Chapter 2
Power Management History and Motivation

In the early days of personal computing, under the DOS and CP/M operating systems, there was no power management—computers either used 100 percent of their power requirements or were switched off. Personal computer power management history dates back at least to1989, when Intel shipped processors with technology to allow the CPU to slow down, suspend, or shut down part or all of the system platform, or even the CPU itself, to preserve and extend battery life. The increasing importance of power management reflects the increasing number of PCs in use and their transition to tools that go wherever people go.

Three interlocked tracks guided the development of PC power management—one for power management specifications, one for system design specifications, and one for application specifications. Figure 2.1 shows the timeline for key power management and system design specifications; application program power management specifications are defined by the underlying ACPI operating system. You'll learn later in this chapter how problems in early power management technology, particularly the *Advanced Power Management (APM) Specification*, directly influenced the design of the Peripheral Component Interconnect (PCI) and ACPI specifications.

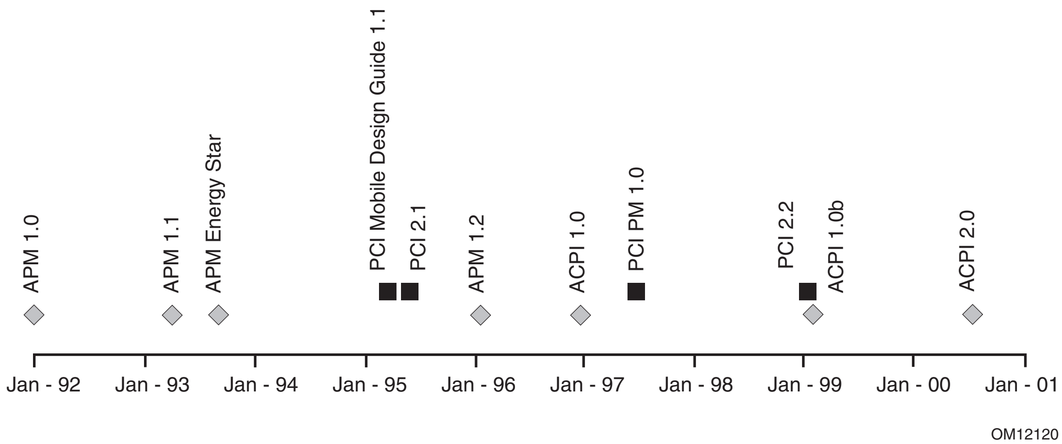

Figure 2.1 PC Power Management Specification Timeline

Responding to Requirements

Mobile computers started the drive for PC power management. Early attempts, including the Osborne computer and "luggable" PCs with cathode ray tubes, drew so much power they had no alternative but to plug into external power sources. The development of low cost, reliable liquid crystal displays (LCDs) made battery powered laptops possible in the late 1980s, and once hardware technology crossed that line, a sequence of hardware and software improvements began that have combined to increase the performance and on-battery lifetime of laptops.

The monumental number of PCs operating worldwide creates other requirements for PC power management. Because there are hundreds of millions of PCs in operation, the installed base of computers worldwide consumes tens of gigawatts for every hour of operation. Even small changes in average desktop computer power consumption can, on the global scale, save as much power as generated by a small power plant.

Mobile and Battery Powered Computers

Table 2.1 shows the historical mobile platform power road map, detailing the ever-increasing power required for increased performance in all subsystems. Mobile computers use the least power of any PC platform, both to extend battery life and reduce heat dissipation. However, the increasing power demands from the constant drive to increase mobile computing performance to near-desktop levels mean even these limits will go down over time.

Table 2.1 Historical Platform Power Road Map

	1999 Power (Watts)		2000 Power (Watts)
	3D Game	**MPEG 2 Movie**	**3D Game**
CPU & L2 Cache	9.5	7.9	9.5
Memory Controller	1.2	0.9	1.6
System Memory	1.4	1.4	1.3
Graphics Subsystem	2.4	2.4	2.4
I/O Subsystem	0.5	0.6	0.6
Audio	1.5	1.5	1.6
Modem			0.4
Hard Drive	0.7	0.0	1.3
DVD Drive	1.4	3.0	1.4
1394 Controller	0.0	0.0	0.0
USB	0.0	0.0	0.0
CardBus	0.1	0.1	0.2
LAN	0.4	0.4	0.4
Power Supply	2.0	2.0	2.6
Charging	0.1	1.0	0.0
Cooling	0.5	0.5	0.5
Other	0.3	0.3	1.0
TOTAL	22.0	22.0	24.8

Energy Star Guidelines

The United States Environmental Protection Agency (EPA) publishes Energy Star guidelines that suggest ways to reduce power consumption. Addressing the power consumed by the many computers in use, the EPA notes:

> A typical computer, monitor, printer, fax machine, and medium speed copier cost about $185 annually to operate, not including the cost of paper. Operating similar equipment that meets the Energy Star criteria costs about $97. Turning off the equipment at night can cut the annual energy costs even further.

The Energy Star label identifies energy efficient products that save money by eliminating unnecessary energy use. Energy-efficient office equipment saves energy by powering down and going to sleep during periods of inactivity, but retains all the performance features of standard equipment. On a monthly or annual basis, this equipment uses about half as much energy as standard equipment, saving owners millions of dollars in electricity costs.

The EPA is interested in power management for another important reason: using power management in office equipment offers huge potential for dramatically cutting air pollution associated with electricity use.

The Energy Star guidelines also have been adopted outside the United States. Although power costs overall have declined in the U.S., they have increased in other parts of the world. Increased power costs underscore the benefits of using Energy Star qualified products.

The Energy Star guidelines introduced in 1993 required computers to consume less than 30 watts when in their reduced-power state. Revised guidelines effective as of January 1, 1998 imposed the requirements in Table 2.2, making it possible for workstations to qualify for Energy Star compliance..

Table 2.2 1998 Energy Star Guidelines

Equipment	Specifications for Energy Star Label
Computers	Automatically enter a low-power "sleep" mode after a period of inactivity. Efficiency specifications based on the maximum continuous output power rating of the power supply. (Many computers use a 200-watt power supply for which the specification is 30 watts or less.)
Monitors	Automatically enter two successive low-power modes of 15 watts and 8 watts or less after 15-30 minutes of inactivity.

A clarification to the standard published in April 2000 notes that the operating system is typically a key component for entering and recovering from sleep modes, and requires that all necessary software to enable sleep modes be provided by manufacturers. The Energy Star web site (http://www.energystar.gov) provides the latest updated specifications.

The Energy Star guidelines for computers require sleep state power consumption of 15 watts or less. In addition, computers must be able to maintain communication and wake up on demand.

Hardware Improvements and Advanced Power Management

Energy Star is a set of voluntary energy-efficiency guidelines, not a technology or design guide. The computer industry first introduced the Advanced Power Management technology in 1992 to reduce power consumption below the 60 to 80 watt requirements of DOS-based systems. APM was invented at Intel as a power management technology for mobile systems. The APM 1.0 specification was written by Intel, Microsoft, and IBM with the goal of coordination between BIOS-driven power management and the operating system to prevent unpredictable behavior caused by the two operating independently. Desktops used APM primarily to allow the operating system to reduce CPU power when idle. The APM interface passed events between the operating system and the BIOS, such as the BIOS telling the operating system events such as "the platform is going to suspend now," or the operating system posting a "suspend request event." Later versions of APM added handshaking, but the underlying goal remained to coordinate power management actions between the operating system and the platform.

APM defined five power states:

- Full on.
- APM enabled—the processor clock can slow or stop.
- APM standby—the PC is idle, so the processor and device power states are reduced. Clocks may be stopped, but return to full operation is fast.
- APM suspend—the PC is idle, so the power level is reduced to minimum. More time may be required to resume operation.
- Power off.

APM provided calls into the BIOS so application software could interoperate with APM to provide the operating system with some knowledge of device power requirements according to the application needs.

The most significant hardware addition giving the BIOS control over power consumption was the processor Stop Clock pin. This pin, when asserted, causes the processor to stop execution on the next instruction boundary, stop prefetch, empty all the internal pipelines and write buffers, terminate bus cycles, and stop the internal clock. The processor logic effectively terminates operation at that point, drastically reducing power consumption.

Microsoft first delivered APM capabilities in Windows† 3.11. APM version 1.0 was targeted at notebook computers to extend battery operation. APM version 1.1 enhanced the original specification for use with Windows 95; APM version 1.2 incorporated features for power management on desktop computers.

APM BIOS implementations watched all the standard interfaces on motherboards for activity, including the keyboard and mouse, but lacked the ability to monitor software or key interfaces not on motherboards. For example:

- APM could not determine if a communications port was in use by software, and therefore could not determine if a lack of transmitted characters was because the port was idle or because there was simply no traffic being exchanged.
- APM could not determine what the user was doing with the computer, leading to endless numbers of presentations interrupted by screen savers.

APM defined no standard user interface, and no application-level user interface. Users commonly had to reboot their computers to invoke the

BIOS setup program to change power management settings on APM-based computers, a dangerous procedure for naive users. The means to change power settings was invariably different from one BIOS and computer manufacturer to the next, leading to further user confusion.

In the end, APM was too simple, lacking the ability to make state changes robustly and reliably. Powering down systems in the middle of communication sessions left computers in a state unable to respond to external events such as an incoming LAN packet or phone call. After APM version 1.0, many features such as notifying the operating system before suspend and control of CPU throttling when idle, were optional, and there was no standard way to describe a system's power management capabilities. The mixed set of features available on an APM system made it hard to get the APM driver right. Failure to notify the operating system or application software of major state changes left systems unable to prevent inappropriate state changes, and unable to recover once full operation was restored. The resulting problems, along with a number of implementations that simply failed to work, led most users to disable APM functions.

Advanced Configuration and Power Interface Specification

Hindsight being what it is, APM taught engineers what they needed to know to get power management right, and lead to the *Advanced Configuration and Power Interface (ACPI) Specification*. Version 1.0 of the ACPI specification was released in 1997, and offers four key improvements over APM:

- All components in the system can be power managed.
- Power management decisions are made by the operating system.
- No executable power management code is required in the BIOS.
- There is a standard way to describe the power management capabilities of a system.

The combination of these improvements gives an ACPI operating system the ability to reduce system power consumption below that of APM for both reduced power and sleep modes, and to do so with robust performance so users will actually run the technology. The combination of technologies in ACPI gives system designers the ability to reduce suspended state power consumption to less than 3 watts while maintaining communication capabilities. Power can be reduced to almost zero when hibernation or the suspend-to-disk state is used.

The first ACPI system was demonstrated in 1997, followed by prototypes from Siemens, Fujitsu, NEC, and other manufacturers over the next two years. Essentially all PCs manufactured today implement ACPI.

Getting the System Configuration Right

The inclusion of configuration functions in ACPI supports reliable power management, because it provides a well defined way to discover and control all the components in the system. For example, using the configuration data supplied by ACPI, the system knows which controllable power supplies and power planes must be enabled for each device to work. ACPI configuration information lets all unused devices be placed in a low power state or switched off completely, and makes that possible through an open standard. Without a system configuration standard, system designers had to provide proprietary configuration models for power management, and had to repeat the configuration work for each different system configuration. ACPI also allows the operating system to discover legacy and on-board devices, and their configuration, without using time-consuming and error-prone techniques like probing.

Making the Right Decisions at the Right Time

The premise designed into APM, i.e., that the BIOS could make power management decisions based solely on watching the hardware, was naive because people use computers in ways that don't involve constant interaction with the hardware. The application software must understand the functions the user is running, and only the operating system has the perspective to collect that data from multiple running programs and combine it with knowledge of what devices are active and how they are being used.

By combining the understanding that a presentation graphics program is delivering a slide show with knowledge that the external video output port is enabled and batteries are supplying power, a laptop computer could slow the processor and dim the LCD illumination to conserve power, and disable the screen saver at the same time.

Enabling Robust Power Management Implementations

Perhaps the worst problem with APM is that it was unreliable, because both the BIOS and operating system had control of devices, and each changed states of those devices without notification to the other one.

ACPI solves the problem of requiring BIOS and system manufacturers to deliver robust power management code by moving the executable power management code out of the BIOS and into the operating system. All power management interfaces in the BIOS are implemented through data tables in ACPI, and although there may be interpretable command sequences in the data tables, code in the operating system has the responsibility to execute those sequences safely. The ACPI interface for controlling power management is not what it was in APM—power management functionality resides in the operating system, explicitly supported by the BIOS and device drivers.

Creating New Power Management Opportunities

The reduced standby power levels made possible by ACPI create new opportunities for how computers operate. With standby power levels of just a few watts, it's reasonable to leave computers on all the time, having them remain asleep until they have work to do. ACPI makes it possible to let LAN adapters, modems, and other devices wake the computer from its sleep state, and extends the function of the power switch to command either power off or sleep states.

Hardware implementations are keeping up with the new opportunities created by ACPI. At first, implementing wake on LAN functionality required an additional cable from the motherboard to a PCI LAN adapter carrying the signal to wake the computer and supply standby power. Those signals were added to the PCI bus connector in revision 2.2 of the PCI bus specification, eliminating the requirement for the auxiliary cable and simplifying system integration. The *PCI Power Management Specification* further refined PCI to support power management, defining and explaining additional registers added to support power management. Using power management registers defined in the ACPI specification, an operating system going through its wake up sequence determines the source of the wake up event while enumerating the cards on the bus.

The USB and IEEE 1394 peripheral buses and power management specifications extend power management capabilities outside the system box to external devices like cameras and speakers. These buses allow the operating system to lower power to individual bus devices, portions of the bus, or the entire bus and bus controller itself.

Intel's Instantly Available PC Initiative and Microsoft's OnNow Initiative

"Initiative" is an industry term for a bid to introduce new technology, such as wanting to reduce PC power consumption. Initiatives are the first step in realizing goals; they begin the process of convincing people the ideas are worthwhile. The cooperation of other companies is usually required to realize an initiative's goal—for example, reducing PC power consumption with ACPI required Intel and other chipset developers to provide management capabilities in the hardware, Microsoft to implement functionality in Windows, motherboard designers to use the ACPI chipsets and provide the related support, power supply suppliers to implement dual mode supplies, and driver writers to support power management functions.

The beginnings of ACPI were in Intel's Instantly Available PC (IAPC) initiative and Microsoft's OnNow initiative. IAPC defined a hardware architecture for efficient power management on the desktop. OnNow provided a broad system architecture for power aware operation.

Instantly Available PC and OnNow complement each other. IAPC technologies include suspend to RAM (STR) with split power planes, dual mode power delivery, PCI power management, and thermal control. IAPC made resume possible in less than 5 seconds from a power state consuming less than 5 watts, and did so with a low cost implementation. Features common between IAPC and OnNow include ACPI, bus power management, and device class power management specifications.

Summary

PC power management responds to the need to reduce system power consumption in support of longer battery run times and global power conservation. Inventors of early power management technologies failed to understand the importance of centralized decision-making in the power management architecture, but were successful in creating the technology foundation for ACPI, in which power management decisions are made by the operating system.

You'll read about the key power management concepts underlying ACPI in the next chapter.