Introduction to AGP-8X

Allan Chen
Jessie Johnson
Angel Suhrstedt

ADVANCED MICRO DEVICES, INC.
One AMD Place
Sunnyvale, CA 94088

Background

Tremendous efforts have been made to improve computing system performance through increases in CPU processing power; however, I/O bottlenecks throughout the computing platform can limit the system performance. To eliminate system bottlenecks, AMD has been working with industry leaders to implement innovative technologies including AGP, DDR SDRAM, USB2.0, and HyperTransport™ technology. This white paper will provide an overview of the latest graphics subsystem innovation: AGP-8X.

Introduction

The Accelerated Graphics Port (AGP) was developed as a high-performance graphics interface. It alleviates the PCI graphics bottleneck by providing high bandwidth throughput and direct access to system memory. The new AGP3.0 specification adds an 8X mode definition, which doubles the maximum data transfer rate from the previous high reached with AGP-4X by doubling the amount of data transferred per AGP bus cycle.

Figure 1 shows the graphic interface bandwidth performance evolution from PCI to AGP-8X. In this figure, AGP3.0 refers to the specific AGP interface specification. AGP-1X, AGP-2X, AGP-4X and AGP-8X represent the data transfer speed mode.

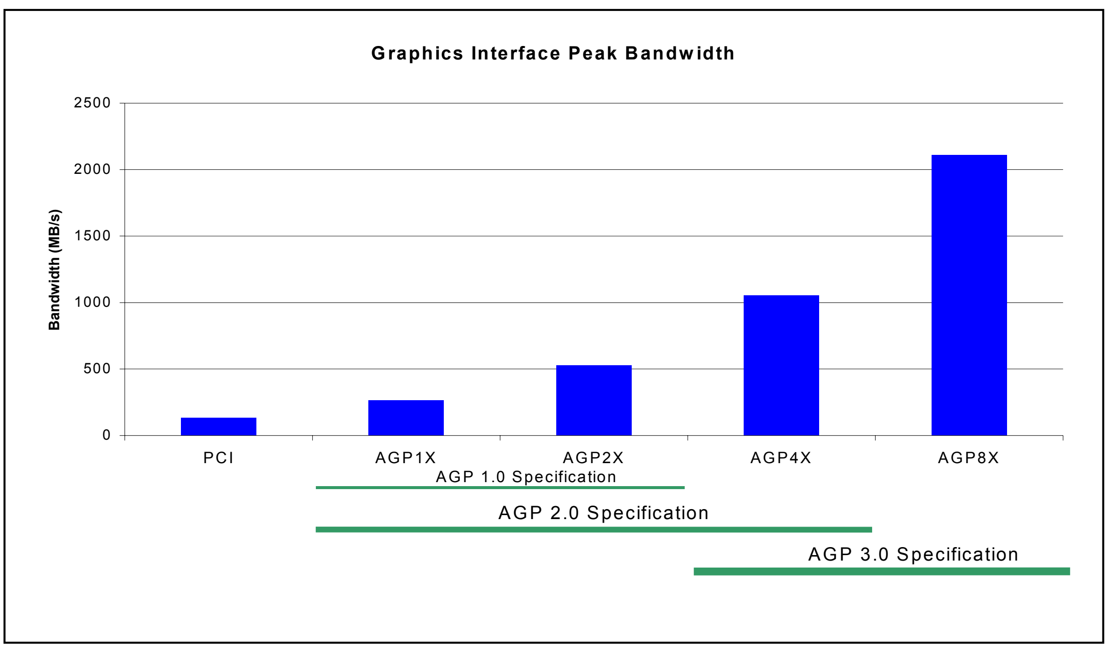

Figure 1: **Available bandwidth of different graphic interfaces.**

Benefits and Motivations

Laying the foundation at the platform level to harness the power of future graphics processors, AGP-8X delivers a peak bandwidth of 2.1GB/s to the graphics subsystem. This increase over AGP-4X enables graphics intensive applications to transfer more data to bandwidth-hungry graphics processors to enable the rendering of improved 3-dimensional images via enhanced vertex complexity and texturing, which in turn expands the end user's visual experience with more immersive, cinematic visual realities.

AGP3.0 Interface Specification Highlights

- New 8X data transfer mode, doubling the bandwidth to 2.1GB/s.
- New signaling scheme with few inverted signals and a lower voltage swing.
- Side band addressing to achieve higher data bus utilization.
- Power-on calibration scheme to optimize signal quality.
- Dynamic Bus Inversion to reduce the SSO noise.
- Isochronous mode operation to enable continuous data transfer for applications such as video streaming.

Compatibility to AGP-4X

- AGP-8X is compatible with the AGP-4X interface.
- Same AGP-4X signal pins with a few additional signals to support the AGP-8X interface.
- Utilizes the same connector as the AGP-4X interface.
- Compatible to AGP-4X and AGP Pro power delivery scheme.
- Motherboard can be designed to support both AGP-4X and AGP-8X.

For more details about AGP-8X features, please refer to the AGP3.0 specifications. For hardware board level implementation guidelines and recommendations, please reference the AGP Design Guide.

System Architecture with AGP-8X

Figure 2 is an example of balanced system architecture. An 8th Generation AMD Athlon™ processor links to the system through the AMD-8151™ HyperTransport AGP3.0 Graphics Tunnel. The 6.4GB/s total bandwidth from the CPU's HyperTransport interface enables AGP-8X and other system I/O peripherals to achieve optimal performance.

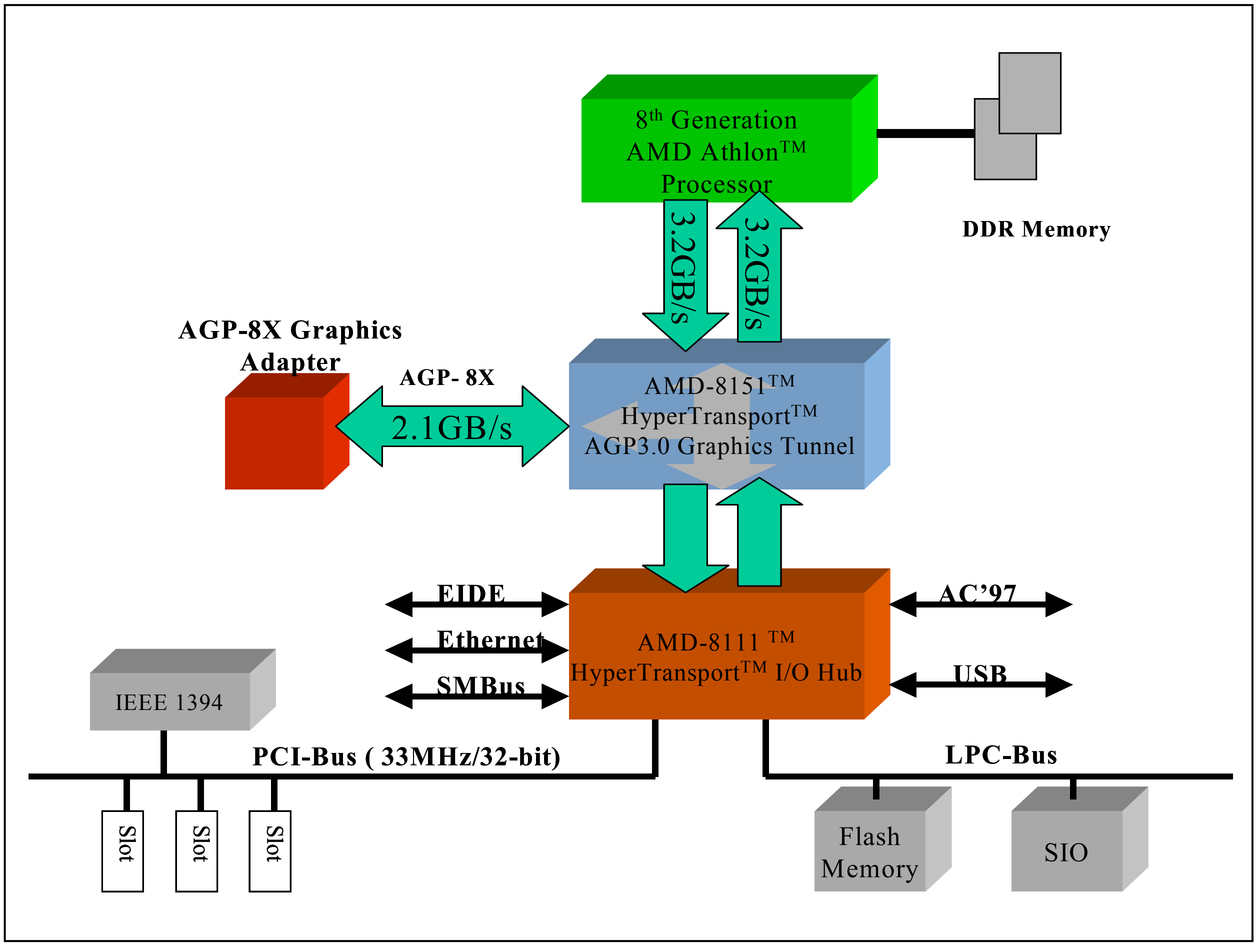

Figure 2: **Balanced system architecture featuring the AGP-8X graphics adapter.**

PC Graphic Subsystem Performance Evolution

To understand the advantages and benefits of AGP graphics, it is necessary to understand the issues that AGP technology has resolved. Figure 3 shows an architectural diagram of a generic PCI bus-based graphics subsystem. In this architecture, the graphics subsystem resides on the PCI bus. Note that the PCI bus graphics adapter embeds its own local memory on the adapter card. While this architecture performed well in the past, several issues developed that motivated the need for AGP:

1. Upgrading graphics memory is expensive, as memory modules must be added to the graphics card, or the graphics card must be replaced entirely
2. Since graphics data such as textures is stored in main memory, the PCI bus-based graphics card must access main memory via the PCI bus. These accesses can occur frequently, particularly if the graphics adapter has a small amount of local memory. Unfortunately, the graphics card must compete with other PCI bus peripherals for PCI bus bandwidth
3. If the graphics adapter must make frequent PCI bus accesses, other PCI bus peripherals may become starved for PCI bus bandwidth

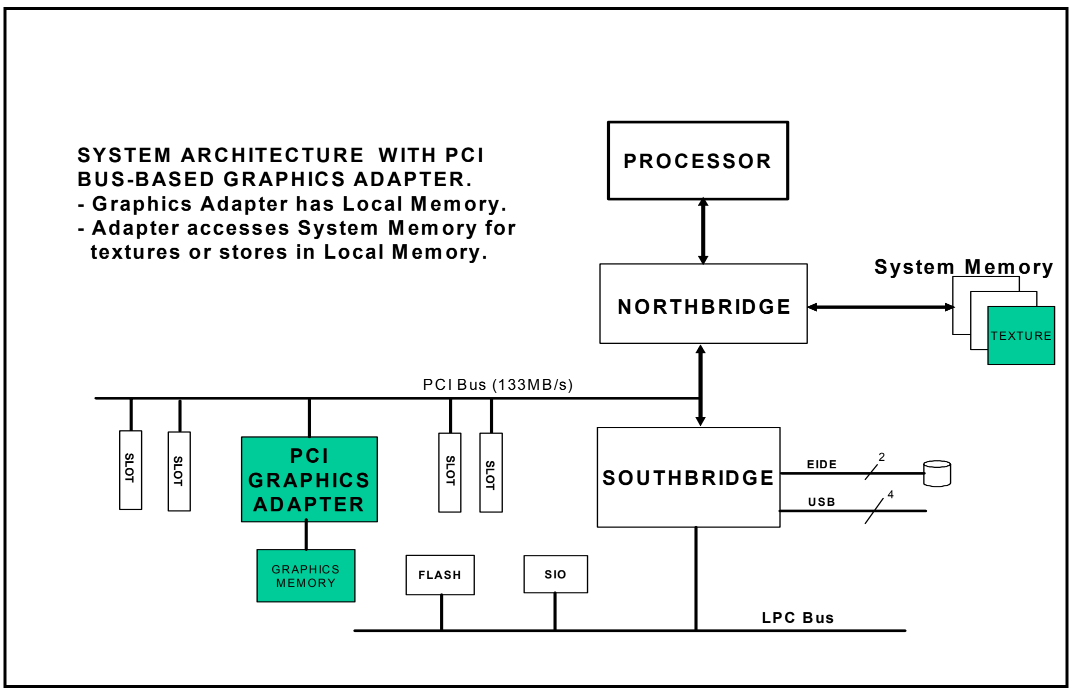

Figure 3: **Older system architecture showing PCI bus-based graphics.**

Figures 4 and 5 illustrate how AGP technology elegantly resolves the issues facing PCI bus graphics architecture. The system controller embeds the system's AGP graphics interface. The AGP interface utilizes the 66MHz PCI bus protocol in tandem with a side-band addressing (SBA) bus for concurrent posting of commands from the graphics card to the AGP logic embedded in the Northbridge. The Northbridge embeds read/write and command queues (buffers) to allow full-speed data and command transport between the AGP device and the system controller, and concurrent full-speed data transport between the system controller and the DDR memory subsystem.

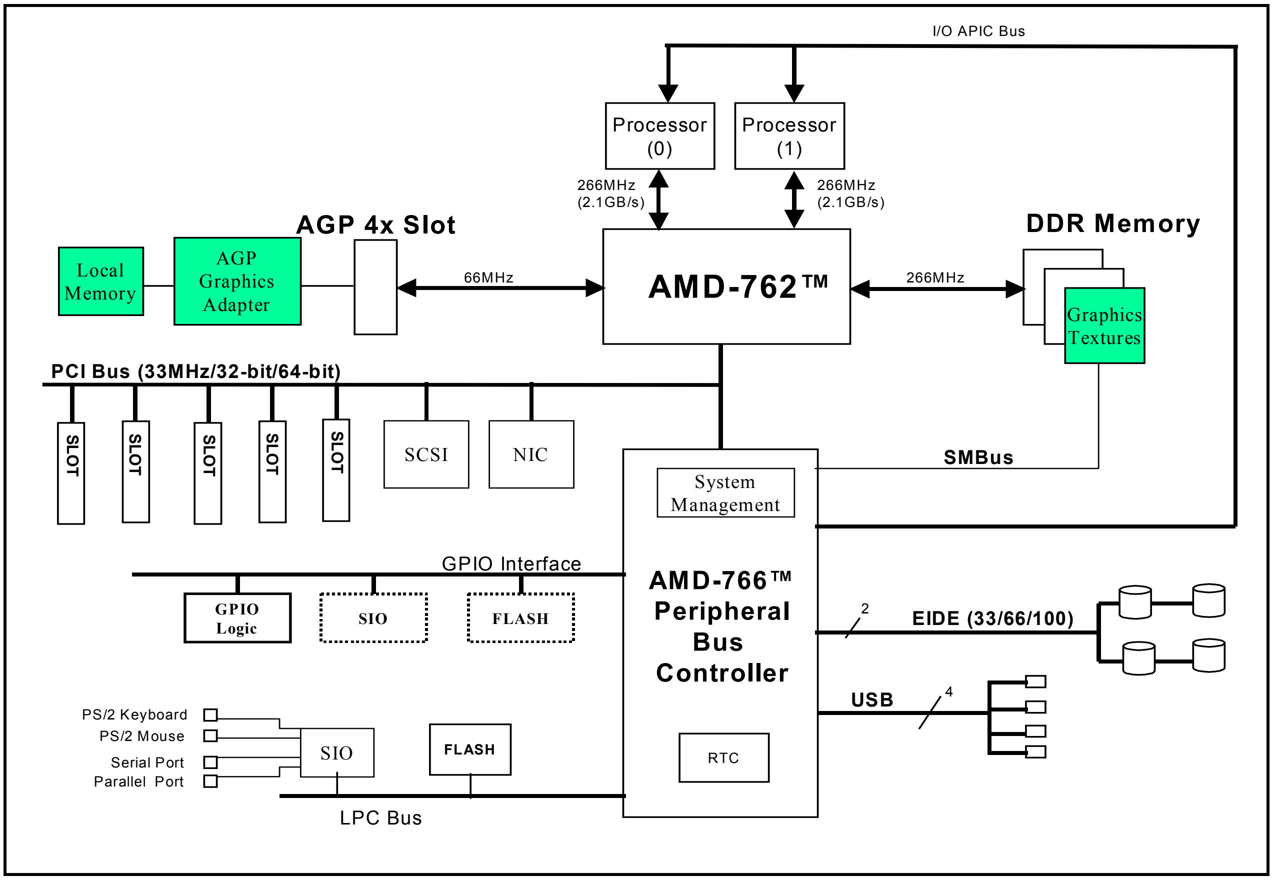

Figure 4: **AMD-762™ system controller AGP graphics architecture showing use of system memory for graphics operations.**

The architectures shown in Figure 4 and Figure 5 produce the following benefits:

- The native architecture of the AGP graphics subsystem offers significant raw performance improvement over PCI bus-based graphics subsystems
- The AGP architecture allows the AGP graphics subsystem to view and use main memory just like its own local memory—meaning the AGP graphics card shares system memory. The AGP graphics card cannot distinguish between system memory and local memory, as it all appears as local memory. To the end-user, graphics performance can be enhanced by increasing system memory (inexpensive), rather than by adding expensive graphics memory.
- The graphics subsystem no longer has to compete for PCI bus bandwidth to access data in system memory. This benefit allows the graphics subsystem to run at full speed with minimal interruption from other components in the system. It also increases system concurrency—meaning that the processor, AGP graphics subsystem, and PCI bus device can run independently and concurrently, thus increasing system performance.
- PCI bus devices no longer have to compete with the graphics adapter for PCI bus bandwidth. PCI bus availability has been increased with the removal of the graphics subsystem from the PCI bus.

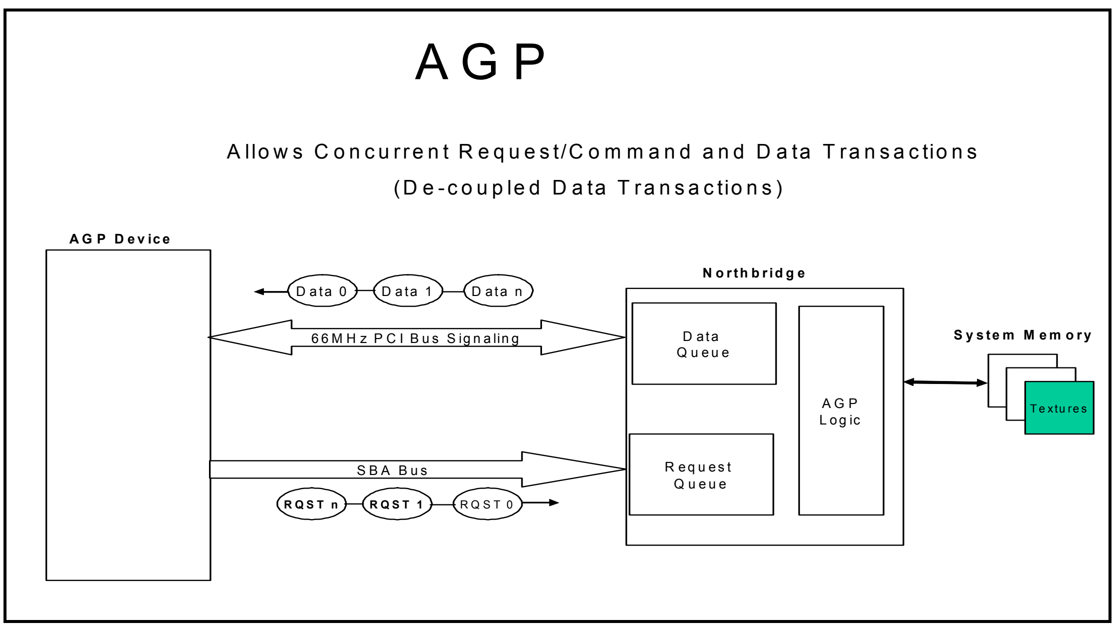

Figure 5: **High-level AGP interface diagram showing bus architecture and embedded Northbridge components**

Over time, the AGP graphics subsystem has scaled to increasing levels of performance. As shown in Table 1, there are several modes (data transfer rates) that have evolved over time. Analogous to gears on a sports car transmission, first gear is the original AGP-1X mode, offering a data transfer rate of up to 264MB/sec. Second gear is AGP-2X mode, doubling the data transfer rate to up to 528MB/sec. The third gear is AGP-4X, offering a transfer rate of up to 1GB/sec. Finally, the fourth gear, AGP-8X offers the highest performance data transfer rate of up to 2.1GB/sec. (The notation 2X, 4X and 8X are relative the original AGP-1X mode).

Table 1: **AGP Modes and Corresponding Peak Bandwidths.**

AGP Graphics Mode	Peak Bandwidth (data transfer rate)
AGP-1X	Up to 264MB/sec.
AGP-2X	Up to 528MB/sec.
AGP-4X	Up to 1GB/sec.
AGP-8X	Up to 2.1GB/Sec.

Summary

The Accelerated Graphics Port (AGP) AGP-8X mode is the next step forward in the evolution of high-performance graphics interface technology. By virtue of its 100% increase in peak bandwidth over AGP-4X offered to the graphics subsystem, AGP-8X provides the graphics I/O performance headroom necessary to meet the requirements of future graphics processors in order to enable further increases in delivered system performance to the end user.

Information and Links

- AGP Specifications/Design Guide: www.agpforum.org
- AGP System Architecture by Dave Dztko, MindShare Inc: www.mindshare.com
- HyperTransport™ Technology Consortium: www.hypertransport.org
- AMD Core Logic Product Briefs, Data Sheet, Design Guide: www.amd.com

AMD Overview

AMD is a global supplier of integrated circuits for the personal and networked computer and communications markets with manufacturing facilities in the United States, Europe, and Asia. AMD produces microprocessors, flash memory devices, and support circuitry for communications and networking applications. Founded in 1969 and based in Sunnyvale, California, AMD had revenues of $3.9 billion in 2001. (NYSE: AMD).

Cautionary Statement

This White Paper includes forward-looking statements that are made pursuant to the safe harbor provisions of the Private Securities Litigation Reform Act of 1995. Forward-looking statements are generally preceded by words such as “expects,” “plans,” “believes,” “anticipates,” or “intends.” Investors are cautioned that all forward-looking statements in this white paper involve risks and uncertainties that could cause actual results to differ from current expectations. Forward-looking statements in this white paper about the AMD Opteron™ processor and AMD Athlon™ processor involve the risk that AMD will not be able to produce the processor in the volume required by customers on a timely basis; that AMD may not be successful in developing an infrastructure to support the processor; that third parties may not provide infrastructure solutions to support the processor; that the processor will not achieve customer and market acceptance; and that software applications will not be optimized for use with the processor. We urge investors to review in detail the risk and uncertainties in the company’s Securities and Exchange Commission filings, including the most recently filed Form-10K.