



Architecting the Future

Michael Doggett

Architect

August 2, 2005

Overview

- Xenos – XBOX360 GPU
 - Unified shader
- GPU realities
- Graphics APIs
- GPU research
- Beyond the horizon



ATI - Driving the Visual Experience Everywhere

- Products from cell phones to super computers



Gaming



Gaming Console



Integrated



Embedded Display




Notebook



Digital TV



Multimedia



Workstation



Color Phone Display



Multi Monitor Display




Xenos

- Unified Shader
 - One hardware design that performs both Vertex and Pixel shaders
 - A revolutionary step in Graphics Hardware
 - Generalization of vertex and pixel shaders
 - Vertex processing power
 - GPU based vertex and pixel load balancing

Vertex Shader

Pixel Shader



Unified Shader

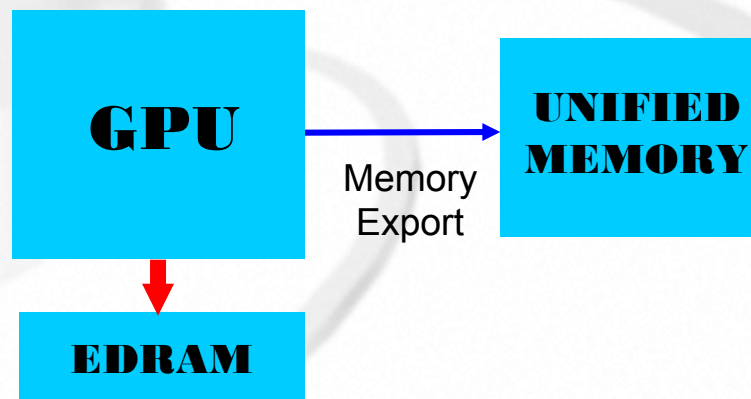
4comp+1scalar/shader

16 shaders/SIMD

3 SIMD

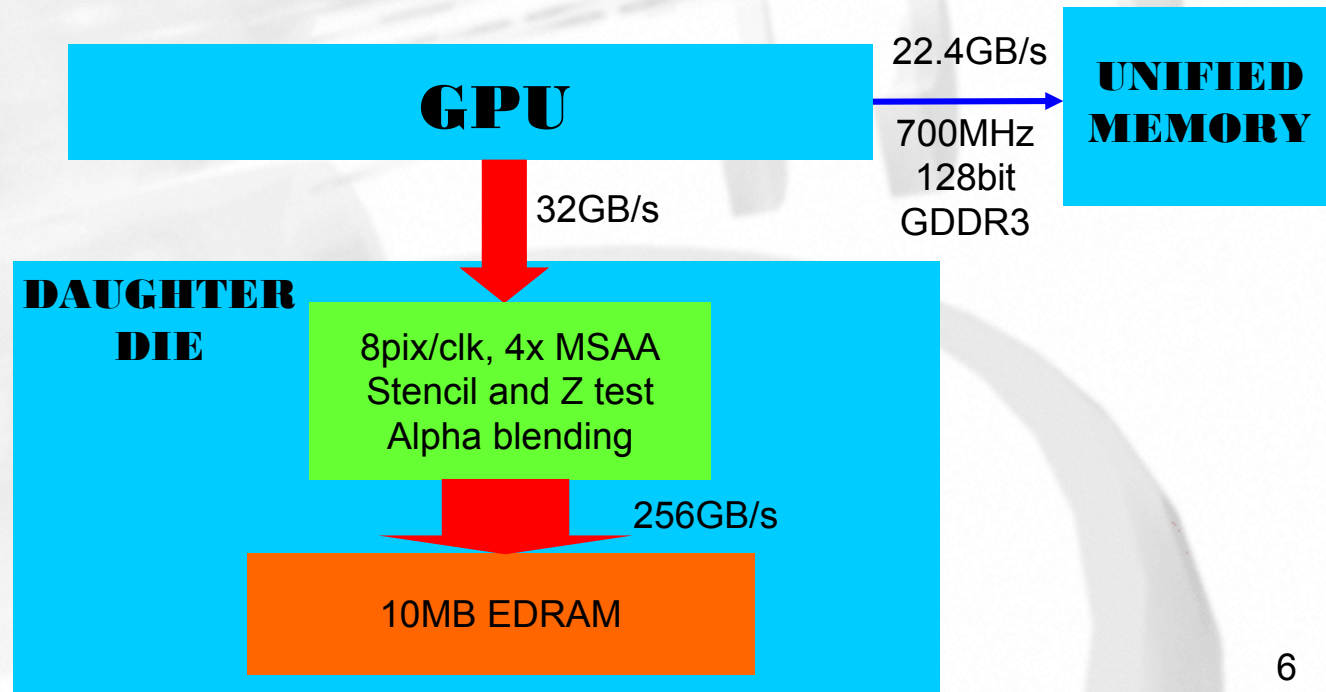
Xenos

- Memory Export
 - Shader output to a computed address
 - Scatter write
 - Randomly update data structures from Vertex or Pixel Shader
 - Ray tracing acceleration structures
 - Physical simulation – GPGPU
 - Virtualize shader resources
 - Enabling exploration for the future

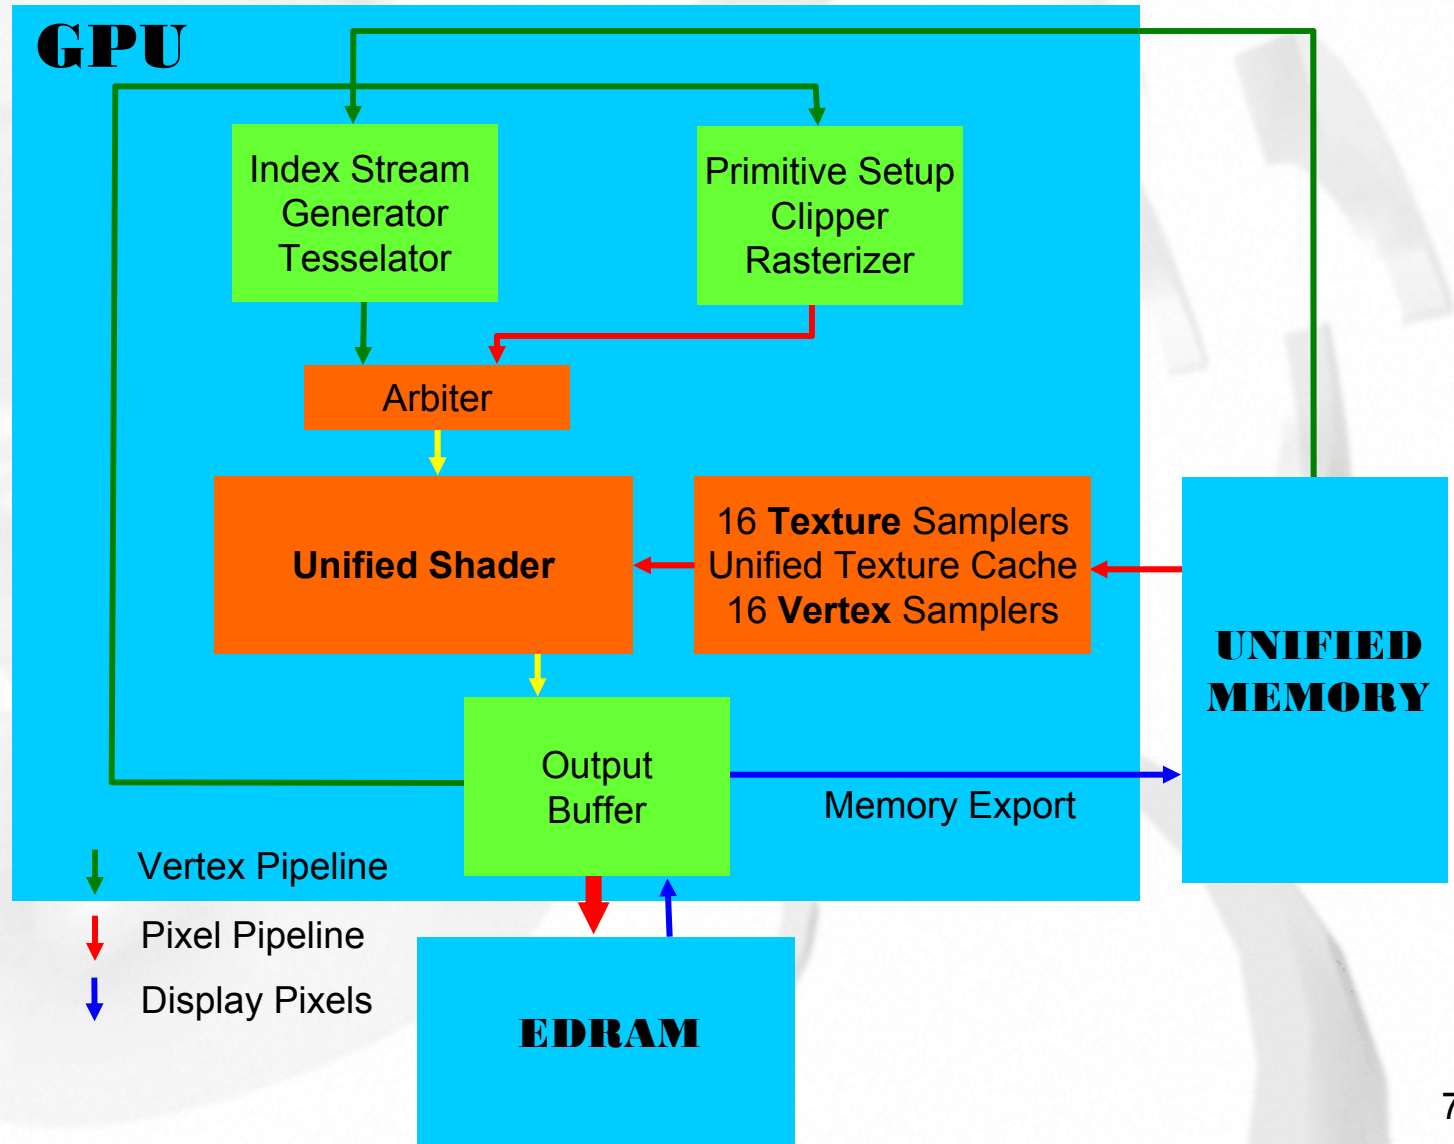


Xenos

- Daughter Die
 - 8 pix/clock @ 4xMSAA bandwidth Daughter Die
 - Separate Framebuffer and texture/main memory interface



Xenos



GPU realities

- Reaching the limit of GPU die sizes
 - Lower yields/Higher cost
 - One die, multiple products
- Decreasing technology, increasing power
- How to cool large chips ?
- Ensure design is scalable

GPU realities

- Need refinements and enhancements of the current pipeline
 - Unified shader
 - Greater flexibility of existing features
 - Replace fixed function with programmable
 - Fully programmable configured to do graphics rendering
 - RAW, Chen GH05
- Suggested small modifications get adopted

Graphics APIs

- Windows Vista
 - Virtual Memory
 - Improved state change efficiency
- Windows Graphics Foundation 1.0
 - ClearType
- Windows Graphics Foundation 2.0
 - Unified shader programming model
 - Geometry Shader
 - Access to entire triangle and adjacent vertices
 - Output to array of render targets, cube maps
 - Stream output from Geometry Shader
- OpenGL extensions
 - BOF Wed, Aug 3, 6-8pm, Wilshire Grand

GPU research

- Improved multi-GPU performance and antialiasing
 - CrossFire
 - Multi-chip, multi-core
- Improved texture compression
 - 3Dc normal compression

GPU research

- Higher Order Surfaces
 - Subdivision surfaces, NURBS
 - *Geometry on GPUs*, Thursday 8.30
- Order Independent Transparency
 - Depth peeling
 - Multi-pass, existing hardware
 - Use Multisample sample mask to control transparency
 - Alpha to coverage
 - Foliage, Chicken wire
 - Single pass hardware

GPU research

- Shadows
 - Performance enhancements for stencil shadows and shadow buffers
 - In the future
 - Global illumination – Ray tracing
- Hierarchical rendering
 - User low resolution shadow map to find areas that require detail rasterization
 - Chan, Durand EGSR04
 - Stencil shadow volumes using 8x8 pixel tiles
 - Aila, Akenine-Möller GH04
 - Improving Z, stencil
- Physical simulation on the GPU
 - General Purpose GPU (GPGPU)

Beyond the horizon

- Ray tracing on GPUs
 - Uniform grid
 - Purcell et. al. SIG02
 - KD-Tree, performance improvement greater than hardware improvement
 - Foley et. al. GH05
 - EarlyZ Volume rendering ray casting
 - J. Krüger, R. Westermann VIS03
 - Build and update acceleration structures
 - Xenos shader scatter write

Beyond the horizon

- Ray tracing hardware
 - FPGA based ray tracing design
 - RPU, Woop et. al. SIG05
 - Evolution towards ray tracing GPUs

Conclusion

- Trend towards unified shaders
- Consideration for physical limitations
 - Yield
 - Power
 - Heat
- Performance from more shaders, multi-chip, multi-core
- Features for increased visual realism and APIs

ATI Booth 1101



- ATI FireGL products
- XBOX360 – Ruby: The Assassin