Latest Reviews

- NVIDIA Fermi GPU and Architecture Analysis
- ATI Cypress Gaming Performance Analysis
- ATI Cypress GPU and Architecture analysis
- ATI RV740 GPU and Architecture Analysis
- NVIDIA GT200 GPU and Architecture Analysis

Latest Articles

- Travelling in Style: Beyond3D's C++ AMP contest - WINNER ANNOUNCED
- Dave Kirk - NVIDIA Tesla Launch
- AMD's John Bridgman on Radeon, Linux and Open Source 3D
- A speculative look on the Wii U GPU
- Andy Keane Interview & Tesla Adoption/Deployment
- Diving into Anti-Aliasing
- Was Harry Potter actually into rendering? Beyond3D's first ever book review
- Q&A with Visceral's Technical Art Director Doug Brooks on Dead Space 2
- E3 2011: Behind Closed Doors - Witcher 2 Xbox 360... and a Nurse
- Ian Buck - NVIDIA Tesla Launch

ATI Xenos: Xbox 360 Graphics Demystified - Page 1

Published on 13th Jun 2005, written by Dave Baumann for Consumer Graphics - Last updated: 21st Mar 2007

Foreword

Those that have followed the development of 3D graphics over the past ten years or so will have seen a continual development of the capabilities of the processors, but fundamentally following the path of OpenGL pipeline model. 3dfx really ignited the market with their "Voodoo Graphics" add-in boards, which were not much more than just a raster engine: it utilised one chip for texture sampling and another for pixel processing (a simple Render Output unit - ROP); 3dfx further evolved that by adding an extra texture unit, allowing for slightly more complex effects in the raster pipeline. And so it was that this model was followed for a number of years with the main developments being the number of pixel pipelines and textures supported per pipeline, until NVIDIA took the step of moving further forward on OpenGL pipeline and giving accelerated support to the Transformation and Lighting process with GeForce 256. Whilst graphics processors had varying degrees of the geometry process, from clipping to setup, handled in hardware, adding a T&L engine was a significant step up the OpenGL pipeline, but didn't really fundamentally change our thinking of graphics processors.

At the same time as the graphics vendors started giving us T&L engines the pixel processors gradually increased in flexibility as well, up until the point that "programmable shader architectures" were all anyone could talk about. The pixel pipelines became more flexible such that they had limited programmability, as did vertex processing, with vertex shaders operating in parallel with T&L engines. Nowadays the level of programmability of both vertex and pixel shaders has increased significantly with each vertex shaders enveloping the T&L processors entirely and pixel shaders consuming the texture processors. However, despite an increasingly important onus being placed on the arrangement and capabilities of the shader Arithmetic Logic Units (ALU's) in this programmable era, the designs of contemporary graphics processors still bear the fundamental similarities to their forebears: vertex processing up one end of the pipeline, pixel processing down the other and still very much aligned with multiples of pixel pipelines.

Conceivably there is no reason why this development model couldn't continue to exist in the PC space and it certainly seems like it will from all vendors for at least the next year. However, ATI have multiple design teams working on different architectures concurrently, so whilst their PC processors may follow a fairly familiar lineage other parts of the company have been talking this shader era with a completely fresh perspective in order to consider the needs of a "Programmable Graphics Processor" and extract as much of the potential of the ALU's as possible by trying to minimise the

wasted cycles. In doing so they will force us to reconsider how we think of the overall pipeline and make initial performance assessments based upon "pipelines" alone.

Introduction

Ever since the announcement that ATI were working with Microsoft on "Future XBox technologies" the rumour mill has been working overtime as to the graphics behind it. Some of the messages since the announcement of the XBOX 360, the eventual console ATI's work will appear in, have not necessarily been reflective of the actual operation and even a little contradictory from representatives directly from ATI. With strict NDA's and designs being built for two different competitive consoles, very tight controls of what could be talked about had to be implemented within ATI, and the XBOX group operated very much within their own silo; it wasn't until Microsoft lifted the NDA's that ATI could even speak of it on a wider internal basis, let along externally, and even then there is a lot of information to gather.

Since XBOX 360's announcement and ATI's unleashing from the non disclosure agreements we've had the chance to not just chat with Robert Feldstein, VP of Engineering, but also Joe Cox, Director of Engineering overseeing the XBOX graphics design team, and two lead architects of the graphics processor, Clay Taylor and Mark Fowler. Here we hope to accurately impart a slightly deeper understanding of the XBOX 360 graphics processor, how it sits within the system, understand more about its operation as well as give some insights into the capabilities of the processor. Bear in mind that we are under NDA for some of the operational details of the graphics processor to gain an understanding of how it differs from current platforms however some of the specifics won't be revealed in full detail in this article.

Click for a bigger version

Throughout this article we'll attempt to piece together the operation of the graphics processor based on our conversations with ATI and some developers who have already had some knowledge of XBOX 360's capabilities, however we'll also offer some opinions on certain elements. Sections typed in blue indicate Beyond3D's suppositions and have not been directly indicated to us by ATI.

Page Navigation

Page 1 ▼ ▷

Site width adjust

Reset width

Privacy Policy

Latest Reviews

- NVIDIA Fermi GPU and Architecture Analysis
- ATI Cypress Gaming Performance Analysis
- ATI Cypress GPU and Architecture analysis
- ATI RV740 GPU and Architecture Analysis
- NVIDIA GT200 GPU and Architecture Analysis

Latest Articles

- E3 2011: Behind Closed Doors - Witcher 2 Xbox 360... and a Nurse
- A speculative look on the Wii U GPU
- Q&A with Visceral's Technical Art Director Doug Brooks on Dead Space 2
- Ian Buck - NVIDIA Tesla Launch
- AMD's John Bridgman on Radeon, Linux and Open Source 3D
- Andy Keane Interview & Tesla Adoption/Deployment
- Was Harry Potter actually into rendering? Beyond3D's first ever book review
- Dave Kirk - NVIDIA Tesla Launch
- Travelling in Style: Beyond3D's C++ AMP contest - WINNER ANNOUNCED
- Diving into Anti-Aliasing

ATI Xenos: Xbox 360 Graphics Demystified - Page 2

Published on 13th Jun 2005, written by Dave Baumann for Consumer Graphics - Last updated: 21st Mar 2007

Xbox 360 System Overview

The "XBOX 360" console was officially unveiled at a show on MTV the week prior to E3 2005, and at the unveiling Microsoft revealed a few technical details of the platform. The primary specifications for the system are:

- **3.2GHz Custom IBM Central Processor**
  - Three CPU Cores
  - Two Threads Per core
  - VMX Unit Per Core
  - 128 VMX Registers Per Thread
  - 1MB L2 Cache (Lockable by Graphics Processor)
- **500MHz Custom ATI Graphics Processor**
  - Unified Shader Core
  - 48 ALU's for Vertex or Pixel Shader processing
  - 16 Filtered & 16 Unfiltered Texture samples per clock
  - 10MB eDRAM Framebuffer
- **512MB System RAM**
  - Unified Memory Architecture (UMA)
  - 128-bit interface
  - 700MHz GDDR3 RAM

Of these core components obviously we are going to be most concerned with the graphics processing element. Whilst the graphics processor is different from others seen before in the PC space, and is very different from even ATI's impending new PC graphics components, it will be interesting to take a look at the graphics processor for the very reason that it doesn't directly correspond to any current graphics processor but also because we feel that this will give hints as to the architectural direction ATI are likely to be taking in the future for PC and other applications.

ATI C1 / Xenos

A name that has long since been mentioned in relation to the graphics behind Xenon (the development name for XBOX 360) is R500. Although this name has appeared from various sources, the actual development name ATI uses for Xenon's graphics is "C1", whilst the more "PR friendly" codename that has surfaced is "Xenos". ATI are probably fairly keen not to use the R500 name as this draws parallels with their upcoming series of PC graphics processors starting with R520, however R520 and Xenos are very

distinct parts. R520's aim is obviously designed to meet the needs of the PC space and have Shader Model 3.0 capabilities as this is currently the highest DirectX API specification available on the PC, and as such these new parts still have their lineage derived from the R300 core, with discrete Vertex and Pixel Shaders; Xenos, on the other hand, is a custom design specifically built to address the needs and unique characteristics of the game console. ATI had a clean slate with which to design on and no specified API to target. These factors have led to the Unified Shader design, something which ATI have prototyped and tested prior to its eventual implementation (*with the rumoured R400 development ?*) , with capabilities that don't fall within any corresponding API specification. Whilst ostensibly Xenos has been hailed as a Shader Model 3.0 part, its capabilities don't fall directly inline with it and exceed it in some areas giving this more than a whiff of WGF2.0 (Windows Graphics Foundation 2.0 - the new name for DirectX Next / DirectX 10) about it.

Click for a bigger version

Click for a bigger version

The Xenos graphics processor is not a single element, but actually consists of two distinct elements: the graphics core (shader core) and the eDRAM module. The shader core is a 90nm chip manufactured by TSMC and is currently slated to run at 500MHz*, whilst the eDRAM module is another 90nm chip, manufactured by NEC and runs at 500MHz* as well. These two chips both exist side by side, together on a single package, ensuring a fast interlink between the two. The main graphics chip, the parent core, could be considered as a "shader core" as this is one of its primary tasks. The eDRAM module is a separate, daughter chip which contains the elements for reading and writing color, z and stencil and performing all of the alpha blending and z and stencil ops, including the FSAA logic. We'll explore the capabilities and operations of both these chips in greater detail throughout the article.

(*) Note: We understand the clockspeeds for the shader core and daughter die are target clockspeeds at present and there may be some room for small movement either way on both dies dependant on yields. As Microsoft have now announced 500MHz speeds it is more likely that these will be the eventual release speeds.

One element that has been reported on is the number of 150M transistors in relation to the graphics processing elements of Xenon, however according to ATI this is not correct as the shader core itself is comprised from in the order of 232M transistors. It may be that the 150M transistor figure pertains only to the eDRAM module as with 10MB of DRAM, requiring one transistor per bit, 80M transistors will be dedicated to just the memory; when we add the memory control logic, Render Output Controllers (ROP's) and FSAA logic on top of that it may be conceivable to see an extra 70M transistors of logic in the eDRAM module.

Click for a bigger version

Update: *We've recently been given an image of the Xenos graphics chip package (above) that highlights the dual die nature, with the parent die quite clearly to the centre of the package and the daughter over to the left. While the 232M transistor figure for the parent was given to us by ATI we are still trying to establish a more official figure for the daughter (even though these types of transistor counts are very much estimates anyway). We've speculated that the 150M figure that appeared when XBOX 360 was first announced may just relate to daughter die, however another figure that has arisen is 100M - judging from the die sizes the daughter die doesn't have more than half the area of the parent, which would giv indications towards the 100M side although 80M of those transistors are DRAM which may be more dense than the logic circuitry that will dominate the parent die. We are trying to get further clarification.*

One of the mistakes that Microsoft made with the original XBox was to contract their component providers into supplying entire chips with, evidently, no development path - at least, this was the case with NVIDIA NV2A graphics processor, which resulted in Microsoft and NVIDIA going through a legal arbitration process. Although the components in the XBOX 360 in its initial form are hardly low cost, the cost of the unit over the course of its lifetime is one that has quite obviously been addressed with contracts that pay via royalties for chips sold and with Microsoft in charge of ordering the chips from the various Fabs, however the original semiconductor manufacturers are likely to still be in charge of further developments in terms of putting the cores on to smaller processes and we believe that this is part of the contract that ATI has with Microsoft. An obvious area for cost reduction of the Xenos processor is by merging the shader and daughter die on to a single core - we suspect that this will not happen until there is a process shrink available (that can also cater for both the complex logic and eDRAM) as two cores on 90nm mitigate some of the yield risks of a single, large die on 90nm.

Page Navigation

◁ Page 2 ▽ ▷

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS
Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS
Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS
Sync your system to PACER to automate legal marketing.