



Understanding Your Personal Computer



Understanding Your Personal Computer

Understanding Your Personal Computer

Note: Before using this information and the product it supports, be sure to read the general information under Appendix, "Notices and trademarks" on page 34.

First Edition (January 2000)

The following paragraph does not apply to the United Kingdom or any country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This publication could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time.

This publication was developed for products and services offered in the United States of America. IBM may not offer the products, services, or features discussed in this document in other countries, and the information is subject to change without notice. Consult your local IBM representative for information on the products, services, and features available in your area.

Requests for technical information about IBM products should be made to your IBM reseller or IBM marketing representative.

© **Copyright International Business Machines Corporation 2000. All rights reserved.**

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

About this book	v
Chapter 1. Microprocessors and memory	1
Microprocessors	1
Microprocessor speed	1
Microprocessor cache	1
Microprocessor bus	2
Microprocessor instructions	2
Memory	2
Main memory	2
CMOS memory	3
Flash memory	3
Cache memory	4
Memory organization	5
Physical and virtual memory	6
Control logic	6
Chapter 2. Expansion bus	7
PCI bus	7
PCI interface	7
PCI performance	8
Expansion bus slots	8
Plug and Play adapter configuration	8
Chapter 3. Video	9
Monitors	9
Resolution and dot pitch	9
Scanning characteristics	9
DDC protocol	10
Power management	10
Selecting a monitor	11
Using monitor controls	11
Using screen savers	11
Video control	11
Accelerated Graphics Port	12
Dynamic display memory video controller	12
SVGA mode	12
Special characters and languages	13
Chapter 4. I/O features	14
Audio ports	14
Serial ports	15
ADSL Modems	16
Universal Serial Bus ports	16
Parallel port	16
Keyboard port	18
Keyboard layout	18
Languages	21
Mouse port	22
Monitor port	22

Chapter 5. Storage devices	23
Diskette drives	23
Hard disk drives and IDE	23
CD-ROM drives	24
CD-RW drives	24
DVD drives	25
Iomega Zip drives	25
Chapter 6. Security	26
Information backup	26
Erased files	27
Computer viruses	27
Chapter 7. Software	28
Software layers	28
POST	28
Configuration/Setup Utility program	29
Operating systems	29
Device drivers	29
Latest-level device drivers	30
Device-driver types	30
Video device drivers	30
Advanced Power Management	30
Advanced Configuration Power Interface (ACPI)	31
Application programs	31
Protecting your software	31
Chapter 8. Networking and system management features	32
Desktop Management Interface (DMI)	32
DHCP and RPL	32
Remote Administration	33
Alert on LAN	33
Wake on LAN	33
Wake on Ring	33
Appendix. Notices and trademarks	34
Notices	34
Trademarks	34

About this book

This book provides information about your IBM® Personal Computer. It includes information about microprocessors, memory, bus architectures, input/output features, video, ports, storage devices, security features, and software.

Not all features and capabilities described in this book are available on all models. For specific information about features and instructions on how to set up, operate, maintain, or install options into your computer, see the user guide publication that comes with your computer.

Chapter 1. Microprocessors and memory

The microprocessor and memory subsystem consists of components on the system board that perform logical operations and calculations, control memory, and manage data transfer operations.

The devices and features that make up the microprocessor and memory subsystem are:

- Microprocessors and buses
- Memory and control logic

Microprocessors

The microprocessor controls most of the activity in your computer. Almost all information passes through it, whether it is a keyboard stroke, data from a disk drive, or information from a communication network. The microprocessor reads, changes, processes, and reroutes information as needed.

Your computer comes with an Intel® microprocessor that has a 64-bit internal data path and integrated L2 cache. In addition, the microprocessor incorporates Intel MMX™ technology and Streaming Single Instruction Multiple Data (SIMD) Extensions technology. These technologies enhance the performance of the microprocessor in processing graphic, video, and audio data, thereby enhancing the performance of multimedia and communication software. For information about the type and speed of the microprocessor in your computer, view the System Summary screen that is displayed by the Configuration/Setup Utility program.

The microprocessor includes *system management mode (SMM)*, which enables the microprocessor to control power used by peripheral devices and other components. This makes processing more energy efficient and allows the system to run cooler.

Microprocessor speed

The microprocessor operating speed, referred to as *clock speed*, is measured in *megahertz (MHz)*¹. The design of the microprocessor determines the maximum clock speed at which it can operate. A quartz crystal on the system board generates a pulse to the microprocessor, causing it to operate at a specific speed. A *clock cycle* is the time that the microprocessor takes to perform one instruction at a given clock speed. Clock cycles are measured in *nanoseconds (ns)*.

Microprocessors can operate at two clock speeds: an *internal speed* for operations within the microprocessor and an *external speed* for transferring data into and out of the microprocessor.

Microprocessor cache

Both *level-1 (L1) cache* and *level-2 (L2) cache* are integrated into the microprocessor on the system board. Cache is high-speed memory that stores information most often used by the microprocessor. Integrated cache provides a performance increase over cache on the system board. See “Cache memory” on page 4 for more information.

¹ MHz denotes only the internal clock speed of the microprocessor, not application performance; many factors affect application performance.

Microprocessor bus

The microprocessor has an external bus that connects it with main memory and control circuits. This pathway, which is also called the *processor bus* or *local bus*, has the same bus width as the microprocessor and operates at the same external speed.

Another computer bus, called the *I/O bus* or *expansion bus*, carries data and instructions between the microprocessor bus and the peripheral devices. The width of the I/O bus is 32 bits. With advanced bus technologies, the speed of the I/O bus might approach that of the microprocessor bus. With standard bus technology, however, the speed of the I/O bus is much slower than that of the microprocessor bus. See Chapter 2, “Expansion bus” on page 7 for additional information.

Microprocessor instructions

Data and instructions are necessary for each processing operation that the microprocessor performs. Data and instructions are copied from memory into *registers* within the microprocessor. Registers are also used to store the data that results from each processing operation until the data is returned to memory.

The set of instructions that the microprocessor can perform determines whether the computer can run a particular program. For example, programs written for 32-bit computers require a microprocessor capable of performing 32-bit instructions.

Memory

Your computer uses several types of memory to store information. This section explains memory concepts, types of memory, and how the types of memory are used.

Although a computer is a complex machine, the method it uses to store information is quite simple. All information (data and instructions) is stored in a coded format made up of 0s and 1s. Memory is a series of switches, with a 0 representing an open switch and a 1 representing a closed switch. Each switch represents the smallest unit of computer storage, a *bit*; eight consecutive bits of storage equals a *byte*. Memory is allocated in *kilobytes (KB)*, *megabytes (MB)*, and *gigabytes (GB)*, where 1 KB equals approximately 1 000 bytes, 1 MB equals approximately 1 000 000 bytes, or 1 000 KB, and 1 GB equals approximately 1 000 000 000 bytes, or 1 000 MB. In 1 MB of memory, the computer can store over 1 000 000 characters of information. The computer organizes its memory by assigning an *address* to each byte as a point of reference. The first address is 0, and the addresses increase sequentially. When information is written to memory, each encoded character is placed at a specific address.

Main memory

Main memory (also known as *system memory*) is a temporary workspace that is active only while your computer is on. Any information placed in main memory is lost when you turn off your computer. Therefore, if you create or modify information, you must save the data to permanent storage, such as a diskette or hard disk.

IBM computers use two types of memory modules for main memory: *dual inline memory modules (DIMMs)* and *Rambus inline memory modules (RIMMs)*. To find out which kind of memory modules are used in your computer, see the user guide publication that is provided with your computer.

The DIMMs contain a type of memory known as *synchronous dynamic random access memory (SDRAM)*. SDRAM is used for temporary storage of data and instructions. SDRAM is *volatile* memory, which means

that it must be constantly refreshed by an electric current. While the computer is turned off, no current is supplied to the modules, so no data is retained in SDRAM.

RIMMs contain a type of memory known as *Rambus DRAM (RDRAM)* modules. RDRAM was created to help close the processor-memory performance gap at an affordable cost. The various features of RDRAM enable extremely high peak and effective bandwidth. If your computer uses RIMMs for main memory, you can get more information from the user guide publication that is provided with your computer.

Both ECC (error checking and correction) and non-ECC SDRAM are supported in most models for both DIMMs and RIMMs. (Some microprocessors do not support ECC. For more information, see the user guide publication.) If a mix of ECC and non-ECC is used, the memory modules will function as non-ECC.

The module connectors are located on the system board. For information about the type and amount of main memory in your computer, view the System Summary screen displayed by the Configuration/Setup Utility program.

CMOS memory

Your computer also uses some *nonvolatile RAM (NVRAM)*, also called *complementary metal-oxide semiconductor (CMOS)* memory, for storing configuration and setup information. Powered by a small battery, CMOS retains its contents while your computer power is off. CMOS maintains information about:

- Date and time
- Security features
- Power-management settings
- Storage devices
- Keyboard and mouse
- Plug and Play configuration information
- Port assignments
- I/O addresses and interrupts
- Other selectable features

Flash memory

To store programs for startup procedures and other internal operations, some computers use *read-only memory (ROM)*. These programs are encoded in ROM modules on the system board. ROM is *nonvolatile* memory, which means that it retains its contents when the computer is turned off. Generally, the contents of a ROM module cannot be modified. However, *electrically erasable programmable ROM (EEPROM)* modules (also referred to as *flash memory*) can be reprogrammed while they are in the computer. Your computer stores various system programs and data in *flash memory* so that they can be updated whenever enhancements are made. Flash memory contains:

- Basic input/output system (BIOS)
- Power-on self-test (POST)
- Configuration/Setup Utility program
- Vital product data (VPD)
- Microprocessor update code

Cache memory

Microprocessors can be so fast that main memory cannot respond to read and write requests as quickly as the microprocessor can send them. In some cases, main memory imposes one or more *wait states* on the microprocessor when it reads data from or writes data to memory.

A wait state is a period of time (one microprocessor clock cycle) during which the microprocessor suspends processing and waits for memory to respond to a read or write operation (a *memory I/O operation*). The speed of a memory I/O operation is measured in microprocessor clock cycles, so the microprocessor clock speed determines the minimum time required for a memory I/O operation. Wait states cause the computer to operate less efficiently than if the microprocessor were able to continue processing data during memory I/O operations.

A method of reducing the need for wait states is by using *cache memory*, which improves system performance by temporarily storing frequently used data and instructions in a *cache*. A cache is a buffer between the microprocessor and main memory.

For information about the cache memory in your computer, view the *System Summary* screen displayed by the Configuration/Setup Utility program.

Level-1 cache

The type of *level-1 (L1)* cache is determined by the microprocessor installed in your computer. The L1 cache contains high-speed memory, known as *static random access memory (SRAM)*, that can respond to memory I/O operations without imposing wait states on the microprocessor.

L1 cache memory is used to store the information most often used by the microprocessor. This allows the microprocessor to handle information faster than if it had to use the system memory each time it needed new information. During processing, the cache controller copies other data and instructions into the cache, replacing data and instructions that are no longer needed.

Performance is improved each time the microprocessor finds what it needs in the cache (a *cache hit*). If it does not find what it needs (a *cache miss*), the cache controller must locate the data or instruction in memory and copy it into the cache, while one or more wait states are imposed on the microprocessor. The cache controller manages the use of the cache so that the number of cache hits far exceeds the number of cache misses.

Level-2 cache

Your computer has *level-2 (L2)* cache memory integrated into the microprocessor. L2 cache complements L1 cache to increase the probability of cache hits. If the microprocessor cannot find what it needs in L1 cache, it searches L2 cache. If the data or instruction is not in either cache, the cache controller locates it in main memory and copies it into both caches.

Cache mode

The microprocessor frequently updates cache memory with changed data. Caches, in turn, pass these changes to main memory.

When updating cache memory, your computer uses the *write-back mode*. In write-back mode, the microprocessor updates the cache, then goes on to perform other functions while the cache controller updates main memory.

Memory organization

Operating systems are responsible for allocating memory space, assigning addresses, and performing many other tasks associated with memory management.

Memory management

Operating systems such as Microsoft® Windows® 95, Windows 98, Windows NT®, and Windows 2000 Professional® have very sophisticated memory management systems. These operating systems use a technique called *disk swapping* or *disk paging*, in which memory is organized into small page fragments. The operating system can mark pages as being present in main memory or as available on the disk. If your computer does not have enough memory to meet the needs of your active programs, these operating systems transfer the least-used information from memory to the hard disk to make more memory available. When the information on the hard disk is needed, it is exchanged with other information in memory.

Microsoft Windows 95, Windows 98, Windows NT, and Windows 2000 Professional also use memory in another way: These operating systems assign blocks of memory as *virtual DOS machines (VDM)*. Each VDM runs independently of the others, providing the same function as a separate computer running DOS. With very few exceptions, these operating systems are fully compatible with application programs written for DOS.

If you are using any of these operating systems, adding memory can increase the overall performance.

Traditional memory management

Your operating system organizes memory into the following types:

- Conventional
- Extended
- Expanded

Conventional memory is the first 1 MB of memory address space. Of this, the first 640 KB is available for use by your operating system and application programs. Memory from 640 KB through 1 MB is reserved as a work space for hardware devices and the BIOS.

Extended memory starts above the first 1 MB, appended to conventional memory. Use of this space is dependent on the operating system and application programs. Not all programs can use extended memory.

Expanded memory is controlled through an expanded memory specification (EMS) device driver, such as the one provided with your operating system. The EMS device driver is commonly known as a *memory manager*. The EMS device driver uses part of the reserved area of conventional memory as a work space and the memory above the first 1 MB as a storage area. The memory above the first 1 MB is not written to or read from directly. It is broken into 16 KB pieces called *pages*, or 64 KB pieces called *frames*, and is moved in and out of the reserved area as needed. When a page or frame is moved into the reserved area, it can be read from and written to as any other address in conventional memory.

Physical and virtual memory

Physical memory consists of all the addressable memory locations in the computer. Physical memory is used to store such items as the operating system, video data, and instructions and data the computer uses to run programs.

Virtual memory is memory that appears to be allocated to application programs. The operating system uses a portion of the hard disk as virtual memory, swapping data and instructions between the hard disk and physical memory.

Virtual memory makes *multitasking* possible. In a multitasking session, the memory requirements of all the programs that might be running in the system at the same time can far exceed the amount of physical memory that is available. The operating system allocates virtual memory to meet the total memory requirements of each program and then manages the available physical memory to meet the actual requirements. Thus, the amount of virtual memory that is allocated can be much greater than the amount of physical memory that is installed in the computer.

Control logic

The control logic consists of the modules on the system board that control access to main and cache memory by the microprocessor and I/O devices. Control logic includes the following devices and functions:

- Audio-port controller and interface
- Bus-mastering Integrated Drive Electronics (IDE) interface
- Counters and timers
- Direct memory access (DMA) controller
- Diskette-port controller and interface
- Interrupt controller
- Keyboard and mouse port controller and interface
- Memory (DRAM) controller
- Microprocessor interface control
- Parallel-port controller and interface
- Peripheral Component Interconnect (PCI) bus interface
- Power-management controller
- Serial-port controller and interface
- Universal Serial Bus (USB) port controller and interface

Chapter 2. Expansion bus

This section gives an overview of the expansion bus and explains how advanced buses can improve performance.

A computer *bus* is a pathway of wires and signals that carries (or transfers) information inside the computer; information includes data, addresses, instructions, and controls. The microprocessor has an external bus, called the *microprocessor bus* or *local bus*, that carries information between the microprocessor and main memory. The microprocessor bus has the same bus width (64 bits) as the microprocessor and operates at the same external speed.

Another computer bus, the *expansion bus*, carries information between the microprocessor or memory and peripheral I/O devices. While microprocessor-bus performance has improved rapidly, improvements in expansion-bus performance have not equaled those in microprocessors and some peripheral devices, such as video and disk controllers. Regardless of how fast the microprocessor and other components are, data transfers between them must pass through the expansion bus.

PCI bus

Peripheral component interconnect (PCI) is an advanced expansion-bus standard developed by the computer industry to keep up with performance improvements of processor buses and advanced peripheral devices. Although advanced designs can match the performance of the microprocessor bus only up to a point, they do achieve higher throughput by speeding up the expansion bus and widening its data path.

The PCI bus is an answer to the need for a higher-performance expansion bus. PCI architecture offers many features that improve expansion-bus performance, such as:

- Microprocessor independence
- Industry-standard compatibility
- Wider data path (32 bits)
- Faster data-transfer rates
- More efficient data-transfer methods
- Enhanced peripheral-device performance
- Automatic configuration

PCI interface

The PCI bus connects to the microprocessor bus through a buffered bridge controller. A *bridge* translates signals from one bus architecture to another. PCI devices get all their data and control information through the PCI controller. The PCI controller looks at all signals from the microprocessor bus and then passes them to the peripheral devices connected to the PCI bus. However, the PCI bus is not governed by the speed of the microprocessor bus. The PCI bus can operate at speeds as fast as 33 MHz, slow down, or even stop if there is no activity on the bus, independent of the microprocessor operations. Microprocessor independence also makes the PCI bus adaptable to various microprocessor speeds and families and enables consistency in the design and use of PCI peripheral devices across multiple computer families.

Expansion bus

PCI performance

The PCI bus has a 32-bit data path, allowing for high data-transfer rates. Depending on the mode of operation and computer components used, the PCI bus can transfer data at speeds up to 132 MB per second.

Expansion bus slots

If you want to add new capabilities to your computer, you can do so by installing optional *adapters*. Your computer provides PCI-bus expansion slots that you can use to connect high-performance devices to your computer, such as graphics, SCSI, or LAN adapters.

There are a set of dedicated PCI expansion slots in your computer. These slots are grouped together on the system board. Some models might include support for ISA slots or a riser card. See your user guide publication for further details.

The width of the expansion bus determines the type of adapters that the computer supports. The dedicated PCI slots accept only 32-bit PCI adapters. The width of the bus does not affect software compatibility.

Plug and Play adapter configuration

A method of configuring adapters and devices called *Plug and Play* is used on your computer to make expansion an easy task. A Plug and Play adapter comes with built-in identification and configuration specifications that are set in the adapter to be provided to the computer during startup.

Adapters that are designed for PCI slots are Plug and Play devices and are automatically detected by operating systems, such as Windows 98 or Windows 95, that support Plug and Play technology. If a Plug and Play adapter is installed in a computer that has an operating system, such as Windows NT 4.0, that does not support Plug and Play, the adapter must be configured using the device-driver software that comes with the adapter. If you want to install a Plug and Play adapter in a computer whose operating system does not support Plug and Play technology, see the documentation that comes with the adapter for instructions.

This information in the Plug and Play adapter memory is interpreted by the *basic input/output system (BIOS)*, which supports Plug and Play technology. The BIOS routines automatically configure the adapter if the required resources are not already allocated to other devices. See the user guide publication for more information on Plug and Play devices and resource conflicts.

Chapter 3. Video

This section describes the video features of your computer, including monitors, the graphics controller, video modes, and other video features.

Monitors

Monitors come in two major classes: the traditional cathode-ray tube (CRT) monitors and flat-panel monitors. CRT monitors use cathode-ray tubes with advanced phosphor coatings for the display. Flat panel monitors use liquid crystal display (LCD) screens, which are usually smaller and lighter. Although there are various types of monitors, the most common distinguishing characteristics are resolution, dot pitch, scanning rates, screen size, and color.

Resolution and dot pitch

The entire image on a monitor screen is made up of many tiny dots. These dots are called *picture elements* or *pixels*. The monitor *resolution* (also called addressability) is described by the number of pixels that fill the screen in an array of horizontal rows and vertical columns. For example, the pixels might be arranged in a horizontal-by-vertical array of 640 x 480 pixels to make up a typical screen image. Higher resolutions, such as 800 x 600 or 1024 x 768 pixels, mean more dots on the screen. The benefit of higher resolutions is either displaying more information at one time or displaying the same information with a sharper image.

Dot pitch refers to the space between the individual dots or pixels. The inside of the CRT monitor screen is coated with light-emitting materials called *phosphors*. Color monitors use three phosphors: red, green, and blue. The phosphor coating consists of dots arranged in a red-green-blue pattern. The distance from the center of one phosphor dot to the center of the next phosphor dot of the same color is the dot pitch. Dot pitch is measured in millimeters, such as 0.28, 0.31, 0.39, and 0.41. In general, the monitor with the highest resolution and finest dot pitch produces the sharpest image. A monitor with a fine dot pitch, such as 0.28, produces sharper, better-defined characters than a similar monitor that uses a coarser dot pitch, such as 0.41.

Monitor characteristics can be very important. A computer that is used extensively for detailed graphics might need a monitor that shows more colors, at a higher resolution and finer dot pitch, than one that is used exclusively for word processing.

Scanning characteristics

The phosphors on the inside of the CRT monitor screen hold their brightness and color for a very short time. Therefore, the image has to be scanned (redrawn) many times each second to refresh the phosphors. The *vertical refresh rate* is the speed at which the image on the screen is scanned. Refresh rates are specified in *hertz (Hz)*.

The image is scanned from top to bottom (or vertically). By the time the bottom pixel rows are drawn, the top rows are starting to fade. There are two ways of scanning the image: *interlaced* scanning and *noninterlaced* scanning. The interlaced method draws every other row from top to bottom, then starts at the top to fill in the rows that were skipped on the first pass. The noninterlaced method draws the complete image on each pass.

The vertical refresh rate affects the stability of the image on the screen. Monitors that have a fast, noninterlaced refresh rate provide a very stable image with little or no detectable *flicker*. Monitors with

Video

slower refresh rates tend to have more noticeable flicker. Flicker can be particularly noticeable with white and other light-colored backgrounds. You can minimize flicker by changing the colors you work with. Many programs enable you to select colors.

Sensitivity to flicker varies from person to person. Noninterlaced flicker occurs when the screen is redrawn at a rate that is slow enough to be detected by the human eye. Interlaced flicker, or jitter, occurs when the pattern of pixels in adjoining lines is very different or when an object on the screen is made up of alternating bright and dim lines.

Some application programs cannot reproduce all colors; they use a technique called *dithering* to combine two colors. For example, light red might be produced by alternating lines of white and red. As the image is continually refreshed using the interlaced method, the red lines are drawn on one pass and the white lines on the next pass. As a result, the light-red object might appear to flicker while the rest of the image on the screen remains stable.

DDC protocol

Potential problems with monitor flicker are reduced by operating a monitor at its highest refresh rate. To simplify the job of setting the refresh rate, some monitors offer Plug and Play capability using the *VESA Display Data Channel (DDC)* protocol. DDC is an industry standard for passing configuration information between a monitor and a computer. Both the monitor and the computer must have DDC capability.

Your computer supports the DDC2B protocol. With DDC, the computer can receive configuration information sent from a DDC monitor. When you turn on the computer with a DDC monitor connected, the monitor immediately starts sending its identification information to the computer. This information, called *Extended Display Identification Data (EDID)*, includes a description of the monitor and its operational limits and timings. Power-on self-test (POST) programs use the information from the monitor to automatically configure the video controls to operate the attached monitor in its optimum mode. Thus, when the monitor is turned on, it is already set to the highest refresh rate that the monitor and graphics controller can support.

Power management

Power management reduces the power consumption of your computer when you are not using it. A monitor with *VESA Display Power Management Signaling (DPMS)* support can save a significant amount of energy and have a longer useful life. In DPMS-compliant monitors, the power can be reduced in three modes: *Standby*, *Suspend*, and *Off*. The stages that can actually be used, however, depend on the power-management capabilities of your computer.

In Standby mode, the monitor (and other devices) are switched to a reduced power state if you do not use your mouse or keyboard for a specified period. The devices remain in a state of readiness and return to full power if you press a key or move the mouse. Suspend mode further decreases monitor power, and Off mode reduces monitor power to its lowest level. Refer to “Advanced Power Management” on page 30 for more information on power management.

Check the documentation that comes with your monitor to see if your monitor is enabled for power management. If you find no information about power management, your monitor probably does not support this feature.

Selecting a monitor

The graphics controller that is installed in your computer is a high-performance, high-resolution graphics accelerator that includes many performance-enhancing features for your operating-system environment. It has a built-in monitor connector. See “Video control” on page 11 for more information. The graphics controller supports various Super Video Graphics Array (SVGA) monitor resolutions. If you connect an SVGA monitor to your computer, you can take advantage of SVGA resolutions such as 1024 x 768. However, the default resolution might be different for certain software packages, especially if they are not SVGA-ready. Also, you might have to install device drivers for certain software packages. Refer to the device-driver installation instructions that come with the software on the *Product Recovery CD* for a list of programs and the device drivers that they require. If you did not receive a *Product Recovery CD*, some device drivers might be available on a separate CD.

To take full advantage of the video capabilities provided by your computer, use a noninterlaced monitor that is capable of supporting high resolution and a fast refresh rate. Also, the type of monitor you select depends on the software you use. Some applications are less demanding than others.

Using monitor controls

A monitor has brightness and contrast controls for adjusting the image on the monitor screen for maximum viewing comfort. On a new monitor, set the controls near their center positions. Setting these controls to their maximum positions can decrease the useful life of the monitor. As the monitor is used more, the image tends to become dimmer. You can adjust the brightness and contrast controls to compensate for this condition. Monitors also have additional controls for fine-tuning the screen image. Check the documentation that comes with your monitor for more information on controls.

Using screen savers

Leaving the monitor turned on for prolonged periods without changing the image can harm the monitor. Depending on what is on the screen, the image might be permanently burned into the phosphor coating or cause darkening of some areas of the screen.

With DPMS-compliant monitors, you can avoid this problem by using the power-management features of your computer. Also, you can use a *screen saver* program, which turns off the video signal when the keyboard or mouse has not been used for a defined period. This action results in a dark, blank, or animated screen. When you are ready to work with the computer again, pressing a predefined key or a password restores the image to the screen. Several companies produce screen savers. For more information, contact your IBM reseller or IBM marketing representative.

Video control

Video control is the process of manipulating information into the desired format and sending it to the monitor screen. Video control can be handled by the microprocessor or a separate video or graphics controller. For graphics and window environments, video control involves operations that handle a high volume of data and computations. These types of operations typically require high speed. Moving a window from one location to another, dragging an icon across the screen, and scrolling through a document are examples of such operations.

Your computer uses either an Accelerated Graphics Port or a dynamic display memory video controller. Read the section that is appropriate for your computer.

Accelerated Graphics Port

The Accelerated Graphics Port (AGP) provides a full bus-mastering subsystem dedicated to graphics operations. The AGP provides a dedicated graphics subsystem that works independently of the microprocessor and provides direct access to main memory (if your computer supports AGP).

On some models the AGP subsystem is integrated with the system board and provides a monitor connection on the back of the computer. Other models come with an AGP adapter installed in an expansion slot.

AGP effective bandwidth

The AGP provides high-bandwidth capabilities and operates at a minimum clock rate of 66 MHz, compared to the 33 MHz speed of the PCI bus. The AGP has different standards of effective bandwidth. The AGP 1X operates at a clock rate of 66 MHz and has an effective bandwidth of 266 Megabytes per second (MBps). The AGP 2X operates with an effective bandwidth of 533 MBps and the AGP 4X operates with an effective bandwidth of 1 066 MBps. The higher bandwidth allows for faster graphics in applications that use 2D and 3D design.

Digital Visual Interface

The Digital Visual Interface (DVI) is a connector that you can use to attach an LCD flat panel monitor to some computers. DVI comes in several standards, including:

- DVI Digital (DVID), which you can use to attach a digital display to your computer
- DVI Integrated (DVII), which you can use to attach an analog or digital display to your computer

For more information about connecting your monitor using a DVI, see the user guide publication.

Dynamic display memory video controller

Models with integrated video include the dynamic display memory (DDM) video controller that is integrated into the graphics memory controller hub (GMCH) on the system board. It has a PCI-like interface dedicated to high-performance 2D and 3D graphical display applications. DDM provides dynamic allocation of system memory for graphics functions resulting in an increase in overall system performance.

The graphics controller is supported by device drivers that are provided on the *Product Recovery CD* that comes with your computer. Instructions for installing the device drivers are also provided in README files on the *Product Recovery CD*. If you did not receive a *Product Recovery CD*, some device drivers might be available on a separate CD. If your computer has IBM-preinstalled software, the video device drivers have already been installed on the hard disk. However, the information that is included in the device-driver installation instructions will be helpful if you ever need to reinstall or update the device drivers, or if you want to change the video resolution or color depth.

SVGA mode

Super video graphics array (SVGA) is a video standard for displaying text and graphic images. Like other video standards, SVGA supports a variety of *video modes*. Video modes are different combinations of resolutions, refresh rates, and colors defined by a video standard for displaying text or graphics. SVGA provides higher-resolution graphics, faster refresh rates, and more colors than previous video standards, such as video graphics array (VGA). At the same time, SVGA video supports all standard VGA modes.

Your computer supports SVGA modes on a variety of monitors. Also, the SVGA modes on your computer are in accordance with the specifications of the *International Organization for Standardization (ISO) 9241*

Part 3, Visual Display Requirements. SVGA supports reduced-flicker operation when used with noninterlaced monitors that meet ISO standards.

Some application programs require video device drivers to use SVGA modes. The device drivers for your computer and instructions for installing them (README files) are provided on the *Product Recovery CD* that comes with your computer. If you did not receive a *Product Recovery CD*, some device drivers might be available on a separate CD. If your computer has IBM-preinstalled software, the device drivers have already been installed on the hard disk of your computer.

Some application programs provide their own video device drivers to take advantage of SVGA modes. These device drivers are installed through a setup program built into the application program. Setting up these applications to operate in SVGA modes might be confusing, because terminology for these modes is not consistent throughout the industry. If you are uncertain about the terminology used by your application program, refer to the documentation that comes with the application for additional information.

Special characters and languages

The video controller can display a variety of characters and languages. The language that is supported depends on the *code page* loaded by your operating system.

You can find additional information about code pages in the documentation that comes with your operating system.

Chapter 4. I/O features

Input/output (I/O) devices provide you with many ways of getting information into and out of your computer. You can connect a variety of standard and optional devices to your computer, including a monitor, keyboard, mouse, and printer.

This section includes information about I/O device connectors and also explains the characteristics of some of the I/O devices.

The system board of your computer includes *connectors* (or *ports*) for attaching I/O devices. An I/O connector is the physical connector on the computer that you can use to connect an external I/O device. These connectors are also called ports because they provide a communication pathway into the computer for an I/O device. Each connector is identified by a number and an address that distinguishes it from other connectors. No two connectors can have the same address. The terms *I/O connector* and *I/O port* are often used to mean the same thing.

Note: Your computer model might not include all of the I/O devices or connectors described in this section. Refer to the user guide publication for more specific information.

Following are the system board I/O connectors that might be available on your computer:

- Audio: microphone in, line out, line in, and headphone out
- Serial
- Universal Serial Bus (USB)
- Parallel
- Keyboard
- Mouse
- Monitor

Audio ports

The audio controller is integrated on the system board and supports SoundBlaster, SoundBlaster Pro, and Windows Sound System. The controller has three audio ports. Audio ports provide the ability to capture and play back sound and music, as well as the ability to play sounds that are part of multimedia application programs. You can enable or disable audio support by using the Configuration/Setup Utility program to enable or disable these ports.

The ports are labeled:

Mic in	Connection for a microphone
Line in	Connection for external audio devices, such as a CD player
Line out / Headphone out	Connection for headphones or powered external speakers

Note: When headphones or speakers are connected to the headphone jack, the internal computer speaker is disabled.

There might also be an internal speaker and connections on the system board for attaching the speaker to the CD-ROM drive.

Serial ports

Your computer comes with one or two serial ports. The serial ports are used to communicate with devices such as printers, plotters, and modems designed for the serial interface. Some modems are installed inside the computer; others are attached externally using a serial port connector. For example, you can connect an external modem to a serial port to communicate with other computers over telephone lines.

The serial port transfers data one bit at a time at speeds ranging from 300 to 115 000 bits per second (bps). The serial ports are 16550-UART (universal asynchronous receiver/transmitter) compatible so that they can support high-speed modems.

You can increase the number of serial ports by installing a serial adapter in one of the expansion slots.

Operating systems and application programs often designate serial ports with the name *COM* (*communication*). Many programs use *COM1*, *COM2*, *COM3*, and *COM4* in their setup commands or menus to identify serial ports. When setting up your applications, you might have to identify the COM port to which a particular device is connected. Also, if you are adding a serial adapter or modem, you might have to configure the adapter or modem for a specific COM port. This might involve setting switches on the adapter or modem as well as changing settings in the Configuration/Setup Utility program.

The serial connector in your computer can use any of four available port settings, provided that a different setting is used for each. The settings include the port *address* (in hexadecimal) and the *IRQ* (*interrupt request line*), which determines how the microprocessor responds to an interrupt from the serial port.

The four available port addresses are (in hexadecimal):

- 3F8
- 2F8
- 3E8
- 2E8

The available IRQs are:

- IRQ 4
- IRQ 3
- IRQ 9
- IRQ 10
- IRQ 11
- IRQ 14
- IRQ 15

There is no direct relationship among the port connectors, the available port settings, and the COM numbers. When you start the computer, the power-on self-test (POST) assigns COM numbers to the port addresses that are actually in use at the time. POST goes down the list of addresses sequentially to assign COM numbers to each address in use by a serial device. If an address is not in use, a COM number is not assigned to it. POST assigns the next available COM number to the next address in use, if any.

You can view the connector addresses and IRQ settings for serial ports using the Configuration/Setup Utility program. You do not see the COM numbers on the Configuration/Setup Utility program screen. To see port settings and COM assignments for all serial ports, you can use one of the diagnostic programs, such as IBM Enhanced Diagnostics, which is available with the IBM-preinstalled software package.

Generally, you do not have to change serial-port settings. However, if you add a serial adapter or use special hardware or software that causes a configuration conflict, you might have to change port settings.

I/O Features

If so, you can change the settings for the serial port using the Configuration/Setup Utility program. Refer to the user guide publication for more information about using the Configuration/Setup Utility program.

To connect a device to your serial port, use a serial cable (purchased separately) to connect the device to the serial connector at the back of the computer. Refer to the documentation that comes with your I/O device for specific installation instructions.

ADSL Modems

ADSL (Asynchronous Digital Subscriber Line) modems vastly improve Internet connectivity with ease of use and connection speed. The difference with ADSL over traditional analog modems is the way the telephone lines are used. ADSL divides your telephone line into three different channels: voice, download and upload.

The channels are divided based on the copper wires in the standard RJ-11 telephone jack. Typically four or six copper wires are connected to the telephone. For ADSL, the middle two wires carry voice transmissions at 4MHz and below. The other two wires, or pairs of wires if there are six total, are reserved for data download from your Internet Service Provider (ISP) and data upload to the Internet from your computer. The download channel carries between 2.0 - 8.0 megabits per second (Mbps). The upload channel carries from 640 kilobits per second (kpbs) - 960 kbps.

The benefits of ADSL connectivity include:

- Continuous connectivity. Your Internet connection can remain constant. Once the connection is established with your local phone service, no dial-up or disconnect is needed. Your computer can remain attached to the Internet 24 hours a day, much like an Ethernet connection.
- Plain Old Telephone Service (POTS). Because the telephone line is divided into three channels, normal voice telephone usage is available on top of your Internet connection. You can receive and make telephone calls or faxes while still maintaining your Internet connection in the background.

Universal Serial Bus ports

Universal Serial Bus (USB) is a high-speed serial interface standard. The USB port is a connector for devices that previously used serial, parallel, keyboard, mouse, and game ports. USB technology uses Plug and Play features to determine what device has been connected to the port. Each USB device is accessed by a unique USB address. A device called a *hub* is used to convert a single USB connector into multiple attachment points. A hub has multiple ports where devices can be attached. USB operates at up to 12 megabits per second (Mbps) and supports a maximum of 127 devices.

Parallel port

Your computer has a parallel port connector as a standard, built-in feature. The parallel port is most often used to communicate with a parallel printer, but it can also be used with a variety of other parallel devices, such as external Iomega Zip drives and Internet cameras.

The parallel port transfers data 8 bits at a time, compared to the serial port, which transfers data 1 bit at a time. The parallel port on the computer supports extended, high-speed modes, which means that it can transfer data up to 10 times as fast as a standard parallel port.

You can increase the number of parallel connectors by installing a parallel adapter in one of the expansion slots.

Operating systems and application programs often designate parallel ports with the name *LPT* (*line printer*). Many programs use *LPT1*, *LPT2*, and *LPT3* in their setup commands or menus to identify parallel ports. When setting up your applications, you might have to identify the LPT port to which a particular device is connected. Also, if you are adding a parallel adapter, you might have to configure the adapter for a specific LPT port. This might involve setting switches on the adapter as well as changing settings in the Configuration/Setup Utility program.

Each parallel connector or adapter on your computer can use any of three available port settings, provided that a different setting is used for each. The settings include the port *address* (in hexadecimal) and the *IRQ* (*interrupt request line*), which determines how the microprocessor responds to an interrupt from the parallel port.

The three available port addresses are (in hexadecimal):

378
278
3BC

The two available IRQs are:

IRQ 7
IRQ 5

There is no direct relationship among the three available port settings and the three LPT numbers. When you start the computer, the POST assigns LPT numbers to the port addresses that are actually in use at the time. POST goes down the list of addresses sequentially to assign LPT numbers to each address in use by a parallel device. If an address is not in use, an LPT number is not assigned to it. POST assigns the next available LPT number to the next address in use, if any.

You can view the connector address and IRQ setting for the built-in parallel port using the Configuration/Setup Utility program. However, you do not see the LPT number on the Configuration/Setup Utility program screen. To see port settings and LPT assignments for all parallel ports, you can use one of the diagnostic programs, such as IBM Enhanced Diagnostics, available with the IBM-preinstalled software package.

Generally, you do not have to change the setting for the built-in parallel port. However, if you add a parallel adapter or use special hardware or software that causes a configuration conflict, you might have to change the parallel-port setting. You can change the setting using the Configuration/Setup Utility program. Refer to the user guide publication for more information about using this program.

The parallel port can operate in five modes. One is a *standard*, unidirectional mode; the other four are *extended*, bidirectional modes that provide additional function and higher performance. Refer to the documentation that comes with printers and other parallel devices to determine the appropriate parallel mode to use and the required device drivers.

Standard	This is the default mode. In this mode, the parallel port is limited to writing information to the device that is attached to it. You can use this mode with most IBM-compatible parallel printers. This mode is sometimes referred to as AT-compatible mode.
Bidirectional	This is a bidirectional mode that is used for data transfer to other computers and supported devices. This mode is sometimes referred to as PS/2-compatible mode.
ECP	The extended capabilities port (ECP) mode is a high-performance, bidirectional mode that uses direct memory access (DMA) for data transfer to a high-speed printer or other devices.

I/O Features

EPP The enhanced parallel port (EPP) mode is a high-performance, bidirectional mode that has capabilities similar to ECP mode; the main difference is that EPP data transfers are processor-initiated instead of DMA. EPP supports the connection of up to eight external devices, such as hard disk drives, CD-ROM drives, tape drives, diskette drives, and a printer, to the parallel port. These devices can be connected to each other in a *daisy chain* arrangement, or they can be connected through an external multiplexer. The attachment of multiple devices requires device drivers that are supplied by the device manufacturers.

ECP/EPP This mode combines the capabilities of the ECP and EPP modes. Select this mode if you have both ECP and EPP devices connected to the parallel port.

Note: ECP, EPP, and ECP/EPP modes are available only if the port address is hex 378 or hex 278.

You can set these modes of operation for the parallel port using the Configuration/Setup Utility program. Refer to the user guide publication for more information about using this program.

To connect a parallel I/O device, such as a printer, use a parallel cable (purchased separately) to connect the device to the parallel connector at the back of the computer. Refer to the documentation that comes with your I/O device for specific installation instructions.

Keyboard port

The keyboard is one of the devices that you use to communicate with your computer; the keyboard connects to the keyboard connector on the back of the computer. Use your keyboard for the following:

- Typing information
- Performing application functions
- Performing computer operations
- Typing numbers with the numeric keypad

Keyboard layout

The illustrations in this section show the 104-key keyboard. The keyboard includes the following groups of keys:

- Standard typing keys
- Function keys
- Special computer keys
- Cursor keys
- Numeric keypad
- Function keys for Windows interface



Typing keys

The layout of the typing keys is similar to that of a typewriter. The standard typing keys are *typematic*; that is, a key repeats typing the character as long as you hold it down. The home row contains raised identifiers for the touch-typist. The 1(one) and l(lowercase L) characters are not interchangeable as they are on some typewriters.


The Shift key is used with the typing keys to produce uppercase characters and symbols. When you press the Caps Lock key, a light turns on at the upper-right corner of your keyboard. When the Caps Lock light is on, the alphabet keys you type appear on the screen in capital letters. Caps Lock does not affect the number or punctuation keys.

The function of the Enter key is controlled by the program you are using. For example, a word-processing program uses the Enter key to move the cursor to the beginning of the next line. A spreadsheet program uses it to enter information into a specific cell or field. Your operating system uses it to accept a command.



Function keys


Function keys are arranged across the top of your keyboard. Use them to send instructions to the software you are using. Think of them as shortcuts for doing routine tasks on your computer.



Function key tasks can vary from one application program to another. For example, you press the F1 key in some application programs to display a help message; you might press the same key in another program to display a list of files.

Special computer keys

The keyboard has several keys that are not on a typewriter. These are Escape (Esc), Control (Ctrl), Alternate (Alt), Insert (Ins), Delete (Del), Print Screen, Scroll Lock, and Pause/Break.



I/O Features


Your software controls these keys. It can disable them or modify their purpose. The following list describes the most common uses of these keys.

Esc	<i>Escapes</i> from an operation, window, message, or program and returns you to what you were doing previously.
Ctrl and Alt	Used with other keys to perform a specific operation defined by the program in use. When two or more keys are used together to perform a certain function, they are called <i>combination keys</i> . To use combination keys, you simply hold them down at the same time. For example, pressing Ctrl+Pause/Break stops the operation of most programs. Combination keys might work differently, depending on your software.
Insert	Activates the insert mode, which enables you to <i>insert</i> information between two existing characters. When you press Insert again, you return to the replace mode and write information over existing characters.
Delete	<i>Deletes</i> a character or field.
Print Screen	Sends information on your screen to a printer. If the screen contains graphic images, they might not be printed correctly.
Scroll Lock	Determines how you move through information. When you press the Scroll Lock key, a light turns on at the upper-right corner of the keyboard. With scroll locking on, the cursor can move through the information on the screen. With scroll locking off, the cursor remains stationary, and the information on the screen moves. Not many programs use this function.
Pause/Break	Stops a program temporarily.

Some application programs use the Ctrl, Alt, and Shift keys in combination with the function and the typing keys to perform specific operations and to define additional operations. Refer to your operating system and application program documentation for information about the key assignments and for the definitions of the function keys.

Cursor keys


The Cursor keys are used to move the *cursor* (a flashing short line or small box) to various positions on the screen.



Arrow keys	Move the cursor up, down, right, or left on the screen, one space at a time. Like the standard typing keys, the arrow keys are <i>typematic</i> ; that is, a key repeats its function as long as you hold it down.
Home	Moves the cursor to the beginning of a line or document.
End	Moves the cursor to the end of a line or document.
Page Up	Moves the cursor up one screen of information.
Page Down	Moves the cursor down one screen of information.

Numeric keypad

The numeric keypad is arranged like a calculator keypad. It is useful for entering numeric information.




When the Num Lock light is on, the numbers marked on the keys are active. When the light is off, the symbols or functions marked *below* the numbers are active.

You can use the numeric keypad to:

- Activate the Num Lock key to type numeric information.
- Deactivate the Num Lock key to control cursor movement.
- Type special characters by using the numeric keys with the Alt key. This is done by holding down the Alt key while typing a numeric value on the keypad. The special characters are determined by the *code page* used by your operating system. Code pages vary by language.

Function keys for Windows interface

The 104-key keyboard provides three keys that are supported by Microsoft Windows NT Workstation, Windows 2000 Professional, Windows 98, and Windows 95.



The two keys with the Windows logo are supported by the operating system and provide system-level functionality. The third key, with an arrow pointing to a menu item, provides a context menu when you press it, although other functions might be available for different applications.

Languages

The operating system you use determines what characters can appear on your screen. Your operating system provides programs that you can use to type characters from other languages. Keyboard templates are also available to show you which keys on your keyboard correspond to the characters of other languages. Refer to your operating-system documentation for additional information.

Mouse port

Your computer has a mouse port connector on the back panel for connecting a mouse. A mouse is a pointing device that provides an easy method of moving the cursor or pointer around the screen to make selections. Not all application programs support a mouse. Graphics or object-oriented application programs and operating systems, such as Microsoft Windows NT, Windows 2000 Professional, Windows 98, and Windows 95 are much easier to use if you have a mouse.

Some mice are designed for the round mouse connector on the back of your computer, while others are designed for the D-shaped serial connector or USB connector. Either type of mouse will work, provided you have the correct device driver installed.

Monitor port

You can connect a variety of *video graphics array (VGA)* or *super video graphics array (SVGA)* monitors to your computer monitor port. Refer to “Selecting a monitor” on page 11 for information on monitors, and to “SVGA mode” on page 12 for supported video modes.

Chapter 5. Storage devices

This section contains information about some of the data-storage devices available for your computer. You can use these devices to expand the capacity of your computer.

The common types of storage devices that can be used in the computer are:

- Diskette drives
- Hard disk drives
- CD-ROM drives
- CD-RW drives
- DVD drives
- Iomega Zip drives

Each drive has a letter assigned to it. For example, if your computer has two diskette drives, one is called *drive A*, and the other is called *drive B*. If you have only one hard disk drive, it is called *drive C*. If you partition the hard disk, the partitions are named as if they were separate drives: *drive D*, *drive E*, and so on.

The type and number of devices that you can have in your computer vary by computer model. You can add internal storage devices if expansion bays are available. Refer to the user guide publication for more information on expansion space.

Diskette drives

Your computer comes with a 1.44 MB diskette drive that is designed to use only 3.5-inch diskettes. Optional diskette drives that support different sizes and capacities of diskettes are available from your IBM marketing representative or IBM reseller.

Hard disk drives and IDE

Your computer comes with a hard disk drive. The capacity of the drive varies by model. An expansion bay is provided for installing an additional drive. For information about the hard disk drive in your computer, refer to the user guide publication.

Hard disk drives have different types of interfaces, or controllers. The main function of the interface is to transmit data to and from the hard disk drive. The different interface types provide various degrees of performance and differ in the speed with which they can transfer data between the hard disk drive and other system components. The most common interface specifications used in IBM computers are discussed in the following sections.

Your computer has a PCI-to-IDE interface that enables the hard disk drive in your computer to communicate with the PCI bus. This interface includes the original Integrated Drive Electronics (IDE) interface with extensions for Ultra DMA-66. The original IDE interface specification refers to a hard disk drive with a built-in controller (as opposed to a hard disk drive with a separate controller). Two IDE connectors are provided on the system board. One connector is designated as primary and the other as secondary. Your computer system board can accommodate up to two IDE drives attached to each connector.

When you connect two hard disk drives to the same ribbon cable, the two drives could potentially send and receive data at the same time (because each drive has its own controller). To prevent the drives from

Storage devices

competing for priority, you must designate one drive as the primary (master) drive, and the other as the secondary (subordinate) drive. This is determined by switch or jumper settings on each hard disk drive. If only one hard disk drive is attached to either the primary or secondary IDE connector, it must be designated as the master drive.

CD-ROM drives

An *IDE compact disc-read only memory (CD-ROM)* drive uses a removable compact disc (CD) that stores various formats of audio and video information. A CD-ROM drive can play back or read from a CD, but it cannot write information to it. The information on the CD can be heard or viewed, but not changed or updated as it can on a diskette, hard disk, or tape.

CD-ROM technology offers quick access to enormous amounts of information. Diskettes, hard disks, and tapes are magnetically recorded. CDs are optically recorded at a much higher density. CD-ROM drives use industry-standard 12 cm (4.75 in) CDs with formatted capacities of approximately 650 MB of data. CDs are used for music, voice, animated graphics, multimedia, education, software distribution, and many other applications.

To hear music or voice from a CD, you can use the speaker built into the computer or, if your computer has an audio controller, you can connect headphones or powered speakers. For more information, see the user guide publication.

Some computer models come with an IBM-preinstalled CD-ROM drive.

CD-RW drives

CD-RW, CD-rewritable, is a technology that allows CDs to be recorded, erased, then re-recorded. CD-RW drives differ from traditional CD-ROM drives in the laser and the recordable media in the disc.

The recordable media in CD-RW discs is a compound of silver, indium, antimony and tellurium. This compound has different properties when heated to different temperatures. When a CD-RW drive plays a CD, the laser operates at its lowest intensity, below 200°C, and the compound remains stable. When the laser heats the compact disc to a temperature between 200° - 500°C, the compound returns to its original crystalline structure without data. When the laser heats the compact disc to still higher temperatures, 500°C - 700°C, the compound records data. The process of recording data is called "burning a CD".

CD-RW technology complies with ISO 9660, a file system standard called Universal Disc Format (UDF), to allocate data to a compact disc. A method called *packet filing*, sending data to the compact disc in allotments, allows you to record data to a CD-RW over several sessions. Each session produces a *Virtual Allocation Table (VAT)*, to tell the CD-RW drive where to find a particular file on a compact disc.

CD-RW drives can read traditional CDs, but many older CD players cannot read CD-RWs. Their light reflective properties are about one third that of CDs.

You can record and erase CD-RWs again and again, up to about 1 000 times. Each CD-RW holds about 640 MB of data.

DVD drives

Digital Versatile Disk (DVD) technology differs from traditional CD and even CD-RW technology in several ways. These differences are possible with refinements in focusing and targeting the laser beams used for DVDs over the lasers used for CDs.

DVD recording tracks are closer together and are not as deep as CDs, consequently the recording layer is thinner. All this condensing allows more data storage room on a side than traditional CDs.

On each side, DVDs have two recording layers. The top is translucent; the inner layer is opaque. Both layers can hold data. The data in this inner layer can be read from either direction, from the inside out as well as the outside in. When it's time to read the inner layer, the laser simply refocuses and begins reading where it is, without wasting time resetting to the beginning of the DVD.

At 0.6mm, DVD media are about one half the width of traditional CDs. DVDs are literally two disks sandwiched back-to-back to create one DVD. The bonus of this is both sides of a DVD are usable, compared to one usable side on a CD.

Depending on the format, DVDs range in storage capacity from about 2 GB up to 17.2 GB. DVD drives can play CDs but CD drives cannot read or record DVDs.

DVDs are used for movies, animated graphics, multimedia, voice, music, education, software distribution, and many other applications.

lomega Zip drives

Some computers come with an internal lomega Zip drive. A Zip drive is a high-capacity disk drive that uses a removable disk that is slightly larger and about twice as thick as a conventional 3.5-inch diskette. The Zip disk has a capacity of 100 MB of data. Zip drives have become very popular for backing up hard disks or transporting large files.

Chapter 6. Security

Security is important for protecting your computer and the data stored in it. If you use your computer in a public environment, such as an office, you might want to protect the computer and data by using one or more of the security features available to you. Not all features are available on all computers. See the user guide publication provided with your computer for more specific information.

Some of the features might include:

- A lockable cover using a customer-supplied padlock to help secure computer hardware components
- Use of a U-bolt and a lockable cable to secure the computer to a desk or table
- A power-on password to help prevent unauthorized use of the computer
- An administrator password to authorize computer setup changes
- An unattended start mode to enable the computer to operate without input from the keyboard or mouse
- A selectable startup sequence to prevent use of unauthorized programs that bypass security controls
- Selectable disabling of diskette and IDE drives to prevent unauthorized use
- I/O port control to disable input and output functions of the serial and parallel ports and their attached devices .Asset ID serial numbers on key system components: CPU, DIMMs, FDD, CDROM, HDD, and Planar. These numbers are stored in the EEPROM. If the computer is connected to a network, the serial numbers are downloaded to the network administrator.

The use of some features also depends on your operating system and network software. For detailed information about using these security features, see the user guide publication.

While the security features can offer considerable protection for your computer and data, there are limitations, depending on the operating environment. For example, these features cannot authenticate or identify the adapter that attaches your computer to a LAN, and they cannot prevent outside persons or computers from gaining access to your computer if your computer is equipped with a modem that is set to the automatic answering mode.

A major aspect of your security plan is protecting the information (data) stored in your computer. Some of the considerations of data security are:

- Backup and recovery
- Erased files
- Computer viruses

Information backup

Most of the information that is stored on a hard disk is contained in *data files*. Data files contain the information you put in the computer, such as letters and reports.

You can lose data files from your hard disk for a variety of reasons. Security violations, improper handling during a move, or computer viruses can destroy stored information. To avoid losing information, make copies of the data files by using storage devices such as diskette drives, tape drives, or Iomega Zip drives. Also, if you are connected to a LAN, you might be able to back up information to another computer on the LAN. Backup copies can save you time and money recovering files accidentally erased from a hard disk or lost when a hard disk drive fails.

Back up your hard disks on a regular basis (daily, weekly, or monthly depending on how often the information changes). Also, back up the hard disk before moving the computer. Backup procedures differ from one operating system to another.

Erased files

Erasing a file from a diskette or hard disk does not necessarily destroy the file. With the proper software tools, someone might be able to restore all or part of an erased file. An erased file, therefore, might be a security exposure if you sell your computer or give someone a diskette that once contained confidential information.

The standard operating-system format commands do not remove all information from a hard disk. Low-level format programs are used for this purpose and are available from a number of software manufacturers.

Diskettes are slightly different. An operating-system format command writes a repeated pattern over the entire surface of the diskette. As a result, any information that was on the diskette becomes unreadable.

Depending on the type of information stored, you might require additional safeguards.

Computer viruses

A computer virus is a program (or instructions hidden within a program) that can “infect” files and programs on your computer. Unlike most other programs, viruses are typically designed to spread themselves. Some viruses can display a message or cause erratic screen behavior. Others are destructive, erasing or damaging files or overloading memory and communication networks.

Viruses are difficult to detect. Many stay inactive until triggered by a specified event such as a date, command, or some other operation. Some are activated when an infected program is started a specified number of times. When the symptoms of the virus appear, it might be difficult to determine whether the problem is a hardware failure, a software problem, or a virus in action.

Viruses are generally spread inadvertently from computer to computer when programs are copied or exchanged. If you do not know where a program came from, be careful. You can take steps to protect your information by using a virus-scan program.

Chapter 7. Software

Software includes a variety of instructions, operating systems, application programs, utility programs, and tools that enable your computer to do what you want it to do.

Software layers

Several layers of software work together to make your computer work.

System programs are the basic layer of software built into every computer. They include the POST, BIOS, and Configuration/Setup Utility programs. *Power-on self-test (POST)* is a set of tests and procedures that are performed every time you turn on your computer. *Basic input/output system (BIOS)* is a layer of software that translates instructions from other layers of software into electrical signals that the computer hardware can understand. Except for a short message on your screen when you turn on your computer, you rarely have much direct involvement with BIOS. You can use the Configuration/Setup Utility program to set up your computer for different features and options.

Other software used by your computer includes *device drivers*, which support specific types of hardware devices and *Advanced Power Management (APM)* or *Advanced Configuration and Power Interface (ACPI)*, which enable your computer to conserve energy or to turn on when it is unattended.

An *operating system* is the layer of software that manages all computer operations by directing the flow of information to and from the various parts of the computer. The operating system interprets instructions from you and your application programs and directs these instructions to the BIOS. Your computer needs an operating system to *run* application programs.

Application programs are the highest layer of software. You can use application programs to do specific work, such as word processing, desktop publishing, or financial accounting. These are the programs you use to get work done on your computer. Some application programs, called *tools* or *utility programs*, are designed for very specific tasks, such as configuring the computer, using a feature, or testing the computer. And, of course, there are games to keep you entertained.

Your software might have been preinstalled for you, or you might have to install it yourself. If not preinstalled, your software package usually contains program diskettes, CDs, and instruction manuals. Refer to the information that is provided with your software for installation and operating instructions.

POST

Each time you turn on your computer, *power-on self-test (POST)* is automatically run. POST programs are stored in flash memory along with the BIOS and the Configuration/Setup Utility program. Functions of POST are:

- Test the computer hardware and subsystems
- Update the computer configuration
- Load the BIOS into system memory
- Load the operating system into system memory
- Report problems

While POST is running, the computer displays an IBM screen with several messages and prompts. A password prompt asks you for a password, if you have set a password with the Configuration/Setup Utility program. Other prompts allow you to access the Configuration/Setup Utility program or escape from

POST. When POST is successfully completed, the computer proceeds to the first screen of your operating system or application program. If problems or configuration changes are found, POST displays error codes and messages. Refer to the problem-solving section of the user guide publication for explanations of error messages and actions you can take to correct problems.

Configuration/Setup Utility program

Most computer devices (storage devices, ports, adapters, and other optional features) have programmable settings. Your computer has a special type of nonvolatile memory (called *CMOS*) where it saves these settings. This memory is kept active by a battery so that the settings are not lost when you turn off the computer. Each time you turn on the computer, a POST routine copies the configuration settings from memory to the various devices to get them ready for operation. If the battery fails, the settings in memory are lost, and the Configuration/Setup Utility program automatically reconfigures built-in devices.

When you add or remove hardware from your computer and restart the computer, the configuration settings must be updated. In most cases, the BIOS detects the change in hardware and *automatically* updates the configuration. In other cases, you might need to use the Configuration/Setup Utility program to *manually* update the configuration. Either way, the computer displays a configuration error message on the screen and gives you the option of accessing the Configuration/Setup Utility program. You must enter the Configuration/Setup Utility program and either *save* the automatic update or manually change the new configuration settings. For more information on using the Configuration/Setup Utility program and saving configuration changes, see the user guide publication.

Operating systems

An operating system is the software that controls the execution of programs, allocates hardware resources (such as memory, microprocessor time, disk space, and peripheral devices), and manages data. The operating system enables a computer to accept information from a keyboard, display it on a monitor, print it, and save it on a disk. The operating system has built-in routines for performing input and output operations, such as reading and writing disk files. When an application program gives a command to the operating system to save a file, the operating system communicates with the BIOS to transfer the file from memory to disk.

Just as the BIOS defines how the hardware works, the operating system defines how software works by establishing a set of rules that all software must follow. For example, the operating system defines how files are named, what commands can be used, and how memory is organized. Your computer is designed to use a variety of operating systems to meet many different needs.

Device drivers

Device drivers are programs that support a specific type of hardware device, such as a printer or high-resolution monitor. These programs provide instructions that allow the computer to interact with the device or take advantage of the special features of a device. Device drivers are necessary because the BIOS cannot be large enough or current enough to contain all the code that is needed to control every hardware device. If your computer has IBM-preinstalled software, the device drivers might be included with your operating system or application programs. Also, some device drivers are located on the *Product Recovery CD*. If you did not receive a *Product Recovery CD*, some device drivers might be available on a separate CD. Optional hardware features also might include diskettes that contain the device drivers that you need to make the optional features work.

Latest-level device drivers

The latest device drivers are always available on the World Wide Web.

For the latest device drivers, visit the IBM support Web page at <http://www.ibm.com/pc/support/>.

Device-driver types

Device drivers fall into two general categories: application-specific and device-specific. Application programs load application-specific device drivers into memory. These device drivers stay in memory while the application is running, and generally are cleared from memory when you exit from the application.

Device-specific drivers are loaded into memory each time you turn on the computer. Some device drivers check for the presence of a device each time you turn on the computer. If the device it supports is not attached or is not turned on, the device driver is not loaded, and the computer might generate an error message. After they are loaded, device-specific drivers stay in memory until you turn off your computer. For more information about installing device drivers, refer to your operating-system documentation and the documentation that comes with your hardware or device drivers.

Video device drivers

Video device drivers are provided on the *Product Recovery CD* that comes with your computer. If you did not receive a *Product Recovery CD*, some device drivers might be available on a separate CD. If your computer has IBM-preinstalled software, video device drivers are already installed on the hard disk of your computer. (However, note that DOS video device drivers might need to be installed for running applications that are primarily used in DOS.)

Advanced Power Management

Conserving energy might be important to you. Your computer provides energy-saving software, collectively called Advanced Power Management (APM) or Advanced Configuration and Power Interface (ACPI), depending on your operating system.

The following are the APM features that are included with your computer:

- **APM BIOS** is important if your computer uses an operating system that supports APM-capable computers. When this setting is disabled, all APM BIOS support is disabled.
- **Automatic Hardware Power Management** reduces or suspends power to components of the computer after a specified period of inactivity has passed (SCSI hard disk drives are unaffected by this feature). For more information on setting and using this feature, see the user guide publication.
- **Activity Monitor** is used to select which devices are monitored to determine whether the computer should enter or exit the low-power state. Not all operating systems support this feature. Refer to your operating system documentation to determine whether the Activity Monitor is supported.
- **Automatic Power On** enables the computer to be turned on when it is unattended. Optional hardware and software are needed for some features of Automatic Power On. For more information about Automatic Power On, see the user guide publication.

Attention: If a device, such as a monitor, does not have power management capabilities, it might be damaged when placed in a reduced-power state.

Advanced Configuration Power Interface (ACPI)

ACPI is a BIOS mode that enables the operating system to control the power management features of your computer. ACPI offers similar features to APM, but is enabled and configured by the operating system. Not all operating systems support ACPI BIOS. Refer to your operating system documentation to determine whether ACPI is supported.

Application programs

Application programs provide the tools necessary to complete a specific type of job, such as editing text or preparing a spreadsheet. A variety of programs are available to aid you in generating reports, working with numbers, using graphics, and performing many other personal and business functions.

Many application programs are commercially available for your computer. After you determine the types of application programs you need, select programs that will work with your computer. Most program packages list the computer equipment you need to run the program. Read the package carefully and look for the following information:

Computer Type	Make sure that the box is labeled for IBM or IBM-compatible computers.
Diskette Size	Your computer has a 3.5-inch diskette drive. Make sure that the package contains 3.5-inch diskettes.
CD-ROM Support	If your computer has a CD-ROM drive, you can purchase application programs that are available on CDs.
Video Support	Your computer has an SVGA graphics controller that also supports video graphics array (VGA) standards.

Protecting your software

Under normal circumstances, there is little that you can do to damage your software; however, accidents can happen. If you establish good habits when using your computer, you are less likely to experience problems.

Your computer might come with the operating system and other software already installed on the hard disk. The operating system might fail if you modify operating-system files incorrectly or if you try to install software that is not designed to operate on your computer. Back up your operating system and other software so that you can more easily recover from an operating-system or hard-disk failure. For instructions on how to back up the operating-system programs, refer to the documentation for your operating system.

Also, make backup copies of any original software diskettes. Using the backup copy can prevent damage to the original diskettes. Your operating-system documentation gives instructions for making backup copies.

Chapter 8. Networking and system management features

This section contains information about the networking features available on your computer.

A network consists of computers linked together so that they can share resources, such as files, programs, and peripheral devices. The simplest form of a network consists of two computers connected through their serial or parallel ports. This type of network might be used by an individual to transfer files between a laptop computer and a desktop computer.

Other networks are larger and typically consist of computers designated as *servers* and *workstations*. Servers provide shared resources (they store and retrieve shared data files) for computers that are linked to the network. They typically have fast microprocessors and large, high-capacity disk drives that operate unattended. In contrast, workstations, often referred to as *clients*, are computers that are used by individuals and do not provide shared resources for an entire network.

Some computers come network-ready with a preinstalled Ethernet adapter. This adapter provides a high-performance network connection with an RJ-45, unshielded twisted pair (UTP) connector. For computers with IBM-installed software, device drivers for this adapter are on the hard disk; for some adapters, device drivers are also provided on the *Product Recovery CD* that comes with your computer. If you did not receive a *Product Recovery CD*, some device drivers might be available on a separate CD. For instructions on connecting a network cable to your computer, refer to the user guide publication. For information on configuring your computer to work with a network, consult your network administrator.

When your computer is linked to a network, you can take advantage of the following features:

- Desktop Management Interface (DMI)
- Dynamic Host Configuration Protocol (DHCP) or Remote Program Load (RPL)
- Remote Administration
- Alert on LAN
- Wake on LAN®
- Wake on Modem
- Wake on Ring

Desktop Management Interface (DMI)

DMI consists of software that network administrators can use to remotely monitor and control your computer in a network environment. DMI can be used to remotely track many types of information about computers, including serial numbers, memory attributes, product-specific characteristics of installed devices, temperature, and operating-system configuration information. This information can be accessed using a DMI browser. DMI browsers are provided by all major operating systems and all major LAN-management packages (including Netfinity Manager® and LANClient Control Manager™).

DHCP and RPL

A network administrator can use Dynamic Host Configuration Protocol (DHCP) and Remote Program Load (RPL) to control your computer. The DHCP feature makes it possible for a DHCP server on your intranet² to assign an internet protocol (IP) address to your computer so that a startup image can be loaded from the server. The DHCP server must be one that supports boot protocol (BOOTP) on your intranet using software such as the Intel LANDesk® Configuration Manager. Your computer must have network management software, such as the LANClient Control Manager.

If you use RPL with software such as LANClient Control Manager, you can use a feature called *Hybrid RPL*. Using LANClient Control Manager, you set up your computer by installing hybrid images (or files) on the hard disk. Then, each time the computer starts from the network, LANClient Control Manager recognizes your computer as a Hybrid RPL client, and a *bootstrap* program is downloaded to your computer. This bootstrap program is small and helps prevent network congestion. Working from the hybrid images, the bootstrap program initiates the startup process from the hard disk of your computer. An advantage of using Hybrid RPL is that the network load associated with standard RPL is avoided.

Remote Administration

When this feature is enabled and you are connected to a network, the system programs in your computer can be updated remotely by a network server. If you have set an administrator password for your computer, it does *not* have to be entered by the server. Consult your network administrator for information on setting up your network server to perform POST/BIOS updates.

Alert on LAN

Alert on LAN is a new hardware and software technology supported on some models of your computer. It immediately alerts network administrators when a computer has a problem, even if the power is off or has no operating system running. For system management, Alert on LAN helps the network administrator to react quickly and effectively when problems occur, rather than waiting for input from the end-user.

For more information about Alert on LAN, go to <http://www.pc.ibm.com/us/desktop/alertonlan/index.html> on the World Wide Web.

Wake on LAN

Some computer models have an installed LAN adapter that is designed to use the Wake on LAN feature, an IBM-developed technology that allows your computer to be turned on remotely by a network server. When the Wake on LAN feature is used with network management software, such as Netfinity (included on the *Software Selections CD* that comes with your computer), functions such as data transfers and software and device-driver updates can be performed remotely on your computer and all other computers on your network that are enabled for Wake on LAN. To avoid interruptions and help increase productivity, these updates can be done after normal working hours or on weekends, when LAN traffic is at a minimum.

Wake on Ring

Your computer might have an external modem attached to the serial port. The Wake on Ring feature enables your computer to be turned on remotely when a modem ring is detected. In the Configuration/Setup Utility program, this is referred to as *Serial Port Ring Detect*. When Wake on Ring is used with network management software, such as Netfinity (included on the *Software Selections CD* that comes with your computer), functions such as data transfers and software and device-driver updates can be performed remotely on your computer. To avoid interruptions and help increase productivity, these updates can be done after normal working hours or on weekends, when telecommunication traffic is at a minimum.

² An intranet is a private network that conforms to the same protocols as the Internet, but is contained within an organization. The intranet contains one or more servers that provide services to the workstations on the private network. Some intranets are also connected to the Internet.

Appendix. Notices and trademarks

This section contains product notices and trademark acknowledgements.

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Subject to IBM's valid intellectual property or other legally protectable rights, any functionally equivalent product, program, or service may be used instead of the IBM product, program, or service. The evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the responsibility of the user.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Trademarks

The following terms are trademarks or registered trademarks of the IBM Corporation in the United States or other countries or both:

AT
Alert on LAN
IBM
LANClient Control Manager
Netfinity
Netfinity Manager
Wake on LAN

Intel, LANDesk, and MMX are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.



Part Number: xxxxxxx

Printed in U.S.A.

xxxxxxx

