Chapter 4
Technology Components of Platform Architecture

This chapter and the two chapters that follow are intended to provide the technology environment and platform context in which the microprocessor of Chapters 2 and 3 functions. The material here in Chapter 4 examines the interactions between the microprocessor and the balance of the platform, including hardware and software relationships, but it continues to be hardware centric. Chapter 5, Platform Memory Technology, examines memory types, memory controllers, and concludes with sections on Rambus and SLDRAM. Chapter 6, Platform Optimizations and Directions, ties together the material in Chapters 4 and 5, and examines performance issues and forces that are driving platform evolution.

This chapter builds on the platform overview provided in Chapter 1 (which we suggest you review at this point) and provides background information and overviews for the important technologies, standards, and initiatives impacting PC platforms. It should provide literacy with respect to the technologies examined and thereby facilitate further in-depth study of each technology. It will also act as an introduction to less-accessible industry standards. This chapter also provides the reader with information on the groups that control and publish the various platform standards and how the standards may be obtained.

This chapter is built around three major sections: PC Design Metastandards, Platform-Level Technology, and Component-Level Technology. The PC Design Metastandards section introduces the on-going series of industry endorsed PC Design Guides and Specifications, and describes the Microsoft logo certification process. The Platform-Level Technology section focuses on technologies, standards, and initiatives that have a major impact on multiple components or subsystems or the system as a whole. The Component-Level Technology section focuses on technologies, standards, and initiatives that tend to be centered in a single or few

components or subsystems. The intended audience for each section in the chapter is indicated in the Road Map below.

ROAD MAP OF CHAPTER 4

Section	Audience
PC Design Metastandards	
Forces Driving Platform Architecture PC Design Guides and Specifications	Students and Faculty
Platform Categories	Those unfamiliar with PC platforms
Platform-Level Technology Enhanced User Experience Appliance-Like Operation Total Cost-of-Ownership Connectivity	Those unfamiliar with PC industry initiatives
Component-Level Technology Processors General-Purpose Buses Device-Specific Buses and Ports	Practitioners, managers, students, and faculty who want to study outside their area of specialization

PC DESIGN METASTANDARDS

PC platform technology is being driven forward by underlying economic and technology forces and by recent PC industry metastandards and certification programs for PC design. The economic and technology forces were present long before the industry standards and certification were in place and would continue to drive the platform forward in the absence of the standards. Nevertheless, one should not underestimate the importance of the present PC Design Specifications. Unlike previous PC industry attempts at metastandards, the present PC Design Specifications are more than just a codification of what PCs already are. They are now co-driving the PC technology evolution by aggressively raising the mandatory baseline functionality that PC platform designers must implement. The underlying economic and technology forces are still there, but the evolution is being driven faster in the short term than it would be otherwise, as a direct result of the PC Design Specifications and the certification process.

FORCES DRIVING PLATFORM ARCHITECTURE

From its inception, the PC platform has been perceived as a very flexible, but basic, programmable machine that the owner or user could customize

via application software. This view was typified by early use of the platform as a near-dedicated word processor, spread sheet, data terminal, instructional aid, or game machine. The continuous but largely uncoordinated efforts at improving the PC platform's diverse roles as a publishing tool, a business machine, a tool for personal productivity, a self-instruction education device, or an entertainment machine, have been long established primary forces evolving the architecture. Another well established primary force behind platform evolution has been the necessity to overcome severe installation, configuration, and reliability problems stemming from the way in which the original I/O architecture was defined and used. More recently, PC platforms in different specialty forms have been emerging as preferred tools and appliances for the creation, collaboration, storage, distribution, retrieval, and presentation of content. The efforts to accelerate the realization of this vision are becoming additional primary forces.

Secondary Forces Driving Platform Architecture

Several secondary forces have driven and continue to drive the PC platform. A first of these is the economics of advancing semiconductor and magnetic-recording technologies, which have provided dramatically increased performance at constant and generally decreasing prices.

Another force is the keen desire of the PC and related industries to continue to grow the PC platform market at a double-digit rate. For at least the semiconductor sector, strong growth is necessary to offset continued decreasing sales prices and to finance the enormous investments required in factory and tooling costs. Strategies for accomplishing such growth are to increase platform affordability, accessibility, and acceptability in the general population, and to capture market share from other recreational pursuits. Platform affordability is improved by providing increased baseline functionality at constant or decreasing prices. Accessibility and acceptability are addressed by making platforms similar to appliances in their ease of operation, convenience, and unobtrusiveness. Market share capture from other recreational pursuits is accomplished by offering superior experiences or comparable experiences at a better value.

Additional secondary forces driving platform architecture include the adoption of the PC as a standard piece of office equipment, keen industry competition, and the transition of the platform to an item of mass production. Keen competition is most prevalent in common applications software, packaged systems, and such components as memory, hard disks, and system logic. Mass production has the expected volume-pricing benefits, but it also has attendant compatibility burdens. These burdens arise because standards with a large installed base are difficult to supplant with new standards, even when the new standards are superior.

PC Platform Evolution

The PC platform has evolved in many ways from the PC/AT archetype. Mainstream users interacted with the original ISA-Bus-based platform predominantly via a 23-row by 80-column character-only display and a keyboard. The keyboard was required for communicating commands and file names to the operating system and applications. The keyboard was also the principal means of moving the data entry point about the display screen. Printer output was also predominantly character oriented. Dial-up communications were available, but modems were expensive and slow.

The Rise of Graphics, Multimedia, and Networking

The later half of the 80's saw the augmentation of the mainstream PC platform with capabilities for graphics, multimedia, networking, and content distribution. The open architecture of the ISA Bus contributed to the availability of affordable peripherals, which in turn facilitated the adoption of new capabilities and the success of the platform as a whole. Following in the footsteps of the Apple Macintosh, graphics features were applied to the control of the operating system and applications and for the creation of documents and other works. These features included the ability to display, create, and print graphics; to move and select screen objects via a hand-held pointer; to use a wide range of typeface styles and sizes; and the ability to imbed both images and text on the same page.

These new graphics features were motivated by a desire to make computers easier to use and to enable *desktop publishing*. Desktop publishing required the capability to create virtually any kind of printed work and to have the final printed form of the work match exactly what was created on the screen.[40] These features required a migration to high-resolution color graphics displays and printers, larger hard-disks, new applications, and support from the operating system.

Multimedia features included the ability to have simultaneous use of text, images, video, and audio and sound effects. Multimedia features were motivated by the desire to exploit, experiment with, and be entertained by a new genre of dynamic and interactive works that was much more effective at communication than the static printed page. Multimedia offered exciting games, attention grabbing education, and compelling business presentations. Multimedia required high data-transfer rates and it was soon realized that the ISA Bus had insufficient bandwidth headroom to provide quality results. The data storage requirements for multimedia works were also high and necessitated the widespread adoption of the CD-ROM as a storage peripheral to augment the hard disk.

---

[40] This ability to have print output exactly match the screen is known as *What You See Is What You Get* (WYSIWIG).

During this same period, networking evolved from a specialty add-on to an integral component of the operating system. In addition, consumer use of modem communications rose PC in tandem with improving speeds and the availability of dial-up services providing consumer oriented content.

The System Resource Crisis

The rise of graphics, multimedia, and networking features frequently resulted in what was perceived to be a shortage of available system resources. These resources included expansion slots, certain classes of control lines, and other system resources[41] generally necessarily assigned to each peripheral for proper operation. The I/O paradigm from the beginning of the PC platform was that each expansion device generally used a dedicated interrupt resource to signal the microprocessor that it needed servicing in a timely fashion. There was no defined way of sharing interrupts that ensured that conflicts would not result, or that each of the devices sharing the interrupt would be serviced within any particular time interval. While expansion slot availability was solved in part by integrating some popular expansion peripherals into the standard logic of the system, the shortage of other system resources continued. The resource shortage resulted in considerable loses in productivity and goodwill toward the PC platform, due to difficulties in installation and maintenance and instabilities and unreliability due to marginal operation. The perception arose that in addition to its performance limitations, the small number of critical resources defined for the ISA Bus's I/O architecture was going to ultimately limit the success of the platform.

Focus on Performance for Graphics and Multimedia

The early 90's saw continued advancements in the performance of the PC platform in many areas and attempts to address the system resource configuration problem. PCs became highly connected, with ubiquitous networking for office PCs and ubiquitous modem communications for home PCs. Graphics acceleration became common for speeding complex screen updates. The PCI Bus was added to provide improved peak bandwidth for high speed peripherals, such as high speed hard disks and the new graphics accelerators. The PCI Bus also provided needed low latency capability and bandwidth headroom for the multiple latency-intolerant-processes associated with multimedia. The view arose that one used the PCI Bus for

[41] The control lines mentioned are associated with the ISA Bus's interrupt request and Direct Memory Access features. The other system resources mentioned included areas in the I/O address space and reserved areas in the memory address space.

high performance peripherals, but one could continue to rely on the ISA Bus for inexpensive low-performance devices.

Early Automatic Configuration Efforts

The system-resource configuration-problem was partially mitigated through an industry initiative known as Plug and Play, which placed new requirements on both hardware and software. Hardware devices, whether integrated in the system logic or on expansion cards, were required to report their resource needs to the operating system and be flexibly configurable under operating system control. Application programs were required to be as flexible as possible and the operating system managed the assignment of resources.

Call for Abandonment of Legacy I/O

In spite of a lessening of configuration problems due to the Plug and Play initiative, the installation, maintenance, and reliability of the PC platform continues to be a serious problem that is blamed on legacy-based I/O. Furthermore, the perception has arisen that as performance requirements have increased, the ISA Bus is not merely inadequate in performance, but the very use of the ISA Bus for low performance peripherals is limiting the performance realizable from the system as a whole.

Recently, new peripheral bus architectures, including the Universal Serial Bus and IEEE 1394, have emerged that redefine the use of system resources and physical connections to the platform in a far more efficient manner than legacy I/O. In the new paradigm of resource allocation, many diverse devices communicate via packet messaging to a host controller in the PC. The communications protocol of these buses are defined to ensure that any device that so requires it may be serviced within guaranteed maximum latencies. With these new buses, individual interrupts per device are neither required nor available.

To ensure completely that no conflicts or instabilities arise due to legacy peripherals, the present PC Specifications require not only the adoption of the new buses, but the complete abandonment of legacy I/O architecture devices and the ISA Bus. This radical departure from legacy hardware compatibility also ensures that the use of legacy I/O devices will no longer limit system performance or create installation and maintenance difficulties. The new peripheral bus architectures are also much more versatile, supporting carefree attachment and removal of peripheral devices. Such events required significant reconfiguration and loss of productivity in the legacy I/O architecture.

Constants of Change

Throughout the evolution of the PC platform, speeds have continually increased for processors, memory, and peripherals, as have memory and storage capacities. These advances have been partially offset by the demands of increasingly sophisticated operating systems and applications. The practical and reliable operation of both types of software would not be possible without the attendant increases in processors and memory speed and the capacity of memory and storage.

PC Design Guides and Specifications

The PC Design Guides and Specifications are a running series of published guides and specifications that have become the overarching standards governing all aspects of PC platform architecture. The PC 98 guide was over 600 pages in length, although a summary appendix provides a condensed requirements checklist of slightly over 30 pages. The specifications consist of the published guides in conjunction with amendments and associated compliance schedules. The PC Specifications are followed closely by all major PC manufacturers.

While the first two guides were authored solely by Microsoft, Microsoft and Intel began co-authoring the guides starting with the *PC 98 System Design Guide*. Recent guides have been initially posted on the Web in draft form and reviewers from the general industry have been encouraged to provide feedback and suggested changes that help mold the final specification. The guides systematically detail requirements for performance, interoperability, ease of use, and ergonometrics. These requirements cover all of the platform's subsystems, buses, and devices. This section examines why the guides and specifications have achieved the significance they have and provides an overview of their organization.

Audio on CD-ROM

Legacy Software Compatibility: The CD-ROM includes a short audio-segment that describes some implicit PC platform requirements.

Logo Certification

Microsoft's "Designed for Windows"

The PC specifications and their associated design guides have achieved their stature as a result of Microsoft's "Designed for Windows" logo certification program. Microsoft permits hardware vendors to use special Microsoft logos in the marketing of products that meet the specification guidelines. The Microsoft "Designed for Windows" logo is a very desirable

marketing feature for system vendors and other Independent Hardware Vendors (IHVs).

Win32 operating systems include the Windows 95, Windows 98, and Windows NT 32-bit operating systems.

Burned by incompatibilities and difficult installations, many consumers and corporate buyers rank products with the logo well above those products that do not have it. At Microsoft's WinHEC '97[42], Microsoft said that 273 PC manufacturers, 348 hardware vendors, and 267 software vendors had been certified to use a Microsoft *Designed for Windows* logo for its Win32 operating systems.

The product-by-product license to use Microsoft's "Designed for" logos and listing of the vendor's product on Microsoft's Windows Hardware Compatibility List (HCL), occurs only after the vendor's product passes qualification by Microsoft's Windows Hardware Qualification Labs (WHQL). The WHQL develops its Hardware Compatibility Tests (HCT) from the requirements in the design guide, amendments to the guide, and a compliance schedule. The amendments and compliance schedule are posted on Microsoft's Web site. Passing the HCT involves significant effort on the part of the vendor, including self pretesting, satisfying the Test Submission requirements (including providing pretest logs and many duplicates of the system or device to be tested), signing the relevant logo license agreements, and payment of test fees.

Each subsequent guide often incorporates, by reference, sections from earlier guides in the series. To design a compliant system and fully understand each guide, you generally need access to more than the most recent guide.

Products submitted for testing on or after July 1st of each year must meet the PC Specifications for that year as detained in the PC Design Guides and as amended online. However, several requirements usually are often not enforced until a later date. This occurs when it is perceived that the requirements are too aggressive for earlier enforcement. The compliance schedules for such postponed enforcement items are also posted online.

Report on CD-ROM

To learn more about Microsoft Windows Logo Certification, see the article by Peter N. Glaskowsky, "WinHEC Shows Road to Memphis," *Microprocessor Report*, Vol. 11. No. 6, May 12, 1997, on the CD-ROM.

Summary of Principal System Requirements

The core system requirements for certification of PCs for several of the past specifications are shown in Table 4.1. The staged phasing-out of legacy devices, the elimination of the ISA Bus, and the adoption of replacement peripheral buses are quite evident. The continued evolution of

[42] WinHEC is Microsoft's Windows Hardware Engineering Conference, which is held twice a year.

increased processor performance, memory size, and graphics capabilities are unmistakable as well. Note that generally the changes between each year's requirements are incremental. Recommendations are used to signal future major design shifts, with major technology components usually being been phased in (or out) over multiple years. The requirement trends for performance and feature attributes that have changed rapidly in the past (e.g., processor performance, main memory size, and graphics resolution) also likely will continue and designers can extrapolate these trends and make expectations for future requirements.

Platform Category Definitions

Figure 4.1 shows a Venn diagram of the subcategories defined by the PC Specifications and the various groupings into which they are associated. The overall consumer category includes the Consumer PC and Entertainment PC subcategories. The overall business category includes the Office PC, Mobile, Workstation, and NetPC subcategories. Each of these subcategories is explained below in the Platform Categories section.

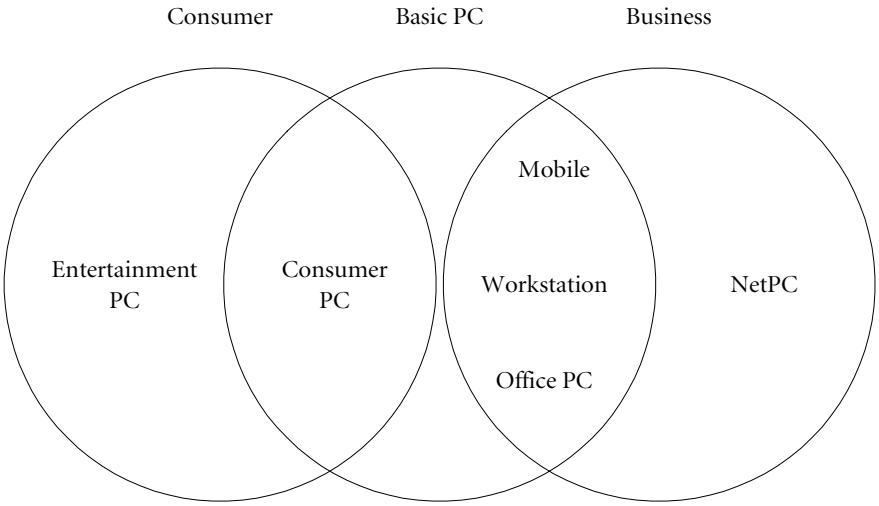

Figure 4.1 PC Platform Categories and Subcategories

Basic PC

A Basic PC meta-category is defined to include the Consumer PC, Office PC, Mobile, and Workstation subcategories. These Basic PC platforms all generally share many common hardware characteristics and more importantly what has become the basic operating paradigm for PC platforms. To the extent that this is so, the Basic PC category serves as a useful abstraction, and acts to distinguish the Entertainment and NetPCs as truly distinct species from conventional PC platforms.

Table 4.1 Comparison of Core System Requirements for Different PC Specifications

Feature[a]	PC 95[b] [c]	Basic PC 97	Basic PC 98	Basic PC 99
Equivalent Processor Performance	80386, 33 MHz 80486 recommended	120 MHz P5	200 MHz P55C (with MMX)	300 MHz PII (with MMX)
Main Memory Size	4 MB, 8MB recommended	16 MB, 32 MB recommended	32 MB, 64 MB at 66 MHz recommended	32 MB, 64 MB at 100 MHz recommended
External Cache	(not mentioned)	recommended	256 KB	256 KB
Graphics	640 x 480 x 8 bits per pixel (bpp) recommended, high-speed bus (not ISA) recommended	800 x 600 x 16 bpp, uses high-speed bus, does not use ISA, multiple display adapters and monitors supported	1024 x 768 x [8,15,16] bpp and lower Video Electronics Standard Association modes, otherwise same as PC 97	same as PC 98, AGP 2.0 recommended
2D or 3D Graphics Acceleration	(not mentioned)	(not mentioned)	if present, 2D/3D rendering must have double buffering for up to 800 x 600 x 16 bpp. 3D rendering must also have a Z buffer [d] and 1.25 MB local texture cache or equivalent AGP capability	same as PC 98, Z buffer must be at least 16-bits
ISA Bus	ISA devices must be Plug and Play compliant, 16-bit I/O decoding required	same as PC 95, but not recommended for any device	no populated ISA slots when sold, not recommended for motherboard devices	No ISA slots or motherboard devices
High-speed Expansion Bus	recommended, PCI must be Revision 2.0 compliant	recommended, PCI must be Revision 2.1 compliant	SCSI or IEEE 1394 recommended, PCI must be Revision 2.1 compliant	IEEE 1394 recommended, PCI must be Revision 2.2 compliant
Universal Serial Bus	(not mentioned)	required, 2 user accessible ports recommended	1 user accessible port required, 2 recommended	2 user accessible ports required
IEEE 1394 Bus	(not mentioned)	recommended	recommended	required, recommended as a secondary host controller
Device Bay	(not mentioned)	(not mentioned)	recommended	recommended

[a] Each of SCSI, IEEE 1394, the Universal Serial Bus, and Device Bay are discussed later in this chapter.

[b] There was no PC 96 guide.

[c] Platform designers must rely on official guides as amended online. See http://www.microsoft.com/hwdev/.

[d] A *Z buffer* is a memory used in 3D rendering to facilitate the relative foreground-to-background ordering of modeled objects.

Consumer PC

The Consumer PC is essentially a desktop platform targeted at and optimized for the consumer market segment. Consumer PC platforms are expected to have Internet connectivity via dial-up modem, be used primarily for entertainment, education, and personal productivity applications, have strong graphics capability, and be digital-TV and audio ready.

Office PC

manageability features

The Office PC is essentially a desktop platform targeted at and optimized for the commercial market segment. Such platforms are connected to a Local Area Network (LAN) and run productivity applications. Each Office PC has manageability features, such as Advanced Configuration and Power Interface support, centralized administration, and is upgradable for remote boot capability. It is intranet, Microsoft Zero Administration Initiative for Windows (ZAW), and conferencing ready. AGPI is discussed in the Appliance-Like Operation section. ZAW is discussed in the Total Cost-of-Ownership section.

Industry Standards

PC Design Guides

The most recent guides are available in most technical bookstores and can be ordered direct from Microsoft Press at URL:

http://mspress.microsoft.com/

Microsoft provides free access to formatted text files and online navigable versions of the guides, along with up-to-date information about logo compliance dates, under the following URL:

http://www.microsoft.com/hwdev/desguid/

Microsoft "Designed for Windows" Logo Qualification Standards

More information about the Hardware Compatibility Tests (HCT) Test Suite and how to order a CD-ROM containing the suite is at URL:

http://www.microsoft.com/hwtest/

Suggested Readings

PC Design Guides

The past design guides can provide additional understanding and collectively they chronicle an important era in the evolution of the PC platform:

1. *Hardware Design Guide for Microsoft Windows 95*, Microsoft Press, 1994.
2. The *PC 97 Hardware Design Guide,* Microsoft Press, 1996.
3. The *PC 98 System Design Guide,* Microsoft Press, 1997.

Platform Categories

This section surveys conventional platform types and market segments, and platform categories as defined by the PC Specifications. The relationships and distinctions between each of these is important to understanding the PC marketplace and our analysis of industry initiatives and platform-and component-level technologies.

Platform Types

New content and connectivity uses constitute new subsegments of platform and market types. The continual redefinition of platform types and market segments is discussed in the discussed in the Contemporary Platform Strategic Issues section of Chapter 6.

A specific PC is designed primarily for use as a desktop, workstation, server, mobile, entertainment appliance, or network PC. These distinctions between platform types are idealizations that are often blurred in practice. For example, mobiles are frequently used as desktop surrogates. Desktops are pressed into service as low-price servers, particularly in small office peer-to-peer networking situations.

Desktops

Desktops (including consumer and office desktops, mini-towers, and towers) are typically used predominantly by one person, for one or a few key applications, such as word-processing, database entry or management, or Web browsing. Desktops are available in a wide range of features, performance, and price. Currently, desktop platforms come in single-and-multiple processor configurations, have a wide range of advanced memory technology, have different levels of I/O integration, and offer multiple types of system and microprocessor buses.

Workstations

Workstations are high-end desktops, virtually always used by professional knowledge workers engaged in some form of critical intellectual property or content creation, be it software development, large desktop publishing projects, elaborate multimedia works, or engineering simulations.

Servers

A *server* is a networked high-performance PC that is a gateway for other PCs to one or more desired resources or concentrations of data. Specialized server categories included *Workgroup Servers, Departmental Servers,* and *Enterprise Servers*. Often a variant of the workstation class, *workgroup servers* may be found distributed throughout offices to provide file and printer sharing, often organized around the teams reporting to first-line managers and the staffs of higher-level management.

workgroup servers

departmental servers

Departmental servers are generally used in organized collections of servers used in divisional or secondary data centers, supporting site intra-

nets, project databases, and departmental services such as e-mail and remote access dial-in and dial-back modem connections. Departmental servers are often packaged for mounting in racks with other servers or networking equipment, and may share switched monitors, keyboards, and mice, among several other servers.

enterprise servers

Enterprise servers are an extension of departmental servers, but designed for use in the principal data centers of large enterprises. These servers support corporate intranets, large transactional databases, and other applications requiring centralized processing or control. They are characterized by the use of tightly coupled multiprocessors within the same platform, loosely coupled platform clusters, ultra high-capacity storage, and very high speed network connections.

Mobiles

portable

mobile

laptop

Mobile PCs (including portables, laptops, systems having integrated digital phones and future systems incorporating Personal Communication Services (PCS) features) are machines that can easily be moved about. A *portable* is a personal computer that is capable of being carried, and in which the motherboard, keyboard, pointing device, and display are generally integral to a single package. A *mobile* is a portable computer that can be used solely on batteries or directly or indirectly on 12V. A *laptop* is a mobile computer that can be used on one's lap. While these distinctions were once important, today the terms are used almost interchangeably because virtually all portable computers sold today meet the above definition of a laptop.

Entertainment PCs

Entertainment PCs are being positioned through industry design initiatives to become the favored centerpiece of consumer free-time activity. An offshoot of consumer desktops, the Entertainment PC was first introduced in PC 97 as the Entertainment Platform (EP), and was described as an "interactive multimedia system optimized for games, education, personal communications, and video playback."

The Entertainment PC is intended to be the ultimate PC for audio/visual and game enthusiasts, with the best graphics, video, and audio of any platform class. It is to be the platform of choice for cutting-edge applications, characterized by engaging rich detail and realism. The Entertainment PC will be the vehicle used to enhance TV via interactivity and better image quality, and it provides advanced e-mail, voice mail, and telephony. These latter applications require that this platform be as easy or easier to use than present consumer appliances found in the home.

Network PC

The PC Specifications require that Net PCs have hard disks. However, the hard disks are not used for resident storage of either system or application software or data, but for data caching to improve performance and reduce network traffic.

Network PCs or NetPCs are an emerging category envisioned as providing the lowest cost-of-ownership for massive commercial PC installations while addressing the typical needs of most knowledge and service workers. These *thin clients* rely on networked servers for most or all of their data storage and are designed specifically for remote configuration and management.

AUDIOS ON CD-ROM

Additional information about each Platform Type is included on the CD-ROM:

- Desktops
- Workstations
- Servers
- Mobiles
- Entertainment PC
- Network PC

Market Segments

Market segments are a device for abstracting the buying patterns and decisions of groups of PC buyers who tend to behave in a like manner. The PC Market was once simply conceptually divided between the Consumer and Commercial market segments. New platform uses have grown the market steadily to encompass new segments. For example, today the Commercial market may be considered to include Enterprise Computing, Workgroup Computing, Internet Service Provider, and Small Business and Emerging Markets Computing Segments. Across the various segments there is a wide variation in the scale of infrastructure for computing, storage, networking, and communications.

Consumer Segment

The *Consumer* segment has always been centered on providing entertainment or amusement, self-paced education (especially for school-age children), and enhancing personal productivity, while meeting acute cost sensitivities. Important subsegments are the gaming and Small Office/Home Office (SOHO) markets. Future subsegments are anticipated for home servers, computing appliances, and hybrid TV/PCs.

Commercial Segment

Enterprise Segment

The segments of the *Commercial* Business market are centered on enhancing the profits or service of companies and organizations via cost-effective personal-productivity and information-management tools. The *Enterprise Segment* characteristically has a client-server galaxy managed and linked

via a large routing and switching network. The network is used to logically integrate application and data resources and provide distributed and remote computing services including Internet and intranet access.

Workgroup Segment

The *Workgroup Segment* focuses on peer-to-peer networking within a local site, while providing connectivity to the centralized computing resources of a parent Enterprise. The *Internet Service Provider* Segment focuses on managed storage servers and Telephone Company ("Telco") access. The *Small Business and Emerging Market Segment* focuses on ease of use and the Internet access.

Platform-Level Technology

This major section and the one that follows (Component-Level Technology) examine the key standards, initiatives, and technologies that define present and emerging platform architecture. We distinguish platform-level technologies from component-level technologies by their scope. Platform-level technologies have a major impact on multiple components or subsystems or the system as a whole. Component-level technologies have a focus that is generally limited to a single type of component or subsystem.

As will come out in the individual discussions, many of the standards and initiatives have attributes that fall in multiple categories.

It is also the case that we view the platform-level technologies as representing technology directions. This contrasts with the component-level technologies that we view as generally representing specific technology destinations. Accordingly, we examine a number of platform-level standards, initiatives, and technologies organized under four categories that are our abstraction of the most significant thrusts in platform hardware. These four thrusts are the deliverance to the platform user or owner of an *Enhanced User Experience*, *Appliance-like Operation*, reduced *Total Cost-of-Ownership (TCO)* and *Connectivity* to networks and other platforms.

Enhanced User Experience

This section focuses on the efforts by the PC industry to improve the experiences that users accrue. These experiences are key to the continued growth of the PC market. Manufacturers are counting on users to be enthusiastic proponents of the platform to potential users, they want users to continually upgrade, and they want users to consider buying multiple PCs for diverse purposes, including supplanting other consumer entertainment electronics.

Providing enhanced user experiences is progressing along several fronts. The advent of Windows 3.X and Windows 95, brought great strides to the PC platform through a greatly improved user interface (UI) technology (discussed below) compared to the earlier command-line-oriented operating systems. The improved UI provided user ease of use and presentation quality, at levels previously unseen in the PC platform. Ease-

of-use features serve to improve PC usability and productivity for users. This greatly minimizes the intimidation and frustration experienced by novice users and enhances productivity for all users for applications that are just being learned or which are used infrequently. Presentation quality features enable users to rapidly generate business-quality documents and presentations.

Enhanced user experience features go beyond ease of use and presentation quality. New technologies are being incorporated to engage the user with dazzling and sophisticated audio, video, and rich 3D graphics effects. Plug and Play technology is reducing the very negative experiences that users have had in PC platform configuration and maintenance. New interconnect technologies (e.g., USB and IEEE 1394) will further this by eliminating legacy resource limitations that had been a source of much user configuration frustration. Reliability, compatibility, and performance are being improved through software initiatives including the introduction of the Win32 Driver Model (WDM) and the DirectX API. Also essential to improved performance and reliability is the eventual elimination of legacy ports.

Summary of Key User Interface Technology

These fundamental sections are intended for those who are relatively new to PC platforms. PC veterans will want to skip ahead.

The user interface paradigms and enabling technologies reviewed in this section promote ease of use and presentation quality.

Graphical User Interface (GUI)

Today's principal operating systems (OSs, such as Windows 95 and Unix) and associated applications for PC platforms use a *Graphical User Interface (GUI).* GUIs are typified by the use of on-screen menus and icons (stylized graphical objects) to represent programs, program controls, and data objects. The objects are activated by a button-press while using a handheld pointer (typically a mouse) to position a cursor over the object. This provides a *point-and-click interface* for *navigation* (locating and selecting programs and data) and primary control (activating, opening, copying, moving).

Windows[43]

PC platform OSs and applications also manage complexity by organizing the display as arbitrarily sized and opaquely overlapping rectangular display regions or *windows,* each generally associated with an individual pro-

[43] A reference in this chapter simply to "Windows" refers to the generic windows described in the section associated with this footnote. A reference to an operating system product will give a specific product name (e.g., Windows NT) or will include the vendors name (e.g., Microsoft) in close proximity.

gram, data file, status message, or control selection. The user is generally permitted wide latitude in the placement and organization of these windows. Windows promote awareness by the user of the computer's multiple ongoing activities and permit quick shifts in the user's focus to any of the activities.

High-Resolution Color BitMaps

PC platform OSs and applications also generally use a *bitmap* to permit the definition of window boundaries and other screen graphics with resolution that is conceptually the same as the pixel resolution of the display screen. High-resolution color monitors, printers, and scanners for high-quality bit-mapped images greatly augment and enhance GUIs, by providing displays and hard-copy output with graphics, drawings, and pictures.

When using analog displays (e.g., a CRT monitor) the frame buffer resolution is typically somewhat higher than the realizable resolution of the display screen. See the section on Graphics and Video Adapters later in this chapter for information on graphics standards.

Historical Comments

Bit-Mapped Graphical User Interfaces (GUIs) with Windows

Bit-mapped Graphical User Interfaces (GUIs) with Windows of today's PCs share many similarities with concepts investigated in Ted Nelson's, multi-decade long, 1965-initiated Project Xanadu and further developed in the Xerox Palo Alto Research Center (PARC) Alto personal computer in 1973. Apple founder Steve Jobs was so impressed with a demonstration of the Alto that he hired away several key designers to develop similar machines for Apple. Apple subsequently introduced the unsuccessful Lisa personal computer in 1982 and the wildly successful Macintosh in 1984.

Portable Document Format (PDF)

PDF is a standard file format for distributing presentation-quality formal documents having both images and text, such as datasheets, manuals, technical reports, and journal papers. Via the PDF standard, creators of documents continue to compose works with their favorite content creation tools (typically desktop publishing and graphics programs). The work is then converted using a special writing tool and distributed in the.pdf format. Users are not required to have the same creative tool used by the publisher in order to view and print the work, yet the document appears identical across all output devices.

INDUSTRY STANDARDS

Portable Document Format (PDF)

The PDF is a proprietary standard of the Adobe Corporation. The Adobe Acrobat Reader, which permits viewing and printing of .pdf files that will look in the final reproduction as the creator intended, is available free to download over the Internet. Adobe benefits by selling the writing tools necessary to create the .pdf files.

The Acrobat Reader is included on the companion CD-ROM. Generating PDF files requires purchasing the full Adobe Acrobat package.

See http://www.adobe.com/prodindex/acrobat/ for information about Adobe Acrobat and the Adobe Acrobat Reader.

Hypertext Documents

Hypertext Documents are Windows-GUI documents containing point-and-click activated text or graphics objects. The reader activates these objects to move within the same document, to go to a new machine-readable document, to activate file transfers and content delivery, activate an embedded program, or activate other arbitrary functions.

HISTORICAL COMMENT

Hypertext

Hypertext has similarities with the "memex" information system described in a 1945 article in *The Atlantic Monthly*, by Vannevar Bush, entitled *As We May Think*. Doug Englebart demonstrated a system using hypertext at the 1968 Fall Joint Computer Conference.

Interactive Assistance and Automation

Sophisticated interactive assistance and automation is common in today's applications for desktop publishing, word processors, presentations, spreadsheets, databases, and design graphics. These features enable individual workers to more efficiently generate, manipulate, and communicate information.

Audios on CD-ROM

Additional tutorial information about several of these Key User Interface Technologies is included in several audio clips on the CD-ROM.

- Graphical User Interface (GUI)
- High-Resolution Color Bitmaps
- Hypertext Documents
- Interactive Assistance and Automation

Multimedia Technology

Multimedia is the ability to augment a static program display with music, sound effects, informational audio messages, video clips, or dynamic graphics and animation, and particularly the simultaneous combination of these. The information presented by the additional media enhances the user experience by providing background information, different perspectives, and other supplemental material, and by grabbing and keeping the user's attention through its dynamic presentation. Multimedia applications are facilitated by the data density of CD-ROMs, or by more expensive high-bandwidth data communications network connections. Multimedia requires additional hardware, such as sound-cards, external speakers and microphones, and interfaces to audio and video sources. Multimedia processing is characterized by the use of intensive, recurring, computations and data-independent recurring memory-accesses on large volumes of small native data types.

CD-ROM and DVD-ROM

CD-ROMs permit volumes of locally held content to be streamed through the platform to the user at data rates that if held remotely would require network connections that are presently prohibitively expensive for many users. These characteristics make this inexpensive removable optical media ideal for software and data distribution via mail and archival storage. The read-only CD-ROM drive is widely available with good performance at low prices and has become a standard feature in all platforms, including the mobile platform. The DVD-ROM is expected to supplant CD-ROM over time, but is being hindered by the existence of multiple competing formats and slow acceptance by some sectors of the entertainment industry.

Streaming Media

Digitization and digital signal processing are discussed in the next subsection.

Audio and video interact with the user as streams of information in analog form. PC platforms employ any of a number of popular digital audio and video standards for capturing and digitizing analog sources for later reconstruction prior to output. The digitized data also exists as streams of information, this time in sequential bit patterns. More generally, streaming data could be any serial sequence of digitized analog signals, such as telemetry from an instrumented process or event. Such a sequence might originate from a CD-ROM drive or other mass storage device, from a video digitizer coupled to a television-like source, or from a high-speed network. Digitized data streams may be compressed, edited, filtered, or otherwise processed using digital signal processing techniques.

Signal processing on streaming data may be done either by the primary processor or dedicated audio and video hardware accelerators (data pumps). Video processing requires input and output peripherals and processing power of significantly higher bandwidth, sophistication, and expense, than for analog. Trade-offs frequently must be made between fidelity and bandwidth.

The processing of streaming data by the processor requires special considerations compared to general-purpose data processing. Unlike other types of data processing, the processing ordering of the data streams generally matters. To prevent acute performance degradation, special care should be taken to prevent the transient streaming data from causing the rapid displacement of other (non-streaming) data from the caches. The most straightforward solution to this problem is to treat streaming data as noncacheable.

DEFINITION

Real-Time Data Processing

Real-time processing of streaming data requires information capture, storage or transmission, processing, or timely generation of responses, for events that are impractical, undesirable, or impossible for the system to temporarily stop or run at a slower rate. A system is effectively a real-time system for a given processing task if it is fast enough to perform its desired function, while processing information at least as fast as it is being naturally generated (or needed) by the evolving event.

Capture, playback, and live conferencing must be performed such that noticeable delays, lack of synchronization, or lost data are minimized. Capture must always be performed in real time. Coding to achieve efficient file sizes involves compression techniques that for video often cannot be performed in real time, especially if dedicated accelerators are

not used. Playback, including digital decoding, generally must be performed in real time. For interactive or multimedia applications, real-time processing should not introduce any processing artifacts noticeable to the user, such as unexpected delays, gaps or abrupt changes in the output, or out-of-sync multimedia.

Digitization Digital Signal Processing

PCM

Pulse Code Modulation (*PCM*) and derivative techniques perform digitization of real-world analog signals; intermediate transmission, storage, and signal processing in the digital domain; and eventual analog reconstruction. The arbitrary processing (filtering or other manipulation) of signals in the digital domain is known as Digital Signal Processing (DSP).

digitization
sampling
coding

Digitization is the sampling and coding of an analog signal into a digital word. *Sampling* uses a sample-and-hold circuit to repetitively capture analog domain (continuous) voltage levels at discrete time intervals. *Coding* maps the continuous voltage level of each sample to one of the discrete voltage levels representable by the fixed bit-width of the digital word output. The mapping is frequently and intentionally non-linear. The Analog-to-Digital conversion (ADC) and front-end signal processing is called *Coding* and is done by a Coder (sometimes Encoder). Sometimes the sampling step is integral (or implied to be) to the coding step. The inverse Digital-to-Analog conversion (DAC) with back-end signal processing is called *Decoding* and is done by a Decoder.

decoding

Codec

A Coder/Decoder, or *Codec*, is a method, a hardware device or subsystem, or signal processing software that combines both a Coder and Decoder. Historically and in most dedicated or high-performance applications, the Codec is implemented in special-purpose hardware. However, depending on the pass-band frequencies of interest, processors can execute Software Codecs to perform the compression and decompression functions in software.

Software Codecs

Software-based Codecs and other DSP routines can consume considerable processor power, particularly for digital video. If DSP functions are used simultaneously with other major applications, or on a regular basis, the platform may need to increase processor speeds, make use of appropriate MMX or AMD-3D instruction sets, or resort to the appropriate acceleration hardware, in order to deliver acceptable performance.

Broadly speaking, a Codec is a standard for implementing digital audio (e.g., AC-3, AC 97) or video (e.g., Indeo, Cinepak). Each standard defines applicable attributes such as the bit-width of the digital words, the sample-rate (and thereby the effective analog pass-band), stereo/mono capability, full/half duplex capability, and the characteristic "law" or "curve" used to map (generally non-linearly) between the analog continuous voltage values and the digital discrete data values.

Digital Audio

Digital audio is music, voice, and sound effects that are at least partially created, transmitted, stored, or received, using PCM techniques. Digital audio also encompasses the use of Musical Instrument Digital Interface (MIDI) and three-dimensional audio. These are each discussed in turn.

Music and Sound Effects Synthesis

FM Synthesis

Sound Blaster is a series of sound cards originally produced by Creative Labs. Its popularity resulted in its becoming a register-level hardware standard for audio cards. Sound Blaster and compatible products relied on *FM Synthesis* to emulate musical instruments and generate sound effects by artful combination and control of many analog oscillators, mixers, and special wave-shaping analog signal processing circuits.

Wavetables

Digital PCM-based Wavetable Synthesis, already popular for portable and compact musical instruments, is supplanting analog FM synthesis methods in the PC platform as well. "*Wavetables*" are sets of digitized samples of musical instruments and other real-world sounds that are stored in ROM or downloaded into RAM. High-fidelity music and audio are then reconstructed from the wavetable samples. Commercial providers can offer wavetables digitized under ideal conditions, sampling notes over the entire range of a perfectly tuned, studio-quality instrument.

Musical Instrument Digital Interface (MIDI)

Musical Instrument Digital Interface (*MIDI*) is a protocol and interface standard for the flexible control and operation of music synthesizers. MIDI files can control both the FM and Wavetable Synthesizer used by PCs to generate music. MIDI-support is a common standard or optional feature on PC sound cards. MIDI has been commonly used for quite some time by portable and "compact" musical instruments and in PC-based games. It is used on Web pages as a low-bandwidth means to add musical content.

MIDI files are the modern day enhanced equivalent of a player-piano song paper-roll. Music is represented by a sequence of digital words. The digital values specify "voice" (choice of a particular mode on the target musical instrument) and musical notes with associated attributes (amplitude, start, duration, and optional special effect). The current General MIDI (GM) specification calls for support for 128 voices and 24 simultaneous notes.

Three-Dimensional Audio

3D Audio, 3D Sound, or Positional Audio, describes multi-channel audio systems that employ signal processing to enhance or emphasize positional or dimensional effects for simulating reflections, depth and spaciousness, and directionality of sounds. Positional audio directional effects include pinpoint sounds that can be moved arbitrarily around the listener.

Dolby Surround AC-3

Dolby Surround AC-3 (for Audio Coding 3), also known as Dolby Surround Digital, or simply AC-3, is the 3D-audio codec used by DVDROMs. It is the Dolby Labs specification for a "5.1" channel PC implementation of Dolby's Surround Sound movie-theatre sound technology. The 5.1 channel designation refers to AC-3's five 3Hz-20kHz channels, for rear-left-side "surround", left, center, right, and rear-right-side "surround;" and one 3-120Hz channel for an arbitrarily placed sub-woofer

Plug and Play Configuration and Maintenance

"We have been successful in spite of ourselves." – Andy Grove, Chairman of Intel, in reference to the poor ease of use of PCs.

PC platforms are noted for their relatively cheap peripherals that are sometimes trivial to install and configure. Once set up, PCs are usually stable and do not require significant maintenance. In contrast, adding new software or hardware is fraught with setup sensitivities, interactions, and inconsistencies. Making a seemingly minor change in software or hardware can wreak havoc, making a PC partially or wholly unusable. One reason the PC platform has been so successful is that PC novices usually don't find out how bad things really are until sometime after they have made a sizable investment in time and money with their system.

Small-business owners and consumers must either maintain their own system, or pay for others to do so. Yet, configuration problems may take hours to resolve for someone who is a senior EE/CS-type and a PC veteran user. Novice PC users without deep pockets and PC-savvy friends may try in vain until they give up in despair and disgust. Small-business owners who have developed a reliance on their PC are usually forced to seek (and pay) for outside help. Such experiences can leave them bitter and hostile toward the PC industry. Tenacious veteran users will resort to exhaustive trials of various combinations (often relieving the tension by cursing a certain software magnate).

The expansion of the PC into both consumer and commercial markets will continue to be limited until this situation is resolved. Potential purchasers are (rightly) scared by the configuration horror stories they hear. They are not completely blind to the fact that such configuration difficulties add to the cost of ownership and subtract from the promised productivity gains.

The problematic nature of PC configuration is often blamed on lack of foresight in early PC architecture decisions leading to acute resource limitations. However, shortcomings in addressing installation and service needs are also to blame. Devices often have a bewildering number of poorly documented configuration parameters. There is typically a near total absence of diagnostic information to aid in isolating what device (or software routine), device parameter, or incompatible interaction between devices, is at the root of the problem. Furthermore, it is generally difficult or time-consuming to get competent support either via telephone, Internet, or computer bulletin-board systems (BBSs).

Requirements Summary for PC Platform Generic Devices

- Every device driver[a] meets requirements specific to its device-class
- Provided companion applications are Win32 compliant
- Every device meets power management requirements specific to its device-class, including class specific support for OnNow and wave-up events[b]
- Provided device drivers meet installation requirements, including unattended installation and help files for special parameters
- Device connectors are labeled with defined icons
- Every device meets Plug and Play requirements specific to its device-class, including unique device ID, auto resource assignment and dynamic-disable[c]

[a] A *device driver* is low-level software that interfaces device independent I/O routines in the OS to a specific type or specific instance of a peripheral device. Short of a custom written application that directly accesses the device, the device cannot be used with the OS without an appropriate device driver.

[b] See the *OnNow and ACPI* section of this chapter.

[c] See the *Plug and Play Configuration and Maintenance* section of this chapter.

Requirements Summary for PC Platform Audio

- Generic device requirements apply
- Only required in Entertainment platform
- Does not use ISA Bus
- Device drivers support Win32 Driver Model[a]
- Full-duplex support for mono or stereo, at 2 data widths, and 6 sample rates
- Provides externally accessible I/O ports
- Reports sample position for stream synchronization with 1ms accuracy
- Meets specific frequency response, dynamic range, distortion, voltage level, and cross-talk requirements for playback, recording, and end-to-end analog
- If DVD Video present, audio specs must be comparable to stand-alone DVD player
- If PCI Bus, must be digital ready (capable of routing final output to OS using bus master[b] transfers for mixing, streaming, sampling rate conversion, and transfer to Universal Serial Bus or IEEE 1394 devices)[c]

[a] See the *Win32 Driver Model (WDM)* section in this chapter.

[b] See the *Backplane Bus — PCI* section in this chapter.

[c] See specific sections in this chapter for discussions of the PCI Bus, Universal Bus, and IEEE 1394.

Plug and Play and Elimination of Legacy Resource Limitations

Plug and Play (Plug-n-Play or PnP) and similar initiatives are intended to ease the frustration and complexity that is typical of installing and configuring hardware for PC platforms.

Plug and Play-compliant devices must have the capability to uniquely identify themselves and their device drivers, state the services they provide and the resources they require, and permit themselves to be programmatically configured by system software. These capabilities, in conjunction with Plug and Play-compliant drivers and OSs, enable the management of peripheral cards for the optimal system assignment of I/O addresses, Interrupt Request (IRQ) select line, and Direct Memory Access[44] (DMA) channels.

It is now possible to buy "Plug and Play" and Microsoft "Designed for Windows" logo devices that supposedly only require insertion (plug), answers to on-screen configuration questions, and the new device is ready for use (play). Such simplicity is designed to reduce returns, minimize customer support costs, and reduce end-user cost of ownership. Plug and Play generally makes installation of new devices easier. However, the installation of devices requiring legacy IRQ and DMA assignments in fully loaded systems is still often difficult. Until all legacy dependencies are removed, the goal of ease of configuration and maintenance remains elusive.

Industry Standards

Plug-N-Play Specifications

Plug and Play is actually a family of specifications for the system BIOS and various device classes, including buses and ports. A compilation of links to many of the PnP specifications is found at URL:

http://www.microsoft.com/hwdev/respec/pnpspecs.htm

Suggested Readings

Plug and Play

The book by Tom Shanley, *Plug and Play System Architecture*, Addison-Wesley, 1995, gives a clear and detailed explanation of the concepts behind Plug and Play.

[44] DMA is described in the section on Legacy Direct Memory Access (DMA).

Mechanical Design

As we discussed in the Mechanical and Electrical Considerations section of Chapter 1, the ATX motherboard standard has introduced a new size and orientation and component placement. Figure 4.2 provides an example ATX motherboard layout. Figure 4.3 shows an example connector placement at the rear of the ATX motherboard. Relative to the Baby AT board form factor shown in Figure 1.3 on page 37, the ATX motherboard is intended to reduce system assembly time and costs for manufacturers, increase ease of use for end users, and lower technical support needs. It does this by attention to detail with respect to the placement of the microprocessor and connectors.

ATX places the microprocessor near the power-supply. The ATX power-supply fan blows air onto the microprocessor, instead of out of the system unit as in an AT-style chassis. This eliminates the need for a fan-heatsink on the processor. In addition, the microprocessor is out of the path of any full-length expansion boards that might obstruct access to the microprocessor. This ensures that processor upgrades do not require removal of any expansion boards.

The floppy and ATA/IDE connectors are at the front of the board and closer to the peripherals, permitting reduced cable lengths. Integrated I/O connectors on the rear of the board obviate the need for cabling between the motherboard and connectors mounted on the system unit. The power-supply connector is also a single piece, instead of the two-piece connector used in an AT-style chassis.

Reports on CD-ROM

The following AMD technical documents on the CD-ROM are related to ATX motherboards and electrical and thermal system design issues:

- *ATX Reference Design for AMD-640 Chipset*, Users Guide Publication 21265, June 1997.
- *ATX Reference Design for AMD-640 Chipset,* Technical Specification Publication 21264, June 1997.
- *AMD-K6 MMX Enhanced Processor Power Supply Design*, Document 21103, June 1997.
- *AMD-K6 MMX Enhanced Processor Thermal Solution Design*, Document 21085, June 1997.

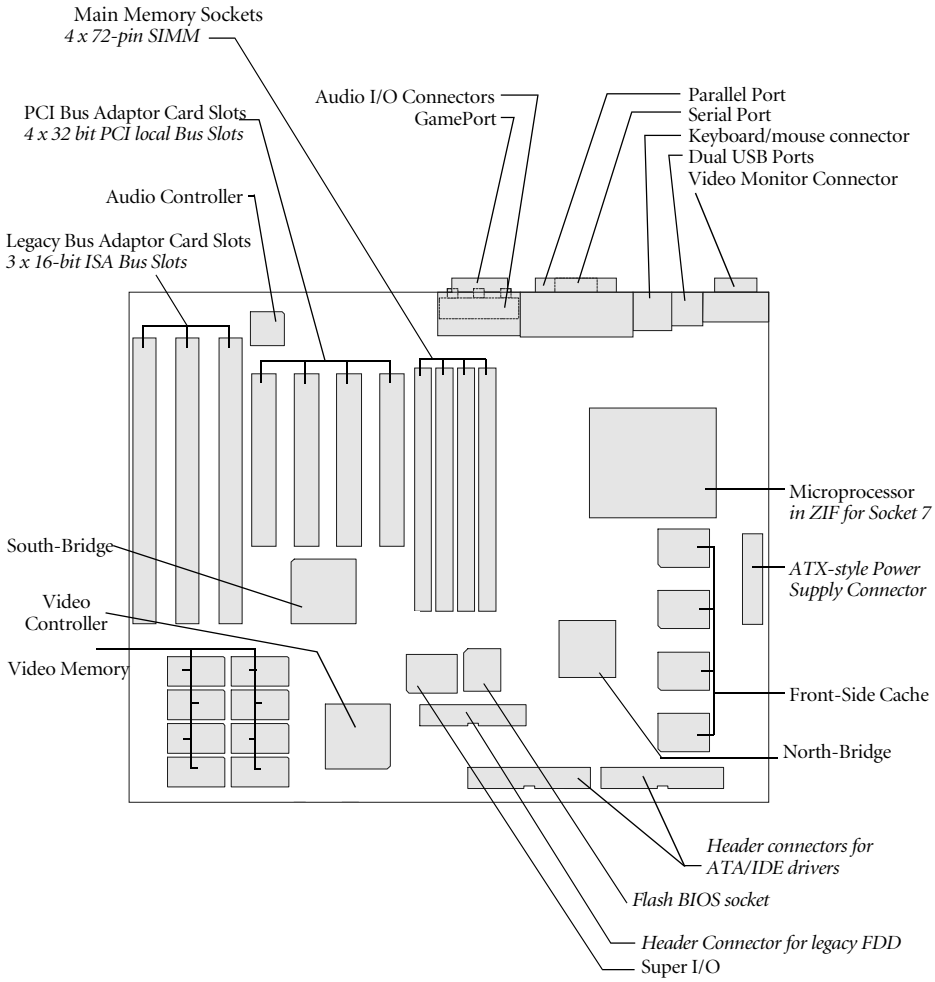

Figure 4.2 EXAMPLE ATX MOTHERBOARD LAYOUT

Adapted with permission of Advanced Micro Devices Inc., from *AMD ATX Reference Design Technical Specification,* Copyright 1997.

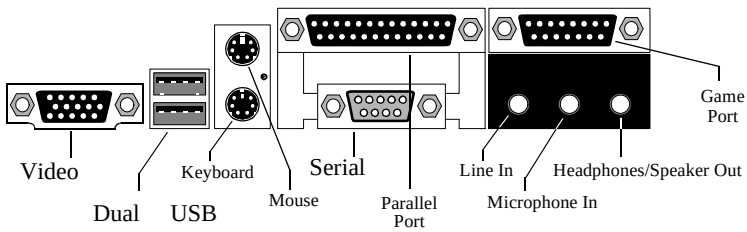

Figure 4.3 EXAMPLE ATX CONNECTOR PLACEMENT VIEWED FROM REAR OF BOARD

Adapted with permission of Advanced Micro Devices Inc., from *AMD ATX Reference Design Technical Specification,* Copyright 1997.

STANDARDS ON CD-ROM

The following Intel authored motherboard and system specifications are included on the CD-ROM:

- *NLX Motherboard Specification.*
- *ATX Specification Version 2.01.*
- *microATX Motherboard Interface Specification Version 1.0.*
- *SFX Power Supply Design Guide Version 1.0 Release*

Win32 Driver Model (WDM)

Driver problems are considered the most significant contributor to poor consumer experiences with the PC. Problems arise from bugs, compatibility problems, and unintended side effects from driver upgrades. The Win32 Driver Model (WDM) is a planned Microsoft architecture for device drivers, which hierarchically splits device drivers into OS-provided device-class drivers and IHV-provided minidrivers. The WDM interface resides above bus and device drivers, and below PnP, power management, and I/O executives.

WDM is intended to improve device driver quality and consistency, while reducing the burden of device driver development by IHVs. Late '90s Windows 95 and Windows NT releases will be based on the NT 5.0 Kernel and will ship with common device-classes drivers. Earlier drivers will not be WDM compliant, but Windows 98 will be backward compatible. WDM-compliant device drivers updates are envisioned as being distributed via the Web.

WDM provides a common set of I/O services and binary-compatible drivers for Windows 98 and Windows NT 5.0. WDM supports Windows NT Plug and Play and power management, USB, IEEE 1394, OnNow power management, Streams (video, video camera, CD-ROM or DVD sourced), Audio, still image (camera or scanner) capture, and Human Interface Device (HID – keyboard, mouse, game pad).

DirectX and ActiveX

DirectX is the name of low-latency and high-performance Win32 APIs, which give the performance effect of writing directly to hardware, while maintaining register-level independence. DirectX supports "transparent" hardware acceleration, wherein APIs automatically make use of any available hardware acceleration, without any special effort by the application programmer. Applications that are written using these APIs, instead of

being written for a specific hardware register set, will run on any hardware that provides the user with drivers appropriate to the user's OS. DirectX has been described as game extensions to Microsoft Windows.

DirectX Classes

Class Name	Associated Functions
Direct3D	3D graphics accleration
DirectDraw	2D graphics acceleration, foundation for more complex video services, including Direct3D
DirectInput	input devices, including keyboards, mouse-like pointing devices, and joysticks
DirectPlay	communications for multi-PC and multi-player games
DirectSound	audio, including mixing
DirectVideo	full-motion video overlays
DirectMPEG	MPEG full-motion video codec

Industry Standards

DirectX

Information about DirectX, including Software Development Kits (SDKs) and Device Driver Kits (DDKs), can be found at:

http://www.microsoft.com/directx/

Information on "Meltdown," an annual software/hardware compatibility testing event, can be found at:

http://www.microsoft.com/hwdev/meltdown.htm

ActiveX is the name of Microsoft time-based APIs used for creation, coordination, and management of synchronized multimedia streaming data. ActiveX is intended to find use particularly in Web-based applications over networks (including intranets and the Internet), where latencies and bandwidth are unpredictable and time-varying. ActiveX includes support for the reusability, coexistence, and interoperability of remote and local files and objects. ActiveX also provides support for unified browsing. Built on top of the DirectX APIs, ActiveX has been described as Internet extensions to Microsoft Windows.

Appliance-Like Operation

Appliance-like operation describes a PC platform that is unobtrusive, always ready yet energy efficient, and extremely simple to install, operate, maintain, and upgrade. Simply Interactive Personal Computer (SIPC) is the umbrella initiative that defines the vision of "simple, convenient, and approachable" appliance-like operation. Device Bay, OnNow, and ACPI are key standards and initiatives that will enable SIPC to eventually be a reality.

Simply Interactive Personal Computer (SIPC)

Simply Interactive Personal Computer (SIPC) is a broad Microsoft system-oriented vision of future PCs that encompasses the OS, applications, platforms, and peripherals. The SIPC proposal integrates a number of Microsoft and industry initiatives and standards.

SIPC machines are to have the ease of use of established consumer devices, such as a TV or VCR. This will be accomplished in part by virtually grafting a consumer electronics user interface onto SIPC platforms. Microsoft's vision is that a user can do the most common PC tasks envisioned for the future without any prior experience. Such tasks include "playing a game; watching a movie; writing an email, letter, or invitation; browsing the Internet, hooking up a device (such as a digital camera or camcorder); and listening to voice messages."[45]

Like other consumer devices, SIPC platforms should be sealed-case devices as far as the user is concerned. This means that peripherals and communications devices can be hooked and unhooked, and basic upgrades can be performed, without opening the case. This is to be accomplished by reliance on USB, IEEE 1394, and Device Bay. Instant-on capability is also required, to enable users to interact with the machine with more spontaneity, yet conserve power when not active. SIPC power features are addressed by OnNow and ACPI technologies. SIPC is also consumer-entertainment-focused, calling for large screens, DirectX-and ActiveX-based multimedia, DVD-ROM-based movies, and AC-3-based three-dimensional audio.

Industry Standards

Microsoft's Simply Interactive Personal Computer (SIPC) Initiative

Information on the SIPC initiative can be found at URL:

http://www.microsoft.com/hwdev/desinit/sipc.htm

[45] *Making The PC An Appliance*, presented by Bill Gates at WinHEC '96.

Report on CD-ROM

To learn more about the SIPC initiative, see Peter N. Glaskowsky's article, "PCs Head Toward Appliance Status," *Microprocessor Report*, Vol. 10, No. 6, May 6, 1996, on the companion CD-ROM.

Device Bay

Device Bay is an initiative intended to make installation of new peripherals "as easy as inserting a game cartridge." Hard-disk drives, modems, network adapters, CD drives, DVD drives, and other electronic devices can be Plug and Play hot-swapped into a "sealed-case" PC with the same convenience that floppy disks and other removable media presently offer. Figure 4.4 illustrates the Device Bay concept.

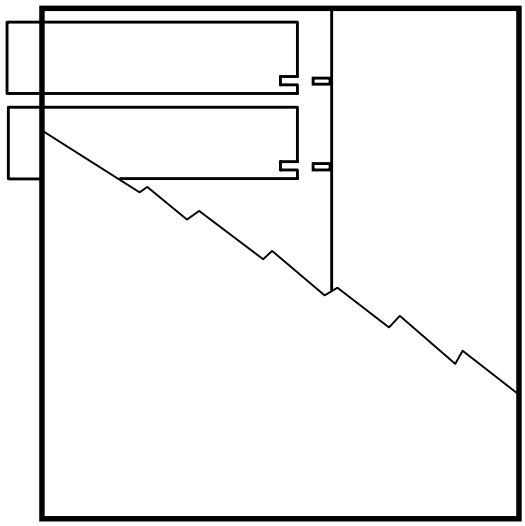

Abstract Cut-Away of System Unit with Ejected Device Bay Devices

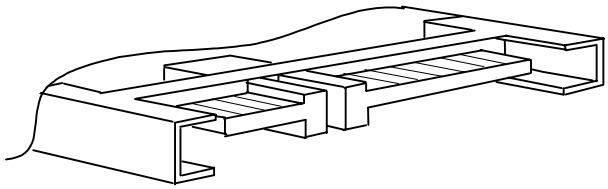

Recessed Connector in Rear of Removable Device

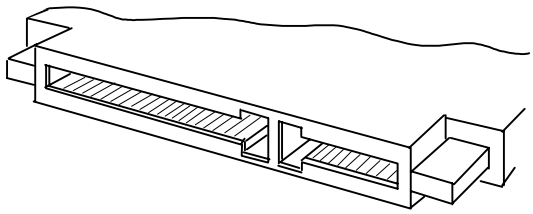

Connector in System Device Bay

Figure 4.4 Device Bay Conceptual Drawings

On the left is shown a cut-away view of a system unit with two ejected Device Bay devices. All Device Bay compatible devices will have a standardized recessed connector and connector placement. The mating connector is correspondingly mounted within the system-unit. These self-aligning connectors are drawn to the right of Figure 4.4. Platform customization of peripherals should also prove faster and easier. Its promoters claim that Device Bay will provide platform scalability over the next 5-10 years via triv-

ial peripheral additions and upgrades. Support for Device Bay is scheduled for Windows NT 5.0 and in a secondary release of Windows 98.

A Device Bay mobile hard disk will provide the advantages offered today by removable media in easily moving working files between machines. Typical examples include disk transfers for commuting between home and office or for business trips between laptop and workstation.

Electrically, Device Bay relies on a hybrid combination of the USB and IEEE 1394 buses. Each bay includes multi-voltage DC power, a USB port, and an IEEE 1394 port. Presently the 1394 port operates at 400 Mbps,[46] but later will be upgraded to 1 Gbps when devices based on the 1394B standard are available. The Device Bay specification defines the physical and electrical interfaces, desktop and mobile form factors (size), and OS response to device insertion and removal.

INDUSTRY STANDARDS

The Group Known as Device Bay

The Device Bay Specification is being developed and controlled by Compaq, Intel, and Microsoft. The Device Bay group maintains a Web site, with downloadable copy of the draft specification and a variety of useful technical documents. The main URL is:

http://www.device-bay.org/

The draft Device Bay specification is downloadable free from URL:

http://www.device-bay.org/tech/spec.htm

OnNow and ACPI

An OnNow PC is a platform that appears to be turned off, yet immediately responds when called upon by the user (or other devices). It is also an industry initiative and standard for enabling OS awareness of power use and requirements for all system components, as well as power control of the components. OnNow initiative goals include reducing typical power consumption to one-third of present levels and achieving a 5-second or less "warm-up" delay. The OS will have the ability to dynamically enable subsystems, whether on the motherboard or in attached peripherals. This requires "OnNow awareness" by platform hardware, expansion peripherals, applications, the OS, and the BIOS.

The Advanced Configuration and Power Interface (ACPI) is a specification and language that enables OS power control over installed devices. The OS is made aware of the power functions and requirements of installed devices through ACPI-compliant device drivers. Intel, Microsoft,

[46] We use "bps" for bits per second and "Bps" for bytes per second.

and Toshiba created ACPI in support of the OnNow initiative. ACPI is compliant with both VxD and WDM device driver models.

Industry Standards

OnNow Specifications

The OnNow area of Microsoft's Web site includes device-class power management specifications, ACPI Source Language Assembler (Asl.exe), the Microsoft Windows Hardware Compatibility Tests for ACPI, and information about fast-boot BIOS and Simple Boot Flag specification. Power management reference specifications are available for each device class and bus. The main URL is:

http://www.microsoft.com/hwdev/onnow.htm

Requirements for OnNow system designs were originally published in the *PC 97 Hardware Design Guide*. Information about obtaining this guide was included earlier in the chapter. Clarifications and changes subsequent to the release of the PC 97 guide are found at URL:

http://www.microsoft.com/hwdev/desguid/onnowpc97.htm

Total Cost-of-Ownership

Total Cost-of Ownership (TCO) of PC installations is an important facet of the commercial market segment. Reducing TCO is viewed as an enabler of increased commercial sales. TCO includes capital costs; administration; technical support; and end-user training, applications development, data management, supplies, and other miscellaneous costs. This section surveys two management initiatives held up as solutions to reduce TCO. Specifically, these initiatives are designed to help increase platform control, improve interoperability, increase user productivity, reduce the need for user support, and reduce capital and administrative costs.

The Desktop Management Initiative defines a very general operating-system-and device-independent means to interface management (control) applications with controllable devices. Management applications include network, telecom, Microsoft Windows kernel processes, DMI processes, and other unspecified applications that may need to control a device. The Microsoft Zero Administration Initiative for Windows (ZAW) is Microsoft's comprehensive management vision that integrates a number of management technologies and initiatives, including DMI.

INDUSTRY STANDARDS

Microsoft's Total Cost-of-Ownership (TCO) Initiative

Information on the TCO initiative can be found at URL:

http://www.microsoft.com/windows/platform/info/tco.htm

Desktop Management Initiative (DMI)

The Desktop Management Initiative (DMI) specification requires that an ASCII text file database be set up for all controllable devices. The files contain the "manageable attributes" of the devices in a language called *Management Information Format (MIF)*, which has a grammar and syntax. The files are called MIF files. A continuously running background task called the Service Layer manages the MIF database, compiling MIF file data, and servicing database queries from management applications. The DMI enables users and technical support personnel to retrieve a wealth of information about a PC, including processor type, installation dates, installed peripherals, power sources, and maintenance history. DMI 2.0 incorporated Remote Procedure Calls (RPC) to permit technical support to retrieve the management data remotely over a network. The Co*mmon Interface Model (CIM)* is an object-oriented version of DMI that permits cross-platform interoperable management over intranets and the Internet.

INDUSTRY STANDARDS

Desktop Management Task Force (DMTF)

The Desktop Management Task Force sees itself as "*Driving industry standards for systems management to reduce total cost of ownership. (DMTF) is the industry consortium chartered with development, support and maintenance of management standards for PC systems and products, including DMI, the most-widely used management standard today.*" DMTF membership is $2,500 or $10,000 per year, depending on level of participation. The DMTF maintains the following Web site:

http://www.dmtf.org/

DMI Specification

The DMI specification and other DMTF standards are downloadable for free at the following URL:

http://www.dmtf.org/tech/specs.html

The Zero Administration Initiative for Windows (ZAW)

Microsoft's Zero Administration Initiative for Windows (ZAW) is an umbrella initiative for other Microsoft management initiatives, including Web-Based Enterprise Management (WBEM), Microsoft Windows Management Interface (WMI), Net PC, and Hydra (a Microsoft Windows terminal). In general these initiatives attempt to reduce TCO via reduced needs for user support, increased centralized but flexible control, automation of administrative tasks, while maintaining or increasing user productivity.

The WBEM Initiative supports the sharing of management data across network, desktop DMI, telecom, and Microsoft Windows applications. WBEM incorporates DMI and has provisions for power management, full system monitoring, and remote configuration. WMI defines a low-level instrumentation layer for efficient development of management-instrumented drivers.

NetPC platforms implement OnNow power management and perform auto detection and configuration of all buses and devices via ACPI protocols. Net PC systems perform automatic network configuration and then implement centralized configuration of the Net PC desktop via pre-configured system policies, user profiles, and setup scripts. "Roaming" and easy machine replacement are made possible by maintaining all application and data storage (centralized state) on the network server.

Users are not allowed access to system files and features, cannot install unapproved applications, and are otherwise prevented from performing end-user operations that may require technical support to correct. "Task Oriented" workers (e.g., clerks or bank tellers) are configured into *TaskStation Mode*, in which the Net PC machine boots into a single management-specified dedicated application, such as a Web browser or a business application. "Knowledge" workers are configured into *AppStation Mode*, in which the user may run any one of multiple management-selected business applications.

TaskStation Mode

AppStation Mode

INDUSTRY STANDARDS

Microsoft's Zero Administration Initiative for Windows

Information about ZAW and other Microsoft Windows Management Initiatives can be found at the following URL:

http://www.microsoft.com/management/

Information about ZAW and other Microsoft Windows Total Cost of Ownership Initiatives can be found at the following URL:

http://www.microsoft.com/windows/innovation/

Hydra describes a client that is thinner than the Net PC. All applications run on the server. (HID) inputs are uploaded to the server, which downloads each display screen to a terminal. The user sees a Window 95 User Interface

Connectivity

Connectivity describes the interconnection of PC platforms (particularly with servers) via data communications and network technologies. Interconnection may be point-to-point with other platforms, but is increasing via a network, including workgroups, intranets, and the Internet. The first subsection on *Data Communications and Networking* surveys fax, "data," Integrated Services Digital Network, and Asymmetric Digital Subscriber Line modem technologies; Ethernet and derivative networks; and compares bandwidths for a variety of channels and links. The second subsection, on the *Internet*, provides a short technical overview of Internet access and Internet-based applications and uses.

Data Communications and Networking

Reducing file sizes via compression prior to transmission facilitates transfer (over any connection) of large files at low data rates. Received files are restored to original form via decompression prior to use. These complementary processes can be done in express steps on the entire file, respectively performed before sending and after receipt, using zip and unzip utilities. Communications protocols and publishing standards are increasingly incorporating compression and decompression as integral components of the end-to-end transfer. The PDF format discussed earlier incorporates text and image compression.

Medium-speed (14.4kbps-54kbps) fax and dial-up data communications hardware, software, and related services are now integral platform components for home and business. On the hardware side, "high-speed" serial ports and progressively faster high-speed data and fax modems support the trend toward faster connectivity. Increased connectivity has resulted in additional home and business phone lines to avoid the unavailability of normal voice use that would otherwise occur from high dial-up usage.

Data Communications and FAX and Data Modems

Data modems[47] provide access to a variety of data communications services via the public telephone-switching network ("on-line" or "dial-up" services) to Internet Service Providers (ISPs), consumer online services, and subscription news and database providers. Generally these services insure that their subscribers have all the necessary software for HTML browsing or for accessing the services proprietary content.

Modems supporting common fax standards are available for minimal additional cost over that of non-fax modems. Such fax-modems typically are bundled with supporting software, enabling the associated platform to replace a conventional fax machine for receiving faxes, while providing

[47] Modem is short for modulator-demodulator, a device that encodes a digital data stream into analog tones for transmission over an analog communications link and decodes received tones back to digital form. The term is also loosely and widely used for any adapter between a computer and a communications link regardless of whether a modulator-demodulator is present or not.

additional features only available in very expensive standalone fax machines, or only available on a PC. Fax-equipped PCs are commonly used to permit viewing with optional printing, archiving of fax images, delayed fax sending, maintenance of address books, and broadcasting to distribution lists. For sending faxes, either an optional scanner must be used, or messages must be composed on the computer.

Industry Standards

International Telecommunications Union (ITU)

Modem modulation standards are promulgated by and only available from the International Telecommunications Union (ITU). The modem standards are known as *ITU-T Series V Recommendations*. The "-T" stands for Telecommunications. Series V covers all ITU-T standards for data communication over the telephone network.

The ITU-T Series V Recommendations can be downloaded in zipped Word or PostScript formats via a secure credit-card payment server. Prices, in French francs, vary with each standard. The V.34 standard was roughly $6 as of this writing. The standards, listed by number and title, may be ordered from the following URL:

http://www.itu.int/itudoc/itu-t/rec/v.html

Recent important standards include:

ITU-T Rec. V.17 for 14,400 bps class 1 & 2 FAX modems.

ITU-T Rec. V.32 bis for 14,400 bps data modems.

ITU-T Rec. V.34 for 33,600 bps data modems.

ITU-T Rec. V.90 for 56,000 bps data modems.

Requirements Summary for PC Platform Modem Devices

- Generic device requirements apply
- Required, if no network adapter[a]
- Devices use the Universal Modem Driver (Unimodem), which uses vendor supplied INF files to characterize the device to the OS
- Hayes an early, popular modem compatible command set
- 33.6 Kbps minimum using V.34, V.42, and V.42bis
- Pulse Code Modulation modems (e.g., X2, K56flex) are V.90 compliant
- Fax modems support 14.4 Kbps (V.17) with Class 1 command set
- Controller supports the diagnostic command; has a 60 character minimum command buffer; reports connection status including rate, error control, and data compression; and is software upgradable
- Devices on power-managed buses support wake-up on incoming ring
- ISDN modems (discussed next) have additional special requirements

[a] See the *Ethernet and Derivatives* section in this chapter.

Other Data Communications Technologies

Integrated Services Digital Network (ISDN) is the name of a family of switched (dial-up) digital telephone services that use existing telephone wiring. Some of these services include burglar alarm monitoring and virtual PBC telephone services. For PC platform applications, the most common type of ISDN is the Basic Rate Interface (BRI), which consists of two independent 64-Kpbs (56-Kbps in some places) "B" channels for data and a "D" channel for control.

ISDN has established a reputation for being limited in general availability to sites close (18,000 feet) to telephone company Central Offices in major metropolitan areas, difficult and expensive to install and configure, and costly to use. It is however an established technology, including built-in support in the most recent releases of Microsoft Windows. Furthermore, short of leasing a dedicated frame relay or T1 phone line, as of this writing it is the fastest way (other than satellite modem, discussed next) to connect to the Internet for major metropolitan users. It is estimated that there were 1.7 million ISDN subscribers in 1997.[48] ISDN may soon become obsolete in face of the newer technologies discussed below.

Microsoft Windows Support for ISDN

Microsoft maintains an area on its Web site dedicated to ISDN. General information on ISDN, geographic availability, and a facility for requesting ISDN service, are found at URL:

http://www.microsoft.com/windows/getisdn/

Consumer satellite modem systems have been made recently available. These use a data modem (e.g., V.34 or ISDN), for low-or medium-speed (28.8-128 Kbps) client-to-provider data over conventional telephone lines, coupled with a special satellite downlink data modem for medium-speed (400 Kbs) provider-to-client data. Such systems are targeted at Web browsing, where such asymmetric bandwidth is typical. These systems find particular application in areas where other high-speed services are not available.

Cable Modems

Cable modems use specially conditioned segments of the existing cable TV infrastructure to provide bidirectional high-bandwidth communication channels for Internet connectivity. Cable modems are currently undergoing trials in several metropolitan areas in the

[48] Mo Krochmal, "ISDN Dukes It Out in Bandwidth Battle," *TechWeb*, January 25, 1998.

U.S., with roughly 100,000 households participating.[49] At present, many competing and incompatible cable-modem interface standards are being used. As a result, local cable companies with cable modem service require the subscriber to lease a cable modem. This ensures that the cable modem is interoperable with the modulation methods and interface protocols chosen for the cable company's headend servers. Because volumes of these diverse modem types is low, their costs have stayed high ($300-$1500), and installation and service rates are not yet to levels affordable by the masses. Standardization efforts are attempting to address this problem.

Industry Standards

The Data Over Cable Service Interface Specifications (DOCSIS)

The Data Over Cable Service Interface Specifications (DOCSIS), also known as the MCNS specifications, is a family of interoperability certification standards for cable modems that are based on TCP/IP[a] protocols. MCNS stands for Multimedia Cable Network System (MCNS) Partners, L.P., a group consisting of four major cable providers that originally started the effort. The project, which is now managed by CableLabs discussed below, is implementing a certification program for cable modems so that cable operators can be assured that a given modem will be interoperable with DOCSIS-compliant headend equipment. The project also plans to work with the Society of Cable Telecommunications Engineers (SCTE) to eventually submit a DOCSIS-based standard for approval by ANSI and the ITU. The DOCSIS Project has a web site. An overview, answers to frequently asked questions, and member links, are available at URL:

http://www.cablemodem.com/

Public releases of DOCSIS technical specifications may be freely downloaded at URL:

http://www.cablemodem.com/public/pubtech.html

Cable Television Laboratories, Inc. (CableLabs) is a research and development consortium of cable television system operators. Vendors of cable equipment, other telecommunications providers, and the general public are not eligible. Member companies pay dues based on their subscriber base. Cable Labs manages the DOCSIS Project, discussed below. CableLabs operates a web site. Tutorial publications on cable modems and the cable industry in general are available at URL:

http://www.cablelabs.com/Publications.html.

[a] TCP/IP, or Transmission Control Protocol/Internet Protocol.

[49] Seth Schiesel, "Three Giants of PC World Turn Focus to Speed," *New York Times*, January 26, 1998.

While present cable modem designs are coupled to PCs via 10BaseT interfaces, having 10Mbps nominal rates, entry-level service plans start with advertised asymmetrical rates of 500Kbps downstream and 100Kbps upstream for $100 per month. Symmetrical rates or higher bandwidths (e.g., 1Mbps advertised) cost more, $300 and $500 per month, respectively. It is hoped that competition with Digital Subscriber Line (DSL) connectivity (discussed next) will rapidly lower these monthly charges.

Like any other network, effective data rates are actually a function of loading on the local cable operator's network and the cable operator's connection to the Internet. Under light loads some users have reported effective rates within the same local cable network limited only by the capabilities at which contemporary PCs can process TCP/IP protocols, said to be approximately 4Mbps.[50] Conversely, skeptics have expressed concerns as to whether there is enough bandwidth headroom on the local nets if a significant number of cable subscribers sign up. Also, like getting service from any other form of Internet Service Provider, users are advised to inquire as to what bandwidth the cable service provider has to the Internet backbone, and how many users are sharing that bandwidth.

Emerging DSL and ADSL Modems

Digital Subscriber Line (DSL) modems promise to offer fast (1.5 Mbps-8 Mbps) connections to the home in the near future. It is estimated that there may be 1 million DSL subscribers by the year 2000.[51] A number of DSL variants exist and are undergoing development and field trials. Asymmetric DSL (*ADSL*) is one variant that may be the first to see widespread use. ADSL downstream and upstream rates are different, due to cost-benefit compromises in engineering a low-cost, high-data-rate solution for the masses. Maximum rates are said to be 9 Mbps downstream and 800 Kbps upstream, depending upon line length and conditions. [52]

ADSL

In one ADSL implementation being proposed by the Universal ADSL Working Group (UAWG), a consortium led by Microsoft, Intel, Compaq, and several Baby Bells, downstream rates are 1.5 Mbps and upstream rates are 384 Kbps. Since by far most users will consume much more data than they generate, and since the upstream rate will still be roughly an order of magnitude improvement over analog modems, its promoters do not perceive the disparity in upstream and downstream bandwidths as being a problem.

---

[50] www.cnet.com/Content/Features/Techno/Cablemodems/

[51] Mo Krochmal, "ISDN Dukes It Out in Bandwidth Battle," *TechWeb* (http://www.techweb.com/), January 25, 1998.

[52] ADSL Forum.

Universal ADSL

Spliterless ADSL
Lite ADSL

In addition to improved performance, the UAWG claims several advantages of the technology will speed its global acceptance. Universal ADSL is also known as a *Spliterless ADSL* or *Lite ADSL* version of DSL. ADSL does away with a splitter device that other versions of ADSL require to be installed at the point of entry into each home, reducing time and costs for installation. Microsoft promises to incorporate simplified setup for Universal ADSL in future versions of Microsoft Windows. Finally, Universal ADSL is being proposed as an "always-on" service. "Always-on" connectivity eliminates connection delays and will enable continuous Web services and e-mail delivery.

INDUSTRY STANDARDS

Asymmetric Digital Subscriber Line (ADSL) Forum

The *ADSL Forum* is a 200-member group formed in 1994 to educate and promote ADSL technology to telephone companies and their suppliers. Membership costs $1500-$5000 depending on company size and level of participation in Forum activities. The ADSL Forum maintains a Web site. Information on joining and tutorial papers are available at URL:

http://www.adsl.com/

The Universal ADSL Working Group (*UAWG*), composed of PC industry, networking, and telecommunications companies, is developing a compatible extension to the present ANSI standard T1.413 for ADSL. Their proposal is to be submitted to the International Telecommunication Union (ITU) G.Lite subcommittee. The UAWG has set up a Web site at URL:

http://www.uawg.org/index.html

ADSL Specifications

ANSI T1.413-1995, *Telecommunications - Network and Customer Installation Interfaces - Asymmetric Digital Subscriber Line (ADSL) Metallic Interface*, $135, must be ordered directly from ANSI. ANSI standards can be ordered through their online catalog, searchable at URL:

http://www.ansi.org/catalog/search.html

Networking and Local-Area Networks

This section overviews Ethernet and its derivative technologies (Fast Ethernet and Gigabit Ethernet) and considers network applications enabled by Ethernet. Such applications include workgroup computing, the client-server paradigm, departmental and Enterprise servers, and intranets.

Ethernet and Derivatives

Media Access Control (MAC)

Ethernet is the most popular of several *Media Access Control (MAC)* methods used for LANs. *Media* refers to the physical infrastructure that comprises the communications network. *Access Control* refers to how the media is shared among multiple nodes that request to use the media, often at the same time. Developed at Xerox PARC in 1976.

Carrier Sense Multiple Access with Collision Detection

Ethernet uses a MAC method known as *Carrier Sense Multiple Access with Collision Detection (CSMA/CD)*. Essentially, each node with data to transmit checks the status of the media and begins transmitting if the media is not in use, or waits if the media is already busy. Each node also checks to make sure that a *collision* (simultaneous transmission attempt) did not occur with another node. If a collision is detected, then both nodes wait for random intervals before starting over with checking the media status before transmitting. In the absence of a collision, the transmitting node gets the exclusive use of the media to transmit its data. Ethernet later became the IEEE 802.3 (dot3) standard. Ethernet has a number of variations. The topology for three of these variations is shown in the left half of Figure 4.6 on page 359.

10Base 5
Thick Ethernet

The initial physical implementation of the 802.3 standard was known as 10Base-5, which is a reference to its 10Mbps rate, reliance on baseband (non-frequency shifted) data signaling, and 500-meter range. *10Base-5* is also known as *Thick Ethernet*, due to its use of thick 50-ohm coaxial cable. Thick coax is expensive and its coiling radius makes it difficult to work with inside small clearance areas. Thick Ethernet is prevalent at many sites, but newer Ethernet variants have supplanted it for new installations. In 1984, a variant known as *10Base-2*, was standardized. 10Base-2 (having a 185-meter range) is also known as *Thin Ethernet,* due to its use of thin 50-ohm coaxial cable and BNC-type connectors. Figure 4.5 compares the diameter of the two coax types.

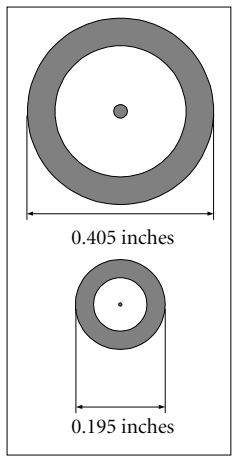

Figure 4.5
COMPARISON OF THICK AND THIN COAXIAL CABLES USED FOR ETHERNET

10Base 2

Thin Ethernet

10Base T
Twisted-Pair Ethernet

In 1990, another variant known as 10Base-T was standardized as IEEE 802.3i. 10Base-T is also known as *Twisted-Pair Ethernet* due to its use of Unshielded Twisted Pair (UTP) and RJ-45 connectors. UTP is a type of telephone-wiring-like cabling that is available in a range of performance categories. More recently, a 100Mbps extension, known as *100Base-T*, or *Fast Ethernet*, was standardized as IEEE 802.3u. Variants include 100Base T2, T4, and TX. IEEE 802.3z is a proposed 1000Mbps extension, known as *Gigabit Ethernet*.

100Base T
Fast Ethernet

Gigabit Ethernet

A Medium Attachment Unit (MAU, also known as a transceiver) couples a LAN controller in the Data Terminal Equipment (DTE, e.g., computers or data terminals) to the network. The MAU and LAN controller combination comprise a *network adapter*. For 10Base-5 and 10Base-2, the MAU couples the DTE to a *trunk* coaxial cable, which function as a distributed data bus. For 10Base-5, each trunk can be no more than 500 meters in length and have no more than 100 nodes attached. Five trunks can be coupled using four repeaters for a maximum length of 2,460 meters, although source and destination nodes must be no more than two repeaters apart.

Medium Attachment Unit (MAU)
transceiver

network adapter

A *repeater* extends a logical LAN segment[53] by regenerating transmissions observed on a first subnetwork onto a second subnetwork. For 10Base-2, each trunk can be no more than 186 meters in length and have no more than thirty nodes attached. Five trunks can be coupled using four repeaters for a maximum length of 910 meters, although again no more than two repeaters may exist between source and destination nodes. In both coaxial Ethernet forms, each far-end must have a 50-ohm *terminator* (termination resistor). The terminator reduces unwanted signal reflections on the trunk and ensures proper operation.

repeater

For 10Base-T, the MAU couples the DTE via transmit and receive twisted-pairs to a *hub* (also known as a *concentrator*), which functions as a logical data bus. Workgroup-class hubs commonly come with up to twelve ports. Up to twelve Workgroup-class hubs may be connected to an Intermediate-class central hub.

hub
concentrator

In all three Ethernet variants, LAN segments are usually limited to significantly less than the maximum number of nodes to avoid network congestion. Multiple LAN segments are then bridged, routed, or switched to form larger departmental or enterprise networks. A *bridge* selectively links two network segments and will forward Ethernet transmissions to a second segment only if the destination address specified by the transmission and the second segment match. A *router* selectively couples many network segments and can perform sophisticated filtering and best-path selection in addition to simple address matching.

bridge

router

The MAU for each Ethernet type is shown in the right half of Figure 4.6. The MAU consists of a network interface (either to coax or twisted-pairs), an *Attachment Unit Interface (AUI)*, a transceiver for coupling signals between the two interfaces, and some means to provide DC isolation between the two interfaces. The AUI includes separate differential signal-pairs for transmit data (TD) receive data (RD) and transmit collision status (COL). The MAU passes received data from the network to the LAN controller and broadcasts transmit data from the LAN controller onto the network. It also detects collisions between its transmit data and activity on the network and notifies the LAN controller of any detected collisions.

Attachment Unit Interface (AUI)

[53] A network segment is a portion of a network having a unique address. Different network segments necessarily have different addresses.

transceiver cable

When collisions occur, the LAN controllers involved in the collision reschedule their transmissions after a random interval. For 10Base-2 and 10Base-T expansion cards designed for the PC platform, the MAU is usually integrated with the LAN controller on the card. However, if the MAU is external, an *AUI cable* or *transceiver cable* couples the DTE to the AUI connector of the MAU. The MAU transceiver has an associated power supply run off 12 volts supplied over the AUI cable or integrated interface.

For Thick Ethernet, a *tap connector* frequently attaches the MAU onto the thick coax without need to interrupt network operation to cut the cable and install connectors. Figure 4.7(a) has end-on views of male and female *Type-N* connectors, which may be used instead of a piercing tap. Figure 4.7(b) shows a segment of thick coax with a male Type-N connector that has been screwed onto a Type-N barrel (double-female) connector. The barrel is required to re-couple segments previously cut in two for installation of a MAU using connectors.

AUI connector

Figure 4.7(c) and (d) are an end-on and side view, respectively, of the 15-pin male *AUI connector* (also known as a DIX-type connector). The male connector is essentially a common D-shell-type connector with added retaining posts. The female AUI has a sliding clamp used to retain the posts of the male connector. Figure 4.7(e) and (f) are views of the female AUI connector with the clamp in the unlocked and locked positions, respectively.

Figure 4.7(g) is a view of male (top) and female (bottom) BNC connectors, used for Thin Ethernet. Figure 4.7(h) is a view of a BNC T-connector with a segment of thin coax connected on the left of the T and a BNC 50-ohm terminator shown on the right of the T. Thus, the illustration corresponds to a node at one of the far-ends of a Thin Ethernet network. The center of the T would connect to a MAU.

Figure 4.7(i) is a view a UTP cable with the outer insulation removed to show the four twisted-pairs contained within. Some variations of Ethernet use all four pairs per node. Figure 4.7(j) and (k) are views of an 8-pin modular RJ-45 plug connector and its mating socket connector, respectively.

INDUSTRY STANDARDS

Ethernet and Fast Ethernet Manufacturers

A plethora of technical information and white papers is available from the numerous manufacturers of Ethernet (10-Mbps) and Fast Ethernet (100-Mbps) products. The main URL is:

http://www.gigabit-ethernet.org/

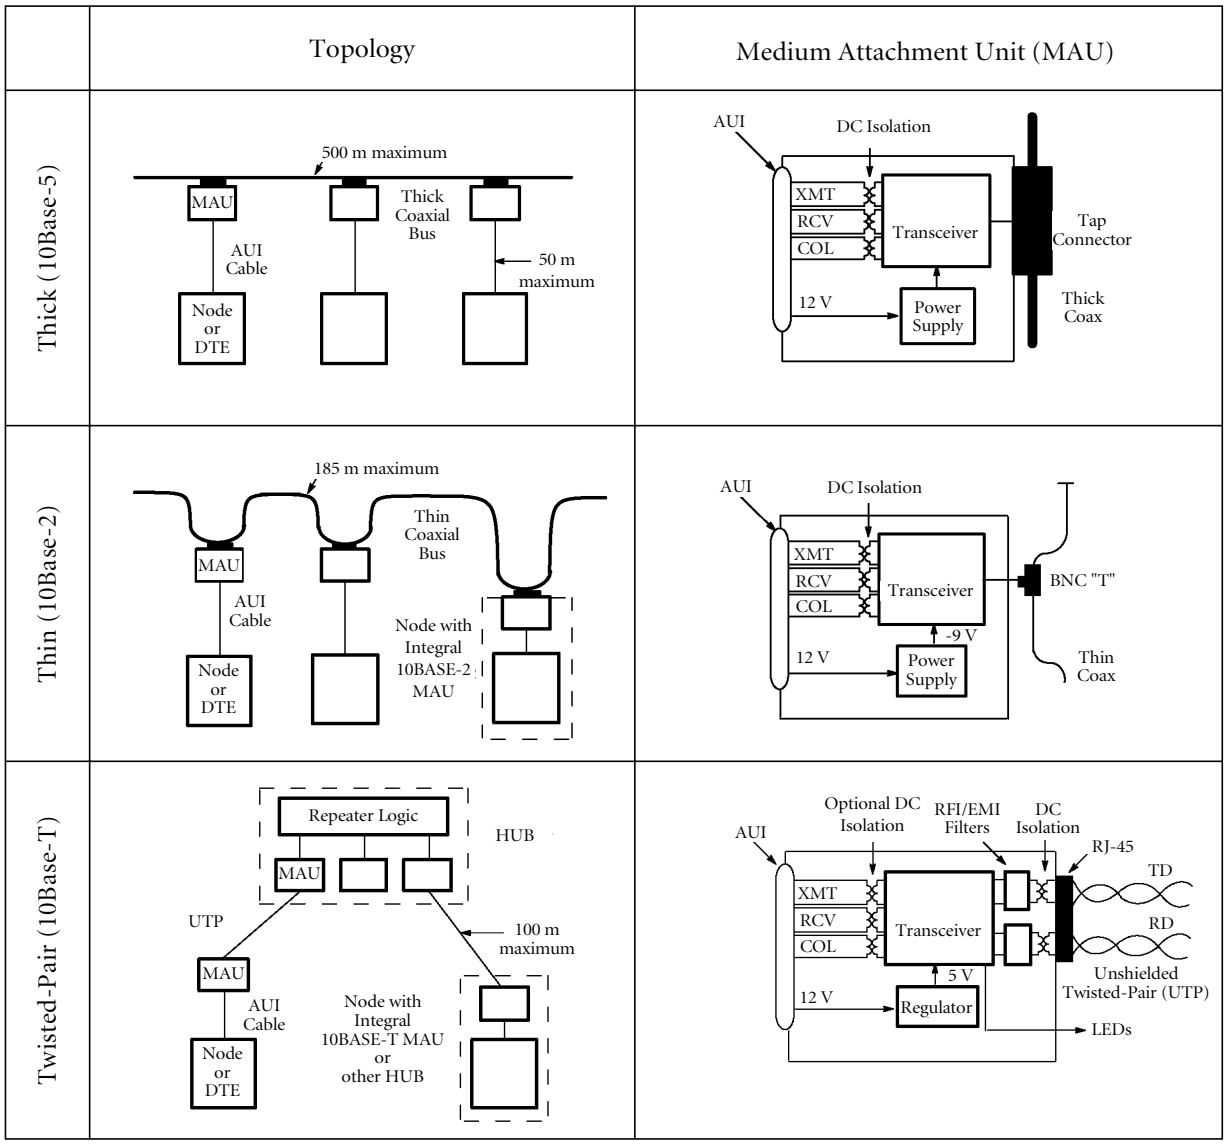

Figure 4.6 Comparison of Ethernet Types

Adapted with permission of Advanced Micro Devices, Inc., from "AM79C960 PCnet-ISA Technical Manual," Copyright 1992,

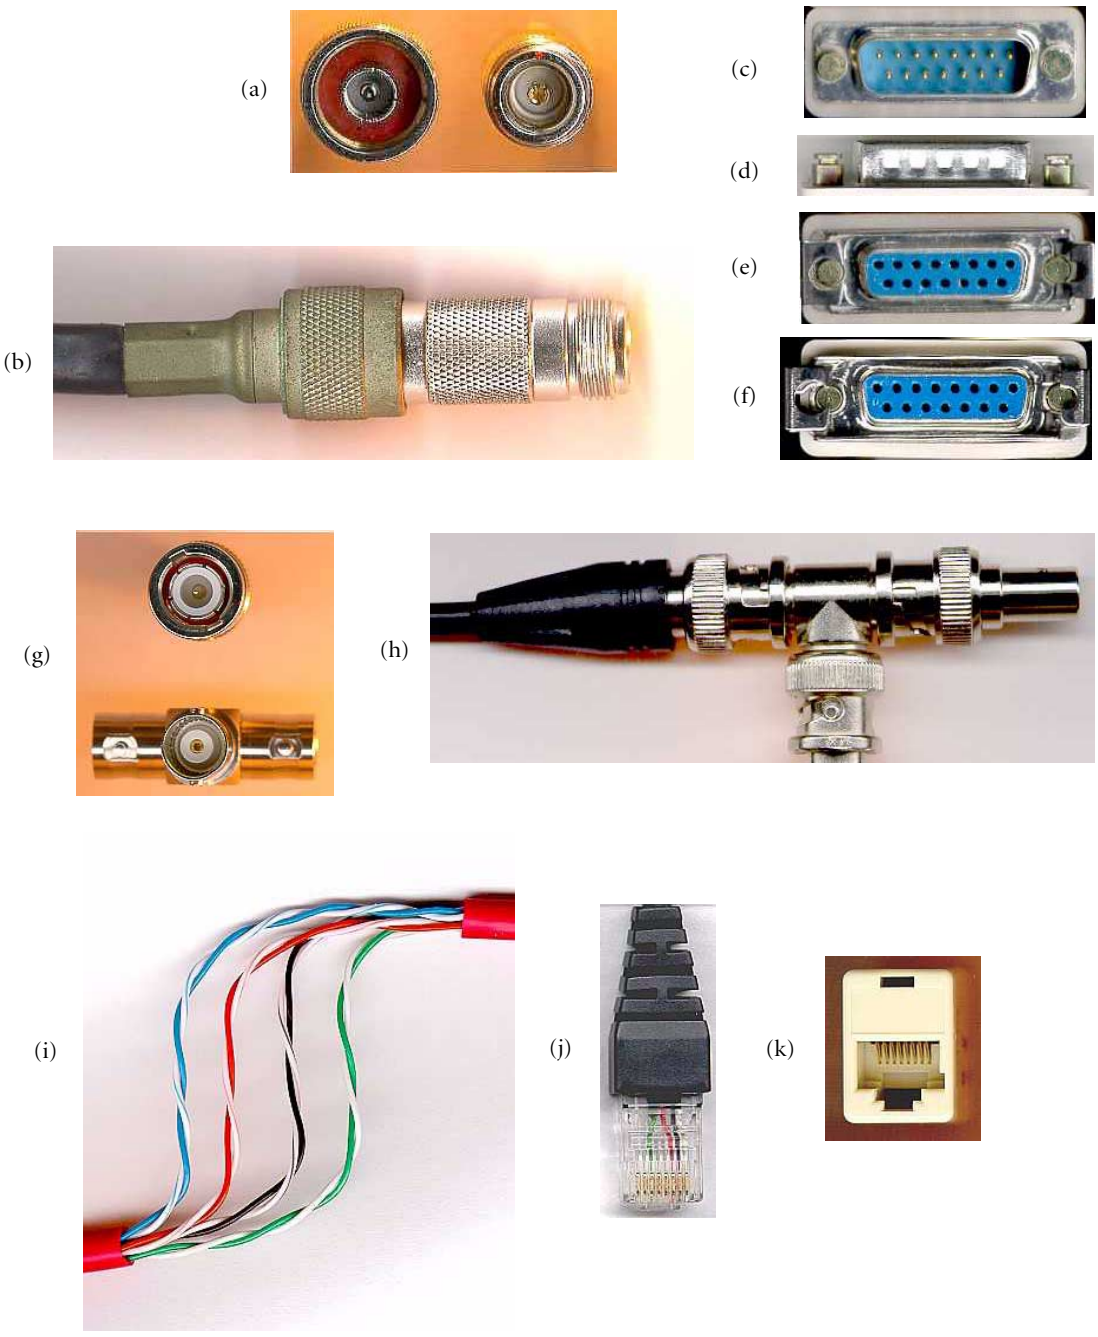

Figure 4.7 COMMON ETHERNET CONNECTORS AND CABLES

IEEE Standards on CD-ROM

IThe 802 Architecture and Overview Specification provides an overview to the 802 family of Standards. It describes how the 802 family relates "*to the Open Systems Interconnection (OSI) Basic Reference Model... and explains the relationship of these standards to higher layer protocols...*"

The 8802.2 standard defines a *Logical Link Control (LLC)* protocol. Logical links are *virtual circuits* that establish a communication session between two nodes. The standard defines various types of communication services over logical links. The LLC protocol applies to any Media Access Control method, not just Ethernet.

The 8802-3 standard defines the implementation of CSMA/CD methods for different media types including coax, twisted-pair, and fiber-optic cables.

A copy of several Ethernet Standards documents standard are included on the CD-ROM.

- *802-1990 IEEE Standards* for Local and Metropolitan Area Networks: Overview and Architecture Information Technology—Telecommunications and information exchange between systems—Local and Metropolitan area networks—Specific requirements.
- *8802-2: 1994* (ISO/IEC) [ANSI/IEEE 802.2, 1994 Edition] Part 2: Logic Link Control.
- *8802-3: 1996* (ISO/IEC) [ANSI/IEEE Std 802.3, 1996 Edition] Part 3: Carrier Sense Multiple Access with Collision Detection (CSMACD) Access Method and Physical Layer Specifications.

Requirements Summary for PC Platform Network Devices

- Generic device requirements apply
- Device driver compliant with Network Driver Interface Standard (NDIS) version 5.0, which implements much of the Media Access Control (MAC) protocol in the OS. The vendor supplies only a miniport driver tailored to the device, only makes NDIS defined library calls to the kernel, and uses new INF files to characterize the device
- Works correctly with Microsoft's implementations of the popular TCP/IP, IPX/SPX, and NetBEUI protocols
- Driver supports modes of operation to enable network monitoring by administrative applications
- Automatic sensing and appropriate configuration for full-duplex operation, transceiver type, and connection to network
- To prevent the adapter buffer from overburdening software with packet alignment, the adapter performs byte alignment on transmission and performs quadword alignment, or smaller, on receive
- Adapter supports push technology via hardware filtering (match detection) of a minimum of 32 multicast addresses
- If used to install the OS on a new PC, must support the Dynamic Host Configuration Protocol (DHCP) standard.
- Device supports wave-up on match to specified network address
- ISDN, ADSL and other adapters have additional special requirements

Applications of Networks

LANs, such as *Ethernet*, provide a means to electronically transfer files at high speeds (commonly 10 Mbs migrating to use of 100 Mbs) between interconnected PC platforms within the confines of small offices or workgroups. The file transfer can be of an explicit nature, such as when electronic mail is used, a disk backup is performed, or when someone wants to copy or move files from one networked platform to another. The file transfer can also be of an implicit nature, such as when a disk drive at a remote node is mounted for use over the network.

workgroup computing

LANs have enabled *workgroup computing*, which includes the ability to send electronic mail to colleagues on the LAN; to selectively share data files, directories, or entire disks; and to share expensive peripherals. The shared peripherals may include high-end color printers, high-capacity backup resources, and high-speed access ports (gateways) to other networks, including the Internet. Electronic mail can be simple bulletins and messages, or it can include *attachments*, of logically separate files, such as graphics, word processing files, or programs.

attachments

Platforms on LANs are frequently characterized as *servers, clients,* or *peers*. A platform that has a shared resource attached is referred to as a server for that resource. For example, one speaks of disk or file servers and printer servers. The platforms that connect to the server to make use of the server's resource are *clients*. The larger the client population, the more important and economically pragmatic it is that the server be a high-end platform equipped with the most superlative features available. In small office environments, the distinction between clients and servers is usually blurred, and peer-to-peer networking between desktop platforms is the norm.

clients

peer-to-peer networking

Multiple LANs can be coupled and interconnected via a variety of techniques to encompass an entire site or multiple sites at a large company. Such an interconnection is often referred to as an *Enterprise Network*. Within an isolated site, networking addressing and protocols must be compatible, but can be largely arbitrary. Beyond an isolated site, different sites may be linked together via either private or public high-bandwidth communications channels. Networking more than one site requires destination addresses to be coordinated and properly assigned and a common communications protocol to be used.

Enterprise Network

intranet

An *intranet* is an Enterprise Network that uses the Internet standard Transmission Control Protocol/Internet Protocol and related HTTP addressing standard. By supporting the same Web browsers and other tools used for Internet access, software license and training costs are minimized. Intranets are used for intra-company departmental and corporate-headquarters Web sites, enabling easy company-wide access to standards, publications, and services.

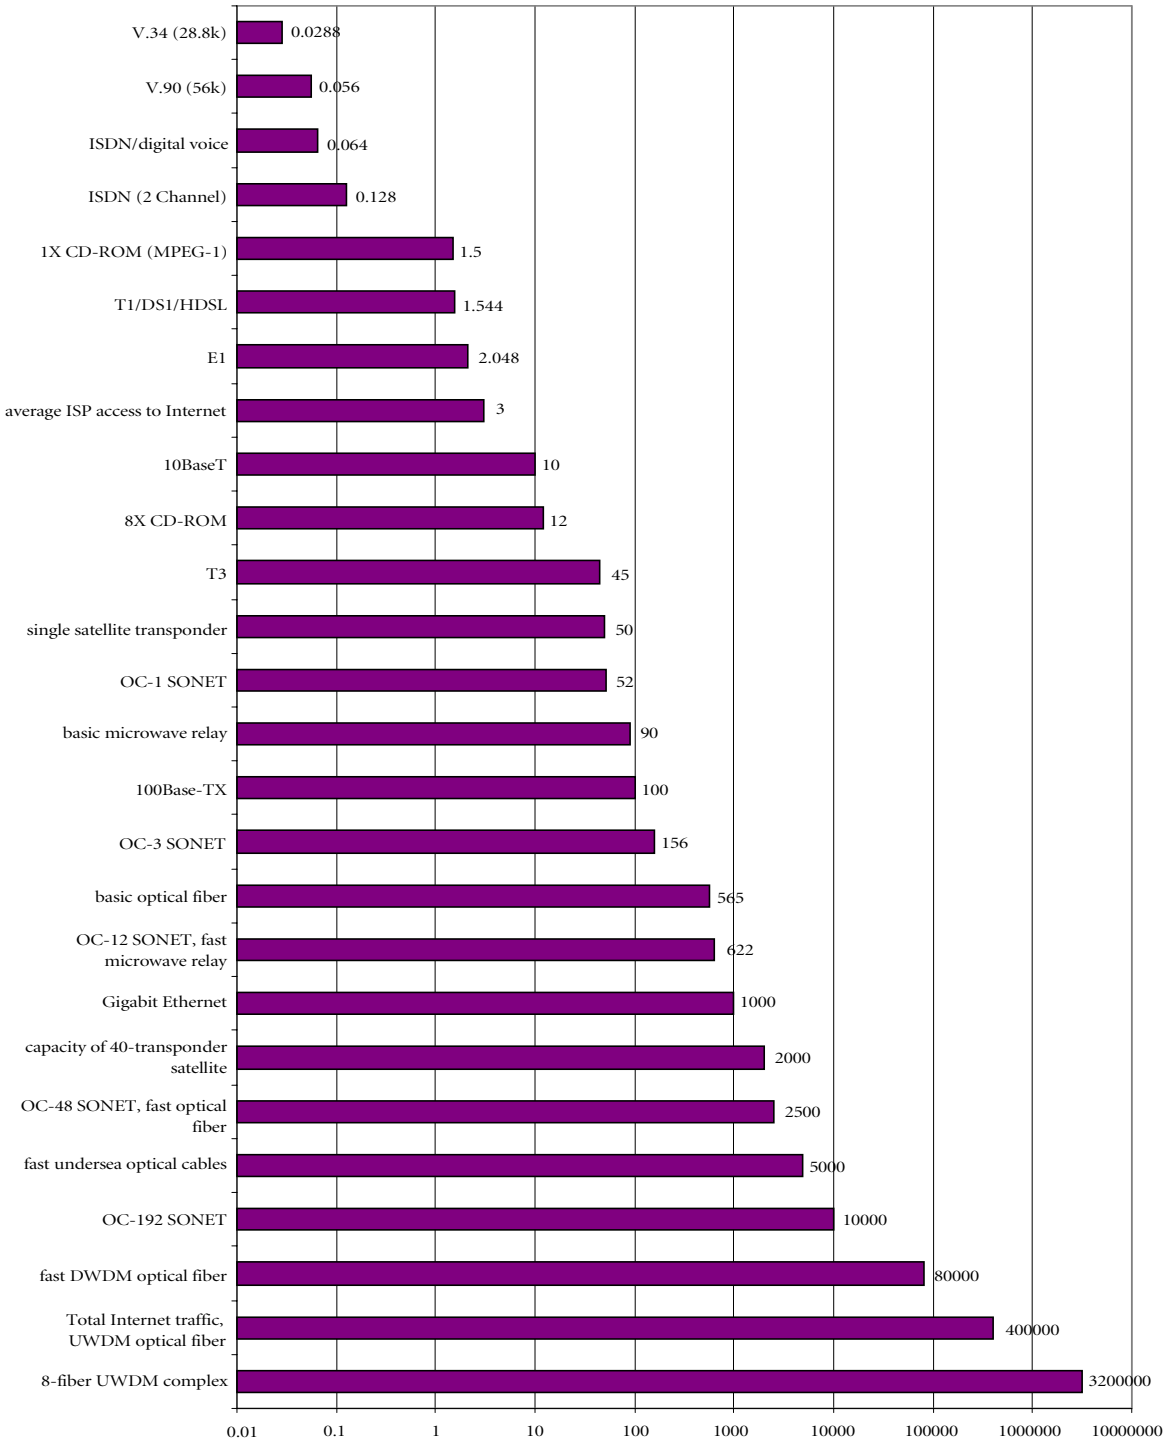

Figure 4.8 BANDWIDTH OF CHANNELS AND LINKS (BANDWIDTH IN MBITS PER SECOND, LOGARITHMIC SCALE)

Comparison of Communication and Network Bandwidths

Figure 4.8 on page 363 provides a comparison of bandwidths (shown in Mbps) for a variety of common channels and links.[54] A number of interesting observations can be made from this figure, where orders of magnitude differences in bandwidth are readily apparent. The present total Internet traffic is shown at 400 Gbits per second. This is also the capacity of an Ultra Dense Wave Division Multiplexed (UWDM) single optical fiber. *Wave Division Multiplexing* (WDM) techniques are being used to substantially increase the capacity of many existing installed fibers, by modulating multiple light sources of different wavelengths. The 8-fiber UWDM complex represents the state-of-the-art in optical network systems at the time of this writing.[55] With continued reference to Figure 4.8, we will next examine the different Ethernet variations against the various data communications rates found in the telephone network and against CD-ROM transfer rates.

Wave Division Multiplexing

10BaseT

Note that 10BaseT is only roughly five times slower than the T3 leased phone line rate at which most Internet backbones operate. 10BaseT is also roughly three times faster than the average rate at which Internet Service Providers connect to the backbone. Medium and small-size businesses often are connected to the Internet via dedicated T1 leased phone lines, which are then coupled to individual office PCs on multiple 10BaseT segments. For light loads on both the LAN and the T1, each PC should enjoy basically the equivalent of their own T1 connection. Under today's normal business-user loads of light Web browsing, periodic e-mail transmissions, and occasional large file downloads, a single T1 line could support over fifty simultaneous users with V.34 (28.8k) equivalent performance. However, video is another matter. A single MPEG-1 bit stream from a 1X CD-ROM is roughly the same as the bandwidth of the T1. A modern 8X or better CD-ROM can deliver more bandwidth than 10BaseT.

100BaseTX

A single satellite transponder has sufficient bandwidth to support a single commercial-quality color TV broadcast. A basic microwave relay has sufficient bandwidth to support a thousand voice channels. 100BaseTX has comparable bandwidth to both of these.

---

[54] Many of the more obscure link names in this figure are described in the following section on the Internet.

[55] *Lucent Technologies delivers record-breaking optical networking capacity; five times greater than current systems*, Lucent press release, January 26, 1998. Dense Wave Division Multiplexing (DWDM) was a precursor to UWDM.

Gigabit Ethernet

Gigabit Ethernet has roughly half the bandwidth of an entire 40-transponder satellite. Thus its raw bandwidth compares to that of 20 color TV channels. Or, you could compare it to ten-thousand voice channels! Gigabit Ethernet also exceeds the bandwidth of the OC-12, the fastest rate planned for near-term upgrades to most Internet backbones.[56]

The Internet

Internet Access

The *Internet* is the name given to an interconnected conforming system of networks used for global file transfer and data communications. The Internet enables PC platforms to run applications that support global e-mail, Web browsing, and audio and video conferencing. The Internet in its current form was enabled by the advent of the HTML document (Web page) publishing standard, the HTTP Web navigation standard, and the proliferation of HTML readers (Web browsers). The number of people with access to the Internet is growing rapidly, supplanting the use of both electronic bulletin boards systems (BBSs) and other limited/proprietary access services. Responsible for this growth is the perception of readily available quality content and the now widespread availability of low-cost Internet connections.

In the U.S., the Internet is a meta-network interconnection of more than a dozen major national high-speed networks and another dozen major regional high-speed networks independently operated by competing communications companies (network providers). The major national and regional networks consist of data links and packet routers that are logically structured around one or more 45 Mbs (known as *T3*), or higher, *backbone* links, which couple, ring, or span, major metropolitan areas. The backbones of the national and regional providers are interconnected at multiple high-bandwidth junctions called peer interconnect points. In the U.S. there are a dozen major peer interconnect points.

backbone

Depending on the provider and usage, a 45-Mbps T3 leased line costs between $20,000-60,000 per month at the time of this writing, while a 1.544-Mbs T1 leased line costs between $1,400-3,000 per month.[57] These costs do not include setup and required interface equipment. Nearly all backbone providers have migration plans that include backbones at 155 Mbs (*OC-3*), and at 622 Mbs (*OC-12*). In addition to the backbones, these

[56] Optical Carrier Level One (OC-1) is the basic building block channel capacity for the Synchronous Optical Network (SONET) standard. Each OC level is corresponding multiple of the OC-1 rate of 51.84 Mbps.

[57] All statistics in this section are based on data from the *Boardwatch Magazine Directory of Internet Service Providers*, Fall 1996.

networks will also have dozens to hundreds of 1.544-Mbs (*T1*)[58] links to serve the metropolitan areas surrounding the major metropolitan nodes of the backbones. Both the backbones and the lesser links may be owned by the backbone providers or leased from the major long-distance exchange carriers and the local telephone companies.

The national and regional network providers lease dedicated 58 Kbps-45 Mbps Internet access to major corporations, institutions, and over 3000 independently operated Internet Service Providers. The three national providers MCI, Sprint, and UUNET provide backbone access to 79% of the ISPs. ISPs on average have access to a backbone at 3 Mbs. The ISPs in turn provide Internet access to businesses and consumers via over 11,000 "points of presence" (POP), to which dedicated leased lines are connected, or as needed dial-up connections may be made.

Internet-Based Applications and Uses

The Internet uses the standardized TCP/IP (Transmission Control Protocol/Internet Protocol) for all communications links, routers, connected platforms, and compatible communications applications programs. TCP/IP requires messages to be logically partitioned into multiple "packets." Multiple data links are connected to routers, which attempt to relay the packets closer to their destination, based on an associated address. The packets may pass through an arbitrary number of data links and routers on their path from source to destination. Furthermore, not all packets from the same source message need take the same path. TCP/IP permits point-to-point electronic messaging, file transfers, and remote access log-in, wherever such privileges are enabled, throughout an arbitrarily sized network.

Commonly used Internet-based application programs and associated TCP/IP-compliant sub-protocols include: electronic mail (e-mail) using Simple Mail Transfer Protocol (SMTP), Post Office Protocol 3 (POP3) and newer protocols; interactive teletypewriter-like network access to remote computer systems via TELNET; file transfer using File Transfer Protocol (FTP); exchange of Users Network (USENET) "newsgroups" (electronic conferences) via Net News Transfer Protocol (NNTP); and Web publishing using Hyper Text Markup Language (HTML) and Web navigating using Hyper Text Transfer Protocol (HTTP).

---

[58] *E1* is a European variant of T1.

Industry Standard

Internet Protocol Standards

TCP/IP and other Internet-related protocol standards are developed by the Internet Engineering Task Force (IETF) and recommended by the Internet Engineering Steering Group (IESG), both subsidiary groups of the Internet Architecture Board (IAB). Information on the IETF can be found at

http://www.ietf.org/

RFC 1920, "Internet Official Protocol Standards," is an overview of Internet standards and the standards process. Links to this overview (contained in the file "std1.txt") and the Internet protocol standards themselves can be found in the /in-notes/std/files subdirectory at URL:

http://info.internet.isi.edu/.

Other Internet protocols are found in the /pub/WWW/Protocols/ subdirectory at URL:

http://www.w3.org/.

Suggested Readings

Data Communications and Networking

Two texts that provide good tutorial overviews of data communications and networking technology and standards are:

1. Tom Sheldon, *LAN TIMES Encyclopedia of Networking*, Osborne McGraw-Hill, 1994.
2. John G. Nellist, *Understanding Telecommunications and Lightwave Systems, an Entry Level Guide,* 2nd Edition, IEEE Press, 1996.

Component-Level Technology

We have defined component-level technologies as those standards, initiatives, and technologies that have a focus generally limited to a single type of component or subsystem. Component-level technologies are often enabling technologies for the platform-level technologies just examined. For example, the new USB and IEEE 1394 buses, surveyed in this section, are key to the Device Bay and Plug and Play initiatives previously discussed.

Our survey of component-level technologies is carried out in four subsections: *Processors*, *Storage Devices*, *General-Purpose Buses*, and *Device-Specific Buses and Ports*. Of course, each of these technologies has

its own field of study to which volumes of texts have been devoted. We examine them here to provide a broad view from an overall platform architecture perspective and to ensure that the reader is literate with respect to the most important components and subsystems used in platforms. The detailed structure of this section is as follows.

Suggested Readings

Additional Resources: Mindshare and Annabooks are two companies that write books and sponsor seminars on various aspects of platform technology. Each operates a web-site. The URLs are:

http://www.annabooks.com/

http://www.mindshare.com/index1.html

Processors

Processor performance is generally not an issue for basic word-processing and data-entry applications.

Computation performance is important to application and system software emphasizing presentation quality and ease-of-use features, new 3D graphics applications, software-based DSP, and many scientific and engineering applications. Presentation quality and ease of use are now baseline features of PC platform application and system software. These features greatly enhance user accessibility and productivity, as outlined earlier, but add significant computational overhead to the platform's primary processor.

Real-time processing requirements or extensive data filtering and manipulation often characterize 3D graphics, DSP, and other scientific and engineering applications, leading to increased processor performance requirements. In these areas, the time until results are obtained is often directly proportional to processor performance. In some applications this can be merely a matter of productivity. For 3D graphics and DSP human interface applications, processor performance is usually directly proportional to the quality of the result. Higher quality results mean more realistic simulations that do a better job of engaging the user.

To support the initiatives for *Enhanced User Experience*, discussed earlier in this chapter and as a reflection on trends in PC platform applications, processor performance is being increased through Instruction Set Architecture, physical design, and microarchitecture enhancements. Particular x86 Instruction Set Architecture, enhancements include the MMX and AMD-3D instruction sets for communications, multimedia, and 3D geometry functions. Processor clock rates are being increased steadily due to continual advances in circuit design and semiconductor processing, as illustrated in Figure 4.9.

The number of instructions per clock cycle (IPC)[59] has also increased as advanced microarchitectural techniques (both new and old) are applied to microprocessors. These techniques attempt to increase resource concurrency and utilization, reduce the frequency of pipeline stalls and flushes, and attempt to exploit the full extent of instruction level parallelism (ILP) in the program being executed.

Techniques used to increase processor performance in the K6 3D are discussed in detail in Chapters 2 and 3.

Articles on CD-ROM

Michael Bekerman's and Avi Mendelson's article,, "A Performance Analysis of Pentium Processor Systems," *IEEE Micro*, pp. 72-83, October 1995, describes a performance study made by Intel on the P5 architecture and speculates on how to improve it.

The rate of continued increases in microarchitectural performance (in terms of SPECint92 benchmark results normalized by clock speed) is shown in Figure 4.10.

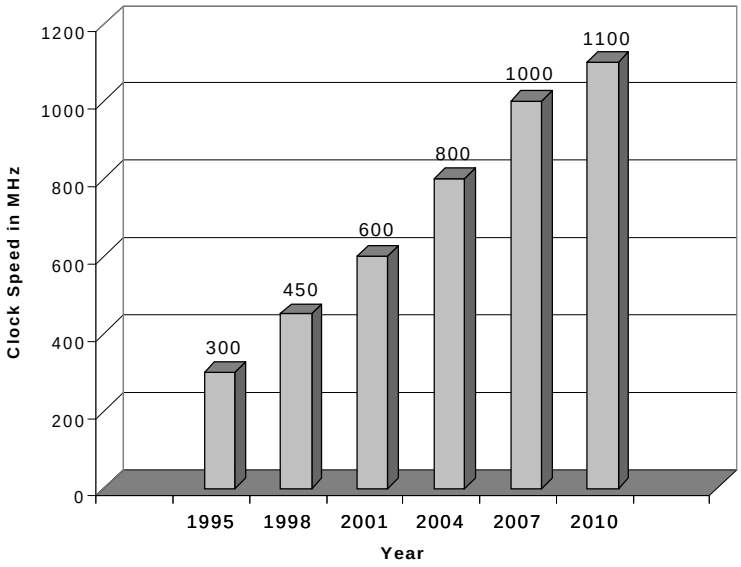

Figure 4.9 Expected Growth in Microprocessor Clock Speed
Based on Semiconductor Association (SIA)Forecast

The graph suggests that the rate of microarchitectural performance growth is slowing and is presently in a range of 10% growth per year. Overall future platform performance growth will thus likely come mostly from frequency and system architectural improvements, the latter

[59] The reciprocal of IPC, Cycles per Instruction (CPI), is also a frequently used metric.

including optimized connectivity and increased bandwidth to other processing and data storage resources.

.

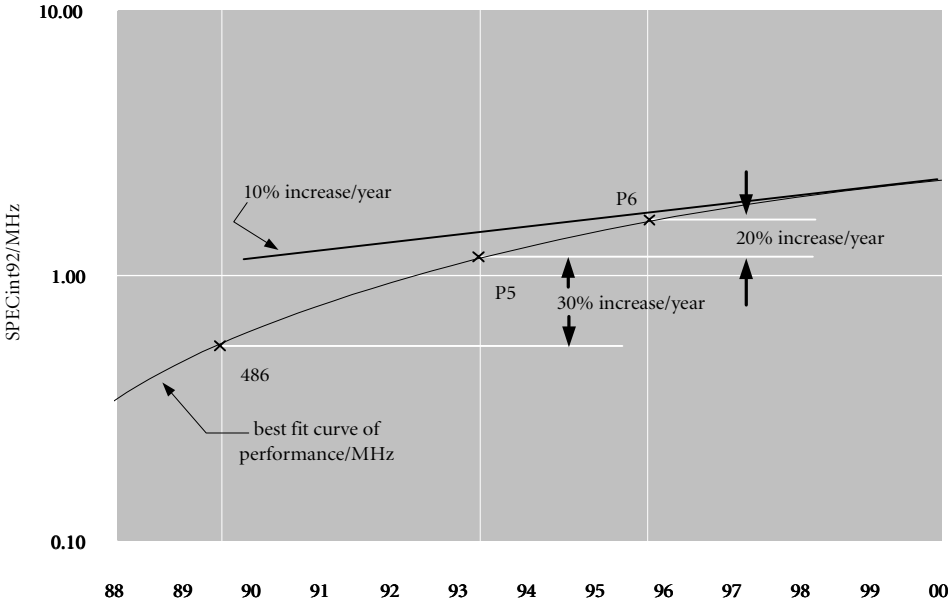

Figure 4.10 MICROARCHITECTURAL PERFORMANCE VS TIME

Adapted with permission of Advanced Micro Devices Inc., from *Competing with Intel*, Copyright 1997.

STORAGE DEVICES

For mobile devices, hard-disk power consumption and management is also a key concern.

In conventional PC platforms, the hard disk is absolutely crucial to system operation and performance. The hard disk is the repository for system and application software and data sets, and holds virtual memory pages not presently in main memory. Fortunately, recent history has delivered hard disks with increasing density and capacity and decreasing costs.

The raw performance of hard disks, in terms of both rotational latency and sustained throughput, is proportional to the rotational speed of the platters. Also significantly impacting hard-disk performance is the chosen electrical interface, of which there are several. Because each of these interfaces can be used for a number of different device types besides hard disks, we have chosen to treat the interfaces separately. Readers are encouraged to separately examine the interface sections on SCSI, ATA/IDE, and the emerging USB, IEEE 1394, and Device Bay standards, according to their interests.

From a system perspective, hard-disk performance is optimized through data staging techniques, in particular the use of a disk cache, and the management of the disk cache using write-back techniques. Hard-disk performance is also optimized by careful attention to the block size and

per block latencies, for large file transfers. Finally, the system as a whole is performance optimized when hard disks are managed using a bus-mastering controller. These optimization techniques are discussed in a more general context in the Overall System Architecture Performance Optimization of Chapter 6.

Besides the hard disks discussed above, other common platform storage devices include floppies, tape drives, CD-ROMs and DVD-ROMS, and a number of removable media technologies (removable hard disks). The section that follows will examine CD-ROM and DVD-ROM technology in more detail. .

Requirements Summary for PC Platform Storage Devices

- Generic device requirements apply
- Device and controller support bus master operation
- Removable media supports media status notification
- System BIOS or Option ROM support for Int 13H Extensions (includes high capacity drives and consistent drive-letter mapping across OS operating modes), and in CD-ROM or DVD systems, the controller must support the El Torito standard for the CD-ROM or DVD installation process
- Device and file system run in protected mode[a] following installation
- Driver for partitioned media supports all Windows partition types
- IDE hard drives must be compliant with Self-Monitoring, Analysis, and Reporting Technology (SMART) IOCTL API Specification, Version 1.1, or higher

[a] A privileged, or system, mode of the processor that protects unprivileged or user mode execution threads from each other.

Requirements Summary for PC Platform Optical Storage Devices

- General storage requirements apply
- Rewriteable drives meet SFF 8070i specification
- A CD-ROM drive uses CD-Enhanced compatible support to mount multisession CD-ROM discs; is compatible with CD Red, Yellow, White, and Blue Book logical format standards; meets SFF 8020i, Version 1.2 or higher (requirements for ATAPI devices); and provides 8x (12Mbits/s throughput) or higher performance
- A DVD (if present) uses a high-speed expansion bus; is compatible with CD Red, Yellow, White, and Blue Book logical format standards; supports the Universal Disk Format Specification (UDFS), Version 1.02, or higher; its device driver meets the SFF 8090 specification; and it supports CSS copyright protection to enable decryption and prevent duplication of CSS protected media

Industry Standard

Small Form Factor (SFF) Committee Specifications

The Small Form Factor (SFF) Committee was formed in 1990 to quickly promulgate standards for small devices, such as those typically found in PC platforms. It is administered by ENDL Publications. Observers may joint for $300/year and full membership is $1,800/year. There is a SFF related FTP-site. The URL is:

ftp://fission.dt.wdc.com/pub/standards/SFF/specs/

General-Purpose Buses

This section surveys interconnect designs that support communication between many arbitrary devices. These buses are distinguished from the next subsection on *Device-Specific Buses and Ports*, which generally couple only a very few (often, just two) devices or subsystems, and frequently implement device-specific protocols.

Distributed Peripheral Buses

The buses in this section are designed to couple multiple devices over distributed and user-extensible paths that are long compared to conventional system expansion buses. These buses are a synthesis of the conventional expansion buses and LANs. The emerging USB and IEEE 1394 buses, and their hybrid combinations in the form of Device Bay and the Enhanced Video Connector (EVC), are key technologies enabling the Plug and Play and SIPC initiatives discussed in the platform-level technology section. SCSI is an established technology with ongoing enhancements. SCSI plays a key role in high-performance hard disks, particularly for server applications.

Universal Serial Bus (USB)

The Universal Serial Bus (USB) is a serial protocol bus intended to support trivial hookup of low-to-medium-bandwidth peripherals. Typical devices include input devices, control functions, telephony/modems, audio, scanners, and printers.

USB interconnections employ a daisy-chained multi-ported hub star topology, superficially similar to a twisted-pair Ethernet connection between a PC and an Ethernet hub. Each hub in the daisy-chain represents a "tier," of which a maximum of 6 tiers are allowed with no more than 5 meters between segments. Signals are repeated across all tiers of the network without any storage delay. Up to a maximum of 127 simultaneous devices are possible in such a "tiered-star topology." Figure 4.11 on page 373 shows a PC system in which a 5-port external hub is connected to a *host controller* (resident in each PC) and 4 other peripherals. The host also acts as a hub, connecting with the external hub and to the keyboard.

host controller

Each segment is a bidirectional, half-duplex link, with error detection and recovery. The connectors are 4-pin, and the cables are 4-wire. 3.3V signaling differential signaling is used with NRZI coding and bit stuffing. 5V at 100 to 500ma (depending on the hubs) is available to supply low-power peripherals. Figure 4.12 on page 373 gives a close-up of a USB connector and a pair of USB sockets. Figure 4.13 is an abstract drawing of the internal detail of the connector and port.

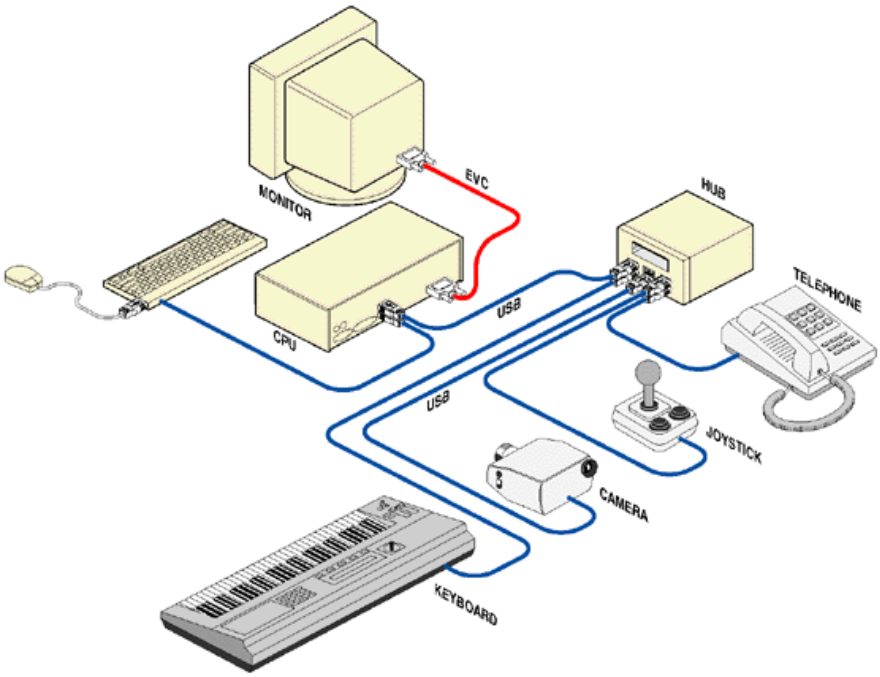

Figure 4.11 EXAMPLE USB INTERCONNECT USE IN A SMALL SYSTEM

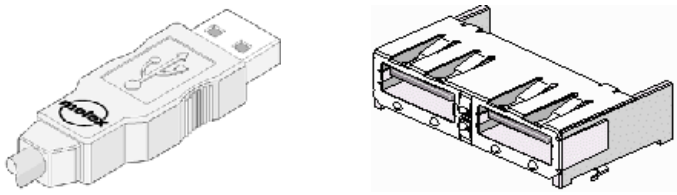

Figure 4.12 USB CONNECTOR AND PORT

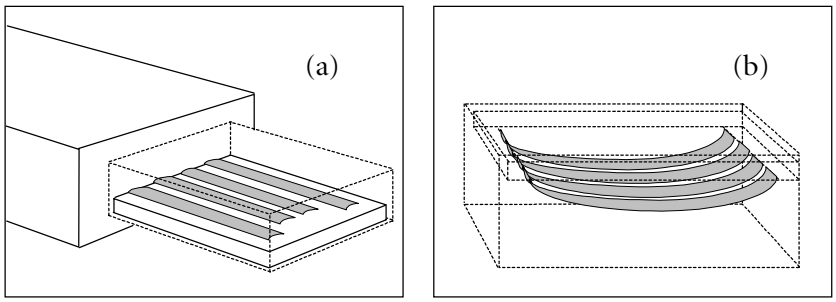

Figure 4.13 DRAWING OF USB CONNECTOR AND PORT INTERNAL DETAIL

Microsoft recommends that any peripheral that consumes more than 30% of the available USB bandwidth should use IEEE 1394 instead.

The host controller manages the bus in many ways. It issues tokens with unique device addresses, to which addressed devices may respond by presenting read data or accepting write data. Transfer types include bulk, isochronous, interrupt, and control transfers. The host controller detects the attachment of new devices, queries the device, determines its capability, assigns it an address, invokes the OS to load the driver corresponding to the device, configures the device, and signals the OS to start the device driver. It performs dynamic reconfiguration in response to any insertion or removal activity. The host controller performs power management for the bus by issuing global suspend/resume commands. It also dynamically allocates the bus bandwidth, including the assignment of isochronous use of the bus, with an aggregate maximum rate of 12 Mbps.

The ability to communicate with 127 devices through only a single connection on a PC platform is a key aspect of the appeal of USB. Such leverage is particularly appealing for laptops, but it also permits a reduced footprint[60] for desktops. Other key aspects include the ability to perform hot insertion and removal of connections, and to perform the connections on the outside of a sealed-case PC.

Industry Standards

Universal Serial Bus Implementers Forum (USB-IF)

The USB-IF describes itself as "is a support organization formed by the seven promoters of USB to help speed development of high quality compatible devices using USB." Intel administers the USB-IF. USB-IF membership is $2500/year.

The USB-IF maintains a Web site. Useful information on joining, white papers, conference presentations, and free download access to the USB specification and clarifications and enhancements is best found in the technical area, under URL:

http://www.usb.org/developers/

Report on CD-ROM

To learn more about Universal Serial Bus (USB), see the article by Michael Slater titled, "Universal Serial Bus to Simplify PC I/O," *Microprocessor Report*, Vol. 9, No. 5, April 17, 1995, on the CD-ROM.

[60] Footprint refers to the amount of furniture-top surface area required by a PC.

Industry Standards

VESA Enhanced Video Connector (EVC)

The Enhanced Video Connector is a hybrid combination of IEEE 1394, USB, digital audio, and both analogy and digital display monitor drive signals. The EVC permits peripheral devices mounted in the monitor to attach to these individual buses yet only requires a single cable between the monitor and the system unit. The following standards define EVC:

- *VESA Enhanced Video Connector (EVC) ™ Pinout and Signal Standards, Version 2* establishes the electrical pinout and signals for the EVC connector.
- *VESA Enhanced Video Connector (EVC) ™ Physical Connector Standard, Version 1, Revision 2* establishes the physical features of the EVC connector.

VESA standards are only available to VESA members. A complete list of VESA standards with a brief abstract of each is available at URL:

http://www.vesa.org/ve00013.html

The Video Electronics Standards Association (VESA) is discussed in more detail in the section on Graphics and Video Adapters later in this chapter.

Suggested Readings

USB

Don Anderson's book, *Universal Serial Bus System Architecture*, MindShare, Inc., Addison-Wesley, 1997, gives a clear and detailed explanation of the concepts behind USB.

Requirements Summary for PC USB Devices

- Required
- Generic device requirements apply
- All devices compliant with USB Specification, Version 1.0
- The host controller must be compliant with the Open Host Controller Interface (OpenHCI) or Intel's Universal Host Controller Interface (UHCI)
- The host controller can wake the system

IEEE 1394

IEEE 1394 is a PC 98 required high-speed serial protocol bus that provides easy hookup of high-bandwidth peripherals, such as video cameras,

VCRs, video or audio conferencing channels, hard disks, or CD-ROMs and DVD-ROMs. It is cited by many as being a replacement for the present ATA/IDE interface as well as small LANs.

IEEE 1394 offers speeds of 100, 200, and 400 Mbps on the same cable with dynamic speed negotiation between source and destination. Transfers are strictly between source and destination and do not require the involvement of a host. Transfers can be made with guaranteed bandwidth and known latency.

Dynamically assigned device identifiers let each bus support up to 63 devices, connected in a point-to-point fashion. Splitters permit branching at any point, in any direction. Intermediate devices need not be powered and all connections are hot-swappable. Removal or insertion of a device results in transparent dynamic reconfiguration of addresses. Any non-loop topology is permitted, provided that node-to-node distances are a maximum of 4.5m and no more that 16 hops exist between any two nodes. Bus identifiers permit bridging of up to 1023 buses. Figure 4.14 shows an IEEE 1394-1995 external port and plug connector.

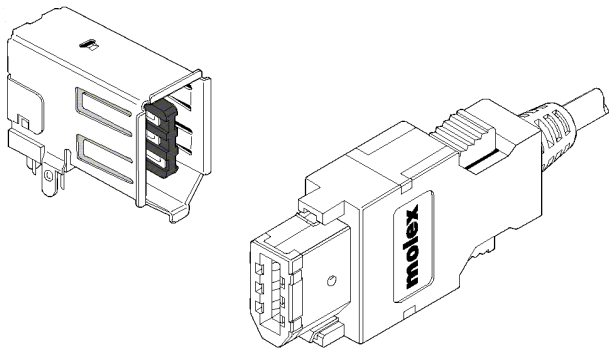

Figure 4.14 Six-Wire External Port and Plug Connector

These are 6-wire connectors, consisting of one pair of power conductors and two pairs of signal conductors. Low-power devices can take power from the cable, up to a cumulative 60 watts maximum. Another external connector pair (not shown) is a 4-wire Audio-Visual (AV) connector that does away with the power conductors. There are also internal device connectors, suitable for motherboard mounting for the mating of either daughterboards or cables.

IEEE Standard on CD-ROM

A copy of the IEEE 1394-1995 standard is included on the CD-ROM.

INDUSTRY STANDARDS

IEEE 1394 Trade Association (1394 TA)

In its own words, "The 1394 Trade Association actively promotes the proliferation of the IEEE 1394 Serial Bus standard technology into the computer, consumer, peripheral, and industrial markets to enable a truly interoperable, standardized, universal I/O and backplane interconnect." Membership fees in 1394 TA are $4000 or $8000/year, based on gross revenues. The 1394 TA maintains a Web site with information on joining, a mailing list, and links to conference presentations and papers. The main URL is:

http://www.1394ta.org/

TECHNICAL PRESENTATION ON CD-ROM

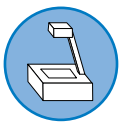

Two IEEE 1394 technical presentations are included on the CD-ROM. The first gives numerous quality illustrations of the various connector types used for IEEE 1394 and its hybrid forms, including the Enhanced Video Connector discussed earlier. The second is an outstanding presentation that gives a detailed overview of the IEEE 1394 protocol.

- Max Bassler, *IEEE 1394 Interface Technology*.
- Michael D. Johas Teener, *IEEE 1394-1995 High Performance Serial Bus*.

REPORTS ON CD-ROM

To learn more about IEEE 1394, see the following reports on the companion CD-ROM:

- Michael Teener, *New Technology in the IEEE P1394 Serial Bus—Making it Fast, Cheap and Easy to Use,* March 31, 1998.
- Curtis P. Feigel, "FireWire Brings Fast Serial Bus to Desktop," *Microprocessor Report*, Vol. 8, No. 3, March 7, 1994.
- Michael Teener, *A Bus on a Diet—The Serial Bus Alternative: An Introduction to the P1394 High Performance Serial Bus*, December 14, 1993.

Requirements Summary for PC Platform IEEE 1394 Devices

- General device requirements apply
- Devices support a peak data rate of 400 Mb/s, minimum
- Controllers must be compliant with the Open Host Controller Interface (OpenHCI) for 1394, which defines standard register addresses and functions, data structures, and DMA models and promotes performance, security, and error handling
- Recommended for all devices needing high-bandwidth and Plug and Play connectivity
- The controller must exhibit robust response to topology faults, including surprise removal, safe removal (device is switched off prior to removal from bus), more than 16 cable hops, or more than 63 devices
- Devices capable of initiating peer-to-peer communications must support a programming language that enables them to be remotely controlled by the platform
- A configuration ROM compliant with a number of special requirements is required to enable the PC to uniquely identify each device on the bus, determine and configure its requirements, and load the correct device driver
- Devices and controllers must comply with the Cable Power Distribution Specification and the IEEE 1394 Specification for Power Management
- Devices must keep their physical interface (PHY) powered at all times when otherwise powered down to enable the device to perform the repeater function
- Devices that consume or source power must report such characteristics and must notify the bus power-manager if the device is switched off
- AC-powered platforms must source power to the bus
- Self-powered devices propagate the power bus through each connector
- Each source supplies a minimum of 20 volts at 15 watts
- Cable and self-powered devices must allow a bus manager to control their power state
- End user cabling purchases and decisions are reduced by generally requiring devices to use a standard 6-pin connector
- Single-port, leaf-node, devices may use the 4-pin connector intended for hand-held devices
- Multiple port devices use all 6-pin connectors, all connectors propagate the power bus, all ports support a common peak data rate, and all cabling supports 400 Mb/s

Articles on CD-ROM

The first article in the following set provides yet another perspective on the IEEE 1394 architecture. The last two articles describe how consumer video applications can make use of IEEE 1394.

- Stephen L. Diamond, "IEEE 1394: Status and Growth Path," *IEEE Micro*, Vol. 16, No. 3, June 1996.
- Alan T. Wetzel and Michael R. Schell, "Consumer Applications of the IEEE 1394 Serial Bus, and a 1394 DV Video Editing System," *ICCE*, June 1996.
- Adam J. Kunzman and Alan T. Wetzel, "1394 High Performance Serial Bus: The Digital Interface for ATV," *IEEE Transactions on Consumer Electronics,* Vol. 14, No. 13, August 1995.

Small Computer System Interface (SCSI)

The Small Computer System Interface (SCSI) has been a popular interface for peripheral devices and particularly so in high-end platforms. SCSI is second only to the IDE/ATA interface in popularity for attachment of hard disks. Variants (and their theoretical peak bandwidths) include SCSI-1 (5 MBps), SCSI-2 (5 MBps), Fast SCSI-2 (10 MBps), Wide SCSI-2 (20 MBps), Ultra SCSI or SCSI-3 (20 MBps), and Wide Ultra SCSI (40 MBps). SCSI is frequently used for high-performance hard disks, CD-ROMs, other media, and scanners. SCSI bus mastering controllers were widely available earlier than with IDE/ATA. Consequently, SCSI has been adopted for some time as the preferred hard disk interface in file servers, where the ability to have multiple simultaneous file-transfers and other execution processes is essential. SCSI shares attributes with LAN technologies, in that node devices (up to 8, or up to 16 for Ultra SCSI) are coupled over an extendable distributed bus, often at least partially external to the host platform. Figure 4.15 shows commonly used SCSI connectors.

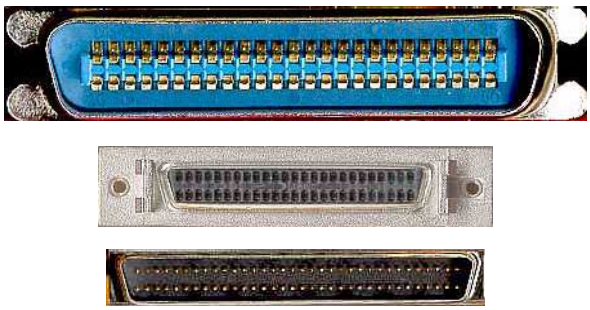

50-pin Centronics-style male

50-pin high-density female

68-pin high-density male

Figure 4.15 Common SCSI Connectors

The 50-pin Centronics-style connector was used for SCSI-1 devices. The 50-pin high-density connector is used for SCSI-2 and 8-bit Ultra SCSI devices. The 68-pin high-density connector is used for Wide SCSI-2 and Wide Ultra SCSI. The PC Specifications plan to supplant the SCSI interface with IEEE 1394, which offers enhanced expansion capabilities and performance.

Industry Standards

SCSI Trade Association

The SCSI Trade Association is involved in promoting and explaining the technology. The SCSI Trade Association operates a Web site with information on joining and downloadable white papers and technical presentations. The main URL is:

http://www.scsita.org/

Requirements Summary for PC SCSI Devices

- General device requirements apply
- Bootable controllers must support the El Torito standard for the CD-ROM installation process
- Controllers must support virtual DMA services (including scatter/gather)
- Controllers must include support for SCSI-3 defined DIFFSENS, which senses and configures for either differential or single-ended peripherals, and also support SCSI-3 automatic termination, which permits external devices to be added without opening the case
- External connectors must be a high-density connector according to the SCSI-2, or higher, specification
- Connectors must clearly identify the bus as SCSI and external connectors must display the SCSI icon connectors
- Peripherals must implement on the parity SCSI bus
- Cabling must meet clause 6 of the SCSI-3 spec, shielded device connectors must meet SCSI-2 or higher spec, users must be prevented from incorrect cabling of either internal or external devices, and external devices must use automatic termination or an externally accessible switch
- hardware supports the SCSI-2 spec defined STOP/START UNIT command used by software to conserve power or resume normal operation

Industry Standards

SCSI Specifications

Approved SCSI-related standards are available only by purchase from the American National Standards Institute (ANSI). The available standards include:

SCSI-2 CAM Transport and SCSI Interface Module X3.232:1996

SCSI-3 Architecture Model (SAM) X3.270:1996

SCSI-3 Controller Commands (SCC) X3.276:1997

SCSI-3 Fast-20 Parallel Interface (Fast-20) X3.277:1996

SCSI-3 Interlocked Protocol (SIP) X3.292:1997

SCSI-3 Parallel Interface (SPI) X3.253:1995

SCSI-3 Primary Commands (SPC) X3.301:1997

Small Computer System Interface - 2 (SCSI-2) X3.131:1994

SSA SCSI-2 Protocol (SSA-S2P) X3.294:1996

ANSI maintains a Web site with a searchable online catalog and electronic ordering of specifications. Prices of each specification vary. The main URL is:

http://www.ansi.org/

Secondary SCSI Technical Documents

Draft SCSI-related specifications are available from the T10 Web site. The main URL is:

http://www.symbios.com/x3t10/

The ISA Bus

DEFINITION

Industry Standard Architecture Bus (ISA Bus)[a]

The ISA Bus is the 1984 PC/AT bus that was a 16-bit extension of the 8-bit bus used in the 1981 IBM PC. The IBM PC bus was a near derivative of the processor bus signals generated by an Intel 8088[b] in conjunction with an Intel 8288 bus controller. Similarly, the ISA Bus is a near derivative of the processor bus signals generated by the Intel 80286 used in the original PC/AT. Like these Intel processors, the ISA Bus supports a memory address space and a separate I/O address space,[c] which are nominally used for memory and I/O peripherals, respectively.[d] In addition to basic cycles for accessing memory and I/O, the ISA Bus architecture is considered to include a number of other facilities associated with the programmer visible features of the PC/AT architecture. These features include its interrupt controller, DMA, RTC, CMOS Memory, Keyboard and mouse controller, and system timers.

[a] The derivation of the ISA Bus name was given in Chapter 1.

[b] The Intel 8088 has the same 16-bit internal microarchitecture as an Intel 8086, but has an 8-bit external data bus.

[c] While PC platforms make extensive use of the peripherals, partially or wholly memory-mapped peripherals are also common.

[d] The MEMR# and MEMW# signals are asserted for memory space reads and writes, respectively. The IOR# and IOW# signals are asserted for I/O space reads and writes, respectively.

Figure 4.16 shows the 8-bit portion of an ISA slot, just above the nickel. The ISA edge connector detail is visible on an 8-bit ISA card that is lying on its side above and adjacent to the slot and aligned with it. There are thirty-one contact positions on each side of the 8-bit edge connector, for sixty-two total bus lines. These include 20-bits of address and 8-bits of data, 6 interrupt request (IRQ) lines, and 3-pairs of DMA handshake lines. The full 16-bit ISA slot extended this with an additional eighteen contact positions on each side of another edge connector segment. The additional thirty-six bus lines brings the total number of 16-bit ISA slot bus lines to ninety-eight. The 16-bit extension included the following additional signals: 4-bits of address, 8-bits of data, five additional interrupt lines, and three additional DMA pairs.[61] The new edge connector segment is

[61] A fourth additional DMA pair appears on the motherboard, but not on the ISA bus.

collinear with the original 8-bit edge connector in a manner that permits the reliable use of 8-bit ISA cards in the 16-bit ISA slot.

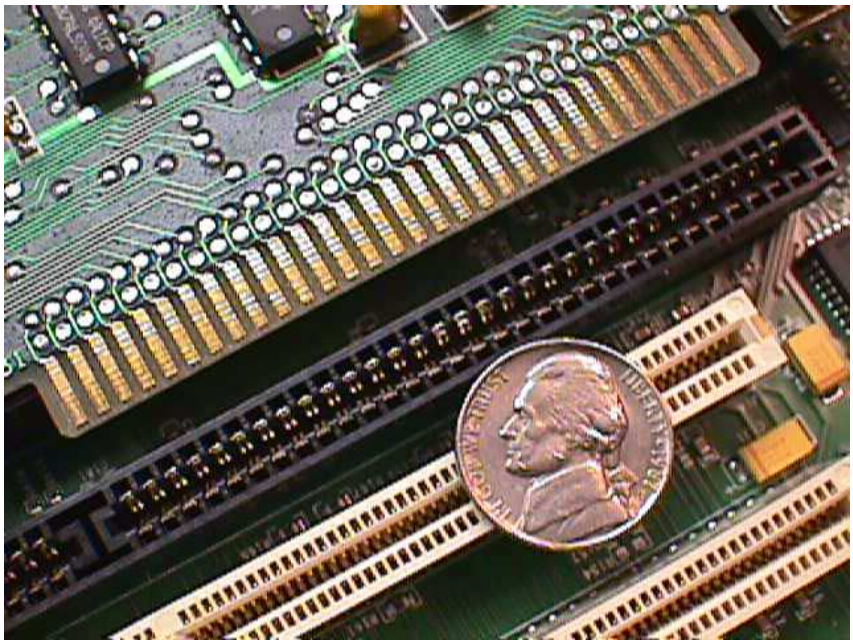

Figure 4.16 VIEW OF 8-BIT ISA CARD EDGE ADJACENT TO 8-BIT PORTION OF ISA CARD SLOT

DEFINITION

The X-Bus (The Extended ISA Bus)

The ISA Bus per se connects only to the ISA expansion slots. However, motherboard implemented 8-bit-wide discrete memory and I/O devices may use the *X-Bus*, or *Extended ISA Bus*, a buffered extended variation of the ISA Bus.

In the PC/AT, X-Bus discrete motherboard peripherals included the ROM BIOS, the RTC, the CMOS Memory, the Keyboard and mouse controller, the DMA controller, and the interrupt controller. Today, the Super I/O module, itself often integrated into the South-Bridge, integrates all of these former X-Bus functions save for the ROM BIOS. (The backup-battery for the RTC and the CMOS Memory also still remains on the motherboard.) Moving the contents of the ROM BIOS into DRAM on power-up initialization is now one of the last remaining uses of the X-Bus on the motherboard.

Protocol

DEFINITION

Bus Protocols

A *bus protocol* is the method the bus uses to communicates address, data, command, and status information between source and destination devices. Two fundamental attributes of bus protocols are whether they are *clocked* and whether they make use of any *handshakes* between the source and destination. A clocked bus includes a clock signal among the bus signals to provide a reference for the sequential logic associated with the bus interfaces at both the source and destination. The buses used for expansion cards on PC platforms are clocked. Peripheral buses are often not clocked. Handshakes are additional status signals frequently included among the bus signals to pace the progress of transfers over the bus. The extent of handshaking varies considerably with different bus protocols. Whether or not a bus handshaken is totally independent of whether it is clocked or not, and vice versa.

clocked
handshakes

The terms *synchronous* and *asynchronous* are also often used to describe bus protocols. Unfortunately, neither of these terms is used consistently in the literature. The reader must generally deduce what is meant from the context in which these terms appear. Synchronous may mean clocked, and asynchronous may mean not clocked. This usage is consistent with general usage in logic design.[a] However, in bus protocols, asynchronous may be used to mean handshaken, independent of whether the bus is clocked or not. Likewise synchronous may be used to mean a bus that is not handshaken, independent of whether the bus is clocked or not.

synchronous
asynchronous

A *semi-synchronous protocol* is one that makes use of optional handshakes. When the optional handshakes are not used, the semi-synchronous bus is said to use *passive termination*, and operates in an entirely synchronous (i.e., not handshaken) fashion using the defined default bus cycle timings.

semi-synchronous protocol

Clocked buses perform transactions in integer multiples of the reference clock cycle. Generally, different transaction types have different *bus-cycle timings*, in that the number of cycles varies with the transaction type. For handshaken buses, each transaction type has a nominal-case[b] bus cycle timing corresponding to the case when both source and destination are ready to proceed at the nominal rate for that transaction. This nominal-case timing can be lengthened by the introduction of additional bus clock cycle multiples via the bus's handshake mechanism. Each additional bus clock cycle introduced beyond the nominal-case timing is referred to as a wait state or a wait cycle.

[a] Unfortunately, synchronous and asynchronous may also be used to describe the extent to which two clocked signals maintain frequency or phase relationships of particular interest.

[b] The nominal-case timing is usually the best-case timing. This is usually greater than a single bus clock period.

An 8.33MHz bus clock (BCLK) is the reference for ISA Bus timing. The ISA Bus has two types of control signals that may be respectively used to optionally shorten or lengthen default bus cycle timings. Thus, the ISA Bus uses a clocked semi-synchronous protocol.

Reports on CD-ROM

ISA Bus Detailed Operation

Two data sheets included on the CD-ROM have valuable supplemental information on the ISA Bus.

From the AMD-645 Peripheral Bus Controller Data Sheet:

- Section 2.2 on pp. 2-3 through 2-5 gives an overview of the ISA Bus controller functions of the South-Bridge.
- Section 4.2 on pp. 4-4 through 4-10 describe each of the ISA bus signals.
- Section 4.4 on pp. 4-13 and 4-14 describe each of the X-Bus signals.
- I/O Read/Write cycles and timing diagrams for ISA Slaves[a] are discussed in Section 5.2.3 on pp. 5-3 through 5-5.
- Memory Read/Write cycles and timing diagrams for ISA Slaves are discussed in Section 5.2.4 on pp. 5-4 through 5-14.
- Memory Read/Write cycles and timing diagrams for DMA ISA Masters[b] are discussed in Section 5.4.1 on pp. 5-18 through 5-10.
- Memory and I/O Read/Write cycles and timing diagrams for non-DMA ISA Masters are discussed In Section 5.4.2 on pp. 5-20 through 5-22.
- Table 9-7 through 9-11 and figures 9-3 through 9-8, on pp. 9-8 through 9-17. These tables and figures provide detailed ISA interface timing diagrams and associated parameters for ISA Masters, ISA 16-bit Slaves, ISA 8-bit Slaves, and ISA Master-to-PCI Accesses.

[a] Slaves, or targets, are devices that respond to transactions initiated by other devices.

[b] Masters are devices that initiate transactions to which other devices must respond.

REPORTS ON CD-ROM (CONT.)

ISA Bus Detailed Operation

- Tables 9-12 through 9-14 and figures 9-8 through 9-10, on pages 9-18 through 9-23. These tables and figures provide detailed DMA interface timing diagrams and associated parameters for DMA read, DMA write, and Type F DMA cycles.
- Table 9-15 and figure 9-11, on pages 9-24 and 9-25. This table and figure provide a detained X-Bus interface timing diagram and associated parameters for X-Bus operation.

From the ATX Reference Design for AMD-640 Chipset Technical Specification:

- Table 2-7, on pp. 29-30 maps the ISA Bus signals to the specific pin locations on the slot connectors.

Performance

Typical transfer-rates using the ISA Bus are well below the theoretical maximum of 8.33 MBytes per second.

Using fast 16-bit peripherals, the fastest ISA transaction transfers two bytes in two BCLK cycles. Thus, 8.33 MBytes per second is the theoretical maximum transfer rate for the ISA Bus. However, the default timing for 16-bit transfers is three BCLK cycles (5.55 MBytes per second). Furthermore, many peripherals are only 8-bits wide. The best timing for 8-bit transfers is three BCLK cycles (2.77 MBytes per second) and the default timing is six BLCK cycles (1.38 MBytes per second). Note that every datum transferred over the ISA Bus requires an associated address broadcast. The more efficient transfer of multiple data items for a single address is discussed in the section below on PCI.

The rate of continued increases in microarchitectural performance (in terms of SPECint92 benchmark results normalized by clock speed) is shown in Figure 4.10. The graph suggests that the rate of microarchitectural performance growth is slowing and is presently in a range of 10% growth per year. Overall future platform performance growth will thus likely come mostly from frequency and system architectural improvements, the latter including optimized connectivity and increased bandwidth to other processing and data storage resources.

Legacy Direct Memory Access (DMA)

DEFINITION

Direct Memory Access (DMA)

DMA controller

Direct Memory Access (DMA) uses dedicated DMA *controllers* to transfer blocks of data between memory and peripherals. Once initialized and initiated by the processor, DMA transfers are carried out by the DMA controller, *theoretically* freeing the processor for other activities. To use DMA, The processor may program the DMA controller with a start address and the number of bytes to transfer. The DMA controller will execute the transfer and notify the processor via interrupt when the operation is complete.

As originally implemented, Legacy DMA required the processor to relinquish the processor's local bus in order for the DMA controller to access main memory. This meant that concurrency between the DMA controller and the processor was severely limited. The processor could only proceed to the extent that it could execute instruction out of its prefetch buffers and internal cache. In a later section on Legacy I/O Issues, pointers are provided to materials that describe how Legacy DMA is implemented in current PCI Bus systems.

Suggested Readings

ISA Bus

The following texts are extensive references on the ISA Bus. The first is a very well organized and clear presentation of all basic concepts. The second is a very comprehensive discussion of low-level design issues.

1. Tom Shanley and Don Anderson, *ISA System Architecture*, 3rd Edition, Addison-Wesley, 1995.
2. Edward Solari, *ISA & EISA Theory and Operation*, Annabooks, 1993.

Elimination of the ISA Bus

Generally, the history of the PC platform has been one of absolute compatibility with legacy software and hardware. Yet, as we indicated in the Legacy Issues section of Chapter 1, the legacy ISA Bus is being eliminated from new mainstream platform designs. In this instance, compatibility has been subordinated to the ability to provide a system that is substantially higher performing, better behaved, and easier to support.

As we will see in the discussion of the PCI Bus and the AGP, the bus-width, bus-timings, and effective rate of transfer for the ISA Bus are but a small fraction of what is available with the newer buses. These new buses

were developed because the ISA Bus is totally insufficient to satisfy the needs of emerging 3D graphics acceleration and other high performance peripherals.

Equally important, the problems frequently caused by ISA cards are not limited to the card or the application using the card, but may adversely affect the entire system. ISA-based legacy peripherals may cause performance degradation to all threads of execution on today's highly multi-tasked systems. Additionally, the system as a whole may develop obscure and difficult to diagnose problems due to ISA cards that do not have Plug-n-Play automatic resource configuration or power-management capability.

The issue can be also be easily viewed from a support context. The relationship between the use of ISA cards and such systemic problems is neither obvious nor easily explainable to the naïve end-user. Furthermore, asking an end-user to remove an installed ISA card to help debug a problem, or asking him to forgo use of an ISA card that the user is presently dependent on, are not practical solutions.

While clearly a drastic move, eliminating ISA slots altogether does away with the aforementioned performance, poor-behavior, and support problems. The burden to the user of excising the ISA slots is mitigated by the integration of much of the common peripheral functionality into the Super I/O component of the South-Bridge. The burden will be further mitigated by the growth of peripherals making use of the USB and IEEE 1394 buses and the Device Bay initiative.

Backplane Bus — PCI

Through the 80's and early 90's, PC/AT memory subsystems were operated off the processor's increasingly faster local bus, while display adapter cards and other expansion peripherals were generally operated off the ISA/EISA Bus. To insure compatibility with earlier peripherals, the speed of the ISA Bus was left unchanged. Bus transceivers isolated the processor local bus and the ISA Bus, coupling the two only when necessary. This permitted certain programs or program segments that were not heavily I/O oriented to run much faster as processor and memory speeds increased. By the early 90's however, growing interest in multimedia applications, which are heavily I/O oriented, was creating additional performance demands beyond increased processor speed. Multimedia applications need an expansion bus with high effective data-transfer rates, support for burst transfers, support for concurrent subsystems, support for bus-mastering peripherals, and guaranteed low-latency access to the bus.

DEFINITIONS

Bus Features for Multimedia Applications

High effective data-transfer rates are essential for full-motion video due to the need to move high frame rate, large size, high-resolution color bitmaps, or related data and commands, to the display adapter. While the requirement is most acute for the display adapter, graphics-related data needs to be moved about between various other system components, including the processor, main memory, hard disks, and communications adapters.

Support for large numbers of data transfers per address, or long burst mode transfers, is key to high effective data-transfer rates. Burst mode transfers realize the highest data-transfer rates and make the most efficient use of the available bus bandwidth. Burst mode writes for transferring data into the video memory are an important performance feature, as the video memory holds both the frame buffer and command information for the display adapter.

Support for concurrent subsystems is inherent in the very nature of presenting the user with multiple simultaneous types of media. Furthermore, concurrent subsystems are necessary to provide multimedia augmentation without compromising the performance of the fundamental data processing and control flow requirements of the application program.

Support for bus mastering peripherals is key to concurrent subsystems. Bus mastering peripherals may autonomously carry out transfers, permitting the processor to carry out another task while the bus-master transfer is on going. Bus mastering peripherals are frequently, but not limited to, peripherals with bus-mastering DMA controllers. Bus mastering DMA is an important performance feature for display adapters, hard drives, and high-speed networking.

Guaranteed low-latency access is necessary because devices on the expansion bus generally cannot be held off the bus indefinitely or else noticeable visual or aural anomalies will be experienced by the user and data being recorded may be irretrievably lost.

Historical Comment

The VESA Local Bus (VL Bus)

In the early 90's, the display adapter was the peripheral subsystem where it was most apparent that peripheral speeds needed to be improved over those possible with ISA Bus. Interest arose is adding slots to the processor local bus for display adapter cards, and in 1992, the Video Electronics Standards Association (VESA) announced the VESA Local Bus (VL Bus, or VLB) standard. This formally defined slots for adding display adapter cards and other high-speed peripherals to a bus that was designed electrically to be a compatible derivative of the Intel 80486 (486) processor local bus, then used in most new PC platforms.

Definitions were included for a type "A" nonbuffered single VL Bus slot version and a type "B" buffered version for up to three slots. Physically the VL Bus standard called for two collinear connectors, one specific to the expansion bus (i.e., ISA, EISA, or Micro Channel) and another specific to the VL standard. This permitted VL Bus cards to optionally have separate connections to both buses. A typical 33-MHz PC platform might have one 8-bit ISA slots, five 16-bit ISA slots, and two VL Bus/ISA Bus hybrid slots for eight slots total. The hybrid slots could also be used for regular ISA cards.

A one slot VL Bus at 33-MHz was capable of an transfer rate of 106 MBytes per second for 16-byte burst reads[a]. However, this was only possible if the bytes were cacheable, aligned on a cache-line boundary, and the VL Bus card supported the 486 cache-line burst mode. Non-burst Reads and Writes to 32-bit devices can be performed at 66 MBytes per second. Since writing into the video memory is a dominant display-adapter use, the general transfer rate for display adapters on the VL Bus was likely closer to the non-burst transfer rate than the burst read rate.[b]

While the VL Bus provided needed higher bandwidth, it did not address the other needs of multimedia applications, as presented above. It also suffered from a number of other weaknesses. It placed significant loading on the processor local bus, due to their tight coupling. This meant that only two VL Bus slots were possible for 33-MHz machines, while 66-MHz machines were limited to one slot. Furthermore, the buffered multi-slot implementation required additional wait states over the single-slot version.[c] As a result, a multi-slot 33-MHz VL Bus was likely to have an transfer rate of no more than 44 MBytes per second for Non-burst Reads and Writes, or 88 MBytes per second for burst reads.

[a] 16-bytes in 5-cycles at 33MHz.

[b] Because data is being transferred to the display adapter only a fraction of the time, the overall throughput for a display adapter card is a corresponding fraction of the available transfer rate.

[c] Tom Shanley and Don Anderson, *PCI System Architecture,* 3rd Edition, Mindshare, Inc., 1995, p. 27.

HISTORICAL COMMENT (CONT.)

The VESA Local Bus (VL Bus)

Several factors unique to its time also seriously compromised the VL Bus's chances for long term survival. First, any VL Bus compatible product needed to be designed solely for use in one of the three popular expansion-buses then in use. Second, VL Bus products designed for the more prevalent 33-MHz machines were not necessarily compatible with the emerging 60 and 66-MHz machines. Finally, products designed for a 486-based VL Bus bus, were not directly compatible with a different processor local bus, such as that of the highly anticipated Pentium.

VESA has implemented a number of important standards for PC Platforms. The Section on Graphics and Video Adapters gives a brief discussion of some of these standards and a pointer to the VESA organization.

The PCI Bus has recently gained strong acceptance in several RISC architectures, enabling them to exploit the wide availability of relatively inexpensive high-performance peripherals designed for the higher-volume PC platform.

HISTORICAL COMMENT

Factors Leading to the Success of PCI

The PCI Bus was introduced in mid "93 and was rapidly adopted as the primary expansion bus for post-486 PC Platforms. Its success is attributable to a number of factors:

- Many PC buyers were keenly interested in multimedia applications and were willing to pay for new high-performance hardware,
- the ISA Bus was clearly inadequate for high performance multimedia and the higher performance of the then new Intel Pentium processor further highlighted the need for a new primary systems bus,[a]
- the roll-out of the PCI Bus coincided fairly closely with the roll-out of the 486-socket-incompatible Pentium and buyers needed a new motherboard anyway,
- PCI was an open architecture with wide industry backing, and
- PCI addressed the needs of multimedia applications better than did the VL Bus.

[a] The 64-bit processor local bus of a 66 MHz Pentium had a peak burst bandwidth of 528 MBytes per second and a transfer rate of 264 MBytes per second for non-burst reads and writes.

The PCI Connector

Figure 4.17 shows two views of the 32-bit PCI edge card connector. The right close-up view provides details of both the connector and the matching card edge of a PCI card. The card is on its side but is otherwise aligned with the connector, as it would be inserted. Figure 4.18 is a composite image[62] that provides a same scale comparison of slots for the ISA, VL, and PCI Buses, and AGP. The AGP is a variation of the PCI Bus that has been customized specifically for display adapter use. The AGP is discussed briefly later in this section and in the later section on Graphics and Video Adapters. A reference line representing the rear edge of a motherboard has been provided toward the right edge of the image. The slots are placed horizontally with respect to the reference line, in a manner that is representative of their relative alignment on a motherboard. A 16-bit ISA Bus slot is at the upper right with a collinear 32-bit VL Bus slot on the upper left. A 32-bit PCI Bus slot is in the lower middle. An AGP slot is at the bottom of the figure.

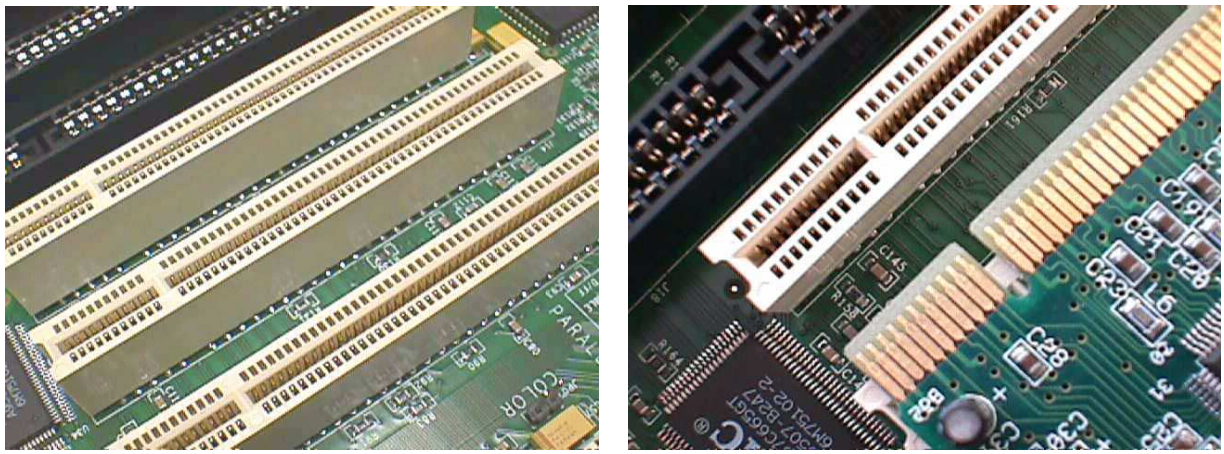

Figure 4.17 Views of PCI Card Slots and PCI Card Edge

[62] While theoretically possible, VL Bus slots normally are not present in systems with PCI.

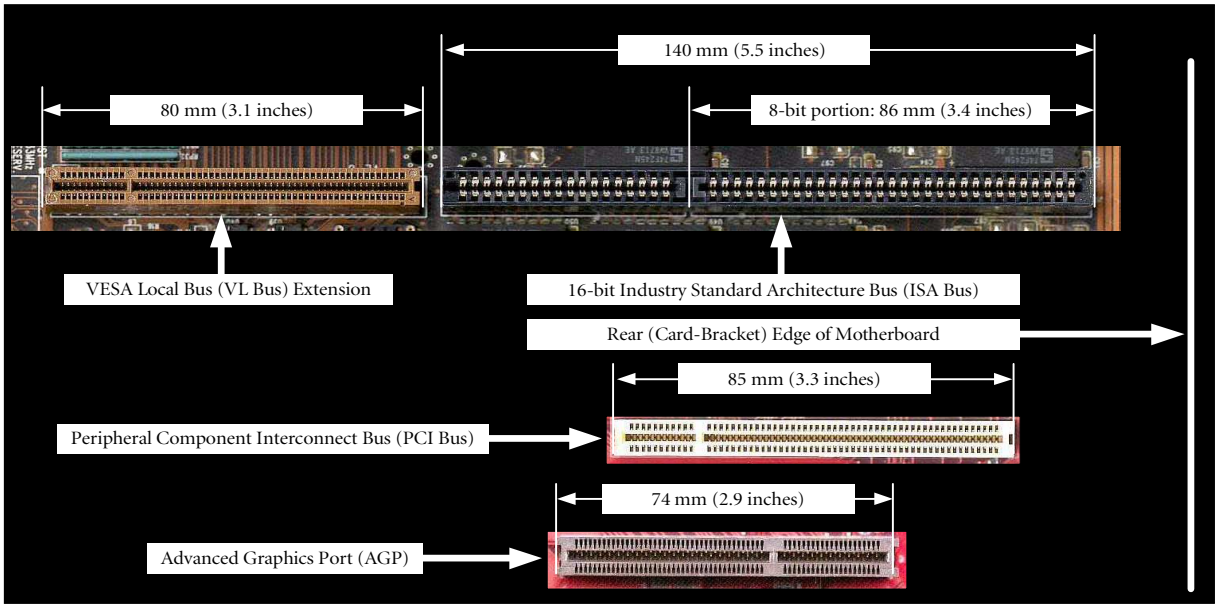

Figure 4.18 COMPARISON OF POPULAR CARD SLOTS

This is a same scale comparison. Horizontal placement is representive of relative alignment on a motherboard. Dimensions are approximate.

Key Features of the PCI Bus

The PCI Bus provides good support for multimedia applications via high-bandwidth, long burst transfers, well-done support for multiple bus-masters, guaranteed low-latency access. Additionally, PCI offers a number of other outstanding features including a processor independent bus architecture, scalable performance via bridged multiple buses, low power operation via Reflected-Wave Switching technology, and automatic configuration support for the Plug-n-Play initiative.

High-Bandwidth and Long Burst Transfers

A 32-bit PCI bus operating at 33MHz is capable of peak transfer rates of 132 MBytes per second. A 64-bit PCI bus at 33MHz, has a 264Mbyte per second peak transfer rate. Finally, a 64-bit PCI bus at 66MHz, has a 528 MByte per second peak transfer rate. Sustained rates are also more effectively realizable, at roughly 60-75%[63] of the peak rate, because the PCI bus is capable of long burst-mode transfers.

---

[63] Anthony Cataldo, "PC makers say Intel's 440LX chip set won't deliver sufficient performance—Compaq, Micron roll core logic for workstations," *EE Times TechWeb News*, August 18, 1997.

DEFINITIONS

Burst Cycles and Burst Transfers

Burst cycles, or burst-mode transfers, transfer multiple data items for a single address. Burst cycles may be viewed as being composed of a l*ead-off*, or initial, transfer and subsequent *burst transfers*. The address and first data item is sent during the lead-off. Burst transfers consist of contiguous data transfers sent without their associated address information. Generally, each data item sent in a burst transfer requires only a single bus clock cycle, while the lead-off transfer requires a larger number bus clock cycles. As a result of their compressed timing, burst transactions generally realize much greater effective data transfer rates than non-burst transactions and thus make more efficient use of available bus bandwidth.

For relatively short bursts, such as cache line fills, burst transfer performance is often described using notation of the form *W-X-Y-Z*. W is the number of bus clock cycles for the lead-off transfer, and X, Y, and Z are the number of bus clock cycles for the subsequent burst transfers.

lead-off
burst transfers

Processor Independent Bus Architecture

The PCI Bus is designed as a processor independent expansion bus. Neither the processor's local bus nor its speed directly affects the operation of the PCI Bus. A bus bridge[64] is customized for each local bus architecture to selectively couple the processor local bus with the PCI Bus for *crossing-transfers*. Unlike the simple transceivers used to isolate the processor local bus and the ISA Bus, this bridge has separate bus mastering controllers that implement the bus protocols unique to the processor local bus and the PCI Bus. The bridge may also be augmented with bi-directional (dual) first-in-first-out (FIFO) command and data buffering, which further reduces the coupling duration as perceived by the faster processor local bus.

crossing-transfers

Concurrent Subsystems, Bus Masters, and Guaranteed Low-Latency Access

The processor-to-PCI bus-bridge provides default isolation between the processor and peripherals on the PCI Bus. It also implements a PCI Bus arbiter, selectively granting the PCI bus to bus masters that request it. The PCI Bus permits any peripheral[65] to be a bus master, enabling peripherals to

[64] This is the North-Bridge of Figure 1.3.

[65] This is in contrast to the ISA architecture where a centralized DMA controller is generally the only bus master other than the processor.

readily carry out transfers on the isolated PCI Bus while the processor is carrying out other tasks on the processor local bus. PCI bus-masters are designed such that burst transfers that have exceeded a predetermined length can be momentarily suspended and subsequently resumed, should another bus master need to perform a transfer. This mechanism enables long burst transfers to attain high effective transfer rates yet ensures that every bus master will gain access to the bus within a desired interval.

Bus-Mastering DMA

Bus-Mastering DMA controllers reside on peripherals attached to the expansion bus and are dedicated to use by the associated peripheral. PCI Bus-mastering DMA is far superior to the centralized Legacy DMA mechanism. For bus-master DMA transfers between two PCI devices, the processor continues to have access to its external cache, to main memory, and to the AGP. This means that truly concurrent independent operation can take place by the Bus-master DMA controller and the processor. Examples include transfers between two hard disks, transfers between a hard disk and a network adapter, and transfers between a video capture device and either a hard disk or a display adapter. For bus-master DMA transfers between a PCI peripheral and main memory or AGP, the processor continues to have access to its external cache. While limiting, this is still greatly improved over having access only to the prefetch buffers and the internal cache as was the case with the original implementation of Legacy DMA. Because Bus-master DMA controllers are dedicated to a peripheral, the peripheral need not contend for scarce, shared DMA resources. Bus-mastering DMA controllers also generally have reduced microprocessor configuration requirements, because they can be custom-tailored to the needs of the associated peripheral and never need to be reconfigured for use with a different peripheral.

Scalable Performance via Bridged Multiple PCI Buses

Multiple instances of the PCI Bus may be bridged together via *PCI-to-PCI Bridges*. This provides scalable PCI performance. Standard PCI Bus implementations are generally limited to three or four slots. PCI-to-PCI Bridges permit servers and workstations to have larger numbers of expansion devices or greater bandwidth capacity. Since the available bandwidth of each bus must be shared among all slots, multiple bridged PCI buses permits greater flexibility in allocating the total available bandwidth. Specifically, a peripheral that has unusually high bandwidth requirements can be used in a lightly loaded bus, which is normally isolated from heavily loaded buses populated with multiple peripherals having lower bandwidth requirements. While generally the number of PCI Buses bridged together has been limited to a few, there are development efforts to bridge together as many as 256 PCI Buses.[66]

---

[66] Terry Costlow, "Startup stretches PCI bandwidth," *Electronic Engineering Times*, Nov. 10, 1997, p. 16.

Bridged Special Purpose PCI Variants

CardBus and AGP are two special cases of bridging different types of PCI Buses. Each of these is a variant of PCI for special applications. CardBus is an PCI Bus variation that is used in laptops as a second expansion bus dedicated to removable PC Cards. The CardBus is bridged to a standard PCI Bus for non-CardBus peripherals.

AGP is also a PCI Bus variation that is dedicated to a single high bandwidth peripheral, the display adapter. This permits the display adapter to have essentially the entire bandwidth of an enhanced PCI Bus for 3D graphics texture data traffic between main memory and the display adapter. This is traffic that in a sans AGP single-PCI platform would otherwise dilute the bandwidth available to other PCI peripherals. AGP is discussed again in the section on Graphics and Video Adapters.

Reflected-Wave Switching

Electrically, the PCI Bus employs a switching technique called *Reflected-Wave Switching*. The signal traces, which act as transmission lines at the frequencies involved, are carefully laid out and purposefully *unterminated* to cause significant *reflected* waves to occur along the traces. A carefully designed driver creates an *incident* wave of half the desired switching voltage. The incident and reflected waves add together to create the desired switching voltage, which is sampled by the destination devices on the clock cycle following the driving of the incident wave. In contrast, the ISA Bus and VL Bus used a traditional *Incident-Wave Switching* propagation technique, in which bus drivers drive the signal traces to the final switching voltage and the traces are terminated and laid out to minimize reflected waves. Reflected-Wave Switching reduces driver size and peak currents by half, for the same performance as Incident-Wave Switching.

Plug and Play

PCI Bus devices are required to have automatic configuration registers. These registers appear in a PCI configuration address space that is separate from the memory and I/O address spaces. The registers are programmed on system initialization at each power up, by a Plug and Play compliant System BIOS. The Plug and Play BIOS largely relieves the user or system administrator from the burden of avoiding conflicts in the assignment of system resources. The BIOS assigns each PCI card with the system resources it requires. Generally, such system resources would include a DMA channel, an interrupt request line, I/O addresses, and memory addresses. This also eliminates the need for manually changing jumpers or switches on the card. Plug and Play was discussed in the section on Plug and Play Configuration and Maintenance earlier in this chapter.

PCI Bus Protocol

Initiator
Target

The PCI Bus is a clocked bus using handshakes to pace the progress of each bus cycle. PCI uses the term *Initiator* for bus master devices and the term *Target* for slave devices. Initiator Ready (IRDY#) and Target Ready (TRDY#) are handshakes that if not asserted will introduce wait cycles. Unlike the ISA and VL-Bus, the PCI bus multiplexes its address and data on a common set of signal lines. This is a factor contributing to the small size of the PCI connector.

Reports on CD-ROM

PCI Bus Detailed Operation

Two data sheets included on the CD-ROM have valuable supplemental information on the PCI Bus.

From the AMD-640 System Controller Data Sheet:

- Section 4.2 on pp. 4-5 through 4-7 describe each of the PCI Bus signals.
- The introduction to Section 5.4 on pp. 5-30 and 5-31 gives an overview of the PCI Bus controller functions of the North-Bridge.
- Sections 5.4.1 through 5.4.7 on pp. 5-31 through 5-52 discusses bus cycles and timing diagrams for PCI Bus Initiator transactions, PCI Target transactions, arbitration, configuration, accesses by other Initiators, fast back-to-back cycles.
- Section 7 describes the Configuration Registers for the North-Bridge. Section 7.1 describes the PCI Configuration Mechanism used to access the Configuration Registers.
- The two data sheets discussed earlier in the ISA Bus Section also have valuable information on the PCI Bus.

From the ATX Reference Design for AMD-640 Chipset Technical Specification:

- Table 2-6, on pp. 26-28 maps the PCI Bus signals to the specific pin locations on the slot connectors.

PCI Card Types

There are a number of different configurations for PCI cards. Figure 4.19 is a matrix of PCI card edge connector silhouettes. A pair of 32-bit and 64-

bit card types exist for each of the 5V and 3.3V signaling voltages. In addition, a pair of 32-bit and 64-bit Universal card types exist that are compatible with either signaling voltage. Beneath each card silhouette are the slot types that the card is compatible with. The 5V PCI and 3.3V PCI cards have a *key* cut into the edge connector that corresponds to a keyway, or *key segment*, in their associated slot. This insures that a card can only be inserted into a bus that operates with the signaling voltage to which the card was designed. The 5V and 3.3V 32-bit slots (but not the 64-bit slots) are physically identical except for orientation. Universal cards are cards that have two keys in the edge connector in order that they will fit into either a 5V or a 3.3V slot.

A 64-bit PCI slot will accept either 64-bit or 32-bit cards having signaling voltage keying compatible with the slot. In a nicely done forward compatibility feature, a 32-bit slot will accept either 32-bit or 64-bit cards[67] with compatible signaling voltage keying. All 32-bit cards identify themselves as such during transfers, which are dynamically adjusted to match the card bit-width, regardless of the slot bit-width. All 64-bit cards in 32-bit slots configure themselves for 32-bit transfers.

It is not possible to mix 5V and 3.3V slots on a given PCI Bus segment. All cards on a given segment must use the same signaling voltage. Universal cards use I/O buffers that can be compliant with either signaling voltage, detect the type of slot into which they are plugged, and adapt their signaling voltage accordingly. The signaling voltage used is independent of the power-supply power rail voltages used by card components. 5V, 3.3V, and Universal cards all have connections defined to both 5V and 3.3V power rails, and are permitted to have any mix of components powered from the two rails.

Use of lower operating voltages provides advantages in energy savings and are generally required by present day higher density IC processes. The intent is for all PCI cards to eventually be 3.3V cards. While all this was known at the time the PCI specification was developed, it was also the case that the state of the industry was such that PCI chipsets were only going to be initially available with 5V signaling. The scheme of Figure 4.19 was devised to enable a migration from 5V technology to 3V technology.

All initial PCI cards and slots were designed for 5V signaling. Vendors are encouraged to use Universal cards for all new designs. While Universal cards are increasingly available, presently 5v-only cards are still the norm. Once Universal cards are widely available, it is envisioned that 3.3V slots will begin to replace 5V slots in new designs.

---

[67] 32-bit slots have associated "keep-out areas" to permit unobstructed insertion of 64-bit cards.

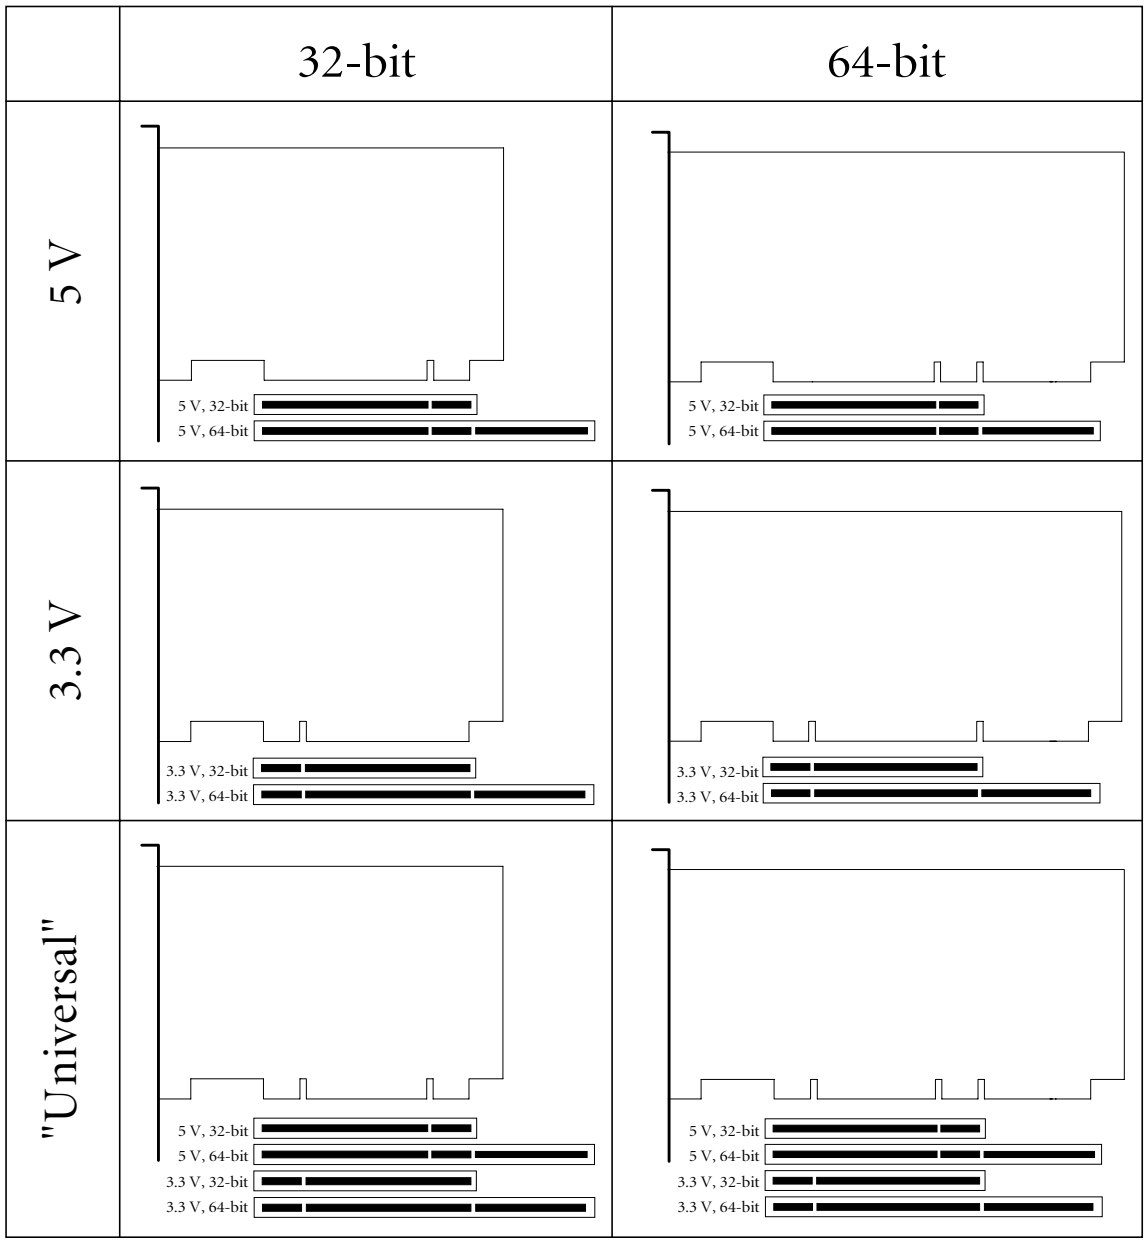

Figure 4.19 MATRIX OF PCI CARD TYPES AND ASSOCIATED COMPATIBLE SLOTS

The initial version of PCI was designed for operation at 33 MHz. Revision 2.1 of the PCI specification extended PCI buses with 3.3V signaling (but not 5V signaling) to optional operation at 66 MHz. If a 66 MHz PCI card is installed in a 33 MHz bus (theoretically including 66 MHz Universal cards in 5V slots), the device will configure itself to operate at 33 MHz. If a

33 MHz PCI card is installed in an otherwise 66 MHz bus, the bus will configure itself to operate at 33 MHz. 66 MHz PCI buses thus necessarily use 3.3V slots. One or two 66 MHz 3.3V slots are beginning to appear in high-end workstations and servers. These are augmenting existing 5V slots that are still needed for compatibility with the large installed base of 5V PCI cards. This is accomplished by using multiple bridged PCI bus segments, one segment for the 66MHz 3.3V slots and one or more segments for 33 MHz 5V slots.

Legacy I/O Issues

One thing that the PCI specification does not directly address is the handling of legacy I/O. Legacy software that assumes that the ISA Bus is present, expects to be able to use legacy locations in the DMA, Interrupt, I/O, and memory spaces. As these peripherals are integrated into the Super I/O module residing off the PCI Bus, and the ISA Bus eliminated, legacy references must be claimed by PCI devices. See the section on PCI-based Ports in this chapter for examples of PCI-based Legacy I/O implementations. The following documents describe in greater detail how Legacy I/O on the PCI Bus is done to ensure backward software compatibility.

Technical Presentation on CD-ROM

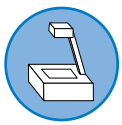

The following presentation overviews the migration of legacy peripherals from the ISA Bus to the PCI Bus and introduces the Distributed DMA and Serialized IRQ approaches to legacy I/O: Richard Wahler, *Preserving Legacy DMA's and IRQ's on the PCI Bus*.

Standards on CD-ROM

The full text of three industry standards for implementing legacy I/O on the PCI Bus are included in their entirety on the CD-ROM:

- *Common Architecture,* Desktop, PC/AT systems.
- *Distributed DMA Support for PCI Systems.*
- *Serialized IRQ Support for PCI Systems.*

Recent and Future PCI Developments

Hot Plug capabilities will permit PCI cards to be inserted and removed without incident while a system is running, and is directed mainly toward mission-critical servers, where downtime is costly. PCI power manage-

ment will permit PCI cards to power down when not in use and power up on demand. Vital Products Data (VPD) will permit the storage of manufacturer, warranty, and serial number data for use by system administrators and support personnel.

PCI is expected to be the primary expansion bus in desktops for some time. Nevertheless, in the next five years it may be supplanted in high-end Symmetric Multi-Processing (SMP) server clusters by multi-Gigabit per second serial buses currently under development.[68] One or more of the serial buses are proposed as links between processor, memory, and I/O and between multiple processors both in a local cluster and over significant distances.

Industry Standards

Peripheral Component Interconnect Special Interest Group (PCI SIG)

The PCI SIG describes itself as "an unincorporated association of members of the microcomputer industry set up for the purpose of monitoring and enhancing the development of the Peripheral Component Interconnect (PCI) architecture." Membership is $2,500 per year.

The PCI SIG maintains a Web site with information on joining, a mailing list, and how to order the specifications it controls. The main URL is:

http://www.pcisig.com/

PCI-related specifications are available only by mail from the PCI SIG. A bundled hard-copy and Acrobat copy of the main specification is sold for $25. Ordering information is found at URL:

http://www.pcisig.com/specs.html

Suggested Readings

PCI Bus

The following books are valuable as additional references for the PCI Bus. The first gives extensive low-level design details. The second gives a clear and well-organized explanation of the principal concepts.

1. Edward Solari and George Willse, *PCI Hardware and Software Architecture & Design*, 2nd Edition, Annabooks, 1995.
2. Tom Shanley & Don Anderson, *PCI System Architecture*, 3rd Edition, Addison-Wesley, 1995.

[68] Rick Boyd-Merritt, "Gigabit Bus Carries Intel Into Communications Territory," *EE Times*, March 9, 1998.

Requirements Summary for PC PCI Bus Devices

- General device requirements apply
- Devices and controller must comply with PCI 2.1 specification
- PCI chips sets implement general DMA controllers according to SFF 8020i, which defines the hardware and software requirements for ATAPI devices, while PCI-based ATA controllers implement Ultra ATA[a]
- Buses, devices, and functions define a three dimensional coordinate space in which multiple functions per card are OK, but *ghost cards* (single cards that appear at multiple coordinates) are not allowed
- PCI-to-PCI bridges that don't implement internal memory or I/O regions must proactively close (inhibit decoding to) any I/O or memory windows by configuring the *Base Address Registers (BARs)* according to the PCI to PCI Bridge Specification
- Controller requirements include that multifunction devices generally do not share writable configuration space bits, all expansion cards may be bus masters, and the Plug and Play write-data configuration-port-address must be propagated during power up and system reset to all ISA Buses that may contain Plug and Play cards (to enable Plug and Play configuration during the boot process)
- Plug and Play specific requirements include special configuration space register and configuration space requirements; supplemental subsystem device identifiers; interrupt routing must be ACPI compliant; BIOS must configures the boot device IRQ to a PCI-based IRQ; and device hot swapping must be compliant with an ACPI-based insert/remover notification mechanism
- All devices comply with PCI Bus Power Management Specification

ghost cards

Base Address Registers (BARs)

[a] See the *IDE/ATA* section in this chapter for a discussion of ATAPI, ATA and Ultra ATA.

Device-Specific Buses and Ports

The interconnects studied in this subsection generally have a minimum number of devices connected, typically just two devices or subsystems, and frequently implement device-specific protocols between the two devices. We will examine in turn the Socket 7 Processor bus, so-called Super I/O ports coupled to the PCI Bus, and graphics and video adapters. Our discussion of graphics and video adapters will include an examina-

tion of the Advanced Graphics Port (AGP) and its impact on the graphics subsystem.

PCI-based Ports

See the Legacy I/O Issues section for a discussion of how Legacy I/O should be implemented on the PCI Bus.

A number of ports are associated with controllers that are coupled to the PCI Bus. With reference to Figure 1.2 on page 36 in Chapter 1, a Super I/O block is shown coupled to the PCI Bus. The Super I/O typically has dedicated ports for legacy devices, including one or two PS/2 pointing-device/keyboard ports, floppy disk controller port, IEEE 1284 parallel port, legacy parallel ports, legacy serial ports, and ATA/IDE interfaces. A CardBus Bridge and a Legacy Bus (ISA or EISA) Bridge may also be coupled to the PCI Bus. Figure 1.2 on page 36 shows a South-Bridge containing all of these previously discussed components. Actual implementations may vary, with one or more of the CardBus Bridge, Super I/O, and Legacy Bus Bridges being implemented as individual chips.

Report on CD-ROM

To learn more about PCI-based Ports, see the following data sheets on the CD-ROM:

- *AMD-645 Peripheral Bus Controller Data Sheet*.
- *128 Pin Ultra I/O with ACPI Support and Infrared Remote Control*, FDC37B78x, Rev. December 8, 1997.

The PC Specifications allow legacy PS/2, serial (COM), and parallel (LPT) ports implemented on the PCI Bus. However, pointing devices and modems may not be connected using the legacy serial ports. No device except printers may be provided that are connected to the legacy serial and parallel ports. The preference is for the mouse (and keyboard) to be connected via USB, the modem via USB or Device Bay, and the printer via USB or IEEE 1394.

Requirements Summary for I/O Devices
• General device requirements apply • Required external devices include serial, parallel, keyboard, and pointing device ports • Multiple keyboards attached to mobile devices do not conflict with each other • game control devices and drivers must support the USB Human Interface Device Class Specification, Version 1.0 or higher • Legacy serial port specific requirements include that the mouse must not use a legacy serial port and any legacy serial port must use a 16550A Universal Asynchronous Receiver Transmitter (UART) and support up to 115.2 Kbaud[a] • Legacy parallel port specific requirements include that the Enhanced Parallel Port (EPP) must support flexible I/O addressing; ports must support IEEE 1284-1984 defined compatibility and nibble modes as well as Enhanced Capabilities Port (ECP) protocols; and Ports must, at a minimum, use IEEE 1284-I defined connectors

[a] A baud is a measure of signal changes per second. The baud rate is usually less than the effective bit per second rate due to the use of signal encoding techniques.

IEEE Standard on CD-ROM

The IEEE 1284 Standard covers asynchronous, fully interlocked, bidirectional parallel communications between hosts and printers or other peripherals. It is included in its entirety on the companion CD-ROM.

> 1284-1994 IEEE Standard Signaling Method for a Bidirectional Parallel Peripheral Interface for Personal Computers.

The IEEE 1284.1 Standard defines a vendor-independent standard protocol for the control of printers. It includes provisions for automatic identification and configuration of printers and messaging from the printer to a monitoring host computer.

> 1284.1-1997 IEEE Standard for Information Technology--Transport Independent Printer/System Interface (TIP/SI).

IDE/ATA

The AT Attachment (ATA) Interface and its variants (e.g., IDE, ATA-2, ATAPI, E-IDE, Fast ATA, all discussed below) are presently the most common method of interfacing storage media to the PC Platform. Integrated

Drive Electronics (IDE) is a tradename for ATA products made by Western Digital. The name was intended to emphasize that most of the drive's low-level electronics were integrated in the drive in contrast to earlier drive interfaces. The distinction between ATA and IDE is often lost, and IDE is more prevalent in the PC platform literature.

An ATA Interface is implemented between minimal control logic and two drives. The control logic resides on an expansion bus and operates under the control of the system's microprocessor. The interface standard includes definitions for a bus, a register set within the drive, a command set executed by the drive, and a physical connector and cable for coupling the drive to the controller. The bus is a variant of the ISA Bus. It is a 16-bit data bus using single-ended TTL compatible drivers. The bus is clockless. Seven bus control signals generated by the controller select which of roughly two dozen I/O registers are selected for read and write transfers of commands, data, status, and various drive-address attributes (drive, cylinder, head, and sector number). Roughly 40 commands may be given to the drive. A 40-pin connector and a 40-conductor non-shielded flat ribbon cable are used.

DEFINITION

Programmed I/O (PIO)

Programmed I/O, or *PIO*, refers to peripheral reads and writes carried out by the processor. This is a fundamental method of I/O transfer and is the preferred method for relatively small byte-count transfers. In a general sense it refers to either memory or I/O address space accesses. As applied to hard disks it usually literally refers only to I/O space accesses.

PIO is also often done for transfers of large data blocks between memory and a peripheral using the processor as an intermediary. PIO is generally performed via program loops. *String*[69] instructions may be used, if the peripherals are memory-mapped. As discussed in Chapter 6, the effective data transfer performance may be substantially reduced by overhead delay associated with the transfers. PIO Overhead delays take two forms. The first is the inner loop delay of the routine that actually performs the transfer. Thus, it is very important to programmed I/O block transfer rates for the program loop executing the transfer to be tight. The second overhead delay is the turnaround time between executions of the transfer routine. If interrupts and associated interrupt service routines, or other significant

[69] String instructions are complex instructions that transfer data byte sequences between source and destination memory locations. Instruction operands include the source and data locations and the number of bytes to transfer.

processor activity intervenes between executions of the transfer routine, the turnaround time can be large.

The ATA specification (as opposed to the later variants) includes the definition of three *Programmed I/O (PIO) modes and one Multiword DMA mode*.[70] In all modes, single word control and status transfers are common, but data transfers between the drive and the controller occur in *sectors*, which are groups of 512 bytes. Data transfers are normally 16-bits. Once a sector transfer is initiated the data transfer is not handshaken and proceeds at a rate determined solely by the host processor. In the PIO modes, the system microprocessor is responsible for accessing the drive's 16-bit wide data register to transfer data to or from the drive's sector buffer using I/O instructions. PIO Modes 0, 1, and 2, define data transfer cycles of 600, 383, and 240 ns, respectively. These correspond to peak transfer rates of 3.3, 5.2, and 8.3 MBps, respectively.

sectors

Transactions at a procedural-level are handshaken, in that the system's I/O routines are paced by status signals from the drive. The following is an example of the PIO protocol used on reads. The microprocessor initializes the drive by writing the appropriate parameters for the desired read. The microprocessor then sends a command to initiate the drive to process the read request. The drive asserts BSY (Busy) and begins reading the data and loading the drive's sector buffer. Once a sector is available in the drive's sector buffer, it asserts DRQ (Drive Request), lowers BSY, and asserts Interrupt Request (INTRQ) to interrupt the microprocessor. Once the sector is read by the host, INTRQ is deasserted, then DRQ is deasserted. An optional multiple sector mode permits the transfer of blocks of 2, 4, 8, or 16 (and possible greater) sectors before the drive interrupts the microprocessor. In order to minimize the overhead delay associated with sector transfers, it is clearly desirable to maximize the number of sectors transferred per interrupt.

In the Multiword DMA mode, a bus-mastering DMA controller in the ATA control logic is initialized by the microprocessor prior to initiating a multi-sector transfer command. The sectors are then transferred by the DMA controller and the processor is not interrupted until all sector transfers are completed. Multiword DMA Mode 0 defines a data transfer cycle of 480 ns, corresponding to a peak transfer rate of 4.2 MBps. The DMA mode theoretically allows the microprocessor to do other work instead of performing the data transfers and significantly reduces the interrupt servicing overhead for large sector count transfers.

ATA-2

Subsequent to the ATA standard, ATA-2 defines two additional PIO modes and two additional Multiword DMA modes. PIO Modes 3 and 4, define data transfer cycles of 180 and 120 ns, respectively. These correspond to

[70] Singleword DMA modes also exist, but are generally not used.

peak transfer rates of 11.1 and 16.6 MBps, respectively. In addition, these modes require that an IORDY handshake be implemented to permit the drive to pace the data transfers. Thus, PIO Mode 3 and 4 data transfers are optionally handshaken with passive termination. Multiword DMA Modes 1 and 2, define data transfer cycles of 150 and 120 ns, respectively. These correspond to peak transfer rates of 13.3 and 16.6 MBps, respectively. ATA-3 is and extension of ATA-2, but no additional PIO or DMA modes were added. ATA-2 and ATA-3 devices are marketed under the tradenames Enhanced-IDE (E-IDE) and Fast ATA. In the fastest modes, the maximum effective rate realizable in PC platforms was 7.8 MBps. The reduction was due to a confluence of a system intra-read latency and disk rotational latency effects that resulted in missed sector-read opportunities that effectively increased the data transfer cycle time. This became the limiting factor in ATA-based disk subsystem performance as the raw data rate capability of drives now exceeds 10MBps.[71]

Ultra ATA

Ultra ATA (also known as Ultra DMA/33) is an ATA extension that defines an additional multiword DMA mode. Multiword DMA Mode 3 defines an effective data transfer cycle of 60 ns, corresponding to a peak transfer rate of 33.3 MBps. This rate is temporarily more than raw drive data rates and it avoids the disk rotational effects that plagued the ATA-2 DMA modes. Multiword DMA Mode 3 actually uses the 120 ns cycle of Multiword DMA Mode 2, but redefined the signaling to transfer data on both edges of the data strobes instead of just the rising edge as on previous modes. In addition, the data strobe is sent by the data source instead of always by the drive controller as on previous modes. Such *source synchronous signaling* reduces clock skew between the timing reference (here the strobe) and the data and contributes to the reduced cycle time. Ultra ATA also adds a Cyclic Redundancy Check (CRC) to all transfers between the drive and drive controller to insure data integrity.

Use of ATA DMA Modes

From the advent of the PC/AT until just recently, the PIO modes have been exclusively used for ATA devices in PC platforms. The use of an ATA DMA mode requires a drive that supports it, compatible hardware, and software to manage the hardware and interact with the application programs that want to access the associated peripherals. The hardware necessary includes a compatible bus-master DMA controller in the ATA controller and a high-performance bus that supports bus-mastering DMA. The ISA Bus did not support bus-mastering DMA and the legacy DMA associated with the ISA Bus was limited to a peak transfer rate of

[71] *Ultra DMA, A Quantum White Paper*, Copyright 1998, Quantum Corporation.

2MBps. Thus, legacy DMA does not support ATA DMA mode 0 and later modes.

It has always been possible to use custom bus-master DMA hardware and a custom written application that directly controls the DMA hardware. However, to make the DMA modes transparently usable by all applications requires generally available hardware and support from the operating system and a device driver tailored to the ATA controller. Until the advent of Ultra DMA, the ATA Modes inherently did not offer any performance improvement over the PIO modes and there was no demand for either hardware or software support. Also, unless the OS is capable of multi-tasking and applications with multiple execution threads are used, there is no apparent performance benefit to the use of the DMA modes.

Ultra ATA is now the preferred mode to use on an ATA interface in new platforms. All the necessary components to both enable and motivate its use are present. The enabling components include available bus-mastering DMA chipsets, the PCI Bus, Microsoft Windows, and device drivers. The motivation comes from multimedia and other graphics intensive applications that have multiple real time threads of execution.

AT Attachment Packet Interface

The AT Attachment Packet Interface (ATAPI) is a compatible an extension to the ATA protocol. It defines a packet protocol for CD-ROM and Tape devices, which permits them to coexist on and ATA interface with traditional ATA devices.

Requirements Summary for IDE/ATA Devices

- Generic device requirements apply
- Controller is ATA-2 compliant
- Bootable controllers must support the El Torito standard for the CD-ROM installation process
- Any sharing of hardware between the channels of a Dual channel IDE adapter is transparent to all system software
- System BIOS and devices support the Logical Block Address (LBA) standard for accessing high capacity drives
- PCI Bus Mastering DMA requirements include that the DMA's register set and the drives must comply with SFF 8038i and must support Ultra ATA
- Controller and peripherals must use keyed and shrouded connectors and cables must have pin 1 designated, to avoid incorrect cabling
- Devices support the ATA STANDBY command to spin up devices after being placed in a power saving state

Graphics and Video Adapters

Graphics adapters and Advanced Graphics Port (AGP) have been covered fairly extensively in other parts of the text. Specifically, see the Display Adapters section in Chapter 1 and *Performance Optimization Specifics in a Contemporary 3D Graphics Platform* in Chapter 6. As covered in both of the cited sections, AGP is a system topology wherein the memory controller (within the North-Bridge) incorporates an additional port (the AGP) in the group of ports that are arbitrating for main memory access. The AGP is used for a direct high-bandwidth connection between the memory controller and the graphics/video adapter. The AGP supports higher throughputs between the graphics/video adapter and main memory than possible over the PCI Bus, while simultaneously eliminating major traffic on the PCI Bus, thus increasing its bandwidth availability to other platform devices. To some extent, AGP reduces the amount of video memory that would otherwise be required, as the fast connection to main memory can be thought of as effectively increasing the available video memory. Its primary motivation however, is to enable 3D rendering processes to exploit a ready availability of a wealth of textures that can be preloaded and manipulated in main memory for on-demand transfer to the 3D graphics accelerator, while keeping the texture traffic between memory and the display adapter off of the PCI Bus.

There are currently three defined AGP modes: AGP (1-mode), 2 AGP (also 2-mode or 2x mode), and 4 AGP (4x mode). 2 AGP increases transfer rates to 533 MBps, 4AGP offers 1GB/s transfers.

Figure 4.21 on page 409 illustrates the 124-pin single-signaling-voltage AGP connector. The AGP card edge is seen close up in the right of the figure. Depending on the orientation of the key segment, the connector is used for AGP cards with 3.3V signaling or for AGP cards with 1.5V signaling. A 132-pin universal voltage AGP connector also exists that will accept AGP cards designed for either 1.5V or 3.3V signaling. AGP card types and their associated compatible connectors are illustrated in Figure 4.20.

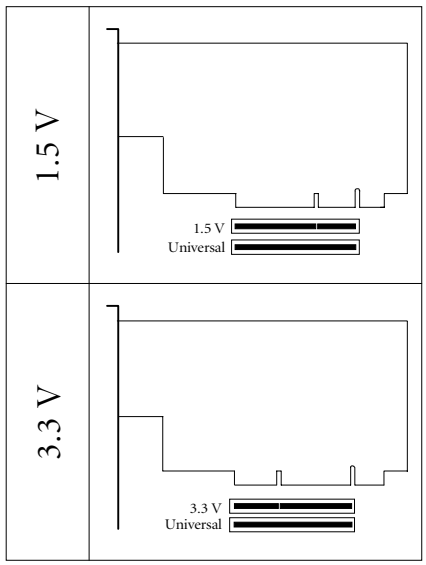

Figure 4.20 AGP Card Types and Associated Compatible Connectors

Report on CD-ROM

To learn more about AGP, see the article by Yong Yao, "AGP Speeds 3D Graphics," *Microprocessor Report*, Vol. 10, No. 8, June 17, 1996, on the CD-ROM.

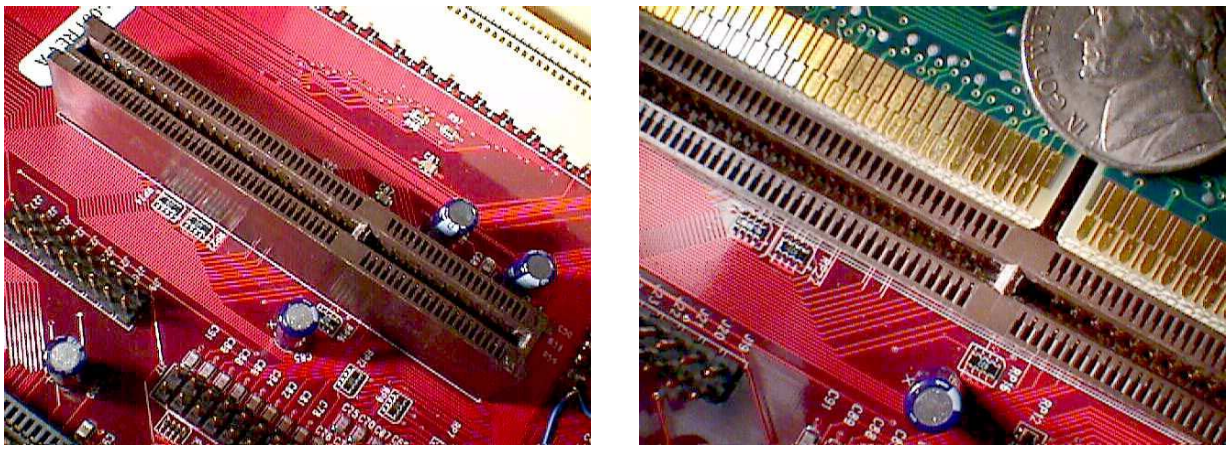

Figure 4.21 Views of AGP Slot and AGP Card Edge

Requirements Summary for Graphics Adapter Devices

- Adapter uses high-speed bus
- System exploits processor *write combining* (dynamic opportunistic batching by the processor of multiple individual byte writes within a single larger write), when available
- Primary adapter works with Window's default VGA mode driver
- Adapter and driver support multiple adapters and multiple monitors
- Adapter supports all VESA standard timings for refresh and for resolutions up to 1024 x 768 x 16 bpp
- All supported color depths are enumerated by the driver and its associated INF file
- Only *relocatable* (arbitrarily reassignable to new locations) registers should be used for dynamic graphics operations
- Multiple-adapter/multiple-monitor use requires all extended (non-VGA) resources to be dynamically relocatable after system boot and VGA resources to be capable of being disabled
- Adapter supports downloadable RAMDAC entries, which are used for transition effects, palette switches, and gamma correction for color matching between the display and output devices
- Adapter supports Display Data Channel Standard (DDC), Version 2.0, in support of software management of the display
- Additional requirements apply to video playback, 2D and 3D hardware acceleration, television output, BIOS and option ROMs, AGP, PCI-based adapters, and device drivers and installation

INDUSTRY STANDARDS

Video Electronics Standards Association (VESA)

VESA is an organization that supports and sets industry-wide interface standards for the PC, workstation, and computing environments. VESA promotes and develops timely, relevant, open standards for the video electronics industry, ensuring interoperability and encouraging innovation and market growth. Annual membership dues run $500 to $5000, depending on company structure and revenue. The main URL is:

http://www.vesa.org/

VESA Monitor Standards

VESA standards are only available to VESA members. Some of the more important video standards include:

- *Discrete Monitor Timings (DMT) Standard 1.0*: Timing standards for display monitors for resolutions from 640 x 350 to 1280 x 1024 with refresh rates from 60 Hz up to 85 Hz.
- *VESA Display Data Channel (DDC) ™ Standard, Version 3*: This standard defines a communications channel between a computer display and a host system for configuration and display control information and for uses as a data channel for peripherals connected to the host via the display.

A complete list of VESA standards with a brief abstract of each is at URL:

http://www.vesa.org/ve00013.html

Processor Bus — Socket 7

Another motherboard configuration, seen in earlier 486 platforms, is to provide an empty socket for an "upgrade" processor, obviating the need to remove the original processor. Intel calls its upgrade processors OverDrive Processors.

Desktop motherboards for Pentium-class and higher performance x86 processors are generally manufactured and originally sold with empty ("unpopulated") sockets for the processor and DRAM main memory. Socketing, while expensive and bulky, provides great flexibility for both manufacturers, system vendors, technicians, and end users. Sockets permit the same motherboard to be used in a range of system performance determined largely by the speed of the chips chosen to populate the sockets. Socketing facilitates system diagnosis and repair and decouples the motherboard manufacturer and vendor from the volatile and generally rapidly decreasing prices for processors and DRAM. Sockets are the only practical way to enable a consumer end user to replace large pinout chips as part of a processor upgrade.

Pin-Grid-Array (PGA)

Socketed x86 processors use plastic and ceramic *Pin-Grid-Array_(PGA)* packages with hundreds of pins, suitable for through-hole board permanent mounting. PGAs, turned pins up, are reminiscent of a "bed-of-nails." These flattish, generally square, packages have a processor die in the center cordoned by regularly spaced rows of nested pins (perpendicular to the package) arranged in horizontal and vertical (grid-like) pattern. The grid can be standard or interstitial. In a *standard grid*, all pin rows/columns are collinear. In an *interstitial grid*, alternating pin rows/columns are collinear, while adjacent pin rows/columns are offset by half the pin-spacing. A PGA using an interstitial grid is also referred to as a *Staggered PGA (SPGA)*. Figure 4.22 shows a 321-pin SPGA.

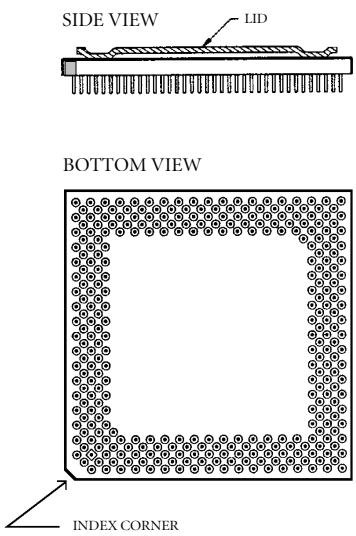

standard grid

interstitial grid

Figure 4.22 321-PIN SPGA PACKAGE

Adapted with permission of Advanced Micro Devices Inc., from *AMD-K6 Processor Data Sheet*, Copyright 1997.

Staggered PGA (SPGA)

Due to the force multiplication from hundreds of pins, the flexibility of socketing the processor is best realized by the use of specially designed and relatively expensive *Zero-Insertion-Force (ZIF) sockets*. ZIF sockets generally have a lever (handle), which must be manipulated to socket the processor. When the lever is moved to the "off" (up or open) position, there is no force constraining the processor pins, and easy insertion or replacement is possible. When the lever is moved to the "on" (down or closed) position, reliable physical and electrical contact is achieved, and the system is ready for normal operation. Figure 4.23 shows an empty ZIF socket with the lever in the off position

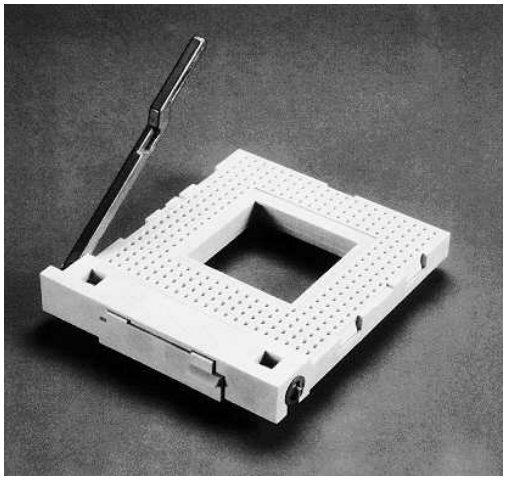

Figure 4.23 AMP ZIF

Reprinted with the permission of AMP.

Industry-Standard Sockets

Due to their popularity, x86 sockets and their associated pin assignments (pinouts) and electrical buses become de facto industry standards. The moniker "Socket 7" (or 4, or 5) has come to more generally signify the collective of physical and electrical interfaces associated with the socket. Socket 4, Socket 5, and Socket 7 are three related Pentium-bus standards associated with three notable classes of Pentium processors.[72] These three sockets are illustrated in Figure 4.24.

[72] Socket 6 defines a 3V 486DX4.

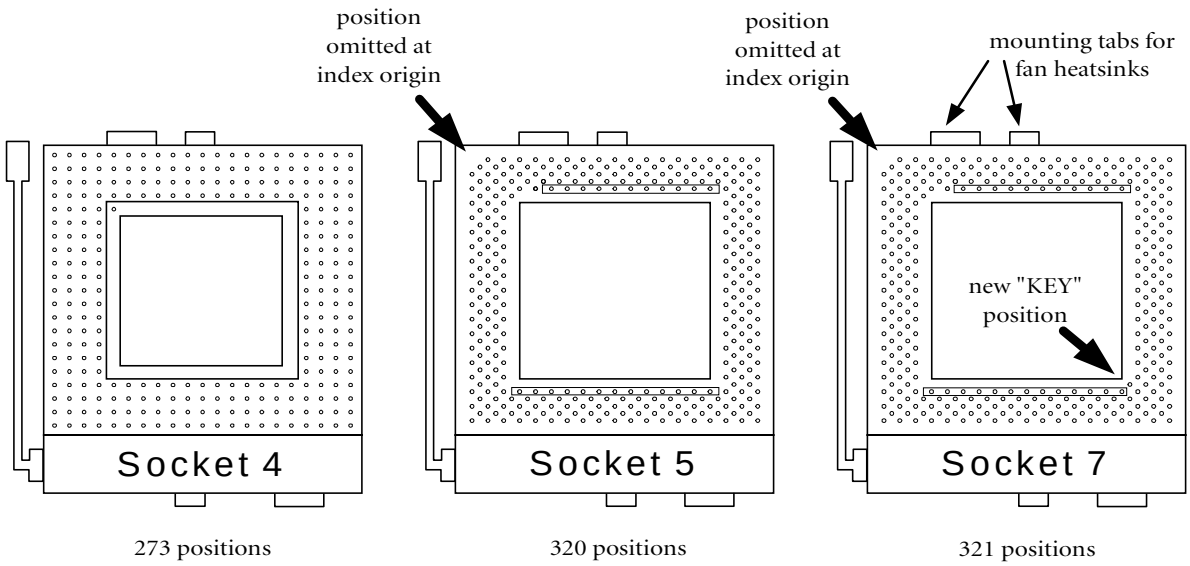

Figure 4.24 SAME SCALE COMPARISON OF SOCKETS 4, 5 AND 7

DESIGN NOTE

Distinguishing Socket 4, 5 and 7 Devices

You can tell apart the major Pentium-bus packaged device pinouts at a glance by examining the nested rows of pins on the underside of the package. 273-pin Socket 4 devices align pins on a regular grid. Socket 5 (296-pin and 320-pin) and Socket 7 ("296-pin" and 321-pin) devices align pins on an "interstitial" grid that staggers alternate rows. A 296-pin device has 4 nested pin rows on two opposite sides, and 5 nested pin rows on the other two opposite sides. 320-pin and 321-pin devices have 5 nested pin rows on all four sides. Along the major diagonal opposite the package's index corner (the corner with the 45-degree chamfer, establishing the origin that pin assignments are referenced from), a Socket 7 device has six pins (includes the electrically not-connected KEY pin), compared with five pins (no KEY pin) for a Socket 5 device.

While thinking of these Pentium-bus Sockets as industry standards is a useful abstraction, there have been numerous variants within each socket type. Manufacturers have used a wealth of combinations, including variations in bus frequency, core frequency, minimum timing and frequency

spreads, unified versus split core and I/O power supplies, core operating voltages and ranges, boundary-scan support, and multi-processor (MP) support. Not all combinations are offered in each socket type or by all manufacturers. Beyond the bus electrical and functional variations discussed in the main text, manufacturers may market further variants based on chip-package type, processor revision (tooling and process stepping), level of testing, OEM shipping trays versus consumer cartoning (so-called *boxed processors* often bundled with fan/heatsinks), and instruction-set support.

Socket 4

Socket 4 was used for Intel's CPU Family 5, Model 1, also designated by the part number substring 80501, and widely known by the (once internal) Intel code name *P5*. Intel has marketed P5 processors with core frequencies of 60 and 66 MHz. The Pentium Overdrive Processor at a core frequency of 133 MHz was designed as a Socket 4 upgrade for the 60-and 66-MHz P5 processors. Socket 4 uses a 273-pin-position ZIF. Socket 4 specifies a (non-staggered) PGA having pins arranged on a 21 x 21 grid with 0.10-inch standard grid spacing. The power supply pins for the Socket 4 bus are defined for a single supply, of roughly 5V. Different P5-compatible parts have used a variety of voltage ranges including 4.75-5.25V, 4.90-5.25V, 4.90-5.40V, and 5.15-5.40V. Pentium Overdrive Processors at core frequencies of 63, 83, and 133 MHz were designed as Socket 4 upgrades for the 60-and 66-MHz P5 processors.

Socket 5

Socket 5 was used for Intel's CPU Family 5, Model 2, also designated by the part number substring 80502, and widely known by the Intel code name *P54C*. Intel has marketed P54C processors with core frequencies of 75, 90, 100, 120, 133, 150, 166, and 200 MHz. Strictly speaking, P54C describes only the Socket 5, 0.5-micron BiCMOS, Pentium processors (at core frequencies of 75, 90, and 100 MHz), while P54CS describes Socket 5, 0.35-micron BiCMOS, Pentium processors (those with core frequencies of 120 MHz and above). The additional "S" reportedly being a mnemonic for a process "shrink." However, the distinction is not generally made, as the term P54C is loosely used for the P54CS parts as well. There are no substantive bus changes between the P54C and P54CS. Pentium Overdrive processors, Intel code named the *P54CT*, at core frequencies of 125, 150, and 166 MHz were designed as Socket 5 upgrades for the 75-,90-,and 100-MHz P54C processors. Socket 7 later became the recommended upgrade socket for new platforms designed around the same processors.

Socket 5 uses a substantially different layout than Socket 4. Socket 5 uses a 320-pin-position ZIF. Socket 5 specifies an SPGA having pins arranged on a 19 x 19 grid with a 0.05 x 0.10-inch interstitial grid spacing. Intel's P54C processors for new installations used 296 of Socket 5's 320 pin

positions. As a result, Socket 5 is often misdescribed as having 296 pins. However, Intel's Pentium OverDrive processors use all of Socket 5's 320 pin-positions.

The new Socket 5 pins are used in Socket 7, as well, and are individually discussed in detail later in the main text.

Most of the forty-seven new pins on Socket 5 were provided for future expansion options and unassigned in the P54C. Six of the new pins are for Dual Processor support, one for Functional Redundancy Checking, and three for an Advanced Programmable Interrupt Controller module (APIC). (Two of Socket 4's interrupt-related signals are dynamically redefined, when a Dual Processor system is present.) One new pin is used to put the processor in a "stop clock" low-power mode that is cache coherency aware. One Socket 5 pin was also defined to program a clock multiplication ratio, setting the ratio of the frequencies of operation for the processor core and the external bus. Four instruction-tracing-related signals defined for Socket 4 were deleted in Socket 5.

Socket 5 processors are marketed in a number of power supply variations and designations. Socket 5 processors marketed primarily for the desktop are unified-plane processors. *Unified-plane processors* have a single positive power supply to which both the core and the I/O circuitry are commonly connected. Unified-plane Socket 5 processors with an Intel *Standard (STD) operating voltage* designation have a 3.3V nominal voltage, but are guaranteed to operate over a range from 3.135 to 3.6V (3.3V -5%, 3.5V +100mV). Unified-plane Socket 5 processors with an Intel *Voltage Regulated (VR) operating voltage* designation have a reduced voltage specification of 3.300-3.465V (3.3V -0%, +5%). Unified-plane Socket 5 processors with an Intel *Voltage Regulated Extension (VRE) operating voltage* designation have a reduced and shifted voltage specification of 3.40-3.60V (3.5V -/+ 100mV).[73] Intel was able to produce profitable yields and increased clock speeds of VR and VRE parts in advance of STD parts, permitting OEMs to get systems to market earlier and at less cost than would otherwise have been possible. All Socket 5 processors have operating voltages reduced from 5V, but are required to have 5V-tolerant clock inputs.

Split-plane processors have separate core and I/O positive power supplies. These generally permit compatibility with standard 3.3V I/O chip-sets, while ensuring reliable operation for submicron core circuits and permitting reduced power consumption. The power dissipation of the core, which dominates overall chip power dissipation, is proportional to the square of the core operating voltage, so the payback from reducing the core operating voltage is dramatic. Care should be taken that split-plane Socket 5 parts not be used in motherboards configured for a single power

[73] The 3.135-3.6V range for STD, and the 3.40-3.60V range for VRE, are true for Intel's C2-step and later Pentium processors. B-step parts had a STD range of 3.135-3.465V (3.3V -/+ 5%), and a VRE range of 3.45-3.60V(3.5V-50mV,+100mV).

supply, as such operation overstresses the core, generally reducing reliability and life expectancy.

The positive power supply pins in unified-plane Socket 5 processors are labeled with V_{CC}. In split-plane Socket 5 and Socket 7 processors, the core pins are labeled V_{CC2}, while the I/O pins are labeled V_{CC3}. Socket 5 assigns the core (V_{CC2}) pins and I/O (V_{CC3}) pins generally on opposite halves of the socket. The "unified" and "split" terminology arises because the common (unified) V_{CC} power plane used on motherboards for single-supply processors is (fairly neatly) partitioned (split) into two power planes for the dual-supply processors. This is apparent in Figure 4.25, which illustrates the separate power planes for V_{CC2} and V_{CC3}, and associated decoupling capacitors, in an under-the-socket layout for a Socket 5 processor printed circuit board. Socket 5 processors have no means to identify to the motherboard whether they are unified or split-plane processors.

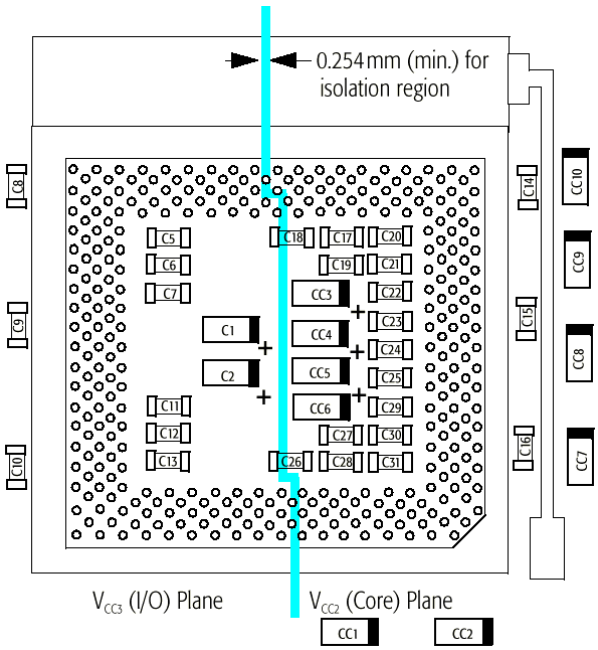

Figure 4.25 Split Power Planes for Dual-Supply Processors

Split-plane Socket 5 processors were initially marketed primarily as mobile processors, where power conservation is especially critical. Split-plane processors with an Intel *Voltage Reduction Technology (VRT) operating voltage* designation have 3.3V I/O and 2.9V core nominal voltages. Split-plane processors with an Intel *3.1V operating voltage* designation are VRT processors with 3.3V I/O and 3.1V core nominal voltages. Both classes of

VRT processors have 25 V_{CC2} pins, 28 V_{CC3} pins, 53 V_{SS} pins, and 296 pins total.

In addition to the split-plane technique, *mobile* Socket 5 processors generally have other significant variations from the *desktop* Socket 5 parts. These mobile parts generally do not support the Upgrade, Dual Processing (DP), Functional Redundancy Checking (FRC), and the Advanced Programmable Interrupt Controller (APIC) features, and do not support two optional signals included primarily for L2 caches.

OverDrive processors also require a +5V supply for their required fan heatsink.

Intel targeted its Pentium OverDrive processors to be consumer upgrade products. Other than the 200Mhz part, these parts are compatible with a Socket 5 pin out with a single 3.3V power supply, typical of most older Pentium desktop motherboards. The 200Mhz OverDrive processor requires Socket 7. The Socket 5 OverDrive parts work just as well in the Socket 7 upgrade socket found in newer boards. The Intel Pentium OverDrive processors at 125, 150, and 166 MHz are unified-plane processors with STD operating voltages. Pentium OverDrive processors with MMX Technology (Intel code-named the *P54CTB*) at 125, 150, 166, and 180 MHz are internally split-plane designs that work in unified-plane motherboards, because of an *internal* voltage regulator powering the core from the external unified supply.

Unlike other Socket 5 processors with 296 pins, the 320-pin Pentium OverDrive processors have twenty-four pins in positions near the center of the Socket 5 SPGA grid. Four of these extra pins are No Connects (NC). Twenty of the extra pins are used for extra power supply pins. Additionally, the Pentium OverDrive uses two power supply pins that are defined as No Internal Connection (NIC) pins on the 54C. In total there are sixty V_{CC} and sixty-eight V_{SS} pins. This is twenty-two power supply pins more than the fifty-three V_{CC} and fifty-three V_{SS} pins of the P54C.

Socket 7

Intel's P55C processors used only 296 of Socket 7's 321 pin positions. As a result, Socket 7 is often misdescribed as having 296 pins. The AMD-K6 uses all of Socket 7's 321 pin positions.

Socket 7 uses a substantially similar layout to Socket 5. Socket 7 also specifies an SPGA having pins arranged on a 19 x 19 grid with a 0.05 x 0.10-inch interstitial grid spacing. However, Socket 7 uses a 321-pin-position ZIF. Socket 7 standardizes for the first time the placement of the ZIF's integral clips (tabs) for the mounting of fan heatsinks. Designed as a superset of the various Socket 5 variations, Socket 7 is backward compatible with most Socket 5 processors.

Socket 7 adds one new pin, KEY, required of all Socket 7 processors. The KEY pin prevents Socket 7 devices from being inserted into a Socket 5 socket. Yet, Socket 5 processors may be inserted in Socket 7's socket without difficulty.

Socket 7 defines all positive power-pins using split-plane definitions in a manner that compatibly merges the power supply definitions used for the 296-pin split-plane Socket 5 VRT processors and the 320-pin unified-plane Socket 5 OverDrive processors. Specifically, in 296-pin Socket 7

implementations, the definitions for the core (V_{CC2}) and I/O (V_{CC3}) pins previously used in the 296-pin split-plane Socket 5 VRT processors are adopted. As discussed earlier, such processors have 25 V_{CC2} pins, 28 V_{CC3} pins, and 53 V_{SS} pins.

Although optionally implemented, Socket 7 also defines the twenty-two extra power pins used previously in the 320-pin Socket 5 OverDrive processors, but with new distinction made for which are for the core (V_{CC2}) and which are for I/O (V_{CC3}). The new definitions continue the assignment of core pins and I/O pins generally on opposite halves of the socket. A full 321-pin implementation of Socket 7 has 28 V_{CC2} pins, 32 V_{CC3} pins, and 68 V_{SS} pins.

Compatible with standard chipset I/O, V_{CC3} is 3.3V. V_{CC2} in the initial Socket 7 processors was roughly 0.5V below V_{CC3}, but V_{CC2} is being reduced with each processor version implemented in smaller geometry processes. The 0.25-micron K6 has a nominal V_{CC2} of 2.2V. Socket 7 revises the specification for the clock inputs (CLK and PICCLK), requiring them to be driven only by 3.3V clock drivers.

Design Note

VCC2DET

In both the K6 and the Pentium processor with MMX Technology (a.k.a. P55C), VCC2DET is always logic 0.

The most significant change over Socket 5 is that Socket 7 standardized a means for optional automatic detection of unified-plane versus split-plane processors and configuration of motherboard supplies for the core supply and I/O supply pins. This enables motherboards designed with Socket 7 to be configurable to accept either unified-plane (e.g., P54C or current Pentium OverDrive processors) or split-plane processors. This was accomplished by redefining a previous Socket 5 No-Connect (NC) pin as a V_{CC2} Detect (VCC2DET) output pin that is to be externally tied to a pull-up resistor and to the control input of a selectable-output voltage regulator. VCC2DET indicates whether the core and I/O supplies are to be unified (logic level 1) or split (logic level 0).

Flexible motherboards and VRMs are discussed in more detail in Intel's AP-579 (order 243187-001) Pentium Processor Flexible Motherboard Design Guidelines.

A Socket 7 motherboard capable of being configured for either split-plane or unified-plane processors is called a *flexible motherboard*. Flexible motherboards may use VCC2DET with automatic configuration, may be configured at manufacturing build or assembly time (by installation of particular voltage regulator components, including jumpers), may use a *Voltage Regulator Module (VRM)*, which is a modular DC-to-DC converter having an integral *header* connector, or combinations of the above.

If a Socket 7 device sinks VCC2DET low, a flexible motherboard with VCC2DET capability can automatically configure the supplies for a split-

plane processor. If a device lets VCC2DET go high, as all pre-Socket 7 devices will, the supplies are configured for a unified-plane processor. Desktop Socket 5 processors, the majority of Socket 5 processors, will have their supplies properly configured when used in a motherboard having VCC2DET and automatic configuration capability. However, the "mobile" split-plane Socket 5 devices will not.

All new-installation Socket 7 processors to date have been split-plane processors. Socket 7 is *required* for the AMD-K6 MMX Enhanced Processor. AMD has marketed AMD-K6 processors at bus frequencies of 60, 66, 83, and 100 MHz and core frequencies of 166, 200, 233, 250, 266, and 300 MHz. Socket 7 is also *required* for Intel's CPU Family 5, Model 4, also designated by the part number substring 80504, and widely known by the Intel code name *P55C*. Intel has marketed P55C Pentium processors with MMX technology at core frequencies of 166, 200, and 233 MHz.

Newer platforms designed around Pentium processors (P54C) at 75, 90, 100, and 120 MHz were *recommended* to use Socket 7 as the upgrade socket. Pentium processors (P54C) at 133, 150, 166, and 200 MHz were *required* to use Socket 7, in order to accept their intended Pentium OverDrive processors with MMX technology.

The split-plane AMD-K6 MMX Enhanced Processor has a nominal I/O voltage of 3.3V, with a range of 3.135-3.465V (3.3 V -/+ 5%), and a nominal core voltage of 2.9V, with a range of 2.755-3.045V (2.9V -/+ 5%). Split-plane Intel desktop and mobile Pentium processors with MMX (P55C) have a nominal I/O voltage of 3.3V, with a range of 3.135-3.60V (3.3V -5%, 3.5V +100mV). The desktop processors have a nominal core voltage of 2.8V, with a range of 2.7-2.9V (2.8V -/+ 100mV). The mobile processors have a nominal core voltage of 2.45V, with a range of 2.285-2.665V (2.45V-165mV, +215mV).

Performance

The Socket 7 bus is designed for high performance and supports a peak throughput 8 bytes every cycle. Thus at speeds of 66, 75, 83.3, and 100 MHz, the bus offers a peak throughput of 528, 600, 664, and 800 MBps, respectively. These peak rates are only approached when the system can exploit the bus cycle pipeline of the bus and burst cache line fill features.

fractional speed bus

The Socket 7 bus is known as a *fractional speed bus*. This means it supports core clock multiplication controls, such that the core can operate at fractional multiples greater than the bus speed. This enables the processor core to execute significantly faster than the bus, while keeping the motherboard low cost and easy to design due to its slower speed.

The transfers on the Socket 7 data bus consist of 64-bit (8-byte) wide objects. Because the 8086 used 16-bit registers, the x86 architecture has historically referred to a 16-bit object as a "word" and a 32-bit object as a "doubleword." Thus, a 64-bit object transferred on the Socket 7 bus is called a "quadword."

Basic Bus Operation

A dedicated 64-bit wide data bus is organized as 8 selectively enabled, byte wide data buses, for optimal use with 8-byte wide memory systems. Separate 32-bit addressing is provided by 8 byte-enables and 29 address-bits (A31..A3). The 8 byte-enables provide the same address range as 3-bits of address, but with the capability of selecting multiple arbitrary bytes within an 8-byte group. Cache line block fills of 32-bytes are supported, which are filled using 4, 8-byte bursts, that only require a single start address.

INDUSTRY STANDARD

Big and Little Endian Byte and Bit Ordering

In multiple-byte data objects, the relative significance of individual bytes that compose the object is an arbitrary characteristic of the processor's architecture. The bit-ordering within bytes is generally consistent with the byte-ordering convention. The two standard conventions use monotonic orderings that are known as the Little and Big Endian orders.

The K6 and other X86 processors use Little Endian ordering, in which the address of the object corresponds to the least significant byte, and bytes of greater significance within the object correspond to higher byte addresses. In a Big Endian ordering, the address of the object corresponds to the most significant byte, and bytes of less significance within the object correspond to higher byte addresses.

The terms Little and Big Endian were suggested by D. Cohen in an analogy drawn to the argument in Gulliver's Travels over whether an egg should be opened from the "little end" or the "big end." See "On holy wars and a plea for peace," by Cohen, *IEEE Computer*, October 1981, pp. 48-54.

The *M/IO! (Memory/IO!)* signal is output by the processor to distinguish accesses to the memory address space versus the I/O address space. I/O address space is intended for system I/O ports and is the target address space whenever I/O instructions are executed. All other instructions target the memory address space. The *D/C! (Data/Control! or Data/Code!)* signal is output by the processor to distinguish access to data versus program code (instructions). The *W/R! (Write/Read!)* signal is output by the processor to distinguish write accesses versus read accesses.

Platform Chip-sets use M/IO! for decoding between the memory and I/O address spaces and W/R! to command the addressed memory and I/O devices to accept or provide data. In addition to the generic purposes described individually for these signals, two reserved combinations of M/IO!, D/C!, and W/R!, are used to define the Interrupt Acknowledge and Special Cycles. Chip-sets must employ decoders to monitor for the reserved combinations in order to appropriately respond to such cycles.

Socket 7 can perform read and write cycles as either *single transfers* or *burst transfers*. In both cycle types, the processor outputs only a single address at the beginning of the cycle. Single transfers end after only one transfer. Burst transfers, without additional addresses, end only after four transfers have occurred. Due to their efficient bus usage, burst transfers are the preferred method of performing transfers, but are made only under select conditions. The four transfers of a burst occur according to a predetermined address order, or burst order.

Support for burst transfers by the external memory system is *mandatory*, as Socket 7 directly associates burst transfers with cacheable data or code, and vice versa. If data or code is cacheable, then it is *only* brought into the processor (read or prefetched) via a burst transfer. If data cached in the processor has been modified and needs to be written back to the Level 2 cache, then it is *only* output from the processor (written) via a burst transfer. Burst transfers are not used for non-cacheable data or code, for Interrupt Acknowledge, or I/O Cycles.

The *CACHE!* signal output serves different purposes, depending on whether a read or write is indicated. Assertion of CACHE! on prefetches or reads indicates that the processor considers the code or data to be cachable and is prepared to perform a four transfer burst *cache-line fill*.

Caches are organized as linear arrays of entries, known as "cache-lines." Socket 7 implicitly calls for cache-lines of 32-bytes in width (8 doublewords, or 4 quadwords). K6 also organizes every two cache-lines to share a common tag address, the collection being referred to as a cache "sector."

INDUSTRY STANDARDS

Burst Orders

critical-item

critical-item first burst order

linear burst order

round-robin burst order

486-compatible burst order

Cache-line reads are initiated when a request is made to any data location that falls within the bounds of the cache-line. The requested location is called the critical-item. Frequently the critical-item is not the first location within the cache-line. Socket 7 processors using *critical-item first burst orders*, retrieve the critical-quadword first, passing it immediately to the CPU, permitting the processor to continue execution in parallel with the retrieval of the remaining (3) cache-line locations. In a strict *linear burst order*, each quadword in the cache line is fetched in sequential order, generally low-to-high, without regard for the location of the critical-quadword.

Critical-item-first burst orders include the round-robin variant of the linear burst order and the 486-compatible (Intel) bust order. A *round-robin burst order* begins with the critical quadword and then proceeds sequentially, modulo the cache-line width, generally from low-to-high, but always sequencing in the same direction.

The *486-compatible (Intel) burst order* is a round-robin burst order that was used in the 486, where the sequence direction is a function of the even/odd status of the critical quadword. If the critical quadword is even (A3 is zero), the burst order sequences up, modulo the cache-line width. If the critical quadword is odd (A3 is one), the burst order sequences down, modulo the cache-line width.

The 486-compatible order arguably has advantages for use with certain memory banks organizations that access two 64-bit banks (representing one-half of a cache-line) in parallel. The 486-compatible burst order will sequence such that each-half of the cache-line need only be retrieved once.

Whether such bank organizations are actually used or not is largely irrelevant, if the processor supports only the 486-compatible burst order. The chip-set must use a compatible (identical) burst order to the processor. Some chip-sets only support the 486-compatible burst order.

For code prefetches or data reads, deassertion of CACHE! indicates that the processor will perform a single data transfer that will not be cached. This occurs when the L1 cache is disabled, the paging unit has determined the code or data address to be non-cachable based on the cache disable page attribute bits, TLB replacements, and locked cycles. CACHE! is likewise deasserted and only single transfers are used for Interrupt Acknowledge, Special, or I/O Cycles.

KEN! (Cache Enable) is an input to the processor, of consequence only for reads for which CACHE! is asserted. KEN! is ignored for writes and for reads when CACHE! is deasserted. KEN! is generated collectively by the L2 cache controller and *Non-Cachable Address (NCA)* logic within the chipset. KEN! is asserted on hits to L2, and on misses when the address is decoded to fall outside one or more programmable non-cachable address regions. On reads, if CACHE! and KEN! are both asserted, the final determination of cachability, the present cycle will be a four transfer burst. If the addressed location is a miss and is decoded by the NCA logic as non-cachable, KEN! will be returned deasserted, and the cycle will consist of only a single data transfer that will not be cached.

In support of a *write-back cache management protocol*, on writes, CACHE! indicates that the processor will perform a burst transfer modified-cache-line write-back. The write-back may be prompted by the *replacement process* used to manage the Level 1 Cache. The write-back may also be necessitated by the *MESI shared-memory coherency protocol*.

The *ADS! (Address Status or Address Strobe)* output is asserted by the processor to indicate the first clock of a new active bus cycle, and is kept deasserted otherwise. ADS! assertion within a given clock period signals the validity within the same clock period of the following signals (some of which are yet to be discussed): A31-A3, AP, BE7!-BE0!, CACHE!, LOCK!, M/IO!, W/R!, D/C!, SCYC, PWT, PCD.

The *BRDY! (Burst Ready)* input is generated by the external memory system to indicate that the external system considers the current transfer complete. For reads, BRDY! signals to the processor that the data being provided to the processor by the external system is valid. For writes, BRDY! signals to the processor that the data being provided by the processor has been accepted by the external system. The processor monitors for the assertion of BRDY! in each clock cycle for which there is an on-going transfer. Hence there is one assertion of BRDY! in single transfer bus cycles, and four assertions for burst transfer bus cycles. The assertion of BRDY! on the last transfer of a bus cycle is considered to terminate the current cycle and defines the validity of the following signals sampled during the same clock cycle: KEN!, WB/WT!, NA!, EWBE!, data and data parity.

The absence of a BRDY! assertion constitutes a *wait-state* request. For each contiguous clock cycle in which the external system cannot complete

the current transfer within the clock-cycle, BRDY! is kept deasserted. The processor inserts wait-states as required, suspending the termination of the bus cycle until the external system asserts BRDY! for each transfer of the bus cycle. This wait-state mechanism means Socket 7 is a semisynchronous bus, which normally operates with the speed of a synchronous bus.

The use of wait-states is common in PC platforms, arising due to disparities between the period of the bus clock and the timing of affordable memories. There are number of potentially different memory timings of interest, of which at least one usually requires wait-states. The different timings arise because the external memory system consists of a relatively fast SRAM-based Level 2 Cache, used whenever possible, and a slower DRAM-based main memory, used on L2 cache misses and for non-cachable data. Furthermore, each of these memory types has generally different read and write access times, and typically different timings for lead-off and burst operation. Finally, the timing for I/O space accesses is unrelated to the memory system timings and may vary with each I/O device.

If any of these memory timings is not less than the bus clock period by a comfortable margin, one or more wait-states must be added as required to insure a valid transfer. The memory and cache controllers of the external system collectively must dynamically generate the BRDY! timing according to the different memory timings and as a function of the hit/miss status of the L2 cache, the read/write status of the bus cycle, and whether the transfer is a lead-off transfer or a subsequent burst transfer. The different memory timings are usually programmed into non-volatile storage on the motherboard by a setup program run after the memory is installed.

address pipelining

bus-cycle pipelining

Under certain conditions, Socket 7 permits *address pipelining*, the overlapping of (at most) two bus cycles. The external system must request such *bus-cycle pipelining*, via the assertion of the processor's *NA! (Next Address)* input. During overlapped bus cycles, ADS!, which defines the validity of the address and transfer type and the beginning of the next cycle, is asserted before the last assertion of BRDY! in the current bus-cycle, which defines the current bus-cycle's completion.

When bus cycles are not overlapped, ADS! can be asserted no sooner than the first clock-cycle after the last BRDY! assertion in the previous bus-cycle. This results in an idle state between bus cycles. Memory controllers that support (by latching the address and transfer type) and request bus cycle pipelining, increase the efficiency of the bus by reducing the number of idle states between bus cycles.

Reports on CD-ROM

More detailed descriptions regarding Socket 7 are available in the following items on the companion CD-ROM. The first document includes a detailed description of every bus signal. In particular see Chapter 5, *Signal Descriptions,* and Chapter 6, *Bus Cycles*. The second document describes models that are available to simulate accurately the I/O drivers in circuit simulations of systems built around the processor. The third document provides additional insight into interfacing with a Socket 7 bus processor. In particular see section 5.1, *Processor Interface*. In addition, Section 4.1, *Processor Interface Signals*, describes the processor bus signals in summary form.

- *AMD-K6 Processor*, 20695H/0-March 1998.
- *Application Note AMD-K6 MMX Enhanced Processor I/O Model*, Document 21084, June 1997.
- *AMD-640 System Controlle*r, 21090C/10-June 1997.

Suggested Readings

Processor Local Buses

The following books provide detailed information on processor local buses and associated sockets and slots:

1. Tom Shanley, *Pentium Pro Processor System Architecture*, Addison-Wesley, 1997.
2. Don Anderson & Tom Shanley, *Pentium Processor System Architecture*, Mindshare Press, 1993.
3. Tom Shanley, *80486 System Architecture, 3rd Edition*, Addison-Wesley, 1995.
4. *Understanding x86 Microprocessors*, Ziff-Davis Press, 1987.

Chapter Summary

This chapter has presented the standards, metastandards, platform-level technologies, and component-level technologies that play important roles in PC platform architecture. Chapter 5 extends the examination of component-level technologies to memories and their control.