

TECHNOLOGY BRIEF

November 1998

Compaq Computer Corporation

ISSD Technology Communications

CONTENTS

Introduction	3
Architecture Overview	3
Dual Memory Controllers	4
Dual-Peer PCI Buses	6
Dual Channel SCSI Buses	6
AGP Support	7
Optimized Multiprocessing Support	8
Dual Independent Bus Architecture	8
System Concurrency	9
Alternative Architectures	10
Typical NT/X86 Architecture	10
Concurrent Multiport Architecture	11
Crossbar Switch Architecture	12
Unified Memory Architecture	13
Conclusion	13

Highly Parallel System Architecture in the Compaq Professional Workstation SP700

As memory-intensive applications for financial analysis, computer-aided design (CAD), computer-aided engineering (CAE), and digital content creation (DCC) place growing demands on system resources, system bandwidth continues to be a critical business issue. As a technology leader, Compaq anticipated the growing need for increased bandwidth and introduced the first-generation Highly Parallel System Architecture in the Compaq Professional Workstations 5100, 6000, and 8000. To provide even greater bandwidth and performance than before, Compaq has implemented an improved, second-generation Highly Parallel System Architecture in the Compaq Professional Workstation SP700. The SP700 is specifically designed for users requiring uncompromising levels of performance and scalability to run demanding applications.

This technology brief describes the second-generation Highly Parallel System Architecture and differentiates it from other architectures used in X86 systems.

Please direct comments regarding this communication to the ISSD Technology Communications Group at this Internet address: TechCom@compaq.com

COMPAQ

NOTICE

The information in this publication is subject to change without notice and is provided "AS IS" WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OF THIS INFORMATION REMAINS WITH RECIPIENT. IN NO EVENT SHALL COMPAQ BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE OR OTHER DAMAGES WHATSOEVER (INCLUDING WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION OR LOSS OF BUSINESS INFORMATION), EVEN IF COMPAQ HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

This publication does not constitute an endorsement of the product or products that were tested. The configuration or configurations tested or described may or may not be the only available solution. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal state or local requirements.

Compaq, Contura, Deskpro, Fastart, Compaq Insight Manager, LTE, PageMarq, Systempro, Systempro/LT, ProLiant, TwinTray, ROMPaq, LicensePaq, QVision, SLT, ProLinea, SmartStart, NetFlex, DirectPlus, QuickFind, RemotePaq, BackPaq, TechPaq, SpeedPaq, QuickBack, PaqFax, Presario, SilentCool, CompaqCare (design), Aero, SmartStation, MiniStation, and PaqRap, ProSignia, Concerto, Vocalyst, and MediaPilot are registered with the United States Patent and Trademark Office.

Change is Good, Compaq Capital, Colinq, Armada, SmartQ, Counselor, CarePaq, Netelligent, Smart Uplink, Extended Repeater Architecture, Scalable Clock Architecture, QuickChoice, Systempro/XL, Net1, LTE Elite, PageMate, SoftPaq, FirstPaq, SolutionPaq, EasyPoint, EZ Help, MaxLight, MultiLock, QuickBlank, QuickLock, UltraView, Innovate logo, and Compaq PC Card Solution logo are trademarks and/or service marks of Compaq Computer Corporation.

Microsoft, Windows, Windows NT, Windows NT Advanced Server, SQL Server for Windows NT are trademarks and/or registered trademarks of Microsoft Corporation.

Pentium is a registered trademark of Intel Corporation.

Other product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

©1998 Compaq Computer Corporation. All rights reserved. Printed in the U.S.A.

**Highly Parallel System Architecture in the
Compaq Professional Workstation SP700**

First Edition (November 1998)
ECG067/1198

INTRODUCTION

Companies that run applications for demanding tasks such as financial analysis, CAD, CAE, and DCC have a critical need for greater system bandwidth. As a technology leader, Compaq anticipated the growing need for increased bandwidth and pursued possible solutions. With the Professional Workstations 6000 and 8000, Compaq was one of the first computer companies to implement a unique new architecture that provided the greatest bandwidth possible for X86 systems running such demanding applications under the Windows NT operating system. That high bandwidth was the result of a highly parallel system architecture design.

Based on high-performance Intel Pentium Pro and Pentium II processors, along with the memory and input/output (I/O) subsystems that support them, the first-generation Highly Parallel System Architecture comprised parallel processors, parallel memory controllers, and parallel I/O. Since the June 1997 announcement of the Professional Workstations 6000 and 8000, Compaq engineers have studied, tested, and improved the Highly Parallel System Architecture to provide even greater bandwidth and performance than before. This second-generation technology, implemented in the Compaq Professional Workstation SP700, is based on Pentium II Xeon processors and includes dual memory controllers, dual-peer PCI buses, dual Wide-Ultra SCSI controllers, accelerated graphics port (AGP) support, and support for optimized multiprocessing.

This brief will explain the elements of the second-generation Highly Parallel System Architecture implementation in the Professional Workstation SP700. In addition, it will define the benefits of each of the elements and compare other architectures. For more information about first-generation Highly Parallel System Architecture, see the technology brief *Highly Parallel System Architecture for Compaq Professional Workstations 5100, 6000, and 8000*, document number ECG066/1198.

ARCHITECTURE OVERVIEW

The Highly Parallel System Architecture increases overall system bandwidth to improve performance in demanding workstation applications by using:

- *Multiple data paths* that deliver a high degree of parallelism to key subsystems such as memory and I/O. Memory and I/O requests are processed in parallel, reducing system bottlenecks and delivering higher application performance.
- *Large, high-speed data buses* that increase throughput to critical subsystems such as processor, memory, and cache, resulting in uncompromising application performance.
- *Balanced system resources.* Critical system resources reside on separate buses to help balance throughput, improve system efficiency, and increase performance.

The Professional Workstation SP700 implements the second-generation Highly Parallel System Architecture and addresses the need for even greater overall system performance by increasing bandwidth to more subsystems than before. For example, the dual memory controllers featured in the SP700 implementation deliver a peak aggregate memory bandwidth of 1.6 GB/s, up from the first-generation bandwidth of 1.07 GB/s. The SP700 also supports dual-peer PCI buses that provide an aggregate PCI bus bandwidth of up to 267 MB/s and dual Wide-Ultra SCSI controllers that deliver an aggregate bandwidth of 80 MB/s. The AGP support in the SP700 improves graphics performance by providing a dedicated path to system memory with an effective data transfer rate of 533 MB/s. In addition, optimized multiprocessing support gives the SP700 the ability to handle eight simultaneous transactions and uses the Pentium II Xeon dual independent bus architecture. The block diagram in Figure 1 depicts how these subsystems are carefully crafted into a standards-based, high-performance architecture. The following sections discuss each of these subsystems in detail.

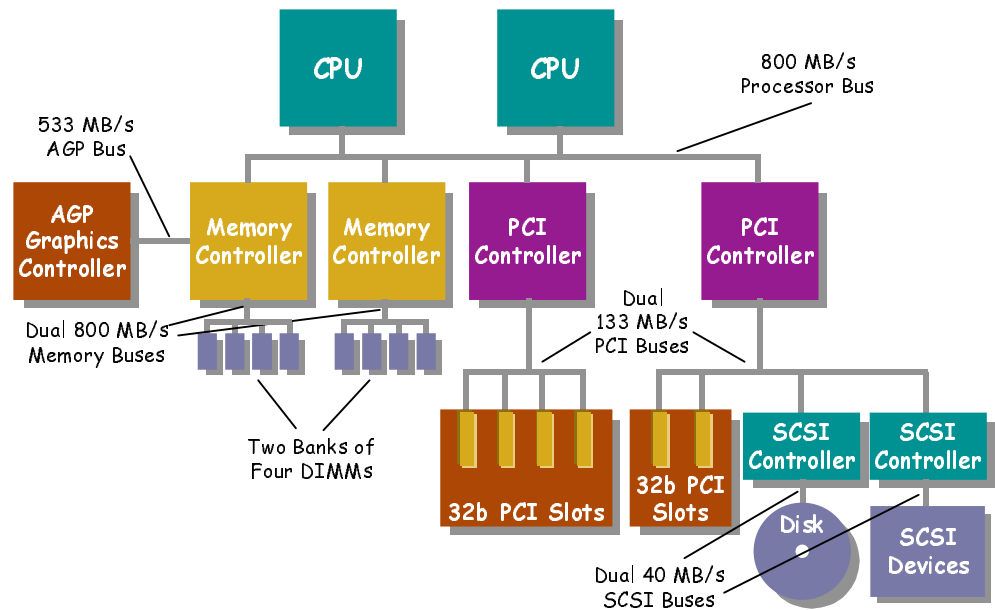


Figure 1. Highly Parallel System Architecture as implemented in the Compaq Professional Workstation SP700.

DUAL MEMORY CONTROLLERS

Like the first-generation Highly Parallel System Architecture, the second-generation implementation in the SP700 also uses two independently operating memory buses. The dual memory controllers employed in the Highly Parallel System Architecture provide twice the memory bandwidth of the single memory controller subsystem used in other Windows NT/X86 architectures. Dual memory controllers also enable users to optimize performance of the SP700's memory subsystem by distributing the DIMMs equally across both controllers.

The Compaq Professional Workstation SP700 uses industry-standard, 100-MHz Registered ECC SDRAM DIMMs. The DIMMs connect to the memory controller, which is connected to the processor by the front side bus. The front side bus in the second-generation Highly Parallel System Architecture runs at 100 MHz instead of the 66 MHz used in the previous generation. This allows the dual memory controllers to deliver uncompromising performance with an aggregate peak bandwidth of 1.6 GB/s (800 MB/s per controller). The greater bandwidth results in improved system responsiveness, especially in memory-intensive applications such as Pro/Engineer and NASTRAN. Additionally, since the memory bus runs at the same speed as the front side bus, it delivers a true pipeline for processor-to-memory transfers and alleviates any bottlenecks that a system with a slower memory bus would have.

The Registered DIMMs used in the SP700 allow the Highly Parallel System Architecture to offer greater memory capacity. Registered DIMMs use synchronous buffer logic chips on their control lines to reduce loading on the system board and eliminate performance degradation typically encountered with unregistered DIMMs. The buffering action increases the maximum number of memory chips on a DIMM. Unregistered DIMMs do not use buffer logic chips, which results in a reduction in the maximum number of memory chips on the same system board. Registered memory allows better memory expandability—up to 512 MB per DIMM, twice that of unregistered DIMMs—allowing Compaq to support memory-intensive applications even better than before.

DIMM – Dual Inline Memory Module

ECC – Error Checking and Correcting

SDRAM – Synchronous Dynamic Random Access Memory

In addition, the Highly Parallel System Architecture provides customers with more flexible memory configurations. Typical architectures support only up to four DIMM sockets, compared to the eight-socket configuration in the Compaq Professional Workstation SP700. More sockets allow users to reach larger memory capacities without having to use more expensive memory technologies. For example, the SP700 could be configured with 1 GB of RAM by using eight 128-MB DIMMs instead of four 256-MB DIMMs. Currently, one 256-MB DIMM costs 50 percent more than two 128-MB DIMMs. Compaq provides even more flexibility in the SP700 by using SDRAM DIMMs, which allow users to upgrade memory in single DIMM increments rather than in pairs.

These technologies allow the Compaq Professional Workstation SP700 to support up to 4 GB of system memory (using eight 512-MB DIMMs)—twice the memory expandability of other systems. Larger memory expandability gives Compaq workstations greater ability to run memory-intensive applications.

The Highly Parallel System Architecture is designed to improve memory responsiveness in other ways as well. For example, the SP700 supports CL2 memory. CL2 memory allows the Highly Parallel System Architecture to reduce the initial memory access by one clock. The CL2 support provides for a timing of 9-1-1-1, where initial 8-byte reads take nine clocks and the remaining 8-byte reads take only one clock. Other architectures use a 10-1-1-1 timing. This reduction in the number of clocks required for the initial read can decrease total access time by eight percent and provide a more efficient overall system.

Additionally, Compaq included a new paging architecture in the SP700's dual memory controllers that increases the number of memory pages that can be open simultaneously. Increasing the number of pages that can be open increases the chances that memory requests can be serviced faster. The first-generation Highly Parallel System Architecture allowed a maximum of two pages to be open at once. Typical SDRAM systems with a single memory controller allow a maximum of 32 pages of data to be open at once. Allowing 64 pages of memory to be open simultaneously, the Professional Workstation SP700 doubles the chance that a memory request can be serviced from the open memory pages.

Realizing that data integrity is just as critical as performance, Compaq includes hardware-based ECC memory scrubbing in the Highly Parallel System Architecture. With memory scrubbing, single-bit errors are repaired in the memory controller, so that subsequent reads from that location are correct. In addition, the system management software will alert the user of the problem and continue memory scrubbing until the problem can be eliminated.

The new memory architecture in the Professional Workstation SP700 has several benefits over the previous EDO DRAM technology. Key benefits include:

- *Increased performance.* Customers will experience greater processor responsiveness with 100-MHz SDRAM technology. New generations of SDRAM support new high-speed Pentium II Xeon 450-MHz technologies that EDO DRAM cannot easily support.
- *Faster bus speeds.* SDRAM can run up to 100 MHz, while the maximum bus speed that EDO DRAM can run is 66 MHz. SDRAM is, therefore, better suited to handle the increasing complexity of graphics programs.

For additional information on SDRAM technology, please read the technology brief *Industry Shift Toward Synchronous DRAM Technology*, document number ECG059/1098.

*EDO – Extended Data
Out*

DUAL-PEER PCI BUSES

Some workstation applications require large I/O bandwidth. For example, NASTRAN requires significant amounts of both I/O and memory bandwidth. Another example includes video editing, which requires significant disk and I/O bandwidth. Such applications can take full advantage of the new Highly Parallel System Architecture.

The new architecture features two independently operating PCI buses (that is, peer PCI buses), each running with a peak throughput of 133 MB/s. Together they provide a peak aggregate I/O bandwidth of 267 MB/s.

Because each PCI bus runs independently, it is possible to have two PCI bus masters transferring data simultaneously. In systems with two or more high-bandwidth peripherals, optimum performance can be achieved by splitting these peripherals evenly between the two PCI buses.

The new architecture also includes an I/O cache that improves system concurrency, reduces latency for many PCI bus master accesses to system memory, and makes more efficient use of the processor bus. The I/O cache is a temporary buffer between the PCI bus and the processor bus. It is controlled by an I/O cache controller. When a PCI bus master requests data from system memory, the I/O cache controller automatically reads a full cache line (32 bytes) from system memory at the processor transfer rate (800 MB/s) and stores it in the I/O cache. If the PCI bus master is reading memory sequentially (which is very typical), subsequent read requests from that PCI bus master can be serviced from the I/O cache rather than directly from system memory. Likewise, when a PCI bus master writes data, the data is stored in the I/O cache until the cache contains a full cache line. Then the I/O cache controller accesses the processor bus and sends the entire cache line to system memory at the processor bus rate. The I/O cache ensures better overall PCI utilization than other implementations, which is important for high-bandwidth peripherals such as disk controllers.

In addition to doubling the I/O bandwidth, the dual-peer PCI buses also allow greater system expandability by supporting up to 12 PCI devices—twice the number supported on single-bus implementations. This allows the Professional Workstation SP700 to deliver six available PCI-based I/O expansion slots while also integrating other PCI components, such as the SCSI and network controllers, on the system board.

Though other workstations may offer the ability to support more than six PCI devices, they achieve this support through the use of a PCI bridge. Although the bridge does extend the PCI bus and allows more devices to be connected, it still uses a single bus implementation with a maximum bandwidth of 133 MB/s. By extending the bus to accept more devices, this design can actually cause greater traffic on the PCI bus and may degrade performance.

DUAL CHANNEL SCSI BUSES

The Compaq Professional Workstation SP700 uses dual, independent channel Wide-Ultra SCSI controllers with an aggregate bandwidth of 80 MB/s (40 MB/s per bus). This dual bus architecture allows users to balance the disk subsystem workload by dividing high-performance peripherals between the two buses. In addition, performance of a system with a mixture of low- and high-performance devices can be optimized by separating lower performance devices, such as tape backup devices, from high-performance devices, such as 10,000-rpm hard drives and RAID arrays.

• **AGP SUPPORT**

• The Highly Parallel System Architecture in the Professional Workstation SP700 includes an AGP graphics controller that provides support for the AGP 2X specification. AGP is a new industry-standard, high-performance bus interface designed specifically for graphics applications. The specification was developed as a solution to alleviate the strain on the PCI bus caused by continually increasing graphics demands.

• AGP provides a dedicated 66-MHz PCI-based bus connection between the graphics controller, processor bus, and main memory. With AGP, some 3D rendering data structures, including textures, are moved from local graphics memory to main system memory. AGP allows the video processor to access main system memory directly, which frees the PCI bus for other performance-critical peripherals. This lowers graphics memory requirements and reduces overall system cost. The constraint that graphics memory puts on texture size is also eliminated, allowing graphics applications to use larger texture maps to create higher quality and more realistic images.

• AGP 2X uses a technique called double pumping that allows the graphics controller to transfer data on both clock edges, providing a theoretical 533-MB/s graphics bandwidth. Using the bandwidth of a 100-MHz system memory bus more efficiently, AGP 2X transfers up to four times more data per second than PCI at 133 MB/s.

• To access the texture maps in main memory more quickly, AGP uses a technique called Direct Memory Execute. Direct Memory Execute effectively connects main memory to the AGP chip set, allowing speedy access to texture data. Without Direct Memory Execute, the graphics controller can access and manipulate only the texture data that fits in graphics memory. That limited amount of data is insufficient for larger textures. Managing this entire process is the Graphics Address Remapping Table (GART), which is built into the AGP chip set.

• Additionally, the Highly Parallel System Architecture delivers two memory controllers for AGP's use. On the Compaq Professional Workstation SP700, the AGP is directly connected to one of the memory controllers. If the graphics controller requires more memory to store textures than is available on this primary controller, the AGP bus will also use memory from the secondary controller. Although AGP devices can access memory on either memory controller, the fastest performance is achieved when memory is accessed from the primary controller. Furthermore, performance can be optimized by placing at least half the system memory on the primary memory controller.

• The GART and the dual memory controllers allow the Highly Parallel System Architecture to use noncontiguous system memory to map texture data, providing up to 2 GB for texture data, in contrast to 256 MB provided by other architectures. This enables the Professional Workstation SP700 to accommodate rigorous graphical environments with large texture data requirements while concurrently supporting the remaining system memory requirements.

• The second-generation Highly Parallel System Architecture also improves graphics performance by using read merging techniques. Read merging allows the AGP 2X system to combine requests to consecutive memory locations into a single cache line, thus improving graphics performance by reducing the number of reads from memory.

• For more information on AGP and GART, please read the technology brief *Accelerated Graphics Port Technology*, document number ECG081/0898.

OPTIMIZED MULTIPROCESSING SUPPORT

The Compaq Professional Workstation SP700 uses Intel's Pentium II Xeon processor. The new processor and chip set provide a 100-MHz front side bus, faster processor speeds (400-MHz and 450-MHz), and a large, full-speed L2 cache. Pentium II and Pentium II Xeon processor are both based on Intel's 0.25 micron processor technology, but the Pentium II Xeon processor-based workstations offer more processing power and scalability than comparable Pentium II and Pentium Pro processor-based workstations.

To gain the most from Pentium II Xeon processors, the Highly Parallel System Architecture is specifically engineered for optimized multiprocessing support. While each processor can issue up to four transactions on the processor bus, some architectures limit the processors to four concurrent transactions. In contrast, the Compaq Professional Workstation SP700 permits the processors to post up to eight simultaneous transactions (up to four for each processor). By handling more simultaneous transactions, the Professional Workstation SP700 provides higher utilization and scalability than competing workstations.

While most multiprocessing implementations take advantage of this support simply by adding an additional processor to an already existing desktop design, the Highly Parallel System Architecture takes multiprocessing to the next level by enhancing memory and I/O bandwidth as well. Multiprocessor systems designed without the Highly Parallel System Architecture will bottleneck as the multiple processors try to access other system resources that have not been enhanced to accommodate the additional data traffic. The Highly Parallel System Architecture significantly reduces these bottlenecks by incorporating enhanced subsystem resources, such as dual memory controllers and dual-peer PCI buses, to accommodate the increased data traffic from the multiple processors.

Dual Independent Bus Architecture

Complementing the Highly Parallel System Architecture, the Pentium II Xeon uses Intel's Dual Independent Bus architecture to provide two independent buses: a processor-to-cache bus and a processor-to-memory bus. The processor-to-memory bus runs at the front side bus speed of 100 MHz. The processor-to-cache bus speed depends on which processor is being used. With the Pentium II Xeon used on the Compaq Professional Workstation SP700, this bus runs at the processor speed (that is, 400 MHz in a 400-MHz processor).

This design delivers significantly more bandwidth than a single bus architecture processor because the buses can work independently, increasing throughput. And, as processor speeds increase, so will the speed of the processor-to-cache bus, which will allow performance to scale with speed.

SYSTEM CONCURRENCY

The Highly Parallel System Architecture in the Compaq Professional Workstation SP700 delivers a high degree of system concurrency. The biggest concurrency issue is parallel graphics and I/O. A total bandwidth of 1.33 GB/s (800 MB/s processor + 533 MB/s AGP) is required when the processor and AGP read from memory at the same time. A typical system with only one memory controller would encounter a bottleneck using the maximum 800-MB/s bandwidth (Figure 2). The dual memory controllers in the Professional Workstation SP700 deliver a 1.6-GB/s bandwidth, which more than accommodates the 1.33 GB/s required for AGP and processor concurrency (Figure 3).




Figure 2. A typical system with a single 800-MB/s memory controller cannot handle the 1.3-GB/s bandwidth required for concurrent processor and AGP transactions.




Figure 3. The Highly Parallel System Architecture in the Compaq Professional Workstation SP700 has more than enough memory bandwidth (1.6 GB/s) to accommodate the 1.3 GB/s required for concurrent processor and AGP transactions.

Another concurrency issue occurs when two processors attempt to access memory simultaneously. A typical system with only one memory controller and a bandwidth of only 800 MB/s would encounter a bottleneck. Again, Compaq’s dual memory controllers deliver the 1.6-GB/s bandwidth required by the processors.

The dual-peer PCI buses in the Professional Workstation SP700 also allow concurrent processor-to-PCI requests. For example, a typical system that has only one PCI bus can only handle one PCI request at a time. With the Highly Parallel System Architecture in the SP700, two processor-to-PCI requests can be processed simultaneously.

Compaq has taken concurrency even further and ensured that the Highly Parallel System Architecture includes concurrent dual PCI buffers. Because the PCI buses allow a bandwidth of only 133 MB/s each, bottlenecks sometimes occur. For example, when the processor—with an 800-MB/s bandwidth—makes a PCI request, a bottleneck would occur because the processor has to wait for the PCI bus to transfer the data. The PCI buffers help to alleviate some of the

bottleneck. Compaq has ensured that the concurrent buffers are adequate to accommodate PCI traffic, and they are independent of other buffers.

ALTERNATIVE ARCHITECTURES

Compaq is implementing this second-generation Highly Parallel System Architecture because other architectures do not provide equal levels of bandwidth, performance, expansion, and cost effectiveness. This section describes four of these alternative architectures.

Typical NT/X86 Architecture

Most workstations in the NT/X86 market support two processors to process instructions concurrently (Figure 4). Overall system bandwidth is limited in such systems, however, because each processor must compete with the other for access to subsystems such as memory and disk. The Highly Parallel System Architecture reduces these limitations by incorporating enhanced subsystem resources, such as dual memory controllers and dual-peer PCI buses, to accommodate the increased data traffic from the multiple processors.

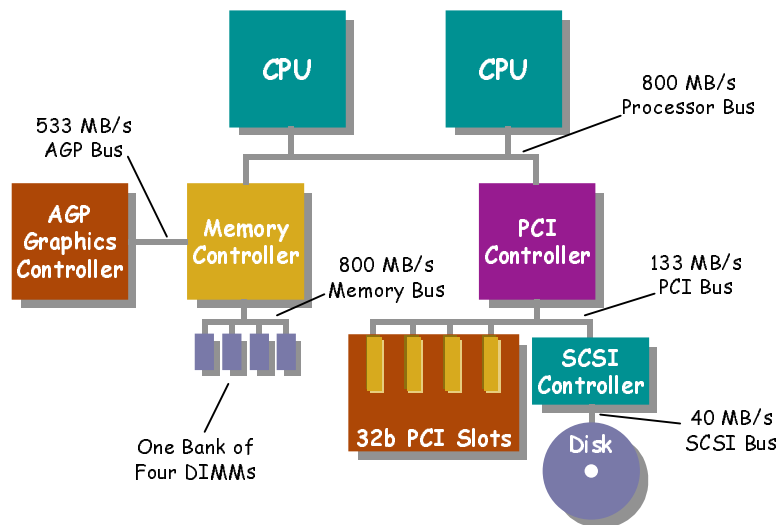


Figure 4. Typical architecture for an X86 computer running Microsoft Windows NT.

Additionally, even though other NT/X86 systems support up to two Pentium II or Pentium II Xeon processors, the Highly Parallel System Architecture is specifically engineered for optimized multiprocessing support. While typical NT/X86 architectures limit the processors to four concurrent transactions, the Professional Workstation SP700 can post up to eight simultaneous transactions. As a result, it improves scalability by permitting both processors to operate at peak utilization in a multiprocessing environment.

Traditional memory architectures use a single memory controller through which all memory requests are processed. The maximum bandwidth of these memory subsystems is only 800 MB/s. Compaq's Highly Parallel System Architecture, on the other hand, employs dual memory controllers that can process memory requests in parallel. This design allows memory bandwidth to reach up to 1.6 GB/s—twice the bandwidth of other NT/X86 systems.

Moreover, the memory configuration in the Compaq Professional Workstation SP700 provides customers with more flexibility and better expandability than traditional memory configurations. While typical systems support only up to four DIMM sockets, the SP700 delivers an 8-socket

configuration. As a result, the Highly Parallel System Architecture can support up to 4 GB of system memory—twice that of traditional memory architectures.

The Highly Parallel System Architecture also improves memory responsiveness in other ways. For example, the SP700 retrieves data from registered DIMMs in three clock cycles compared with four clocks for traditional architectures. Also, the Highly Parallel System Architecture doubles the number of pages that can be cached—64 pages compared with 32 pages for traditional architectures.

Typical NT/X86 systems use a single PCI controller that delivers an I/O bandwidth of 133 MB/s. By using dual-peer PCI controllers, the Highly Parallel System Architecture delivers twice the I/O bandwidth (267 MB/s) of most other NT/X86 architectures. The dual controllers support as many as four more PCI devices than other architectures, without the use of performance-degrading PCI bridge chips. For example, the Compaq Professional Workstation provides six PCI slots, in addition to three embedded PCI devices.

The SP700's AGP support surpasses that of typical systems, as well. While most traditional architectures support AGP 2X, the Highly Parallel System Architecture provides up to 2 GB for texture memory, in contrast to the 256 MB allowed by other implementations.

Concurrent Multiport Architecture

Intergraph's new Concurrent Multiport Architecture (CMA) is amazingly similar to the Highly Parallel System Architecture used in the Compaq Professional Workstation SP700. Both architectures use the second-generation Reliance supporting chip set, which provides dramatically improved system throughput and application performance. The Reliance chip set provides optimized buffers and delivers new levels of system concurrency.

Like the Highly Parallel System Architecture, CMA incorporates dual memory controllers that deliver a peak bandwidth of 800 MB/s per controller. However, the Highly Parallel System Architecture is engineered to deliver better performance and higher throughput. Intergraph's architecture allows a maximum of 48 pages of memory to be open at once, while Compaq's architecture permits up to 64. Furthermore, Intergraph's CMA supports a maximum memory configuration of only 3 GB, while the Highly Parallel System Architecture can support up to 4 GB of system memory.

Another notable difference is Intergraph's use of the Pentium II processor in CMA instead of the more powerful Pentium II Xeon found in the SP700. With a 100-MHz front side bus, faster processor speeds, and a large, full-speed L2 cache, Pentium II Xeon processor-based workstations offer more processing power and scalability than comparable Pentium II processor-based workstations.

Crossbar Switch Architecture

A crossbar switch architecture provides multiple, independent paths to system memory. As Figure 5 illustrates, individual paths can be established to memory from each processor or I/O bus. Thus, a crossbar switch can avoid contention of multiple memory requests on a given bus. This style of crossbar switch is used in the Sun Microsystems Ultra Port Architecture (UPA).

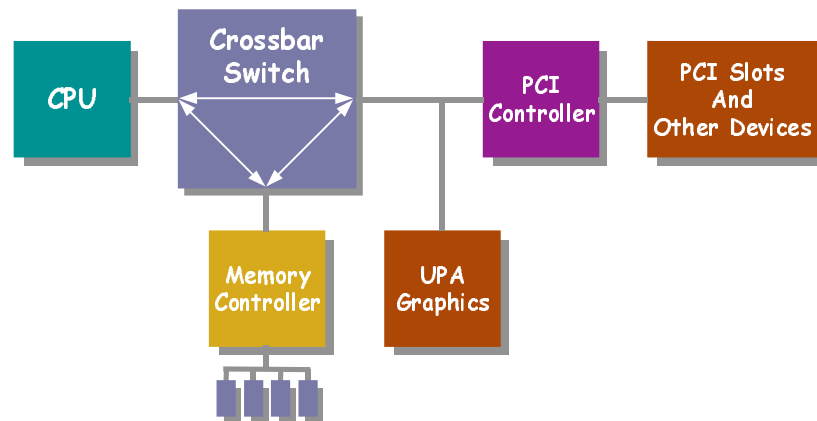


Figure 5. Crossbar switch architecture.

Like Intergraph's CMA, Sun's UPA has many similarities with the Highly Parallel System Architecture in the SP700. They both have processor, memory, PCI, and graphics ports, but Sun's UPA provides a peak memory bandwidth of 1.9 GB/s. However, this is implemented with a single memory controller using older EDO memory. EDO memory only supports one open memory page at a time and may actually perform slower than the SDRAM used in the SP700, which supports up to 64 open pages of memory. Also, Sun's UPA is limited to 2 GB of system memory, half that of the SP700.

By allowing separate paths to system memory, this style of crossbar switch can improve performance of both I/O traffic and processor cycles. A crossbar switch for a single processor bus, memory bus, and I/O bus can be implemented cost effectively. However, a crossbar switch is an expensive solution in a system with several buses. The reason is that all the buses must go into a single chip that has sufficient pins for each bus. This requires a large and expensive chip when several buses are implemented in the crossbar switch. A crossbar switch supporting the processor bus, two PCI buses, a graphics bus, and two memory buses is not cost effective with today's silicon technology. Sun avoids this problem by forcing the PCI buses and graphics bus to all share the same bus into the crossbar switch. While this reduces the cost, it negates many of the intended benefits of the switch.

Compaq implemented a crossbar switch to give the AGP controller a direct path to main memory. AGP is frequently the highest bandwidth peripheral and is better served with a direct path to main memory rather than sharing bandwidth with PCI as is done with Sun's UPA. However, Compaq chose not to use a crossbar switch in any other subsystem because a better architectural solution was possible. First, the Pentium II Xeon processor bus is capable of processing up to eight transactions simultaneously. Since the I/O cycle can be run simultaneously with processor cycles on the shared processor bus, independent paths to memory are not as critical. Furthermore, the I/O cache significantly reduces the amount of bandwidth required by each PCI bus on the processor bus. Second, having dual-peer PCI buses and two memory buses can produce significant performance increases. The new architecture using dual-peer PCI buses and dual memory controllers on a shared processor bus gives better performance than a crossbar switch with a single memory controller and a single PCI bus. Moreover, the higher performance comes

at a price point well below that of a crossbar switch with dual memory controllers and dual PCI buses.

Unified Memory Architecture

Silicon Graphics, Inc. touts the Unified Memory Architecture (UMA) used in their O2 workstation. Although UMA provides a cost-effective system, it does so by sacrificing performance. With SGI's UMA, the processor and graphics controller share one memory pool that is connected by a single bus with a peak bandwidth of up to 2.1 GB/s (Figure 6).

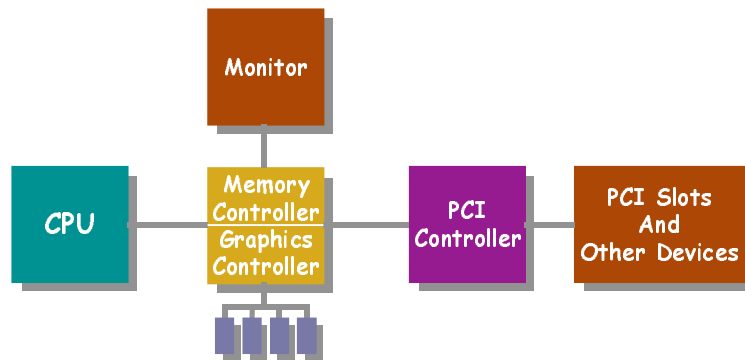


Figure 6. Unified Memory Architecture used in the O2 workstation from Silicon Graphics, Inc.

Because the processor, the graphics controller, and the monitor compete for access to memory, however, this architecture does not deliver as much actual throughput as the new Highly Parallel System Architecture. For example, refreshing the monitor at 85 Hz at a screen resolution of 1280 x 1024 true color requires that data be transferred to the monitor at a rate of 334 MB/s. The large amount of memory bandwidth consumed by monitor refreshing is not available to the processor and graphics controller for other tasks. In contrast, the second-generation Highly Parallel System Architecture incorporates a graphics controller that supports AGP, providing a dedicated path to system memory with a theoretical 533-MB/s bandwidth.

CONCLUSION

As business-critical applications become more demanding, the need for more bandwidth becomes increasingly important to customers. The second-generation Highly Parallel System Architecture implemented in the Compaq Professional Workstation SP700 improves overall system performance by delivering increased bandwidth to critical subsystems, including memory and I/O. Multiple data paths, large high-speed data buses, and balanced system resources make the Highly Parallel System Architecture the best choice for delivering uncompromising performance in today's demanding applications. Compaq once again demonstrates technology leadership in providing innovative computing solutions to meet the needs of customers.