LAPTOP COMPUTERS

A BRIEF HISTORY

Laptops and notebook computers are part of a category called mobile or portable computers. A portable computer is defined as any computer designed to be carried from one place to another. This category consists of laptop computers, notebooks, personal digital assistants (PDAs), handhelds, and the original transportables. In addition, other similar devices that have the same technologies as portables, such as pen-based tablets or notepads, are commonly lumped in the same category as portable computers.

Since the mid 1980s, portable computers have become smaller, lighter, and significantly more powerful every year to the point where portable computers now replace desktop computers as the primary computer for general purpose use in most cases. Although this article focuses primarily on the technology related to notebooks and laptops, most of the technology discussed, the tradeoffs in the design, and the features are similar across the entire spectrum of portable computing devices described earlier. Portable computer technologies continue to advance at an astronomical rate. About every 6 months or so, portable computers get closer to offering similar features to that of desktop computers.

The first portable computers were called transportables or luggables. They consisted of desktop computer components in a single plastic case with a keyboard tethered to it by a cord. Contained in this unit was a central processing unit (CPU) subsystem (the computer base), a display, the storage subsystem, and expansion slots. The display contained in this unit consisted of a standard monochrome cathode ray tube (CRT) monitor that was small enough to fit inside this ~30 to 40 lb. unit. The main computer subsystem in this class of computer design was actually the same subsystem used in traditional desktop computers of that time. Because the same components were used in both designs, both performed similarly and both used an ac power source. In addition, because of the size and weight of the unit, the portable computer was not really considered portable enough to be used by everybody as a general purpose computer. Although some variant of this type of portable computer is still found, it has changed from the general purpose, mainstream category to the special vertical market category of industrialized applications.

In the late 1980s, the next class of portable computers, the clam-shell-based laptop was developed. Advances in technologies helped pave the way for designing actual computers weighing around 12 to 15 lb. Some of the technologies that led to this technological feat were liquid crystal display (LCD) panels that replaced CRT tubes, smaller devices, such as fixed disk drives, and the first wave of semiconductor integration.

One additional and very important feature of this laptop generation was the fact that these were actually the first portable computers to operate on batteries. Even though the battery life of these computers was actually about 30 min to 1 h, it was still significant because it indicated great things to come.

Over the next few years, laptops evolved from monochrome to color displays, from floppy-disk-based units to units containing internal 2.5 in. fixed disk drives, and from units containing one battery to units containing two batteries that actually gave an operating life of at least 3 h between charges.

In the early 1990s, the next wave of portable computers, notebook computers, became a reality. Today this is the most common type of portable computer. Notebooks have advanced to the point where they contain features and components similar to those found in desktop computers, but, the technology gap has narrowed to the point where some technologies are being introduced at the same time for both market segments, such as processors and other associated silicon.

Over the next several years, portable computers will continue to advance and will begin to branch off into other areas of functionality, such as more sophisticated and smaller PDA designs, subnotebooks smaller than those existing today, handhelds, and eventually, pocket computers.

Today's notebooks have the following characteristics. They weigh between 4 and 8 lb., have LCDs that are at least 8.5 in. and go up to 14.4 in. diagonally, have processors that are at least 150 MHz and approach 300 MHz in speed, contain at least 32 Mbyte of random access memory (RAM), have full-size keyboards, built-in CD-ROMs, and fixed disk drives with capacities starting at 2.1 Gbyte. In the past, desktop computers maintained a technology lead that created a gap in which the desktop led the notebook in performance and features. This gap was at least 2 to 3 years before notebooks had as much computing power as mainstream desktops. Now this gap has narrowed to ~6 months in most cases. In some areas, technologies are being introduced at the same time, such as processors and core system chips. In other areas, notebooks have led the way. These areas include power management, plug and play, communications, and the integration of both components and subsystems.

NOTEBOOK TECHNOLOGY

The technologies used in notebook designs are similar to the architecture and technologies used in traditional desktop

J. Webster (ed.), *Wiley Encyclopedia of Electrical and Electronics Engineering.* Copyright © 1999 John Wiley & Sons, Inc.

computers. Although this article focuses primarily on notebook architectures based on the X86 architecture, some of the technologies are architecturally agnostic and cross all segments of the portable computing segment. Where this is not the case, it is pointed out in the appropriate section.

Notebooks can be subdivided into several categories each designed to serve a particular market segment. The categories consist of subnotebooks, PDAs, full-featured notebooks, and desktop replacement notebooks. Creating notebooks for a particular market segment involves significant design tradeoffs. Subnotebooks may lack an integrated flexible disk drive or CD-ROM, may have smaller LCD panels, have smaller keyboards, and use smaller technology, such as fixed disk drives, to achieve the desired size. In most cases subnotebooks generally weigh no more than 4 lb. Full-featured and desktop replacement notebooks are fully portable but have technologies that rival desktop computers and are designed for maximum performance and functionality although they still deliver the promise of portability. In some cases of desktop replacement notebooks, their claim of portability can be questioned.

The typical clam-shell notebook contains two major parts connected by hinges. The first part consists of the display head which contains the LCD and some electronics enclosed in a plastic housing. The second part of the notebook is the main body of the computer typically called the base. Each section has various features and tradeoffs that notebook designers make to optimize the notebook design for the particular market segment or to deliver a certain feature set.

Display Head

As previously stated, the display head is the part of the notebook that encompasses the LCD panel and in some cases may include additional electronics specific to the LCD. One of these is the power inverter board. LCDs require current that is much higher than that of other parts of the notebook. Most notebook display heads usually contain a board that uses the power supply voltages and currents and increases them to the amounts required to power the LCD. Because this is done only to power the display, it is typically designed and enclosed in the display head alongside the LCD when there is sufficient room. Although this is not a mandatory placement of the inverter board, it is much more practical to place it here instead of down at the base where it may interfere electrically with other components.

The LCD is the glass display output device. The most popular display types in notebooks fall into two categories. Super Twist Nematic (STN), also known as passive panels, and TFT which stands for Thin-Film Transistor display and are commonly referred to as active panels. The TFT panel is the preferred panel because of some of the characteristics of its design. Specifically, TFT panels typically have a faster response time enabling multimedia to appear more fluid to the user. In addition, the brightness and wide viewing angles available in TFT panels also make it the preferred panel for use with presentations. Lastly, the richness in colors viewed on a TFT panel is not as washed out as is typical of STN panels. As the technology for STN panels improves, however, the differences will go away. Although the TFT panel is the preferred panel, the STN panel was the primary driver for creating the demand for color notebooks because it is cheaper to manufacture than the TFT panel. The difference in cost is because the process for creating a TFT panel is more complex and expensive. When STN panels were introduced, the cost of the STN panel was around 25% of the cost of a TFT panel. In recent times, as manufacturing efficiencies and volume have gone up, the costs of TFT panels have come down considerably. However, TFT panels may still enjoy a 1.5 to 2 $\times$ price premium over comparable STN panels of the same size/features.

The main features that differentiate the two types of panels are the response time, contrast ratio, and viewing angle. Response time can be generally defined as the time it takes for the pixel to switch between on and off states. The contrast ratio gives the appearance of a brighter image through a wider contrast between on and off states and the viewing angle enables a wider view of the image being displayed. TFT panels typically have a wide viewing angle, while STN panels do not. The viewing angle on STN panels has the advantage of offering a narrower, more private view so that images can only be seen by looking directly to the screen.

TFT panels are based on transistors that are etched on the LCD glass and these transistors are powered by row and column drivers coming from the edge of the display. STN panels don't have transistors. Aside from these differences, the two display devices are similar in size, weight, and support for multiple colors.

STN and TFT panels come in both monochrome and color; however, most notebooks are now all color and the few monochrome panels that exist are predominately STN. The color STN panel is known as CSTN or color super twist. CSTN displays come in both single scan and dual scan models. The basic difference between the two is that in a single scan display, the LCD refreshes the complete screen in one refresh cycle. This cycle is typically a 1/480 duty cycle on displays with 640 $\times$ 480 resolution, which means that the screen is refreshed from the top of the screen to the bottom in one cycle. This duty cycle results in a display that is slow and has the effect of sacrificing the quality of animations and brightness. Single scan displays do not have the quality of other color panels and hence are not too common in notebooks except for really cheap models or in displays smaller than 7.5″ diagonally. These displays are now moving to other types of applications besides notebooks, such as PDAs or handheld devices. The more common CSTN panel used in portable computers is the dual scan CSTN or DTSN display.

On a dual scan display, the display is divided horizontally into two sections that are refreshed independently using odd and even frames. Each section is refreshed at approximately half the full size of the display. For example, in an LCD screen with a 640 $\times$ 480 display resolution, each half is refreshed with a 1/240 cycle. On screens with 800 $\times$ 600 resolution the independent halves are refreshed with the odd frame using a 1/300 duty cycle and the even frame using a 1/328 duty cycle. This has the effect of creating a faster responding display with a picture quality that comes close to that of TFT panels because the screen is refreshed much faster and because the transistors are kept alive longer. Dual scan CSTN was the primary market driver in the notebook marketplace that made it possible for color notebooks to become popular and affordable since the dual scan display is cheaper to manufacture than a TFT display. To give an example, prior to the introduction of the CSTN display, color notebooks cost an average of $5000 for notebooks that contained a TFT panel.

Although these two display types are the most common in notebooks, newer display technologies are being developed that may become standard in portable computers. These displays include plastic-based displays and displays based on other technologies that have a faster response and lower manufacturing cost such as reflective displays that use the light that reflects off the display to eliminate the backlight tube.

One final electronic component that may be found in the display head is a converter board. LCD panels generally have proprietary interface signals that drive the display. These signals differ from one manufacturer to another and from one type of display to another. To accommodate support for a variety of display devices, a board that translates the signals coming from the LCD controller on the computer mainboard into the appropriate signals of the LCD panel is usually designed in the display head. Recent industry initiatives, such as Panelink, have led to standardization of the signals across all of the display types.

Base Subsystem

The lower half of the notebook is the base or CPU subsystem. The base of the computer consists of two pieces of plastic that encompass the mainboard, keyboard, storage subsystem, and expansion ports. In most computer designs, all of the electronics are placed on a single board known as a mainboard or motherboard. This is not too different in notebooks, except that physical design tradeoffs driven by the mechanical design of the notebook may force the electronics to be partitioned onto multiple boards stacked on each other through connectors. For the rest of this section, we assume that the notebook contains one mainboard which contains all of the electronics. Unlike desktops that follow a specific form factor, notebook mainboards are proprietary to the manufacturer and product line.

Mainboard Electronics

Mainboards contain the CPU, the system memory soldered onto the mainboard, the graphics subsystem which drives the LCD, the core logic which controls the memory and peripherals, and connectors used for expansion along with additional necessary circuitry. Figure 1 is a block diagram of a typical current notebook design. Because architectures and technologies are still evolving at a fast rate, this block diagram provides an example of the most common notebook architecture based on the Socket 7 architecture at the time this article is being written.

Most notebook computers today are based on a 32 bit CPU conforming to the x86 architecture. There are two major types of x86 processors. The first one is “socket 7” class processors, such as the Intel Corporation Pentium CPU or the AMD Corporation K6 CPU. Although these CPUs are known as Socket 7 compatible, the CPUs come in various packages designed specifically for different mobile form factors. The desktop CPU is packaged in a pin grid array (PGA) package, which is either a ceramic or plastic package around the die of the processor with 296 pins coming from the bottom of the package. Even though some of the notebooks available are based on this package, the predominant CPU package for notebook designs is the tape carrier package (TCP). In the TCP package, the die is mounted on a small substrate and has leads coming from the side. Then this CPU is mounted either on a small circuit board with pins that is installed on a socket on the motherboard or is soldered directly to the motherboard. In most cases the circuit board option or module is used because it provides flexibility for the computer manufacturer. During manufacturing, it is essential to be able to vary the configuration of a notebook until the last point of the manufacturing process. This flexibility is achieved by using a module because on some of these modules, the designers place other components in addition to the CPU that are specific to the type of CPU to create a standard interface which isolates the mainboard from this module. Some of the components include the voltage regulator, configuration strap settings, the north side core logic, and other necessary components which are placed on this module. Because different CPUs require different voltages and speed settings, putting the other components on the module enables standardizing the pins coming from the module and this means that the interface on the motherboard can be the same, independent of the CPU speed/voltage required by the CPU. An example of this module is the Intel-designed mobile module.

The second class of processors used in notebooks is based on the Pentium II class, also created by Intel. These processors are similar to the Pentium class, but offer the advantage of faster performance and the ability to directly interface to cache memory without having to go through the chipset to do this. Systems developed on these types of processors are typically based on the mobile module or other customer-specific module. The CPU is connected to chips on the mainboard called core logic chip sets. These chip sets are responsible for handling tasks, such as arbitrating memory between the CPU and other devices, generating the various system buses, such as the fixed disk (hard disk) interface, expansion buses such as the industry standard architecture (ISA) and peripheral component interconnect (PCI) bus or other standard computing activities, such as steering interrupts and handling direct memory architecture (DMA) events. Core logic chips are typically divided into north side and south side components where the north side is responsible for communicating with the CPU, the memory subsystem, the south side counterpart and in recent architectures, the graphics subsystem. The south side part is typically responsible for generating the ISA bus and other legacy buses, and is also responsible for waking up the system through interrupts. Connected to the south side part are devices, such as graphics controllers, network controllers, keyboard controllers, and serial and parallel controllers.

Graphics controllers for notebooks are similar to desktop graphic controllers but add additional support for the digital interface of the LCD. Graphics controller technology trails desktop technology in terms of graphics capabilities, such as speed and resolution, but lead their desktop counterparts in areas of memory integration and power consumption. In addition, graphics controller display output to both a monitor and the LCD at the same time, which makes them useful for presentations.

Attached to the south side core logic are controllers that have the functionality of several controllers and have been combined into a single chip called a Super I/O chip. This chip is typically placed on the ISA bus. This chip contains the keyboard, the serial and parallel functionality that is used to connect to input and pointing devices, modems, and printers, respectively. Additionally, other devices, such as graphic

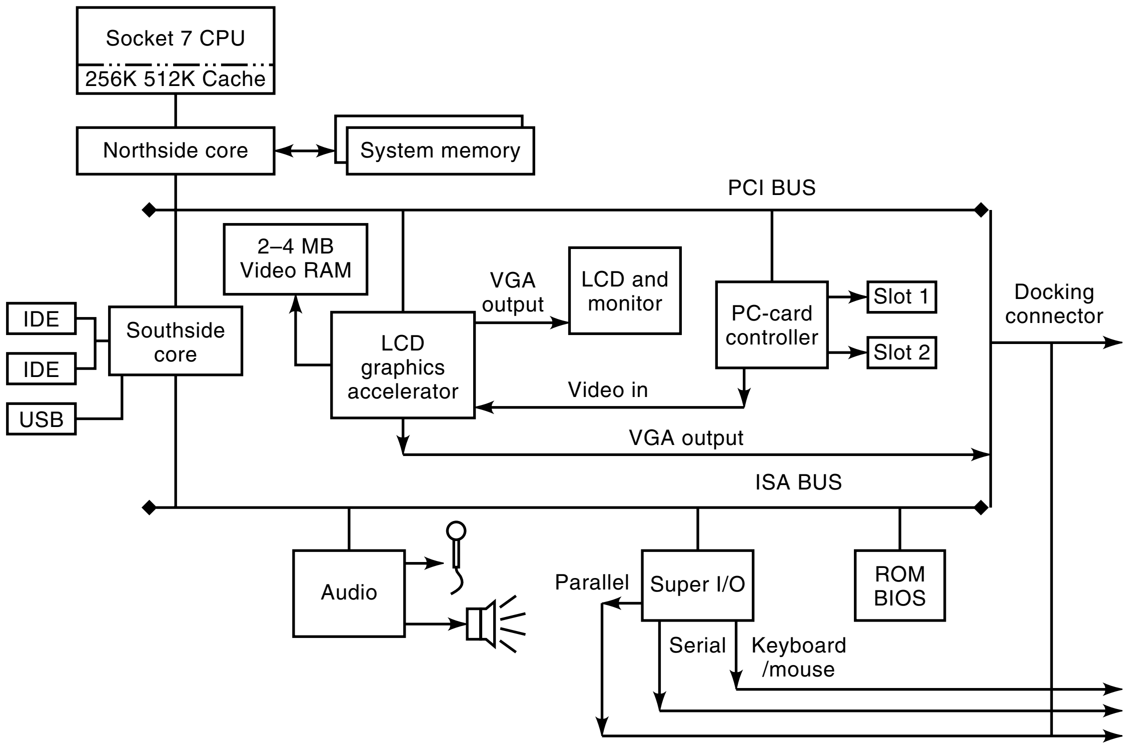

Figure 1. Power point slide.

devices, network interfaces, sound controllers, and expansion controllers, could be either on the PCI or ISA bus. The trend in the industry is to continue reducing the number of chips required for a design. In the near future, all of the functionality of the Super I/O chip will be replaced with other serial-based functionality, such as the universal serial bus (USB) and other fast serial type buses, and will be integrated directly in the core logic chip set.

Since the early 1990s, portable computers were designed with an additional controller called the PCMCIA controller. PCMCIA is a computer industry acronym for Personal Computer Memory Card Interface Association. Connected to this controller are expansion slots that add additional functionality to portables with expansion cards similar to add-in cards available for desktops. Now these expansion cards are called PC-Cards and the interface is called the PC-Card interface because the term PCMCIA was too difficult to memorize and the interface is more than just a memory card interface. Expansion cards are available in three different form factors. These form factors are defined as Type I, Type II, and Type III cards.

Type I cards are used for memory cards, such as flash ram or SRAM. Flash cards are nonvolatile memory cards, and SRAM stands for Static RAM cards. These memory cards are used for storage and memory expansion. Type II cards are the most prevalent and consist of I/O expansion cards, such as sound, disk controller, LAN, modem telephony, and other peripheral type cards. Because of their expanded size, type III cards are typically used for hard disks and other specialty purpose cards, such as wireless modems.

Most notebooks are designed to accommodate either two type II cards or a single type III card. In some cases, the Type II slots on thicker notebooks are stacked to provide support for the type III card.

The PC-Card interface was the first interface designed for a computer that really delivered on the promise of true plug and play. All of the cards installed inside the notebook would be automatically configured for interrupt, DMA, and other memory address usage. The PC-Card interface gives the notebook user the ability to support a variety of additional add-in functionality to notebooks.

Additional components found in most notebooks are microcontrollers that perform several additional functions in the notebook not typically covered by the core logic chip set. One of these functions is to control the power management of the notebook. One of the greatest features and biggest challenges in notebook designs is the ability to control the power of the devices to maximize either the battery life or the performance of the notebook. To perform faster, the notebook uses up more energy. To last longer, the notebook needs to slow down or reduce the amount of power that devices consume. Some examples of power management are starting and stopping the CPU while you are typing on the keyboard. Others are shutting down devices, such as the hard disk, the LCD screen, or the system clocks, after certain periods of inactivity or nonuse. At any time there are devices on and devices off in a notebook, and the user most likely cannot tell the difference if the design is correct. Most notebook designs also enable users to select and configure options to best suit their needs.

An additional function of the microcontroller is to control the charging rate of the battery in the notebook. Battery technologies have evolved at a fast pace to keep up with the advances in notebook design, but some of the biggest complaints about notebooks is that the battery life is insufficient for most needs. As battery technologies improved, the demands of the notebook devices changed considerably. When the first color panels were introduced, the panel consumed around 8 W of power for the backlights and the panel. The CPU consumed around 2 W and the rest of the electronics used around 7 to 8 W. Now, the LCD panels use less than 2 W, the CPU around 8 W, and the rest of the basic electronics use about 5 W. But additional devices, such as CD-ROMs, bigger hard disk drives, internal floppy disks, PCMCIA slots, built-in modems and audio subsystems, designed into notebooks have increased the power requirements considerably.

Common battery types include nickel cadmium (NiCad), and nickel metal hydride (NiMH). Newer battery technologies are based on lithium ion and soon will be based on lithium ion polymer. All of these battery technologies have different characteristics for charging and discharging. This requires that the microcontroller be intimate with the characteristics of the battery cells and that the charging circuitry in the notebook charge and discharge the battery at a rate that is safe for the battery and the user. Charging a battery too fast can damage the battery, reduce its life and, even cause an explosion of the battery pack.

In the past, the power management controller was a separate microcontroller, but this functionality can also be embedded into the core logic chip set or inside the keyboard controller. The controller is also used to control keyboard function hot keys so that the system performs tasks, such as shutting off devices, going to sleep, and switching display output from LCD to CRT and back.

Storage Subsystem

Notebook computers support internal devices, such as fixed disk drives, floppy disk drives, and CD-ROM drives. The fixed disk drives and CD-ROMs are connected to an interface called ATAPI or enhanced IDE interface. This interface is similar to the interface on desktop computers, but some of the signals have been changed mechanically and electrically to support the smaller profile of the disk drive. Disk drives come in two standard form factors. Both are 2.5 inch but have different height requirements. The smaller drive is 11.7 mm thick and is used in subnotebooks or ultrathin portables. The larger 19 mm drive is used in thicker, fuller featured notebooks with larger capacity requirements. The disk drives contain between one and four platters and two to eight heads. The more platters and heads, the bigger the capacity of the drive. One of the major obstacles in creating notebooks is trying to fit all of the components and maintain the notebook's light weight. Board and peripheral stack-up and layout determine how big or thick the notebook will be. Disk drives and the other peripherals play a major role in notebook size.

Keyboard and Pointing Device

Besides the LCD, the other device that defines the overall width of the notebook is the integrated keyboard and pointing device. Just as the LCD has to provide a good output, the keyboard has to provide a very acceptable input to the notebook for the notebook to be accepted. There are several major parts of the keyboard that define how good a keyboard is for typists. Although the keyboard is not as large as a traditional desktop keyboard, it is possible to squeeze all of the functionality of the typical 101 key desktop keyboard into an 84 key keyboard. This is done by assigning multiple functions to a single key and having a special key called Function, labeled Fn, that enables the user to use the other keys. For example, in a desktop keyboard, the numeric pad is separate from the cursor keys. In a notebook keyboard, the numeric pad is integrated inside the other alphanumeric keys.

The keys on a desktop keyboard travel about 22 mm to 24 mm, which is the distance the undepressed key travels until the key hits the bottom of the keyboard to make contact. On a notebook keyboard, the travel is around 19 mm which still gives the user a reasonable amount of "kick" or feedback to indicate that the key was pressed. The keys are also closer together than on the traditional desktop.

The pointing device is included on either the keyboard or the top base of the unit. Several pointing devices are commonly used in notebook computers. There is no clear-cut answer as to which pointing device is better. All the three pointing devices are equally popular. The first device is the trackball, a round ball that the hand rolls to move the cursor on the screen. Trackballs were very popular and come as close as possible to giving the user the same amount of control as a mouse. However, because the ball is exposed, it is much more open to debris from external objects, such as oils from the hand or food, which make the ball unresponsive. The trackball requires regular cleaning for it to function correctly. Trackballs range in size from 11 mm to 24 mm.

Another typical pointing device is the finger stick located between the G, H, and B keys. This stick is similar to a joystick in that moving it pivotally moves the cursor. The pointing stick is commonly viewed as an alternative to the trackball and takes the least amount of space.

The final pointing device available on notebooks is the touch pad. This pad is a small ~2 inch square, flat surface below the space bar that is used to control the mouse with the fingers. The track pad contains sensors underneath the pad that detect when the fingers rub against the pad by measuring the disturbance caused by the electrical current in the finger.

Docking Station Support

Most typical notebook computers can also be expanded by connecting them to another device called a docking station. This expansion comes by way of a single connector that contains upward of 100 pins. To facilitate docking with a docking station, all of the important signals, such as the peripheral signals, monitor and power signals, and bus interface signals, are brought out to a convenient single connector. This connector, called the docking connector, connects the notebook to various types of docking stations. Although docking stations vary from manufacturer to manufacturer and in product families, there are two basic types of docking stations, port replicators and full docks. A port replicator, as the name implies, is a quick connect device that enables the user to quickly connect to the standard array of devices in the office, such as monitor, printer, LAN connection, and power. Some port replicators also include additional PC-Card slots or built-in network devices inside of them. This enables the user to come to the office and quickly attach to the docking station and the peripherals with one connection and fully use the peripherals. Full docks contain all of the functionality of a port replicator and provide additional ISA or PCI expansion slots that can accommodate standard desktop expansion cards and device bays so that peripherals, such as desktop hard disks and CD-ROMS, can be used. In the past, docking stations were specific to the brand and specific model of a notebook computer. However, several factors have led computer manufacturers to create standard docking stations for a particular family and in some cases across their entire mobile product portfolio. These factors include development and support costs, shorter life cycles of the notebook computers (life cycle is defined as the time from which a model is released to the market until it is replaced by a newer model) and support for compatibility with other models, as requested by customers for investment protection.

The combination of a powerful notebook and a full dock comes close to giving the user the typical power of a desktop. This notebook and full dock combination is called a desktop replacement. There are several ways that a notebook docks to a docking station. These methods include Cold Docking, Warm Docking, and Hot Docking. Cold Docking is when a notebook is completely off and is connected to the docking station before the machine is turned on and the operating system is loaded. This was the original method used with notebooks to dock. In order to undock, the system was completely turned off and then ejected from the docking station. Warm Docking is when the system is placed in a sleep state with some of the peripherals and buses turned off and then inserted into the docking station. Hot Docking is when the machine is fully powered up and then inserted into the docking station. This power-up model enables the machine to automatically configure itself and does not require the system to restart and take time to load up.

The same is true with trying to undock the system. The methods to undock follow a similar fashion but are more difficult because if you try to undock without telling it that you are going to do this, then the operating system can crash on you. All of these modes are detected by specific circuitry in both the notebook and docking station that knows which state the machine is in or detects a state change. The operating system then configures and responds accordingly.

Software Support

Software is as much a part of the notebook technology as the hardware. Software has played a pivotal role in redefining the architecture of notebook and laptop devices and also has significantly affected desktop architecture. To give an example, software developed for notebooks to perform tasks, such as putting devices in rest mode or a low-power state, has helped redefine how standard CPUs and core logic devices are designed for desktops and notebooks. Now processors and core logic chip sets support special interrupt service routines to enable power management of the hardware regardless of the operating system in use. Now that these routines are embedded in the hardware, they can be used to perform additional tasks, such as docking/undocking, and most recently are being used to perform virtualization of hardware devices, such as graphics and audio chips.

Firmware

Inside a typical notebook are multiple locations where firmware is embedded in ROM chips. One of these is the standard ROM BIOS that is included in any x86 computer. The second set of firmware is in either a keyboard controller or a special microcontroller. The firmware inside of the ROM BIOS is responsible for initializing devices and for communicating with the operating system and applications. The firmware inside the microcontroller is responsible for controlling the low-level functions of the notebook, such as turning devices on and off and charging the battery pack. This microcontroller performs these functions independently of the operating system used. Because some of these functions are really critical to notebook operation, the microcontroller is usually left on and is always doing these tasks even if the notebook is off or in sleep mode.

In the ROM BIOS is a special section of code that is run only when the system enters the systems management mode (SMM). SMM is a special mode designed in the computer where the CPU starts executing instructions from a specific memory location. This systems management interrupt has been designed into the CPU and the core logic chip set to enable performing tasks at the appropriate time without the aid of the operating system. A good example of this is when a fixed disk drive is automatically turned off after certain periods of inactivity by the user or application. The operating system or application does not know whether the hard disk is on or off. When the hard disk is accessed, the instructions are kept in local memory, the systems management mode is triggered, the power plane or device is turned on, and the command is re-sent to the hard disk all within milliseconds. SMM, also called SMI, is a very high priority interrupt service routine. When a CPU SMI mode is triggered, the CPU automatically jumps to a reserved location in memory and starts executing code there. When the code is processed, the system returns back to the mode it was in and continues to execute commands where it left off. This code is located in ROM at a very high address inaccessible by the operating system.

All standard notebook computers based on the x86 architecture support this SMM. It is also commonly used to support docking during a warm or hot dock operation. The SMI mode helps the operating system know when docking is about to happen. In a typical notebook, the computer contains a certain lever or metal tab that makes contact with the docking station right before the docking connector meets its counterpart on the docking side. When the computer senses this connection, the computer generates an SMI mode call to the processor. At this time, the bus is floated so as not to damage the computer. The docking connection is made and then the system returns to normal use, detects that a docking event occurred and then continues to configure the new peripherals attached to the docking station. This is a typical example of docking, but there may be other examples or methods used by other manufacturers to accomplish the same result.

One of the inventions that enabled portable computers to run off a battery was the concept of power management. Although notebooks use lower power devices than desktop computers, notebooks have to use aggressive power management modes to reduce the power consumption of a device when the device is not in use so as to maximize battery life. In addition to managing the power of individual devices and various power planes on the mainboard, power management enables the computer to enter various power consumption modes during normal use. A notebook computer designed correctly has at least two modes of operation. In one mode, when the notebook computer is docked, the computer keeps the devices at full power to deliver maximum performance, similar to a desktop computer. When the notebook is powered by the battery, the notebook is constantly changing the power states of devices by turning them on and off, as needed, to maximize battery life. Different configurations can be selected by customers to vary the settings for their particular usage models. In general, the more aggressive the power management, the slower the computer is because devices have to be powered up more often.

The computer industry has developed a series of specifications that standardize the various power modes so that applications, the operating system, and system firmware can be developed to support all notebook computers. The first specification is defined as Advanced Power Mode (APM). It is generally available on the Internet. This specification was primarily a firmware/ BIOS specification with some limited operating system support. Although this method proved adequate for most types of usage models, it is not the most ideal implementation of power management. When the operating system does not control the power management, applications cannot reliably determine the power state of the device. Applications might lock up the machine, or the machine and devices might go to sleep in the middle of the application.

In the latest trends, most of the power management is done by the operating system, except for high priority areas, such as battery charging and discharging which are always done in firmware. The newest specification is called Advanced Configuration and Power Interface (ACPI). This specification defines how the operating system performs power management of the system and its resources. It also defines how the operating system configures the resources through Plug and Play. This new method enables applications to ask the operating system if a device is away or asleep and to tell the operating system when a device is in use so that it does not lose power. For example, when a fax application is preparing to send a fax, the operating system keeps the modem on and ready to send the fax.

Now all of this functionality has become common in all portable computing devices and has also become standard in the desktop world. Portable computer technology has led the way for the so-called Green PC which consumes less than 30 W of power and goes to sleep when not in use.

Plug and Play Support

Plug and Play is an industry definition describing the concept of adding a device to a computer and expecting the computer to configure the device automatically and start using it. Notebook computers were the first devices to offer this capability in the form of the PC-Card interface and add-in cards. Notebooks have a set of device drivers called Card and Socket services that communicate with the PC-Card sockets and the cards installed in them. Socket services communicate with the controller inside of the system, and Card services configure the cards. In a typical notebook, all of the devices are configured before the operating system load, except for the cards inside a notebook PC-Card slot. When the operating system loads up, it continues to initialize and configure the rest of the devices, including the cards and sockets. When a person removes the card, the operating system releases the resources assigned to that card. When another card is inserted, the free resources are reassigned. This same process is used when a notebook docks. It automatically configures the integral and new devices located in the docking station during docking.

CONCLUSION

It is easy to see that notebook and laptop technologies have evolved considerably during the past 10 years. From an offshoot, second thought derivative of desktops to driving key technologies in areas of reduced power consumption, dynamic configuration, miniaturization, and powerful communications. In the future, portable technologies will be used in every aspect of computing from handheld and pocket computers to standard desktop PCs, from appliances, such as set top boxes that wake up at night to download the next day's programs to other devices that have yet to be invented.

The next wave of notebook computers will continue to improve technologies to make the portable computer more like a full function desktop computer. The major improvements scheduled to appear over the next few years will be in processing power, storage, communications, and mechanical size. In computing power, processors will move to the sixth generation CPU architecture, such as the Pentium II CPU, and graphics will move from the PCI bus to the new Advanced Graphic Port (AGP) bus. The AGP bus, a new bus defined by Intel, consists of a single point to point bus located on the north side core logic chip set. This new bus, which has become standard in the desktop world allows 3-D graphics accelerators to use system memory for storing textures and images for 3-D processing.

The storage subsystem will increase and storage capacities will exceed 10 Gbyte to 20 Gbyte of storage space. New communications capabilities will enable people to connect to the network completely wirelessly. In addition, miniaturization will continue to enable designers to address the need for smaller and lighter portable computers. Lastly, any advancements in desktop architectures will migrate to the notebooks of tomorrow when the technology becomes small enough and can be powered by a computer running off a battery without cannibalizing the battery life of the computer.

RAFAEL HEREDIA