UPA — Sun's High Performance Graphics Connection

Technical White Paper

© 1999 Sun Microsystems, Inc. All rights reserved.

Printed in the United States of America.
901 San Antonio Road, Palo Alto, California 94303 U.S.A.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

TRADEMARKS
Sun, Sun Microsystems, the Sun logo, Ultra, Sun Elite3D, Sun Enterprise, Java 3D, PGX, Java, SBus, and VIS are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the United States and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

OpenGL is a registered trademark of Silicon Graphics, Inc.

THIS PUBLICATION IS PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS PUBLICATION COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THE PUBLICATION. SUN MICROSYSTEMS, INC. MAY MAKE IMPROVEMENTS AND / OR CHANGES IN THE PRODUCT(S) AND / OR THE PROGRAM(S) DESCRIBED IN THIS PUBLICATION AT ANY TIME.

Please
Recycle

Contents

UPA — Sun's High Performance Graphics Connection

Introduction

Designing high-performance systems is a complex balancing act. As faster components are implemented in designs, performance bottlenecks inevitably shift to other areas of the system. For example, a faster processor may cause a bottleneck in performance to shift from the CPU to the system bus if the bus can't feed data to the processor fast enough. A wider bus could cause the bottleneck to shift back again to the processor, or perhaps to some other part of the system. The challenge for system architects is to carefully balance system performance, and to keep costs within the product's targeted goals.

In 1994, engineers at Sun Microsystems faced these challenges in designing a new family of desktop workstations. These new systems — the Ultra™ family of workstations — were designed to use the new high-speed UltraSPARC™ family of processors. At that time, there was no standardized system or graphics bus technology with adequate bandwidth and scalability to meet Sun's demanding performance requirements. For this reason, Sun™ developed the Ultra Port Architecture (UPA) interconnect bus architecture, which was first introduced at the heart of the Ultra 1 workstation in 1995.

The UPA design has proven to be a superior interconnect architecture, even when compared to other bus technologies that have since emerged in the market. The purpose of this paper is to describe the UPA interconnect, and in particular, its functionality as a graphics bus. In addition, this paper will compare the UPA interconnect with other graphics bus technologies, and discuss how UPA helps to accelerate application performance.

Analyzing System Design and Application Performance

Design Trade-offs

In designing workstations for technical computing, system architects commonly make trade-offs between cost and performance (including the system's ability to handle tasks concurrently, such as computations and simultaneous graphics display). Trade-offs are made according to the platform's targeted markets and price points — for example, whether the system is aimed at low-end 2D graphics applications or high-end 3D geometry or texture mapping applications. As an example, a platform targeted at high-end performance graphics will typically perform geometry calculations on a graphics accelerator, freeing the processor for other computations. Other trade-offs include how much local graphics memory will reside on the graphics accelerator card, which may affect the frequency at which data is accessed from system memory. A high-bandwidth graphics bus can facilitate faster memory access, which can accelerate overall application performance.

Sun's approach is to engineer *balanced systems performance* into all of its system designs. Sun's engineers architect a careful balance between major system components — processor, memory, I/O, and graphics (Figure 1). A system interconnect like UPA helps to deliver this balance, especially since UPA scales according to the system clock speed. As processor speeds increase, the UPA interconnect keeps pace, delivering faster data throughput, enabling better performance of attached UPA graphics devices.

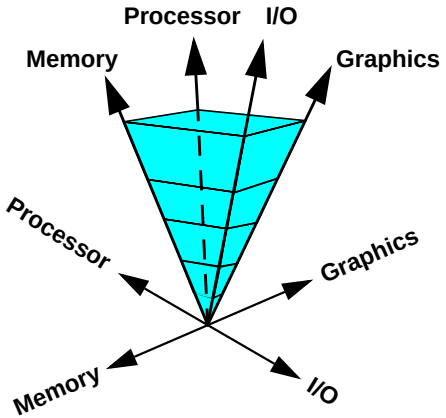

Figure 1 Sun architects systems for balanced performance

An Overview of the UPA Interconnect

As shown in Figure 2, the UPA acts as *a high-bandwidth interconnect between major system components* in Sun systems. Because of its high bandwidth, the UPA interconnect enables fast data transfers between the processor, system memory, and any UPA-attached graphics accelerators. The net effect is faster application performance, especially for demanding graphics applications in Sun's traditional markets: MCAD, MCAE, visualization, simulation, and DCC (Digital Content Creation).

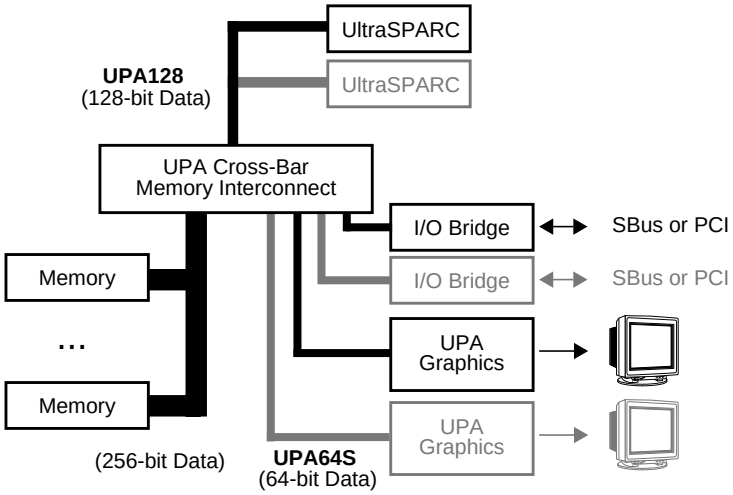

Figure 2 The crossbar design of the UPA interconnect allows simultaneous connections between major system components, unlike circuit-switched buses

The UPA architecture is a crossbar interconnect design, *allowing independent connections on the bus simultaneously* — its packet-switched, crossbar nature is similar to interconnect designs used in many high-end servers and supercomputers.

In looking at Figure 2, note that the UPA interconnect is much more than just a graphics bus. The UPA interconnect is responsible for delivering data to bandwidth-hungry processors, as well as to graphics devices. Because the crossbar design allows independent connections, data transfers can occur at the same time between memory and processors and between memory and I/O devices.

The UPA interconnect is also much more than a peripheral bus (such as PCI or SBus™). As Figure 2 shows, Sun systems typically include PCI or SBus as a peripheral bus, in addition to the UPA interconnect at the heart of the system architecture. (More detailed comparisons of other bus technologies appear later in this paper.)

As shown in Figure 2, different width paths are defined for different system components within the UPA backbone, enabling more balanced system performance. For example, the memory data path is 256 bits wide, while the data path to the processor(s) is 128 bits wide. The wide path to memory helps to deliver data quickly to the UltraSPARC processor(s), facilitating better performance as processor speeds increase.

Data transfers to graphics accelerators occur over the 64-bit wide UPA path known as UPA64S. This 64-bit graphics bus provides high bandwidth — enough to sustain the real-time display of high-resolution video, such as that required for streaming uncompressed HDTV video (1920 x 1080 at 24 frames per second). The UPA64S graphics bus is discussed in more detail later in this paper.

Scalability, High Bandwidth, and Efficiency

In originally designing the UPA interconnect, Sun's engineers anticipated the need for *a scalable design*, one that could be used in a number of systems over an extended period of time. As the clock rate in systems increases, the UPA interconnect scales accordingly, as shown in Figure 3.

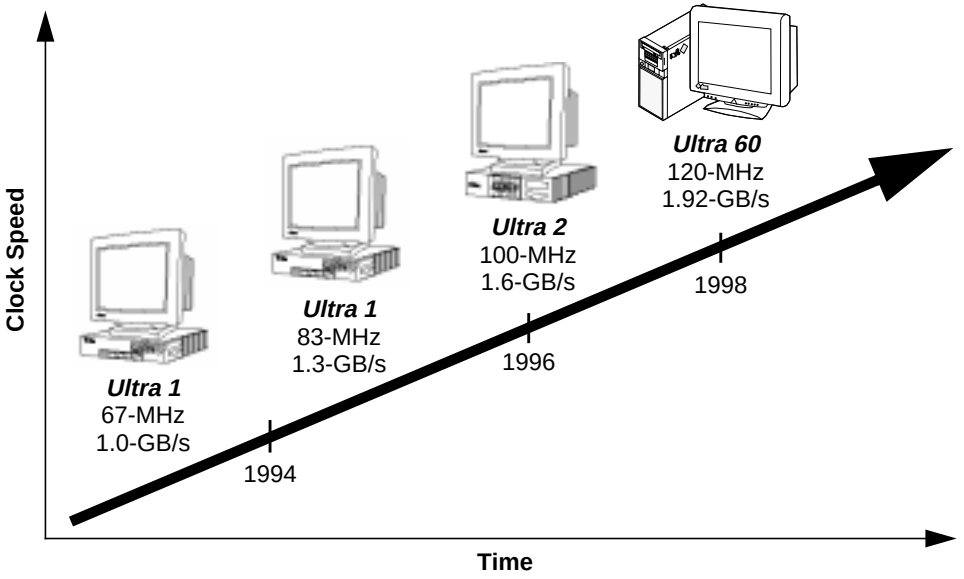

Figure 3 Bandwidth scales with the system clock

When the UPA interconnect premiered in Sun's Ultra 1 workstation, it was clocked at both 67-MHz and 83-MHz, which resulted in peak bandwidths of 1.07-GB/s and 1.32-GB/s. Today, at a clock speed of 120-MHz in Sun's Ultra 60 workstation, *the UPA interconnect can deliver up to 1.92-GB/s to the processor and 960-MB/s to a graphics accelerator.* (For example, to sustain the display of a 1-million pixel truecolor image at 60 frames per second, a sustained bandwidth on the graphics bus of at least 180-MB/s is required. With a peak bandwidth of 960-MB/s, the UPA interconnect provides more than enough bandwidth for today's — and even tomorrow's — UPA-attached graphics accelerators.)

Perhaps more important than raw bandwidth, the UPA interconnect is also extremely efficient, meaning that *its sustained performance is very close to its peak bandwidth capabilities.* This is because the UPA architecture features separate address, control, and data lines, which run at the full UPA speeds.

Other less-expensive bus designs may offer high peak bandwidths, but their sustained performance suffers because address and control lines are carried at slower speeds. Because of the UPA's efficiency, high bandwidth, and tremendous scalability, it offers distinct advantages in Sun systems for performance hungry applications.

The UPA64S Graphics Bus

In Sun systems, high-performance graphics accelerators (like Creator, Creator3D, and the Sun Elite3D™) attach directly to the 64-bit wide UPA graphics bus known as UPA64S. In many PC system designs, graphics cards are treated as peripherals, residing on mezzanine buses such as PCI. *Sun considers high-performance graphics an integral system component.* Graphics accelerators on the UPA interconnect can capitalize on UPA's high bandwidth and the fast data transfers it enables. In the Ultra 10, 60, and 450 architectures, for example, Sun includes 32-bit and 64-bit PCI buses for peripheral expansion capability. PCI buses in these systems allow optional cards to be easily added for extra network or storage connections.

Compared to network and disk devices, graphics devices can consume a large portion of available bus bandwidth. A network device can use roughly 3 to 4-MB/s of bandwidth on a peripheral bus such as PCI, while disk I/O can typically consume anywhere from 8-MB/s to 40-MB/s. As mentioned above, an imaging application that displays a 1-million pixel truecolor image at 60-Hz requires 180-MB/s of bandwidth — less than a fifth of the UPA64S's peak bandwidth. Furthermore, when graphics devices compete directly with network and disk devices for bandwidth, this competition can negatively impact certain applications (such as animation or streaming video). This can be a problem for shared bus designs like PCI.

The UPA64S graphics bus is slave-only, meaning that a graphics device cannot initiate data transfers. This is not a limitation in Sun systems because the graphics accelerator is not in competition with other devices (such as network and storage devices). This is in contrast to AGP-based architectures, where the lack of bus mastering can be a limitation for DMA transfers. UPA's slave-only architecture works well when bus traffic is PIO-intensive, which is the case for most technical computing applications, as discussed in the next section.

UPA64S also offers several electrical and mechanical advantages as a graphics bus. Sun designed the UPA64S bus to deliver sufficient power to drive sophisticated graphics chips, taking into account the robust connector and form factor requirements of state-of-the-art graphics accelerators like Sun Elite3D.

UPA graphics can also scale — from an Ultra 10 workstation with a single graphics display to an Ultra 60 with two graphics ports. (Note that the Sun™ Enterprise™ 6500 server can actually support eight graphics displays, although each operates on an individual UPA64S bus.) The AGP specification defines a path between only two devices — the graphics chip and the AGP chipset — allowing only a single graphics port. The crossbar nature of UPA (and the more than sufficient bandwidth of UPA64S), enables the support of multiple ports, allowing multiple graphics displays to be used in complex visualization and simulation environments.

Characterizing Application Performance and Bus Traffic

When analyzing application performance, graphics operations can be characterized into one of two categories:

- Those which produce short bursts on the graphics bus, such as immediate mode graphics. Immediate mode graphics result in Programmed I/O (PIO) transactions on the graphics bus.
- Those that produce long bursts, such as transfers of large display list primitives (such as triangle mesh strips) or texture maps. These are usually copied to the graphics accelerator via Direct Memory Access (DMA) operations.

When running technical computing applications (such as MCAD, MCAE, etc.), the majority of graphics operations are actually immediate mode graphics, in spite of an increasing trend towards display list APIs like Java 3D™ and OpenGL®. As a result, the traffic on the graphics bus today can be characterized as predominantly short bursts of PIO transactions.

Sun designed the UPA interconnect for high bandwidth, which optimizes PIO performance. With PIO operations, the processor “pushes” data to the graphics device, using the available bandwidth. DMA operations, on the other hand, require a certain amount of overhead — the processor typically places data in memory and then tells the graphics device where to find it.

The UPA interconnect also features separate address, control, and data lines for lower latency. Multiple PIO transactions complete quickly because one PIO does not need to wait for address and control lines before the next PIO is started. As a result, there is relatively little overhead for short PIO transactions. With UPA, longer DMA transfers are simply a block copy at the peak PIO rate. Compared to other bus technologies, UPA64S offers excellent sustained PIO and DMA transfer rates — 720 to 820-MB/s for PIO and a maximum of 820-MB/s for DMA transfers, compared to a peak bandwidth of 960-MB/s.

UPA64S and Other Graphics Bus Technologies

This section introduces PCI and AGP bus technologies, describes their PIO and DMA throughput capabilities, and shows how the throughput of UPA currently surpasses that of these other graphics bus designs.

Comparing PCI and UPA

In many system architectures, the PCI bus acts as a local system interconnect between peripheral components and the rest of the system, as shown in Figure 4. Until recently, most personal computers included PCI graphics cards, and data was transferred as necessary from main memory to the graphics card via the PCI bus.

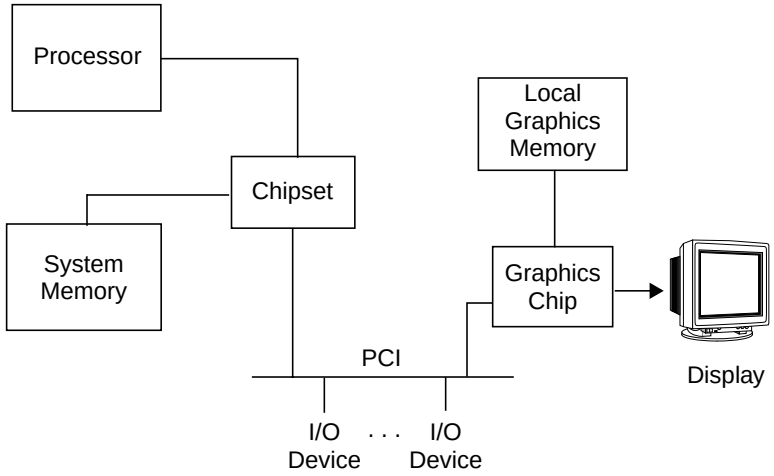

Figure 4 The PCI bus functions as a peripheral interconnect

In contrast, in Sun workstations and servers, the PCI bus is used primarily as a peripheral or mezzanine bus, just like the SBus in previous generations of Sun systems. Sun offers a variety of PCI cards for optional networking and storage options. For 8-bit entry-level graphics and windowing applications (where high bandwidth and performance are not required, and commodity components can be leveraged), Sun offers a PCI graphics card, called Sun's PGX™ card. For the most demanding graphics applications, graphics accelerators attach directly to the UPA interconnect in Sun systems.

The advantage of incorporating the PCI bus into system architectures is that it provides a standardized means of supporting peripheral devices. The drawback for PCI graphics cards is that the PCI bandwidth must be shared with other devices. This is one reason that Sun includes multiple PCI buses in its system designs.

Implementations of the PCI bus may operate at either 33-MHz (Rev. 2.0) or 66-MHz (Rev. 2.1), and support 32-bit or 64-bit data transfers. As shown in Figure 5, this results in peak bandwidths of 133-MB/s for PCI/33 (32-bit), 266-MB/s for PCI/66 (32-bit), and 533-MB/s for PCI/66 (64-bit). Although the PCI implementations provide higher bandwidths than SBus (used in earlier Sun workstations), UPA's bandwidth to graphics of 960 MB/s is significantly higher — 3 times higher than that of 32-bit PCI/66!

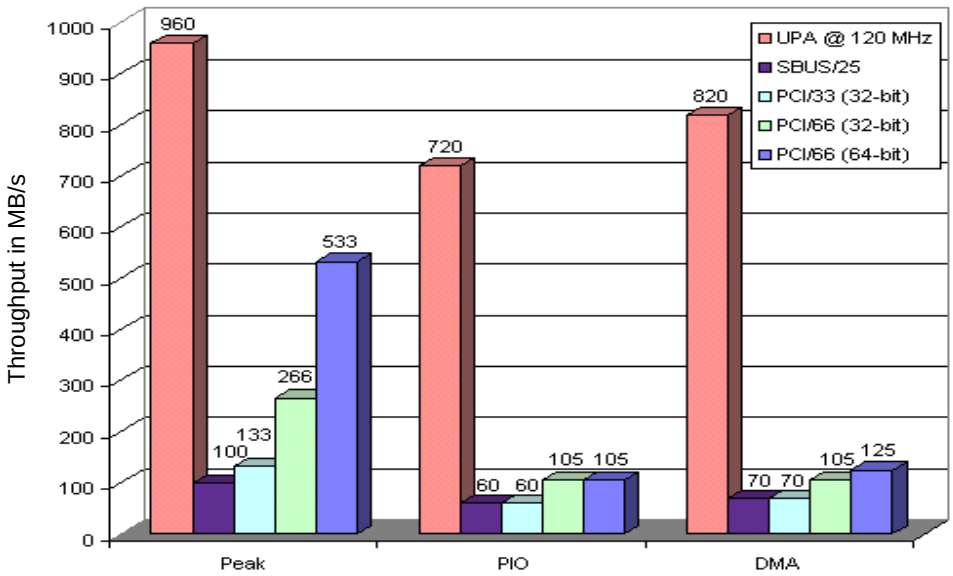

Figure 5 Comparison of UPA, SBus, and PCI bandwidths

Figure 5 shows the sustained throughput for PIO and DMA traffic for each of the bus technologies. For demanding graphics applications, the UPA64S bus provides significantly higher PIO and DMA throughput than PCI, even when compared to 64-bit PCI/66.

Comparing AGP and UPA

AGP (Accelerated Graphics Port) is a specification for a 32-bit graphics bus layered on top of PCI/66. Although AGP is a 32-bit specification, several newer graphics devices advertise architectures based on 64 bits or even 128 bits. In spite of these advanced architectures, the AGP bus still feeds these devices with a 32-bit pipe.

In 1996, Intel developed the AGP specification to mitigate certain PC architecture limitations, primarily for games developers. Many game applications require multiple texture maps (and sometimes large texture maps) to achieve realistic display effects. Because of the low bandwidth and shared nature of the PCI bus, some configurations with small amounts of graphics memory cannot move large texture maps from main memory to graphics memory quickly. In designing AGP, Intel defined a high-speed data channel between memory and graphics (Figure 6), similar to the high-bandwidth connection afforded by UPA64S. AGP, however, relies on PCI/66 for address and control signals, unlike UPA which carries these signals on separate lines.

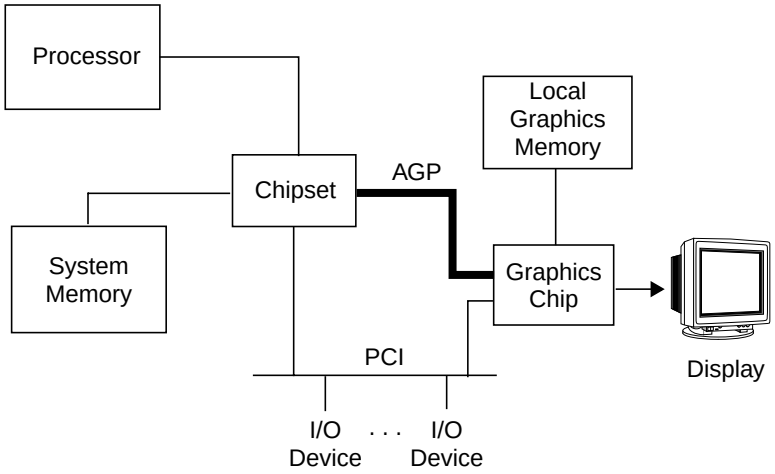

Figure 6 System architecture diagram showing the AGP bus

How Does AGP Accelerate Texture Mapping?

If an application requests a large texture that is not in graphics memory, AGP can facilitate access to the texture in one of two ways:

- The texture can be copied from main memory to local graphics memory over the AGP bus. The transfer can be faster because the data moves over the AGP bus rather than over the lower bandwidth PCI bus. However, true DMA copy operations over AGP won't be available until Windows NT provides AGP bus mastering, which won't occur until Windows NT 5.0 (Microsoft Windows 2000).
- The graphics chip can directly access the texture in main memory via the AGP bus. This method is referred to as DIME (Direct Memory Execute), and requires the operating system to map addresses in AGP memory to main memory. DIME support for remapping this addresses is incorporated in Microsoft Windows NT 5.0 and Microsoft Windows 98.

AGP is clearly optimized for large DMA transfers to facilitate faster texture memory access.

Comparing UPA, PCI, and AGP Throughput

Version 1.0 of the AGP specification defines two timing modes — AGP 1x and 2x — which provide peak bandwidths of 266-MB/s and 533-MB/s respectively (Figure 7). The 2x mode transfers data at both the rising and falling edges of the bus's 66-MHz clock, allowing AGP to achieve twice the throughput at the same clock rate. (Even though the AGP bus is actually clocked at 66-MHz, AGP 2x is sometimes referred to as "AGP at 133-MHz."

AGP 1.0 graphics devices and drivers rely heavily on the 32-bit PCI/66 foundation layer. (For this reason, 32-bit PCI/66 is shown for comparison in Figure 7.) All PIO operations, for example, cannot take advantage of the AGP extensions to PCI. Instead, PIO data transfers revert back to 32-bit PCI/66 transactions, restricting the PIO performance of most system configurations to 32-bit PCI/66 rates (indicated by the bar in Figure 7). Configurations that implement AGP 2.0's "fast PIO write" option (described next) can achieve moderately better PIO performance (indicated by the shaded area).

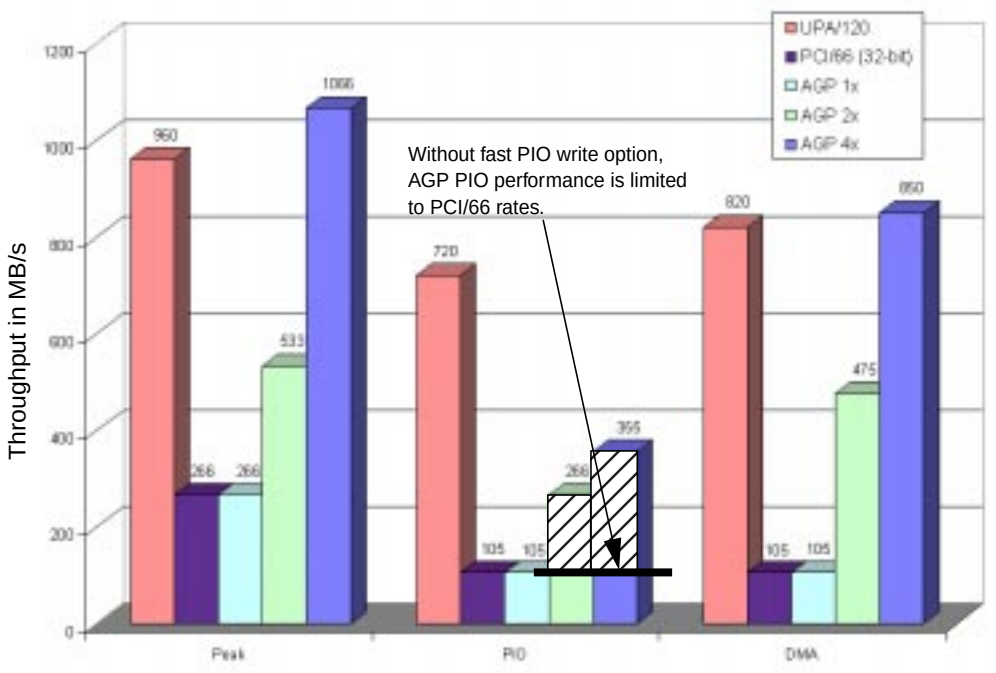

Figure 7 Comparison of UPA, PCI/66 (32-bit), AGP 1x, 2x, and 4x bandwidths

Version 2.0 of the AGP specification defines a “fast PIO write” option, which is not yet widely supported in PC motherboard chips. As a matter of fact, *fast PIO is defined in the AGP specification only as an option*. The fast PIO write option will move data at AGP rates for PIO, but will still run address and control transactions at 32-bit PCI/66 rates and protocols. When fast PIO is more widely implemented, the PIO performance of AGP 2x should improve somewhat (as indicated by the shaded area in Figure 7). However, even with fast PIO, AGP’s sustained PIO throughput will still fall below that of UPA.

Version 2.0 of the AGP specification also defines another timing mode — AGP 4x — which doubles the peak bandwidth of AGP 2x to 1066 MB/s. Although AGP 1x and 2x cards are available today, AGP 4x cards will not be available for quite a while. AGP 4x cards will be more difficult for vendors to implement because they must operate at lower voltages (1.7-volts instead of 3.3-volts). This voltage change will result in totally new designs that will be incompatible with today’s AGP 1x and 2x cards on an electrical level. For example, the graphics chips implemented in today’s 3.3-volt designs will not function in the

lower voltage AGP 4x cards. For this reason, the transition to AGP 4x will be slow. Only when AGP 4x cards are finally available will the peak bandwidth of AGP catch up with what is offered by UPA today. Sustained PIO performance (which is more important than peak bandwidth) will still be below that of UPA.

In the second version of the AGP specification, Intel also addressed many shortcomings of the original AGP specification. Some of these changes were packaged in an addendum to the Version 1.0 specification, called AGP Pro. AGP Pro addresses several electrical shortcomings in AGP 1.0, including the issue of the AGP connector providing more power, which had prevented some vendors from using more advanced graphics chips. Graphics cards that adhere to AGP Pro won't be available until much later in 1999.

Comparing Application Performance — AGP Shows Little Gain over PCI

Although AGP 2x provides about twice the bandwidth of PCI/66 and four times that of PCI/33, *most applications today see at most only a small increase in application performance using AGP* (versus PCI) graphics cards.[1]

There are several reasons why AGP may show little performance gain over some PCI cards:

- Many applications do not request textures larger than what fits in local graphics memory, which is when AGP shows a performance advantage over PCI. Indeed, benchmarks showing tremendous performance increases with AGP are typically fringe cases — for example, transferring a texture of over 6.1-MB, which would be unlikely to fit in even an 8-MB graphics card (note that some graphics memory is allocated for Z-buffering and the framebuffer itself, in addition to memory for textures). Today's technical computing applications typically reflect PIO performance, and AGP is optimized for DMA and not PIO.

1. This result has been shown in an independent benchmark. See the article on the web site *http://www.tomshardware.com/agp-perf.html*). The 3D WinBench 97 benchmark discusses the results of transferring an extremely large (6.1 MB) texture.

- Without bus mastering, AGP cannot do a DMA copy of textures from main memory to local graphics memory, so applications that request large textures can't really take full advantage of AGP's higher bandwidth. Bus mastering will be available with Microsoft Windows NT 5.0.
- Most professional applications (such as MCAD, MCAE, simulation, visualization, and Digital Content Creation) currently rely on Programmed I/O (PIO) performance. AGP is aimed at large data transfers such as those required to move large texture maps for games. Even when AGP's fast PIO option is more widely available in high-end PC workstations, it will result in only a modest improvement in sustained PIO throughput.

UPA64S Delivers Performance and Scalability Today

Today AGP 1x and 2x deliver only marginally better performance for applications than current PCI graphics cards. UPA, on the other hand, offers significantly higher overall bandwidth than both PCI and AGP, as well as better PIO and DMA throughput, as shown in Figure 7. Even AGP 4x technology with fast PIO, which is still quite some time away from shipping, provides a level of PIO throughput well below that of UPA at today's clock speeds — and, of course, Sun could possibly introduce new products (with faster clock speeds and higher UPA bandwidths) in the course of the next year.

Applications written with standard APIs, such as OpenGL and Java 3D, will scale transparently on Sun systems as clock speeds increase and faster processors are available. The demands of technical computing applications — intense computations, quick memory access, and fast graphics display — will continue to be met by UPA64S and the balanced design of Sun systems.

The Future

What bus technologies will Sun use in future products? For entry level graphics, Sun will continue to leverage volume components to minimize costs. This means that Sun may use PCI graphics chips in motherboard designs where appropriate, and continue to offer PCI/33 graphics cards (e.g., like Sun's PGX card) as entry-level designs warrant. Most Sun systems that support PCI also include at least one 64-bit PCI/66 bus, which would allow higher-bandwidth 64-bit PCI/66 devices to be supported.

For mid and high-performance systems, Sun sees no reason to pursue AGP technology, especially since UPA is a truly scalable design, capable of scaling well beyond today's performance levels. Sun will continue to enhance UPA's performance, maintaining Sun's status as a leading vendor of technical/graphical computing workstations. The UPA interconnect is capable of supporting a graphics accelerator that can generate up to 125 million triangles per second. Today, the Sun Elite3D accelerator can generate a peak of 5.9 million triangles per second (25-pixel, Gouraud-shaded), so there is plenty of UPA64S bandwidth available as Sun introduces faster graphics accelerators in the future.

As time goes on, Sun's future bus designs will continue to capitalize on the strengths of UPA:

- Wide data paths
- High bandwidth
- Low latency
- Flexible power and connector features
- Multiple display support
- Tremendous scalability

Finally, as Sun continues to innovate — as it has done with Java™, UltraSPARC, VIS™, and UPA technologies — it will outpace other vendors in providing balanced system design and outstanding application performance.

References

Sun Microsystems posts product information in the form of data sheets, specifications, and white papers on its Internet Web page at *http://www.sun.com*.

Look for abstracts on these and other Sun technology white papers:

- *Creator Graphics Technology,* Sun Microsystems, Inc., 1998.
- *Elite3D Graphics,* Sun Microsystems, Inc., 1998.
- *Ultra 1 Architecture,* Sun Microsystems, Inc., 1996.
- *Ultra 5 & Ultra 10 Architecture,* Sun Microsystems, Inc., 1998.
- *Ultra 60 Architecture,* Sun Microsystems, Inc., 1998.

For a detailed technical paper describing the UPA interconnect, please also see: *http://www.sun.com/microelectronics/whitepapers/wp95-023.html*

In addition, the following industry web sites contain more detailed information about AGP, and how AGP's performance compares to PCI graphics:

- *http://developer.intel.com/technology/agp/tutorial/index.htm*
- *http://www.tomshardware.com/agp.html*
- *http://www.tomshardware.com/agp-perf.html*
- *http://www.tomshardware.com/practicalagp.html*
- *http://newmedia.com/NewMedia/98/12/labreport/Big_League.html*
- *http://newmedia.com/NewMedia/98/12/labreport/agp_potential.html*

Sun Microsystems, Incorporated
901 San Antonio Road
Palo Alto, CA 94303 USA
650 960-1300
FAX 650 969-9131
http://www.sun.com

Sales Offices

Argentina: +54-1-317-5600
Australia: +61-2-9844-5000
Austria: +43-1-60563-0
Belgium: +32-2-716-7911
Brazil: +55-11-5181-8988
Canada: +905-477-6745
Chile: +56-2-638-6364
Colombia: +571-622-1717
Commonwealth of Independent States:
+7-502-935-8411
Czech/Slovak Republics:
+42-2-205-102-33
Denmark: +45-44-89-49-89
Estonia: +372-6-308-900
Finland: +358-9-525-561
France: +33-01-30-67-50-00
Germany: +49-89-46008-0
Greece: +30-1-680-6676
Hong Kong: +852-2802-4188
Hungary: +36-1-202-4415
Iceland: +354-563-3010
India: +91-80-559-9595
Ireland: +353-1-8055-666
Israel: +972-9-956-9250
Italy: +39-39-60551
Japan: +81-3-5717-5000
Korea: +822-3469-0114
Latin America/Caribbean:
+1-650-688-9464
Latvia: +371-755-11-33
Lithuania: +370-729-8468
Luxembourg: +352-491-1331
Malaysia: +603-264-9988
Mexico: +52-5-258-6100
Netherlands: +31-33-450-1234
New Zealand: +64-4-499-2344
Norway: +47-2218-5800
People's Republic of China:
Beijing: +86-10-6849-2828
Chengdu: +86-28-678-0121
Guangzhou: +86-20-8777-9913
Shanghai: +86-21-6247-4068
Poland: +48-22-658-4535
Portugal: +351-1-412-7710
Russia: +7-502-935-8411
Singapore: +65-438-1888
South Africa: +2711-805-4305
Spain: +34-1-596-9900
Sweden: +46-8-623-90-00
Switzerland: +41-1-825-7111
Taiwan: +886-2-514-0567
Thailand: +662-636-1555
Turkey: +90-212-236 3300
United Arab Emirates:
+971-4-366-333
United Kingdom: +44-1-276-20444
United States: +1-800-821-4643
Venezuela: +58-2-286-1044
Worldwide Headquarters:
+1-650-960-1300