Prep Standardizes PowerPC Systems

Questions and Answers on New Specification from IBM, Motorola

by Linley Gwennap

To assist third-party developers of PowerPC systems, IBM and Motorola have developed the PowerPC Reference Platform (Prep) specification. For these developers, the document details how to build products compatible with IBM's forthcoming RISC PCs from its Power Personal Systems group *(see **071601.PDF**)*. For their competitors, it gives a sneak peak at the configuration of these future systems. For other processor vendors, it offers insights into how Motorola and IBM are establishing an open system standard.

Why Have a Specification at All?

The most successful and, by many measures, most open system standard is the so-called IBM PC or PC/AT standard for x86-based personal computers. This success has occured despite the lack of any specification or standards document. Originally, third-party vendors built clone systems essentially identical to IBM's PC/AT; they used the same CPU, similar chip sets, and the same ISA expansion bus.

Over time, the PC standard has evolved away from identical copies of IBM's systems to compatibility with DOS, Windows, and the body of applications for these operating systems. The advent of Windows, which uses device drivers rather than hard-coded I/O routines, allows hardware vendors to use a variety of video and peripheral devices in their systems.

Almost all vendors continue to offer a strict DOS-compatibility mode, however, despite some added costs for redundant components. Because DOS applications (and, to a certain extent, Windows software) access the hardware directly, compatibility requires all systems to support the same memory map and the same set of registers at the same I/O addresses. This restricts the ability of system designers to differentiate their products and makes it difficult to improve certain aspects of the 12-year-old PC design.

Without a controlled standards procedure, the PC standard has moved out of IBM's hands and into Microsoft's and, to a lesser extent, Intel's. As proof, the latter two companies are driving many of the recent innovations to the PC architecture, such as Plug-and-Play ISA and the PCI local bus. Even these companies, however, cannot make radical changes to the standard due to the weight of the installed software base.

By creating a documented standard instead of a de facto one, IBM and Motorola hope to avoid some of these problems. The Prep specification defines a buffer, called the machine abstraction layer (MAL), between the hardware and the software, eliminating the need for low-level compatibility. This allows system designers much more flexibility, enabling more differentiation than is found in current PCs. By adding unique value to their systems, Prep vendors might obtain higher margins than the paper-thin profits gained by most x86 PC vendors while delivering better products to their customers.

The Prep document also provides a way of evolving the standard over time as technologies advance. Right now the document is maintained by IBM, but the company says it will donate the material to an appropriate industry body as soon as one is identified. This is a lesson that Big Blue learned from its experience with Micro Channel, which the company kept too close to its chest, scaring off many potential adopters. If IBM follows through with its plans, a commonly owned standard could be maintained in a fair and intelligent fashion.

What Does the Specification Do?

The Prep specification defines a set of PowerPC systems that is compatible with the same shrink-wrapped operating system and application software. This means that applications that run on one Prep system will run on any other Prep system without modification. The companies hope that this guarantee of compatibility will encourage software vendors to develop applications for the Prep platform. (This is not an issue for Apple, which sets de facto standards as the only vendor of Macintosh systems.) This software base is essential for PowerPC to become a dominant, or even viable, architecture outside of Apple and IBM proprietary systems.

The Prep standard is fairly broad—it is applicable to notebooks, desktop systems, and even some servers. It does not cover smaller devices, such as PDAs, and it is not really suitable for large MP servers, although it could be extended in that area later. The primary focus of Prep is on the desktop, although it does include power-management features that are useful for either notebooks or Green PCs.

For More Information

To receive a copy of the Prep specification, contact IBM Microelectronics at 800.PowerPC or 512.838.2992, or contact your local Motorola sales office.

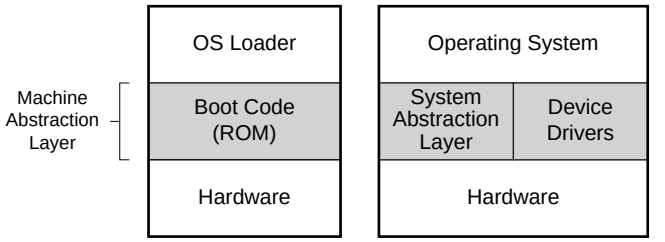

Figure 1. Prep's machine abstraction layer provides a hardware-independent interface for operating systems and applications.

The Prep document is divided into two parts. The Prep standard defines system parameters required for compatibility and is not tied to any particular OS or bus structure. The document also includes a standard reference design that specifies a particular set of components, buses, and register definitions used by a particular PowerPC system implementation.

Some third-party vendors may simply clone the reference design; for these companies, IBM Microelectronics will provide design kits with all required design information and components (processors and system-logic chip sets). Alternatively, IBM will supply motherboards implementing the reference design. In either case, clone vendors can use the same MAL code as IBM. (While the company says it will sell the design kits for a small up-front fee and no per-unit royalty, it may charge a per-unit license fee for its MAL code.)

Other vendors may take advantage of the flexibility of the Prep specification by creating their own compatible system designs. These vendors may combine some IBM system logic with their own or third-party chips to create a differentiated design. In this case, the system vendor is responsible for creating new MAL code that handles the unique features of its hardware.

How Does Machine Abstraction Work?

The machine abstraction layer actually has three parts, as shown in Figure 1. The first is the boot code, which is typically kept in ROM on the system board. It is responsible for finding and loading the operating system, and must initialize and test any portion of the hardware required for this operation. The boot code is machine-specific but not OS-specific; that is, it should be able to load any OS that is compatible with the specification.

To enable loading of multiple operating systems, the Prep document defines a data structure that must be included on bootable storage devices. The boot code selects a device based on information in nonvolatile RAM (NVRAM) and uses the data structure to locate and load the operating-system image. Since the boot code is machine-specific, it has a full understanding of the hardware configuration required to perform this task. Once the OS (or typically, a smaller OS loader) is in memory, the boot code transfers control to that software.

	Minimum Required	Recommended Portable	Recommended Desktop
Processor	PowerPC	PowerPC	PowerPC
L2 Cache	Optional	Optional	Optional
Memory	8M	up to 32M	up to 32M
Hard Drive	80M	240M	240M
Floppy Drive	Optional	1.44M	1.44M
CD-ROM	Optional	Optional	XA
Keyboard	Required	Required	Required
Mouse	Required	2-Button	2-Button
Audio	16-bit stereo, 44.1 kHz	Multivoice, compression	Multivoice, compression
Graphics	$640 \times 480 \times 8$	$640 \times 480 \times 8$	$1024 \times 768 \times 16$ w/ accelerator
NV Memory	4K	4K	4K
Real-Time Clock	Required	Required	Required
Serial Port	19.2 kbps	≥19.2 kbps	≥19.2 kbps
Parallel Port	Centronics	P1248	P1248
SCSI	Optional	Optional	SCSI-2
Network	Optional	Ethernet	Ethernet

Table 1. Prep defines minimum system requirements along with a set of typical configurations for notebook and desktop systems.

The second part of the MAL is the system abstraction layer (SAL), also known as the HAL in Windows NT terminology. This code allows the operating system to perform hardware-specific tasks such as reading the NVRAM, reading the timer, or flushing the cache. The SAL consists of functions that are used by the OS to perform these tasks. The definitions of these functions are specific to each OS. The Prep document describes the set of hardware features that can be accessed by the SAL.

To facilitate distribution of the SAL code, it is stored on disk rather than in ROM. Because the SAL is specific to both the hardware implementation and the operating system, it may be provided by the OS vendor or the system vendor. The SAL code may be included on the CD-ROM with the OS, or it may be preloaded onto the hard drive by the system vendor. (Microsoft has said it will include HAL code for the Prep reference design on its Windows NT CD-ROM.) Alternatively, the OS loader may read the SAL code from a floppy disk before loading the rest of the operating system from the boot device.

The third type of machine abstraction is performed by the device drivers. This mechanism is used by most modern operating systems to prevent the OS from having hardcoded interfaces for all possible peripheral devices. Instead, the OS defines functions that perform basic operations (such as read block and write block) for a peripheral. The device driver then executes these functions as appropriate for a specific device. Like the SAL calls, the device-driver functions are usually implemented as a run-time library.

Device drivers are specific to both the peripheral hardware and a particular OS. These drivers are typically supplied by the system vendor or, for add-in devices, by the peripheral vendor.

Workplace OS at the Center of IBM's Plans for PowerPC

Ever since the great Microsoft/IBM rift, the two companies have been pursuing competitive strategies for operating systems. For x86-based systems, IBM's OS/2 2.x has competed, with limited success, against Windows 3.x. While OS/2 has some technical advantages and is preferred by some large accounts, the momentum clearly has been with Windows.

For non-x86 systems, however, the picture is different. Microsoft's portable operating system is Windows NT, which—at least currently—requires much more memory than Windows 3.x and is therefore not suitable as an upgrade for the installed base or as a high-volume, mainstream operating system.

IBM's portable operating system is called Workplace OS. Based on a Mach microkernel, Workplace OS will be something of a chameleon, supporting a variety of operating-system "personalities." The kernel, which provides personality-neutral services, will be matched with different sets of application programming interface (API) layers for different environments. The initial version will include an OS/2 API that will be very similar to today's OS/2 2.x API. It can be thought of as the portable version of OS/2, although it has a significantly different internal design. IBM does not plan to port the existing OS/2 code to other architectures.

The beta release of Workplace OS with the OS/2 API is scheduled for 1Q94. Later in 1994, the OS/2 API will be extended with Taligent's application framework, providing a richer, object-oriented environment for new applications. Ultimately, Workplace OS will add Taligent's interface as a separate personality. At that point, the system could be shipped (depending on each user's needs) with the OS/2 personality, the Taligent personality, or both. Taligent will use the Workplace OS kernel for its own operating-system offerings.

A UNIX personality, following the COSE standard, will also be available. Because it will operate with the opposite byte order and will not support some AIX-specific functions, AIX applications will require recompilation and varying degrees of modification.

IBM expects that Workplace OS with just the OS/2 personality will be comparable in size to the current version of OS/2, which requires 8M of system memory—half the size of Windows NT. As a result, Workplace OS could be more suitable than NT for low-cost systems. IBM is positioning Workplace OS primarily as a desktop operating system, while NT initially will be most popular in servers. IBM's Power Personal Systems group, aiming to sell hardware, will support NT and AIX as well as Workplace OS.

Workplace OS will support x86 Windows applications via the Win-OS/2 subsystem, just as in today's OS/2, using 486 emulation software developed at IBM. This subsystem uses Windows 3.0 code, which is available to IBM under its agreement with Microsoft. This agreement does not apply to any subsequent Microsoft products, including Windows NT and Windows 4.0 ("Chicago"). Nevertheless, IBM is confident that it will be able to follow Microsoft's API evolution—at least through the Win16 and Win32s APIs. It is less clear whether Workplace OS will support the Win32c (Chicago) API or the full Win32 (NT) API. Note that the Wabi technology, which IBM has licensed from SunSoft, is used only in AIX, and IBM does not plan to use it in Workplace OS.

IBM plans to port Workplace OS to the x86 and to other RISC platforms; these products could begin emerging in 1995. IBM's focus for the x86 is on the current OS/2, however, at least for the near term.

—MS

What Does the Specification Require?

Table 1 summarizes the system requirements of the Prep specification. While the "minimum" column shows the lower limits of the standard, most systems will have more robust configurations, as shown in the columns for notebook and desktop systems.

A Prep-compatible system starts with any PowerPC processor, including the 601, 603, and future versions. Level-two caches are optional and can be included to improve performance. The minimum main memory is 8M, expandable to at least 32M. A hard-disk capacity of at least 240M is recommended. Most systems will have 3.5" floppy drives and, at least on the desktop, CD-ROM drives. Any expansion bus may be used.

At minimum, Prep systems must have CD-quality stereo audio; most systems will support multivoice audio with compression. Portable systems may use VGA-quality graphics, but desktop systems are expected to deliver $1024 \times 768 \times 16$ resolution. Most systems will have a fast serial port and an enhanced bidirectional parallel port (P1248). Desktop systems are expected to support the Fast SCSI-2 protocol.

The document also defines the typical configuration for three other types of systems, which are all similar to the desktop configuration. A medialess system has no CD-ROM, floppy, or hard drive. A technical system includes a level-two cache for higher performance. Finally, a server has a level-two cache and larger hard drive, but simpler video and audio capabilities.

The specification creates a basic model for power management, consisting of five states: On, Enabled, Standby, Suspend, and Off. State changes are governed by the operating system; the document does not cover devices with hardware power management. Device drivers should allow the OS to change the power-management state of the device. A special device driver must be pro-

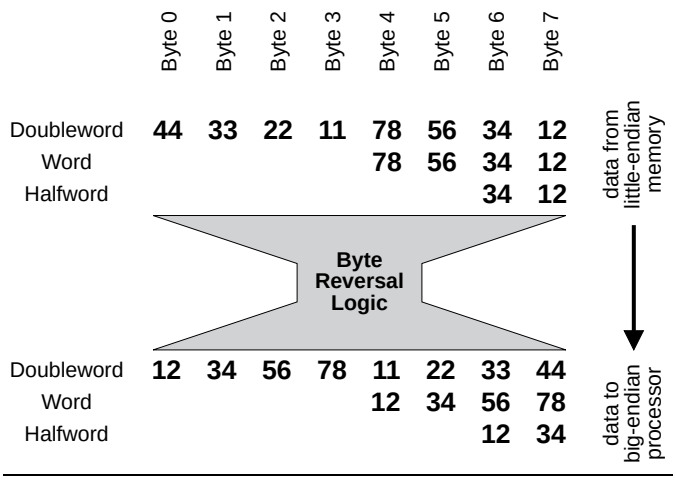

Figure 2. Byte-reversal logic must convert little-endian data into big-endian format for the PowerPC CPU.

vided to manage CPU power.

The Prep standard allows for multiple processors in a single system using a shared-memory model. The PowerPC architecture includes a set of instructions for ordering memory accesses and synchronizing processors in an MP configuration.

How Can a System Be Bi-endian?

All Prep systems must be able to execute instructions stored in either big- or little-endian byte order. Similarly, these systems must access data stored in either format. Current PowerPC processors (601, 603, and 604), however, assume that information is always stored in big-endian mode. Thus, compliant systems must provide some external logic to handle little-endian data.

Current PowerPC processors support an address-translation mode for little-endian operation. In this mode, all addresses generated by the processor are modified in the lower bits to obtain the proper little-endian address. For example, for a byte access, the lower three bits are inverted. For doubleword (64-bit) accesses, no translation is required. This translation does not work for unaligned data accesses; these are trapped and handled in software.

Address translation alone does not solve the problem, however. Once the correct group of bytes has been selected, the byte order within the group must be reversed to convert a little-endian data structure into big-endian form. Figure 2 shows how some little-endian groups of bytes are reversed. In hardware, this reversal function can be implemented as a set of multiplexers. Placing this circuitry between the processor and the rest of the system means it need be implemented only once.

To simplify the problem, Prep assumes that endianness is selected at boot time, and all data in the system (CPU, memory, and peripherals) uses the same format. There are no provisions for a mixed environment, such as a little-endian application running with a big-endian OS.

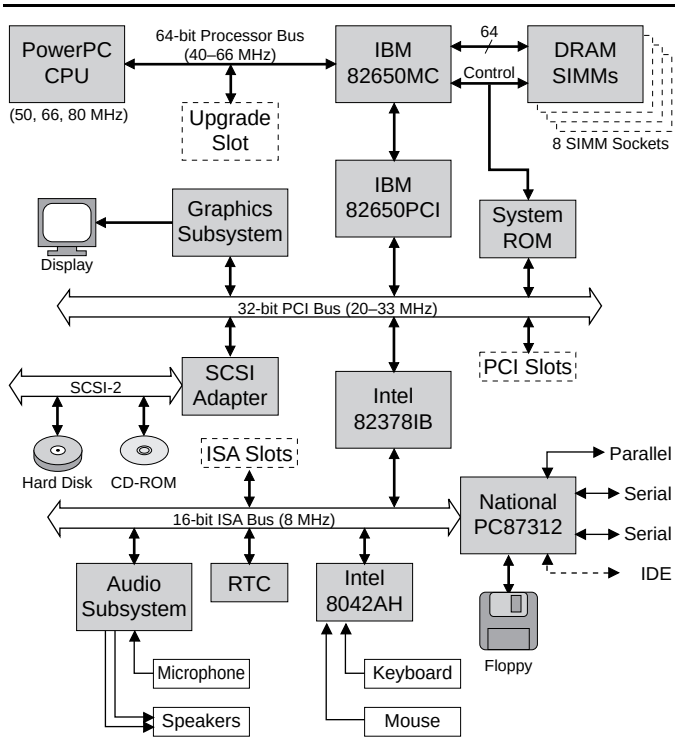

Figure 3. The IBM reference design uses PCI for high-speed peripherals and ISA for standard PC devices.

It is likely that future PowerPC processors (possibly the 620 and others) will integrate the byte-reversal logic, simplifying the system design. Until then, it will probably be included in the system-logic chip set. For example, IBM's 82650 PCI chip set *(see **071601.PDF**)* implements this function. The circuitry itself is not that complicated and is included today in bi-endian processors such as the MIPS R4x00 and HP PA-7100LC.

What's in the Reference Implementation?

The Prep document also describes a reference implementation. This is the specific system configuration shown in Figure 3. The document details the system memory map, location and use of registers, cache timing, and a method of switching to little-endian mode. Systems that implement all features of the reference implementation will be able to use standard MAL code for Prep-compliant operating systems, eliminating the need for the system vendor to develop its own firmware.

The reference system uses a PowerPC 601 processor and the aforementioned 82650 chip set; in the future, other chip sets compatible with the 82650 could also be used. The CPU can run at up to 80 MHz. Eight SIMM slots allow memory sizes of 8M to 256M using 8M or 32M modules. The system uses industry-standard, 72-bit-wide SIMMs.

The PCI bus is used for high-speed peripherals, specifically SCSI-2 and graphics. The former uses an

NCR 53C810 PCI adapter, and the latter can be implemented with either a Weitek Power 9000 or S3 928PCI accelerator chip. The graphics subsystem is contained on a PCI add-in card for modularity; a second PCI connector is provided for additional expansion. A flash ROM (or EPROM) is also attached to the PCI bus but is controlled directly by the 82650. This 512K memory contains power-on selftest and boot code.

Intel's 82378IB (not IBM's 82374) chip bridges the PCI and ISA buses and provides basic system logic, such as DMA, interrupts, and timers. The ISA bus is used for the audio subsystem (Analog Devices 1848KP or Crystal Semiconductor 4321) and other low-speed I/O interfaces (National 87312 Super I/O chip). Three ISA slots are included for expansion.

The system board (a "planar" in IBM-speak) measures 9" × 13" and fits in a standard PC enclosure. IBM Microelectronics will market this board to OEMs and will also provide a design kit with all specifications (Gerber tape, bill of materials, etc.) required to manufacture the design. IBM's Power Personal Systems group is expected to market systems using this design. No pricing has been released for any of these products, but they are expected to be shipping in volume by mid-1994.

What Goes in the Upgrade Slot?

The reference system contains a 200-pin socket connected to the processor bus. IBM calls this the upgrade slot, its answer to Intel's OverDrive socket. Unlike Intel's socket, it can be used to upgrade the system in two ways, with either an external cache or a second processor.

Since the basic reference design has no external cache, the upgrade slot can be filled with a cache card containing SRAM and a cache-control ASIC. (Neither the 601 nor the 82650 chip set directly controls an external cache.) The Prep document states that IBM Microelectronics will market a chip for this purpose, but the company has not yet announced such a part. Typical cache sizes are 256K and 512K. The add-in cache can be either write-through or write-back.

The upgrade slot also can be used to add a second processor to the system. A card with an upgrade processor could optionally contain an external cache for that processor. The second processor can automatically disable the original CPU, or they can operate in tandem. Since the PowerPC 604 is expected to roughly double the performance of the 601, this chip would be the most likely candidate for a processor upgrade.

What's the Prognosis?

The Prep specification offers third-party vendors two choices for building systems compatible with IBM's own PowerPC products. The simplest option is to copy the reference design exactly, using components from IBM Microelectronics, and license the necessary MAL code from IBM. For vendors wishing to differentiate their products, the specification defines areas in which they can innovate while still retaining compatibility via the MAL. These vendors would have to either write their own MAL code or license and modify IBM's.

By publishing the specification a half-year before it actually ships systems, IBM is inviting other vendors to join the PowerPC party at the beginning. If the Prep platform is a rapid success, the initial vendors could profit handsomely before the competition becomes too fierce. But if, as seems likely, the ramp is longer, these vendors could be stuck in a low-volume market with the giant, newly aggressive IBM as their main competitor.

The Prep specification itself appears to be a solid technical definition of a low-cost RISC personal computer. But then again, so was the ARC specification. For Prep to succeed where ARC did not, other major PC vendors must join IBM in backing PowerPC, encouraging ISVs to port their applications. This success depends less on the merits of the specification than on the persuasiveness of IBM and Motorola. ♦

