

CS4803DGC Design Game Consoles

Spring 2010

Prof. Hyesoon Kim



Real-time Rendering Ch 18


**Georgia
Tech**



College of
Computing




Graphics Pipeline Stages





Graphics Processor Trend

- Fixed pipeline processor
 - More and more CISC style processors
 - OpenGL API \sim One instruction
- Programmable parallel processors
 - Scatter/Gather operations
 - Scatter: $A_{out}[B[i]] = A_{in}[i]$
 - Gather: $A_{out}[i] = A_{in}[B[i]]$
 - Enable GPGPU

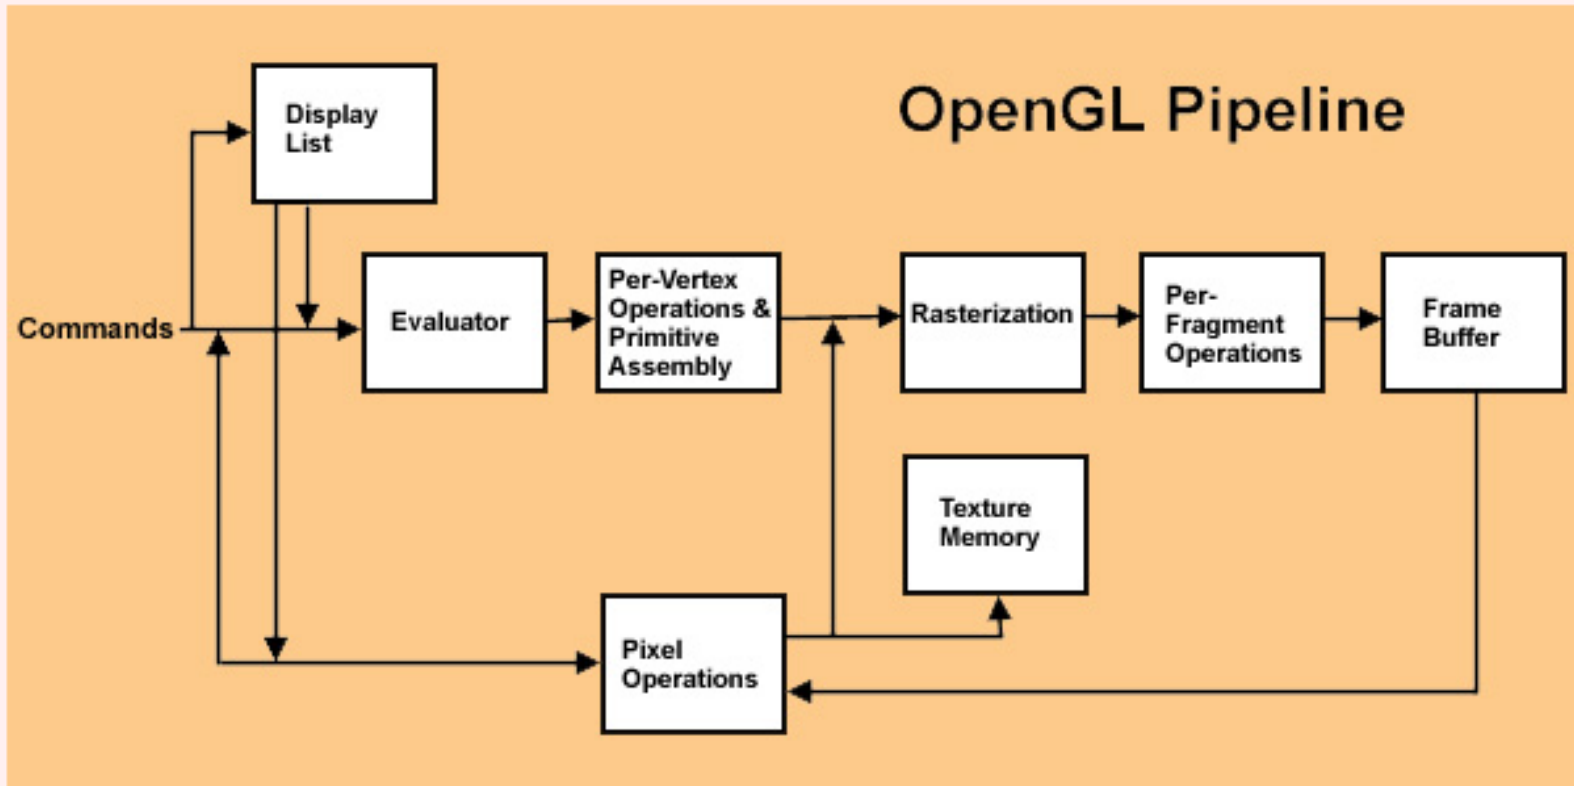




Fixed Function vs. Programmable Function


- Fixed function stage:
 - Fixed vertex stage: Clip-space vertex computations, per-vertex normal, many other per-vertex operations (color material, texture coordinate, normal transformation)
 - Fixed fragment stage: many basic lighting models and effects etc.
- Programmable function pipeline
 - Custom computations

OpenGL pipeline(1997)



OpenGL 3


OpenGL 3 Modern Buffer-centric Processing Model



Fixed Function Scheme

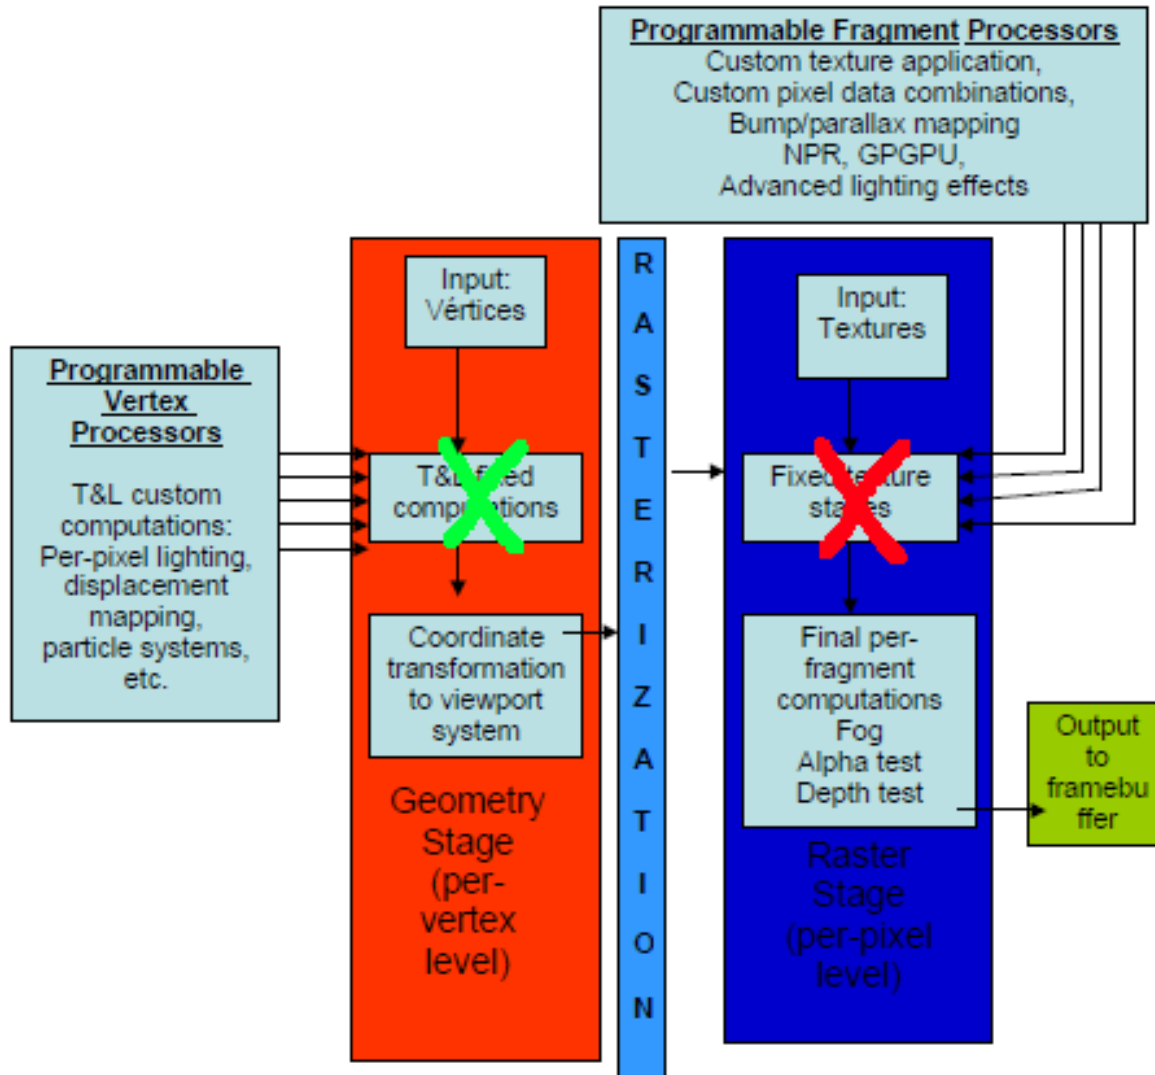


Fixed Function Scheme





Programmable Function Scheme



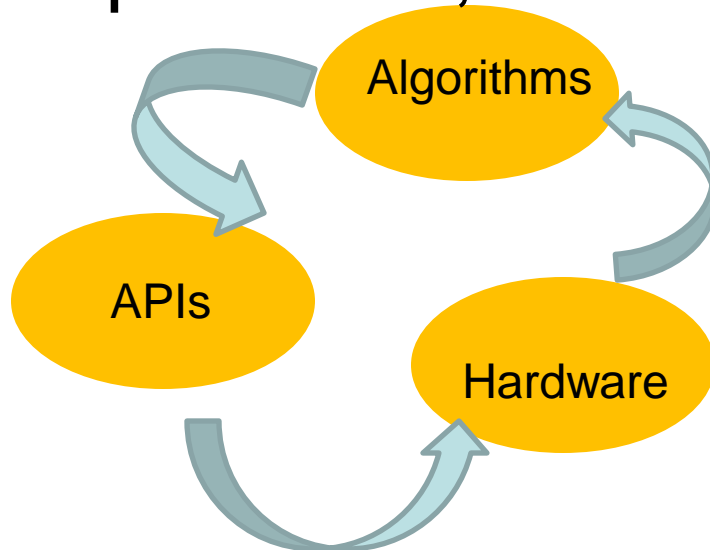


Hardware Design Issues in Fixed Pipeline

- Start from straightforward mapping
- Efficient Instruction mapping
 - Better compiler is good
- Register Bank issues
- Texture cache access
 - Prefetching
- Memory bandwidth
- Better optimizations to reduce the amount of computations,

Hardware Design Issues in Programmable GPUs

- Efficient instruction packing
 - SIMD issue width
- Memory bandwidth saving
- Better optimizations to reduce the amount of computations, increasing parallelism



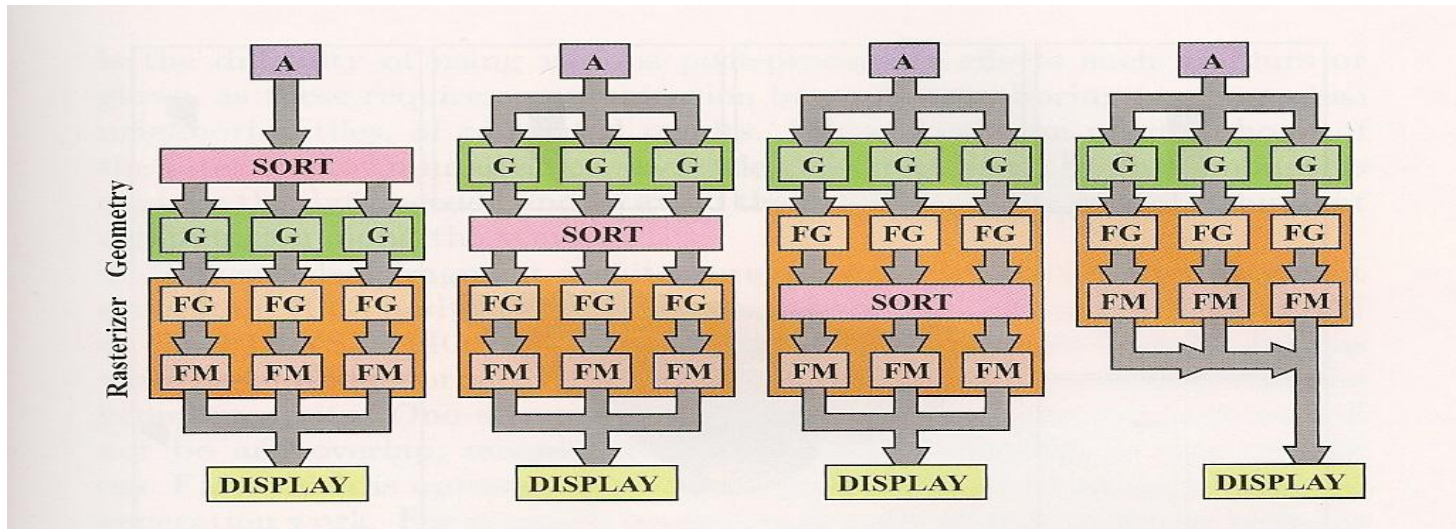


CPU vs GPU

- GPU
 - Many pipeline stages (400-600 stages)
 - Massive data parallelism
 - Memory access patterns are regular

Taxonomy of Parallel Graphics Architectures

- Sort-first, sort-middle, sort-last fragment, sort-last image
- FG: Fragment generation
 - The actual locations inside a primitive
- FM: Fragment merge
 - Merge the results using Z-buffer





The RASTERIZER Z-buffering

- The Z-buffer (aka depth buffer)
- Idea:
 - Store z (depth) at each pixel
 - When scan-converting a triangle, compute z at each pixel on triangle
 - Compare triangle's z to Z-buffer z -value
 - If triangle's z is smaller, then replace Z-buffer and color buffer
 - Else do nothing
- Can render in any order

Z-value




Z-value: distance from the viewer

Z- value




Viewer





Alpha Blending

- *Alpha Blending* is used to render translucent objects.
- The pixel's alpha component contains its *opacity*.
- Read-modify-write operation to the color framebuffer
- $\text{Result} = \text{alpha} * \text{Src} + (1 - \text{alpha}) * \text{Dst}$



Memory Architecture

- Xbox: unified memory architecture
- Xbox 360: hybrid
 - CPU and GPU share the same bus and interface to the system memory (texture memory)
 - GPU-exclusive memory
- Playstation3: GPU dedicated memory


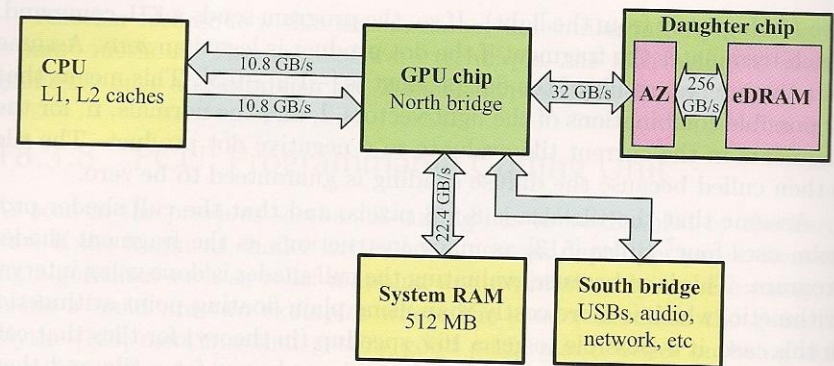


Figure 18.13. The memory architecture of the first Xbox, which is an example of a unified memory architecture (UMA).





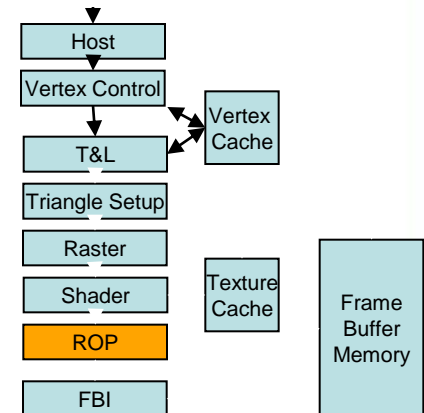
Announcement

- Lab #3 due on Friday
- Friday: Prof. Hsien-Hsin Lee will give a lecture about more detailed graphics techniques



ROP (from Raster Operations)

- C-ROP performs frame buffer blending
 - Combinations of colors and transparency
 - Antialiasing
 - Read/Modify/Write the Color Buffer
- Z-ROP performs the Z operations
 - Determine the visible pixels
 - Discard the occluded pixels
 - Read/Modify/Write the Z-Buffer
- ROP on GeForce also performs
 - “Coalescing” of transactions
 - Z-Buffer compression/decompression





Stencil Buffer

- Per pixel operation
- Compares reference value to pixel's stencil buffer value
- Same spatial resolution as color and depth buffers
- Usually 8-bits'
- Used to hold vales related to elements being written into frame buffer

Stencil Maintains the Floor



Clear stencil to zero.

Draw floor polygon with stencil set to one.


Only draw reflection where stencil is one.



CASE STUDY: XBOX 360

Xbox 360 Memory Architecture

- embedded DRAM(eDRAM): frame buffers,
- 10MB
- Daughter chip: AZ: all alpha and Depth testing.



Xbox 360 System Block Diagram





Figure 2. Xbox 360 system block diagram.

XBOX 360 GPU Block Diagram






Figure 5. Xbox 360 CPU die photo (courtesy of IBM).

GPU Parent Die




Figure 7. Xbox 360 GPU "parent" die (courtesy of Taiwan Semiconductor Manufacturing Co.).

EDRAM


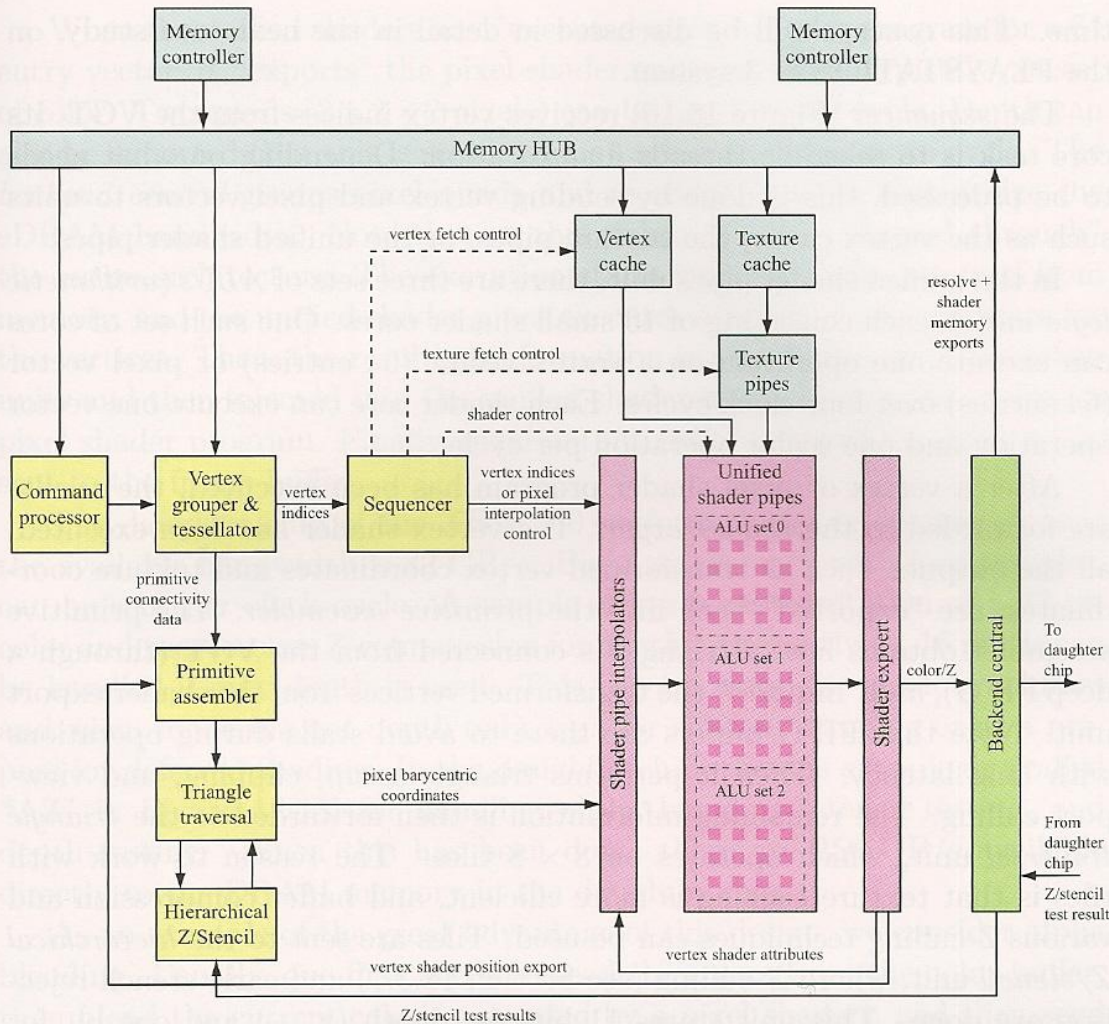


Figure 8. Xbox 360 GPU EDRAM ("daughter") die (courtesy of NEC Electronics).

Xbox 360 Graphics Processors



Unified shader
Command processor:
reads commands from
memory

64 vertices or pixels are
operated together (SIMD)

32 vertex threads or 64
pixel threads can be active

24,576 registers

ALU: 16 small shader cores
32Kb texture cache
Texture pipes: 16 bilinear
filters per cycles


Figure 18.16. Block diagram of the Xbox 360 graphics processor.



- Command processors: reads commands from memory
- VGT (Vertex grouper & tessellator):
 - Receives a group of vertex indices
- Sequencer: schedule threads
- Shader export: (we need to use the shader again!) FIFO for buffer
- Shader pipe interpolators
- Backend central → send data to daughter chip

Xbox 360 Daughter Chip

- Performs merge operation
- GPU-AZ : bandwidth 32GB/s
- 8 pixels * 4 samples can be sent per clock
 - Sample: 32bit color + lossless z-compression for depth
- 16 pixels if only depth test
- AZ logic: alpha blending, stencil testing, depth testing
- AZ-eDRAM: 256 GB/s



Playstation 3 GPU: The RSX


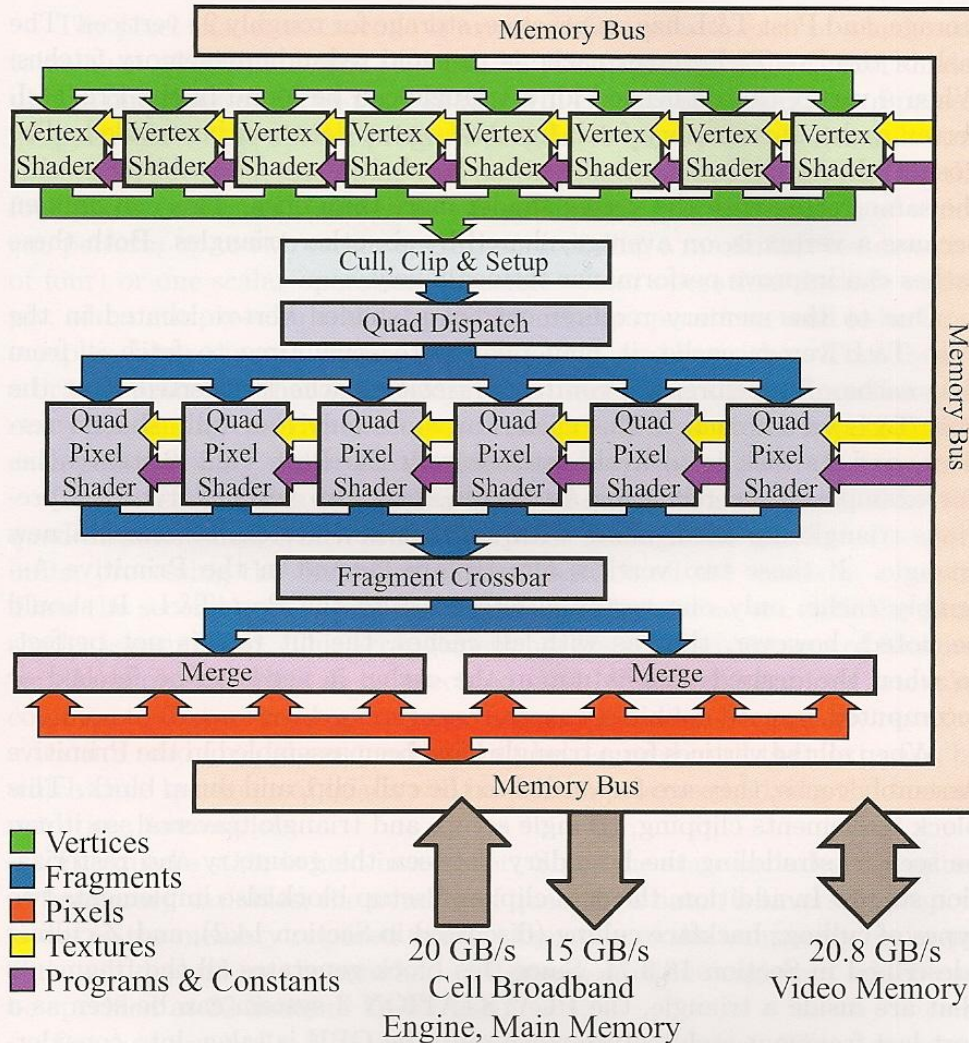


Figure 18.17. PLAYSTATION 3 architecture. (Illustration after Perthuis [1002].)

Video memory bandwidth is lower:

Playstation3 GPU Architecture



Modified version of
GeForce 7800

8 vertex shader

1 Pre-T&L vertex cache
3 Post- T&L vertex caches

Figure 18.18. RSX architecture.

Headsup: Playstaion3 GPU: Cell

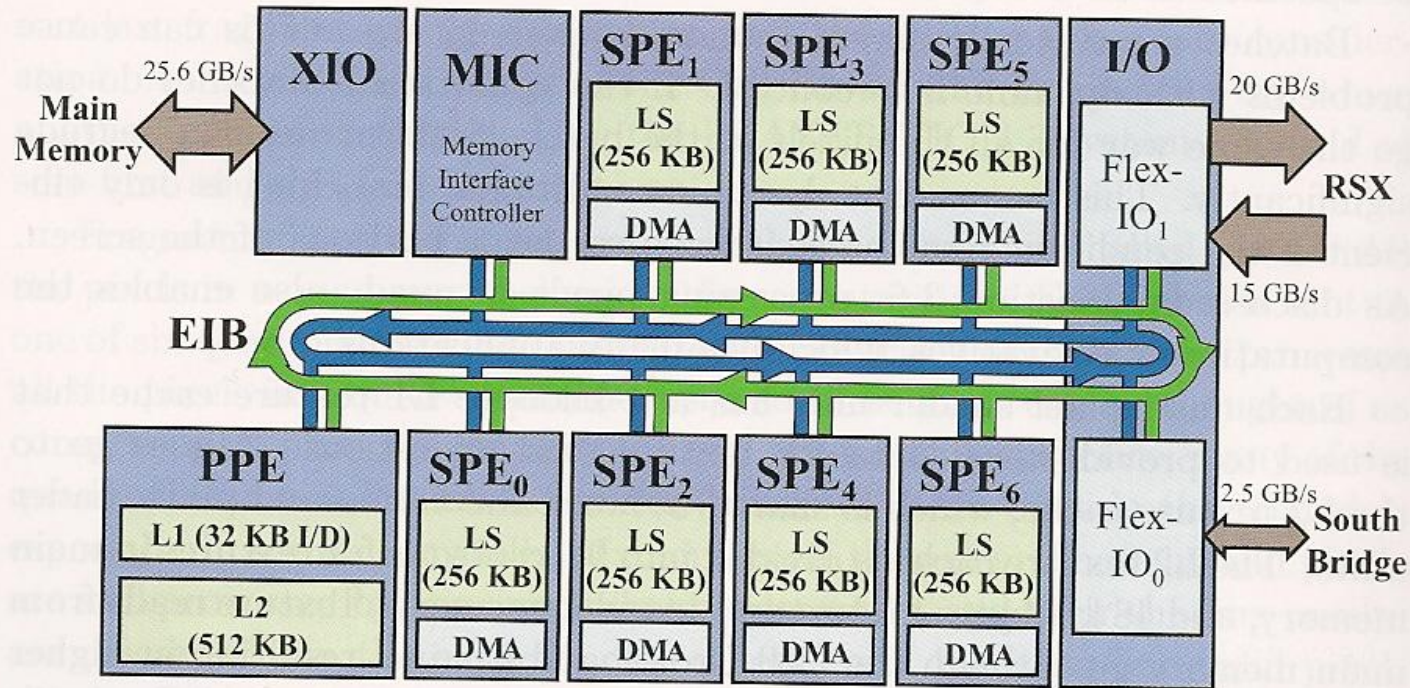


Figure 18.19. Cell Broadband Engine architecture. (Illustration after Perthuis [1002].)



Wii's GPU: Hollywood

- Designed by AMD's ATI Technologies
- Almost no public information
 - (no clock frequency, pixel pipelines or shader units)
 - **The 'Hollywood' is a large-scale integrated chip that includes the GPU, DSP, I/O bridge and 3MBs of texture memory," a studio source told us. From Rage3D**
 - Will: Generally more I/O support than peak CPU or GPU performance
- Fixed graphics function.

