

TECHNOLOGY BRIEF

November 1998

Compaq Computer Corporation

ISSD Technology Communications

CONTENTS

Introduction	3
Architecture Overview	3
Advanced SMP	3
Pentium Pro Processor	4
Pentium II Processor	4
Dual Memory Buses	5
Dual-Peer PCI Buses	8
Multiple Drives	8
Alternative Architectures	9
Typical NT/X86 Architecture	9
Unified Memory Architecture	9
Crossbar Switch Architecture	11
AGPset Architecture	12
Conclusion	13

Highly Parallel System Architecture for Compaq Professional Workstations 5100, 6000, and 8000

As critical applications for financial analysis, computer-aided design (CAD), computer-aided engineering (CAE), and digital content creation (DCC) place growing demands on system resources, increasing system bandwidth becomes a critical business issue. After evaluating available system architectures, Compaq determined that only a new, highly parallel system architecture could provide the required levels of performance, processor and I/O expandability, and bandwidth to satisfy the needs of workstation users. Compaq is therefore implementing a new architecture that delivers the greatest bandwidth available today for systems running such demanding applications under the Microsoft Windows NT operating system.

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



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.

NetWare and Novell are registered trademarks and IntranetWare, NDS, and Novell Directory Services are trademarks of Novell, Inc.

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 for Compaq Professional Workstations 5100, 6000, and 8000

Second Edition (November 1998)
ECG066/1198

INTRODUCTION

Achieving greater system bandwidth is a critical issue to businesses running applications for such demanding tasks as financial analysis, computer-aided design (CAD), computer-aided engineering (CAE), and digital content creation (DCC). As a technology leader, Compaq anticipated the growing need for increased bandwidth and pursued possible solutions. As a result of its R&D efforts, Compaq is one of the first computer companies to implement a unique new architecture that provides the greatest bandwidth possible today for X86 systems running such demanding applications under the Microsoft Windows NT operating system. The high bandwidth results from a highly parallel system architecture design.

The Highly Parallel System Architecture consists of parallel processors, parallel memory controllers, and parallel input/output (I/O). This new architecture is based on high-performance Intel Pentium Pro and Pentium II processors along with the memory and I/O subsystems that support them. It includes a new standards-based memory architecture that provides significantly greater bandwidth than traditional memory architectures by using dual memory controllers. These memory controllers independently process memory requests in parallel, thereby effectively doubling the memory bandwidth of traditional architectures. The new architecture also supports dual-peer PCI buses that double I/O bandwidth and allow for more I/O expandability. The new architecture is being implemented initially in two Compaq products, the Compaq Professional Workstation 6000 and the Compaq Professional Workstation 8000.

This technology brief provides an overview of the new architecture and explains how it differs from other architectures used in X86 systems.

ARCHITECTURE OVERVIEW

Unlike any previous architecture used in X86 systems, the new architecture being implemented by Compaq incorporates a highly parallel system design. As the block diagram in Figure 1 indicates, the new architecture offers parallelism in multiple ways. It supports symmetric multiprocessing (SMP) with up to four Pentium Pro processors or two Pentium II processors. It supports dual-peer PCI buses to provide an aggregate PCI bus bandwidth of up to 267 MB/s. It also features dual memory controllers that deliver a peak aggregate memory bandwidth of 1.07 GB/s.

This high degree of hardware parallelism for critical subsystems such as processor, memory, and I/O maximizes system bandwidth to improve performance in demanding, resource-intensive applications. Hardware parallelism of this architecture can be increased even more if the system configuration includes an optional Compaq drive array controller that can increase performance in disk-bound applications by accessing data in parallel from multiple disk drives.

ADVANCED SMP

It is widely recognized that SMP has been supported on both RISC/UNIX workstations and servers for a number of years. Most people do not, however, associate Microsoft Windows NT Workstation and the applications that run on it with SMP. Windows NT Workstation does support SMP; and for running many demanding applications in finance, computer-aided design, computer-aided engineering, and digital content creation, workstation users can reap the benefits of SMP. Many simulations and analyses performed on workstations in technical computing environments require tremendous numbers of calculations that can take hours or even days to complete. Fortunately, many of these applications are multi-threaded or otherwise lend themselves well to multiprocessing. Therefore, a single user of a workstation can take full advantage of multiple processors.

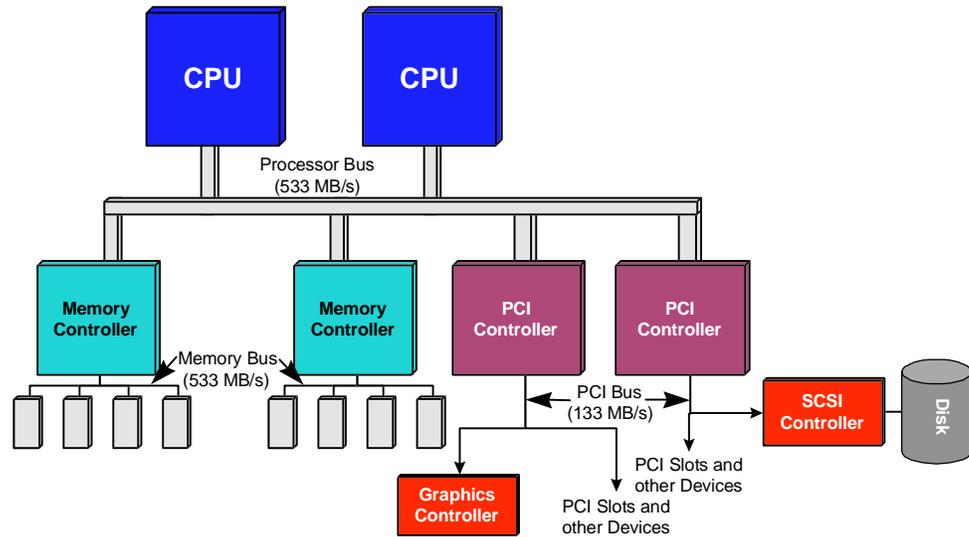


Figure 1. Block diagram of the Highly Parallel System Architecture as implemented in the Compaq Professional Workstation 6000. This architecture supports up to four processors in the Compaq Professional Workstation 8000.

The new Highly Parallel System Architecture implemented by Compaq provides excellent scalability for increased system performance. It supports the use of multiple Intel Pentium Pro processors or Intel Pentium II processors. Each of these processors has advantages for specific computing environments; therefore, Compaq is using each of these processors in new products being developed to meet the widest possible range of customer needs.

Pentium Pro Processor

The Pentium Pro processor family is Intel's current generation of processors for high-end desktops, workstations, and servers. This product family consists of processors running at clock speeds of 150 MHz to 200 MHz with level 2 (L2) memory cache sizes of either 256 KB or 512 KB. The Compaq Professional Workstation 8000 will use the 200-MHz Pentium Pro processor with an integrated 512-KB L2 cache that runs at the core processor speed of 200 MHz. The high-speed processor and cache provide top performance. Because the larger 512-KB L2 cache holds more instructions, it provides a higher cache hit rate than the 256-KB cache. The higher cache hit rate reduces memory bus traffic and therefore increases performance and scalability of the system.

Pentium II Processor

The Pentium II processor is the next generation P6 processor from Intel. Formerly code named *Klamath*, the Pentium II processor will be available in 233-MHz, 266-MHz, and 300-MHz versions, each with an integrated 512-KB L2 cache. Compaq is implementing 266-MHz and 300-MHz Pentium II processors in the Compaq Professional Workstation 6000.

The Pentium II processor provides some enhancements over the Pentium Pro processor, but it also has some limitations. The Pentium II processor incorporates Intel's MultiMedia Extensions (MMX) technology. MMX is the name for 57 multimedia instructions that Intel has added to its new generation of processors. This instruction set is expected to significantly improve the performance of processor-intensive multimedia applications that are MMX-aware. MMX is tailored to audio, video, and other multimedia tasks. An MMX-equipped workstation will use only one instruction to execute the same task that a Pentium Pro processor would perform using

up to 16 instructions. Because multimedia operations such as video and audio use a number of redundant instructions, MMX achieves some efficiencies by using a technique called SIMD (single instruction multiple data). SIMD reduces the required number of clock cycles by performing redundant instructions on multiple sets of data.

For its L2 cache, the Pentium II processor uses industry-standard SRAM (static random access memory). This implementation is sometimes referred to as a half-speed cache because the SRAM runs at half the core processor speed. Use of a half-speed L2 cache instead of a full-speed L2 cache improves manufacturability; however, it increases cache access time, which limits scalability and performance. Pentium II systems will support a maximum of two processors and are limited to 512 MB of addressable system memory. The Pentium II processor is capable of caching 512 MB of system memory. Adding more memory will significantly degrade system performance because the additional memory will not be cached. Therefore, to ensure application performance, the Compaq Professional Workstation 6000 will not boot if more than 512 MB of system memory is installed.

DUAL MEMORY BUSES

The new architecture also includes two independently operating memory buses, each running at a peak speed of 533 MB/s. Together they provide a peak aggregate memory bandwidth of 1.07 GB/s—two to four times the memory bandwidth of other X86/NT workstations. This high memory bandwidth is the key to delivering the highest performance levels of both single processor and SMP-aware applications.

Each memory bus is 144 bits wide and consists of 128 bits of data plus 16 bits for Error Checking and Correction (ECC). The new architecture uses buffered 60-ns Extended Data Out (EDO) DIMMs (Dual Inline Memory Modules), and the memory is interleaved. Interleaved memory takes advantage of the sequential nature of program execution to overcome delays resulting from Column Address Strobe (CAS) precharge. *CAS precharge* is the time a microprocessor must wait between back-to-back accesses to the same Dynamic Random Access Memory (DRAM) chip while the DRAM chip charges back up after a destructive read.

When interleaved memory is used, banks of DRAM are divided into two or more physically separate areas. Consecutive addresses are stored in different areas of a bank. This makes it possible for the next sequential read to begin on bank B while bank A is precharging from the previous read, and vice versa. Thus, interleaved memory can significantly increase memory throughput for sequential reads.

EDO memory is capable of transferring data every other clock cycle (30 ns for a 60-MHz bus). Non-EDO (Fast Page Mode) memory, on the other hand, is capable of transferring data every four clocks (60 ns for a 66-MHz bus). With the combination of interleaved memory and EDO memory, 128 bits of data (plus 16 bits ECC) are read at a time. As Figure 2 illustrates, the first 64 bits of the read go to the processor on one clock pulse, and the second 64 bits of data go on the next clock pulse. This enables the memory bus to transfer data at the peak data rate of the processor.

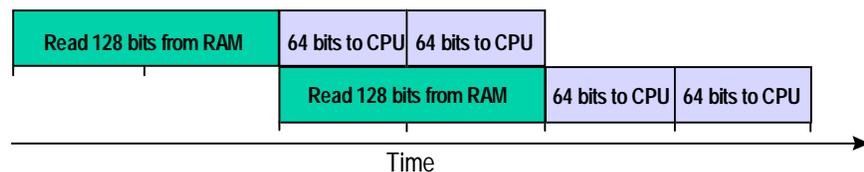


Figure 2. Interleaved EDO memory reads 128 data bits at a time and transfers data every other clock cycle.

Because each of the two memory buses in this architecture is capable of returning data to the processor at the full 533-MB/s speed of the processor bus, it could seem superfluous to have more memory bandwidth than can be sent across the processor bus. One could envision the system looking like a funnel, with the bottleneck being the processor bus. This is not true, however. To understand why, we must look at how DRAM memory works.

Accessing DRAM memory is relatively slow, in part because the DRAM cannot transfer data continuously. In fact, as Figure 3 illustrates, a typical memory cycle transfers data only about one third of the time. First, the DRAM must be precharged. Second, the memory address of the data being sought must be sent to the DRAM. Finally, the DRAM can transfer data.

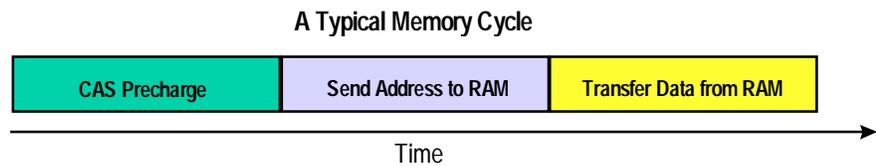


Figure 3. Because a DRAM memory cycle has three components, data is transferred only during about one third of the cycle.

In some cases DRAM can transfer data faster. If two sequential cycles query addresses on the same DRAM page, then the address for the second cycle can be sent during data transfer of the first cycle (Figure 4). Once continuous data transfer is established on both memory buses as illustrated in Figure 4, it is possible to achieve a peak aggregate memory bandwidth of 1.07 GB/s.

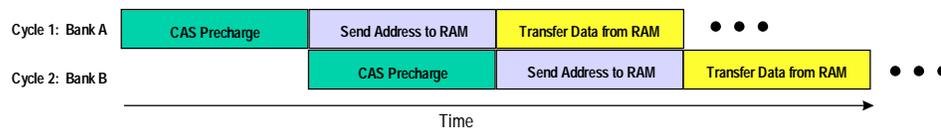


Figure 4. Basic timeline for sequential reads from the same page of DRAM.

While it is fairly common for a single processor to access consecutive memory locations, consecutive cycles in SMP machines are rarely sent to nearby addresses. Individual processors typically run programs and access data from very different areas of system memory. Thus, if processor 1 reads memory at one location and immediately thereafter processor 2 performs a read, it will usually be to a very different address. For this reason, most memory cycles in SMP machines look like the one in Figure 3 and transfer data only about one third of the time, yielding a more typical memory bandwidth of 177.7 MB/s. This is true of all DRAMs and all SMP machines; it is not unique to the new highly parallel architecture. By adding a second memory bus, the new architecture actually doubles typical consumption of the processor bandwidth in SMP machines:

$$177.7 \text{ MB/s memory bandwidth} \times 2 \text{ buses} = 355 \text{ MB/s}$$

To take full advantage of the two memory buses, at least two memory requests must be issued at a time, one on each memory bus. SMP Pentium Pro processors can issue up to eight cycles at a time, which increases the likelihood of having cycles run to both memory controllers. Because a single processor can issue up to four cycles, the dual memory controller can also boost the performance of a single processor.

Memory may be added to either bus individually and the system will continue to operate correctly. For peak performance, however, equal amounts of memory should be added to both buses at the same time. Because memory is interleaved between the two memory channels by pairs of DIMMs, best performance results from using the maximum number of DIMMs for a given

memory size and evenly splitting those DIMMs between the two memory channels. Table 1 is a DIMM configuration guide for optimizing performance of dual memory buses. The table shows matched memory sizes on both banks. Other memory configurations are valid.

TABLE 1: GUIDE FOR CONFIGURING DUAL MEMORY BUSES TO OPTIMIZE PERFORMANCE

Memory Size	Memory Bus 1 DIMMs	Memory Bus 2 DIMMs	Optimization Level *
32 MB	—	2 x 16 MB	1
64 MB	2 x 16 MB	2 x 16 MB	1
64 MB	—	4 x 16 MB	2
64 MB	—	2 x 32 MB	3
128 MB	4 x 16 MB	4 x 16 MB	1
128 MB	2 x 32 MB	2 x 32 MB	2
128 MB	—	4 x 32 MB	3
128 MB	—	2 x 64 MB	4
256 MB	4 x 32 MB	4 x 32 MB	1
256 MB	2 x 64 MB	2 x 64 MB	2
256 MB	—	4 x 64 MB	3
256 MB	—	2 x 128 MB	4
512 MB	4 x 64 MB	4 x 64 MB	1
512 MB	6 x 64 MB	2 x 64 MB	2
512 MB	8 x 64 MB	—	3
512 MB	—	4 x 128 MB	4
1 GB	8 x 64 MB	4 x 128 MB	1
1 GB	4 x 128 MB	4 x 128 MB	2
1 GB	8 x 128 MB	—	3
1 GB	—	4 x 256 MB	4
2 GB	4 x 256 MB	4 x 256 MB	1
2 GB	8 x 256 MB	—	2

* The degree of performance optimization is indicated by a numerical range. Level 1 represents the best performance. Levels 2, 3, and 4 indicate progressively lower performance.

The exact performance increase to be gained by optimizing the memory subsystem is highly dependent upon the application. Some applications will see less than 1 percent benefit, while others may see as much as 33 percent. In general, workstation applications with large data sets (for example, MacNeil Schwendler Corporation's NASTRAN Finite Element Analysis software) make better use of the dual memory controllers than most PC productivity applications with small data sets.

Users should also consider tradeoffs in system cost and memory expansion that may result from optimizing memory performance. In some cases, optimum memory performance can reduce the amount of available memory expansion. For instance, a cost-effective and performance-enhanced 128-MB configuration can be built with eight 16-MB DIMMs. Upgrading such a system to 512 MB, however, would require the addition of 128-MB DIMMs, which currently are not as cost effective as eight 64-MB DIMMs or as replacing some of the 16-MB DIMMs with 64-MB DIMMs.

DUAL-PEER PCI BUSES

Some workstation applications require large I/O bandwidth. For example, NASTRAN requires significant amounts of both I/O bandwidth and memory bandwidth. Other examples include visualization programs that make heavy use of the 3D-graphics controller. 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 at a peak speed of 133 MB/s. Together they provide a peak aggregate I/O bandwidth of 267 MB/s.

Since 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 (533 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 3D graphics.

In addition to doubling the I/O bandwidth, the dual-peer PCI buses can support more PCI slots than a single PCI bus, providing greater I/O expandability.

MULTIPLE DRIVES

The high level of hardware parallelism provided in the new architecture can be enhanced even more by adding multiple disk drives to the system. By using more than one disk drive, certain disk-oriented operations may run faster. For instance, NASTRAN data sets can grow into multiple gigabytes of data. Since this data cannot fit into physical memory, it is paged to the disk drive, which is then continuously accessed by the program as it performs its calculations on the data. To improve performance, a RAID-0 drive array can be used to increase disk performance. A RAID-0 drive array will access multiple drives as a single logical device, thereby allowing data to be accessed from two or more drives at the same time. However, RAID-0 does not implement fault management and prevention features such as mirroring, as other RAID levels do.

ALTERNATIVE ARCHITECTURES

Compaq is implementing this new Highly Parallel System Architecture because other architectures do not provide equal levels of bandwidth, performance, expansion, and cost effectiveness.

Typical NT/X86 Architecture

Most workstations in the NT/X86 market support two processors to process instructions concurrently (Figure 5). 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.

Traditional memory architectures use a single memory controller through which all memory requests are processed. Depending on implementation, the maximum bandwidth of these memory subsystems is either 267 MB/s or 533 MB/s. The actual memory throughput will be limited, however, by the same DRAM constraints identified earlier in the section “Dual Memory Buses.”

The new memory architecture that Compaq is implementing, on the other hand, employs dual memory controllers that can process memory requests in parallel. This design allows memory bandwidth to reach up to 1.07 GB/s—two to four times the bandwidth of other NT/X86 systems. Furthermore, with dual-peer PCI buses, high-bandwidth peripherals can be placed on separate PCI buses.

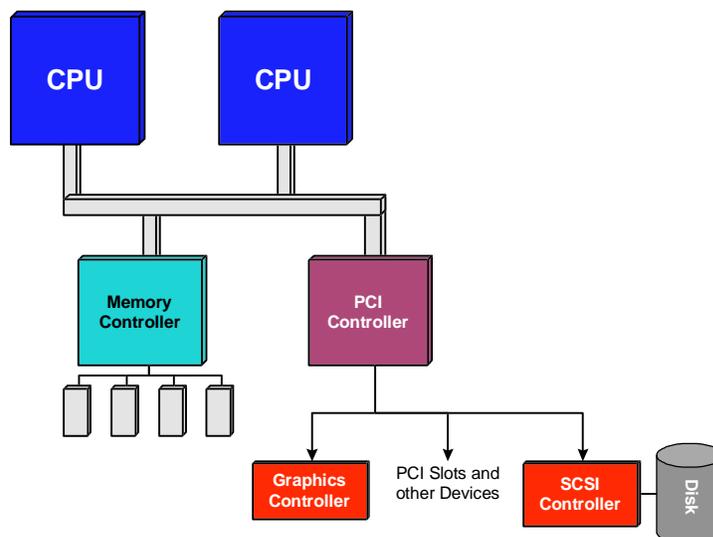


Figure 5. Typical architecture for an X86 computer running Microsoft Windows NT.

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). As noted earlier in the section “Dual Memory Buses,” the actual memory throughput will be limited by DRAM constraints. The graphics controller stores its frame buffer, Z-buffer, and textures in the common memory pool. 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, all graphics cards used in Compaq workstations use dual-ported memory that does not take bandwidth from the graphics controller.

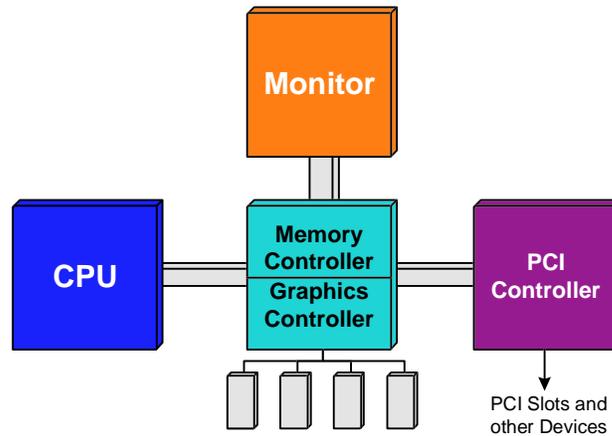


Figure 6. Unified Memory Architecture used in the O2 workstation from Silicon Graphics, Inc.

The new architecture being implemented by Compaq (Figure 7) allows the processor and graphics controller to access separate memory pools concurrently. Furthermore, the ELSA Gloria-L 3D graphics board and the Diamond Fire GL 4000 3D graphics board available with the new Compaq workstations have their own frame buffer and Z-buffer/texture memory. They do not rely on a common pool of memory for most data. Since data is accessed on two separate buses at the same time, the actual bandwidth is 1.07 GB/s for the memory bus plus the bandwidth of the graphics controller memory buses.

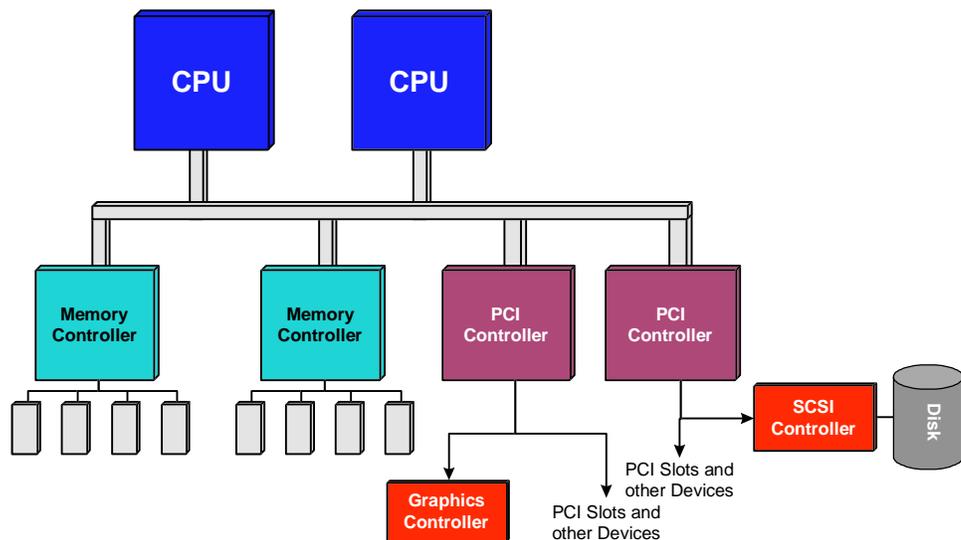


Figure 7. The new Highly Parallel System Architecture allows the processors and graphics controller to access two separate memory buses concurrently.

Crossbar Switch Architecture

A crossbar switch architecture provides multiple, independent paths to system memory. As Figure 8 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 Unified Port Architecture (UPA) and the Compaq TriFlex architectures.

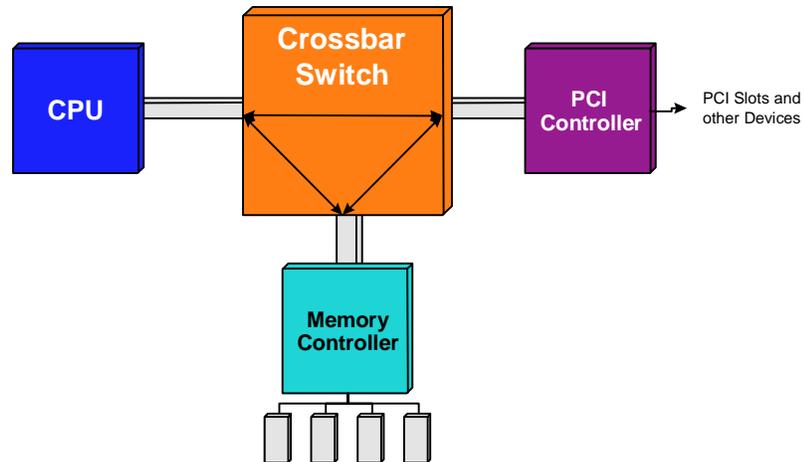


Figure 8. Crossbar switch architecture.

The Sun UPA provides a peak memory bandwidth of 1.2 GB/s in its full implementation. In other Sun implementations, however, memory bandwidth is half that or less. The actual memory throughput of a Sun system with UPA architecture will be limited by the same DRAM constraints identified earlier in the section “Dual Memory Buses.”

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, and two memory buses is not cost effective with today’s silicon technology.

Compaq chose not to use a crossbar switch because a better architectural solution was possible for the following reasons. First, the Pentium Pro processor bus is capable of running 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.

AGPset Architecture

The 440LX AGPset (LX chipset) from Intel is designed primarily for the commercial desktop and consumer desktop markets. However, some workstation vendors are deploying the LX chipset in machines targeted for workstation applications. Key features of the LX architecture include

- Single PCI bus
- Single memory bus
- 66-MHz SDRAM memory
- Advanced graphics port (AGP) bus for graphics cards
- Support for up to two Intel Pentium II processors

Figure 9 illustrates the LX chipset architecture. The microprocessor, memory, and AGP graphics buses operate with a peak bandwidth of 533 MB/s, while the peak bandwidth of the PCI bus is 133 MB/s. This contrasts with the 1.07-GB/s peak memory bandwidth and the 267-MB/s PCI bus bandwidth of the Highly Parallel System Architecture.

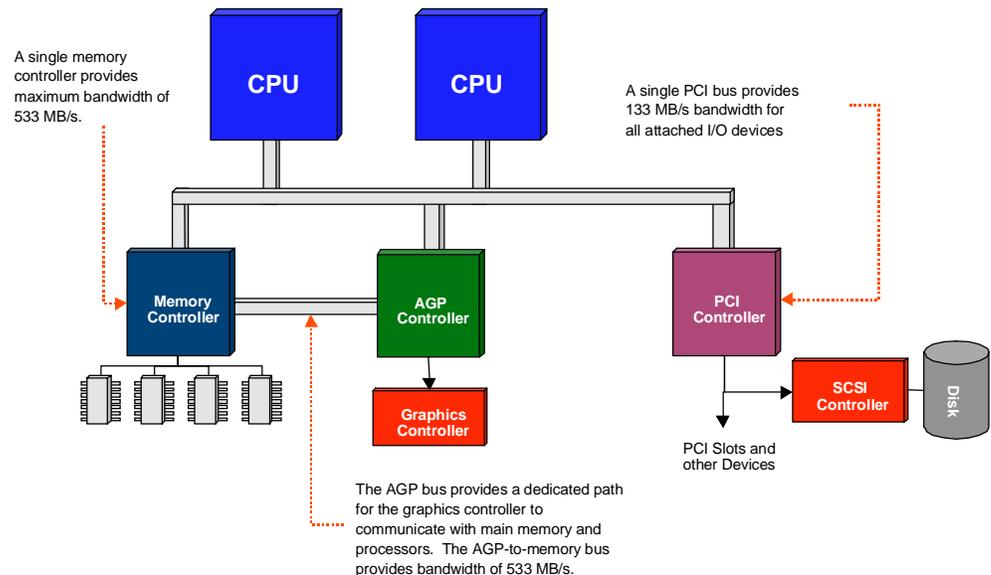


Figure 9. Block diagram of the LX architecture.

The Highly Parallel System Architecture supports industry-standard EDO memory arranged in 2:1 interleaved banks. The LX architecture uses 66-MHz SDRAM technology, which Compaq projects to be a very short-lived technology, based on imminent microprocessor and system architecture advancements. Specifically, Compaq expects that 100-MHz SDRAM will be available as soon as 100-MHz host-bus systems are available—sometime in the first half of 1998. Since dual-processor LX workstations are not expected to be generally available until sometime during the fourth quarter 1997, the 66-MHz SDRAM for workstations may have a useful life of only about four to six months. Large purchases of 66-MHz SDRAM could result in large obsolescence costs as early as next year, when users migrate to higher performance 100-MHz SDRAM technology. This can be a significant expense for customers, as most workstation applications require large memory configurations. The Highly Parallel System Architecture, on the other hand, uses proven EDO memory technology that is supported by other Compaq enterprise products. Customers who plan to migrate to 100-MHz SDRAM next year will be able to protect and fully amortize their investment in EDO memory by redeploying the EDO memory into other systems, such as ProLiant servers and other workstations.

The LX chipset supports an early version of AGP graphics. AGP provides a dedicated 66-MHz PCI bus connection between the graphics card, processor bus, and main memory. Full-AGP graphics systems are able to store texture and Z-buffer information into non-cacheable system memory. Additionally, future versions of AGP will support AGP bus master cycles that will allow the AGP card to transfer data on both clock-edges, providing a theoretical 533-MB/s graphics bandwidth. However, operating system support for AGP bus master cycles will not be available until mid-1998. AGP systems provide a 66-MHz PCI bus connection to the processor, while Highly Parallel System Architecture systems provide only a 33-MHz connection to the PCI buses. This increases the graphics bandwidth from 133 MB/s to 267 MB/s in current AGP systems, but this advantage can be deceiving. As Compaq testing shows, this increased bandwidth represents performance improvement of only about a one percent in most graphics-intensive applications. This very small increase occurs because graphics performance is limited mostly by CPU and graphics processor capability and only rarely by bus bandwidth. Compaq workstation engineering measurements show that the highest-end 3D-graphics cards typically require no more than 20 percent of a PCI bus, with peaks of up to 40 percent. Thus, PCI graphics cards still have headroom to double performance without saturating the PCI bus. The dual PCI buses in the Highly Parallel System Architecture in some cases provide more bandwidth than even a full AGP implementation. For example, in video editing, the dual PCI buses allow the application to manage three video streams, something that would not be feasible using only a single PCI bus implementation.

Compaq Professional Workstations using the Highly Parallel System Architecture support industry-standard PCI-based 2D- and 3D-graphics solutions. In conjunction with dual-peer PCI buses, the Highly Parallel System Architecture graphics solutions provide highly competitive performance for workstation applications.

For a more detailed comparison of the LX chipset and the Highly Parallel System Architecture, please refer to the technology brief *Highly Parallel System Architecture vs. the Intel 440LX AGPset in the Workstation Market*, document number ECG049/1097.

CONCLUSION

As business-critical applications become more demanding, the need for more bandwidth becomes increasingly important to customers. The new Highly Parallel System Architecture provides high performance and expansion capability plus unprecedented levels of bandwidth for X86 systems running Windows NT and very demanding applications. As the first computer company to implement this new architecture, Compaq once again demonstrates its technology leadership in providing innovative computing solutions to meet the needs of all its customers.