



HyperMemory™

*Next-Generation memory management
for PCI Express® graphics*





Introduction.....	3
Background.....	3
Accelerated Graphics Port (AGP)	3
Beyond AGP	4
2004: Introducing PCI Express®	4
Next Generation Operating systems	4
Introducing HyperMemory™	5
Pre-emptive Virtual Memory	5
Increasing Memory Bandwidth with PCI Express®	6
Conclusion	7

Introduction

The availability of significant amounts of high-bandwidth memory is a key component of an efficient, high-performance graphics subsystem. For this reason, high-end graphics subsystems in modern enthusiast and workstation configurations are often seen with: 256 Megabytes of Double Data Rate memory operating on a 256-bit wide memory interface.

While such configurations do allow for high performance by allowing a matched VPU core to always have ready access to a significant pool of working memory to 'feed' it, these configurations also come at significant cost.

This cost is incurred at several levels: the VPU ASIC itself is rendered more expensive by the need for more memory interface silicon die area and by the pins to actually connect to memory devices, and there is the actual cost of large numbers of high-capacity, high speed memory devices.

The graphics subsystem's memory has typically remained distinct from the main system memory because of the need for high-speed and low-latency memory which comes at a cost premium.

A problem thus arises when trying to reduce total system cost: performance must be sacrificed in order to do so. The sacrifice comes in the form of provided less memory, with less bandwidth (less memory devices and/or slower memory).

An intelligent memory management technology outlined in this document can alleviate this sacrifice and provide more performance at a lower total cost.

Background

It is important to review technologies that have previously been put forwards to try and make the best of the performance vs. cost tradeoff. While in recent years memory prices have fallen (for given quantities), a dedicated memory subsystem for graphics still represents a significant cost. This cost has been in place for some time, and the need to tackle it has been ever-present, even as the demand for consumer-class 3D applications has risen, putting upwards pressure on memory requirements.

Given that system memory has traditionally been available at a lower cost (by virtue principally of lower densities and lower speeds), it has long been seen as desirable to somehow use this system memory rather than the more expensive local graphics memory (LGM). A key technology advance proposed a way to do this: the AGP interface.

Accelerated Graphics Port (AGP)

The AGP 1x interface was introduced in 1997 in order to alleviate the need for local graphics memory by providing a high-bandwidth (double that of the fastest interconnect of the time, PCI) interface to the system core logic and system memory. This high bandwidth connection attempted to allow the data required by a 3D VPU to do its work to be hosted in system memory rather than in more expensive local memory. This was in order to facilitate the adoption of 3D capabilities.

One of the key features of AGP is the ability to map a portion of system memory to be addressable by the graphics controller. This AGP aperture benefits from the speed of the AGP bus to act as memory store that, while not local, is readily accessible.

Beyond AGP

The AGP interface has enjoyed success in the PC ever since introduction, largely by virtue of the technology updates that have helped it keep up with the high-bandwidth requirements of 3D applications. The revisions of AGP – 2x, 4x and 8x – have culminated in an available bandwidth of 2.133 Gigabytes/second flowing from the system side (core logic, system memory) to the graphics subsystem (VPU, LGM).

Over the same period of time, the PCI interface itself was seeing updates for speed in the form of PCI-X, though this was primarily for data intensive workstation and server areas (SCSI hard-disk controllers, multi-port Ethernet adapters, Fiber channel, etc).

It became clear that after 10 years of the PCI interface and 6-7 years of AGP, that a new interconnect was required in order to overcome the problems facing system designers as they tried to push the bandwidth of the PC system.

2004: Introducing PCI Express®

The PCI Express interface was developed as a forward-looking solution to the increasing bandwidth requirements of the PC system, while leaving behind the problems that plagued PCI-X and AGP at high speed (fundamentally limitations of a parallel interconnect as these are).

2004 has seen the introduction of the first generation of PCI Express which already delivers substantially more bandwidth per physical pin/line of interconnect and thus promises to improve PC system performance at many levels, while keeping costs and design complexity lower.

Naturally, given how bandwidth intensive 3D applications are, PCI Express is of key use for graphics, and can enable a new class of applications.

Next Generation Operating systems

With the wide availability of 3D capabilities in PC systems (virtually all PCs shipped in the last 5 years have had some level of 3D acceleration), and the increasing processing capabilities of the VPU, developers of operating systems see an opportunity to exploit these capabilities to render richer user experiences and higher graphical quality.

This has the effect of 'promoting' graphics and VPUs to an unprecedented level. Because these operating systems (and the applications that run on them) will be so dependent on the VPU, the need to abstract the VPU into a processor becomes apparent, and for that reason, virtualization of the VPU will take place.

Virtualization means abstracting the physical capabilities of the device (the VPU in this case) and exposing the resources (notably processing time and memory) in a way that is independent of the real capabilities. This ensures that many applications can use the VPU resources simultaneously and the operating system can manage the VPU as a common resource.

This is already done in PC systems for the CPU and memory; the CPU is 'time-sliced' for multitasking, and system memory is also virtualized (hence the term 'virtual memory').

The two key elements of virtualization are thus:

- VPU processing time (cycles), achieved by context switching (swapping VPU state between multiple tasks as they are worked on simultaneously)
- VPU local graphics memory.

While context switching is a technology that will truly shine when Operating systems are able to exploit it, virtual memory is a powerful technology that is of use in the immediate, and serves as the basis for HyperMemory.

Introducing HyperMemory™

ATI has developed a technology that effectively introduces virtual memory for VPU's ahead of its time. Thanks to the high bandwidth, bi-directional PCI Express interface, the VPU's memory can be virtualized. Additionally, the PCI Express interface can be treated by ATI's PCI Express VPU's as an additional memory channel.

HyperMemory is thus a technology with two key components:

- Pre-emptive Virtual Memory and intelligent memory allocation, permitting a VPU to have less local graphics memory, and not suffer the typical performance loss that comes with reducing LGM.
- A memory controller attached directly to the VPU's PCI Express interface, allowing the PCI Express interface to be treated as a local memory resource.

Pre-emptive Virtual Memory

Pre-emptive virtual memory and the intelligent memory allocation it brings are technologies that ATI has readied for next-generation operating systems and that ATI is bringing to customers today because the performance enhancements it can bring to modest and low memory configuration PC systems. In combination with the PCI Express auxiliary memory channel, low memory configurations are given an unprecedented edge.

Pre-emptive virtual memory introduces a totally new layer of intelligence into the use of memory by 3D applications. Previously a VPU's memory manager would allow applications to allocate memory as they saw fit. An intermediate memory manager at the API or OS level will typically try to fit all allocations into local memory first, and into system memory if that fails (see diagram below)

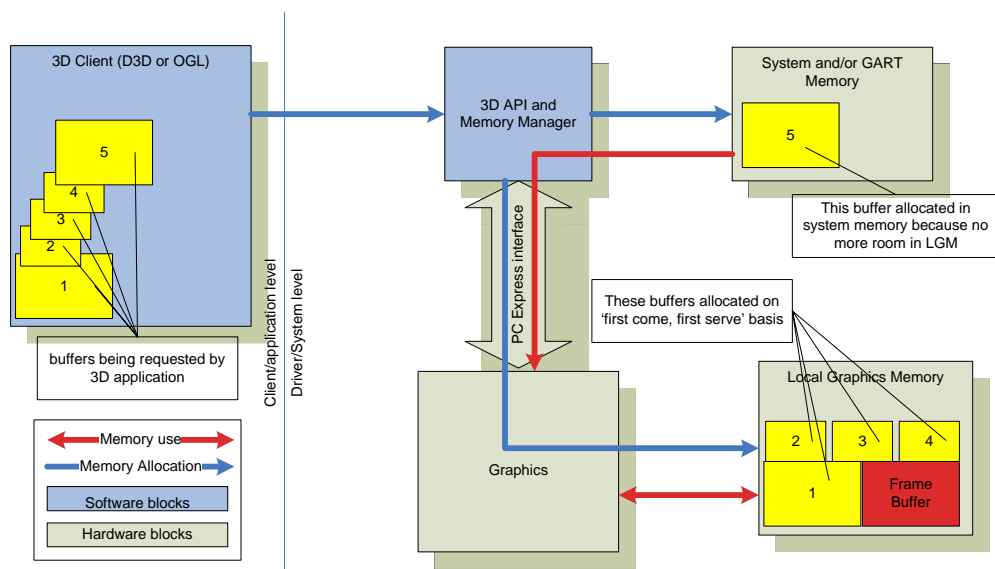


Illustration of what happens in 'typical' memory allocation with no intelligence

Unfortunately, the applications rarely have a good picture of the exact memory allocation at the time they are running, and it is difficult for applications to optimize to every memory configuration even when they do. The result is often a poor use of the precious memory resource.

By using techniques drawn from the field of Artificial Intelligence, pre-emptive virtual memory monitors a number of characteristics of memory allocations as they are requested, and does two key things as a result:

- Preemptively removes unused memory blocks (textures and other allocations) from local graphics memory, when memory is needed for higher-demand uses. This involves 'evicting' them to higher-latency memories (GART memory or even system memory).
- Re-orders memory allocations requested by 3D applications for LGM, making intelligent decisions about where to actually store the data. This works in conjunction with the PCI Express Auxiliary Memory Channel to ensure that all memory channels are used to their best potential.

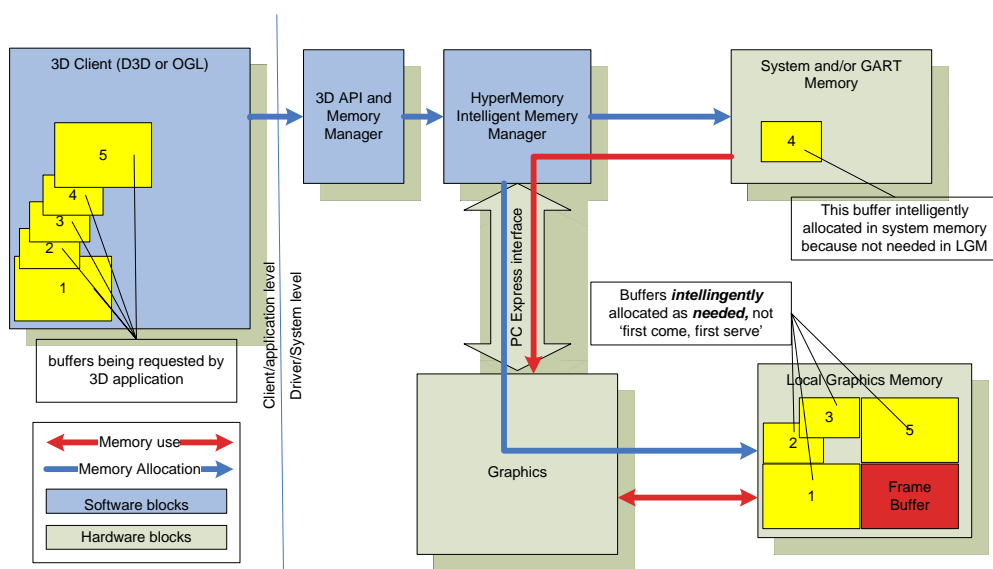


Illustration of memory allocation and use with HyperMemory's intelligent memory manager in place.

Increasing Memory Bandwidth with PCI Express®

With the arrival of a new high-bandwidth interface like PCI Express, ATI not only aggressively pursued implementation of this technology to become the first company to implement, demonstrate and ultimately ship products based on the new interconnect, but also went further with a method to maximize the advantage of PCI Express.

The memory controller on ATI's PCI Express products has been updated to interface to the PCI Express bus and the system memory that lies on the other end of that bus (either directly in cases where the core logic has the system memory controller or where the host CPU assumes this function).

This PCI Express auxiliary memory channel is effectively a 64-bit memory channel with access to system memory. This means that a VPU equipped with a 64-bit local graphics memory bus and a PCI Express auxiliary memory channel has an effective 128-bit memory bus.

It should be noted that two key factors will affect how effective the 64-bit interface of the PCI Express auxiliary memory channel is:

- How saturated the system memory is with other memory transfers (i.e. access from CPU, DMA access from I/O devices etc)
- How much traffic is being sent/received across the PCI Express link in conjunction with the auxiliary memory channel

Below is a diagram showing how the PCI Express auxiliary memory channel fits into a PC system.

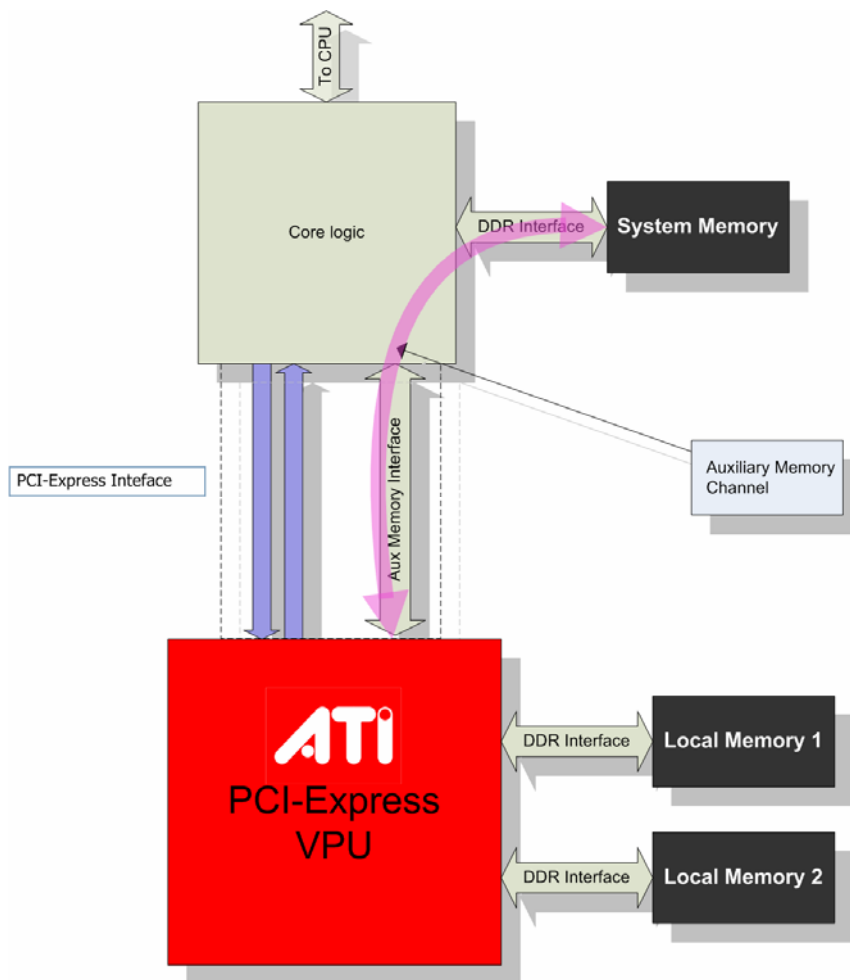


Illustration of graphics subsystem equipped with a PCI Express auxiliary memory channel

Conclusion

It is readily apparent that the two core elements of HyperMemory, pre-emptive virtual memory and PCI Express auxiliary memory channel are complementary and serve to maximize the investment in relatively expensive local graphics memory. This goes to the point of enabling system designers to add fully featured VPUs to their systems without the cost of large memories.

ATI has again proven technology leadership in the area of new device interfaces and maximizing return on investment for customers.

Copyright 2005, ATI Technologies Inc. All rights reserved. ATI and ATI product and product feature names are trademarks and/or registered trademarks of ATI Technologies Inc. All other company and product names are trademarks and/or registered trademarks of their respective owners. Features, availability and specifications are subject to change without notice.

