intel®

Intel Pentium™ Processor Multiprocessor Systems

Technical Backgrounder

Order Number 242169-001

Intel Pentium™ Processor Multiprocessor Systems

Executive Summary

Multiprocessing has traditionally been limited to large, high-end systems available only at high cost. However, recent advances in software and hardware are making multiprocessing practical and available for desktop systems, workstations, and entry-level servers.

- Off-the-shelf, multiprocessing-ready operating systems, including IBM's OS/2*, Microsoft's Windows NT*, Novell's NetWare* and UNIXware*, SCO's MPX*, and SunSoft's Solaris* are available today or will be in the near future.
- Hardware requirements for multiprocessor systems such as cache consistency protocols, multiprocessor interrupt controllers, bus arbitration, write-back caches and wide data buses have been integrated into the Pentium processor and standard components available from Intel and other vendors for ease of design.
- The MultiProcessor Specification (MP Specification) developed by the Intel Architecture Labs in cooperation with leading system manufacturers, operating systems and BIOS vendors brings the same "shrink-wrapped" benefits of the desktop market to the multiprocessing market. MP Specification-compliant operating systems that will be available include IBM OS/2, Microsoft's Windows NT, Novell's NetWare and UNIXware, SCO's MPX and SunSoft's Solaris. System vendors developing MP Specification-compliant systems using the Pentium processor include , ALR, Apricot, AST, AT&T GIS, Compaq, Corollary, Dell, Hewlett-Packard, ICL, Intergraph, Micronics, Olivetti, Unisys, and Zenith Data Systems.
- Multiprocessor systems based on the Pentium processor will be available in two forms: dedicated-cache multiprocessor systems and shared-cache dual-processor systems. Dedicated cache systems provide maximum performance needed by application servers in enterprise environments. The 90- and 100-MHz Pentium processors enable a new class of cost-effective, shared-cache dual processor system targeted at the workstation and entry-level server marketplace.

A New Generation of Operating Systems

The first personal computers contained a single microprocessor and used DOS as the operating system. DOS was limited in that it could only run one application at a time. Changing applications required the user to close one application and open another one.

Windows introduced multitasking, which allowed a user to have two applications open at one time (though only one could be executing). The Windows operating environment increased productivity because it allowed users to switch between applications easily, by clicking on a window or hitting a "hot- key" combination. However, having only one active application or task at a time means waiting for one task to complete before the other can become active. For example, a user must wait for a print command to complete before continuing with word processing.

The next generation of personal computer operating systems (exemplified by Microsoft's forthcoming operating system code-named Chicago) will introduce pre-emptive-multitasking. Pre-emptive-multitasking allows multiple applications to run at the same time. This increases productivity as it allows, for example, the user to continue word processing while the print command is executing. Both tasks may be active at the same time.

Advanced operating systems such as Windows NT, OS/2-MP, UNIX, and Solaris add multiprocessing capability to pre-emptive multitasking. Multiprocessing capabilities enable multiple tasks or applications to run in parallel on different microprocessors. This improves performance, as two tasks running in parallel will finish more quickly than the same two tasks running serially (one after the other).

With new advanced operating systems setting the stage for multiprocessing on smaller, lower-cost systems, all that remains is to design workable architectures based on (or evolving from) currently existing hardware.

Hardware for Multiprocessing

Historically, the traditional PC platform has been a single 80x86 processor in an AT-compatible platform. An AT-compatible platform is not adequate to support multiprocessing. Fortunately, recent advances in hardware technology and system design have made multiprocessing possible without requiring users to completely abandon the large installed base of software designed for the AT platform.

Three capabilities must be added to the AT platform to make it multiprocessing ready: cache consistency; the routing of interrupts to the appropriate processors; and bus arbitration between the processors.

The first requirement for a multiprocessor system is cache consistency. A cache is an area of high-speed memory that acts as a temporary storage place for instructions and data recently used by the microprocessor. Such instructions or data are likely to be used by the processor again, and retrieving them from the cache is much faster than retrieving them from main memory. However, simply adding a cache to a microprocessor can create its own problems for multiprocessing. If two microprocessors are working on the same data, the values for that data held in the chips' caches may not be the same.

Solving this "cache consistency" problem requires a protocol such as the MESI (modified, exclusive, shared, or invalid) protocol. Under the MESI protocol, each line of a cache is marked as either "exclusive" (held by a single cache), "modified" (written-to by a processor), "shared" (by one or more caches), or "invalid" (indicating an old copy of information in a cache that has been changed in main memory). The MESI protocol ensures that all shared data is up-to-date and prevents inconsistencies between the dual processors.

Peripheral devices such as disk drives and network controllers must periodically interrupt the microprocessor to access processor time. In a single-processor system, handling these interrupts is not too complicated, as all interrupts go to the one processor.

Interrupt handling in a multiprocessing system is more complex. In such a system, if all interrupts were to go to only one of the processors, that processor could be completely tied up while other processors stood idle. For the most efficient use of a multiprocessor system's capabilities, interrupts must be intelligently routed to different processors.

In 1992, Intel introduced a multiprocessor interrupt controller called the 82489DX Advanced Programmable Interrupt Controller (APIC), which implemented the Intel APIC architecture. The APIC architecture was designed to handle interrupts in a multiprocessor environment. The APIC architecture provides support for symmetric multiprocessing (in which all processors are equal and have equal access to system resources), while maintaining backward compatibility with the interrupt controller that was previously standard on PC-compatible systems. Interrupt controllers based on the APIC architecture can be programmed to intelligently route interrupts to the appropriate processor in a multiprocessing system.

Effective multiprocessing also requires bus arbitration. Processors in a multiprocessor system share a single bus to memory. If all the processors tried to access the bus at the same time, gridlock would occur, and no processor would get access to the bus long enough to get any work done. Bus arbitration circuitry serves as a sort of "traffic cop" controlling access to the bus. When a processor needs access to the bus, it makes a request to the arbitration circuitry, which then grants bus access to each processor for a specified amount of time.

Approaches to Multiprocessing: Dedicated-Cache versus Shared-Cache

There are a number of ways to design a multiprocessor system. One design approach involves the number and arrangement of caches in the system. In a dedicated-cache system, each processor has its own cache. In a shared-cache system, the processors share access to a single cache. Each of these approaches has its advantages and appropriate applications.

The primary advantage of dedicated-cache multiprocessing systems is their scalability, that is, the ability to increase the number of processors in the system without degrading performance. Degradation in a multiprocessing system is caused by bus saturation—essentially, overcrowding—the inability of a processor to access the bus when it needs to because another processor has control of the bus. The ability to support a number of processors without concurrent loss of performance means that dedicated-cache multiprocessing systems are ideal as high-end servers for applications such as on-line transaction processing (OLTP), supercomputing, and other high-performance applications. However, a multiprocessing system in which each processor has its own dedicated cache requires more hardware and more complex circuitry to coordinate the "division of labor," and is therefore more expensive.

By contrast, a shared-cache system in which the processors share a second-level cache can be developed for minimal cost over that of a standard desktop computer. Shared-cache systems will be lower cost than a dedicated cache system but additional processors will provide smaller performance boosts.

Enter the Pentium™ Processor

The 90- and 100-MHz versions of the Pentium processor enable low-cost, shared-cache dual-processing systems by integrating the three requirements for a multiprocessing system on the chip: cache consistency, multiprocessing interrupt control, and bus arbitration.

All members of the Pentium processor family integrate the MESI cache consistency protocol on-chip. The 90- and 100-MHz Pentium processors integrate a multiprocessor interrupt controller on-chip along with bus arbitration logic. The multiprocessor interrupt controller is based on the APIC architecture and is software compatible with systems based on the 82489DX. The arbitration logic enables two processors to arbitrate for the common bus to the shared cache.

Two other features of the Pentium processor family designed to support multiprocessing are built-in write-back caches and a wide external data bus.

The Pentium processor's write-back caches—one for instructions and one for data—eliminate the need to copy each write to main memory (as happens with less complex "write-through" cache schemes). When two processors are trying to write data to shared memory, an information "traffic jam" can result. Write-back caches reduce the number of writes to a shared-cache resulting in higher performance.

When a microprocessor must go "off chip" for data, the speed of data transfer will be limited by the width of the external bus to main memory. The Pentium processor's 64-bit external data bus allows it to transfer data to and from memory at rates up to 528 Mbytes/second, more than five times that of a 66 MHz Intel486 DX2™ processor (105 Mbytes/second), and powerful enough to enable multiprocessing.

The MultiProcessor Specification: Standard for Hardware and Operating Systems

To facilitate the development of multiprocessing systems, Intel and leading system manufacturers, operating system vendors and BIOS vendors developed the MultiProcessor Specification. The specification is a set of standards for hardware, BIOS, and operating system software that extends the traditional AT architecture and enable new multiprocessing applications while maintaining 100 percent compatibility with the large installed base of PC-compatible software.

The MultiProcessor Specification calls for the system to incorporate one or more processors of at least 486 performance level and one or more interrupt controllers based on the APIC architecture or its equivalent. Systems will also require a MultiProcessor Specification-compliant BIOS and operating system.

Systems built according to the MultiProcessor Specification may incorporate one or more industry-standard buses (such as ISA, EISA, MCI, PCI or vendor-specific buses). The MultiProcessor Specification also defines a multiprocessing configuration table that enables the design of differentiated multiprocessing systems.

Today's advanced operating systems may be made MultiProcessor Specification-compliant with minimal effort. Versions of advanced operating systems that will be MultiProcessor Specification-compliant include IBM's OS/2 for SMP, Microsoft's Windows NT, Novell's NetWare and UNIXware, SCO's MPX, and SunSoft's Solaris.

The result of the MultiProcessor Specification will be a broad range of systems that provide a choice of shrink-wrapped multiprocessing operating systems for applications ranging from desktop systems to high-end workstations and servers. At the same time, these systems will provide improved performance for all of today's PC-compatible operating systems and application software.

Applications and Performance: Your Mileage Will Vary

A pre-emptive multitasking operating system can have more than one application active at a time. These applications can be written in such a way that they are subdivided into threads, which results in higher productivity. A thread is a portion of an application that can run in parallel with, or independent of other parts of the application. For example, in spreadsheet, a print thread could be running while the user is entering new data. In a multiprocessing system, threads can run on different processors.

Server or workstation applications can be threaded or non-threaded. A threaded application will always use the capabilities of a multiprocessor system more efficiently than a non-threaded application, as it can be running on more than one processor at a time. However, power users running non-threaded applications can also benefit from a multiprocessing system. For example, the operating system could run on one processor while the application runs on another.

Power users tend to run more than one application at a time. These users will see significant productivity improvements with multiprocessor systems because they have more processors over which to spread the work.

Applications have different patterns for accessing memory. Database applications (transaction processing workloads) tend to randomly access large areas of memory. These applications perform best on machines with large cache memories and on multiprocessor machines in which each processor has a dedicated cache. Dedicated-cache systems can provide these applications with a performance boost of up to 90 percent with the addition of a second processor.

Workstation applications are more calculation-intensive. They tend to execute tight loops that fit in smaller caches. These applications can benefit from a shared-cache multiprocessor system, which can provide up to a 60 percent performance boost.

Performance on multiprocessor systems will vary widely between different applications and machine configurations. Performance is also affected by a number of variables such as how applications access memory, cache architecture, and the I/O subsystem. Users should not be surprised to see that addition of a second processor may result in up to a 90 percent performance boost for some applications or only a 10 percent boost for others.

Summary

With the creation of the MultiProcessor Specification, the stage is now set for a new generation of multiprocessing-enabled Pentium processor systems. Multiprocessing operating systems and low-cost multiprocessor systems will become widely available beginning in 1994. The MultiProcessor Specification forms a standard that allows multiple operating systems to run on platforms from many vendors. Threaded applications will become more common as multiprocessor systems proliferate.

Multiprocessor systems will find their first applications in the server and workstation markets. These users require the advanced features provided by the new multiprocessor-ready operating systems and demand the highest performance.

Multiprocessing-enabled Pentium processor systems will range from shared-cache dual-processing systems that combine high performance with low cost, to dedicated-cache multiprocessing systems that provide the maximum in performance. High-end applications such as on-line transaction processing, supercomputing, and database servers for very large systems will require a dedicated-cache multiprocessing system, while shared-cache dual-processing systems will be ideal as cost-effective desktop workstations and entry-level servers.

 Printed in USA/0694/11K/NCG/JH/JL