

The AMD-760™ MPX Platform for the AMD Athlon™ MP Processor

MultiProcessing with eXtended Performance

Jessie J. Johnson

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

PID# 25787A

Table of Contents

| | |
|--|----|
| Introduction..... | 3 |
| AMD-760™ MPX Platform Overview..... | 3 |
| Supported Processors | 4 |
| Platform Architecture..... | 5 |
| Smart MP Technology | 8 |
| Dual AMD Athlon™ System Bus | 10 |
| Innovative Bus-Snooping Capability | 17 |
| Optimized MOESI Cache-Coherency Protocol | 19 |
| DDR Memory Subsystem | 21 |
| AGP-4X Graphics Subsystem..... | 22 |
| ATA-100 Storage Subsystem | 23 |
| Primary (66MHz) and Secondary (33MHz) PCI Bus Interfaces | 24 |
| AC-97 Audio Interface | 26 |
| Appendix..... | 28 |
| Glossary of Terms..... | 28 |
| What is Bus Snooping?..... | 28 |
| AGP Graphics Background..... | 29 |
| AMD Overview | 33 |

Introduction

High-performance...scalability...manageability...all are key terms to describe a class of power platforms intended for mission-critical applications required of servers and workstations. The AMD-760™ MPX platform is a high-performance, two-way multiprocessing (MP) system solution designed for server- and workstation-class applications. Building upon its predecessor (the AMD-760 MP chipset), the AMD-760 MPX platform offers **M**ulti**P**rocessor **e**Xtended (MPX) performance beyond the original performance of the AMD-760 MP platform. Consisting of the AMD Athlon™ MP processor and the AMD-760 MPX chipset, this solution offers scalable processing capability, high-bandwidth memory and I/O performance, and sophisticated system management capability to support a wide range of computing infrastructures. This white paper describes the AMD-760 MPX platform, architecture, and underlying technologies.

AMD-760™ MPX Platform Overview

The AMD-760 MPX platform addresses the server and workstation sectors by offering the following high-level features:

- Uniprocessor and two-way symmetric multiprocessing capability
- Dual point-to-point 266MHz AMD Athlon system buses designed to support up to 2.1GB/s transfer rate per system bus
- 266MHz DDR (Double Data Rate) memory interface supporting up to 4GB of memory space using registered PC2100 DIMMs. ECC (Error Correcting Code) memory is also supported
- AGP-4X graphics interface, backwards-compatible with AGP-1X and 2X modes
- Primary 66MHz/64-bit/32-bit PCI 2.2-compliant PCI bus interface
- Secondary 33MHz/32-bit PCI 2.2-compliant PCI bus interface
- AC-97 audio interface

- EIDE storage controller subsystem supporting ATA-33/66/100MB/s data rates
- System management functions

As shown in Figure 1, these features are packaged in a two-chip core logic solution consisting of the AMD-762™ system controller (Northbridge) and the AMD-768™ peripheral bus controller (Southbridge). When implemented with AMD Athlon MP processor technology, these elements combine to deliver outstanding performance to server- and workstation-class systems.

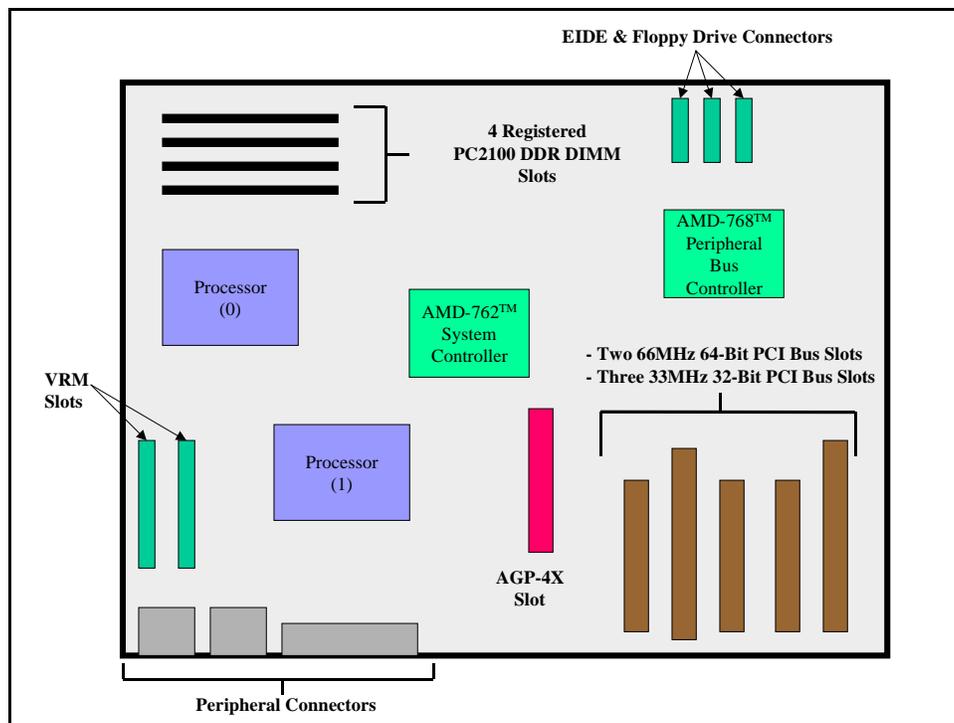


Figure 1: AMD-760™ MPX Motherboard Platform

Supported Processors

AMD manufactures a wide range of microprocessors to support a multitude of market segments and applications. The AMD-760 MPX platform supports AMD processors designed specifically for multiprocessing. The AMD processors shown in Table 1 are supported by the AMD-760 MPX Platform.

Table 1: AMD Processor Front-Side Bus Speed and Corresponding Memory Support

| Processor Model | Processor Clock Speed | Supported Front-Side Bus speeds | Supported DDR SDRAM Memory |
|-----------------|-----------------------|---------------------------------|----------------------------|
| AMD Athlon™ MP | All speed grades | 266MHz | PC2100 |

Two-way multiprocessing is supported only in the configurations shown in Table 2.

Table 2: Supported Two-Way Multiprocessor Configurations

| Configurations | Processor Slot (0) | | | Processor Slot (1) | | |
|------------------------------------|--------------------|-----------|--|--------------------|-----------|--|
| | Model | FSB Speed | Processor Clock Speed | Model | FSB Speed | Processor Clock Speed |
| AMD Athlon™ Configuration 1 | AMD Athlon MP | 266MHz | Same frequency as Processor (1) ¹ | AMD Athlon MP | 266MHz | Same frequency as Processor (0) ¹ |

The multiprocessor configurations allowed are those in which both processor sockets are populated with processors of the same model, front-side bus speed, and processor clock speeds. An AMD Athlon MP processor cannot coexist in the same system with any other processor model other than another AMD Athlon MP nor can different speed grades of the AMD Athlon MP processor coexist in the same system.

Platform Architecture

Servers and workstations leverage similar underlying technologies as fundamental building blocks. Both classes of systems require high-performance computing capability, high-capacity memory elements, and robust I/O (Input/Output) subsystems. These platforms, however, begin to diverge in the areas of scalability, graphics, and I/O control. Workstations, for example, require high-end graphics, mid-level storage capabilities, and limited I/O expansion. On the other hand, servers do not require high-end graphics, but instead require “scalable” high-performance computing capabilities, massive high-performance storage engines, and an I/O subsystem that is robust enough to concurrently support a multitude of high-speed networking and storage subsystems.

¹ All AMD Athlon MP processor speed grades are supported, however, Processor (0) and Processor (1) must operate at the same processor clock frequency. The AMD-760 MPX platform does not support configurations in which Processor (0) and Processor (1) are operating at different processor clock frequencies.

Early in the design process, AMD engineers faced the extreme challenge of developing an architecture that was powerful, robust, and cost-effective, yet flexible enough to support the diverging demands of both server- and workstation-class systems. Their answer: the AMD-760 MPX platform.

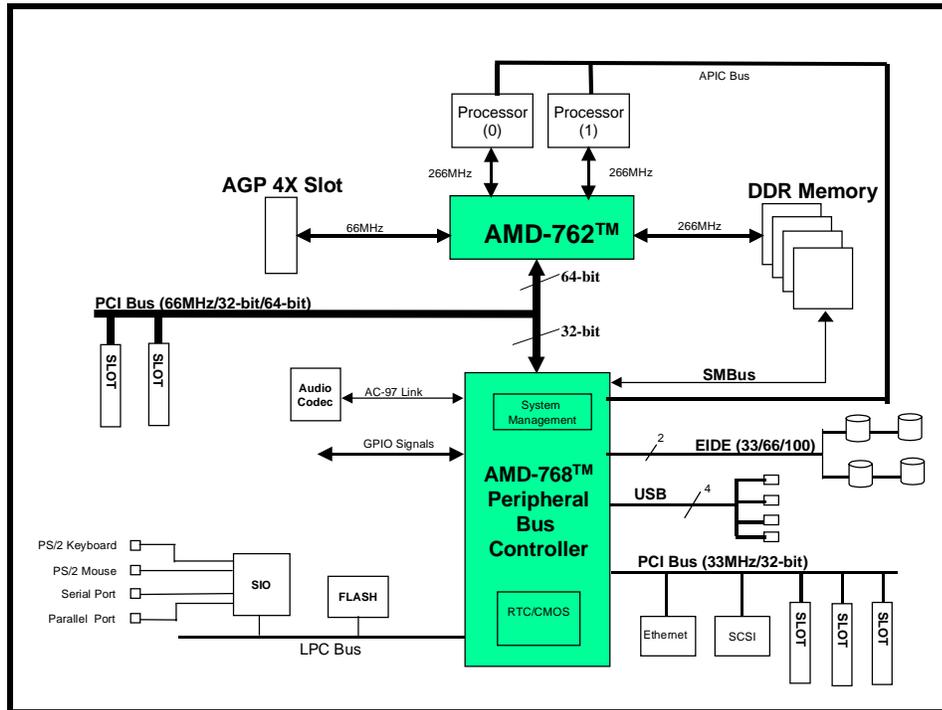


Figure 2: AMD-760™ MP Platform High-Level Architecture

As shown in Figure 2, the AMD-760 MPX platform employs PGA-socketed (Socket A) AMD processor-based technology as the high-performance system engine(s) of choice. The AMD-760 MPX design supports a single processor for workstation-class systems or low-end servers, leaving a second available processor socket for future expansion. Customers will enjoy the security of knowing their original investment is protected with a system that can be scaled as their computing demands grow. Customers can upwardly scale the processing capability of their system in one of several ways:

1. In a uniprocessor configuration, the system can be upgraded by replacing the original single processor with a faster or improved AMD Athlon MP processor.
2. In a multiprocessor configuration, the system can be upgraded by adding a second AMD processor (of the same model and clock speed as the first

processor), thus enabling symmetric multiprocessing and significantly boosting system performance. (Refer to Table 2 for supported multiprocessor configurations.)

3. Replacing the original single processor with a faster AMD processor, and adding a second processor of equivalent model and speed. (Refer to Table 2 for supported multiprocessor configurations.)

In high-end workstation and mid-range server applications, where maximum performance is required from the onset, systems can be shipped with two processors, immediately enabling the benefits of symmetric multiprocessing.

The system logic is partitioned as follows:

The AMD-762 System Controller houses the high-speed elements critical to system performance. Contained in the AMD-762 controller are the following subsystems:

- Dual point-to-point 266MHz AMD Athlon system bus interfaces, supporting up to two processors
- 266MHz DDR memory interface, supporting up to 4GB of memory space using PC2100 registered DDR memory DIMMs. ECC memory is also supported
- AGP-4X graphics interface
- 66MHz/64-bit/32-bit PCI bus interface
- 949-pin plastic CCGA (Ceramic Column Grid Array) package
- 2.5V core

For detailed descriptions, refer to the **AMD-762 System Controller Data Sheet**.

The **AMD-768 Peripheral Bus Controller** compliments the AMD-762 system controller by offering a robust I/O subsystem along with sophisticated system management capability. Embedded are the following features:

- Primary 66MHz/64-bit/32-bit PCI bus interface
- Secondary 33MHz/32-bit PCI bus interface
- AC-97 audio interface
- ATA-33/66/100 EIDE interface
- Four-port OHCI USB host controller
- LPC (Low Pin Count) bus
- System management interface
- SMBus interface
- Real time clock (RTC)
- 256 bytes CMOS memory
- I/O APIC interrupt controller
- Random number generator
- 32-GPIO pins
- 492-pin Ball Grid Array (BGA) package
- 3.3V core and output drivers; 5V tolerant input buffers

For detailed descriptions, refer to the **AMD-768 Peripheral Bus Controller Datasheet**.

Smart MP Technology

GHz, MIPs, MOPs, and GB/s are measures of raw computing power. However, for revolutionary breakthroughs in system performance, raw computing power is not enough – the system has to be “smart” as well. AMD introduces Smart MP technology – a new multiprocessing architecture enabling faster performance beyond traditional multiprocessor system architectures.

Smart MP technology consists of the following architectural features:

- Dual point-to-point high-speed system buses
- Innovative bus-snooping capability
- Optimized MOESI cache-coherency protocol

The AMD-760 MPX platform implements **dual point-to-point high-speed system buses** (shown in Figure 3), which allows two processors to run independently without the overhead of sharing a common system bus (shown in Figure 4). Performance delays caused by bus arbitration and bus ownership transitions are eliminated in this architecture, allowing each processor to perform as if it has a dedicated channel to system resources. The split-transaction nature of the AMD Athlon system bus, combined with its independent data and command channels, delivers a high-speed front-side bus solution for AMD Athlon MP processors.

Bus snooping is a critical mechanism in maintaining a system's data coherency. While one processor is accessing memory, the second processor must snoop or "listen" to bus activity and determine if the current memory access affects its memory space. If so, then appropriate measures must be taken to ensure that all affected processors and bus masters have the most accurate data available.

AMD's Smart MP technology implements a performance-oriented snooping mechanism. The processors leverage the independent processor-to-system, system-to-processor, and data channels of the AMD Athlon system bus to create a "virtual" snooping channel. A processor can transfer data while concurrently receiving snoop information, or a processor can broadcast snoop information while concurrently receiving data. The net effect is that concurrency is achieved, and system performance is enhanced. In some non-split-transactions, shared-bus architectures, snooping activity is "focused" only on the current access occurring on the shared system bus. Hence, there may be less opportunity for concurrent data transfers that are independent of the current snoop activity.

AMD multiprocessing platforms implement the **MOESI cache-coherency protocol**. The MOESI protocol offers a potential performance advantage over systems implementing MESI protocol. The additional "Owner" state allows the processor cache "owning" the data to supply data directly to the second processor requesting access to the cached block. The requesting processor no longer has to wait for the owning processor to write the requested data back to main memory before the data is accessible. Instead, the owning processor supplies the requested data directly to the requesting processor. This scheme reduces memory traffic, and allows faster access to cached data.

Dual AMD Athlon™ System Bus²

Server- and workstation-class platforms require high-performance and scalable compute capability. A key element impacting these requirements is the system bus linking the processor(s) to the system logic. Throughout the industry, this bus is generally termed the “front-side bus.”

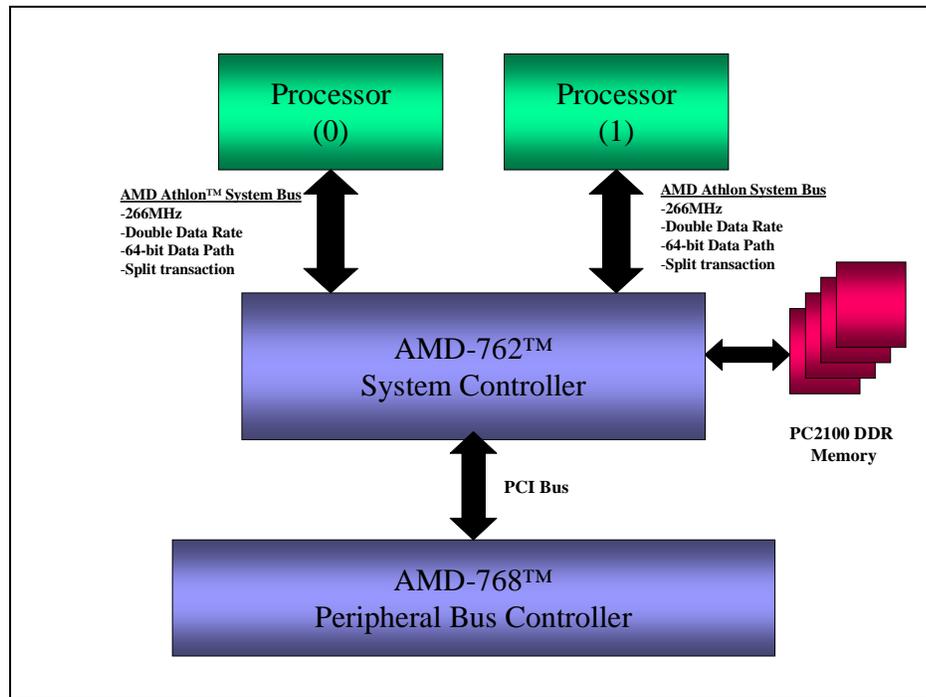


Figure 3: Dual Point-to-Point AMD Athlon™ System Bus Architecture

As shown in Figure 3, the AMD-760 MPX platform employs the AMD Athlon system bus as the high-speed interface between the processor(s) and system elements. Optimized for server and workstation applications, the AMD Athlon system bus offers the following attributes:

- 266MHz³ operation, designed to deliver a peak throughput of up to 2.1GB/s per system bus
- Split-transaction architecture
- Cache-coherency protocol

² The AMD Athlon system bus is also referred to as the AMD system bus.

- ECC capability
- 64-bit data path
- Packetized request transactions

To achieve symmetric multiprocessing capability, AMD chose a dual point-to-point independent bus topology as opposed to the shared bus architecture depicted in Figure 4. The dual point-to-point architecture offers the following advantages over the shared bus architecture:

- Processors on the AMD Athlon system bus can concurrently burst information to and from the system. The shared bus architecture allows only one processor at a time to transfer data.
- Latency is reduced, as processors on the AMD Athlon system bus will not have to arbitrate for system bus control.
- Signal loading, integrity, and termination are much simpler from a design perspective with only one processor per system bus.

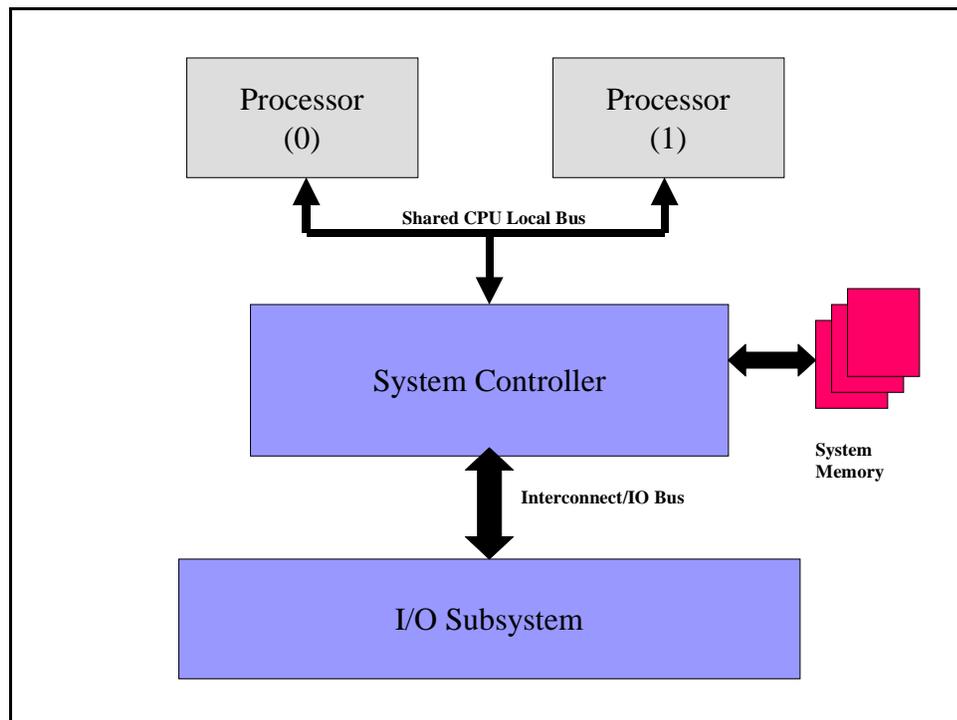


Figure 4: Generic Shared CPU Local Bus Architecture

³ The 266MHz rate refers to a physical clock operating at 133MHz in which information is transferred on each clock edge. Hence, this is calculated as (133MHz clock) x (2 transfers/clock) = 266M transfers/sec = 266MHz.

Split-transaction bus architectures have existed in mini, mid-range, and mainframe computer systems for many years. Only recently have these architectures moved into the desktop market. In essence, split-transaction architecture is a bus transaction technique in which serialized transactions are decoupled and allowed to execute concurrently, or in an overlapping fashion. This technique leverages the latencies in the system to increase overall system performance.

This architecture is different from traditional microprocessor bus architectures in which the address, control, and data buses operate collectively as a single pipe (see Figure 5). The traditional architecture allows only a single transaction to occur at a time, thus creating a serializing effect.⁴

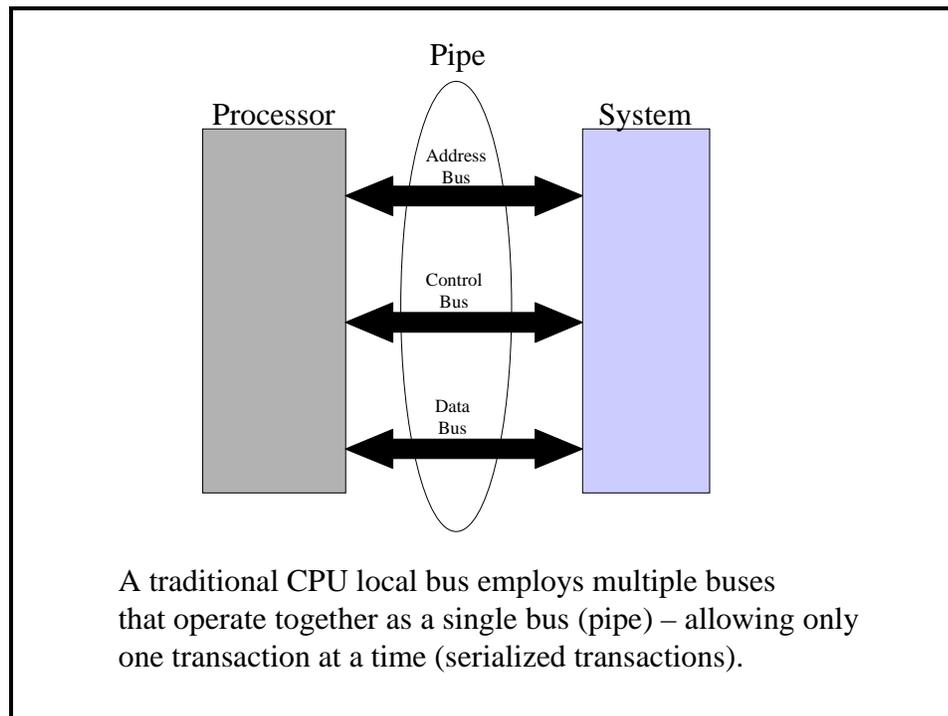


Figure 5: Traditional Microprocessor CPU Local Bus

Figure 6 is an example that illustrates the serialization nature of the non-split transaction type buses. Bus transactions can be decomposed into an address/control phase and a data phase. In the case where the processor is reading data from the system, during the address/control phase, address and control information is broadcast to the system. The

⁴ Techniques such as pipelining and delayed transaction are used to circumvent this type of issue; however, these techniques will be ignored in this white paper to simplify the discussion.

processor must wait for the system to respond with valid data before a new transaction can be launched.⁵ The length of the delay depends on the type of device (memory, disk drive, etc.) being accessed, its associated latency, and its readiness to deliver data. In Figure 6, note how the data phase times can vary from transaction to transaction. In this example, it takes 15 units of time to complete three transactions.

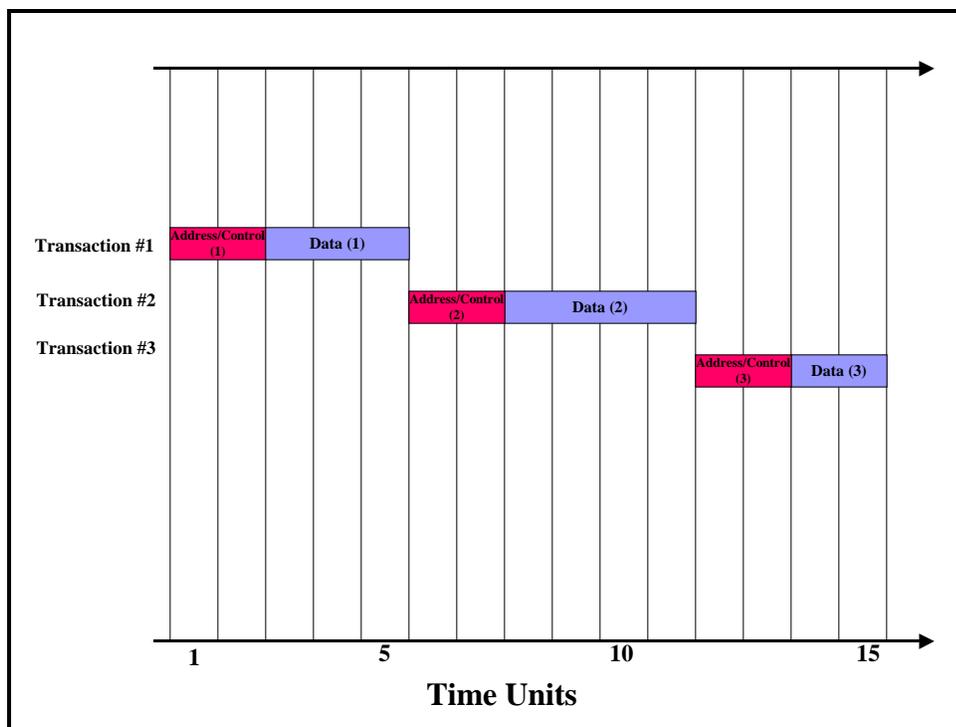


Figure 6: Example of Serialized Transaction Nature of Non-Split Transaction Bus Architectures

Split-transaction architecture improves performance by taking advantage of the data phase latencies, and allowing other transactions to launch as the processor (or system) is waiting for the current data phase to complete. Hence, transactions overlap, concurrency is achieved, and system performance is increased.

⁵ Techniques such as pipelining and delayed transaction are used to circumvent this type of issue; however, these techniques will be ignored in this white paper to simplify the discussion.

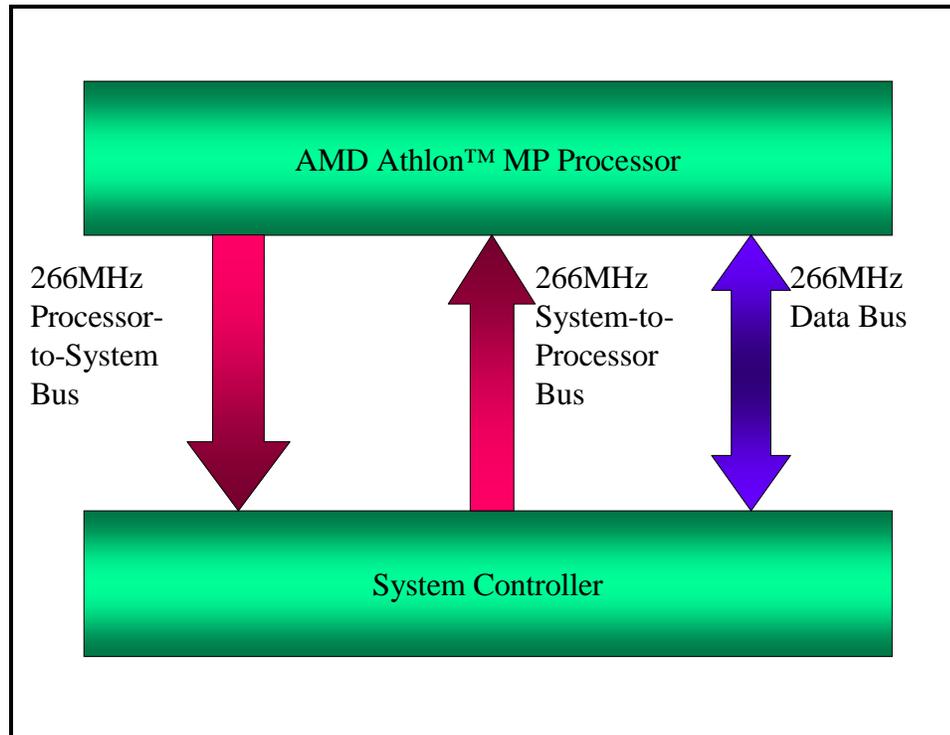


Figure 7: AMD Athlon™ System Bus Components

As shown in Figure 7, the AMD Athlon system bus is composed of three separate buses that operate independently and concurrently:

1. Processor-to-system bus
2. System-to-processor bus
3. Data bus

Note that each bus runs at 266MHz.

The **processor-to-system bus** is the channel in which the processor issues requests/commands (memory read, memory write, etc.) to the system. This is a uni-directional bus controlled only by the processor. Data is not transferred on this bus; only packets containing command, address, and other information are transferred. Data information is transferred over the data bus.

The **system-to-processor bus** is the channel in which the system issues requests/commands to the processor. This is also a uni-directional bus controlled only by the system controller. Similar to the processor-to-system bus, only packets—request and

command information—are transferred on this bus. Data information is transferred over the data bus.

The **data bus** is a bi-directional bus used to transfer data packets between the processor and system elements in response to requests from their respective buses. Each data packet contains an ID tag for association with the corresponding processor-to-system or system-to-processor request.

NOTE: Refer to the AMD Athlon™ System Bus Specification for more details on bus architecture and operation.

As shown in Figure 8, the three independent, high-performance channels (and support logic) enable split-transaction capability. Processor requests (PR) and system requests (SR) can be sent concurrently and independently. System data (SD) and processor data (PD) can be returned to the requestor immediately or at a later time, depending on availability. Ordering is mitigated by instantiating a unique ID associating the request with the corresponding data.

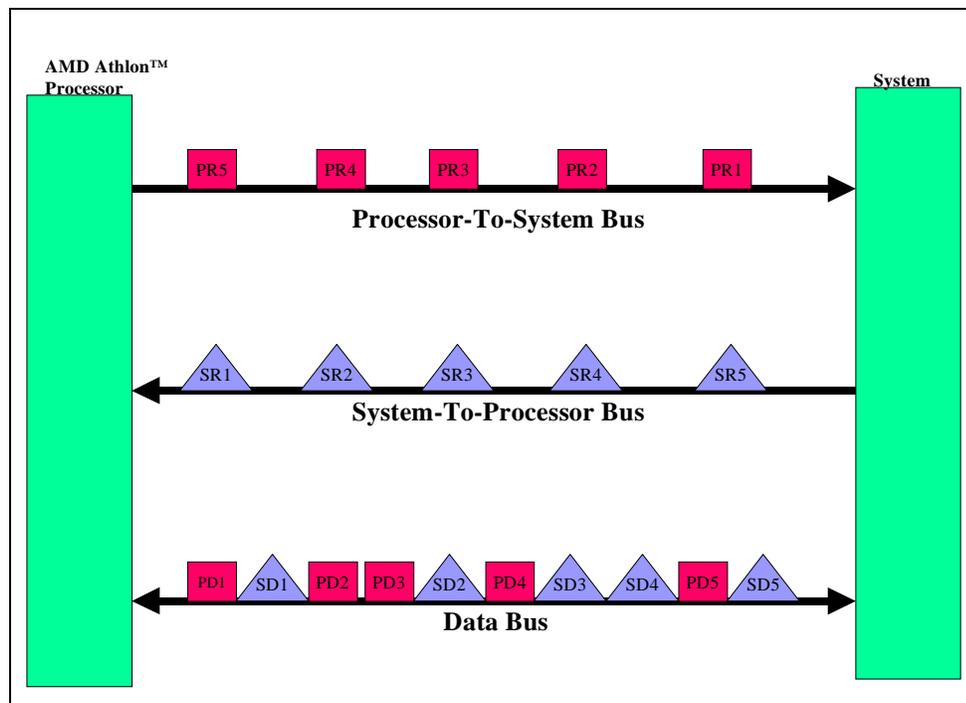


Figure 8: AMD Athlon™ System Bus Split-Transaction Example

The net effect of split-transaction capability is improved performance. As shown in Figure 9, concurrency is achieved by decoupling the processor-to-system transactions from the system-to-processor transactions, and by decoupling the data response from the associated request. Compared to the non-split transaction example shown in Figure 6, the time to complete three transactions has been reduced from 15 time units to 11 time units.⁶

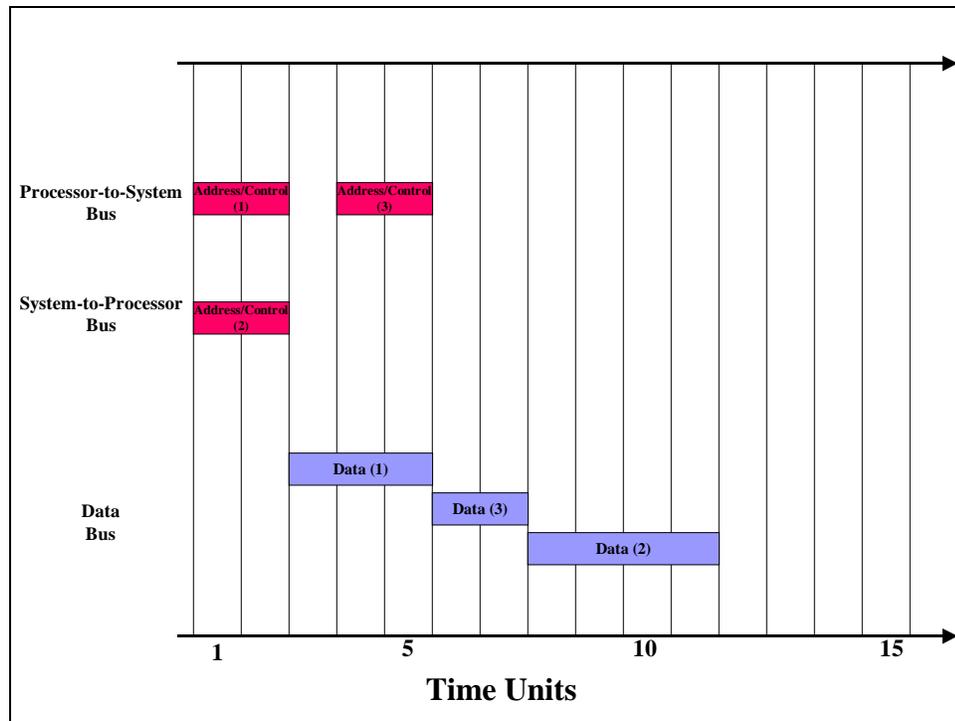


Figure 9: Example Illustrating the Concurrency Offered by Split-Transaction Bus Architecture

The AMD Athlon system bus also supports error checking and correcting across the system bus to evaluate data integrity.

NOTE: When running two processors, both processors must run at the same clock frequencies. Also, the AMD Athlon system bus and DDR memory subsystem are locked in frequency. Whether running one processor or two processors, the AMD Athlon system bus and DDR memory subsystem must operate at the same speed.

In summary, the AMD Athlon system bus offers high-speed communication between the processor and the system. The dual point-to-point multiprocessor topology,

⁶ These performance numbers are arbitrary examples used for concept illustration only.

split-transaction capability, and error checking and correcting features all combine to deliver outstanding performance and protection for mission-critical applications.

Innovative Bus-Snooping Capability

The architecture of the AMD Athlon system bus (and its dual-bus implementation on the AMD-760 MPX chipset) provides an innovative bus-snooping capability that contributes to overall system performance. As memory transactions (memory reads, memory writes...etc.) are requested by each processor, bus “snooping” must occur to determine if the transactions impact another processor or system element. Bus snooping is the process in which the system (processors, core logic, and bus-masters) monitor and track memory requests to ensure data coherency. **(Refer to page 30 for a simple analogy on bus snooping.)**

The AMD-760 MPX platform takes advantage of the independent processor-to-system, system-to-processor, and data channels of the front-side bus to facilitate performance bus snooping.

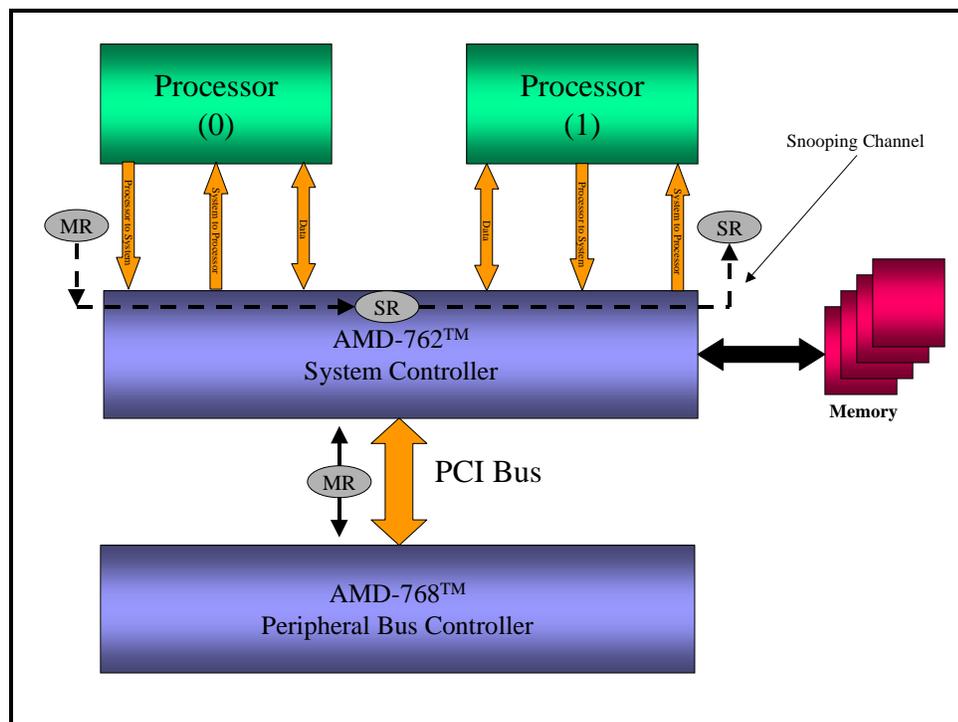


Figure 10: Example 1 – Bus-Snooping Channel

Figure 10 illustrates an example of bus snooping on the AMD-760 MPX platform. Processor (0) launches a **Memory Request (MR)** to the system to obtain data or instructions that are not currently in its cache. The system controller translates this request into a **Snoop Request (SR)** and queries Processor (1) to determine if Processor (1) has the requested data. Processor (1) will respond to the request via messaging on its processor-to-system bus. If Processor (1) has the data, it will return the data to Processor (0), otherwise the system controller will fetch the data from main memory. In essence, a **virtual snooping channel** is created between Processor (0), Processor (1), and the system controller. While the snooping is occurring, note that Processor (0) can concurrently receive messaging and data on its system-to-processor/data buses, and that Processor (1) can concurrently transmit messaging and data over its processor-to-system/data buses. The processors have the potential to concurrently perform transfers that are unrelated to the current snoop activity. This concurrency plays a significant role in improved performance over other multiprocessing architectures.

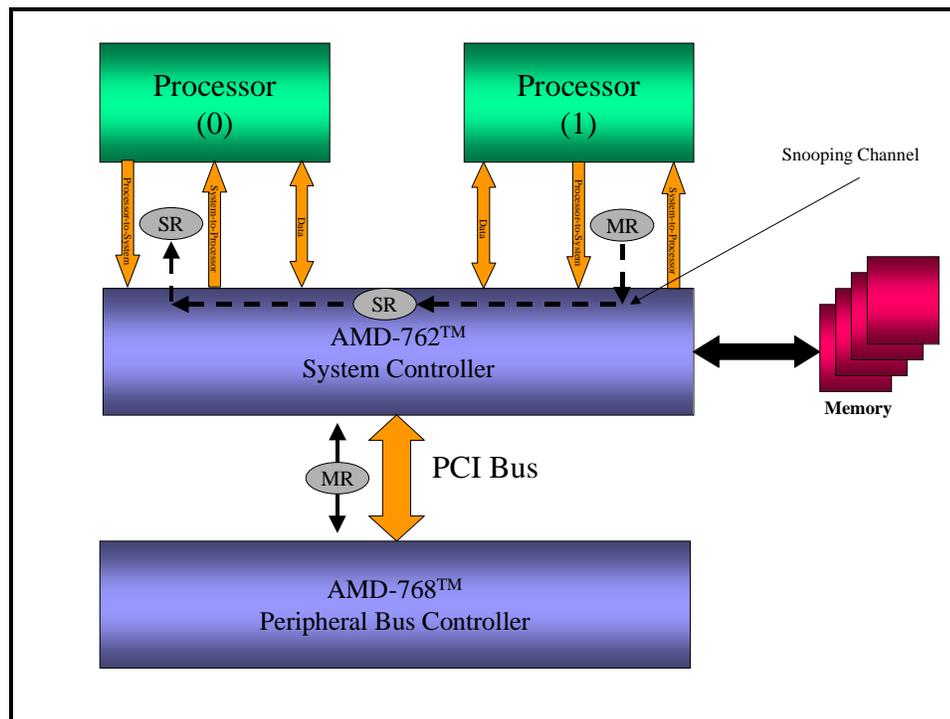


Figure 11: Bus-Snooping Channel Example – Processor (1) to Processor (0)

Figure 11 illustrates a bus-snooping transaction/channel in the reverse direction (from Processor (1) to Processor (0)).

Note also, that the system has to monitor transactions from the PCI bus as part of the snooping process. PCI bus-masters have the potential of accessing memory spaces cached by the processors. Hence, the system controller must snoop PCI traffic and generate snoop requests to the processors. Figure 12 illustrates a virtual snooping channel created between the PCI bus and the processors.

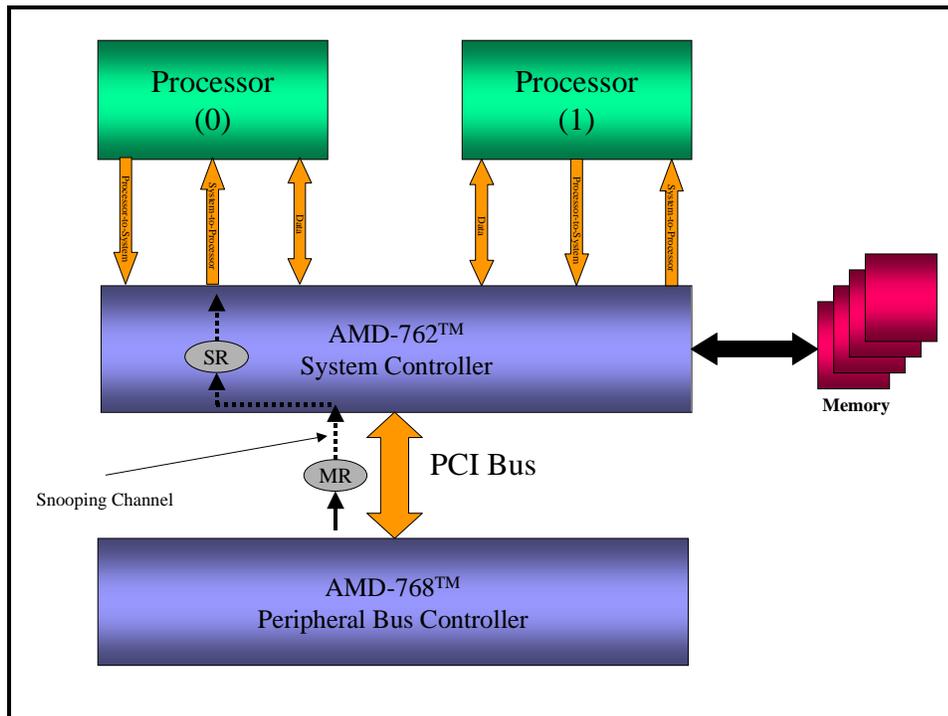


Figure 12: Bus-Snooping Channel Example – PCI Bus-Master to Processor (0)

Optimized MOESI Cache-Coherency Protocol

Although there are several cache-coherency protocols implemented on various platforms, the MESI and MOESI protocols are the most popular in the industry. MESI and MOESI are acronyms representing the cache states in the system.

- MESI – Modified, Exclusive, Shared, Invalid
- MOESI – Modified, Owner, Exclusive, Shared, Invalid

The AMD-760 MPX platform implements the MOESI cache-coherency protocol. The MOESI protocol is similar to the MESI protocol; however, an additional state (Owner) is added, offering increased performance potential over its MESI counterpart. The addition of the owner state allows the processor owning the cache block to directly supply cache data to the requestor. In this respect, the requestor does not have to wait for main memory to be updated to receive the requested data.

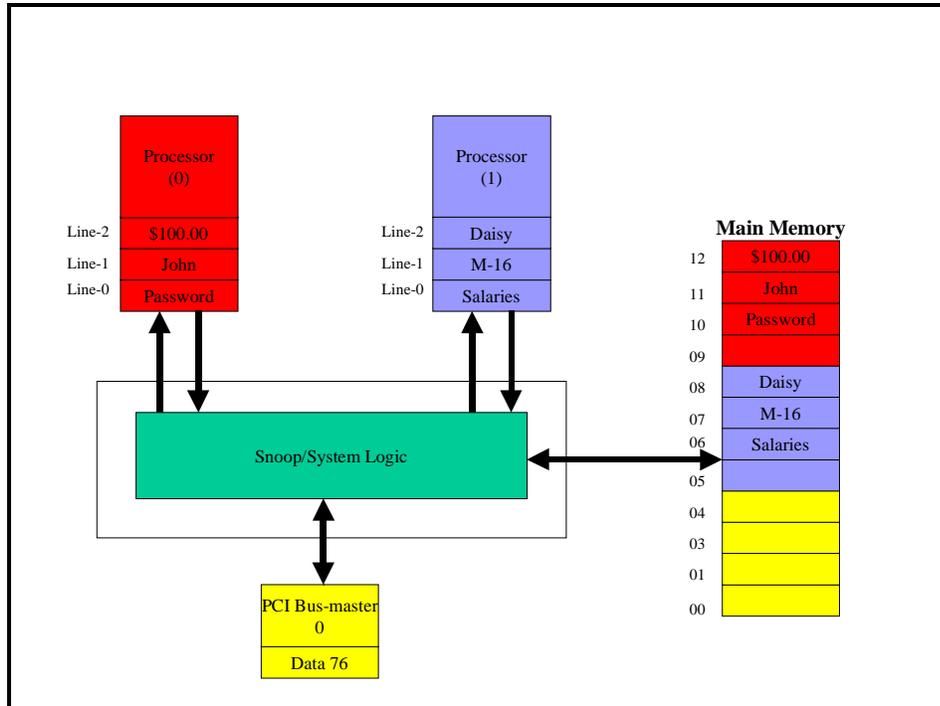


Figure 13: System Cache State Before Transaction

As an example, Figure 13 shows an AMD-760 MPX platform-based architecture with two processors. Each processor/cache (color coded) is mapped to a unique memory space in main memory. In this example, the corresponding cache lines match the corresponding main memory image. Also note, that the PCI bus-master has an assigned memory map in main memory.

Suppose that Processor (1) needed financial information from main memory that is currently “owned” by the cache in Processor (0) (Line-2). Also, suppose that the PCI bus-master concurrently needed to write data to its own allocated memory block in main memory. Because of the MOESI cache-coherency protocol and the AMD-760 MPX architecture, these two operations can occur concurrently. As shown in Figure 14,

because Processor (0) is the “owner,” Processor (0) can supply the information directly to Processor (1) without Processor (1) having to access main memory. Main memory is free to be accessed by the PCI bus-master. Due to the AMD-760 MPX architecture, the transfer of data from Processor (0) to Processor (1), and the data transfer from the PCI bus-master to main memory can take place at the same time. This concurrency increases performance in scenarios where these conditions occur. In some alternative architectures, main memory would have to be updated first, and the PCI bus-master would have to wait for the memory update to complete before it could perform a transaction to its corresponding space in main memory.

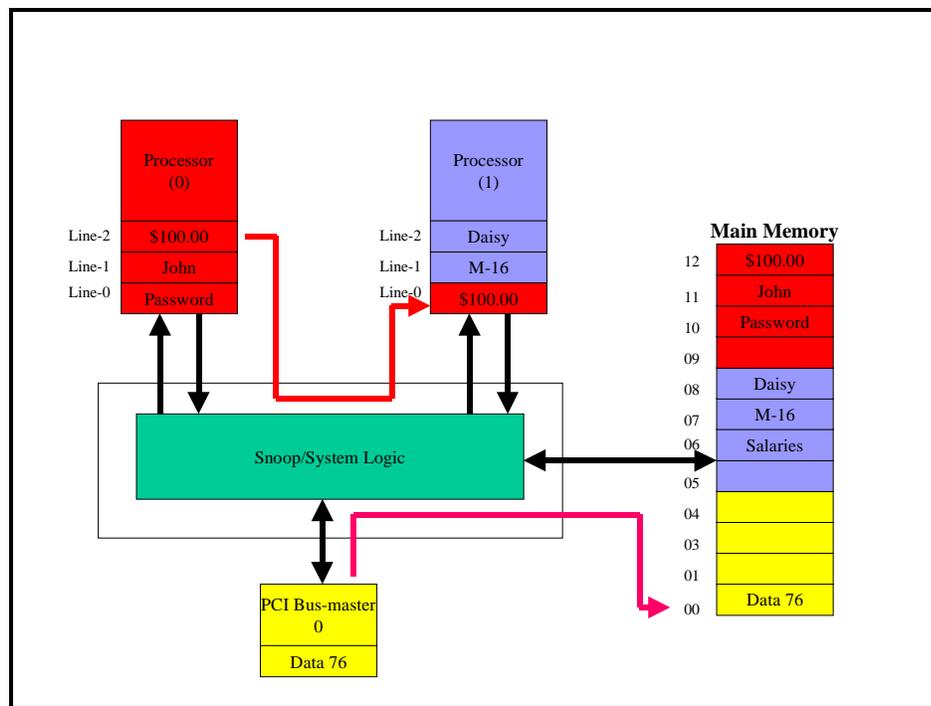


Figure 14: System State After Memory Transaction

Again, the MOESI architecture improves system performance by allowing data owners to supply data directly to the requestor without consuming memory bandwidth.

DDR Memory Subsystem

The AMD-760 MPX platform implements the latest evolution in high-performance memory technology—Double Data Rate (DDR) SDRAM. DDR memory is

a natural extension of current PC100/PC133 SDRAM technology. However, DDR offers higher performance at a competitive cost.

Similar to the AMD Athlon system bus, the DDR memory subsystem is designed to deliver a peak transfer rate of up to 2.1GB/s. To satisfy the large memory requirements of server- and workstation-class systems, the AMD-762 system controller supports up to 4GB of memory space and provides interfacing for four registered DIMM slots.

NOTE: Unbuffered memory DIMMs are not supported.

In applications requiring minimal fault tolerant capability, the AMD-760 MPX DDR memory subsystem supports ECC memory and provides detection and correction of single-bit errors.

AGP-4X Graphics Subsystem

The graphics subsystem is a critical element of high-performance workstations. The AMD-760 MPX platform implements AGP (Accelerated Graphics Port) technology as its graphics subsystem of choice. The intent of the AMD-760 MPX platform is to increase overall system performance by eliminating data-movement bottlenecks, allowing a more efficient match to the compute performance offered by faster AMD processors. Increasing performance without impacting system cost is a difficult design challenge; however, AGP graphics technology offers the perfect balance to satisfy even the most extreme user.

The AMD-762 system controller is equipped with an AGP-4X graphics interface designed to provide powerful graphics capability to the workstation desktop. Optimized to run concurrently with the AMD Athlon system bus and DDR memory interface, high-end graphics adapters can truly utilize up to 1GB/s of bandwidth potential offered by the AGP-4X subsystem.

Server applications typically do not require high-end graphics capability. The AMD-762 system controller's AGP-4X interface is backwards-compatible with lower

performance AGP-1X and AGP-2X modes. This feature allows system designers the flexibility to leverage lower-cost AGP graphics adapters in their system solutions, thus reducing overall platform cost.

ATA-100 Storage Subsystem

Due to the physical/mechanical nature of hard drives, the storage subsystem is one of the slowest elements within all computer systems. This has a significant impact on overall system performance, since all operating system and software applications are initially loaded from the hard drive. Essentially, the storage subsystem “data pipe” becomes a critical bottleneck in the system.

To alleviate this performance issue, the AMD-760 MPX platform implements a storage subsystem controller that offers increased data transfer rates between the EIDE hard drive controller and the actual storage device (hard drive, CD/CD-R/CDRW, DVD, etc.). As shown in Figure 15, the AMD-768 peripheral bus controller embeds an EIDE storage controller that offers data transfer rates of 33MB/s, 66MB/s, and 100MB/s. These data transfer rates are compliant with the ATA/UDMA-33/66/100 standards (shown in Table 3).

Although 5400 RPM drives are supported, optimal EIDE storage subsystem performance is obtained when the ATA-100 mode is used in conjunction with hard drives that operate at spindle speeds of 7200 RPM or higher.

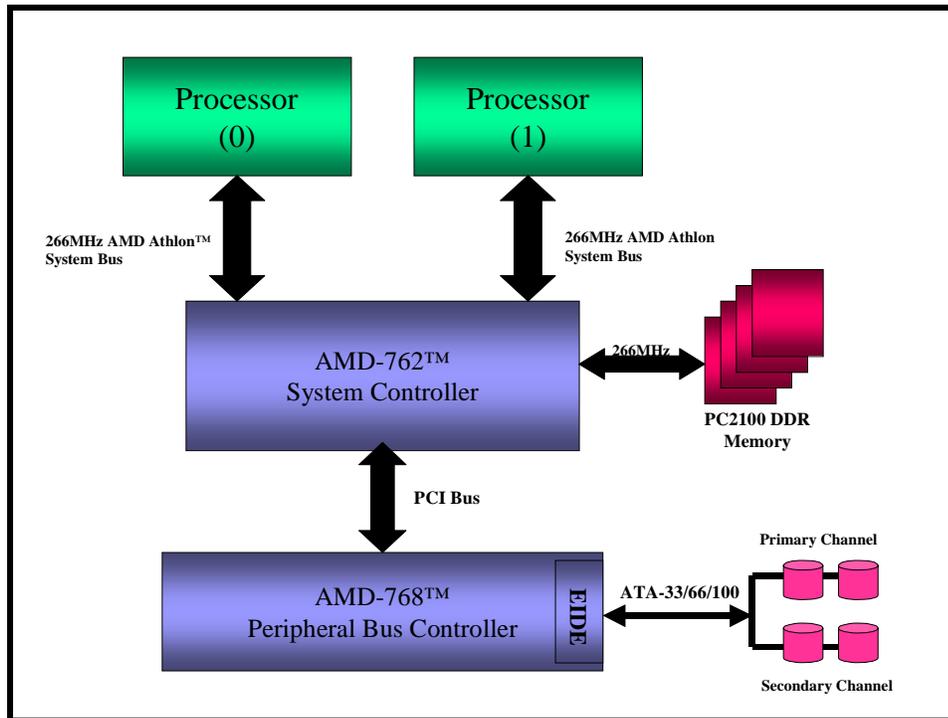


Figure 15: The AMD-760™ MPX Chipset EIDE Controller Supporting ATA-33/66/100MB/s Speeds

Table 3: AMD-760™ MP Chipset EIDE Interface Speeds

| ATA Mode | Data Transfer Rate |
|--------------|--------------------|
| ATA/UDMA 33 | 33MB/sec. |
| ATA/UDMA 66 | 66MB/sec. |
| ATA/UDMA 100 | 100MB/sec. |

Primary (66MHz) and Secondary (33MHz) PCI Bus Interfaces

To increase system performance, the AMD-760 MPX chipset implements a dual PCI bus architecture. As shown in Figure 16, the dual PCI bus architecture consists of a **66MHz/64-bit/32-bit primary PCI bus interface**, and a **33MHz/32-bit secondary PCI bus interface**.

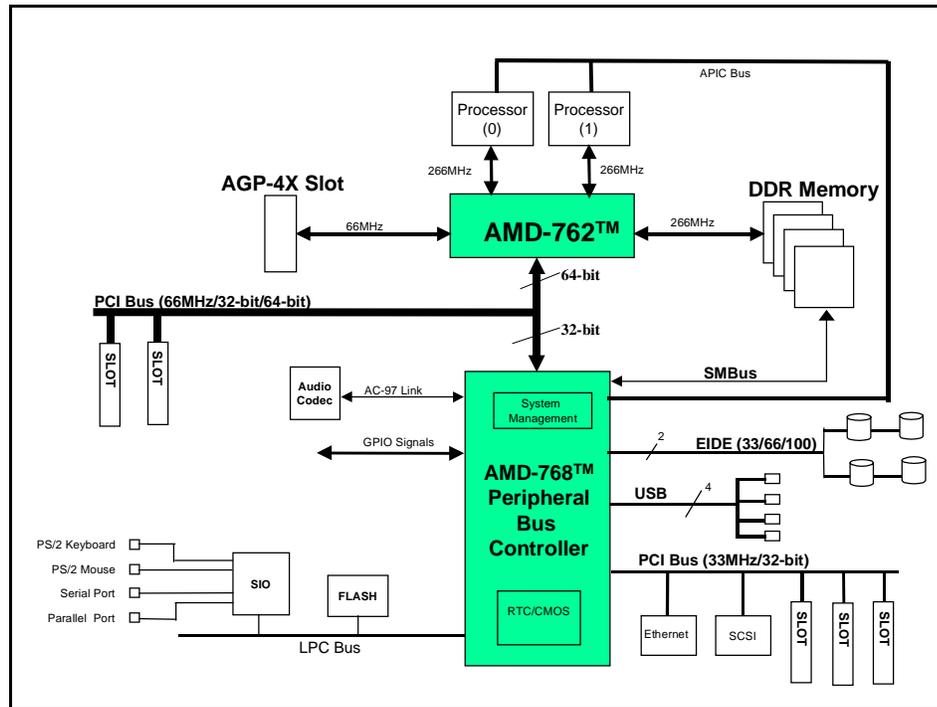


Figure 16: AMD-760™ MXP Architecture Illustrating PCI Bus Connectivity

The **Primary PCI** bus interface serves two functions:

1. High-speed 66MHz/32-bit interface between the core-logic elements (AMD-762 system controller and AMD-768 peripheral bus controller).
2. High-speed expansion card support for 66MHz/64-bit/32-bit PCI bus-master adapter cards.

Although the AMD-762 system controller provides a 64-bit data path on the primary PCI interface, the AMD-768 peripheral bus controller only uses the lower 32-bits, since it is only a 32-bit device (64-bit bandwidth is not required by the AMD-768 peripheral bus controller). However, the AMD-762 system controller can support up to two 66MHz/64-bit/32-bit PCI adapter cards, allowing high-speed devices (such as storage, communication, and networking devices) to stream data quickly.

The **Secondary PCI bus** interface provides a 33MHz/32-bit PCI interface for traditional PCI devices. The AMD-768 peripheral bus controller provides arbitration

capability within its secondary PCI bus to support up to eight PCI devices.⁷ These eight devices can be a mixture of slots and on-board motherboard components.⁸

Table 4 lists the peak data transfer rates associated with each of the PCI buses. These buses allow the system integrator the flexibility to balance cost and performance. For example, in workstation applications, inexpensive 33MHz/32-bit network adapters can be used for connectivity, while higher-performance 66MHz/64-bit adapters can be used to enable a fast SCSI storage subsystem. In server applications, high-performance 66MHz/64-bit adapters can be used to enable a massive RAID storage subsystem or to accommodate several high-speed multi-port networking/communication adapters.

Table 4: PCI Bus 66MHz and 33MHz Peak Data Transfer Rates with 32-Bit and 64-Bit Data Paths

| PCI Bus | Data Path Width | Peak Bandwidth |
|-----------|-------------------------|----------------|
| Primary | 66MHz, 64-bit data path | 528MB/sec. |
| Primary | 66MHz, 32-bit data path | 264 MB/sec. |
| Secondary | 33MHz, 32-bit data path | 132MB/sec. |

AC-97 Audio Interface

The AMD-760 MPX chipset integrates an AC-97 Interface⁹ within the AMD-768 peripheral bus controller to enable two-channel (left and right) audio onto the platform. As shown in Figure 17, a motherboard AC-97 audio CODEC can be attached to this interface (along with the appropriate connectors) to provide an inexpensive base audio solution for workstation applications. This feature offers two advantages.

1. The system integrator doesn't have to suffer the additional cost of a more expensive PCI audio card solution.

⁷ REQ/GNT#[6:0] pairs are allocated for secondary PCI bus slots and on-board devices. A separate pair of signals (PREQ/PGNT#) is dedicated to a priority PCI device.

⁸ AMD designs and tests its reference design motherboards to support three secondary PCI slots. If more slots or additional motherboard devices are desired, the system designer must perform the appropriate analysis and testing to ensure signal integrity and proper operation.

⁹ AC-97 soft modem capability is not supported by the AMD-760 MPX chipset.

2. Since PCI slots are limited by the motherboard form-factor, a valuable PCI slot is not consumed by a PCI audio adapter, allowing the PCI slot to be used to enable a more critical function (such as networking, SCSI storage subsystem, etc).

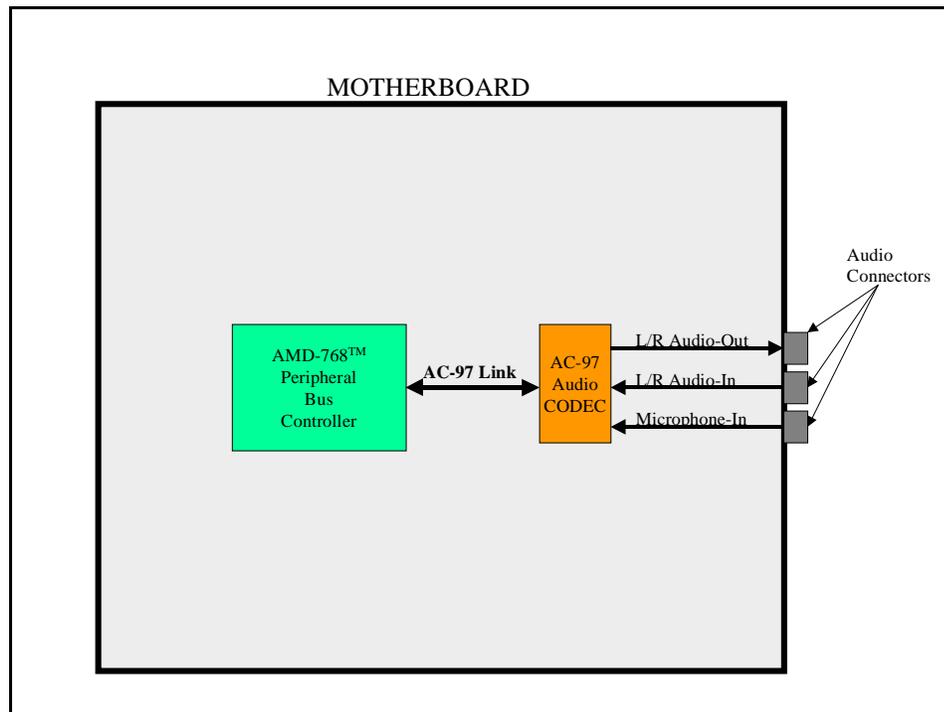


Figure 17: AC-97 Audio Support

As shown, two-channel audio capability includes the following:

- Left and right audio-out
- Left and right audio-in
- Microphone-in

The system designer can choose from a variety of supported CODECS to meet a range of cost and performance objectives.

Appendix

Glossary of Terms

AGP—Accelerated Graphics Port

BGA—Ball Grid Array

CODEC—CODer\DECoder

DDR—Double Data Rate

ECC—Error Correcting Code

I/O—Input/Output

MP—Multiprocessor

SBA—Side-Band Addressing

What is Bus Snooping?

Imagine sharing a checking account with several co-workers. It becomes very important to monitor who made deposits, who made withdrawals, and when the transactions occurred. The objective is to ensure that each person understands the current balance of the account and the pending transactions before a check is written, and hopefully the account will never be overdrafted. Now imagine the process that each of the participants has to perform before he or she writes and commits a check.

1. What is the current balance?
2. What are the pending transactions from myself and the other account participants?
3. Is anyone trying to access the account now, and what type of transaction is occurring?
4. Based on these questions, is the current/pending balance enough to cover the check that is about to be written? Can this transaction be committed without penalty?

Simply, the objective is to understand whether or not the current balance is accurate and to obtain the accurate balance before the transaction is committed. The situation to be avoided is an overdraft, to which a severe penalty will be charged.

This scenario is similar to the issues encountered in multiprocessing systems when several processors and bus-masters must share main memory. The system must ensure that each processor obtains accurate and timely information within a dynamic shared environment. The data obtained by each processor must be “coherent,” before the processors act on the data. Just as an overdraft must be avoided in the shared checking account example, “incoherent” data must be avoided in multiprocessing systems, otherwise errors will result with severe effects.

Bus snooping is the process of monitoring memory transactions in a multiprocessor system to ensure that the memory transaction requestor (typically a processor or bus-master) receives accurate and timely data. The bus-snooping process and mechanisms must ensure that data is coherent throughout the system, otherwise an error will occur. Similar to the thought process steps in the shared checking account example, the processors and core-logic elements must execute a similar process to ensure coherency. If a processor is performing a memory transaction, the following must be determined:

1. Is the data requested contained in the requestor’s cache? If so, is it stale data or is it accurate data?
2. If the data is stale, is the accurate version in main memory or in another processor’s cache?
3. If the data is in another processor’s cache, has the other processor recently changed the data?

All of this analysis and the actions that move and update the data to maintain coherency is part of the bus-snooping process. Bus-snooping and data-coherency maintenance is the system responsibility (i.e., a concerted effort by the processors, core-logic, and participating bus-masters).

AGP Graphics Background

To understand the advantages and benefits of AGP graphics, it is necessary to understand the issues that AGP technology has resolved. Figure 18 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. Although this architecture performed well in its time frame, several issues arose 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 some graphics data (such as textures and other information) are stored in main memory, the PCI-based graphics card must access main memory via the PCI bus. These accesses may 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.

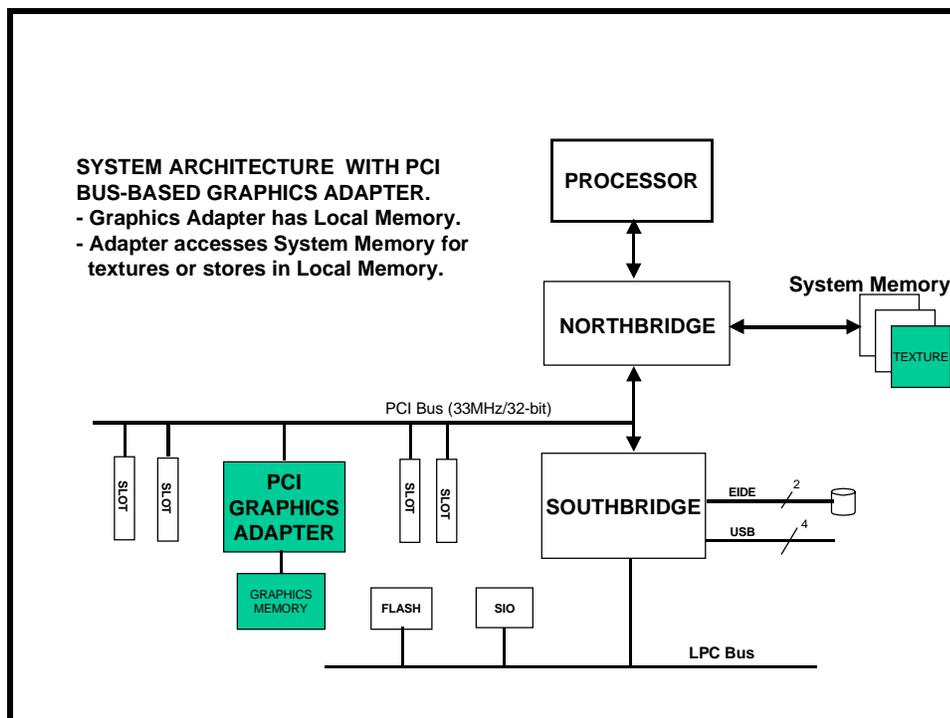


Figure 18: Older System Architecture Showing PCI Bus-Based Graphics

Figure 19 illustrates how AGP technology elegantly resolves the issues facing PCI bus graphics architecture. The AMD-762 system controller (the Northbridge element of the AMD-760 MPX chipset) 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 (See Figure 20). The Northbridge embeds read/write and command queues (buffers) to allow full-speed data and command transport between the AGP device and the AMD-762 system controller, and concurrent full-speed data transport between the AMD-762 system controller and the DDR memory subsystem.

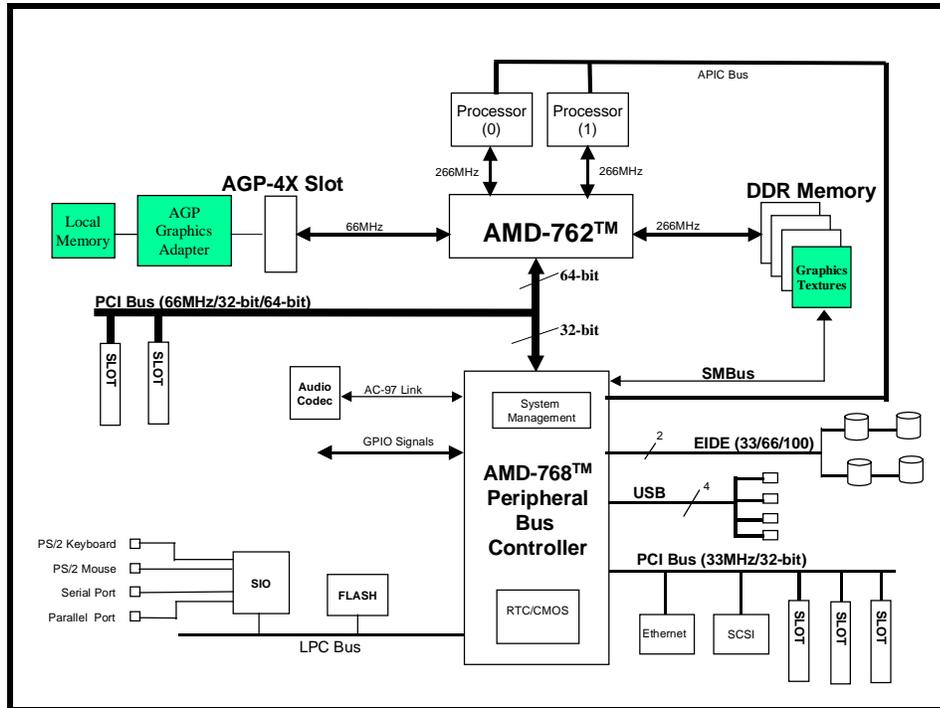


Figure 19: AMD-762™ System Controller AGP Graphics Architecture Showing Use of System Memory for Graphics Operations

The architectures shown in Figure 19 and Figure 20 produce the following benefits:

- The native architecture of AGP graphics subsystem (66MHz PCI bus interface with side-band addressing, and embedded Northbridge AGP logic) 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 that 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.

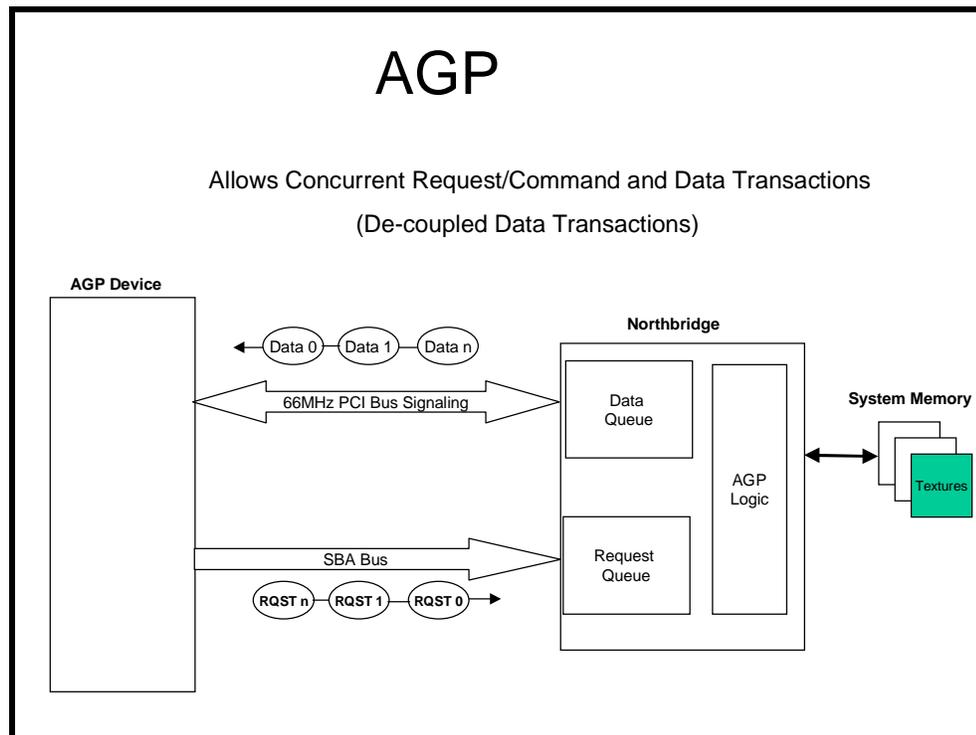


Figure 20: High-Level AGP Interface Diagram Showing Bus Architecture and Embedded Northbridge Components

Over time, the AGP graphics subsystem has scaled to increasing levels of performance. As shown in Figure 20 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. Finally, third gear is AGP-4X, offering the highest performance data transfer rate of up to 1GB/sec. (The notation 2X and 4X are relative to the original AGP-1X mode). In the future, higher performance modes (overdrive and turbo gears, perhaps) may be added.

Table 5: 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. |

NOTE: The AMD-760 MPX chipset is designed to support all modes shown in Table 5.

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 \$4.6 billion in 2000. (NYSE: AMD).

© 2002 Advanced Micro Devices, Inc. All rights reserved.

AMD, the AMD Arrow logo, AMD Athlon, AMD-760, AMD-762, AMD-768 and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.