contain other components as well. For example, the Intel VC820 ATX motherboard uses an SMC (Standard Microsystems Corp.) LPC47M102 Super I/O chip. This chip incorporates the following functions:

- Floppy drive interface
- Two high-speed serial ports
- One ECP/EPP multimode parallel port
- 8042-style keyboard and mouse controller

This chip is typical of recent Super I/O chips in that it has an integrated keyboard and mouse controller. Older Super I/O chips lacked this feature.

One thing I've noticed over the years is that the role of the Super I/O chip has decreased more and more in the newer motherboards. This is primarily due to Intel and other chipset manufacturers moving Super I/O functions, such as IDE, directly into the chipset South Bridge or ICH component, where these devices can attach to the PCI bus (North/South Bridge architecture) or to the high-speed hub interface (hub architecture) rather than the ISA bus. One of the shortcomings of the Super I/O chip is that originally it was interfaced to the system via the ISA bus and shared all the speed and performance limitations of that 8MHz bus. Moving the IDE over to the PCI bus allowed higher-speed IDE drives to be developed that could transfer at the faster 33MHz PCI bus speed.

Newer Super I/O chips interface to the system via the LPC bus, an interface designed by Intel to offer a low-speed connection (up to about 6.67MBps) using only 13 signals. Although technically slower than ISA, it is much more efficient. Because high-speed devices such as IDE/ATA drives are now interfaced through the South Bridge chip or PCI bus, nothing interfaced through the current Super I/O chips needs any greater bandwidth anyway.

As the chipset manufacturers combine more and more functions into the main chipset, and as USB- and IEEE-1394-based peripherals replace standard serial, parallel, and floppy controller–based devices, we will probably see the Super I/O chip continue to fade away in motherboard designs. More and more chipsets are combining the South Bridge and Super I/O chips into a single component (often referred to as a

Super South Bridge chip) to save space and reduce parts count on the motherboard. Several of the SiS chipsets even integrate all three chips (North Bridge, South Bridge, and Super I/O) into a single chip.

Motherboard CMOS RAM Addresses

In the original AT system, a Motorola 146818 chip was used as the RTC and Complementary Metal-Oxide Semiconductor (CMOS) RAM chip. This was a special chip that had a simple digital clock, which used 10 bytes of RAM and an additional 54 more bytes of leftover RAM in which you could store anything you wanted. The designers of the IBM AT used these extra 54 bytes to store the system configuration.

Modern PC systems don't use the Motorola chip; instead, they incorporate the functions of this chip into the motherboard chipset (South Bridge) or Super I/O chip, or they use a special battery and NVRAM module from companies such as Dallas or Benchmarq.

▶▶ For more details on the CMOS RAM addresses, see "Motherboard CMOS RAM Addresses," p. 392.

Motherboard Interface Connectors

There are a variety of connectors on a modern motherboard. Figure 4.49 shows the locations of these connectors on a typical motherboard (using the Intel D845PEBT2 model as the example). Several of these connectors, such as power supply connectors, serial and parallel ports, and keyboard/mouse connectors, are covered in other chapters. Tables 4.45–4.55 contain the pinouts of most of the other different interface and I/O connectors you will find.

▶▶ See "AT Power Supply Connectors," p. 1138.
▶▶ See "Serial Ports," p. 961, and "Parallel Ports," p. 971.
▶▶ See "Keyboard/Mouse Interface Connectors," p. 995.
▶▶ See "Universal Serial Bus," p. 947.
▶▶ See "ATA I/O Connector," p. 508.

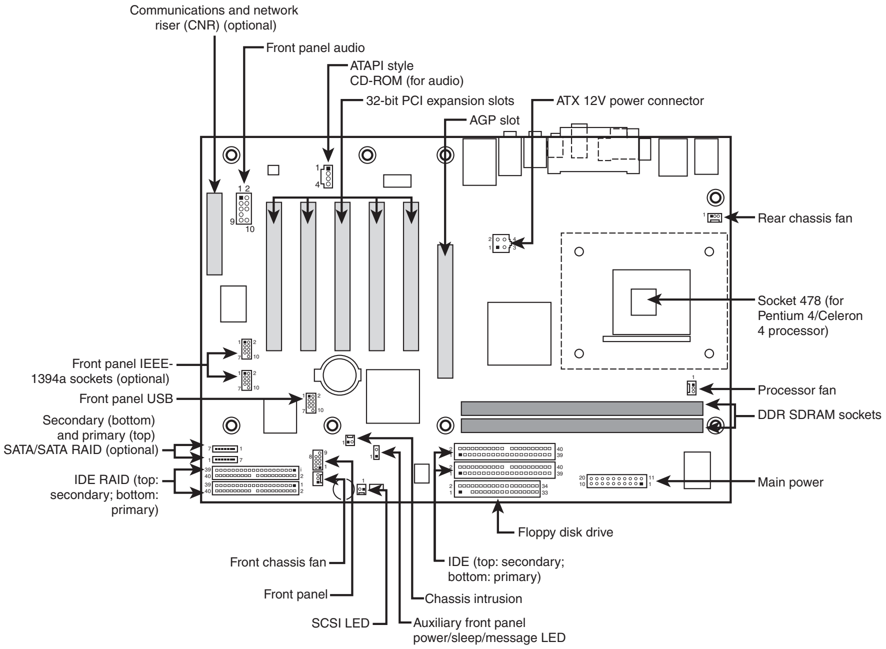

Figure 4.49 Typical motherboard connectors (Intel D845PEBT2 shown).

Table 4.45 Infrared Data (IrDA) Pin-Header Connector

Pin	Signal	Pin	Signal
1	+5V	4	Ground
2	Key	5	IrTX
3	IrRX	6	CONIR (Consumer IR)

Table 4.46 Battery Connector

Pin	Signal	Pin	Signal
1	Gnd	3	KEY
2	Unused	4	+6V

Table 4.47 LED and Keylock Connector

Pin	Signal	Pin	Signal
1	LED Power (+5V)	4	Keyboard Inhibit
2	KEY	5	Gnd
3	Gnd		

Table 4.48 Speaker Connector

Pin	Signal	Pin	Signal
1	Ground	3	Board-Mounted Speaker
2	KEY	4	Speaker Output

Table 4.49 Microprocessor Fan Power Connector

Pin	Signal Name
1	Ground
2	+12V
3	Sense tachometer

Table 4.50 Chassis Intrusion (Security) Pin-Header

Pin	Signal Name
1	Ground
2	CHS_SEC

Table 4.51 Wake on LAN Pin-Header

Pin	Signal Name
1	+5 VSB
2	Ground
3	WOL

Table 4.52 Wake on Ring Pin-Header

Pin	Signal Name
1	Ground
2	RINGA

Note

Some boards have a board-mounted piezo speaker. It is enabled by placing a jumper over pins 3 and 4, which routes the speaker output to the board-mounted speaker. Removing the jumper enables a conventional speaker to be plugged in.

Caution

Do not place a jumper on this connector; serious board damage will result if the 12V is shorted to ground.

Table 4.53 CD Audio Connector

Pin	Signal Name	Pin	Signal Name
1	CD_IN-Left	3	Ground
2	Ground	4	CD_IN-Right

Table 4.54 Telephony Connector

Pin	Signal Name	Pin	Signal Name
1	Audio Out (monaural)	3	Ground
2	Ground	4	Audio In (monaural)

Table 4.55 ATAPI-Style Line In Connector

Pin	Signal Name	Pin	Signal Name
1	Left Line In	3	Ground
2	Ground	4	Right Line In (monaural)

Some of Intel's motherboards, as well as those made by other motherboard manufacturers, use a single row of connectors for the front panel, as shown in Figure 4.50.

Figure 4.50 Typical ATX motherboard front panel connectors (Intel motherboard shown).

Table 4.56 shows the designations for the front-panel motherboard connectors used on some Intel ATX motherboards.

Table 4.56 ATX Motherboard Front Panel Connectors

Connector	Pin	Signal Name	Connector	Pin	Signal Name
Speaker	27	SPKR_HDR		13	HD_PWR +5V
	26	PIEZO_IN	none	12	No connect
	25	Key	IrDA	11	CONIR (Consumer IR)
	24	Ground		10	IrTX
Reset	23	SW_RST		9	Ground
	22	Ground		8	IrRX
none	21	No connect/Key		7	Key
Sleep/Power LED	20	PWR_LED		6	+5V
	19	Key	none	5	No connect
	18	Ground	Sleep/Resume	4	SLEEP_PU (pullup)
none	17	No connect/Key		3	SLEEP
Hard Drive LED	16	HD_PWR	Power On	2	Ground
	15	HD Active#		1	SW_ON#
	14	Key			

System Bus Types, Functions, and Features

The heart of any motherboard is the various buses that carry signals between the components. A *bus* is a common pathway across which data can travel within a computer. This pathway is used for communication and can be established between two or more computer elements.

The PC has a hierarchy of different buses. Most modern PCs have at least three buses; some have four or more. They are hierarchical because each slower bus is connected to the faster one above it. Each device in the system is connected to one of the buses, and some devices (primarily the chipset) act as bridges between the various buses.

The main buses in a modern system are as follows:

- *Processor bus.* Also called the front-side bus (FSB), this is the highest-speed bus in the system and is at the core of the chipset and motherboard. This bus is used primarily by the processor to pass information to and from cache or main memory and the North Bridge of the chipset. The processor bus in a modern system runs at 66MHz, 100MHz, 133MHz, 200MHz, 266MHz, 400MHz, 533MHz, or 800MHz and is normally 64 bits (8 bytes) wide.
- *AGP bus.* This is a high-speed 32-bit bus specifically for a video card. It runs at 66MHz (AGP 1x), 133MHz (AGP 2x), 266MHz (AGP 4x), or 533MHz (AGP 8x), which allows for a bandwidth of up to 2,133MBps. It is connected to the North Bridge or Memory Controller Hub of the chipset and is manifested as a single AGP slot in systems that support it.
- *PCI bus.* This is usually a 33MHz 32-bit bus found in virtually all newer 486 systems and Pentium and higher processor systems. Some newer systems include an optional 66MHz 64-bit version—mostly workstations or server-class systems. This bus is generated by either the chipset North Bridge in North/South Bridge chipsets or the I/O Controller Hub in chipsets using hub architecture. This bus is manifested in the system as a collection of 32-bit slots, normally white in color and numbering from four to six on most motherboards. High-speed peripherals, such as SCSI adapters, network cards, video cards, and more, can be plugged into PCI bus slots.

- *ISA bus.* This is an 8MHz 16-bit bus that has disappeared from recent systems after first appearing in the original PC in 8-bit, 5MHz form and in the 1984 IBM AT in full 16-bit 8MHz form. It is a very slow-speed bus, but it was ideal for certain slow-speed or older peripherals. It has been used in the past for plug-in modems, sound cards, and various other low-speed peripherals. The ISA bus is generated by the South Bridge part of the motherboard chipset, which acts as the ISA bus controller and the interface between the ISA bus and the faster PCI bus above it. The Super I/O chip usually was connected to the ISA bus on systems that included ISA slots.

Some newer motherboards feature a special connector called an Audio Modem Riser (AMR) or a Communications and Networking Riser (CNR). These are dedicated connectors for cards that are specific to the motherboard design to offer communications and networking options. They are *not* designed to be general-purpose bus interfaces, and few cards for these connectors are offered on the open market. Usually, they're offered only as an option with a given motherboard. They are designed such that a motherboard manufacturer can easily offer its boards in versions with and without communications options, without having to reserve space on the board for optional chips. Normal network and modem options offered publicly, for the most part, will still be PCI based because the AMR/CNR connection is somewhat motherboard specific.

Several hidden buses exist on modern motherboards—buses that don't manifest themselves in visible slots or connectors. I'm talking about buses designed to interface chipset components, such as the Hub Interface and the LPC bus. The Hub Interface is a quad-clocked (4x) 66MHz 8-bit bus that carries data between the MCH and ICH in hub architecture chipsets made by Intel. It operates at a bandwidth of 266MBps and was designed as a chipset component connection that is faster than PCI and yet uses fewer signals for a lower-cost design. Some recent workstation/server chipsets from Intel use faster versions of the hub interface. The most recent chipsets from major third-party vendors also bypass the PCI bus with direct high-speed connections between chipset components.

◀◀ See "High-Speed North-South Bridge Connections," p. 277.

In a similar fashion, the LPC bus is a 4-bit bus that has a maximum bandwidth of 6.67MBps; it was designed as an economical onboard replacement for the ISA bus. In systems that use LPC, it typically is used to connect Super I/O chip or motherboard ROM BIOS components to the main chipset. LPC is nearly as fast as ISA and yet uses far fewer pins and enables ISA to be eliminated from the board entirely.

The system chipset is the conductor that controls the orchestra of system components, enabling each to have its turn on its respective buses. Table 4.57 shows the widths, speeds, data cycles, and overall bandwidth of virtually all PC buses.

Table 4.57 Bandwidth (in MBps) and Detailed Comparison of Most PC Buses and Interfaces

Bus Type	Bus Width (Bits)	Bus Speed (MHz)	Data Cycles per Clock	Bandwidth (MBps)
8-bit ISA (PC/XT)	8	4.77	1/2	2.39
8-bit ISA (AT)	8	8.33	1/2	4.17
LPC bus	4	33	1	16.67
16-bit ISA (AT-Bus)	16	8.33	1/2	8.33

Table 4.57 Continued

Bus Type	Bus Width (Bits)	Bus Speed (MHz)	Data Cycles per Clock	Bandwidth (MBps)
DD Floppy Interface	1	0.25	1	0.03125
HD Floppy Interface	1	0.5	1	0.0625
ED Floppy Interface	1	1	1	0.125
EISA Bus	32	8.33	1	33
VL-Bus	32	33	1	133
MCA-16	16	5	1	10
MCA-32	32	5	1	20
MCA-16 Streaming	16	10	1	20
MCA-32 Streaming	32	10	1	40
MCA-64 Streaming	64	10	1	80
MCA-64 Streaming	64	20	1	160
PC-Card (PCMCIA)	16	10	1	20
CardBus	32	33	1	133
PCI	32	33	1	133
PCI 66MHz	32	66	1	266
PCI 64-bit	64	33	1	266
PCI 66MHz/64-bit	64	66	1	533
PCI-X 66	64	66	1	533
PCI-X 133	64	133	1	1,066
PCI-X 266	64	266	1	2,133
PCI-X 533	64	533	1	4,266
PCI Express 1.0 1-lane	1	2,500	0.8	250
PCI Express 1.0 32-lanes	32	2,500	0.8	8,000
Intel Hub Interface 8-bit	8	66	4	266
Intel Hub Interface 16-bit	16	66	4	533
AMD HyperTransport 2x2	2	200	2	100
AMD HyperTransport 4x2	4	200	2	200
AMD HyperTransport 8x2	8	200	2	400
AMD HyperTransport 16x2	16	200	2	800
AMD HyperTransport 32x2	32	200	2	1,600
AMD HyperTransport 2x4	2	400	2	200
AMD HyperTransport 4x4	4	400	2	400
AMD HyperTransport 8x4	8	400	2	800
AMD HyperTransport 16x4	16	400	2	1,600
AMD HyperTransport 32x4	32	400	2	3,200
AMD HyperTransport 2x8	2	800	2	400

Table 4.57 Continued

Bus Type	Bus Width (Bits)	Bus Speed (MHz)	Data Cycles per Clock	Bandwidth (MBps)
AMD HyperTransport 4x8	4	800	2	800
AMD HyperTransport 8x8	8	800	2	1,600
AMD HyperTransport 16x8	16	800	2	3,200
AMD HyperTransport 32x8	32	800	2	6,400
VIA V-Link 4x	8	66	4	266
VIA V-Link 8x	8	66	8	533
SiS MuTIOL	16	133	2	533
SiS MuTIOL 1G	16	266	2	1,066
AGP	32	66	1	266
AGP-2X	32	66	2	533
AGP-4X	32	66	4	1,066
AGP-8X	32	66	8	2,133
ATI A-Link	16	66	2	266
RS-232 Serial	1	0.1152	1/10	0.01152
RS-232 Serial HS	1	0.2304	1/10	0.02304
IEEE-1284 Parallel	8	8.33	1/6	1.38
IEEE-1284 EPP/ECP	8	8.33	1/3	2.77
USB 1.1	1	12	1	1.5
USB 2.0	2	240	1	60
IEEE-1394a S100	1	100	1	12.5
IEEE-1394a S200	1	200	1	25
IEEE-1394a S400	1	400	1	50
IEEE-1394b S800	1	800	1	100
IEEE-1394b S1600	1	1600	1	200
ATA PIO-4	16	8.33	1	16.67
ATA-UDMA/33	16	8.33	2	33
ATA-UDMA/66	16	16.67	2	66
ATA-UDMA/100	16	25	2	100
ATA-UDMA/133	16	33	2	133
SATA-150	1	750	2	150
SATA-300	1	1500	2	300
SATA-600	1	3000	2	600
SCSI	8	5	1	5
SCSI Wide	16	5	1	10
SCSI Fast	8	10	1	10
SCSI Fast/Wide	16	10	1	20

Table 4.57 Continued

Bus Type	Bus Width (Bits)	Bus Speed (MHz)	Data Cycles per Clock	Bandwidth (MBps)
SCSI Ultra	8	20	1	20
SCSI Ultra/Wide	16	20	1	40
SCSI Ultra2	8	40	1	40
SCSI Ultra2/Wide	16	40	1	80
SCSI Ultra3 (Ultra160)	16	40	2	160
SCSI Ultra4 (Ultra320)	16	80	2	320
FPM DRAM	64	22	1	177
EDO DRAM	64	33	1	266
PC66 SDRAM DIMM	64	66	1	533
PC100 SDRAM DIMM	64	100	1	800
PC133 SDRAM DIMM	64	133	1	1,066
PC1600 DDR DIMM (DDR200)	64	100	2	1,600
PC2100 DDR DIMM (DDR266)	64	133	2	2,133
PC2400 DDR DIMM (DDR300)	64	150	2	2,400
PC2700 DDR DIMM (DDR333)	64	167	2	2,666
PC3000 DDR DIMM (DDR366)	64	183	2	2,933
PC3200 DDR DIMM (DDR400)	64	200	2	3,200
PC3500 DDR (DDR433)	64	216	2	3,466
PC3700 DDR (DDR466)	64	233	2	3,733
PC4000 DDR (DDR500)	64	250	2	4,000
PC4300 DDR (DDR533)	64	267	2	4,266
PC2-3200 DDR2 (DDR2-433)	64	216	2	3,466
PC2-4300 DDR2 (DDR2-466)	64	233	2	3,733
PC2-5400 DDR2 (DDR2-500)	64	250	2	4,000
PC2-6400 DDR2 (DDR2-533)	64	267	2	4,266
RIMM1200 RDRAM (PC600)	16	300	2	1,200
RIMM1400 RDRAM (PC700)	16	350	2	1,400
RIMM1600 RDRAM (PC800)	16	400	2	1,600
RIMM2100 RDRAM (PC1066)	16	533	2	2,133
RIMM2400 RDRAM (PC1200)	16	600	2	2,400
RIMM3200 RDRAM (PC800)	32	400	2	3,200
RIMM4200 RDRAM (PC1066)	32	533	2	4,266
RIMM4800 RDRAM (PC1200)	32	600	2	4,800
RIMM6400 RDRAM (PC800)	64	400	2	6,400
RIMM8500 RDRAM (PC1066)	64	533	2	8,533
RIMM9600 RDRAM (PC1200)	64	600	2	9,600

Table 4.57 Continued

Bus Type	Bus Width (Bits)	Bus Speed (MHz)	Data Cycles per Clock	Bandwidth (MBps)
33MHz 486 FSB	32	33	1	133
66MHz Pentium I/II/III FSB	64	66	1	533
100MHz Pentium I/II/III FSB	64	100	1	800
133MHz Pentium I/II/III FSB	64	133	1	1,066
200MHz Athlon FSB	64	100	2	1,600
266MHz Athlon FSB	64	133	2	2,133
333MHz Athlon FSB	64	167	2	2,666
400MHz Athlon FSB	64	200	2	3,200
533MHz Athlon FSB	64	267	2	4,266
400MHz Pentium 4 FSB	64	100	4	3,200
533MHz Pentium 4 FSB	64	133	4	4,266
800MHz Pentium 4 FSB	64	200	4	6,400
266MHz Itanium FSB	64	133	2	2,133
400MHz Itanium 2 FSB	128	100	4	6,400

Note: ISA, EISA, VL-Bus, and MCA are no longer used in current motherboard designs.

MBps = Megabytes per second

ISA = Industry Standard Architecture, also known as the PC/XT (8-bit) or AT-Bus (16-bit)

LPC = Low Pin Count bus

DD Floppy = Double Density (360/720KB) Floppy

HD Floppy = High Density (1.2/1.44MB) Floppy

ED Floppy = Extra-high Density (2.88MB) Floppy

EISA = Extended Industry Standard Architecture (32-bit ISA)

VL-Bus = VESA (Video Electronics Standards Association) Local Bus (ISA extension)

MCA = MicroChannel Architecture (IBM PS/2 systems)

PC-Card = 16-bit PCMCIA (Personal Computer Memory Card International Association) interface

CardBus = 32-bit PC-Card

Hub Interface = Intel 8xx chipset bus

HyperTransport = AMD chipset bus

V-Link = VIA Technologies chipset bus

MuTIOL = Silicon Integrated System chipset bus

PCI = Peripheral Component Interconnect

AGP = Accelerated Graphics Port

RS-232 = Standard Serial port, 115.2Kbps

RS-232 HS = High Speed Serial port, 230.4Kbps

IEEE-1284 Parallel = Standard Bidirectional Parallel Port

IEEE-1284 EPP/ECP = Enhanced Parallel Port/Extended Capabilities Port

USB = Universal serial bus

IEEE-1394 = FireWire, also called i.Link

ATA PIO = AT Attachment (also known as IDE) Programmed I/O

ATA-UDMA = AT Attachment Ultra DMA

SCSI = Small computer system interface

FPM = Fast Page Mode, based on X-3-3-3 (1/3 max) burst mode timing on a 66MHz bus

EDO = Extended Data Out, based on X-2-2-2 (1/2 max) burst mode timing on a 66MHz bus

SDRAM = Synchronous dynamic RAM

RDRAM = Rambus dynamic RAM

DDR = Double data rate SDRAM

DDR2 = Next-generation DDR

CPU FSB = Processor front-side bus

Note that many of the buses use multiple data cycles (transfers) per clock cycle to achieve greater performance. Therefore, the data transfer rate is higher than it would seem for a given clock rate, which allows for an easy way to take an existing bus and make it go faster in a backward-compatible way.

The following sections discuss the processor and other subset buses in the system and the main I/O buses mentioned in the previous table.

The Processor Bus (Front-Side Bus)

The processor bus (also called the front-side bus or FSB) is the communication pathway between the CPU and motherboard chipset, more specifically the North Bridge or Memory Controller Hub. This bus runs at the full motherboard speed—typically between 66MHz and 800MHz in modern systems, depending on the particular board and chipset design. This same bus also transfers data between the CPU and an external (L2) memory cache on Socket-7 (Pentium class) systems. Figure 4.51 shows how this bus fits into a typical Socket 7 PC system.

Figure 4.51 also shows where and how the other main buses, such as the PCI and ISA buses, fit into the system. As you can see, there is clearly a three-tier architecture with the fastest CPU bus on top, the PCI bus next, and the ISA bus at the bottom. Various components in the system are connected to one of these three main buses.

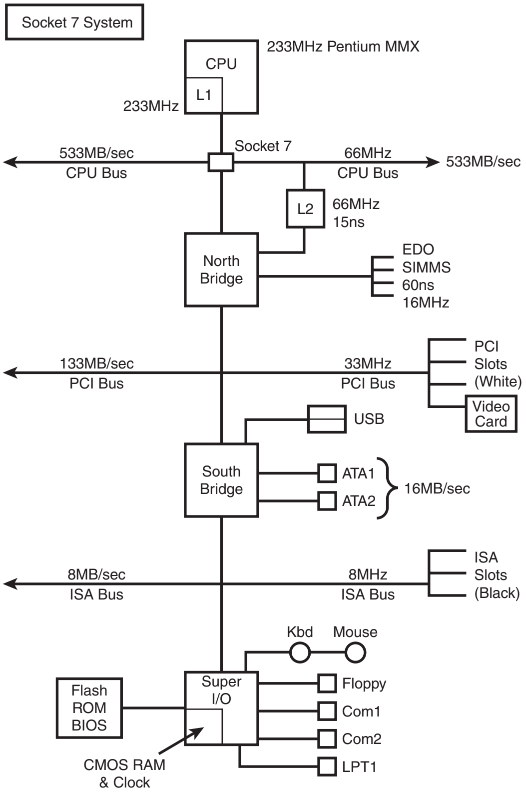

Figure 4.51 Typical Socket 7 (Pentium class) system architecture.

Socket 7 systems have an external (L2) cache for the CPU; the L2 cache is mounted on the motherboard and connected to the main processor bus that runs at the motherboard speed (usually between 66MHz and 100MHz). Thus, as the Socket 7 processors became available in faster and faster versions (through increasing the clock multiplier in the chip), the L2 cache unfortunately remained stuck on the motherboard running at the relatively slow (by comparison) motherboard speed. For example, the fastest Intel Socket 7 systems ran the CPU at 233MHz, which was 3.5x the CPU bus speed of 66MHz. Therefore, the L2 cache ran at only 66MHz. The fastest Socket 7 systems used the AMD K6-2 550 processor, which ran at 550MHz—5.5x a CPU bus speed of 100MHz. In those systems, the L2 cache ran at only 100MHz.

The problem of the slow L2 cache was first solved in the P6 class processors, such as the Pentium Pro, Pentium II, Celeron, Pentium III, and AMD Athlon and Duron. These processors used either Socket 8, Slot 1, Slot 2, Slot A, Socket A, or Socket 370. They moved the L2 cache off the motherboard and directly onto the CPU and connected it to the CPU via an on-chip back-side bus. Because the L2 cache bus was called the back-side bus, some in the industry began calling the main CPU bus the front-side bus. I still usually refer to it simply as the CPU bus.

By incorporating the L2 cache into the CPU, it can run at speeds up to the same speed as the processor itself. Most processors now incorporate the L2 cache directly on the CPU die, so the L2 cache runs at the same speed as the rest of the CPU. Others (mostly older versions) used separate die for the cache integrated into the CPU package, which ran the L2 cache at some lower multiple (one-half, two-fifth, or one-third) of the main CPU. Even if the L2 ran at half or one-third of the processor speed, it still was significantly faster than the motherboard-bound cache on the Socket 7 systems.

Figure 4.52 shows a typical Slot-1 type system, in which the L2 cache is built in to the CPU but running at only half the processor speed. This would also be the same for systems using Slot A. The CPU bus speed increased from 66MHz (used primarily in Socket 7 systems) to 100MHz, enabling a bandwidth of 800MBps. Note that most of these systems included AGP support. Basic AGP was 66MHz (twice the speed of PCI), but most of these systems incorporated AGP 2x, which operated at twice the speed of standard AGP and enabled a bandwidth of 533 MBps. These systems also typically used PC-100 SDRAM DIMMs, which have a bandwidth of 800MBps, matching the processor bus bandwidth for the best performance.

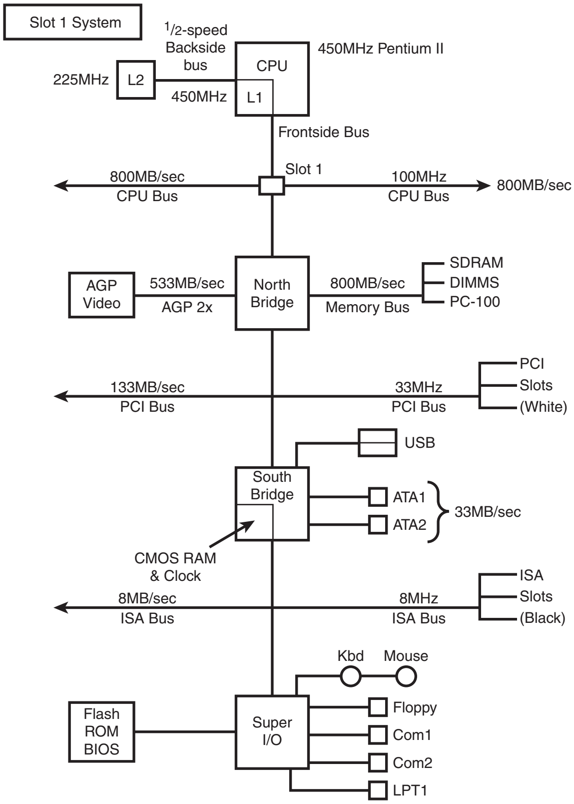

Figure 4.52 Typical Slot-1 (Pentium II class) system architecture.

Slot 1 was dropped in favor of Socket 370 for the Pentium III and Celeron systems. This was mainly because these newer processors incorporated the L2 cache directly into the CPU die (running at the full-core speed of the processor) and an expensive cartridge with multiple chips was no longer necessary. At the same time, processor bus speeds increased to 133MHz, which enabled a throughput of 1,066MBps. Figure 4.53 shows a typical Socket 370 system design. AGP speed was also increased to AGP 4x, with a bandwidth of 1,066MBps.

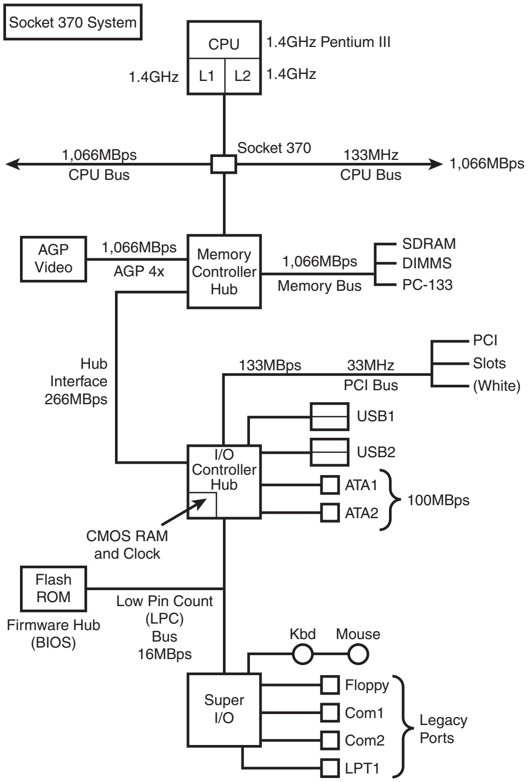

Figure 4.53 Typical Socket 370 (Pentium III/Celeron class) system architecture.

Note the use of what Intel calls "hub architecture" instead of the older North/South Bridge design. This moves the main connection between the chipset components to a separate 266MBps hub interface (which has twice the throughput of PCI) and enables PCI devices to use the full bandwidth of PCI without fighting for bandwidth with a South Bridge. Also note that the flash ROM BIOS chip is now referred to as a Firmware Hub and is connected to the system via the LPC bus instead of via the Super I/O chip as in older North/South Bridge designs. The ISA bus is no longer used in most of these systems, and the Super I/O is connected via the LPC bus instead of ISA. The Super I/O chip also can easily be eliminated in these designs. This is commonly referred to as a *legacy-free* system because the ports supplied by the Super I/O chip are now known as *legacy* ports. Devices that would have used legacy ports must then be connected to the system via USB instead, and such systems would feature two USB controllers, with up to four total ports (more can be added by attaching USB hubs).

AMD processor systems adopted a Socket A design, which is similar to Socket 370 except it uses faster processor and memory buses. Although early versions retained the older North/South Bridge design, more recent versions use a design similar to Intel's hub architecture (see Figure 4.54). Note the high-speed CPU bus running up to 333MHz (2,664MBps throughput) and the use of DDR SDRAM DIMM modules that support a matching bandwidth of 2,664MBps. It is always best for performance when the bandwidth of memory matches that of the processor. Finally, note how most of the South Bridge components include functions otherwise found in Super I/O chips; when these functions are included the chip is called a Super South Bridge.

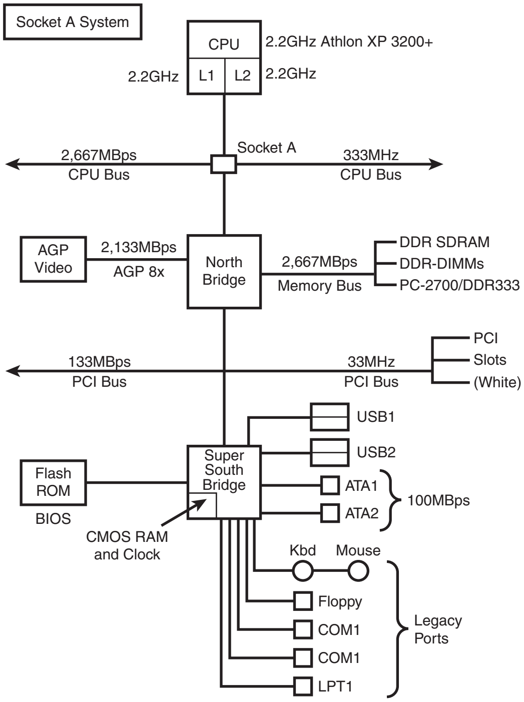

Figure 4.54 Typical Socket A (Athlon/Duron/Athlon XP) system architecture.

The Pentium 4 uses a Socket 423 or Socket 478 design with hub architecture (see Figure 4.55). This design is most notable for including a 400MHz, 533MHz, or 800MHz CPU bus with a bandwidth of 3,200MBps, 4,266MBps, or 6,400MBps. The 533MHz and 800MHz models are currently faster than anything else on the market. In this example, note the use of dual-channel PC3200 (DDR400) SDRAM. A single PC-3200 DIMM has a bandwidth of 3,200MBps, but when running dual-channel (identical pairs of memory) memory, it has a bandwidth of 6,400MBps—which matches the bandwidth of the 800MHz CPU bus models of the Pentium 4 for best performance. Processors with the 533MHz CPU bus can use pairs of PC2100 (DDR266) or PC2700 (DDR333) memory modules in dual channel mode to match the 4,266MBps throughput of this memory bus. It is always best when the throughput of the memory bus matches that of the processor bus.

Because the purpose of the processor bus is to get information to and from the CPU at the fastest possible speed, this bus typically operates at a rate faster than any other bus in the system. The bus consists of electrical circuits for data, addresses (the address bus, which is discussed in the following section), and control purposes. Most processors since the original Pentium have a 64-bit data bus, so they transfer 64 bits (8 bytes) at a time over the CPU bus.

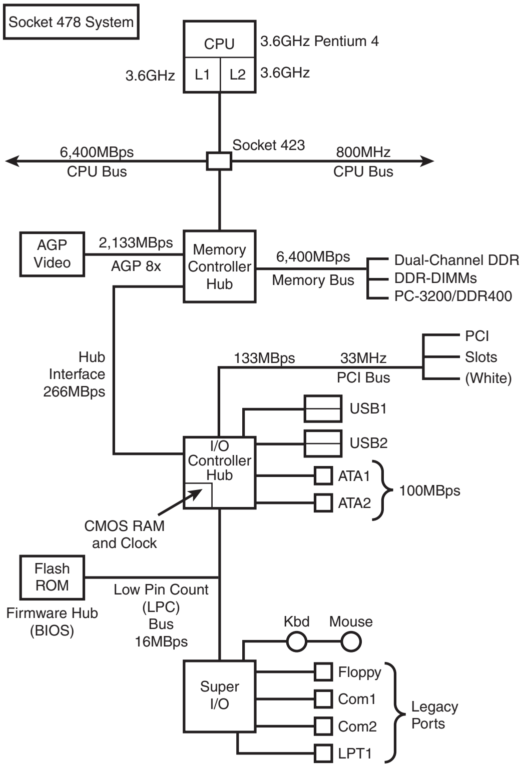

Figure 4.55 Typical Socket 478 (Pentium 4) system architecture.

The processor bus operates at the same base clock rate as the CPU does externally. This can be misleading because most CPUs these days run at a higher clock rate internally than they do externally. For example, an AMD Athlon 3200+ system has a processor running at 2.2GHz internally but only 333MHz externally, whereas a Pentium 4 3.2GHz runs at 3.2GHz internally but only 800MHz externally. In newer systems, the actual processor speed is some multiple (2x, 2.5x, 3x, and higher) of the processor bus.

◀◀ See "Processor Speed Ratings," p. 50.

The processor bus is tied to the external processor pin connections and can transfer 1 bit of data per data line every cycle. Most modern processors transfer 64 bits (8 bytes) of data at a time.

To determine the transfer rate for the processor bus, you multiply the data width (64 bits or 8 bytes for a Celeron/Pentium III/4 or Athlon/Duron/Athlon XP) by the clock speed of the bus (the same as the base or unmultiplied clock speed of the CPU).

For example, if you are using a Pentium 4 3.6GHz processor that runs on an 800MHz processor bus, you have a maximum instantaneous transfer rate of roughly 6,400MBps. You get this result by using the following formula:

800MHz × 8 bytes (64 bits) = 6,400MBps

With slower versions of the Pentium 4, you get either

533.33MHz × 8 bytes (64 bits) = 4,266MBps

or

400MHz × 8 bytes (64 bits) = 3,200MBps

With Socket A (Athlon XP), you get

333.33MHz × 8 bytes (64 bits) = 2,667MBps

or

266.66MHz × 8 bytes (64 bits) = 2,133MBps

or

200MHz × 8 bytes (64 bits) = 1,600MBps

With Socket 370 (Pentium III), you get

133.33MHz × 8 bytes (64 bits) = 1,066MBps

or

100MHz × 8 bytes (64 bits) = 800MBps

This transfer rate, often called the *bandwidth* of the processor bus, represents the maximum speed at which data can move. Refer to Table 4.58 for a more complete list of various processor bus bandwidths.

The Memory Bus

The memory bus is used to transfer information between the CPU and main memory—the RAM in your system. This bus is connected to the motherboard chipset North Bridge or Memory Controller Hub chip. Depending on the type of memory your chipset (and therefore motherboard) is designed to handle, the North Bridge runs the memory bus at various speeds. The best solution is if the memory bus runs at the same speed as the processor bus. Systems that use PC133 SDRAM have a memory bandwidth of 1,066MBps, which is the same as the 133MHz CPU bus. In another example, Athlon systems running a 266MHz processor bus also run PC2100 DDR-SDRAM, which has a bandwidth of 2,133MBps—exactly the same as the processor bus in those systems. In addition, systems running a Pentium 4 with its 400MHz processor bus also use dual-channel RDRAM memory, which runs 1,600MBps for each channel, or a combined bandwidth (both memory channels run simultaneously) of 3,200MBps, which is exactly the same as the Pentium 4 CPU bus. Pentium 4 systems with the 533MHz bus run dual-channel DDR PC2100 or PC2700 modules, which match or exceed the throughput of the 4,266MBps processor bus.

Running memory at the same speed as the processor bus negates the need for having cache memory on the motherboard. That is why when the L2 cache moved into the processor, nobody added an L3 cache to the motherboard. Some very high-end processors, such as the Itanium and Itanium 2, have integrated 2MB–4MB of full-core speed L3 cache into the CPU. Eventually, this should make it down to more mainstream desktop systems.

Note

Notice that the main memory bus must transfer data in the same width as the processor bus. This defines the size of what is called a *bank* of memory, at least when dealing with anything but RDRAM. Memory banks and their widths relative to processor buses are discussed in the section "Memory Banks" in Chapter 6.

The Need for Expansion Slots

The I/O bus or expansion slots enable your CPU to communicate with peripheral devices. The bus and its associated expansion slots are needed because basic systems can't possibly satisfy all the needs of all the people who buy them. The I/O bus enables you to add devices to your computer to expand its capabilities. The most basic computer components, such as sound cards and video cards, can be plugged into expansion slots; you also can plug in more specialized devices, such as network interface cards, SCSI host adapters, and others.

Note

In most modern PC systems, a variety of basic peripheral devices are built in to the motherboard. Most systems today have at least dual (primary and secondary) IDE interfaces, four USB ports, a floppy controller, two serial ports, a parallel port, keyboard, and mouse controller built directly into the motherboard. These devices are usually distributed between the motherboard chipset South Bridge and the Super I/O chip. (Super I/O chips are discussed earlier in this chapter.)

Many add even more items, such as a built-in sound card, video adapter, SCSI host adapter, network interface or IEEE-1394a port, that also are built in to the motherboard. Those items, however, might not be built in to the motherboard chipset or Super I/O chip; they are sometimes configured as additional chips installed on the board. Nevertheless, these built-in controllers and ports still use the I/O bus to communicate with the CPU. In essence, even though they are built in, they act as if they were cards plugged into the system's bus slots, including using system resources in the same manner.

Types of I/O Buses

Since the introduction of the first PC, many I/O buses have been introduced. The reason is simple: Faster I/O speeds are necessary for better system performance. This need for higher performance involves three main areas:

- Faster CPUs
- Increasing software demands
- Greater multimedia requirements

Each of these areas requires the I/O bus to be as fast as possible.

One of the primary reasons new I/O bus structures have been slow in coming is compatibility—that old catch-22 that anchors much of the PC industry to the past. One of the hallmarks of the PC's success is its standardization. This standardization spawned thousands of third-party I/O cards, each originally built for the early bus specifications of the PC. If a new high-performance bus system was introduced, it often had to be compatible with the older bus systems so the older I/O cards would not be obsolete. Therefore, bus technologies seem to evolve rather than make quantum leaps forward.

You can identify different types of I/O buses by their architectures. The main types of I/O buses are detailed earlier in this chapter.

The main differences among buses consist primarily of the amounts of data they can transfer at one time and the speeds at which they can do it. The following sections describe the various types of PC buses.

The ISA Bus

Industry Standard Architecture (ISA) is the bus architecture that was introduced as an 8-bit bus with the original IBM PC in 1981; it was later expanded to 16 bits with the IBM PC/AT in 1984. ISA is the basis of the modern personal computer and the primary architecture used in the vast majority of PC systems on the market today. It might seem amazing that such a presumably antiquated architecture is used in today's high-performance systems, but this is true for reasons of reliability, affordability, and compatibility, plus this old bus is still faster than many of the peripherals we connect to it!

Note

The ISA bus has vanished from all recent desktop systems, and few companies make or sell ISA cards anymore. The ISA bus continues to be popular with industrial computer (PICMG) designs, but it is expected to eventually fade away from these as well.

Two versions of the ISA bus exist, based on the number of data bits that can be transferred on the bus at a time. The older version is an 8-bit bus; the newer version is a 16-bit bus. The original 8-bit version ran at 4.77MHz in the PC and XT, and the 16-bit version used in the AT ran at 6MHz and then

8MHz. Later, the industry as a whole agreed on an 8.33MHz maximum standard speed for 8/16-bit versions of the ISA bus for backward-compatibility. Some systems have the capability to run the ISA bus faster than this, but some adapter cards will not function properly at higher speeds. ISA data transfers require anywhere from two to eight cycles. Therefore, the theoretical maximum data rate of the ISA bus is about 8MBps, as the following formula shows:

8.33MHz × 2 bytes (16 bits) ÷ 2 cycles per transfer = 8.33MBps

The bandwidth of the 8-bit bus would be half this figure (4.17MBps). Remember, however, that these figures are theoretical maximums. Because of I/O bus protocols, the effective bandwidth is much lower—typically by almost half. Even so, at about 8MBps, the ISA bus is still faster than many of the peripherals connected to it, such as serial ports, parallel ports, floppy controllers, keyboard controllers, and so on.

The 8-Bit ISA Bus

This bus architecture is used in the original IBM PC computers and was retained for several years in later systems. Although virtually nonexistent in new systems today, this architecture still exists in hundreds of thousands of PC systems in the field.

Physically, the 8-bit ISA expansion slot resembles the tongue-and-groove system furniture makers once used to hold two pieces of wood together. It is specifically called a *card/edge connector*. An adapter card with 62 contacts on its bottom edge plugs into a slot on the motherboard that has 62 matching contacts. Electronically, this slot provides 8 data lines and 20 addressing lines, enabling the slot to handle 1MB of memory.

Figure 4.56 describes the pinouts for the 8-bit ISA bus; Figure 4.57 shows how these pins are oriented in the expansion slot.

Signal	Pin	Pin	Signal
Ground	B1	A1	-I/O CH CHK
RESET DRV	B2	A2	Data Bit 7
+5 Vdc	B3	A3	Data Bit 6
IRQ 2	B4	A4	Data Bit 5
-5 Vdc	B5	A5	Data Bit 4
DRQ 2	B6	A6	Data Bit 3
-12 Vdc	B7	A7	Data Bit 2
-CARD SLCTD	B8	A8	Data Bit 1
+12 Vdc	B9	A9	Data Bit 0
Ground	B10	A10	-I/O CH RDY
-SMEMW	B11	A11	AEN
-SMEMR	B12	A12	Address 19
-IOW	B13	A13	Address 18
-IOR	B14	A14	Address 17
-DACK 3	B15	A15	Address 16
DRQ 3	B16	A16	Address 15
-DACK 1	B17	A17	Address 14
DRQ 1	B18	A18	Address 13
-Refresh	B19	A19	Address 12
CLK(4.77MHz)	B20	A20	Address 11
IRQ 7	B21	A21	Address 10
IRQ 6	B22	A22	Address 9
IRQ 5	B23	A23	Address 8
IRQ 4	B24	A24	Address 7
IRQ 3	B25	A25	Address 6
-DACK 2	B26	A26	Address 5
T/C	B27	A27	Address 4
BALE	B28	A28	Address 3
+5 Vdc	B29	A29	Address 2
OSC(14.3MHz)	B30	A30	Address 1
Ground	B31	A31	Address 0

Figure 4.56 Pinouts for the 8-bit ISA bus.

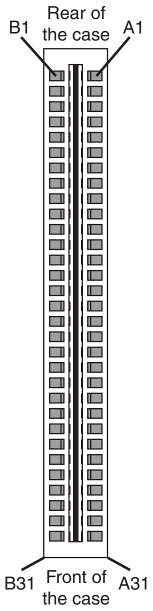

Figure 4.57 The 8-bit ISA bus connector.

Although the design of the bus is simple, IBM waited until 1987 to publish full specifications for the timings of the data and address lines, so in the early days of PC compatibles, manufacturers had to do their best to figure out how to make adapter boards. This problem was solved, however, as PC-compatible personal computers became more widely accepted as the industry standard and manufacturers had more time and incentive to build adapter boards that worked correctly with the bus.

The dimensions of 8-bit ISA adapter cards are as follows:

4.2" (106.68mm) high

13.13" (333.5mm) long

0.5" (12.7mm) wide

The 16-Bit ISA Bus

IBM threw a bombshell on the PC world when it introduced the AT with the 286 processor in 1984. This processor had a 16-bit data bus, which meant communications between the processor and motherboard as well as memory would now be 16 bits wide instead of only 8. Although this processor could have been installed on a motherboard with only an 8-bit I/O bus, that would have meant a huge sacrifice in the performance of any adapter cards or other devices installed on the bus.

Rather than create a new I/O bus, at that time IBM instead came up with a system that could support both 8- and 16-bit cards by retaining the same basic 8-bit connector layout but adding an optional 16-bit extension connector. This first debuted on the PC/AT in August 1984, which is why we also refer to the ISA bus as the *AT-bus*.

The extension connector in each 16-bit expansion slot adds 36 connector pins (for a total of 96 signals) to carry the extra signals necessary to implement the wider data path. In addition, two of the pins in the 8-bit portion of the connector were changed. These two minor changes did not alter the function of 8-bit cards.

Figure 4.58 describes the pinouts for the full 16-bit ISA expansion slot, and Figure 4.59 shows how the additional pins are oriented in the expansion slot.

Because of physical interference with some ancient 8-bit card designs, IBM left 16-bit extension connectors off two of the slots in the AT. This was not a problem in newer systems, so any system with ISA slots would have all of them as full 16-bit versions.

The dimensions of a typical AT expansion board are as follows:

4.8" (121.92mm) high

13.13" (333.5mm) long

0.5" (12.7mm) wide

Two heights actually are available for cards commonly used in AT systems: 4.8" and 4.2" (the height of older PC-XT cards). The shorter cards became an issue when IBM introduced the XT Model 286. Because this model has an AT motherboard in an XT case, it needs AT-type boards with the 4.2" maximum height. Most board makers trimmed the height of their boards; most manufacturers who still make ISA cards now make only 4.2" tall (or less) boards so they will work in systems with either profile.

32-Bit Buses

After 32-bit CPUs became available, it was some time before 32-bit bus standards became available. Before MCA and EISA specs were released, some vendors began creating their own proprietary 32-bit buses, which were extensions of the ISA bus. Fortunately, these proprietary buses were few and far between.

The expanded portions of the bus typically are used for proprietary memory expansion or video cards. Because the systems are proprietary (meaning that they are nonstandard), pinouts and specifications are not available.

Signal	Pin	Pin	Signal
Ground	B1	A1	-I/O CH CHK
RESET DRV	B2	A2	Data Bit 7
+5 Vdc	B3	A3	Data Bit 6
IRQ 9	B4	A4	Data Bit 5
-5 Vdc	B5	A5	Data Bit 4
DRQ 2	B6	A6	Data Bit 3
-12 Vdc	B7	A7	Data Bit 2
-0 WAIT	B8	A8	Data Bit 1
+12 Vdc	B9	A9	Data Bit 0
Ground	B10	A10	-I/O CH RDY
-SMEMW	B11	A11	AEN
-SMEMR	B12	A12	Address 19
-IOW	B13	A13	Address 18
-IOR	B14	A14	Address 17
-DACK 3	B15	A15	Address 16
DRQ 3	B16	A16	Address 15
-DACK 1	B17	A17	Address 14
DRQ 1	B18	A18	Address 13
-Refresh	B19	A19	Address 12
CLK(8.33MHz)	B20	A20	Address 11
IRQ 7	B21	A21	Address 10
IRQ 6	B22	A22	Address 9
IRQ 5	B23	A23	Address 8
IRQ 4	B24	A24	Address 7
IRQ 3	B25	A25	Address 6
-DACK 2	B26	A26	Address 5
T/C	B27	A27	Address 4
BALE	B28	A28	Address 3
+5 Vdc	B29	A29	Address 2
OSC(14.3MHz)	B30	A30	Address 1
Ground	B31	A31	Address 0
-MEM CS16	D1	C1	-SBHE
-I/O CS16	D2	C2	Latch Address 23
IRQ 10	D3	C3	Latch Address 22
IRQ 11	D4	C4	Latch Address 21
IRQ 12	D5	C5	Latch Address 20
IRQ 15	D6	C6	Latch Address 19
IRQ 14	D7	C7	Latch Address 18
-DACK 0	D8	C8	Latch Address 17
DRQ 0	D9	C9	-MEMR
-DACK 5	D10	C10	-MEMW
DRQ5	D11	C11	Data Bit 8
-DACK 6	D12	C12	Data Bit 9
DRQ 6	D13	C13	Data Bit 10
-DACK 7	D14	C14	Data Bit 11
DRQ 7	D15	C15	Data Bit 12
+5 Vdc	D16	C16	Data Bit 13
-Master	D17	C17	Data Bit 14
Ground	D18	C18	Data Bit 15

Figure 4.58 Pinouts for the 16-bit ISA bus.

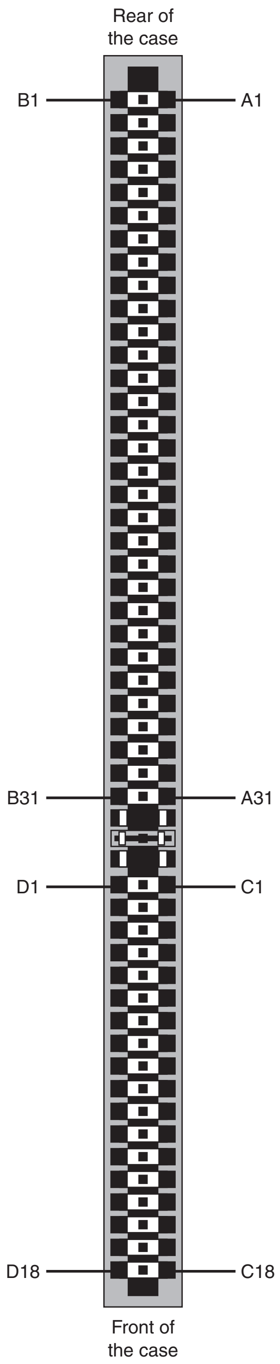

Figure 4.59 The ISA 16-bit bus connector.

The Micro Channel Bus

The introduction of 32-bit chips meant that the ISA bus could not handle the power of another new generation of CPUs. The 386DX chips could transfer 32 bits of data at a time, but the ISA bus can handle a maximum of only 16 bits. Rather than extend the ISA bus again, IBM decided to build a new bus; the result was the MCA bus. *MCA* (an abbreviation for microchannel architecture) is completely different from the ISA bus and is technically superior in every way.

IBM wanted not only to replace the old ISA standard, but also to require vendors to license certain parts of the technology. Many owed for licenses on the ISA bus technology that IBM also created, but because IBM had not been aggressive in its licensing of ISA, many got away without any license. Problems with licensing and control led to the development of the competing EISA bus (see the next section on the EISA bus) and hindered acceptance of the MCA bus.

MCA systems produced a new level of ease of use; they were plug-and-play before the official Plug and Play specification even existed. An MCA system had no jumpers and switches—neither on the motherboard nor on any expansion adapter. Instead you used a special Reference disk, which went with the particular system, and Option disks, which went with each of the cards installed in the system. After a card was installed, you loaded the Option disk files onto the Reference disk; after that, you didn't need the Option disks anymore. The Reference disk contained the special BIOS and system setup program necessary for an MCA system, and the system couldn't be configured without it. To support older PS/2 systems, IBM maintains a library of all its Reference and Options disks at `ftp://ftp.pc.ibm.com/pub/pccbbs`. Check this site if you are supporting any old MCA-based systems and need any of these files.

For more information on the MCA bus, see the previous editions of this book on the included DVD-ROM.

The EISA Bus

The Extended Industry Standard Architecture (EISA) standard was announced in September 1988 as a response to IBM's introduction of the MCA bus—more specifically, to the way IBM wanted to handle licensing of the MCA bus. Vendors did not feel obligated to pay retroactive royalties on the ISA bus, so they turned their backs on IBM and created their own buses.

The EISA standard was developed primarily by Compaq and was intended to be its way of taking over future development of the PC bus from IBM. Compaq knew that nobody would clone its bus if it was the only company that had it, so it essentially gave the design to other leading manufacturers. Compaq formed the EISA committee, a nonprofit organization designed specifically to control development of the EISA bus. Very few EISA adapters were ever developed. Those that were developed centered mainly around disk array controllers and server-type network cards.

The EISA bus was essentially a 32-bit version of ISA. Unlike the MCA bus from IBM, you could still use older 8-bit or 16-bit ISA cards in 32-bit EISA slots, providing for full backward-compatibility. As with MCA, EISA also allowed for automatic configuration of EISA cards via software.

The EISA bus added 90 new connections (55 new signals plus grounds) without increasing the physical connector size of the 16-bit ISA bus. At first glance, the 32-bit EISA slot looks a lot like the 16-bit ISA slot. The EISA adapter, however, has two rows of stacked contacts. The first row is the same type used in 16-bit ISA cards; the other, thinner row extends from the 16-bit connectors. Therefore, ISA cards can still be used in EISA bus slots. Although this compatibility was not enough to ensure the popularity of EISA buses, it is a feature that was carried over into the VL-Bus standard that followed. The physical specifications of an EISA card are as follows:

- 5" (127mm) high
- 13.13" (333.5mm) long
- 0.5" (12.7mm) wide

The EISA bus can handle up to 32 bits of data at an 8.33MHz cycle rate. Most data transfers require a minimum of two cycles, although faster cycle rates are possible if an adapter card provides tight timing specifications. The maximum bandwidth on the bus is 33MBps, as the following formula shows:

$$8.33\text{MHz} \times 4 \text{ bytes (32 bits)} = 33\text{MBps}$$

Figure 4.60 describes the pinouts for the EISA bus. Figure 4.61 shows the locations of the pins; note how some pins are offset to allow the EISA slot to accept ISA cards. Figure 4.62 shows the card connector for the EISA expansion slot.

Lower Signal	Upper Signal	Pin	Pin	Upper Signal	Lower Signal
Ground	Ground	B1	A1	-I/O CH CHK	-CMD
+5 Vdc	RESET DRV	B2	A2	Data Bit 7	-START
+5 Vdc	+5 Vdc	B3	A3	Data Bit 6	EXRDY
Reserved	IRQ 9	B4	A4	Data Bit 5	-EX32
Reserved	-5 Vdc	B5	A5	Data Bit 4	Ground
KEY	DRQ 2	B6	A6	Data Bit 3	KEY
Reserved	-12 Vdc	B7	A7	Data Bit 2	-EX16
Reserved	-0 WAIT	B8	A8	Data Bit 1	-SLBURST
+12 Vdc	+12 Vdc	B9	A9	Data Bit 0	-MSBURST
M-IO	Ground	B10	A10	-I/O CH RDY	W-R
-LOCK	-SMEMW	B11	A11	AEN	Ground
Reserved	-SMEMR	B12	A12	Address 19	Reserved
Ground	-IOW	B13	A13	Address 18	Reserved
Reserved	-IOR	B14	A14	Address 17	Reserved
-BE 3	-DACK 3	B15	A15	Address 16	Ground
KEY	DRQ 3	B16	A16	Address 15	KEY
-BE 2	-DACK 1	B17	A17	Address 14	-BE 1
-BE 0	DRQ 1	B18	A18	Address 13	Latch Address 31
Ground	-Refresh	B19	A19	Address 12	Ground
+5 Vdc	CLK(8.33MHz)	B20	A20	Address 11	-Latch Address 30
Latch Address 29	IRQ 7	B21	A21	Address 10	-Latch Address 28
Ground	IRQ 6	B22	A22	Address 9	-Latch Address 27
Latch Address 26	IRQ 5	B23	A23	Address 8	-Latch Address 25
Latch Address 24	IRQ 4	B24	A24	Address 7	Ground
KEY	IRQ 3	B25	A25	Address 6	KEY
Latch Address 16	-DACK 2	B26	A26	Address 5	Latch Address 15
Latch Address 14	T/C	B27	A27	Address 4	Latch Address 13
+5 Vdc	BALE	B28	A28	Address 3	Latch Address 12
+5 Vdc	+5 Vdc	B29	A29	Address 2	Latch Address 11
Ground	OSC(14.3MHz)	B30	A30	Address 1	Ground
Latch Address 10	Ground	B31	A31	Address 0	Latch Address 9
Latch Address 8	-MEM CS16	D1	C1	-SBHE	Latch Address 7
Latch Address 6	-I/O CS16	D2	C2	Latch Address 23	Ground
Latch Address 5	IRQ 10	D3	C3	Latch Address 22	Latch Address 4
+5 Vdc	IRQ 11	D4	C4	Latch Address 21	Latch Address 3
Latch Address 4	IRQ 12	D5	C5	Latch Address 20	Ground
KEY	IRQ 15	D6	C6	Latch Address 19	KEY
Data Bit 16	IRQ 14	D7	C7	Latch Address 18	Data Bit 17
Data Bit 18	-DACK 0	D8	C8	Latch Address 17	Data Bit 19
Ground	DRQ 0	D9	C9	-MEMR	Data Bit 20
Data Bit 21	-DACK 5	D10	C10	-MEMW	Data Bit 22
Data Bit 23	DRQ5	D11	C11	Data Bit 8	Ground
Data Bit 24	-DACK 6	D12	C12	Data Bit 9	Data Bit 25
Ground	DRQ 6	D13	C13	Data Bit 10	Data Bit 26
Data Bit 27	-DACK 7	D14	C14	Data Bit 11	Data Bit 28
KEY	DRQ 7	D15	C15	Data Bit 12	KEY
Data Bit 29	+5 Vdc	D16	C16	Data Bit 13	Ground
+5 Vdc	-Master	D17	C17	Data Bit 14	Data Bit 30
+5 Vdc	Ground	D18	C18	Data Bit 15	Data Bit 31
-MAKx		D19	C19		-MREQx

Figure 4.60 Pinouts for the EISA bus.

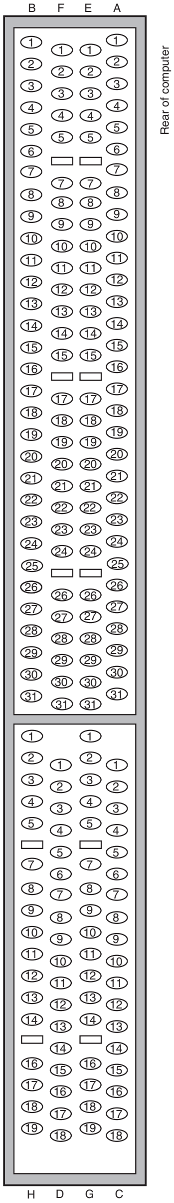

Figure 4.61 Pin locations inside the EISA bus connector.

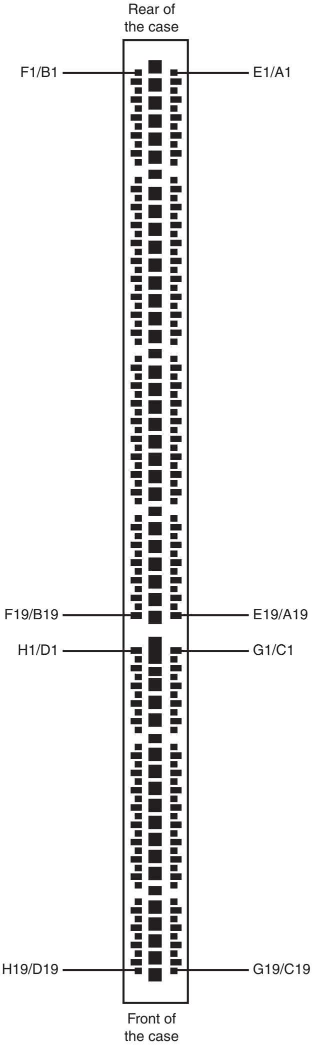

Figure 4.62 The EISA bus connector.

Local Buses

The I/O buses discussed so far (ISA, MCA, and EISA) have one thing in common: relatively slow speed. The next three bus types that are discussed in the following few sections all use the *local bus* concept explained in this section to address the speed issue. The three main local buses found in PC systems are

- VL-Bus (VESA local bus)
- PCI
- AGP

The speed limitation of ISA, MCA, and EISA is a carryover from the days of the original PC when the I/O bus operated at the same speed as the processor bus. As the speed of the processor bus increased, the I/O bus realized only nominal speed improvements, primarily from an increase in the bandwidth of the bus. The I/O bus had to remain at a slower speed because the huge installed base of adapter cards could operate only at slower speeds.

Figure 4.63 shows a conceptual block diagram of the buses in a computer system.

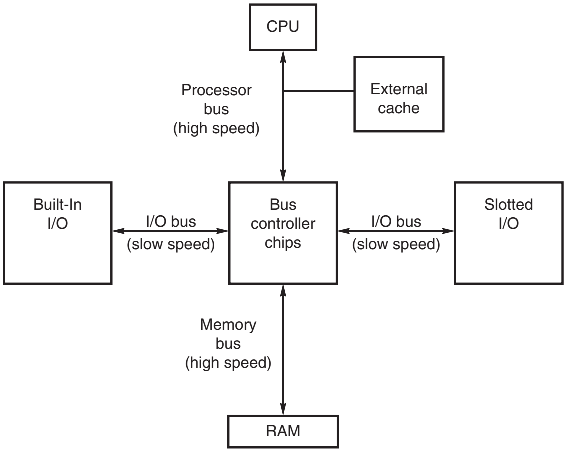

Figure 4.63 Bus layout in a traditional PC.

The thought of a computer system running more slowly than it could is very bothersome to some computer users. Even so, the slow speed of the I/O bus is nothing more than a nuisance in most cases. You don't need blazing speed to communicate with a keyboard or mouse—you gain nothing in performance. The real problem occurs in subsystems in which you need the speed, such as video and disk controllers.

The speed problem became acute when graphical user interfaces (such as Windows) became prevalent. These systems require the processing of so much video data that the I/O bus became a literal bottleneck for the entire computer system. In other words, it did little good to have a processor that was capable of 66MHz–450MHz or faster if you could put data through the I/O bus at a rate of only 8MHz.

An obvious solution to this problem is to move some of the slotted I/O to an area where it could access the faster speeds of the processor bus—much the same way as the external cache. Figure 4.64 shows this arrangement.

This arrangement became known as *local bus* because external devices (adapter cards) now could access the part of the bus that was local to the CPU—the processor bus. Physically, the slosts provided to tap this new configuration would need to be different from existing bus slots to prevent adapter cards designed for slower buses from being plugged into the higher bus speeds, which this design made accessible.

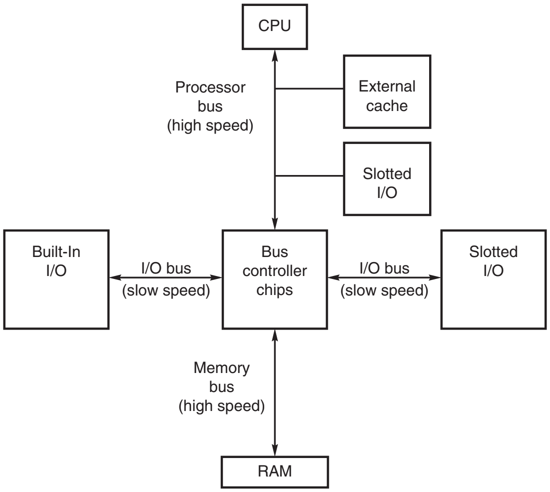

Figure 4.64 How a local bus works.

It is interesting to note that the very first 8-bit and 16-bit ISA buses were a form of local bus architecture. These systems had the processor bus as the main bus, and everything ran at full processor speeds. When ISA systems ran faster than 8MHz, the main ISA bus had to be decoupled from the processor bus because expansion cards, memory, and so on could not keep up. In 1992, an extension to the ISA bus called the VESA local bus (VL-Bus) started showing up on PC systems, indicating a return to local bus architecture. Since then, the peripheral component interconnect (PCI) local bus has supplanted VL-Bus, and the AGP bus has been introduced to complement PCI.

Note

A system does not have to have a local bus expansion slot to incorporate local bus technology; instead, the local bus device can be built directly into the motherboard. (In such a case, the local bus–slotted I/O shown in Figure 4.66 would in fact be built-in I/O.) This built-in approach to local bus is the way the first local bus systems were designed.

Local bus solutions do not necessarily replace earlier standards, such as ISA; they are designed into the system as a bus that is closer to the processor in the system architecture. Older buses such as ISA were kept around for backward compatibility with slower types of adapters that didn't need any faster connection to the system (such as modems). Therefore, until recently a typical system might have AGP, PCI, and ISA slots. Older cards still are compatible with such a system, but high-speed adapter cards can take advantage of the AGP and PCI local bus slots as well. With the demise of ISA slots and the movement of traditionally ISA-based motherboard devices to the LPC interface, today's motherboards essentially use other buses or dedicated interfaces for most of the connections that would have previously used ISA.

The performance of graphical user interfaces such as Windows and OS/2 have been tremendously improved by moving the video cards off the slow ISA bus and onto faster PCI and now AGP local buses.

VESA Local Bus

The Video Electronics Standards Association (VESA) local bus was the most popular local bus design from its debut in August 1992 through 1994. It was created by the VESA committee, a nonprofit organization originally founded by NEC to further develop video display and bus standards. In a similar fashion to how EISA evolved, NEC had done most of the work on the VL-Bus (as it would be called) and, after founding the nonprofit VESA committee, NEC turned over future development to VESA. At first, the local bus slot seemed designed to be used primarily for video cards. Improving video performance was a top priority at NEC to help sell its high-end displays as well as its own PC systems. By 1991, video performance had become a real bottleneck in most PC systems.

The VL-Bus can move data 32 bits at a time, enabling data to flow between the CPU and a compatible video subsystem or hard drive at the full 32-bit data width of the 486 chip. The maximum rated throughput of the VL-Bus is 133MBps. In other words, local bus went a long way toward removing the major bottlenecks that existed in earlier bus configurations.

Unfortunately, the VL-Bus did not seem to be a long-lived concept. The design was simple indeed—just take the pins from the 486 processor and run them out to a card connector socket. So, the VL-Bus is essentially the raw 486 processor bus. This allowed a very inexpensive design because no additional chipsets or interface chips were required. A motherboard designer could add VL-Bus slots to its 486 motherboards very easily and at a very low cost. This is why these slots appeared on virtually all 486 system designs overnight.

Problems arose with timing glitches caused by the capacitance introduced into the circuit by different cards. Because the VL-Bus ran at the same speed as the processor bus, different processor speeds meant different bus speeds, and full compatibility was difficult to achieve. Although the VL-Bus could be adapted to other processors—including the 386 or even the Pentium—it was designed for the 486 and worked best as a 486 solution only. Despite the low cost, after a new bus called PCI appeared, VL-Bus fell into disfavor very quickly. It never did catch on with Pentium systems, and there was little or no further development of the VL-Bus in the PC industry.

Physically, the VL-Bus slot was an extension of the slots used for whatever type of base system you have. If you have an ISA system, the VL-Bus is positioned as an extension of your existing 16-bit ISA slots. Figure 4.65 shows how the VL-Bus slots are oriented on a typical ISA/VL-Bus motherboard. The VESA extension has 112 contacts and uses the same physical connector as the MCA bus.

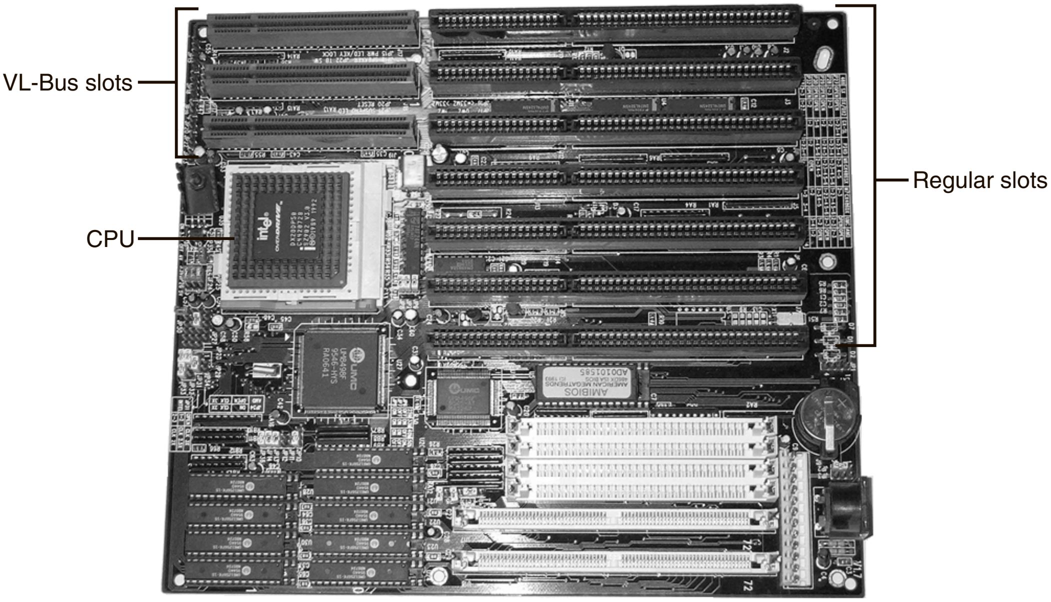

Figure 4.65 An example of a typical motherboard (albeit ancient) with VL-Bus slots.

The PCI Bus

In early 1992, Intel spearheaded the creation of another industry group. It was formed with the same goals as the VESA group in relation to the PC bus. Recognizing the need to overcome weaknesses in the ISA and EISA buses, the PCI Special Interest Group was formed.

The PCI bus specification was released in June 1992 as version 1.0 and since then has undergone several upgrades. Table 4.58 shows the various releases of PCI.

Table 4.58 PCI Specifications

PCI Specification	Released	Major Change
PCI 1.0	June 1992	Original 32/64-bit specification
PCI 2.0	April 1993	Defined connectors and expansion boards
PCI 2.1	June 1995	66MHz operation, transaction ordering, latency changes
PCI 2.2	Jan. 1999	Power management, mechanical clarifications
PCI-X 1.0	Sept. 1999	133MHz operation, addendum to 2.2
Mini-PCI	Nov. 1999	Small form factor boards, addendum to 2.2
PCI 2.3	March 2002	3.3V signaling, low-profile add-in cards
PCI-X 2.0	July 2002	266MHz and 533MHz operation, supports subdivision of 64-bit data bus into 32-bit or 16-bit segments for use by multiple devices, 3.3V/1.5V signaling
PCI-Express 1.0	July 2002	2.5GBps per lane per direction, using 0.8V signaling, resulting in 250MBps per lane; designed to eventually replace PCI 2.x in PC systems

PCI redesigned the traditional PC bus by inserting another bus between the CPU and the native I/O bus by means of bridges. Rather than tap directly into the processor bus, with its delicate electrical timing (as was done in the VL-Bus), a new set of controller chips was developed to extend the bus, as shown in Figure 4.66.

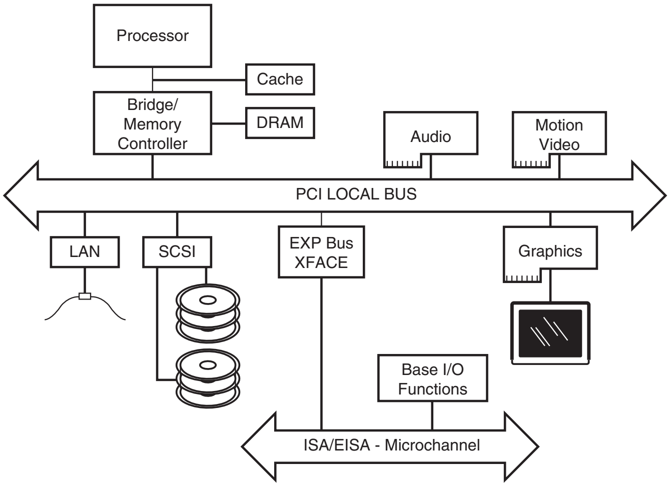

Figure 4.66 Conceptual diagram of the PCI bus.

The PCI bus often is called a *mezzanine bus* because it adds another layer to the traditional bus configuration. PCI bypasses the standard I/O bus; it uses the system bus to increase the bus clock speed and take full advantage of the CPU's data path. Systems that integrate the PCI bus became available in mid-1993 and have since become a mainstay in the PC.

Information typically is transferred across the PCI bus at 33MHz and 32 bits at a time. The bandwidth is 133MBps, as the following formula shows:

33.33MHz × 4 bytes (32 bits) = 133MBps

Although 32-bit 33MHz PCI is the standard found in most PCs, there are now several variations on PCI as shown in Table 4.59.

Table 4.59 PCI Bus Types

PCI Bus Type	Bus Width (Bits)	Bus Speed (MHz)	Data Cycles per Clock	Bandwidth (MBps)
PCI	32	33	1	133
PCI 66MHz	32	66	1	266
PCI 64-bit	64	33	1	266
PCI 66MHz/64-bit	64	66	1	533
PCI-X 64	64*	66	1	533
PCI-X 133	64*	133	1	1,066
PCI-X 266	64*	133	2	2,132
PCI-X 533	64*	133	4	4,266
PCI-Express**	1	2,500	0.8	250
PCI-Express**	32	2,500	0.8	8,000

**Bus width on PCI-X devices can be shared by multiple 32-bit or 16-bit devices.*

***PCI Express uses 8b/10b encoding, which transfers 8 bits for every 10 bits sent and can transfer 1–32 bits at a time, depending on how many lanes are in the implementation.*

Currently, the 64-bit or 66MHz and 133MHz variations are used only on server- or workstation-type boards and systems. Aiding performance is the fact that the PCI bus can operate concurrently with the processor bus; it does not supplant it. The CPU can be processing data in an external cache while the PCI bus is busy transferring information between other parts of the system—a major design benefit of the PCI bus.

A PCI adapter card uses its own unique connector. This connector can be identified within a computer system because it typically is offset from the normal ISA, MCA, or EISA connectors found in older motherboards. See Figure 4.67 for an example. The size of a PCI card can be the same as that of the cards used in the system's normal I/O bus.

The PCI specification identifies three board configurations, each designed for a specific type of system with specific power requirements; each specification has a 32-bit version and a longer 64-bit version. The 5V specification is for stationary computer systems (using PCI 2.2 or earlier versions), the 3.3V specification is for portable systems (also supported by PCI 2.3), and the universal specification is for motherboards and cards that work in either type of system. 64-bit versions of the 5V and universal PCI slots are found primarily on server motherboards. The PCI-X 2.0 specifications for 266 and 533 versions support 3.3V and 1.5V signaling; this corresponds to PCI version 2.3, which supports 3.3V signaling.

Note

The pinouts for the 5V, 3.3V, and universal PCI slots can be found on the DVD-ROM in the Technical Reference section.

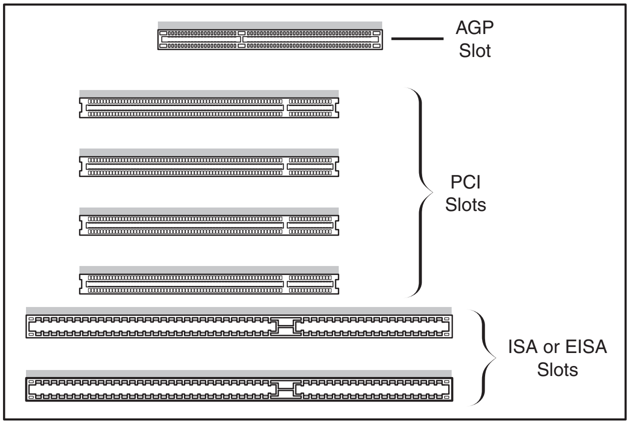

Figure 4.67 Typical configuration of 32-bit 33MHz PCI slots in relation to ISA or EISA and AGP slots.

Figure 4.68 compares the 32-bit and 64-bit versions of the standard 5V PCI slot to a 64-bit universal PCI slot. Figure 4.69 shows how the connector on a 64-bit universal PCI card compares to the 64-bit universal PCI slot.

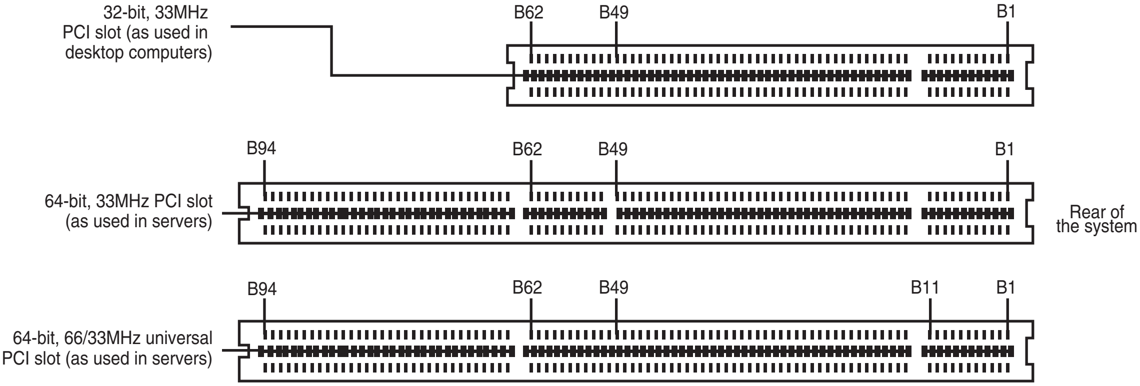

Figure 4.68 A 32-bit, 33MHz PCI slot (top) compared to a 64-bit 33MHz PCI slot (center) and a 64-bit universal PCI slot that runs at 66MHz (bottom).

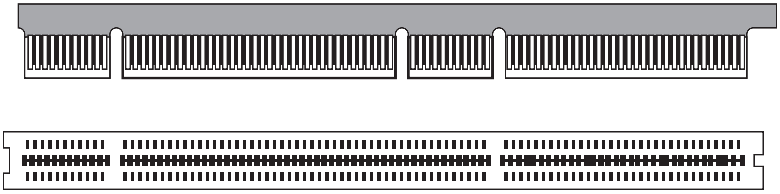

Figure 4.69 A 64-bit universal PCI card (top) compared to the 64-bit universal PCI slot (bottom).

Notice that the universal PCI board specifications effectively combine the 5V and 3.3V specifications. For pins for which the voltage is different, the universal specification labels the pin V I/O. This type of pin represents a special power pin for defining and driving the PCI signaling rail.

Another important feature of PCI is the fact that it was the model for the Intel PnP specification. Therefore, PCI cards do not have jumpers and switches and are instead configured through software. True PnP systems are capable of automatically configuring the adapters, whereas non-PnP systems with ISA slots must configure the adapters through a program that is usually a part of the system CMOS configuration. Starting in late 1995, most PC-compatible systems have included a PnP BIOS that allows the automatic PnP configuration.

PCI Express

During 2001, a group of companies called the Arapahoe Work Group (led primarily by Intel) developed a draft of a new high-speed bus specification code named 3GIO (third-generation I/O). In August 2001, the PCI Special Interest Group (PCI-SIG) agreed to take over, manage, and promote the 3GIO architecture specification as the future generation of PCI. In April 2002, the 3GIO draft version 1.0 was completed, transferred to the PCI-SIG, and renamed PCI Express. Finally in July 2002, the PCI Express 1.0 specification was approved.

The original 3GIO code name was derived from the fact that this new bus specification was designed to initially augment and eventually replace the previously existing ISA/AT-Bus (first-generation) and PCI (second-generation) bus architectures in PCs. Each of the first two generations of PC bus architectures was designed to have a 10- to 15-year useful life in PCs. In being adopted and approved by the PCI-SIG, PCI Express is now destined to be the dominant PC bus architecture designed to support the increasing bandwidth needs in PCs over the next 10–15 years.

The key features of PCI Express are as follows:

- Compatibility with existing PCI enumeration and software device drivers.
- Physical connection over copper, optical, or other physical media to allow for future encoding schemes.
- Maximum bandwidth per pin allows small form factors, reduced cost, simpler board designs and routing, and reduced signal integrity issues.
- Embedded clocking scheme enables easy frequency (speed) changes as compared to synchronous clocking.
- Bandwidth (throughput) increases easily with frequency and width (lane) increases.
- Low latency suitable for applications requiring isochronous (time-sensitive) data delivery, such as streaming video.
- Hot plugging and hot swapping capabilities.
- Power management capabilities.

PCI Express is another example of how the PC is moving from parallel to serial interfaces. Earlier generation bus architectures in the PC have been of a parallel design, in which multiple bits are sent simultaneously over several pins in parallel. The more bits sent at a time, the faster the bus throughput is. The timing of all the parallel signals must be the same, which becomes more and more difficult to do over faster and longer connections. Even though 32 bits can be transmitted simultaneously over a bus such as PCI or AGP, propagation delays and other problems cause them to arrive slightly skewed at the other end, resulting in a time difference between when the first and last of all the bits arrive.

A serial bus design is much simpler, sending 1 bit at a time over a single wire, at much higher rates of speed than a parallel bus would allow. By sending the bits serially, the timing of individual bits or the length of the bus becomes much less of a factor. By combining multiple serial data paths, even faster throughputs can be realized that dramatically exceed the capabilities of traditional parallel buses.

PCI Express is a very fast serial bus design that is backward-compatible with current PCI parallel bus software drivers and controls. In PCI Express, data is sent full duplex (simultaneously operating one-way paths) over two pairs of differentially signaled wires called a *lane*. Each lane allows for about 250MBps throughput in each direction initially, and the design allows for scaling from 1 to 2, 4, 8, 16, or 32 lanes. For example, a high-bandwidth configuration with 8 lanes allowing 8 bits to be sent in each direction simultaneously would allow up to 2,000MBps bandwidth (each way) and use a total of only 40 pins (32 for the differential data pairs and 8 for control). Future increases in signaling speed could increase that to 8,000MBps each way over the same 40 pins. This compares to PCI, which has only 133MBps bandwidth (one way at a time) and requires more than 100 pins to carry the signals. For expansion cards, PCI Express will take on the physical format of a smaller connector that appears adjacent to any existing PCI slots on the motherboard.

PCI Express uses an IBM-designed 8-bit–to–10-bit encoding scheme, which allows for self-clocked signals that will easily allow future increases in frequency. The starting frequency is 2.5GHz, and the specification will allow increasing up to 10GHz in the future, which is about the limit of copper connections. By combining frequency increases with the capability to use up to 32 lanes, PCI Express will be capable of supporting future bandwidths up to 32GBps.

PCI Express is designed to augment and eventually replace many of the buses currently used in PCs. It will not only be a supplement to (and the eventual replacement for) PCI slots, but can also be used to replace the existing Intel hub architecture, HyperTransport, and similar high-speed interfaces between motherboard chipset components. Additionally, it will replace video interfaces such as AGP and act as a mezzanine bus to attach other interfaces, such as Serial ATA, USB 2.0, 1394b (FireWire or iLink), Gigabit Ethernet, and more.

Because PCI Express can be implemented over cables as well as onboard, it can be used to create systems constructed with remote “bricks” containing the bulk of the computing power. Imagine the motherboard, processor, and RAM in one small box hidden under a table, with the video, disk drives, and I/O ports in another box sitting out on a table within easy reach. This will enable a variety of flexible PC form factors to be developed in the future without compromising performance.

PCI Express will not replace PCI or other interfaces overnight. System developers will continue to integrate PCI, AGP, and other bus architectures into system designs for several more years. Just as with PCI and the ISA/AT-Bus before, there will likely be a long period of time during which both buses will be found on motherboards. Gradually, though, fewer PCI and more PCI Express connections will appear. Over time, PCI Express will eventually become the preferred general-purpose I/O interconnect over PCI. I expect the move to PCI Express will be similar to the transition from ISA/AT-Bus to PCI during the 1990s.

Note that I’m including this information on PCI Express well in advance of it actually appearing in PCs. In other words, don’t hold your breath because PCI Express is still in the early design stages, and you won’t see it in PC motherboards for some time yet. The PCI-SIG estimates that the first desktop PCs using PCI Express will begin to emerge in mid- to late 2004. After that, PCI Express is expected to appear in portable devices and low-end servers and workstations by late 2004, and in high-end servers and workstations by late 2005. These are just estimates, of course, and might change according to the dynamics of the industry. The PCI Express Bridge 1.0 and Mini PCI Express Card specifications are designed to help bring PCI Express products into being by using existing PCI technology. These specifications might help shorten the time-to-market for PCI Express products.

For more information on PCI Express, I recommend consulting the PCI-SIG Web site (`www.pcisig.org`).

Accelerated Graphics Port

Intel created AGP as a new bus specifically designed for high-performance graphics and video support. AGP is based on PCI, but it contains several additions and enhancements and is physically, electrically, and logically independent of PCI. For example, the AGP connector is similar to PCI, although it has additional signals and is positioned differently in the system. Unlike PCI, which is a true bus with multiple connectors (slots), AGP is more of a point-to-point high-performance connection designed specifically for a video card in a system because only one AGP slot is allowed for a single video card.

Intel originally released the AGP specification 1.0 in July 1996 and defined a 66MHz clock rate with 1x or 2x signaling using 3.3V. AGP version 2.0 was released in May 1998 and added 4x signaling as well as a lower 1.5V operating capability.

Most newer AGP video cards are designed to conform to the AGP 4X or AGP 8X specification, each of which runs on only 1.5 volts. Most older motherboards with AGP 2X slots are designed to accept only 3.3V cards. If you plug a 1.5V card into a 3.3V slot, both the card and motherboard could be damaged, so special keys have been incorporated into the AGP specification to prevent such disasters. Normally, the slots and cards are keyed such that 1.5V cards fit only in 1.5V sockets and 3.3V cards fit only in 3.3V sockets. However, universal sockets do exist that accept either 1.5V or 3.3V cards. The keying for the AGP cards and connectors is dictated by the AGP standard, as shown in Figure 4.70.

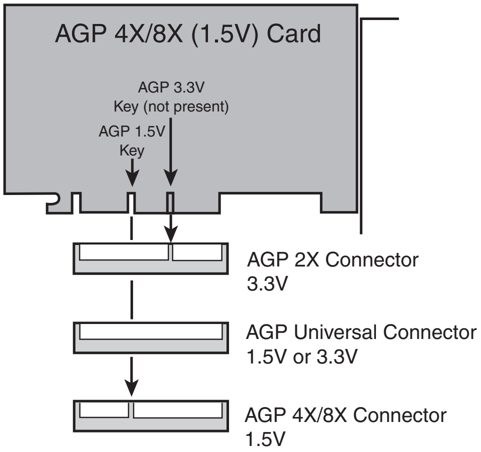

Figure 4.70 AGP 4X/8X (1.5V) card and AGP 3.3V, universal, and 1.5V slots.

As you can see from Figure 4.70, AGP 4X or 8X (1.5V) cards fit only in 1.5V or universal (3.3V or 1.5V) slots. Due to the design of the connector and card keys, a 1.5V card cannot be inserted into a 3.3V slot. So, if your new AGP card won't fit in the AGP slot in your existing motherboard, consider that a good thing because if you were able to plug it in, you would fry both the card and possibly the board as well! In that case, you'd either have to return the 4X/8X card or get a new motherboard that supports the 4X/8X (1.5V) cards.

Caution

Some AGP 4x/8x-compatible motherboards require you to use 1.5V AGP 4x/8x cards only; be sure to check compatibility between the motherboard and the AGP card you want to buy to avoid problems. Some AGP 4x/8x-compatible slots use the card retention mechanism shown in Figure 4.71. Note that AGP 1x/2x slots have a visible divider not present on the newer AGP 4x slot. AGP 4x slots can also accept AGP 8x cards, and vice versa.

Additionally, a newer specification was introduced as AGP Pro 1.0 in August 1998 and was revised in April 1999 as AGP Pro 1.1a. It defines a slightly longer slot with additional power pins at each end to drive bigger and faster AGP cards that consume more than 25 watts of power, up to a maximum of 110 watts. AGP Pro cards are likely to be used for high-end graphics workstations and are not likely to be found in any normal PCs. However, AGP Pro slots are backward-compatible, meaning a standard AGP card will plug in, and a number of motherboard vendors are using AGP Pro slots rather than AGP 4x slots in their latest products. Because AGP Pro slots are longer, an AGP 1x/2x card can be incorrectly inserted into the slot, which could damage it, so some vendors supply a cover for the AGP Pro extension at the rear of the slot. This cover should be removed only if you want to install an AGP Pro card.

The standard AGP 1x/2x, AGP 4x, and AGP Pro slots are compared to each other in Figure 4.71.

The latest revision for the AGP specification for PCs is AGP 8x, otherwise called AGP 3.0. AGP 8x defines a transfer speed of 2,133MBps, which is twice that of AGP 4x. The AGP 8x specification was first publicly pre-announced in November 2000. AGP 8x support is now widely available in the latest motherboard chipsets and graphics chipsets from major vendors. Although AGP 8x has a maximum speed twice that of AGP 4x, the real-world differences between AGP 4x- and 8x-compatible devices with otherwise identical specifications are minimal. However, many 3D chipsets that support AGP 8x are also upgrading memory and 3D graphics core speeds and designs to better support the faster interface.

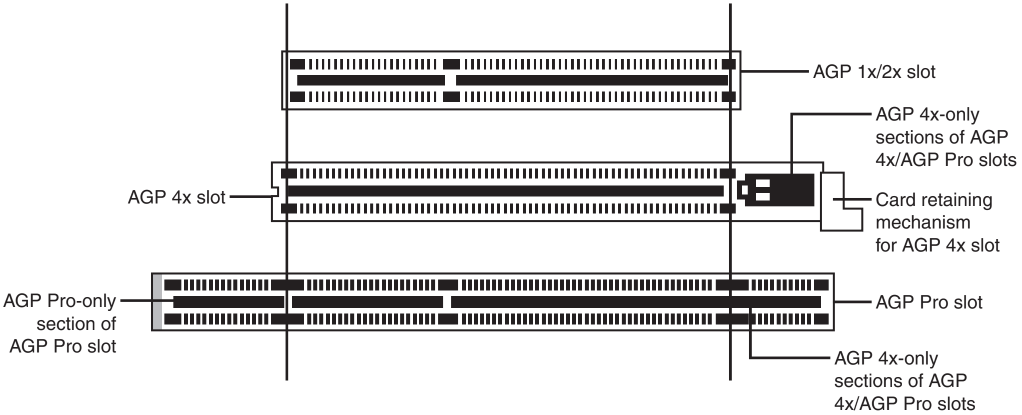

Figure 4.71 AGP standard (1x/2x), AGP 4x, and AGP Pro slots compared to each other. AGP 4x and AGP Pro can accept AGP 1x, 2x, and 4x cards. AGP 4x and AGP Pro slots can also accept AGP 8x cards.

AGP is a high-speed connection and runs at a base frequency of 66MHz (actually 66.66MHz), which is double that of standard PCI. In the basic AGP mode, called 1x, a single transfer is done every cycle. Because the AGP bus is 32 bits (4 bytes) wide, at 66 million times per second it would be capable of transferring data at a rate of about 266MBps! The original AGP specification also defines a 2x mode, in which two transfers are performed every cycle, resulting in 533MBps. Using an analogy in which every cycle is equivalent to the back-and-forth swing of a pendulum, the 1x mode is thought of as transferring information at the start of each swing. In 2x mode, an additional transfer would occur every time the pendulum completed half a swing, thereby doubling performance while technically maintaining the same clock rate, or in this case, the same number of swings per second. Although the earliest AGP cards supported only the AGP 1x mode, most vendors quickly shifted to the AGP 2x mode. The newer AGP 2.0 specification adds the capability for 4x transfers, in which data is transferred four times per cycle and equals a data transfer rate of 1,066MBps. Most newer AGP cards now have support for the 4x standard as a minimum, and the latest graphics chipsets from NVIDIA and ATI support AGP 8x. Table 4.60 shows the differences in clock rates and data transfer speeds (bandwidth) for the various AGP modes.

Table 4.60 AGP Modes Showing Clock Speeds and Bandwidth

AGP Bus Type	Bus Width (Bits)	Bus Speed (MHz)	Data Cycles per Clock	Bandwidth (MBps)
AGP	32	66	1	266
AGP 2x	32	66	2	533
AGP 4x	32	66	4	1,066
AGP 8x	32	66	8	2,133

Because AGP is independent of PCI, using an AGP video card frees up the PCI bus for more traditional input and output, such as for IDE/ATA or SCSI controllers, USB controllers, sound cards, and so on.

Besides faster video performance, one of the main reasons Intel designed AGP was to allow the video card to have a high-speed connection directly to the system RAM, which would enable a reasonably fast and powerful video solution to be integrated at a lower cost. AGP allows a video card to have direct access to the system RAM, either enabling lower-cost video solutions to be directly built in to a motherboard without having to include additional video RAM or enabling an AGP card to share the main system memory. High-performance cards will likely continue the trend of having more and more memory directly on the video card, which is especially important when running high-performance 3D video applications.

AGP allows the speed of the video card to pace the requirements for high-speed 3D graphics rendering as well as full motion video on the PC.

System Resources

System resources are the communications channels, addresses, and other signals hardware devices use to communicate on the bus. At their lowest level, these resources typically include the following:

- Memory addresses
- IRQ (interrupt request) channels
- DMA (direct memory access) channels
- I/O port addresses

I have listed these roughly in the order you would experience problems with them. Memory conflicts are perhaps the most troublesome of these and certainly the most difficult to fully explain and overcome. These are discussed in Chapter 6, which focuses on the others listed here in the order you will likely have problems with them.

IRQs cause more problems than DMAs because they are in much higher demand; virtually all cards use IRQ channels. Fewer problems exist with DMA channels because fewer cards use them, DMA channels are used only by the obsolete ISA standard, and there are usually more than enough channels to go around. I/O ports are used by all hardware devices on the bus, but there are technically 64KB of them, which means there are plenty to go around. With all these resources, you must ensure that a unique card or hardware function uses each resource; in most cases they cannot or should not be shared.

These resources are required and used by many components of your system. Adapter cards need these resources to communicate with your system and accomplish their purposes. Not all adapter cards have the same resource requirements. A serial communications port, for example, needs an IRQ channel and I/O port address, whereas a sound card needs these resources and at least one DMA channel. Most network cards use an IRQ channel and an I/O port address, and some also use a 16KB block of memory addresses.

As your system increases in complexity, the chance for resource conflicts increases. Modern systems with several additional devices can really push the envelope and become a configuration nightmare for the uninitiated. Sometimes under these situations the automatic configuration capability of Plug and Play can get confused or fail to optimally configure resources so that everything will work. Most adapter cards enable you to modify resource assignments by using the Plug and Play software that comes with the card or the Device Manager in Windows 9x and later, thus you can sometimes improve on a default configuration by making some changes. Even if the automatic configuration gets confused (which happens more often than it should), fortunately, in almost all cases a logical way to configure the system exists—once you know the rules.

Interrupts

Interrupt request channels, or hardware interrupts, are used by various hardware devices to signal the motherboard that a request must be fulfilled. This procedure is the same as a student raising his hand to indicate that he needs attention.

These interrupt channels are represented by wires on the motherboard and in the slot connectors. When a particular interrupt is invoked, a special routine takes over the system, which first saves all the CPU register contents in a stack and then directs the system to the interrupt vector table. This vector table contains a list of memory addresses that correspond to the interrupt channels. Depending on which interrupt was invoked, the program corresponding to that channel is run.

The pointers in the vector table point to the address of whatever software driver is used to service the card that generated the interrupt. For a network card, for example, the vector might point to the address of the network drivers that have been loaded to operate the card; for a hard disk controller, the vector might point to the BIOS code that operates the controller.

After the particular software routine finishes performing whatever function the card needed, the interrupt-control software returns the stack contents to the CPU registers, and the system then resumes whatever it was doing before the interrupt occurred.

Through the use of interrupts, your system can respond to external events in a timely fashion. Each time a serial port presents a byte to your system, an interrupt is generated to ensure that the system reads that byte before another comes in. Keep in mind that in some cases a port device—in particular, a modem with a 16550 or higher UART chip—might incorporate a byte buffer that allows multiple characters to be stored before an interrupt is generated.

Hardware interrupts are generally prioritized by their numbers; with some exceptions, the highest-priority interrupts have the lowest numbers. Higher-priority interrupts take precedence over lower-priority interrupts by interrupting them. As a result, several interrupts can occur in your system concurrently, with each interrupt nesting within another.

If you overload the system—in this case, by running out of stack resources (too many interrupts were generated too quickly)—an internal stack overflow error occurs and your system halts. The message usually appears as `Internal stack overflow - system halted` at a DOS prompt. If you experience this type of system error and run DOS, you can compensate for it by using the `STACKS` parameter in your `CONFIG.SYS` file to increase the available stack resources. Most people will not see this error in Windows 9x/Me or Windows NT/2000/XP.

The ISA bus uses *edge-triggered* interrupt sensing, in which an interrupt is sensed by a changing signal sent on a particular wire located in the slot connector. A different wire corresponds to each possible hardware interrupt. Because the motherboard can't recognize which slot contains the card that used an interrupt line and therefore generated the interrupt, confusion results if more than one card is set to use a particular interrupt. Each interrupt, therefore, is usually designated for a single hardware device. Most of the time, interrupts can't be shared.

Originally, IBM developed ways to share interrupts on the ISA bus, but few devices followed the necessary rules to make this a reality. The PCI bus inherently allows interrupt sharing; in fact, virtually all PCI cards are set to PCI interrupt A and share that interrupt on the PCI bus. The real problem is that there are technically two sets of hardware interrupts in the system: PCI interrupts and ISA interrupts. For PCI cards to work in a PC, the PCI interrupts are first mapped to ISA interrupts, which are then configured as non-shareable. Therefore, in many cases you must assign a nonconflicting interrupt for each card, even PCI cards. The conflict between assigning ISA IRQs for PCI interrupts caused many configuration problems for early users of PCI motherboards and continued to cause problems even after the development of Windows 95 and its Plug and Play technology.

The solution to the interrupt sharing problem for PCI cards was something called *PCI IRQ Steering*, which is supported in the more recent operating systems (starting with Windows 95 OSR 2.x) and BIOS. PCI IRQ Steering allows a plug-and-play operating system such as Windows to dynamically map or "steer" PCI cards (which almost all use PCI INTA#) to standard PC interrupts and allows several PCI cards to be mapped to the same interrupt. More information on PCI IRQ Steering is found in the section "PCI Interrupts," later in this chapter.

Hardware interrupts are sometimes referred to as *maskable interrupts*, which means the interrupts can be masked or turned off for a short time while the CPU is used for other critical operations. It is up to the system BIOS and programs to manage interrupts properly and efficiently for the best system performance.

Because interrupts usually can't be shared in an ISA bus system, you often run into conflicts and can even run out of interrupts when you are adding boards to a system. If two boards use the same IRQ to signal the system, the resulting conflict prevents either board from operating properly. The following sections discuss the IRQs that any standard devices use, as well as what might be free in your system.

8-Bit ISA Bus Interrupts

The PC and XT (the systems based on the 8-bit 8086 CPU) provide for eight different external hardware interrupts. Table 4.61 shows the typical uses for these interrupts, which are numbered 0–7.

Table 4.61 8-Bit ISA Bus Default Interrupt Assignments

IRQ	Function	Bus Slot
0	System Timer	No
1	Keyboard Controller	No
2	Available	Yes (8-bit)
3	Serial Port 2 (COM2:)	Yes (8-bit)
4	Serial Port 1 (COM1:)	Yes (8-bit)
5	Hard Disk Controller	Yes (8-bit)
6	Floppy Disk Controller	Yes (8-bit)
7	Parallel Port 1 (LPT1:)	Yes (8-bit)

If you have a system that has one of the original 8-bit ISA buses, you will find that the IRQ resources provided by the system present a severe limitation. Installing several devices that need the services of system IRQs in a PC/XT-type system can be a study in frustration because the only way to resolve the interrupt-shortage problem is to remove the adapter board that you need the least.

16-Bit ISA, EISA, and MCA Bus Interrupts

The introduction of the AT, based on the 286 processor, was accompanied by an increase in the number of external hardware interrupts the bus would support. The number of interrupts was doubled to 16 by using two Intel 8259 interrupt controllers, piping the interrupts generated by the second one through the unused IRQ 2 in the first controller. This arrangement effectively makes only 15 IRQ assignments available, and IRQ 2 effectively became inaccessible.

By routing all the interrupts from the second IRQ controller through IRQ 2 on the first, all these new interrupts are assigned a nested priority level between IRQ 1 and IRQ 3. Thus, IRQ 15 ends up having a higher priority than IRQ 3. Figure 4.72 shows how the two 8259 chips were wired to create the cascade through IRQ 2 on the first chip.

To prevent problems with boards set to use IRQ 2, the AT system designers routed one of the new interrupts (IRQ 9) to fill the slot position left open after removing IRQ 2. This means that any card you install in a modern system that claims to use IRQ 2 is really using IRQ 9 instead.

Table 4.62 shows the typical uses for interrupts in the 16-bit ISA and 32-bit PCI/AGP buses and lists them in priority order from highest to lowest. The obsolete EISA and MCA buses used a similar IRQ map.

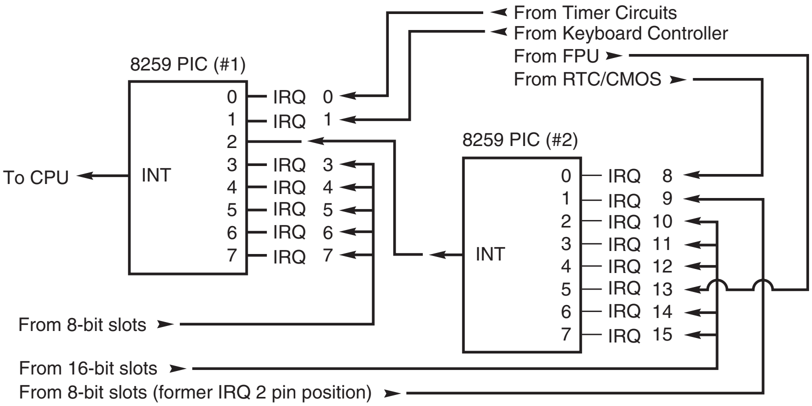

Figure 4.72 Interrupt controller cascade wiring.

Table 4.62 16/32-Bit ISA/PCI/AGP Default Interrupt Assignments

IRQ	Standard Function	Bus Slot	Card Type	Recommended Use
0	System Timer	No	—	—
1	Keyboard Controller	No	—	—
2	2nd IRQ Controller Cascade	No	—	—
8	Real-Time Clock	No	—	—
9	Avail. (as IRQ2 or IRQ9)	Yes	8/16-bit	Network Card
10	Available	Yes	16-bit	USB
11	Available	Yes	16-bit	SCSI Host Adapter
12	Mouse Port/Available	Yes	16-bit	Mouse Port
13	Math Coprocessor	No	—	—
14	Primary IDE	Yes	16-bit	Primary IDE (hard disks)
15	Secondary IDE	Yes	16-bit	2nd IDE (CD-ROM/Tape)
3	Serial 2 (COM2:)	Yes	8/16-bit	COM2:/Internal Modem
4	Serial 1 (COM1:)	Yes	8/16-bit	COM1:
5	Sound/Parallel 2 (LPT2:)	Yes	8/16-bit	Sound Card
6	Floppy Controller	Yes	8/16-bit	Floppy Controller
7	Parallel 1 (LPT1:)	Yes	8/16-bit	LPT1:

Notice that interrupts 0, 1, 2, 8, and 13 are not on the bus connectors and are not accessible to adapter cards. Interrupts 8, 10, 11, 12, 13, 14, and 15 are from the second interrupt controller and are accessible only by boards that use the 16-bit extension connector because this is where these wires are located. IRQ 9 is rewired to the 8-bit slot connector in place of IRQ 2, so IRQ 9 replaces IRQ 2 and, therefore, is available to 8-bit cards, which treat it as though it were IRQ 2.

Note

Although the 16-bit ISA bus has twice as many interrupts as systems that have the 8-bit ISA bus, you still might run out of available interrupts because only 16-bit adapters can use most of the newly available interrupts. Any 32-bit PCI adapter can be mapped to any ISA IRQs.

The extra IRQ lines in a 16-bit ISA system are of little help unless the adapter boards you plan to use enable you to configure them for one of the unused IRQs. Some devices are hard-wired so that they can use only a particular IRQ. If you have a device that already uses that IRQ, you must resolve the conflict before installing the second adapter. If neither adapter enables you to reconfigure its IRQ use, chances are that you can't use the two devices in the same system.

PCI Interrupts

The PCI bus supports hardware interrupts (IRQs) that can be used by PCI devices to signal to the bus that they need attention. The four PCI interrupts are called INTA#, INTB#, INTC#, and INTD#. These INTx# interrupts are *level-sensitive*, which means that the electrical signaling enables them to be shared among PCI cards. In fact, all single device or single function PCI chips or cards that use only one interrupt must use INTA#. This is one of the rules in the PCI specification. If additional devices are within a chip or onboard a card, the additional devices can use INTB# through INTD#. Because there are very few multifunction PCI chips or boards, practically all the devices on a given PCI bus share INTA#.

For the PCI bus to function in a PC, the PCI interrupts must be mapped to ISA interrupts. Because ISA interrupts can't be shared, in most cases each PCI card using INTA# on the PCI bus must be mapped to a different non-shareable ISA interrupt. For example, you could have a system with four PCI slots and four PCI cards installed, each using PCI interrupt INTA#. These cards would each be mapped to a different available ISA interrupt request, such as IRQ9, IRQ10, IRQ11, or IRQ5 in most cases.

Finding unique IRQs for each device on both the ISA and PCI buses has always been a problem; there simply aren't enough free ones to go around. Setting two ISA devices to the same IRQ has never been possible, but on most newer systems sharing IRQs between multiple PCI devices might be possible. Newer system BIOSs as well as plug-and-play operating systems, such as Windows 95B (OSR 2) or later, Windows 98, and Windows 2000/XP, all support a function known as PCI IRQ Steering. For this to work, both your system BIOS and operating system must support IRQ Steering. Older system BIOSs and Windows 95 or 95A do not have support for PCI IRQ Steering.

Generally, the BIOS assigns unique IRQs to PCI devices. If your system supports PCI IRQ Steering and it is enabled, Windows assigns IRQs to PCI devices. Even when IRQ Steering is enabled, the BIOS still initially assigns IRQs to PCI devices. Although Windows has the capability to change these settings, it typically does not do so automatically, except where necessary to eliminate conflicts. If there are insufficient free IRQs to go around, IRQ Steering allows Windows to assign multiple PCI devices to a single IRQ, thus enabling all the devices in the system to function properly. Without IRQ Steering, Windows begins to disable devices after it runs out of free IRQs to assign.

To determine whether your Windows 9x/Me system is using IRQ Steering, you can follow these steps:

1. Open the Device Manager.
2. Double-click the System Devices branch.
3. Double-click PCI Bus, and then click the IRQ Steering tab. There will be a check that displays IRQ Steering as either Enabled or Disabled. If enabled, it also specifies where the IRQ table has been read from.

Note that with Windows 2000 and XP, you can't disable IRQ steering and no IRQ Steering tab appears in the Device Manager.

IRQ Steering is controlled by one of four routing tables Windows attempts to read. Windows searches for the tables in order and uses the first one it finds. You can't control the order in which Windows searches for these tables, but by selecting or deselecting the Get IRQ Table Using check boxes, you can control which table Windows finds first by disabling the search for specific tables. Windows searches for the following tables:

- ACPI BIOS table
- MS Specification table
- Protected Mode PCIBIOS 2.1 table
- Real Mode PCIBIOS 2.1 table

Windows first tries to use the ACPI BIOS table to program IRQ Steering, followed by the MS Specification table, the Protected Mode PCIBIOS 2.1 table, and the Real Mode PCIBIOS 2.1 table. Windows 95 OSR2 and later versions offer only a choice for selecting the PCIBIOS 2.1 tables via a single check box, which is disabled by default. Under Windows 98, all IRQ table choices are selected by default, except the third one, which is the Protected Mode PCIBIOS 2.1 table.

If you are having a problem with a PCI device related to IRQ settings under Windows 95, try selecting the PCIBIOS 2.1 table and restarting. Under Windows 98, try clearing the ACPI BIOS table selection and restarting. If the problem persists, try selecting the Protected Mode PCIBIOS 2.1 table and restarting. You should select Get IRQ Table from Protected Mode PCIBIOS 2.1 Call only if a PCI device is not working properly. To access these settings in the Windows 98 Device Manager, do the following:

1. Open the Device Manager.
2. Scroll down to the System Devices category, and double-click to open it.
3. Select PCI Bus and click Properties.
4. Click the IRQ Steering tab to see or change the current settings.

If IRQ Steering is shown as disabled in Device Manager, be sure the Use IRQ Steering check box is selected. After selecting this and restarting, if IRQ Steering is still showing as disabled, the IRQ routing table that must be provided by the BIOS to the operating system might be missing or contain errors. Check your BIOS Setup to ensure PCI IRQ Steering is enabled. If there is still no success, you might have to select the Get IRQ Table from Protected Mode PCIBIOS 2.1 Call check box, or your BIOS might not support PCI bus IRQ Steering. Contact the manufacturer of your motherboard or BIOS to see whether your board or BIOS supports IRQ Steering.

On systems that have support for IRQ Steering, an IRQ Holder for PCI Steering might be displayed when you view the System Devices branch of Device Manager. This indicates that an IRQ has been mapped to PCI and is unavailable for ISA devices, even if no PCI devices are currently using the IRQ. To view IRQs programmed for PCI-mode, follow these steps:

1. Select Start, Settings, Control Panel, and then double-click System.
2. Click the Device Manager tab.
3. Double-click the System Devices branch.
4. Double-click the IRQ Holder for PCI Steering you want to view, and then click the Resources tab.

I have found this interrupt steering or mapping to be the source of a great deal of confusion. Even though PCI interrupts (INTx#) can be (and are by default) shared, each card or device that might be sharing a PCI interrupt must be mapped or steered to a unique ISA IRQ, which in turn can't normally be shared. You can have several PCI devices mapped to the same ISA IRQ only if

- No ISA devices are using the IRQ.
- The BIOS and operating system support PCI IRQ Steering.
- PCI IRQ Steering is enabled.

Without PCI IRQ Steering support, the sharing capabilities of the PCI interrupts are of little benefit because all PCI-to-ISA IRQ assignments must then be unique. Without PCI IRQ Steering, you can easily run out of available ISA interrupts. If IRQ Steering is supported and enabled, multiple PCI devices will be capable of sharing a single IRQ, allowing for more system expansion without running out of available IRQs. Better support for IRQ Steering is one of the best reasons for upgrading to Windows 98 or newer versions, especially if you are using the original OSR1 release of 95.

Another source of confusion is that the interrupt listing shown in the Windows 9x Device Manager might show the PCI to ISA interrupt mapping as multiple entries for a given ISA interrupt. One entry would be for the device actually mapped to the interrupt—for example, a built-in USB controller—whereas the other entry for the same IRQ would say `IRQ Holder for PCI Steering`. This latter entry, despite claiming to use the same IRQ, does not indicate a resource conflict; instead it represents the chipset circuitry putting a reservation on that interrupt for mapping purposes. This is part of the plug-and-play capabilities of PCI and the modern motherboard chipsets. Windows 2000 and XP can also map multiple devices to the same IRQ, but they don't use the term *IRQ Holder* to avoid confusion.

Note that you can have internal devices on the PCI bus even though all the PCI slots are free. For example, most systems today have two IDE controllers and a USB controller as devices on the PCI bus. Normally, the PCI IDE controllers are mapped to ISA interrupts 14 (primary IDE) and 15 (secondary IDE), whereas the USB controller can be mapped to the available ISA interrupts 9, 10, 11, and 5.

▶▶ See "Universal Serial Bus," p. 947.

The PCI bus enables two types of devices to exist, called *bus masters* (initiators) and *slaves* (targets). A bus master is a device that can take control of the bus and initiate a transfer. The target device is the intended destination of the transfer. Most PCI devices can act as both masters and targets, and to be compliant with the PC 97 and newer system design guides, all PCI slots must support bus master cards.

The PCI bus is an arbitrated bus: A central arbiter (part of the PCI bus controller in the motherboard chipset) governs all bus transfers, giving fair and controlled access to all the devices on the bus. Before a master can use the bus, it must first request control from the central arbiter, and then it is granted control for only a specified maximum number of cycles. This arbitration allows equal and fair access to all the bus master devices, prevents a single device from hogging the bus, and also prevents deadlocks because of simultaneous multiple device access. In this manner, the PCI bus acts much like a local area network (LAN), albeit one that is contained entirely within the system and runs at a much higher speed than conventional external networks between PCs.

IRQ Conflicts

One of the more common IRQ conflicts is the potential one between the integrated COM2: port found in most modern motherboards and an internal (card-based) ISA modem. The problem stems from the fact that true PC card–based modems (not the so-called WinModems, which are software based) incorporate a serial port as part of the card's circuitry. This serial port is set as COM2: by default. Your PC sees this as having two COM2: ports, each using the same IRQ and I/O port address resources.

The solution to this problem is easy: Enter the system BIOS Setup and disable the built-in COM2: port in the system. While you are there, you might think about disabling the COM1: port too because you are unlikely to use it. Disabling unused COMx: ports is one of the best ways to free up a couple of IRQs for other devices to use.

Another common IRQ conflict also involves serial (COM) ports. You might have noticed in the preceding two sections that two IRQs are set aside for two COM ports. IRQ 3 is used for COM2:, and IRQ 4 is used for COM1:. The problem occurs when you have more than two serial ports in a system. When people add COM3: and COM4: ports, they often don't set them to nonconflicting interrupts, which results in a conflict and the ports not working.

Contributing to the problem are poorly designed COM port boards that do not allow IRQ settings other than 3 or 4. What happens is that they end up setting COM3: to IRQ 4 (sharing it with COM1:), and COM4: to IRQ 3 (sharing it with COM2:). This is not acceptable because it prevents you from using the two COM ports on any one of the interrupt channels simultaneously. This was somewhat acceptable under plain DOS because single-tasking (running only one program at a time) was the order of the day, but it is totally unacceptable with Windows and OS/2. If you must share IRQs, you can usually get away with sharing devices on the same IRQ as long as they use different COM ports. For instance, a scanner and an internal modem could share an IRQ, but if the two devices are used simultaneously, a conflict results.

If you need to use serial ports, the best solution is to purchase a multiport serial I/O card that allows nonconflicting interrupt settings or an intelligent card with its own processor that can handle the multiple ports onboard and use only one interrupt in the system. Some older multiport serial cards used the ISA slot, but PCI-slot cards are more common today and have the additional advantages of faster speed and a sharable interrupt.

▶▶ See “Serial Ports,” p. 961.

If a device listed in the table is not present, such as the motherboard mouse port (IRQ 12) or parallel port 2 (IRQ 5), you can consider those interrupts as available. For example, a second parallel port is a rarity, and most systems have a sound card installed and set for IRQ 5 (if it is used to emulate a SoundBlaster Pro or 16). Also, on most systems IRQ 15 is assigned to a secondary IDE controller. If you do not have a second IDE hard or optical drive, you could disable the secondary IDE controller to free up that IRQ for another device.

Note that an easy way to check your interrupt settings is to use the Device Manager in Windows 95/98, Windows NT, or Windows 2000/XP. By double-clicking the Computer Properties icon in the Device Manager, you can get concise lists of all used system resources. Microsoft has also included a program called HWDIAG on Windows 95B; Windows 98 and above feature the System Information program. HWDIAG and System Information do an excellent job of reporting system resource usage, as well as details about device drivers and Windows Registry entries for each hardware component. If you are running Windows XP, a program called MSinfo32 will also give you a report of detailed system information.

DMA Channels

Direct Memory Access (DMA) channels are used by communications devices that must send and receive information at high speeds. A serial or parallel port does not use a DMA channel, but a sound card or SCSI adapter often does. DMA channels sometimes can be shared if the devices are not the type that would need them simultaneously. For example, you can have a network adapter and a tape backup adapter sharing DMA channel 1, but you can’t back up while the network is running. To back up during network operation, you must ensure that each adapter uses a unique DMA channel.

Note

There are several types of DMA in a modern PC. The DMA channels referred to in this section involve the ISA bus. Other buses, such as the ATA/IDE bus used by hard drives, have different DMA uses. The DMA channels explained here don’t involve your ATA/IDE drives, even if they are set to use DMA or Ultra DMA transfers.

8-Bit ISA Bus DMA Channels

In the 8-bit ISA bus, four DMA channels support high-speed data transfers between I/O devices and memory. Three of the channels are available to the expansion slots. Table 4.63 shows the typical uses of these DMA channels.

Table 4.63 8-Bit ISA Default DMA-Channel Assignments

DMA	Standard Function	Bus Slot
0	Dynamic RAM Refresh	No
1	Available	Yes (8-bit)
2	Floppy disk controller	Yes (8-bit)
3	Hard disk controller	Yes (8-bit)

Because most systems typically have both a floppy and hard disk drive, only one DMA channel is available in 8-bit ISA systems.

16-Bit ISA DMA Channels

Since the introduction of the 286 CPU, the ISA bus has supported eight DMA channels, with seven channels available to the expansion slots. Similar to the expanded IRQ lines described earlier in this chapter, the added DMA channels were created by cascading a second DMA controller to the first one. DMA channel 4 is used to cascade channels 0–3 to the microprocessor. Channels 0–3 are available for 8-bit transfers, and channels 5–7 are for 16-bit transfers only. Table 4.64 shows the typical uses for the DMA channels.

Table 4.64 16-Bit ISA Default DMA-channel Assignments

DMA	Standard Function	Bus Slot	Card Type	Transfer	Recommended Use
0	Available	Yes	16-bit	8-bit	Sound
1	Available	Yes	8/16-bit	8-bit	Sound
2	Floppy Disk Controller	Yes	8/16-bit	8-bit	Floppy Controller
3	Available	Yes	8/16-bit	8-bit	LPT1: in ECP Mode
4	1st DMA Controller Cascade	No	—	16-bit	—
5	Available	Yes	16-bit	16-bit	Sound
6	Available	Yes	16-bit	16-bit	Available
7	Available	Yes	16-bit	16-bit	Available

Note that PCI adapters don't use these ISA DMA channels; these are only for ISA cards. However, some PCI cards emulate the use of these DMA channels (such as sound cards) to work with older software.

The only standard DMA channel used in all systems is DMA 2, which is universally used by the floppy controller. DMA 4 is not usable and does not appear in the bus slots. DMA channels 1 and 5 are most commonly used by ISA sound cards, such as the Sound Blaster 16, or by newer PCI sound cards that emulate an older one for backwards compatibility. These cards use both an 8-bit and a 16-bit DMA channel for high-speed transfers.

Note

Although DMA channel 0 appears in a 16-bit slot connector extension and therefore can be used only by a 16-bit card, it performs only 8-bit transfers! Because of this, you generally don't see DMA 0 as a choice on 16-bit cards. Most 16-bit cards (such as SCSI host adapters) that use DMA channels have their choices limited to DMA 5–7.

I/O Port Addresses

Your computer's I/O ports enable communications between devices and software in your system. They are equivalent to two-way radio channels. If you want to talk to your serial port, you need to know on which I/O port (radio channel) it is listening. Similarly, if you want to receive data from the serial port, you need to listen on the same channel on which it is transmitting.

Unlike IRQs and DMA channels, our systems have an abundance of I/O ports. There are 65,535 ports to be exact—numbered from 0000h to FFFFh—which is an artifact of the Intel processor design more than anything else. Even though most devices use up to 8 ports for themselves, with that many to spare, you won't run out anytime soon. The biggest problem you have to worry about is setting two devices to use the same port.

Most modern plug-and-play systems resolve any port conflicts and select alternative ports for one of the conflicting devices.

One confusing issue is that I/O ports are designated by hexadecimal addresses similar to memory addresses. They are not memory; they are ports. The difference is that when you send data to memory address 1000h, it gets stored in your SIMM or DIMM memory. If you send data to I/O port address 1000h, it gets sent out on the bus on that "channel," and anybody listening in could then "hear" it. If nobody is listening to that port address, the data reaches the end of the bus and is absorbed by the bus terminating resistors.

Driver programs are primarily what interact with devices at the various port addresses. The driver must know which ports the device is using to work with it, and vice versa. That is not usually a problem because the driver and device come from the same company.

Motherboard and chipset devices usually are set to use I/O port addresses 0h–FFh, and all other devices use 100h–FFFFh. Table 4.65 shows the commonly used motherboard and chipset-based I/O port usage.

Table 4.65 Motherboard and Chipset-Based Device Port Addresses

Address (hex)	Size	Description
0000–000F	16 bytes	Chipset - 8237 DMA 1
0020–0021	2 bytes	Chipset - 8259 interrupt controller 1
002E–002F	2 bytes	Super I/O controller configuration registers
0040–0043	4 bytes	Chipset - Counter/Timer 1
0048–004B	4 bytes	Chipset - Counter/Timer 2
0060	1 byte	Keyboard/Mouse controller byte - reset IRQ
0061	1 byte	Chipset - NMI, speaker control
0064	1 byte	Keyboard/Mouse controller, CMD/STAT byte
0070, bit 7	1 bit	Chipset - Enable NMI
0070, bits 6:0	7 bits	MC146818 - Real-time clock, address
0071	1 byte	MC146818 - Real-time clock, data
0078	1 byte	Reserved - Board configuration
0079	1 byte	Reserved - Board configuration
0080–008F	16 bytes	Chipset - DMA page registers
00A0–00A1	2 bytes	Chipset - 8259 interrupt controller 2
00B2	1 byte	APM control port
00B3	1 byte	APM status port
00C0–00DE	31 bytes	Chipset - 8237 DMA 2
00F0	1 byte	Math Coprocessor Reset Numeric Error

To find out exactly which port addresses are being used on your motherboard, consult the board documentation or look up these settings in the Windows Device Manager.

Bus-based devices typically use the addresses from 100h on up. Table 4.66 lists the commonly used bus-based device addresses and some common adapter cards and their settings.

Table 4.66 Bus-Based Device Port Addresses

Address (hex)	Size	Description
0130–0133	4 bytes	Adaptec SCSI adapter (alternate)
0134–0137	4 bytes	Adaptec SCSI adapter (alternate)
0168–016F	8 bytes	Fourth IDE interface
0170–0177	8 bytes	Secondary IDE interface
01E8–01EF	8 bytes	Third IDE interface
01F0–01F7	8 bytes	Primary IDE/AT (16-bit) hard disk controller
0200–0207	8 bytes	Gameport or joystick adapter
0210–0217	8 bytes	IBM XT expansion chassis
0220–0233	20 bytes	Creative Labs Sound Blaster 16 audio (default)
0230–0233	4 bytes	Adaptec SCSI adapter (alternate)
0234–0237	4 bytes	Adaptec SCSI adapter (alternate)
0238–023B	4 bytes	MS bus mouse (alternate)
023C–023F	4 bytes	MS bus mouse (default)
0240–024F	16 bytes	SMC Ethernet adapter (default)
0240–0253	20 bytes	Creative Labs Sound Blaster 16 audio (alternate)
0258–025F	8 bytes	Intel above board
0260–026F	16 bytes	SMC Ethernet adapter (alternate)
0260–0273	20 bytes	Creative Labs Sound Blaster 16 audio (alternate)
0270–0273	4 bytes	Plug and Play I/O read ports
0278–027F	8 bytes	Parallel port 2 (LPT2)
0280–028F	16 bytes	SMC Ethernet adapter (alternate)
0280–0293	20 bytes	Creative Labs Sound Blaster 16 audio (alternate)
02A0–02AF	16 bytes	SMC Ethernet adapter (alternate)
02C0–02CF	16 bytes	SMC Ethernet adapter (alternate)
02E0–02EF	16 bytes	SMC Ethernet adapter (alternate)
02E8–02EF	8 bytes	Serial port 4 (COM4)
02EC–02EF	4 bytes	Video, 8514, or ATI standard ports
02F8–02FF	8 bytes	Serial port 2 (COM2)
0300–0301	2 bytes	MPU-401 MIDI port (secondary)
0300–030F	16 bytes	SMC Ethernet adapter (alternate)
0320–0323	4 bytes	XT (8-bit) hard disk controller
0320–032F	16 bytes	SMC Ethernet adapter (alternate)
0330–0331	2 bytes	MPU-401 MIDI port (default)
0330–0333	4 bytes	Adaptec SCSI adapter (default)
0334–0337	4 bytes	Adaptec SCSI adapter (alternate)
0340–034F	16 bytes	SMC Ethernet adapter (alternate)
0360–036F	16 bytes	SMC Ethernet adapter (alternate)
0366	1 byte	Fourth IDE command port
0367, bits 6:0	7 bits	Fourth IDE status port
0370–0375	6 bytes	Secondary floppy controller

Table 4.66 Continued

Address (hex)	Size	Description
0376	1 byte	Secondary IDE command port
0377, bit 7	1 bit	Secondary floppy controller disk change
0377, bits 6:0	7 bits	Secondary IDE status port
0378–037F	8 bytes	Parallel Port 1 (LPT1)
0380–038F	16 bytes	SMC Ethernet adapter (alternate)
0388–038B	4 bytes	Audio - FM synthesizer
03B0–03BB	12 bytes	Video, Mono/EGA/VGA standard ports
03BC–03BF	4 bytes	Parallel port 1 (LPT1) in some systems
03BC–03BF	4 bytes	Parallel port 3 (LPT3)
03C0–03CF	16 bytes	Video, EGA/VGA standard ports
03D0–03DF	16 bytes	Video, CGA/EGA/VGA standard ports
03E6	1 byte	Third IDE command port
03E7, bits 6:0	7 bits	Third IDE status port
03E8–03EF	8 bytes	Serial port 3 (COM3)
03F0–03F5	6 bytes	Primary floppy controller
03F6	1 byte	Primary IDE command port
03F7, bit 7	1 bit	Primary floppy controller disk change
03F7, bits 6:0	7 bits	Primary IDE status port
03F8–03FF	8 bytes	Serial port 1 (COM1)
04D0–04D1	2 bytes	Edge/level triggered PCI interrupt controller
0530–0537	8 bytes	Windows sound system (default)
0604–060B	8 bytes	Windows sound system (alternate)
0678–067F	8 bytes	LPT2 in ECP mode
0778–077F	8 bytes	LPT1 in ECP mode
0A20–0A23	4 bytes	IBM Token-Ring adapter (default)
0A24–0A27	4 bytes	IBM Token-Ring adapter (alternate)
0CF8–0CFB	4 bytes	PCI configuration address registers
0CF9	1 byte	Turbo and reset control register
0CFC–0CFF	4 bytes	PCI configuration data registers
FF00–FF07	8 bytes	IDE bus master registers
FF80–FF9F	32 bytes	Universal serial bus
FFA0–FFA7	8 bytes	Primary bus master IDE registers
FFA8–FFAF	8 bytes	Secondary bus master IDE registers

To find out exactly what your devices are using, again I recommend consulting the documentation for the device or looking up the device in the Windows Device Manager. Note that the documentation for some devices might list only the starting address instead of the full range of I/O port addresses used.

Virtually all devices on the system buses use I/O port addresses. Most of these are fairly standardized, meaning conflicts or problems won't often occur with these settings. In the next section, you learn more about working with I/O addresses.

Resolving Resource Conflicts

The resources in a system are limited. Unfortunately, the demands on those resources seem to be unlimited. As you add more and more adapter cards to your system, you will find that the potential for resource conflicts increases. If your system is fully PnP-compatible, potential conflicts should be resolved automatically, but often are not.

How do you know whether you have a resource conflict? Typically, one of the devices in your system stops working. Resource conflicts can exhibit themselves in other ways, though. Any of the following events could be diagnosed as a resource conflict:

- A device transfers data inaccurately.
- Your system frequently locks up.
- Your sound card doesn't sound quite right.
- Your mouse doesn't work.
- Garbage appears on your video screen for no apparent reason.
- Your printer prints gibberish.
- You can't format a floppy disk.
- The PC starts in Safe mode (Windows 9x/Me) or can start only in Last Known Good Configuration (Windows 2000/XP).

Windows 9x/Me and Windows 2000/XP also show conflicts by highlighting a device in yellow or red in the Device Manager representation. By using the Windows Device Manager, you can usually spot the conflicts quickly.

In the following sections, you learn some of the steps you can take to head off resource conflicts or track them down when they occur.

Caution

Be careful when diagnosing resource conflicts; a problem might not be a resource conflict at all, but a computer virus. Many computer viruses are designed to exhibit themselves as glitches or periodic problems. If you suspect a resource conflict, it might be worthwhile to run a virus check first to ensure that the system is clean. This procedure could save you hours of work and frustration.

One way to resolve conflicts is to help prevent them in the first place. Especially if you are building up a new system, you can take several steps to avoid problems. One is to avoid using older ISA devices. By definition, they cannot share IRQs, and that is the resource most in demand. PCI (and AGP) cards can share IRQs with IRQ Steering and as such are a much better choice.

Tip

The serial, PS/2 mouse, and parallel ports still found in most recent systems are all ISA devices that cannot share IRQs. If you no longer use these ports, you can use these devices' IRQs for other devices if you

- Disable the unused port in the system BIOS.
- Configure the system BIOS to use the IRQ formerly used by the device(s) for PnP configuration; this might be automatic in some systems.

Another way you can help is to install cards in a particular sequence, and not all at once. Modifying the installation sequence often helps because many cards can use only one or two out of a predefined selection of IRQs that is specific to each brand or model of card. By installing the cards in a controlled sequence, the plug-and-play software can more easily work around IRQ conflicts caused by the default configurations of different cards.

The first time you start up a new system you have assembled or done major upgrades on, the first thing you should check is the BIOS Setup. If you have a setting for PnP Operating System in your BIOS, be sure it is enabled if you are running an operating system with plug-and-play support, such as Windows 9x/Me/2000/XP. Otherwise, make sure it's disabled if you are running an OS that is not plug-and-play, such as Windows NT.

On initial startup I recommend a minimum configuration with only the graphics card, memory, and storage drives (floppy, hard disk, CD-ROM, and DVD). This allows for the least possibility of system conflicts in the initial configuration. If your motherboard came with a CD including drivers specific to the chipset or other built-in features of the board, now is the time to load or install them. Complete the configuration of all built-in devices before installing any other cards or external devices.

After the basic system has been configured (and after you have successfully loaded your operating system and any updates or patches), you can then begin adding one device at a time in a specific order. So, you will power down, install the new device, power up, and proceed to install any necessary drivers and configure the device. You'll probably have to restart your system after you are done to fully complete the configuration.

Tip

I sometimes recommend that between installing devices you enter the Device Manager in Windows and print out the resource settings as they are configured at the time. This way you have a record of how the configuration changes during the entire device installation and configuration process.

Here's the loading sequence for additional cards:

1. Sound card
2. Internal or external modem
3. Network card
4. Auxiliary video devices, such as MPEG decoders, 3D accelerators, and so on
5. SCSI adapter
6. Anything else

Normally, using this controlled sequence of configuring or building up your system results in easier integration with less conflicts and configuration hassles.

Resolving Conflicts Manually

In the past, the only way to resolve conflicts manually was to take the cover off your system and start changing switches or jumper settings on the adapter cards. Fortunately, this is a bit easier with plug-and-play because all the configuration is done via the Device Manager software included in the operating system. Although some early plug-and-play cards also had jumper switches or setup options to enable them to be configured manually, this feature was found primarily on ISA PnP-compatible cards.

Be sure you write down or print out your current system settings before you start making changes. That way, you will know where you began and can go back to the original configuration (if necessary).

Finally, dig out the manuals for all your adapter boards; you might need them, particularly if they can be configured manually or be switched to PnP mode. Additionally, you could look for more current information online at the manufacturers' Web sites.

Now you are ready to begin your detective work. As you try various resource settings, keep the following questions in mind; the answers will help you narrow down the conflict areas:

- *When did the conflict first become apparent?* If the conflict occurred after you installed a new adapter card, that new card probably is causing the conflict. If the conflict occurred after you started using new software, the software probably uses a device that is taxing your system's resources in a new way.
- *Are there two similar devices in your system that do not work?* For example, if your modem, integrated serial ports, or mouse—devices that use a COM port—do not work, chances are good that these devices are conflicting with each other.
- *Have other people had the same problem, and if so, how did they resolve it?* Public forums, such as those on CompuServe, Internet newsgroups, and America Online, are great places to find other users who might be able to help you solve the conflict.

Whenever you make changes in your system, reboot and see whether the problem persists. When you believe that you have solved the problem, be sure to test all your software. Fixing one problem often seems to cause another to crop up. The only way to ensure that all problems are resolved is to test everything in your system.

One of the best pieces of advice I can give you is to try changing one thing at a time, and then retest. That is the most methodical and the simplest way to isolate a problem quickly and efficiently.

As you attempt to resolve your resource conflicts, you should work with and update a system-configuration template, as discussed in the following section.

Using a System-Configuration Template

A *system-configuration template* is helpful because remembering something that is written down is easier than keeping it in your head. To create a configuration template, you need to start writing down which resources are used by which parts of your system. Then, when you need to make a change or add an adapter, you can quickly determine where conflicts might arise. You can also use the Windows 9x/Me/2000/XP Device Manager to list and print this information.

I like to use a worksheet split into three main areas—one for interrupts, another for DMA channels, and a middle area for devices that do not use interrupts. Each section lists the IRQ or DMA channel on the left and the I/O port device range on the right. This way, I get the clearest picture of which resources are used and which ones are available in a given system.

Here is the system-configuration template I have developed over the years and still use almost daily:

System Resource Map

PC Make and Model: ____________________
Serial Number: ____________________
Date: ____________________

Interrupts (IRQs):		I/O Port Addresses:
0 -	Timer Circuits ________	040–04B ________
1 -	Keyboard/Mouse Controller ________	060 & 064 ________
2 -	2nd 8259 IRQ Controller ________	0A0–0A1 ________
8 -	Real-time Clock/CMOS RAM ________	070–071 ________
9 -	________	________
10 -	________	________
11 -	________	________
12 -	________	________
13 -	Math Coprocessor ________	0F0 ________
14 -	________	________
15 -	________	________
3 -	________	________
4 -	________	________
5 -	________	________
6 -	________	________
7 -	________	________

Devices Not Using Interrupts:	I/O Port Addresses:
Mono/EGA/VGA Standard Ports ________	3B0–3BB ________
EGA/VGA Standard Ports ________	3C0–3CF ________
CGA/EGA/VGA Standard Ports ________	3D0–3DF ________
________	________
________	________
________	________
________	________
________	________

DMA Channels:

0 - ________
1 - ________
2 - ________
3 - ________
4 - DMA Channel 0–3 Cascade ________
5 - ________
6 - ________
7 - ________

This type of configuration sheet is resource based instead of component based. Each row in the template represents a different resource and lists the component using the resource as well as the resources used. The chart has pre-entered all the fixed items in a modern PC for which the configuration cannot be changed.

To fill out this type of chart, you would perform the following steps:

1. Enter the default resources used by standard components, such as serial and parallel ports, disk controllers, and video. You can use the filled-out example I have provided to see how most standard devices are configured.
2. Enter the default resources used by additional add-on components, such as sound cards, SCSI cards, network cards, proprietary cards, and so on.
3. Change any configuration items that are in conflict. Try to leave built-in devices at their default settings, as well as sound cards. Other installed adapters might have their settings changed, but be sure to document the changes.

Of course, a template such as this is best used when first installing components, not after. After you have it completely filled out to match your system, you can label it and keep it with the system. Whenever you add more devices, the template will be your guide as to how any new devices should be configured if you need to configure devices manually.

Note

Thanks to plug-and-play configuration, the days of fixed IRQ and other hardware resources are receding into the past. Don't be surprised if your system has assigned different IRQ, I/O port address, or DMA settings after you install a new card. That's why I recommend recording information both before and after you add a new device to your system.

You also might want to track which PCI slot is used by a particular card because some systems convert PCI IRQs to different ISA IRQs depending on which slot is used for a card. Also, some systems pair PCI slots or might pair the AGP slot and a PCI slot, assigning cards installed in both paired slots to the same ISA IRQ.

The template shown on the next page is the same template filled out for a typical PC system with a mixture of PCI and ISA devices.

As you can see from this template, only one IRQ and two DMA channels remain available, and that would be *no* IRQs if I enabled the USB on the motherboard! As you can see, interrupt shortages are a big problem in modern systems. In that case, I would probably find a way to recover one of the other interrupts; for example, I am not really using COM1, so I could disable that port and gain back IRQ 4. In this sample configuration, the primary and secondary IDE connectors were built in to the motherboard:

- Floppy controller
- Two serial ports
- One parallel port

Whether these devices are built in to the motherboard or on a separate card makes no difference because the resource allocations are the same in either case. All default settings are typically used for these devices and are indicated in the completed configuration. Next, the accessory cards were configured. In this example, the following cards were installed:

- SVGA video card (ATI Mach 64)
- Sound card (Creative Labs Sound Blaster 16)
- SCSI host adapter (Adaptec AHA-1542CF)
- Network interface card (SMC EtherEZ)

It helps to install the cards in this order. Start with the video card; next, add the sound card. Because of problems with software that must be configured to the sound card, it is best to install it early and ensure that only default settings are used. It is better to change settings on cards other than the sound card.

System Resource Map

PC Make and Model: Intel SE440BX-2
Serial Number: 100000
Date: 06/09/99

Interrupts (IRQs):		I/O Port Addresses:
0 -	Timer Circuits	040–04B
1 -	Keyboard/Mouse Controller	060 & 064
2 -	2nd 8259 IRQ Controller	0A0–0A1
8 -	Real-time Clock/CMOS RAM	070–071
9 -	SMC EtherEZ Ethernet Card	340–35F
10 -		
11 -	Adaptec 1542CF SCSI Adapter (scanner)	334–337[1]
12 -	Motherboard Mouse Port	060 and 064
13 -	Math Coprocessor	0F0
14 -	Primary IDE (hard disk 1 and 2)	1F0–1F7, 3F6
15 -	Secondary IDE (CD-ROM/tape)	170–177, 376
3 -	Serial Port 2 (Modem)	3F8–3FF
4 -	Serial Port 1 (COM1)	2F8–2FF
5 -	Sound Blaster 16 Audio	220–233
6 -	Floppy Controller	3F0–3F5
7 -	Parallel Port 1 (Printer)	378–37F

Devices Not Using Interrupts:	I/O Port Addresses:
Mono/EGA/VGA Standard Ports	3B0–3BB
EGA/VGA Standard Ports	3C0–3CF
CGA/EGA/VGA Standard Ports	3D0–3DF
ATI Mach 64 Video Card Additional Ports	102,1CE–1CF, 2EC–2EF
Sound Blaster 16 MIDI Port	330–331
Sound Blaster 16 Game Port (joystick)	200–207
Sound Blaster 16 FM Synthesizer (music)	388–38B

DMA Channels:	
0 -	
1 -	Sound Blaster 16 (8-bit DMA)
2 -	Floppy Controller
3 -	Parallel Port 1 (in ECP mode)
4 -	DMA Channel 0–3 Cascade
5 -	Sound Blaster 16 (16-bit DMA)
6 -	Adaptec 1542CF SCSI Adapter[1]
7 -	

[1] *Represents a resource setting that had to be changed to resolve a conflict.*

After the sound card, the SCSI adapter was installed; however, the default I/O port addresses (330–331) and DMA channel (DMA 5) used were in conflict with other cards (mainly the sound card). These settings were changed to their next logical settings that did not cause a conflict.

Finally, the network card was installed, which also had default settings that conflicted with other cards. In this case, the Ethernet card came preconfigured to IRQ 3, which was already in use by COM2:. The solution was to change the setting, and IRQ 9 was the next logical choice in the card's configuration settings.

Even though this is a fully loaded configuration, only three individual items among all the cards had to be changed to achieve an optimum system configuration. As you can see, using a configuration template such as the one shown can make what would otherwise be a jumble of settings lay out in an easy-to-follow manner. The only real problems you will run into after you work with these templates are cards that do not allow for enough adjustment in their settings or cards that are lacking in documentation. As you can imagine, you will need the documentation for each adapter card, as well as the motherboard, to accurately complete a configuration table such as the one shown.

Tip

Do not rely too much on third-party DOS-based software diagnostics, such as MSD.EXE, which claim to be capable of showing hardware settings such as IRQ and I/O port settings. Even though they can be helpful in certain situations, they are often wrong with respect to at least some of the information they display about your system. One or two items shown incorrectly can be very troublesome if you believe the incorrect information and configure your system based on it!

Some third-party products such as AMIDiag and CheckIt do a better job, but a much better utility to view these settings is the Device Manager built in to Windows 9x/Me/2000/XP. With plug-and-play hardware, it not only reports settings, but it allows you to change them in some cases (sometimes this requires moving the card to another PCI slot). On older legacy hardware, you can view the settings but not change them. To change the settings of legacy (non–plug-and-play) hardware, you must manually move jumpers or switches and run the special configuration software that came with the card. Consult the card manufacturer or documentation for more information.

Heading Off Problems: Special Boards

A number of devices that you might want to install in a computer system require IRQ lines or DMA channels, which means that a world of conflict could be waiting in the box the device comes in. As mentioned in the preceding section, you can save yourself problems if you use a system-configuration template to keep track of the way your system is configured.

You also can save yourself trouble by carefully reading the documentation for a new adapter board before you attempt to install it. The documentation details the IRQ lines the board can use as well as its DMA-channel requirements. In addition, the documentation details the adapter's upper-memory needs for ROM and adapter.

The following sections describe some of the conflicts you might encounter when you install some popular adapter boards. Although the list of adapter boards covered in these sections is far from comprehensive, the sections serve as a guide to installing complex hardware with minimum hassle. Included are tips on soundboards, SCSI host adapters, and network adapters.

Sound Cards

Sound cards are probably the biggest single resource hog in your system. They typically use at least one IRQ, two DMA channels (in DOS emulation mode), and multiple I/O port address ranges. This is because a sound card is actually several different pieces of hardware all on one board. Most sound cards, including PCI-based models, emulate the Sound Blaster 16 from Creative Labs.

Table 4.67 shows the default resources used by a typical PCI sound card, the Creative Labs SB512. Because sound cards are multifunction devices, each function is listed separately as it appears in the Windows Device Manager.

Table 4.67 Default Resources Used by Creative Labs SB512

Device	IRQ	I/O Port Address	16-Bit DMA Channel	8-Bit DMA Channel
Sound Blaster 16 emulation	5	0220–022F 0330–0331[1] 0388–038B[2]	5	1
Creative Multimedia Interface	N/A	D400–D407	n/a	n/a
Creative Gameport Joystick	N/A	0200–0207	n/a	n/a
Creative SB512	9[3]	C860–C87F	n/a	n/a

1. *Used for MIDI interface*
2. *Used for FM Synthesis*
3. *Varies with system and PCI expansion slot used*

Although other brands of sound cards might use a slightly different configuration, the pattern is the same; Sound Blaster emulation requires a large number of resources. Even if SB emulation isn't used (you might be able to disable it if you no longer play DOS games), the card still uses a single IRQ and several I/O port address ranges. If you take the time to read your soundboard's documentation and determine its communications-channel needs, compare those needs to the IRQ lines and DMA channels that already are in use in your system, and then change the settings of the other adapters to avoid conflicts with the sound card, your installation will go quickly and smoothly. Unfortunately, many vendors no longer provide detailed information on their plug-and-play–compatible cards. That's why you need to install the sound card first and use the system resource map to record which settings the card uses before you install other cards.

Tip

The best advice I can give you for installing a sound card is to put the sound card in before all other cards—except for video. In other words, let the sound card retain all its default settings if you can. Try to change the settings of other adapters when a conflict with the sound card arises. The problem here is that many older programs that use sound are very poorly written with respect to supporting alternative resource settings on sound cards. Save yourself some grief, and let the sound card have its way!

Even the latest sound cards can have problems if they're installed after other cards. I know a user who had to remove all his plug-and-play cards before his system would recognize his plug-and-play sound card.

One example of a potential soundboard conflict is the combination of a Sound Blaster 16 and an Adaptec SCSI adapter, as I noted earlier in this chapter. The Sound and SCSI adapters conflict on DMA 5 as well as on I/O ports 330–331. Rather than changing the settings of the sound card, it is best to alter the SCSI adapter to the next available settings that will not conflict with the sound card or anything else. The final settings are shown in the previous configuration template.

The cards in question (Sound Blaster 16 and AHA-1542CF) are not singled out here because there is something wrong with them, but instead because they happen to be very popular cards of their respective types and, as such, often are paired together in older systems.

Most people would be using PCI versions of these cards today, but they still require the same types of resource settings with the only exception being DMA channels. Unfortunately, it wasn't DMA channels that we were really running out of! The interrupt shortage often continues even with PCI cards because for older real-mode applications or Windows 95 and earlier, they must be mapped to discrete ISA IRQs. The real solutions to the ISA IRQ problem are to use only PCI cards, use Windows 98 or newer (which support IRQ Steering), and have that support in your ROM BIOS. Then, full sharing is possible. Every desktop system sold today is equipped with a motherboard that lacks ISA slots and breaks ties with that bus forever. These motherboards free us of the interrupt restrictions we have been under for so many years.

Tip

The newer PCI sound cards are largely incompatible with older DOS-based software because they don't use DMA channels like their ISA counterparts. If you don't update your software to 32-bit Windows versions, you won't be able to use these newer PCI bus sound cards. Most of the newer PCI cards do include an emulation program that allows the card to work with older DMA-dependent software, but the results are often problematic.

For the best results, use the PC/PCI connector found on some motherboards to connect a patch cable to a PC/PCI-compatible sound card. The PC/PCI connector enables the sound card to use ISA-style DMA channels without clumsy emulation software.

SCSI Adapter Boards

SCSI adapter boards use more resources than just about any other type of add-in device except perhaps a sound card. They often use resources that are in conflict with sound cards or network cards, especially if the card has an onboard BIOS for handling bootable drives. A typical ISA SCSI host adapter with onboard BIOS requires an IRQ line, a DMA channel, a range of I/O port addresses, plus a 16KB range of unused upper memory for its ROM and possible scratch-pad RAM use. Even a simple SCSI host adapter designed for use with scanners still requires an IRQ and a range of I/O port addresses. Fortunately, the typical SCSI adapter is also easy to reconfigure, and changing any of these settings should not affect performance or software operation. PCI-based SCSI host adapters require all of the preceding, except for the DMA channel.

Before installing a SCSI adapter, be sure to read the documentation for the card, and ensure that any IRQ lines, DMA channels, I/O ports, and upper memory the card needs are available. If the system resources the card needs are already in use, use your system-configuration template to determine how you can alter the settings on the SCSI card or other cards to prevent any resource conflicts before you attempt to plug in the adapter card.

Network Interface Cards

Networks are becoming more and more popular all the time, thanks to the rise of easy-to-configure small office/home office networks and the use of network cards to connect to broadband Internet devices such as cable and DSL modems. A typical network adapter does not require as many resources as some of the other cards discussed here, but it requires at least a range of I/O port addresses and an interrupt. Some network interface cards (NICs) also require a 16KB range of free upper memory to be used for the RAM transfer buffer on the network card. As with any other cards, be sure that all these resources are unique to the card and are not shared with any other devices. If your network adapter is built in to the motherboard, it still uses IRQ and I/O port address resources.

Multiple-COM-Port Adapters

A serial port adapter usually has two or more ports onboard. These COM ports require an interrupt and a range of I/O ports each. There aren't too many problems with the I/O port addresses because the ranges used by up to four COM ports in a system are fairly well defined. The real problem is with the interrupts. Most older installations of more than two serial ports have any additional ones sharing the same interrupts as the first two. This is incorrect and causes nothing but problems with software that runs under Windows or OS/2. With these older boards, ensure that each serial port in your system has a unique I/O port address range and, more importantly, a unique interrupt setting.

Many newer multiport adapter cards—such as those offered by Byte Runner Technologies—allow "intelligent" interrupt sharing among ports. In some cases, you can have up to 12 COM port settings without conflict problems. Check with your adapter card's manufacturer to determine whether it allows for automatic or "intelligent" interrupt sharing.

Although most people have problems incorrectly trying to share interrupts when installing more than two serial ports in a system, a fairly common problem exists with the I/O port addressing that should be mentioned. Many of the high-performance video chipsets, such as those from S3, Inc., and ATI, use some additional I/O port addresses that conflict with the standard I/O port addresses used by COM4:.

In the sample system configuration just covered, you can see that the ATI video card uses some additional I/O port addresses, specifically 2EC–2EF. This is a problem because COM4: usually is configured as 2E8–2EF, which overlaps with the video card. The video cards that use these addresses are not normally adjustable for this setting, so you either must change the address of COM4: to a nonstandard setting or disable COM4: and restrict yourself to using only three serial ports in the system. If you do have a serial adapter that supports nonstandard I/O address settings for the serial ports, you must ensure that those settings are not used by other cards, and you must inform any software or drivers, such as those in Windows, of your nonstandard settings.

In many cases, USB or 10/100 Ethernet connections can be used to perform the tasks formerly handled by serial ports. If you no longer need serial ports for any devices in your system, I recommend that you disable the onboard serial ports, remove multiport serial cards, and assign the IRQs used by the serial ports in the BIOS to be available for PCI/AGP Plug and Play assignment. This helps eliminate any IRQ conflicts that could occur.

Universal Serial Bus

USB ports corresponding to USB 1.1 or USB 2.0 are now found on most motherboards, and Windows 98, Windows Me, Windows 2000, and Windows XP provide you with a wide range of operating systems that support them properly. One potential problem is that USB takes another interrupt from your system, and many computers either don't have any free or are down to their last one. If your system supports PCI IRQ Steering, this shouldn't be much of a problem because the IRQ used by your USB controller should be sharable with other PCI devices. If you are out of interrupts, you should look at what other devices you can disable (such as COM or LPT ports) to gain back a necessary interrupt for other devices.

The big advantage of either type of USB from an IRQ or a resource perspective is that the USB bus uses only one IRQ no matter how many devices (up to 127) are attached or how many USB ports are installed in your system. Therefore, you can freely add or remove devices from the USB without worrying about running out of resources or having resource conflicts.

If you aren't using any USB devices, you should turn off the port using your motherboard CMOS Setup so that the IRQ it was using will be freed. As we continue to move to USB-based keyboards, mice, modems, printers, and so on, the IRQ shortage will be less of a problem. As we have already seen, the elimination of the ISA bus in our systems will also go a long way toward solving this problem.

For the best performance—especially with high-capacity removable media, CD/DVD rewriteable drives, scanners, and printers—look for motherboards that have USB 2.0 (High-speed USB) ports. These ports are also completely backward compatible with USB 1.1 devices and provide better performance when several USB devices are in use simultaneously. You can also add a USB 2.0 card to an existing motherboard to provide the same benefits to an existing system.

Miscellaneous Boards

Some video cards ship with advanced software that allows special video features, such as oversized desktops, custom monitors, switch modes on-the-fly, and so on. Unfortunately, this software requires that the card be configured to use an IRQ. If your video card doesn't require an IRQ, I suggest you dispense with this unnecessary software and configure the card to free up the interrupt for other devices. However, keep in mind that many of the latest 3D accelerators must use an IRQ to enable their bus-mastering feature to work correctly. See your video card's manual for details.

Also related to video is the use of an MPEG decoder add-on card that works in addition to your normal graphics adapter. These are used more in specialized video production and editing and in playing

DVD movies; however, they do use additional system resources that must be available. If your CPU runs at speeds above 300MHz and you have an AGP video card, you can probably use your video card for DVD playback and remove the MPEG decoder card. You will need a DVD player program, which might be supplied with your video card. See your video card's manual for details.

Plug-and-Play Systems

Plug and Play (PnP) represents a major revolution in interface technology. PnP first came on the market in 1995, and most motherboards and adapter cards since 1996 take advantage of it. Prior to that, PC users were forced to muddle through a nightmare of DIP switches and jumpers every time they wanted to add new devices to their systems. The results, all too often, were system resource conflicts and nonfunctioning cards.

PnP was not an entirely new concept. It was a key design feature of MCA and EISA interfaces that preceded it by almost 10 years, but the limited appeal of MCA and EISA meant that they never became true de facto industry standards. Therefore, mainstream PC users still had to worry about I/O addresses, DMA channels, and IRQ settings. Early PCI-based systems also used a form of PnP configuration, but because there was no provision for managing conflicts between PCI and ISA cards, many users still had configuration problems. But now that PnP has become prevalent, worry-free hardware setup is available to all computer buyers.

For PnP to work, the following components are desired:

- PnP hardware
- PnP BIOS
- PnP operating system

Each of these components needs to be PnP-compatible, meaning that it complies with the PnP specifications.

The Hardware Component

The *hardware component* refers to both computer systems and adapter cards. The term does not mean, however, that you can't use your older ISA adapter cards (referred to as *legacy cards*) in a PnP system. You can use these cards; in fact, your PnP BIOS automatically reassigns PnP-compatible cards around existing legacy components. Also, many late-model ISA cards can be switched into PnP-compatible mode.

PnP adapter cards communicate with the system BIOS and the operating system to convey information about which system resources are necessary. The BIOS and operating system, in turn, resolve conflicts (wherever possible) and inform the adapter card which specific resources it should use. The adapter card then can modify its configuration to use the specified resources.

The BIOS Component

The BIOS component means that most users of pre-1996 PCs need to update their BIOSs or purchase new machines that have PnP BIOSs. For a BIOS to be compatible, it must support 13 additional system function calls, which can be used by the OS component of a PnP system. The PnP BIOS specification was developed jointly by Compaq, Intel, and Phoenix Technologies.

The PnP features of the BIOS are implemented through an expanded POST. The BIOS is responsible for identification, isolation, and possible configuration of PnP adapter cards. The BIOS accomplishes these tasks by performing the following steps:

1. Disables any configurable devices on the motherboard or on adapter cards
2. Identifies any PnP PCI or ISA devices

3. Compiles an initial resource-allocation map for ports, IRQs, DMAs, and memory
4. Enables I/O devices
5. Scans the ROMs of ISA devices
6. Configures initial program-load (IPL) devices, which are used later to boot the system
7. Enables configurable devices by informing them which resources have been assigned to them
8. Starts the bootstrap loader
9. Transfers control to the operating system

The Operating System Component

The operating system component is found in most modern operating systems, such as Windows 9x/Me/2000/XP. In some cases system manufacturers have provided extensions to the operating system for their specific hardware. Such is especially true for notebook systems, for example. Be sure you load these extensions if they are required by your system.

It is the responsibility of the operating system to inform users of conflicts that can't be resolved by the BIOS. Depending on the sophistication of the operating system, the user then could configure the offending cards manually (onscreen) or turn off the system and set switches on the physical cards. When the system is restarted, the system is checked for remaining (or new) conflicts, any of which are brought to the user's attention. Through this repetitive process, all system conflicts are resolved.

Note

Because of revisions in some of the Plug and Play specifications, especially the ACPI specification, it can help to ensure you are running the latest BIOS and drivers for your system. With the Flash ROM used in most PnP systems, you can download the new BIOS image from the system vendor or manufacturer and run the supplied BIOS update program.

Motherboard Selection Criteria (Knowing What to Look For)

I am often asked to make a recommendation for purchases. Without guidance, many individuals don't have any rhyme or reason to their selections and instead base their choices solely on magazine reviews or, even worse, on some personal bias. To help eliminate this haphazard selection process, I have developed a simple checklist that will help you select a system. This list takes into consideration several important system aspects overlooked by most checklists. The goal is to ensure that the selected system truly is compatible and has a long life of service and upgrades ahead.

It helps to think like an engineer when you make your selection. Consider every aspect and detail of the motherboards in question. For instance, you should consider any future uses and upgrades. Technical support at a professional (as opposed to a user) level is extremely important. What support will be provided? Is there documentation, and does it cover everything else?

In short, a checklist is a good idea. Here is one for you to use in evaluating any PC-compatible system. A system might not have to meet every one of these criteria for you to consider purchasing it, but if it misses more than a few, consider staying away from that system. The items at the top of the list are the most important, and the items at the bottom are perhaps of lesser importance (although I think each item is important). The rest of this chapter discusses in detail the criteria in this checklist:

- *Motherboard chipset.* Motherboards should use a high-performance chipset that supports DDR SDRAM DIMMs or RDRAM RIMMs—preferably one that allows ECC memory as well if you are concerned about catching possible memory errors before they corrupt your data. Also look for

AGP 4x or faster video support and ATA-100 or faster hard drive support. The motherboard chipset is the backbone of a system and is perhaps the single most important part you'll consider. I spend the most time deciding on my next chipset because it affects and influences virtually every other component in the system.

- *Processor.* A modern system should use a socket-based processor with on-die L2 cache. Evaluate the processor choices you have, and try to get the one with the highest speed CPU bus (front-side bus). Don't get too hung up on L2 cache size; a little cache goes a long way. It is more important that the cache run at full core speed (which it will if it is on-die). Current processors such as the Athlon XP, Pentium 4, Celeron 4, and Athlon 64 all meet this criteria. I usually recommend only "boxed" processors as sold by Intel and AMD, which include a high-quality active heatsink as well as installation instructions and a 3-year warranty direct with the manufacturer.
- *Processor sockets*. For maximum upgradeability and performance, you should stick with a system that uses a socket for the CPU. The main sockets in use today on new systems include Socket A (Socket 462) for the Duron/Athlon/Athlon XP, Socket 478 (which has replaced the original Socket 423) for the Pentium 4 and Celeron 4, and Socket 754 for the new Athlon 64. As long as your motherboard has one of these sockets, you should be in good shape.
- *Motherboard speed.* The motherboard typically offers a choice of speeds, including anywhere from 200MHZ to 333MHz for the Duron/Athlon/Athlon XP-based boards, or from 400MHz to 800MHz for the Pentium 4–based boards. Check to ensure the board you are buying runs at the speeds necessary to support the processors you want to install.
- *Cache memory.* All modern systems use processors with integral cache, most of them now having the cache directly on the processor die for maximum speed. As such, there won't be any cache memory on the motherboard in a modern system. The tip is to make sure you are using a processor with full core speed on-die L2 cache because this offers the maximum in performance. All the modern processors now incorporate full-speed on-die L2 cache.
- *SIMM/DIMM/RIMM memory.* SIMMs are obsolete by today's standards, so stay away from any boards that use them. Your board should support either standard or DDR DIMMs or RIMMs that contain SDRAM, DDR SDRAM, or RDRAM, respectively. What you use depends mainly on your motherboard chipset, so choose the chipset and board that accepts the memory type you want to use. Currently, DDR SDRAM and RDRAM are the fastest types of memory available, with RDRAM being by far the most costly.

  Mission-critical systems should use ECC memory and ensure that the motherboard fully supports ECC operation. Note that many of the low-end chipsets from Intel and others do not support ECC and should not be used for mission-critical applications. This is something you should know before purchasing the system.

  Finally, note that most motherboards support either three or four DIMM sockets, or two or three RIMM sockets. Be sure that you populate them wisely so you don't have to resort to removing memory later to add more, which is not very cost-effective.
- *Bus type.* No recent desktop motherboards incorporate ISA bus slots, so if you have any ISA cards, you'll have to discard them or use them in older systems. Instead you should find anywhere from one to five or more PCI local bus slots. Be sure the PCI slots conform to the PCI 2.1 or later revision (primarily based on the chipset). Take a look at the layout of the slots to ensure that cards inserted in them will not block access to memory sockets or be blocked by other components in the case. Systems without onboard video should also feature one AGP 4x or 8x slot for a high-performance AGP video card. Some boards also feature AMR (audio modem riser) or CNR (communications networking riser) slots for special cards that are included with the board to provide sound, modem, or other similar features.

- *BIOS.* The motherboard should use an industry-standard BIOS, such as those from AMI, Phoenix, or Award. The BIOS should be of a Flash ROM or EEPROM design for easy updating. Look for a BIOS Recover jumper or mode setting, as well as possibly a Flash ROM write-protect jumper on some systems.
- *Form factor.* For maximum flexibility, performance, reliability, and ease-of-use, the ATX form factor (including micro-ATX and flex-ATX) cannot be beat. ATX has several distinct performance and functional advantages over Baby-AT and is vastly superior to any proprietary designs, such as LPX. Additionally, the NLX form factor might be a consideration for low-profile or low-cost desktop systems, although flex-ATX is also becoming popular for these designs.
- *Built-in interfaces.* Ideally, a motherboard should contain as many built-in standard controllers and interfaces as possible (except perhaps video). There is a trend toward legacy-free PCs that lack the conventional Super I/O component and therefore have only USB and sometimes IEEE-1394 for external expansion. Legacy-free PCs lack the conventional keyboard and mouse ports, serial and parallel ports, and possibly even the internal floppy controller. Systems that use an integrated Super I/O component have these interfaces.

  Built-in 10/100 Ethernet network adapters are also handy, especially if you are using a cable modem or DSL connection to the Internet. A built-in sound card is a great feature, usually offering full Sound Blaster compatibility and functions, and possibly offering additional features as well. If your sound needs are more demanding, you might find the built-in solutions less desirable, and you might want to have a separate sound card in your system. Built-in video adapters are also a bonus in some situations, but because there are many video chipset and adapter designs from which to choose, generally you'll find better choices in external local bus video adapters. This is especially true if you need the highest performance video available.

  Built-in devices usually can be disabled to allow future add-ons, but problems can result.
- *Onboard IDE interfaces.* All motherboards on the market have included onboard IDE interfaces for some time now, but not all IDE interfaces are equal. Your motherboard should support at least UDMA/66 (ATA-66) speeds, which matches the best real-world performance currently available from IDE drives. UDMA/66, UDMA/100, and UDMA/133 actually exceed the real-world performance of current drives using these standards and thus provide you with headroom for future drives. For even greater speed, consider motherboards with onboard IDE RAID controllers. These motherboards can be configured to perform data striping for extra speed or data mirroring for extra reliability when two or more identical IDE drives are used. These motherboards are based on a variety of standard chipsets with RAID functions added by RAID chipsets from AMI, HighPoint, or Promise. Many recent systems now include Serial ATA drive interfaces, some of which include RAID functions. SATA is even faster than ATA-133 and exceeds the real-world performance of current (first-generation) SATA drives by a wide margin.

Tip

With a never-ending stream of motherboards coming onto the market, finding motherboards with the features you want can be difficult. Motherboard Homeworld's Mobot search engine helps you find motherboards based on your choice of form factor, platform, chipset, CPU type, processor, manufacturer, memory type, slot types, built-on ports, and more. Check it out at `iceberg.pchomeworld.com/cgi-win/mobotGen/mobot.asp`.

- *Power management.* The motherboard should fully support the latest standard for power management, which is ACPI. An Energy Star–compliant system is also a bonus because it uses less than 30 watts of electrical energy when in sleep mode, saving energy as well as your electric bill.
- *Documentation.* Good technical documentation is a requirement. Documents should include information on any and all jumpers and switches found on the board, connector pinouts for all connectors, specifications for other plug-in components, and any other applicable technical information.

You might notice that these selection criteria seem fairly strict and might disqualify many motherboards on the market, including what you already have in your system! These criteria will, however, guarantee you the highest-quality motherboard offering the latest in PC technology that will be upgradable, be expandable, and provide good service for many years.

Most of the time I recommend purchasing boards from better-known motherboard manufacturers such as Intel, Acer, ABIT, AsusTek, SuperMicro, Tyan, FIC, and others. These boards might cost a little more, but there is some safety in the more well-known brands. That is, the more boards they sell, the more likely that any problems will have been discovered by others and solved long before you get yours. Also, if service or support is necessary, the larger vendors are more likely to be around in the long run.

Documentation

As mentioned, documentation is an important factor to consider when you're planning to purchase a motherboard. Most motherboard manufacturers design their boards around a particular chipset, which actually counts as the bulk of the motherboard circuitry. Many manufacturers, such as Intel, VIA, ALi, SiS, and others, offer chipsets. I recommend obtaining the data book or other technical documentation on the chipset directly from the chipset manufacturer.

For example, one of the more common questions I hear about a system relates to the BIOS Setup program. People want to know what the "Advanced Chipset Setup" features mean and what the effects of changing them will be. Often they go to the BIOS manufacturer thinking that the BIOS documentation will offer help. Usually, however, people find that there is no real coverage of what the chipset setup features are in the BIOS documentation. You will find this information in the data book provided by the chipset manufacturer. Although these books are meant to be read by the engineers who design the boards, they contain all the detailed information about the chipset's features, especially those that might be adjustable. With the chipset data book, you will have an explanation of all the controls in the Advanced Chipset Setup section of the BIOS Setup program.

Besides the main chipset data books, I also recommend collecting any data books on the other major chips in the system. This includes any floppy or IDE controller chips, Super I/O chips, and of course the main processor. You will find an incredible amount of information on these components in the data books.

Caution

Most chipset manufacturers make a particular chip for only a short time, rapidly superseding it with an improved or changed version. The data books are available only during the time the chip is being manufactured, so if you wait too long, you will find that such documents might no longer be available. The time to collect documentation on your motherboard is *now*!

Using Correct Speed-Rated Parts

Some vendors use substandard parts in their systems to save money. Because the CPU is one of the most expensive components on the motherboard and motherboards are sold to system assemblers without the CPU installed, it is tempting for the assembler to install a CPU rated for less than the actual operating speed. A system could be sold as a 900MHz system, for example, but when you look under the hood, you might find it's rated for only 600MHz. This is called *overclocking*, and many vendors have practiced this over the last few years. Some even go so far as to re-mark the CPUs, so that even if you look, the part appears to have the correct rating. The best way to stop this is to purchase systems from known, reliable vendors and purchase processors from distributors that are closely connected with the manufacturer. Overclocking is fine if you want to do it yourself and understand the risks, but when I purchase a new system, I expect that all the parts included will be rated to run at the speed to which they are set.

◀◀ See "Processor Speed Ratings," p. 50.

When a chip is run at a speed higher than it is rated for, it runs hotter than it would normally. This can cause the chip to occasionally overheat, which would appear as random lockups, glitches, and frustration. I highly recommend that you check to ensure you are getting the right speed-rated parts you are paying for.

Also be sure to use the recommended heatsink compound (thermal grease). This can improve the efficiency of your heatsink by up to 30%.

This practice is easy to fall into because the faster-rated chips cost more money. Intel and other chip manufacturers usually rate their chips very conservatively. Over the years, I have overclocked many processors, running them sometimes well beyond their rated speeds. Although I might purchase a Pentium III 800 and run it at 1066MHz, if I were to experience lockups or glitches in operation, I would immediately return it to the original speed and retest. If I purchase a 1GHz system from a vendor, I fully expect it to have a 1GHz part, not slower parts running past their rated speeds!

Overclocking has been made more difficult by Intel and AMD, who have both started locking the bus multipliers in their chips to prevent easy overclocking by changing the multiplier setting on the motherboard. This is done mainly to combat re-marking CPUs and deceiving customers, although it unfortunately can also prevent those who want to from hotrodding their chips. Still, you can overclock most chips by increasing the CPU bus (front-side bus) speed within certain tolerances. Many of the motherboards on the market have tweakable CPU bus speeds specifically designed to allow overclocking. Check with your motherboard manual, or download the documentation from the manufacturer's Web site. You might find that your board is capable of things you didn't realize.

If you purchase a processor or system, verify that the markings are the original Intel or AMD markings and that the speed rating on the chip is what you really paid for.

The bottom line: If the price is too good to be true, ask before you buy. Are the parts really manufacturer-rated for the system speed?

To determine the rated speed of a CPU chip, look at the writing on the chip. See Chapter 3 for details on how to interpret the marks to see what the rating on the chip actually is.

Caution

Be careful when running software to detect processor speed. Most programs can only estimate at what speed the chip is currently running, not what the true original rating is. One exception to this is the Intel Processor Frequency ID Utility, which can determine whether an Intel processor is operating at the correct and rated frequency intended. Although it gives only basic information about any Intel processor, it can uniquely identify the original speed ratings of the Pentium III, third-generation Celeron (Coppermine-based), and any newer processors, accurately determining whether they have been overclocked. You can download the Intel Processor Frequency ID Utility from `developer.intel.com/support/processors/tools/frequencyid/download.htm`.

Older system chassis with speed markings or even indicator lights are usually no indication of the actual or rated speed of the processor inside. Those displays can literally be set via jumpers to read any speed you desire! They have no true relation to actual system speed.

Most of the better diagnostics on the market, such as Norton Utilities from Symantec, read the processor ID and stepping information, as well as show current operating (but not rated) speed. You can consult the processor manufacturer or Chapter 3 for tables listing the various processor steppings to see exactly how yours stacks up.