Intel Unveils Multiprocessor System Spec

Based Around APIC, MPS Enables Shrink-Wrapped MP Software

by Michael Slater

Seeking to broaden the use of symmetric multiprocessor systems by making them easier to build, Intel has defined a standard approach to the design of such systems that could eliminate the need for hardware developers to provide customized operating system software. At the heart of the specification is a requirement for register-level compatibility with Intel's advanced programmable interrupt controller (APIC), which is integrated into the P54C version of Pentium *(see **080301.PDF**)*.

Intel is working with Microsoft, Novell, IBM, Santa Cruz Operation (SCO), and SunSoft to support the MP specification with off-the-shelf versions of Windows NT, NetWare, UnixWare, OS/2, SCO Unix, and Solaris. The standard, originally called PC+MP and now officially called simply the MultiProcessor Specification (MPS), is an extension to the standard PC/AT architecture, so systems based on it remain fully PC compatible.

PC makers ALR, AST, Dell, HP, Micronics, and Olivetti announced plans for systems that comply with the MP specification. Phoenix Technologies will provide BIOS support.

Several companies currently making MP systems, including Compaq, Corollary, and Unisys, provided statements saying, in essence, that anything that encourages broader use of multiprocessing is good, but they stopped short of announcing any intent to build conforming systems. Executives at these companies indicated that they initially viewed the specification negatively, but have gradually approached neutrality. Their half-hearted endorsements were probably driven largely by a desire simply to be part of a multiprocessor announcement. Many leading x86 MP vendors, including Sequent, AT&T/NCR, Tricord, and NetFrame, were conspicuously absent from the list of endorsements.

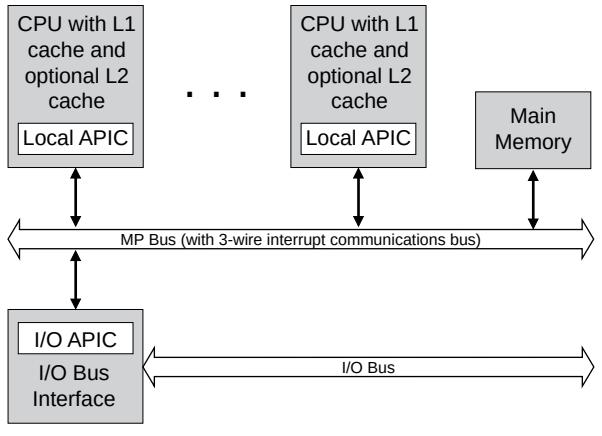

Figure 1. The MP specification requires a local APIC associated with each processor and one or more I/O APICs to accept interrupt inputs. In a simple two-processor system, the MP bus would be the P54C's local bus.

In the past, multiprocessor system designs have required custom hardware and software to support the functions that are not present in a uniprocessor system. A hardware mechanism is required for routing interrupts among the processors, starting up each processor, and handling processor reset. In the past, each system vendor has implemented these functions differently and therefore had to provide customized operating-system software. With Windows NT, this customization is done in the hardware abstraction layer (HAL); other operating systems are not so cleanly divided but are moving toward a HAL-like partitioning.

Intel's MP specification gives system makers a plan to follow for a hardware design that will be supported by most MP operating systems without any additional software work. Until now, the software effort required has kept smaller PC makers, or those without engineering resources, from entering the MP market. The MPS is likely to lead to an increased number of MP system suppliers, including more low-margin players, which should push down the price of such systems. It will also give MP system users more flexibility in choosing system suppliers, since they will be able to run the same software on all compliant systems, and it could increase the number of operating systems supported on each system.

Some MP system developers will continue to develop noncompatible systems, however, because they have existing hardware architectures that don't conform to the specification. For a company that has already invested in developing the OS support needed for its MP hardware, there is little to be gained by supporting the specification. Many high-end servers will require custom software support even if they conform to the specification, because they include features beyond those supported in the specification.

Supporting Hardware Configurations

The MP specification does not require any particular bus or cache configuration or number of processors. It can be applied to a two-processor system with a single, shared, level-two cache or to a higher-end system with a dedicated level-two cache for each processor. The I/O bus

can be ISA, EISA, or Micro Channel, and various multiprocessor interconnect buses can be used.

To supply information about the specific hardware implementation to the operating system, the MPS defines the "MP floating pointer structure." This data structure, usually 16 bytes long, identifies the system as complying with the MPS and provides a pointer to an optional configuration table. Because the PC architecture makes it problematic to assign anything to a fixed address, this structure does not have a fixed location. Instead, the OS must search certain regions of memory for the structure, which includes the ASCII string "_MP_" as an identifier. It also has a checksum, so software can confirm that it is not a random occurrence of this string.

For two-processor systems using any standard PC bus, the specification defines default configurations. For systems that conform to one of the default configurations, the MP floating pointer structure simply identifies the system type with an 8-bit code; no further information is required.

Systems with other configurations use this pointer to provide the address of a more complex data structure, the MP configuration table, that provides detailed information about any number of processors. The MP configuration table specifies the processor types, APIC addresses, and bus types. Some complex MP systems may not be supportable even with the full MP table, requiring custom HALs (or equivalent) for each OS.

APIC at Heart of MP Specification

At the heart of the MP specification is Intel's APIC, which provides the mechanism for distributing interrupts among processors and for enabling one processor to interrupt another. Intel has built the MPS around the register-level features of the APIC. Some MP system designers would have preferred a higher-level specification based on an API-level standard, leaving the register-level details flexible. An API-level standard would have given system designers more options but wouldn't have met Intel's goal of establishing its APIC as a hardware standard.

As Figure 1 shows, the APIC has two parts: the local APIC, which is associated with a particular CPU, and the I/O APIC. The I/O APIC receives interrupt inputs, which can then be routed to any local APIC. There must be one local APIC per CPU. There may be one or more I/O APICs, which are shared by all CPUs.

Intel has integrated the local APIC into its P54C Pentium processors. The I/O APIC can be integrated into system-logic chip sets. VLSI, Opti, Forex, and Symphony will provide chip sets that include the I/O APIC. Intel also offers an APIC chip, the 82489DX, which provides both parts of the APIC and can be used with 486 processors or with the P5 version of Pentium (which does not include the local APIC).

Price & Availability

Intel is distributing the MP specification without charge, and there are no royalties required for its use. Copies of the specification can be obtained from your local Intel sales office or by calling 800.548.4725 and asking for literature packet #242016-001.

Computers and operating systems supporting the MP spec are expected to be available in the second half of this year.

Because the APIC works differently than the PC-standard 8259A interrupt controller, it is not, by itself, compatible with standard PC software. Most systems will therefore include an 8259A as well (typically as part of standard system-logic chips), and the APIC is designed to coordinate with this device.

Intellectual Property Issues

The APIC's utility to Intel goes beyond its mere hardware function: it also may serve as a barrier to non-x86 microprocessors. Intel is licensing the VHDL code for the APIC, enabling system designers to incorporate the APIC in their system logic. But the standard terms of this license require that the APIC be used only with x86 microprocessors, and some sources claim that Intel has proposed licenses limiting use of the APIC VHDL to Intel processors only.

The restrictions on the use of Intel's APIC VHDL have caused some MP system designers to choose not to use the APIC, and at least one large OEM has succeeded in negotiating license terms that allow it to use the APIC with any processor. Sources say that IBM is designing its own reverse-engineered I/O APIC to avoid Intel's licensing restrictions. (There could be an opportunity here for a third-party supplier of an APIC-compatible logic design.) Even if the APIC were reverse-engineered, however, patents could be a problem for system makers lacking an Intel license.

Intel emphasizes that the MP specification itself is fully open, and there are no royalties for its use. The restrictive licensing of the APIC VHDL is in conflict with this openness, however, and Intel may relent and make the APIC design more openly available. System makers are willing to pay a license fee for Intel's APIC design, but they are not willing to accept restrictions on how they can use it.

Intel claims that the intellectual property issues have been blown out of proportion, and perhaps they have, but it is hardly surprising that system makers are wary of Intel's actions in this regard. It does appear, however, that Intel sought to use the APIC to gain a proprietary edge but has gradually backed down in response to vigorous opposition from system makers.

Servers to Dominate MP Use

The MP spec applies to desktop systems as well as to servers, but software limitations will prevent MP systems from becoming popular on the desktop for some time. Having more than one processor increases performance only if there are multiple tasks available to be executed at the same time. In a server, which is typically serving multiple users, this is often the case. But in most single-user desktop systems today, multiple processors won't result in any performance gain, since only one task is typically performed at a time.

Two software advances are needed for MP on the desktop to become widespread. First, there must be a mainstream operating system that provides MP support. Second, the operating system and applications must support multiple threads, which are parallel tasks within a single application. Even then, a benefit is realized from multiple processors only if the task can be decomposed into multiple threads. Without multithreaded applications, the advantage of multiple processors would be gained only when running multiple applications at the same time. Although many users keep multiple applications open on their desktop, rarely is more than one doing any actual work.

The forthcoming "Chicago" version of Windows supports multiple threads but not multiple processors. So although it will not directly support MP systems, it is a step in the right direction. Because Chicago is assured of widespread use, many application developers are likely to invest in making their applications multithreaded. Users desiring more performance could then move to Windows NT, or its successor Cairo, to take advantage of multiple processors.

Mixed Blessings for MP Vendors

For companies that have been making multiprocessor systems for years and have invested in their own interrupt controllers and software support, the MP specification is a decidedly mixed blessing. To comply with the specification, such companies would have to modify their hardware designs, including ASICs. In addition, the shrink-wrapped operating systems promised by the MPS may be inadequate to support the special features of these systems, such as fault tolerance and disk arrays.

If MP system makers license Intel's APIC design to facilitate their designs, they will limit their options to use RISC processors with similar system logic—unless they are able to negotiate better terms with Intel.

The MP specification also makes it easier for other PC makers to compete with the few companies that have invested in MP designs. Multiprocessor servers are one of the few high-margin areas of the PC business, due in part to the engineering resources needed to create them. Now, any PC clone maker that wants to enter this market can do so with relative ease, and its computers will be compatible with a range of operating systems. This will provide system buyers with more choices and lower prices, which is good for them but a threat to the companies that have been profiting from this relatively protected niche of the PC market.

This situation is not as threatening for companies such as Compaq as it might appear at first, however, because the value-added in MP servers goes beyond just the MP support. Features such as disk arrays and fault tolerance are critical for high-end servers, and providing these capabilities—as well as the sophisticated caches and memory systems needed for high-performance MP systems—still requires significant engineering work. Some of these features require OS modifications as well, limiting the usefulness of the standard software enabled by the MPS. MP servers that just follow the MPS and don't provide these added features aren't likely to sell well as enterprise servers, but they could put price pressure on more sophisticated offerings from established MP suppliers.

The MP specification is also a mixed blessing for suppliers of MP chip sets. Corollary, for example, has developed an MP chip set *(see **070503.PDF**)* for Pentium systems that includes an interrupt controller that is different from the APIC, so it is not compatible with the MPS. Corollary has been investing for years in providing MP OS support, and for its OEM customers, the MPS is largely irrelevant: Corollary will provide the same operating systems that are promised to support the MPS. It is also possible to disable Corollary's interrupt controller and use a local APIC in a P54C processor and an I/O APIC in the I/O bus chip set.

LSI Logic also has announced an MP chip set *(see **070503.PDF**)*. LSI's chip set does not include an interrupt controller, so it can conform to the MPS by working with an external APIC chip or a P54C and an I/O chip set with an integrated I/O APIC.

For Intel, the benefits of the MP specification are clear. The server market is one area where RISC processors have a relatively strong position against x86 processors. By enabling more companies to produce MP systems, Intel will increase the competitiveness of its architecture in that market. Using multiple processors also gives Intel a way to counter the performance advantages of RISC processors. At the same time, though, promoting MP encourages a shift to Windows NT, which makes it easier for users to switch to a RISC system.

The MP specification will enable more system makers to produce MP systems, but it won't necessarily do much to increase demand for those systems. For most PC users today, a multiprocessor system just doesn't provide much value beyond a uniprocessor system. The groundwork being laid by the MPS, however, could facilitate desktop multiprocessors moving into the mainstream when the software is ready. ♦

