The IBM RS/6000* 43P Model 260

David Tjon

Introduction

The RS/6000 43P Model 260 is IBM's first 64-bit symmetric multiprocessor (SMP) product that provides outstanding price and performance benefits for customers who have demanding visual computing needs, complex analysis requirements, or entry workgroup server environments. It has been designed to support the MCAD, CAE, Geophysical, and entry technical marketplaces. When the IBM GXT3000P* Graphics Accelerator is installed in the Model 260 running AIX* 4.3.2, the system is an industry leading graphics workstation. The floating point performance of the POWER3 microprocessor makes this an excellent platform for compute-intensive analysis applications. It also positions well as an entry server for Independent Software Vendors (ISV) and customers who want a cost effective platform for development and test of applications that require 64-bit support and will run on larger RS/6000 systems. .

The RS/6000 Model 260 derives its superior computational performance from the IBM's new POWER3 microprocessor. The POWER3 microprocessor has a superscalar architecture that provides concurrent operation of multiple execution units. It provides a high bandwidth interface to a fast level 2 (L2) cache, and a separate high bandwidth interface to memory and other system functions.

Minimizing changes to existing RS/6000 system and application software was a key consideration. The POWER3 microprocessor implements the 64-bit PowerPC Architecture* and is fully compatible with the existing 32-bit PowerPC Architecture. System level design preserves the RS/6000 I/O model. Because the Model 260 conforms to the RS/6000 Platform Architecture (RPA), compatibility is maintained for existing device drivers, other subsystems, and most importantly, applications.

System Overview

The Model 260 system structure is shown in Figure 1. One or two POWER3 microprocessors are connected to an IBM-designed high performance memory/system control chip set. The I/O subsystem is based on industry standard chip sets and the industry standard PCI bus. The Model 260 is a deskside package which fits under 24" high desks, and runs in a quiet office environment.

The Model 260 supports IBM's AIX operating system, Versions 4.3.2 and 4.2.1.

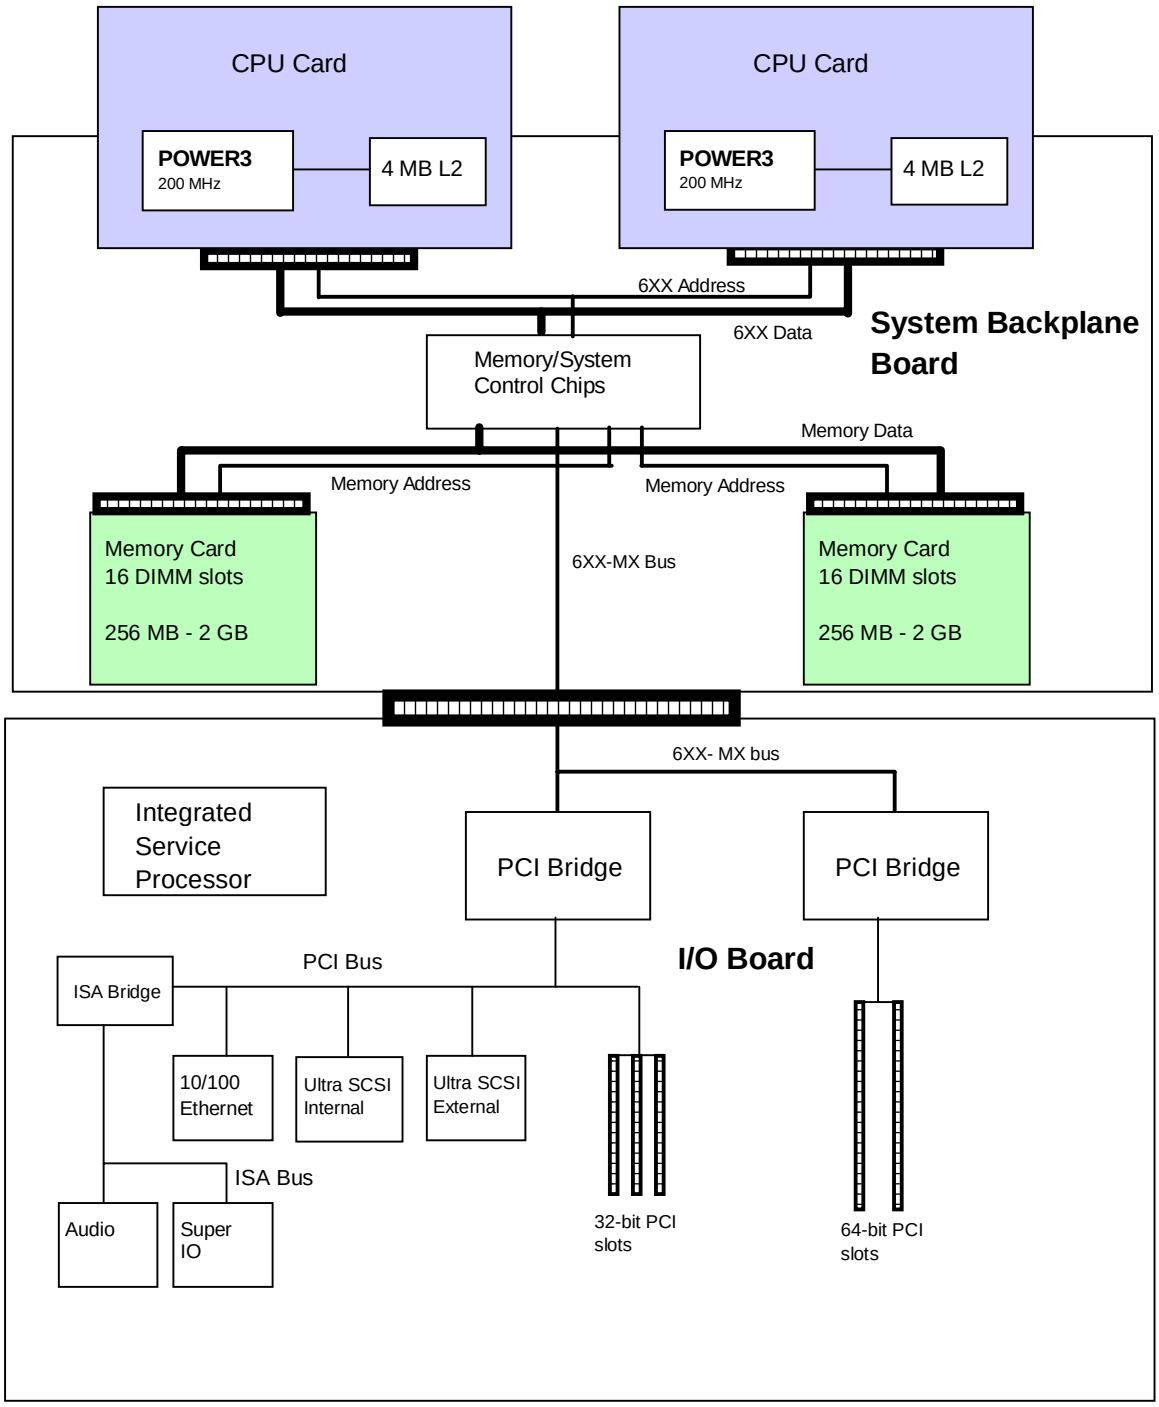

Figure 1. Model 260 Functional Block Diagram

POWER3 Microprocessor

The POWER3 microprocessor is a single chip implemented with 0.25 um CMOS technology. It operates at a 200 MHz clock cycle. The POWER3 microprocessor has eight execution units, and

allows concurrent operation of fixed point instructions, load/store instructions, branch instructions, and floating point instructions. The processor can dispatch up to four instructions at a time, execute them out of order, but is designed to ensure in-order completion and precise interrupts to provide program integrity. There is a 32 KB instruction cache and a 64 KB data cache on the chip, both parity protected. There is 256-bit external interface to a 4 MB L2 cache, which operates at 200 MHz and is ECC protected (Single Error Correction, Double Error Detection). The processor and L2 cache are located on a single, pluggable board for easy upgrade or replacement.

The POWER3 processor was designed to provide high performance floating point computation. For example, there are two floating point execution units, each supporting a 3-cycle latency, 1-cycle throughput Multiply-Add execution rate. This allows the POWER3 to execute four floating point operations per cycle, resulting in a peak throughput of 800 MFlops. The Model 260 has a measured performance of 30.1 SPECfp95.

The design, technology, and performance of the POWER3 microprocessor are more fully described in a companion document [1].

System Bus

The system bus, referred to as the ‘6XX’ bus, connects up to two POWER3 processors to the memory/system control chips. The system bus architecture is an extension to the PowerPC system bus architecture. It operates at a 100 MHz clock cycle. It provides 40 bits of real-address and a separate 128-bit data bus, both of which are parity protected. Peak data throughput is 1.6 GB/second.

The 6XX bus architecture supports pipelining of read/write operations, bus snooping, and error indications, thus maximizing throughput over this interface. For example, in the Model 260, a POWER3 microprocessor can have up to 9 read operations pending at any time. Up to 4 writes can be queued for completion in the memory controller chips. Also a ‘data retry’ function allows forwarding of memory data to the processor concurrently with ECC information.

System Memory

The Model 260 supports 256 MB to 4 GB of 10ns SDRAM. System memory is connected by the memory bus to the memory/system control chips. The memory bus is comprised of an address bus and a 128-bit data bus, both operating at 100 MHz. The memory bus is separate from the system bus, and allows for concurrent operations on the system bus and the memory bus. For example, cache to cache transfers can occur while a DMA operation is proceeding to an I/O device.

There are two pluggable memory cards, each supporting up to sixteen 32 MB or 128 MB memory DIMMs. At least one memory card must be present, and DIMMs are plugged in pairs on the card. System memory is protected with a SECDED (Single Error Correction, Double Error Detection) ECC code. The system is also designed to detect and correct data from a failed memory module on a memory DIMM.

The cpu cards and the memory cards are plugged into a system backplane board. The system backplane also has a pluggable interface to a separate I/O board. This separation of functions onto various pluggable cards allows for increased flexibility in cpu and memory upgrades.

I/O Subsystem

The I/O subsystem is contained on another board, which has a pluggable connection to the system backplane described above. The I/O subsystem interfaces to the cpu and memory through the memory/system control chips, via an IBM bus referred to as the 6XX-MX bus. The MX bus is 64-bits wide and operates at 66 MHz. Peak throughput is 528 MB/second.

The MX bus connects to two PCI bridge chips. One PCI bridge chip provides access to a 64-bit, 50 MHz PCI bus which has two slots. The other PCI bridge chip provides access to a 32-bit, 33 MHz PCI bus, which supports various I/O functions integrated onto the board, and three 32-bit PCI slots.

The Model 260 has several integrated I/O functions in order to reduce the number of separate I/O adapters a user needs to plug in. The system provides 10/100 Mbit Ethernet (10Base5, 100 BaseTX), Ultra SCSI for both internal and external SCSI connections, built in business audio, a 16550 UART interface to a tablet (for MCAD applications), and standard PC I/O such as mouse, keyboard, floppy disk, serial ports. The Model 260 has no ISA slots.

Graphics

The Model 260 has been designed for optimal 3D graphics performance with the IBM GXT3000P Graphics Accelerator. This is a 64-bit PCI adapter and plugs into one of the 64-bit PCI slots in the Model 260. It offers a number of advanced visualization features such as two-sided hardware lighting, 3D texture mapping, video scaling, and stereo-in-a-window. It continues IBM's hardware leadership in native hardware support for multiple application program interfaces (API's), OpenGL 1.2, graPHIGS, and X11R6. The GXT3000P requires AIX 4.3.2.

The GXT3000P design, technology, and performance on the Model 260 are more fully described in a companion document [2].

IPL Firmware and RTAS

The Model 260 has 1 MB of flash memory on the I/O board which contains code executed by the processors during IPL and also during various AIX interactions with the system hardware. The RS/6000 Platform Architecture (RPA) defines a set of code called Run Time Abstraction Software (RTAS), which provides a standard interface to AIX for hardware specific operations such as I/O configuration, time of day functions, power management, various initialization functions, capture and display of hardware indicators and error information, and other system functions. This RTAS code resides in the flash memory with the IPL code.

Integrated Service Processor

There is an integrated service processor which resides on the I/O board. This service processor performs system initialization or re-initialization functions, and its system error recovery and diagnostic functions give the Model 260 a high level of availability. The service processor is designed to save the state of the system to 128 KB of nonvolatile memory in the unlikely event of an error condition, to support subsequent diagnostic and recovery actions taken by other system firmware and the AIX operating system. The service processor also provides remote access to a system administrator or for a product engineer performing diagnostics on the system.

The service processor is a Motorola 68307 microprocessor. Its programs are stored in 512 KB of flash memory, and it uses 512 KB of SRAM memory to execute its programs. It has JTAG (serial scan) access to the latches and registers of the processor, memory/system control chips, I/O control chips, and other key hardware in the system.

Packaging.

Figure 2 shows the packaging layout of the Model 260 without covers. The system provides one diskette bay, two 3.5" full-high hard disk bays, and two 5.25" media bays. A 1.44 MB floppy diskette drive, one 32X CD-ROM, and one hard disk (base is 4.5 GB) come pre-configured with every system, leaving one hard disk bay and one media bay available for customer expansion. Any devices in the media bays are connected by cable to the internal Ultra SCSI chip.

The Model 260 supports the installation of 5 full size PCI adapter cards. Two of the slots are 64-bits.

The operator panel which is behind the front door has a 2 x 16 backlit LCD to provide system status and diagnostic information. Microphone and headphone jacks are built into the operator panel.

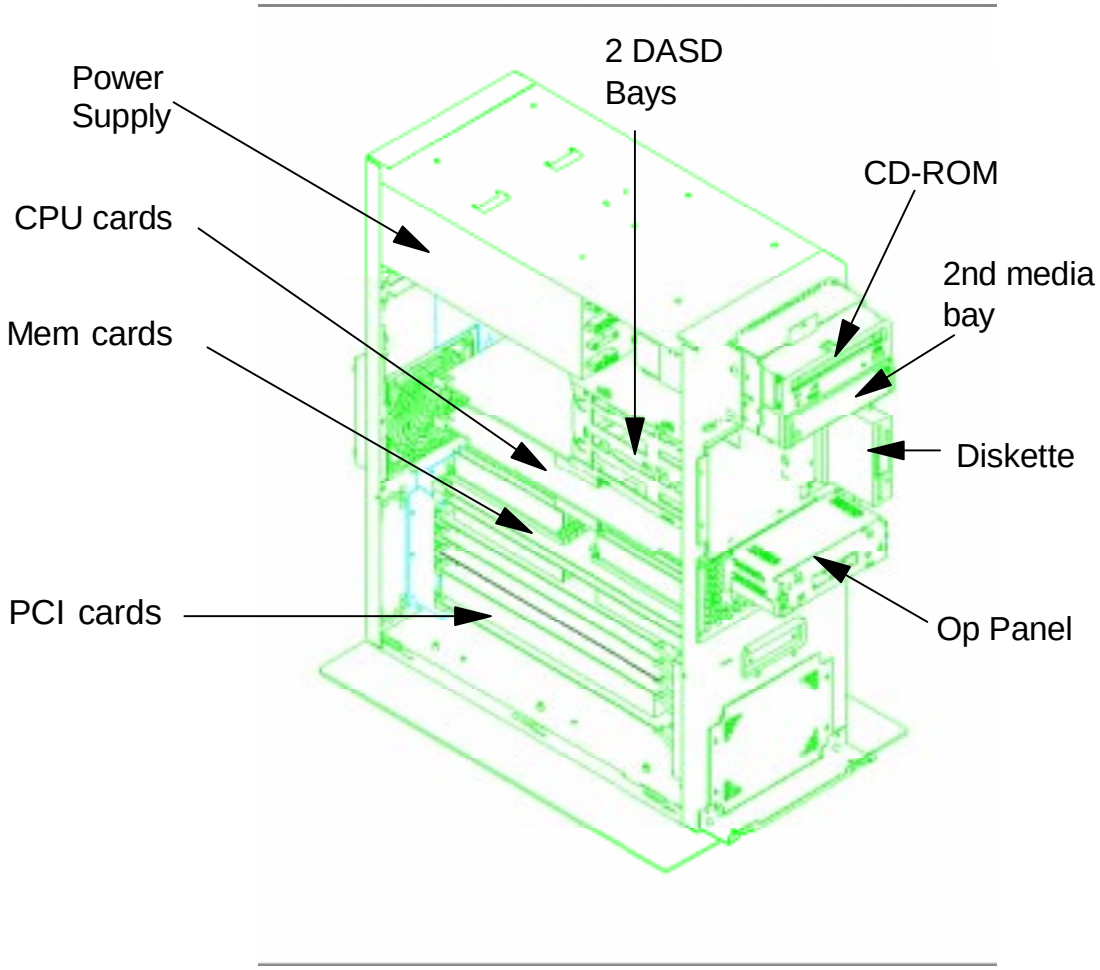

Figure 2. Model 260 Mechanical Package Layout

Summary

The RS/6000 Model 260 combines superior floating point computational power and high bandwidth memory and I/O subsystems which result in industry leading graphics performance and processing power for complex design and analysis. It was designed to high performance, affordable cost, and excellent reliability and availability standards to satisfy the demands of both workstation and entry server users. With its ability to upgrade to additional processors and memory, it can grow as the user's workload grows and will ensure the customer's investment is protected as new technologies become available.

References

[1] POWER3: Next Generation of 64-bit PowerPC Processor Design, IBM, 1998
[2] The GXT3000P Graphics Accelerator. IBM. 1998

Notes

The SPEC benchmarks reflect the performance of the microprocessor, memory architecture, and compiler of the tested system: SPECfp95 component level benchmark measures floating-point performance. Result is the geometric mean of ten tests, all written in FORTRAN, included in the CFP95 benchmark suite.

Biographies

David Tjon is the Product Manager of the RS/6000 Model 260. He is a member of the IBM Server Group, Austin, Texas.