AGP Speeds 3D Graphics

Accelerated Graphics Port Enables New Generation of 3D Accelerators

by Yong Yao

Continuing to set new standards in PC system design, Intel unveiled the details of its proposed accelerated graphics port (AGP) architecture at the recent AGP conference. The new interconnect addresses latency and bandwidth limitations that arise when PCI is used for 3D graphics. AGP is an extension of the basic PCI architecture, but it adds a demultiplexed address bus, pipelined transfers, and a 133-MHz effective transfer rate to improve performance.

This increased performance enables 3D texture data, traditionally kept in the frame buffer, to move to main memory. This change makes memory usage more efficient and frees frame-buffer bandwidth for tasks such as screen refreshing and Z-buffering. Besides helping 3D performance, AGP will also boost 2D performance due to the lower arbitration overhead and the higher data rate among the graphics controller, CPU, and main memory.

To achieve the high transfer rate, AGP is defined as a point-to-point connection, not a bus. While this design simplifies the electrical environment, it implies that the graphics subsystem be connected directly to the system-logic chip set; no other devices in the system can use AGP. A system can include an AGP-based graphics controller on the motherboard, or it can include a single AGP slot for an add-in graphics-accelerator card.

As it did with PCI and USB, Intel is aggressively pushing AGP into the PC market. The company has already signed Microsoft as well as a slew of graphics vendors to back the new standard, which Intel is openly licensing. The first AGP chip sets and graphics cards are due to appear in 1H97; by the end of next year, AGP should be in many high-end systems, becoming widely used within two to three years.

High-performance 3D graphics will be one of the key areas driving PC performance over the next several years. Intel realized that the basic PC architecture of today must be modified to deliver optimal 3D performance and thus has offered AGP as a solution. The company is also leaping into the 3D chip market; its first such product, due next year, will be an AGP-based 3D accelerator.

Resolution	Screen refresh 75 MHz	Color plane (R/W)	Z-buffer	Texture	Other	Total bandwidth required
640 × 480	50	100	100	100	20	370
800 × 600	100	150	150	150	30	580
1024 × 768	150	200	200	250	40	840

Table 1. 3D rendering bandwidth requirements in Mbytes/s. The numbers are based on 16-bit color, 30 frames per second, and average scene complexity (average amount of polygon overlap) equal to three.

Texture Maps Overwhelm PCI

Not long ago, PCI was created to resolve the data-bandwidth problem among peripherals, memory, and the host CPU. Ironically, the bandwidth of the PCI bus has now become the bottleneck in some advanced PC designs for 3D graphics. Rendering in 3D requires a large amount of bandwidth between the graphics accelerator and the frame buffer, as Table 1 shows. To meet these bandwidth requirements, many graphics subsystems use expensive high-performance memory chips, increasing the cost of the system.

One way to ease this problem is to move texture data from the frame buffer to main memory, as Figure 1 shows. Texture data is a good candidate to make this move for several reasons. First, it is generally read only; moving it to main memory does not cause data-consistency problems. Second, texture data must pass through main memory when it is loaded from mass storage; leaving it in main memory actually reduces overhead. Finally, texture size will grow as 3D applications become more complex and deliver better image quality, so leaving textures in the frame buffer will continue to increase the graphics subsystem cost over time.

Moving texture data to main memory reduces the size, and thus the cost, of the frame buffer, but it requires a corresponding amount of storage in main memory. From a system perspective, however, it is better to increase the size of main memory than of the frame buffer, because main memory uses less expensive DRAMs. In addition, storing texture data in main memory improves memory utilization, because application-dependent textures can be returned to the free memory heap when the application completes.

Unfortunately, this move shifts the bandwidth pressure from the frame buffer to the PCI bus. As Table 1 shows, texture data alone will exceed today's PCI bandwidth limit of

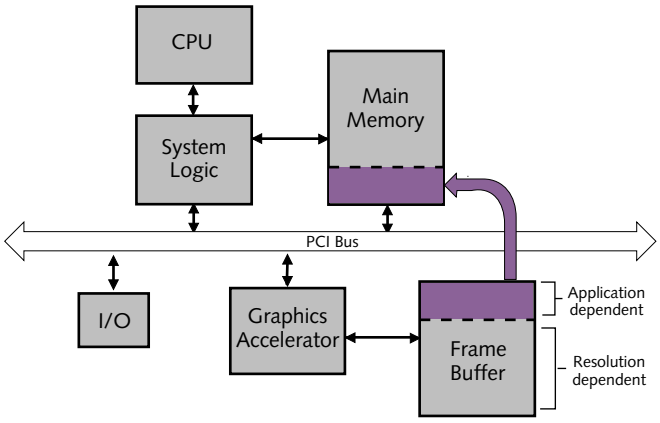

Figure 1. Moving texture data from the frame buffer to main memory can reduce the cost of the graphics subsystem, but the texture data requires more bandwidth than PCI can deliver.

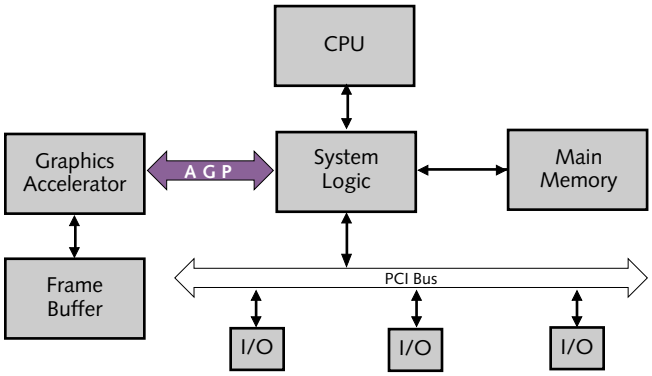

Figure 2. The proposed AGP system architecture adds a high-bandwidth path between the graphics subsystem and the system-logic chip set but has no impact on any other I/O subsystem.

100 Mbytes/s. Thus, a new method of connecting graphics to main memory is required.

AGP Extends PCI Standard

The primary goal of the AGP initiative is to contain the cost of implementing 3D in PCs while enabling performance improvements. Intel has made three major extensions to PCI for the AGP interface:

- Deeply pipelined memory read and write operations
- 133-MHz data-transfer rates
- Demultiplexing of address and data

These changes provide bandwidth and latency improvements well beyond that of the current 32-bit 33-MHz PCI bus and beyond even what the 66-MHz 64-bit PCI specification provides.

Signal name	Function description
IRDY#	Indicates the master is ready to provide all write data for the current transaction.
TRDY#	Indicates the target is ready to provide read data for a whole 32-byte block.
GNT#	Same meaning in AGP as in PCI, but additional information is associated with it.
C/BE[3:0]	Provides command information (different than PCI) when requests are being queued.
PIPE#	Used to queue pipelined request using the AD bus.
SBA[7:0]	Additional bus to queue pipelined requests.
SBE	Enables the SBA bus to be used.
ST[2:0]	Indicates how the AD bus will be used next.
AD_STB	Indicates the AD bus has valid data (2× mode).
SB_STB	Indicates the SBA bus has a valid request (2× mode).
DBF#	Indicates if the master can accept normal-priority read data.

Table 2. The first set of signals have a different meaning in AGP than in PCI; the second set are the 16 new AGP signals that are not part of the PCI specification. In the table, the master refers to an AGP-compliant graphics controller and the target refers to an AGP-compliant system-logic chip set.

Given these differences, Intel attempted to leverage as much of the PCI specification as possible. The three AGP performance enhancements are realized though the use of "sideband" signals, which are separate from the standard PCI signal set. AGP specifically avoids the use of any of the reserved fields, encodings, and pins in the PCI specification. AGP-defined protocols like pipelining are overlaid on PCI in such a manner that a standard PCI-compliant agent would view the bus as idle.

Despite these similarities, AGP and PCI devices cannot coexist. The key reason is that AGP is not a bus; it supports only a single device. Intel could have attempted to build AGP as a precise superset of PCI, putting AGP devices on that bus, but this decision would have either restricted AGP to 33 MHz or forced all PCI devices to support a 66-MHz clock, a needless cost burden for most PCI peripherals. In addition, AGP defines a different connector than PCI; the demultiplexed address signals would not fit in the current PCI connector. Thus, AGP devices need to be off of the PCI bus.

Leveraging the PCI definition allowed Intel to develop the AGP specification more quickly; the very concept of AGP did not exist until six months ago. The similarities to PCI will also ease the task of hardware designers. For the next few years, many graphics chips are likely to sport a dual AGP/PCI interface, allowing the system designer more flexibility. In fact, some the first "AGP" graphics chips may simply be slightly modified PCI chips that do not support many of the optional features of AGP.

AGP Boosts Memory Bandwidth

AGP is defined as a point-to-point connection based on a 3.3-V 66-MHz PCI bus. As Figure 2 shows, at one end of the connection is a device called the AGP-compliant target, which must always be the main-memory controller. The memory controller is part of the system logic in today's PCs.

Cycle name	CBE# status	Cycle description
Read	0000	Read n sequential Q-words, where n=length_field+1*.
Priority read	0001	Same as Read, but the request is queued in the high priority queue.
Write	0010	Write n sequential Q-words, where n=length_field+1*.
Priority write	0011	Same as Write, but indicates that the write data must be transferred from the master within the maximum latency window established for high-priority accesses.
Long read	1000	Same as Read except for access size, in this case n=4* (length_field+1), allowing up to 256-byte transfers.
Priority long read	1001	Same as Priority read except for access size, in this case n=4* (length_field+1), allowing up to 256-byte transfers.
Flush	1010	Similar to read. This command drives all accesses ahead of it to the point that all the results are fully visible to all other system agents, and then returns a single Qwork of random data as an indication of its completion.
Fence	1100	Creates a boundary in a single master's access stream, around which writes may not pass reads. This command does not occupy a slot in the AGP pipeline.

Table 3 The CBE lines provide bus commands for AGP pipeline operations. *The length_field is specified by the SBA bus.

At the other end is the AGP-compliant master, which is a graphics accelerator. The point-to-point connection not only makes the double-edged clocking possible, it also improves data integrity, simplifies the AGP protocols, and eliminates PCI bus arbitration overhead.

The AGP interface closely resembles that of the PCI bus. Most PCI signals are reused in AGP transactions, while others have slightly different semantics, as Table 2 shows. There are 16 new signals, also listed in Table 2, defined by the AGP specification. The PCI FRAME#, DEVSEL#, STOP#, and IDSEL signals are not used during AGP transactions. All AGP devices, however, are also required to respond to PCI transactions, so they must handle these signals as defined by the PCI specification.

Pipelining Improves Read Operations

Pipelining is the major protocol enhancement provided by AGP. Only memory read and write operations targeted to main memory are affected by pipelining. All other bus operations, including reads and writes to the graphics controller, are executed as standard PCI transactions. The request portion of an AGP transaction is signaled differently than a PCI address phase. The information is still transferred on the AD and C/BE# signals of the bus, as is the case with PCI, but is identified with the new signal PIPE# instead of FRAME#.

The system logic can access the graphics chip only with PCI transactions. Only the graphics chip can use pipelined transactions while accessing main memory. Table 3 lists the AGP pipeline operations. Note that, unlike PCI transactions, AGP pipeline transactions are not cache coherent and cannot be retried or terminated after they start.

The pipelined transactions defined by AGP run in a split-transaction fashion. The graphics chip initiates a pipelined transaction with an access request. The system logic responds to the request by initiating the corresponding data transfer at a later time. The graphics chip can then issue its next pipelined transaction while waiting for the previous data to return. This overlap results in several requests (reads or writes) active in the target's request queue at any point in time. The AGP specification does not impose a limit on how many outstanding transactions are allowed; this limit is determined by individual implementations.

As Figure 3 shows, the request queue is divided into high-priority and normal-priority subqueues, each of which deals with respective accesses according to its own rules. Data within a given queue is always returned in order, but data from different queues may be returned out of order. Figure 4 shows a transaction example.

Double Clocking Allows 133-MHz Transfer Rate

Another significant change between AGP and PCI is the double-clocking technique used to achieve a 133-MHz data transfer rate. The AGP interface operates with a 66-MHz clock, but data can be transferred on both clock edges. Most of the AGP electrical interface specifications are based on 66-

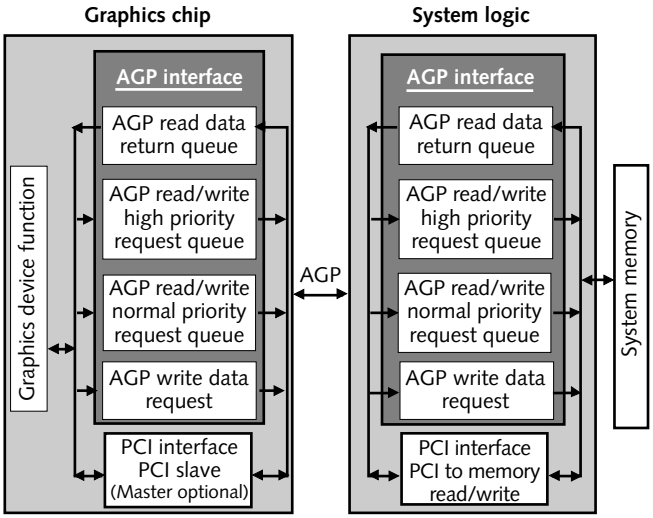

Figure 3. To improve response time, AGP's access queuing model divides transactions into high and normal priorities.

MHz PCI, but AGP has two transfer modes: 1× and 2×. Because it is a point-to-point connection using 3.3-V signaling, the AGP specification for the 1× transfer mode actually relaxes some of the design specifications relative to 66-MHz PCI. The 2× transfer mode, referred to as AGP-133, requires additional interface timing strobes and different signal timings from the 1× mode.

With a 133-MHz effective transfer rate, AGP reaches 533 Mbytes/s of peak bandwidth. In addition, its sustained bandwidth is significantly improved over PCI's by pipelining and demultiplexing.

Sideband Port Demultiplexes Address

To further increase the efficiency of the AGP interface, especially for random memory accesses, Intel added a sideband address (SBA) bus. This bus demultiplexes the address of an AGP transaction from the data. It is used exclusively to transmit AGP access requests. Therefore, it is always driven from

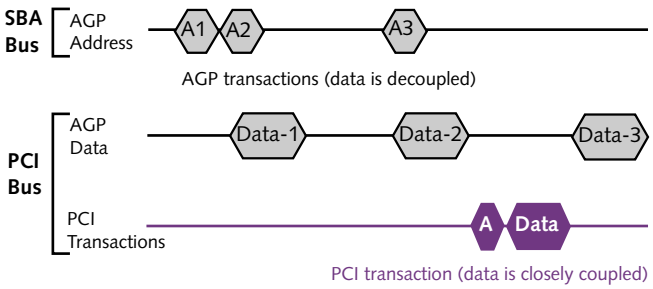

Figure 4. Sample AGP transactions show the PCI address phase and its corresponding data are always together. For a pipelined AGP transaction, the address phase and data phase are decoupled. In the figure, the PCI transaction is running out-of-order with the third AGP transaction. Note that pipelined transactions can also be performed without the optional SBA bus.

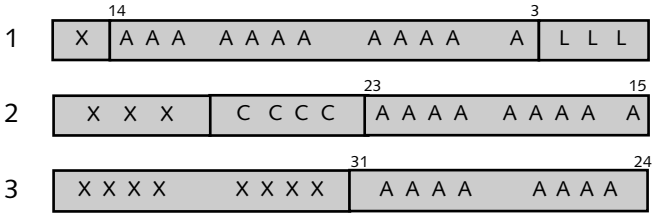

Figure 5. Address (A), command (C), and length (L) bits are transferred across the sideband address bus in three groups, as shown. Each group requires two 8-bit transfers.

the master to the target. Since use of the SBA bus indicates a pipelined transaction, there is no need for the PIPE# signal. Supporting SBA is optional for the master, but an AGP-compliant chip set is required to support this bus. If an AGP master uses the SBA bus, it will always issue demultiplexed transactions.

The sideband address bus (SBA[7:0]) is only 8 bits wide, reducing pin count. To use this bus, a complete AGP request is broken into three parts: low- order address bits and length, middle address bits and command, and high-order address bits, as Figure 5 shows. Each of these three parts takes two 8-bit transfers, or one cycle when double clocking. The latter two parts of the request need be transmitted only if they have changed since the previous request, which exploits potential locality among requests to minimize traffic over the SBA bus.

AGP Requires Address Remapping Table

Texture maps that are stored in main memory must be contiguous from the point of view of the application and the 3D controller. A typical 256 × 256 pixel 16-bit texture map is 128K in size, but Windows cannot allocate that much contiguous physical memory. Intel's solution is to have Windows allocate space for texture maps in the application's virtual address space, then lock the pages into physical memory. This method results in each texture map occupying many 4K pages distributed throughout physical memory.

Intel's AGP chip set contains a graphics address remapping table (GART) that mimics this virtual-to-physical address translation. Accesses from the graphics controller to a special address range are translated by the GART into physical addresses, making the textures appear to be contiguous. The GART is potentially quite large. A 4M window requires a 1024-entry GART, with each entry occupying 4 bytes.

The AGP specification does not describe how the GART is to be implemented by a chip set vendor. Instead, the vendor must provide a chip set driver that manages the GART according to a GART Services API that will be defined in the AGP specification. This definition should allow vendors other than Intel to design AGP chip sets that are compatible with Intel's at the operating-system level. The detailed definition of the GART and the API is not yet available, and this may delay third-party chip-set designs slightly.

Intel will be the first to market with an AGP chip set. This chip set, known as the 440LX, is designed for P6 (primarily Klamath) systems. Intel says it will not develop AGP chip sets for Pentium (including P55C) systems, although the P55C will still be the major volume player when AGP debuts; this decision is apparently a ploy to motivate home users to move to Klamath. We expect other vendors to add AGP to their Pentium and P6 chip sets, although they may lag Intel to market by three to six months.

A well-designed system-logic chip set is critical to bring the AGP architecture into full play. The chip set must bring the graphics accelerator logically closer to the host CPU and handles concurrent operations among the CPU, the graphics accelerator, main memory, and PCI masters. Even when multiple AGP chip sets are available, they may be significant performance differences among them.

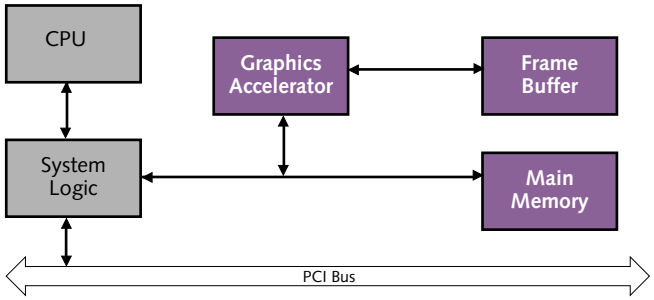

Figure 6. VESA has defined a version of its UMA design that places the graphics accelerator and its frame buffer on the main-memory bus.

UMA Offers Product Differentiation

The Video Electronics Standards Association (VESA) had also foreseen the increase in bandwidth needed for efficient 3D performance. As part of its UMA standard, VESA defined an optional method of connecting the graphics controller, frame buffer, main-memory controller, and main memory all across a high-bandwidth memory bus, as Figure 6 shows. This design allows the graphics controller to freely access either the frame buffer or main memory without requiring any intervention from the system logic; a simple handshaking protocol controls access to the memory bus.

The VESA proposal has some technical advantages:

- It adds only a couple of control lines to the chip set and the graphics accelerator, reducing the cost of manufacturing the chip set.
- It provides a 64-bit connection to memory, twice the width of AGP's. The interface speed will grow with memory technologies. For instance, today's fast SDRAM can operate at 100 MHz or higher.
- The graphics accelerator can directly access main memory without a chip set in between.

AGP, however, has its own technical advantages:

- The AGP interface is independent of the memory technology used for main memory.
- The chip set sees every access to main memory and can

For More Information

For more information on AGP, contact the AGP Implementors Forum at 503.264.9222; fax 503.264.5959 or access the Web at *www.teleport.com/~agfxport*.

allocate valuable memory bandwidth most efficiently.

• AGP does not impart an additional load on the system-memory interface.

• Unlike the VESA proposal, AGP can support both add-in graphics cards and motherboard implementations.

The VESA proposal illustrates a different technical approach to the bandwidth problem, but ultimately the important differences between the two proposals are the industry backers. The backing of Intel and Microsoft, along with a large number of graphics vendors, ensures AGP's success. Just as PCI outmuscled VESA's VL-bus, AGP will even more easily brush aside this alternative proposal.

AGP Headed for Rapid Acceptance

AGP is an open specification, in that it will be licensed to all interested parties under royalty-free reciprocal terms, similar to USB. Unlike USB, the definition of AGP is controlled by Intel alone. Because of this, the definition of AGP has been completed much more quickly than with USB or PCI. All major 3D graphics companies, such as 3Dfx, 3Dlabs, ATI, Cirrus Logic, Rendition, S3, and Trident, plan to support AGP. Intel expects to ship its first AGP system-logic chip set in 1H97. In the same time frame, some of the early adopters will launch their AGP-compliant graphics accelerators.

AGP is not an all-purpose solution. It does not, for example, solve the problem of making PCI faster; for most nongraphics devices, however, the current PCI definition is quite adequate. In fact, graphics was the key driver behind the 66-MHz 64-bit PCI standard; with AGP, it is now unlikely that this PCI extension will be implemented. Another issue is cost. The double-edge transfers require a more precise clock signal. Because of this, both the graphics chip and the system logic may need built-in PLL circuitry, which increases the manufacturing cost of the two devices.

Another issue is software support. To take advantage of an AGP (or UMA) design, either the device drivers or the operating system must support dynamic memory allocation. Windows 95 does not have this capability today, but Microsoft plans to fix this problem by the end of this year.

With its broad industry support, there is no doubt AGP will do well. Intel projects AGP penetration will reach about 90% by the year 2000, a ramp rate similar to that of PCI. With both Intel and Microsoft backing the idea, and a slew of graphics makers on board, AGP has enough momentum to be a sure winner. Graphics and system-logic vendors should plan accordingly. M

Intel to Enter 3D Chip Market

Intel has teamed with Lockheed-Martin to develop a 3D graphics accelerator based on AGP. Intel is not currently in the graphics chip market. The new device is based on Lockheed's 3D aviation simulations; both Intel and Lockheed will sell the new device.

Slated for introduction in the second half of 1997, the Intel 3D graphics accelerator will be an AGP-compliant master for Klamath systems. Theoretically, any AGP graphics accelerator is CPU independent, because it is decoupled by the AGP port, but of course it requires AGP-compatible system logic. Its performance is very chip-set dependent, while the chip set is CPU specific. The jointly developed accelerator will probably integrate 2D graphics, a RAMDAC, and video acceleration along with 3D functions.

We expect the device will rely on the host processor to do geometry calculations such as rotating, scaling, translation, lighting, clipping, and culling. The accelerator will handle setup and rendering. Additional features expected in the Intel graphics accelerator include:

- A single-chip solution based on the unified buffer architecture, with texture maps in main memory
- Support for Z-buffer, texture MIP mapping, perspective correction, Gouraud and Phong shading, anti-aliasing, alpha blending, stencil and scissor testing, and stipple masking
- Optimizations for the Klamath processor, Intel's Klamath (440LX) chip set, and Microsoft's Direct3D API
- A performance target of 30–50 million pixels per second and 500,000–1,000,000 50-pixel triangles per second, where pixels are textured and bilinear filtered

The graphics market is a natural extension of Intel's business. Intel is a giant semiconductor manufacturer with tremendous fab capacity for its advanced microprocessors. As process technology improves, however, CPU production moves to the latest fabs, freeing capacity in older processes. Using this capacity for chip sets and graphics controllers can maximize the return on its fab investment. In addition, high-performance system logic and graphics accelerators can spur demand for Intel's more profitable CPU products.

In the past few years, Intel has become the leading vendor of both system-logic chip sets and motherboards for PCs. It would not be surprising to see Intel become the top seller of graphics chips three years from now. Even if this change occurs, there will still be plenty of room for other vendors in the market, particularly if Intel chooses to focus on leading-edge 3D performance. But with too many vendors already in this market, some consolidation is inevitable.