# ECE4893A/CS4803MPG: Multicore and GPU Programming for Video Games

## Architectural Comparison: Xbox 360 vs. Playstation 3


Prof. Aaron Lanterman

School of Electrical and Computer Engineering

Georgia Institute of Technology

Georgia Institute of Technology

# Xbox 360's Xenon vs. Playstation 3's Cell

Both chips clocked at a 3.2 GHz

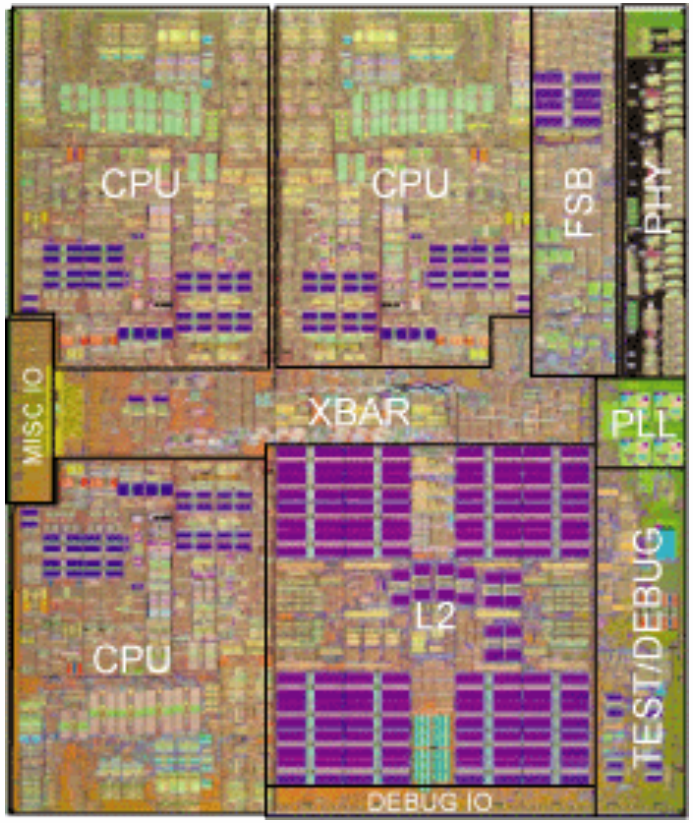


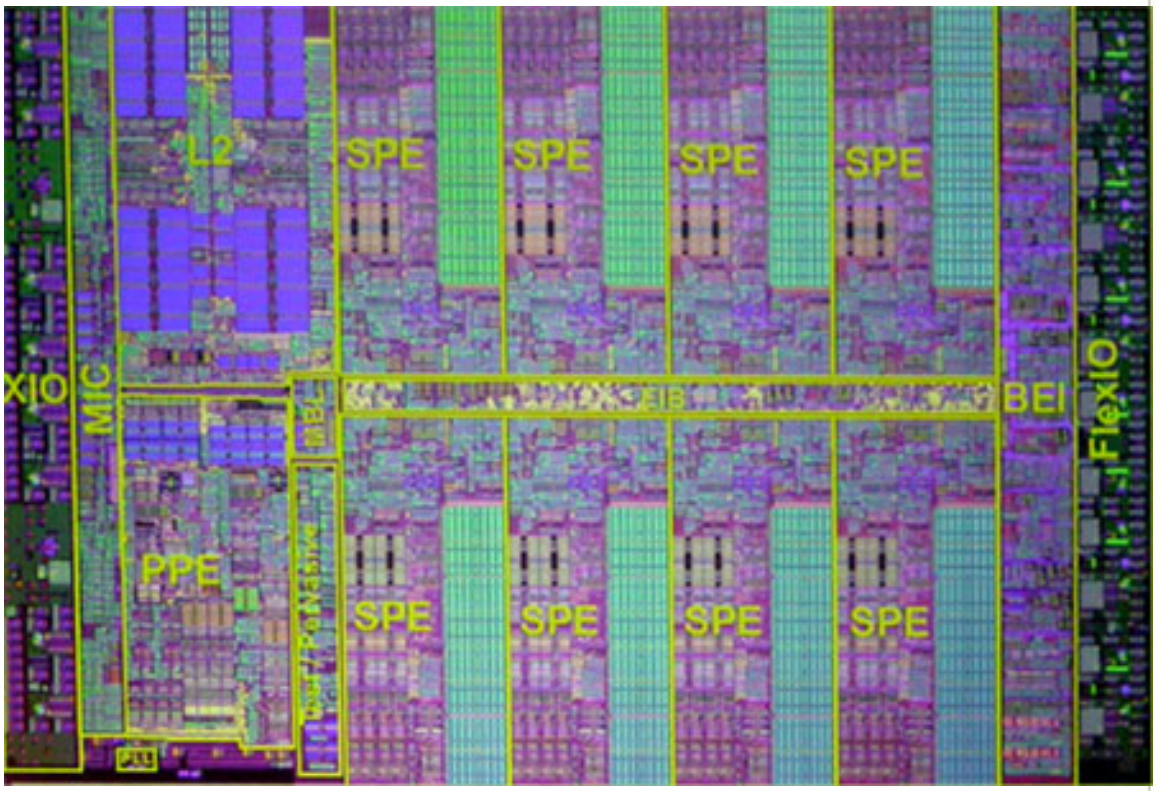


**Images not to scale**

Xenon CPU image from "The Microsoft Xbox 360 CPU story"
www-128.ibm.com/developerworks/power/library/pa-fpfxbox

Cell processor image from "IBM's Cell Processor: Preview to Greatness?"
www.pcstats.com/articleview.cfm?articleid=1727

Georgia Institute of Technology

# Memory: Xbox 360 vs. Playstation 3

- Xbox 360 - 512 MB, 700 MHz, GDDR3, shared by CPU and GPU
- CPU accesses memory through the GPU!
- GPU has 10 MB RAM embedded frame buffer

- PS3 - 512 MB total
- 256 MB 3.2 GHz XDR main RAM for the CPU
- 256 MB 700 MHz GDDR3 video RAM for the GPU

# Xbox 360 high-level architecture

Custom-designed XMA (Xbox Media Audio) decoder for on-the-fly decoding of compressed audio streams

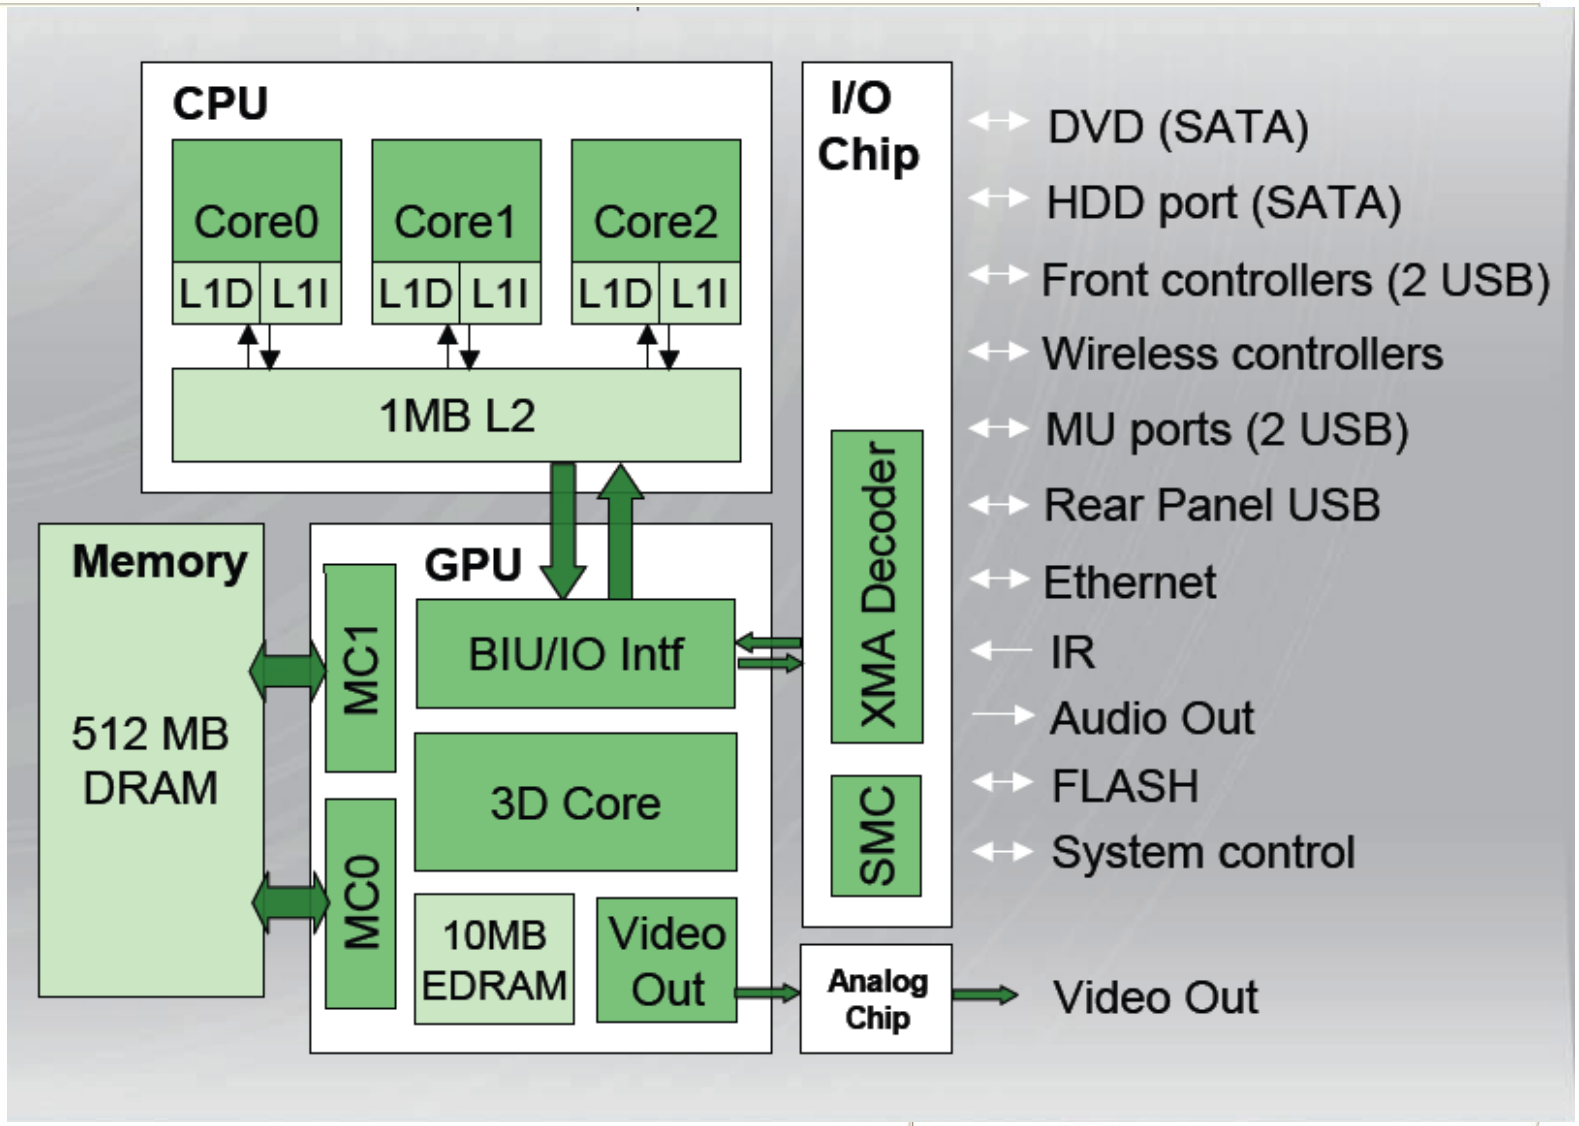


Image from J. Andrews and N. Baker, "Xbox 360 System Architecture," *Hot Chips* Presentation

# Xenon architecture

Front Side Bus runs at 10.8 Gbit/sec read/ write

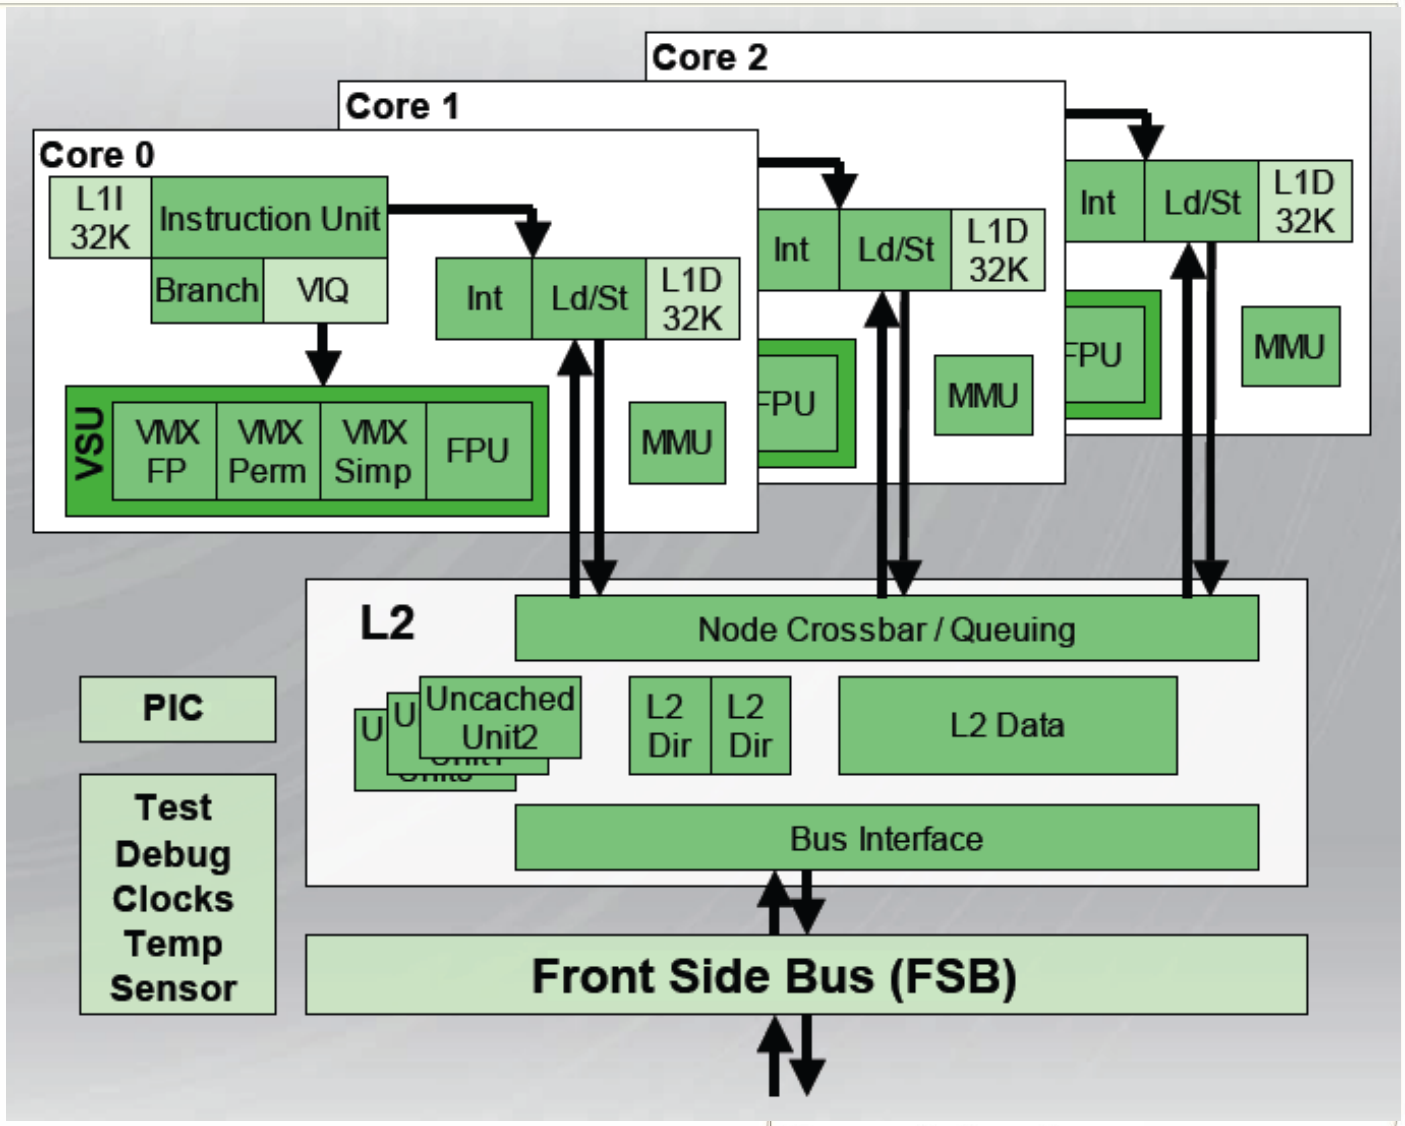


Image from J. Andrews and N. Baker, "Xbox 360 System Architecture," *Hot Chips* Presentation

# Cell BE architecture

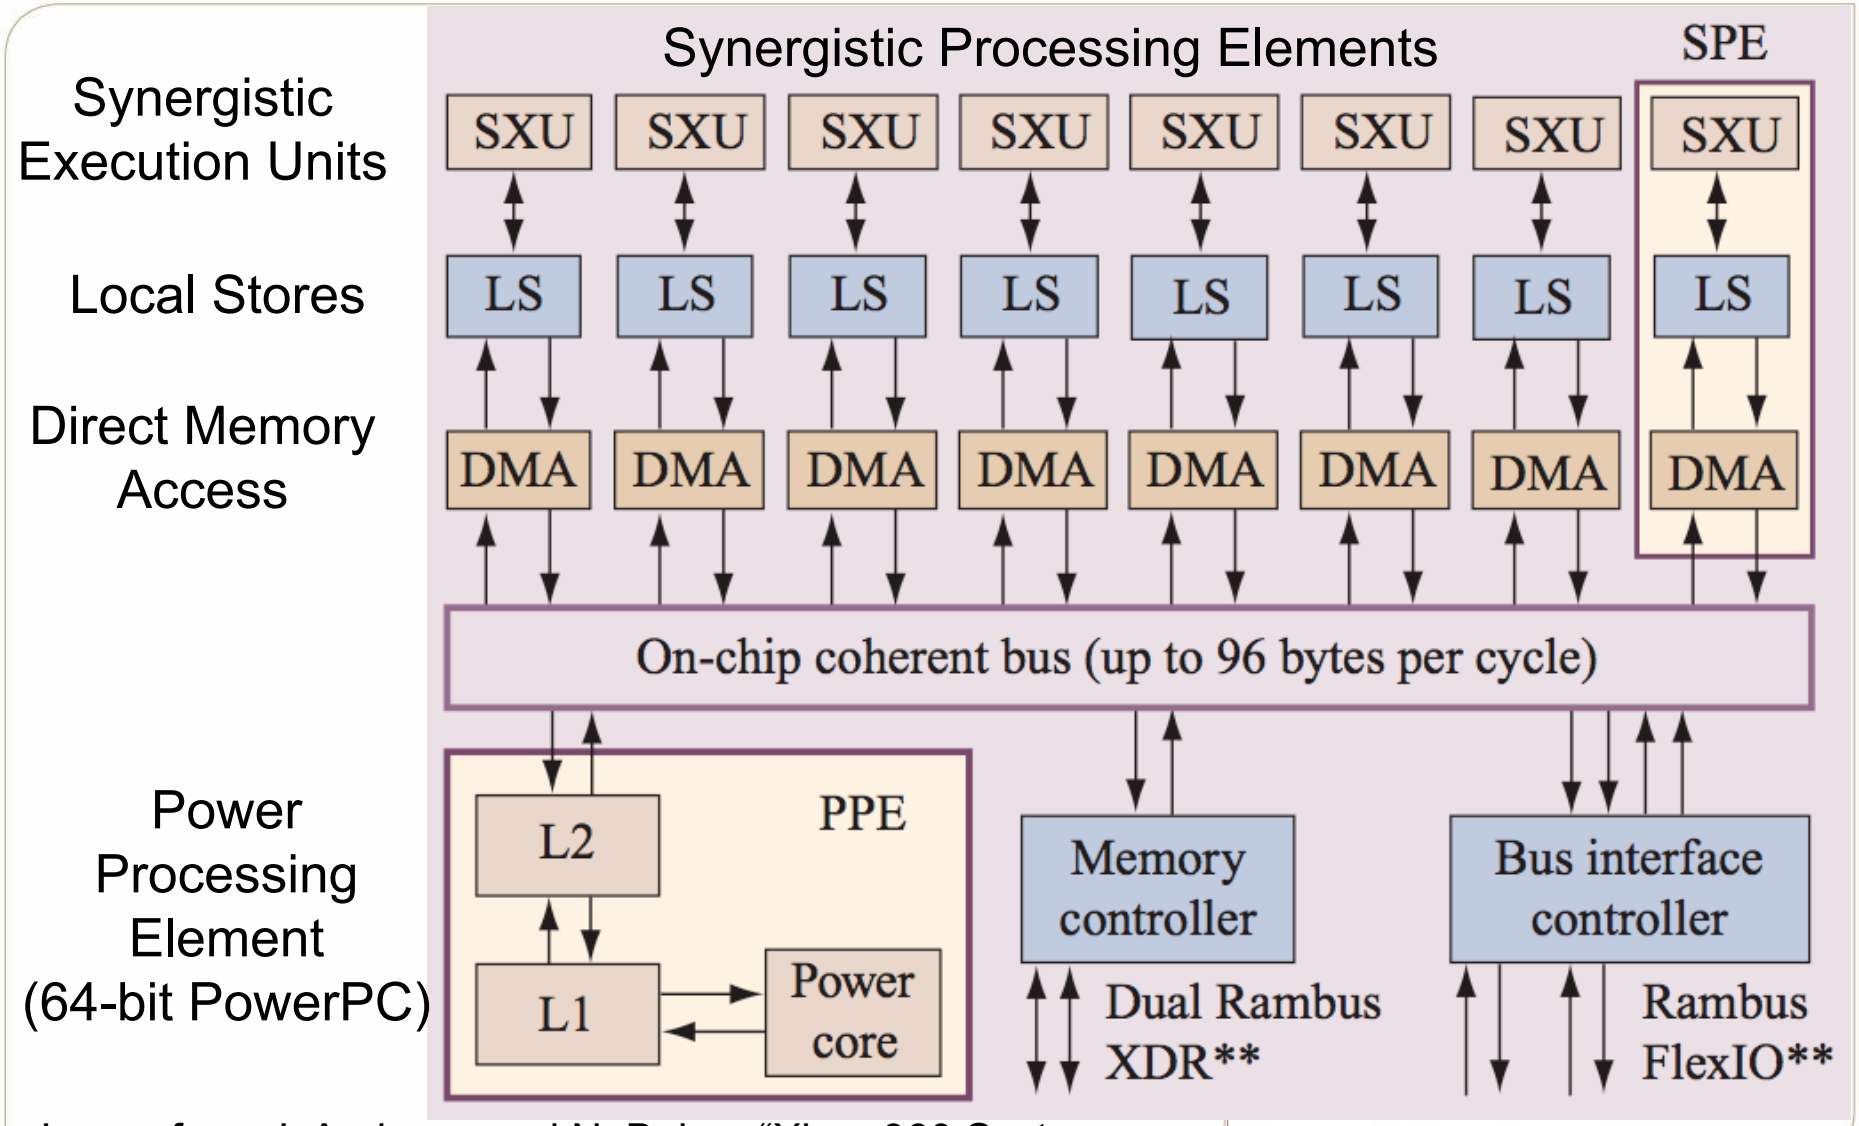


Image from J. Andrews and N. Baker, "Xbox 360 System Architecture," *Hot Chips* Presentation

# What the PowerPC cores have in common

PPE on Cell, and each
core on the Xbox 360 have:

- 64-bit PowerPC architecture
- Two symmetric multithreading (SMT), fine-grained hardware threads (6 total in Xbox 360)
- Integer arithmetic, single and double precision floating point, single cycle for most instructions
- VMX128 “Altivec” vector processor

Information from Andrews & Baker and Kahle *et al.*

# VMX128 “Altivec” vector processor

- 128, 128-bit registers (4-element single-precision) per hardware thread
  - 6 total Altivec-style register files in Xbox 360
- Floating point arithmetic, dot product, permute
- On Xbox 360, CPU can convert 3D data to Direct3D compressed data formats before storing in L2 cache or main memory
  - Typically 50% in bandwidth and memory footprint

Information from Andrews & Baker and Kahle *et al.*

# Set-associative caches

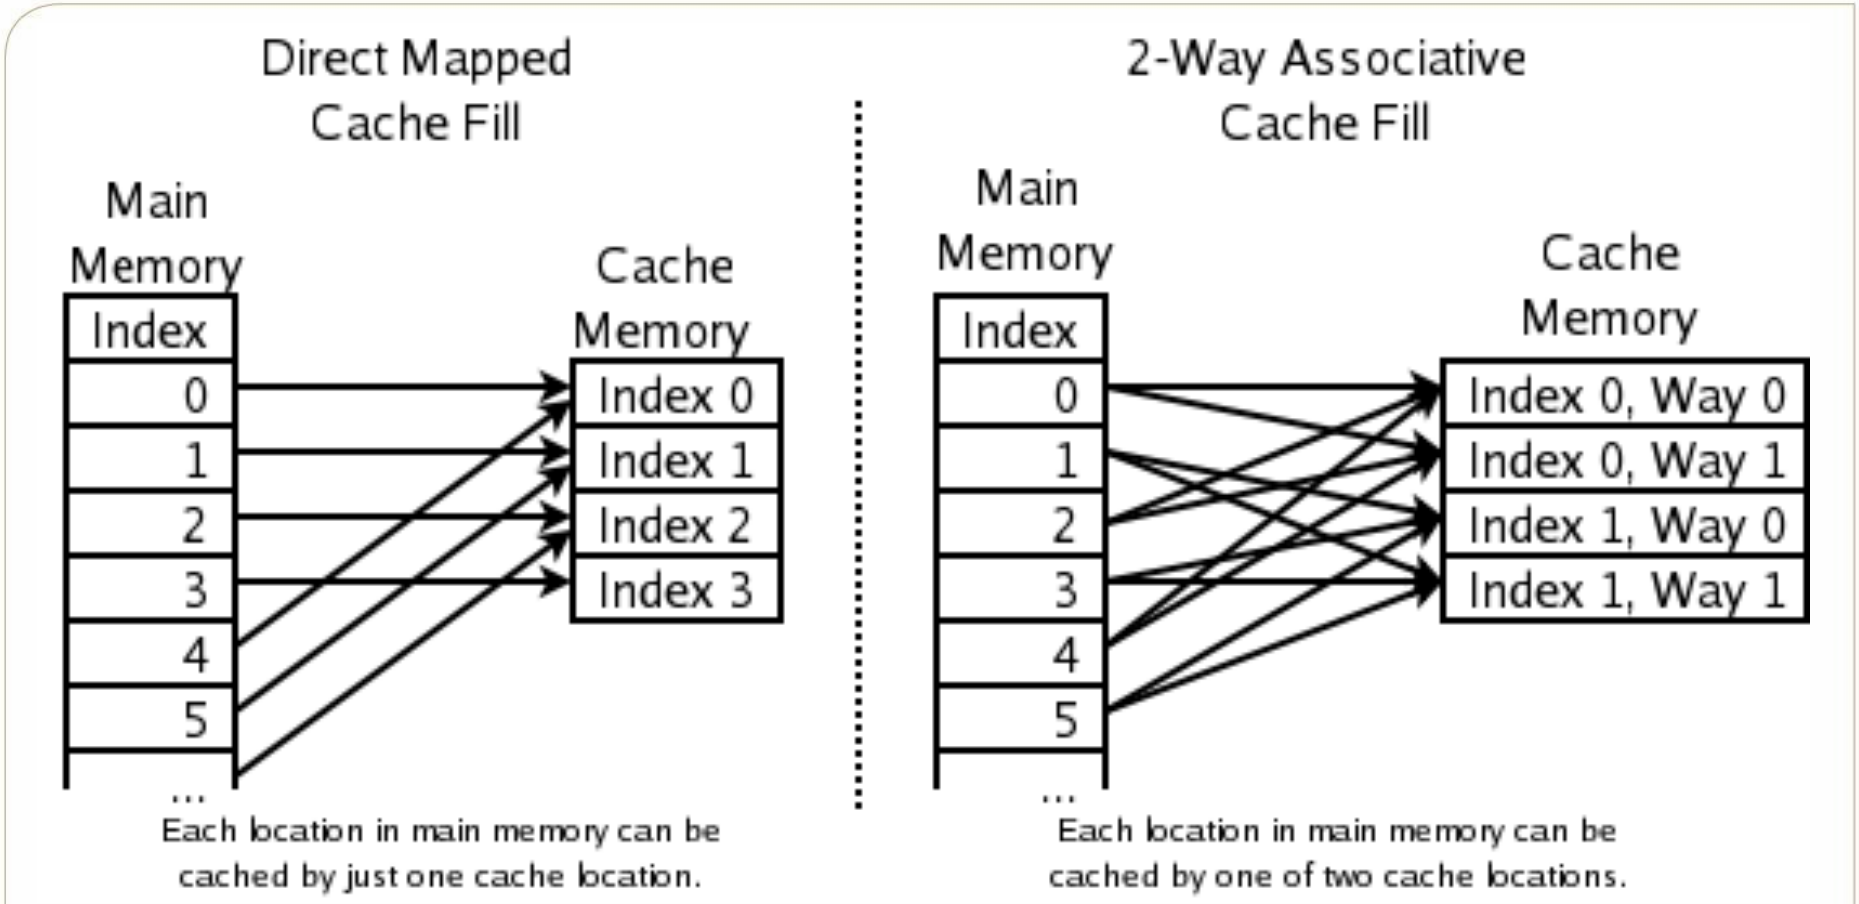


From Wikipedia entry on “cache algorithms”

# Caches

- *Each* PowerPC core on Xenon has:
  - 32 KB L1 two-way set-associative instruction cache
  - 32 KB L1 four-way set-associative, write-through data cache
  - xDCBT “extended data cache block touch” instruction for prefeching data direct into L1 cache, but not L2 cache as usual; avoids thrashing L2 cache
- PowerPC core on Cell has:
  - 32 KB L1 instruction cache
  - 32 KB L1 data cache
  - 512 KB L2 cache

Information from Andrews & Baker and Kahle *et al.*

# Xenon's L2 cache

- All three PowerPC cores share a 1 Megabyte, 8-way set-associative L2 cache
- Cache set locking: "common in embedded systems, but not PCs"
- Lets the cores dynamically allocate L2 usage
- Facilitates communication between cores
- GPU can read directly from the L2 cache

Information from J. Andrews and N. Baker, "Xbox 360 System Architecture," *IEEE Micro*, March-April 2006, pp. 25-37.

# Xenon's L2 cache architecture

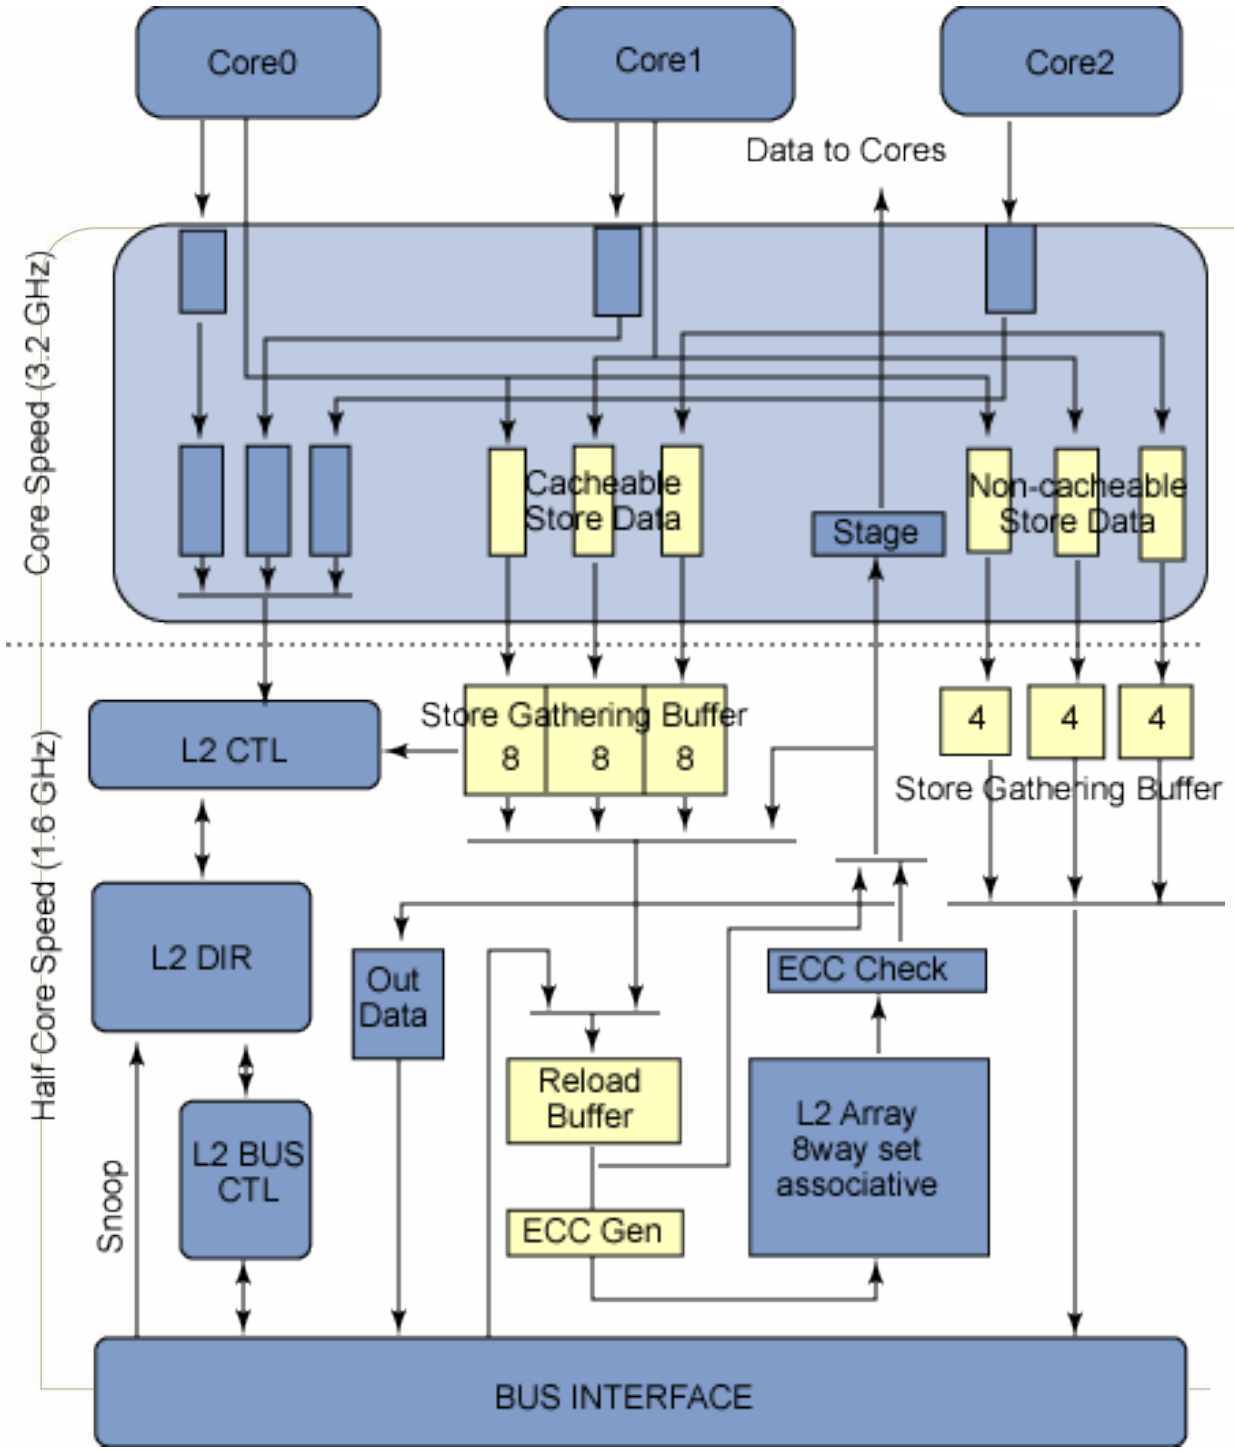


Image from J. Brown, "The Microsoft Xbox 360 CPU story" www-128.ibm.com/developerworks/power/library/pa-fpfxbox

# Xenon core architecture

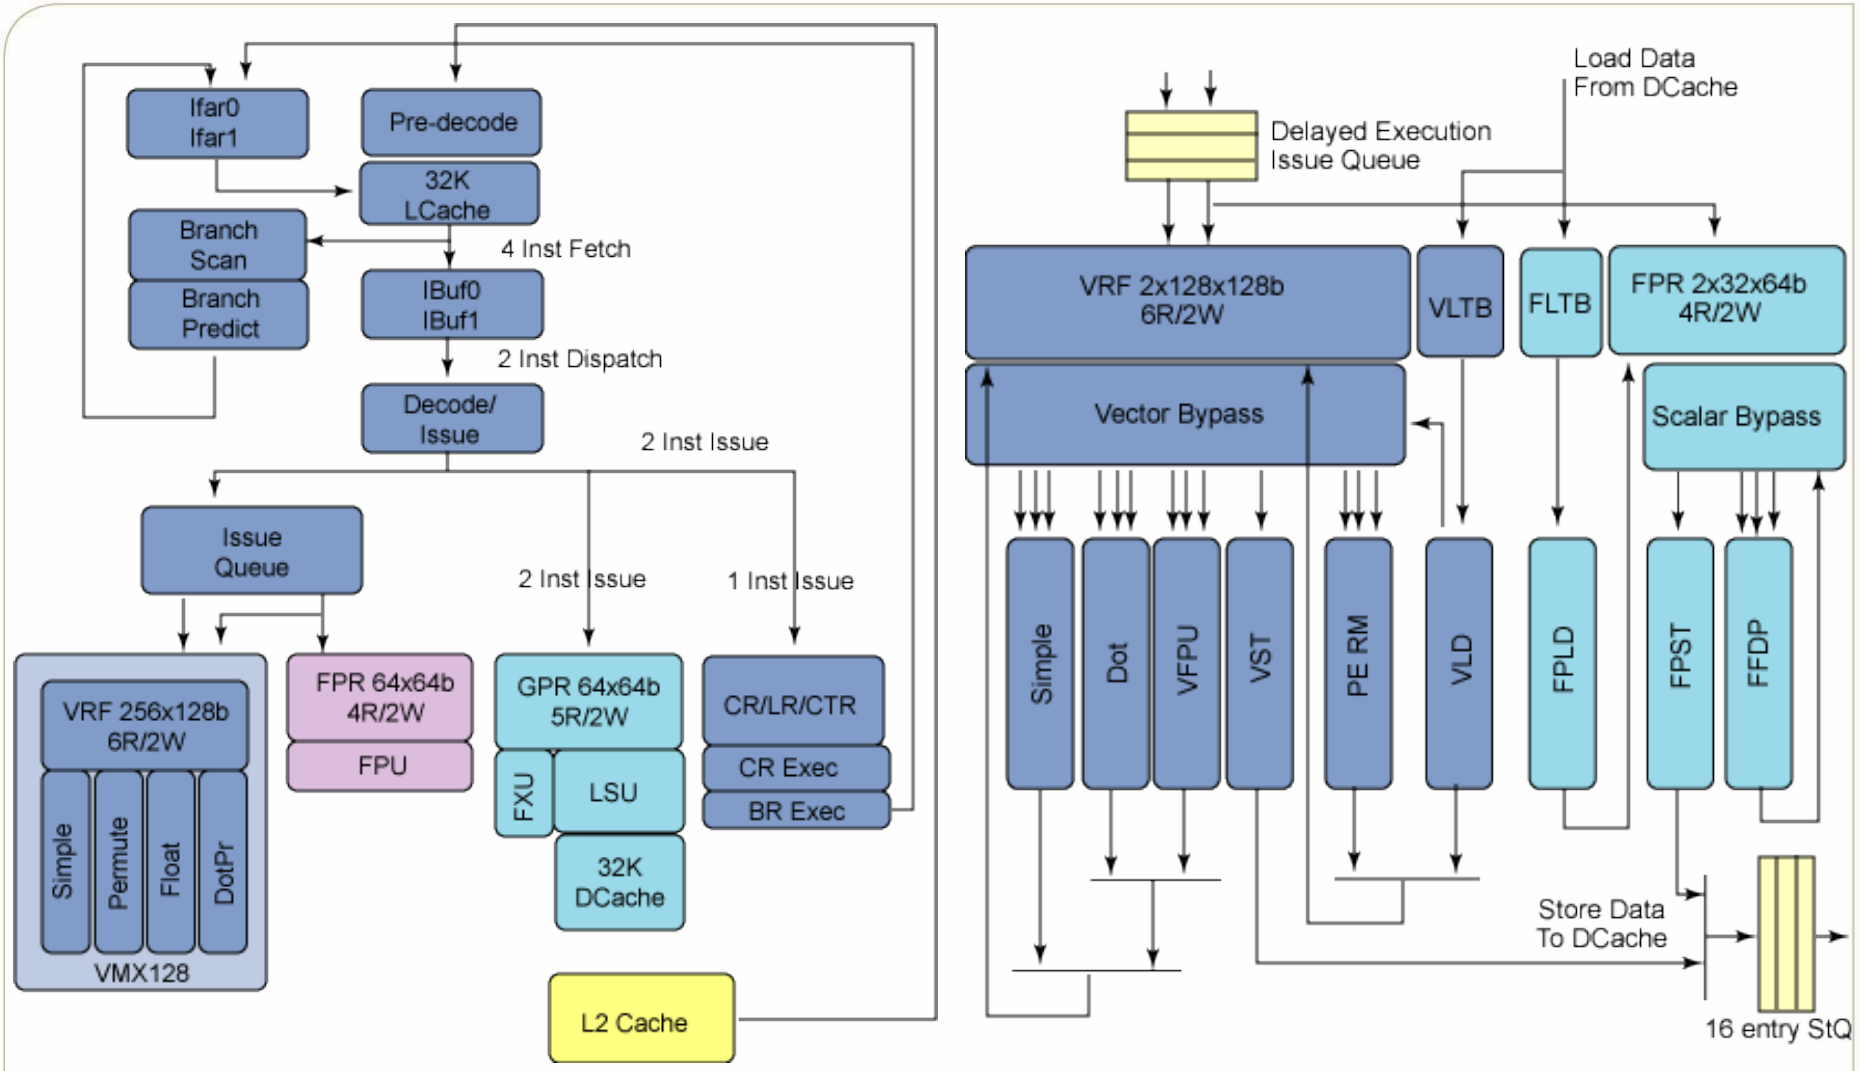


Image from J. Brown, "The Microsoft Xbox 360 CPU story"
www-128.ibm.com/developerworks/power/library/pa-fpfxbox

Georgia Institute of Technology

# Cell PPE architecture

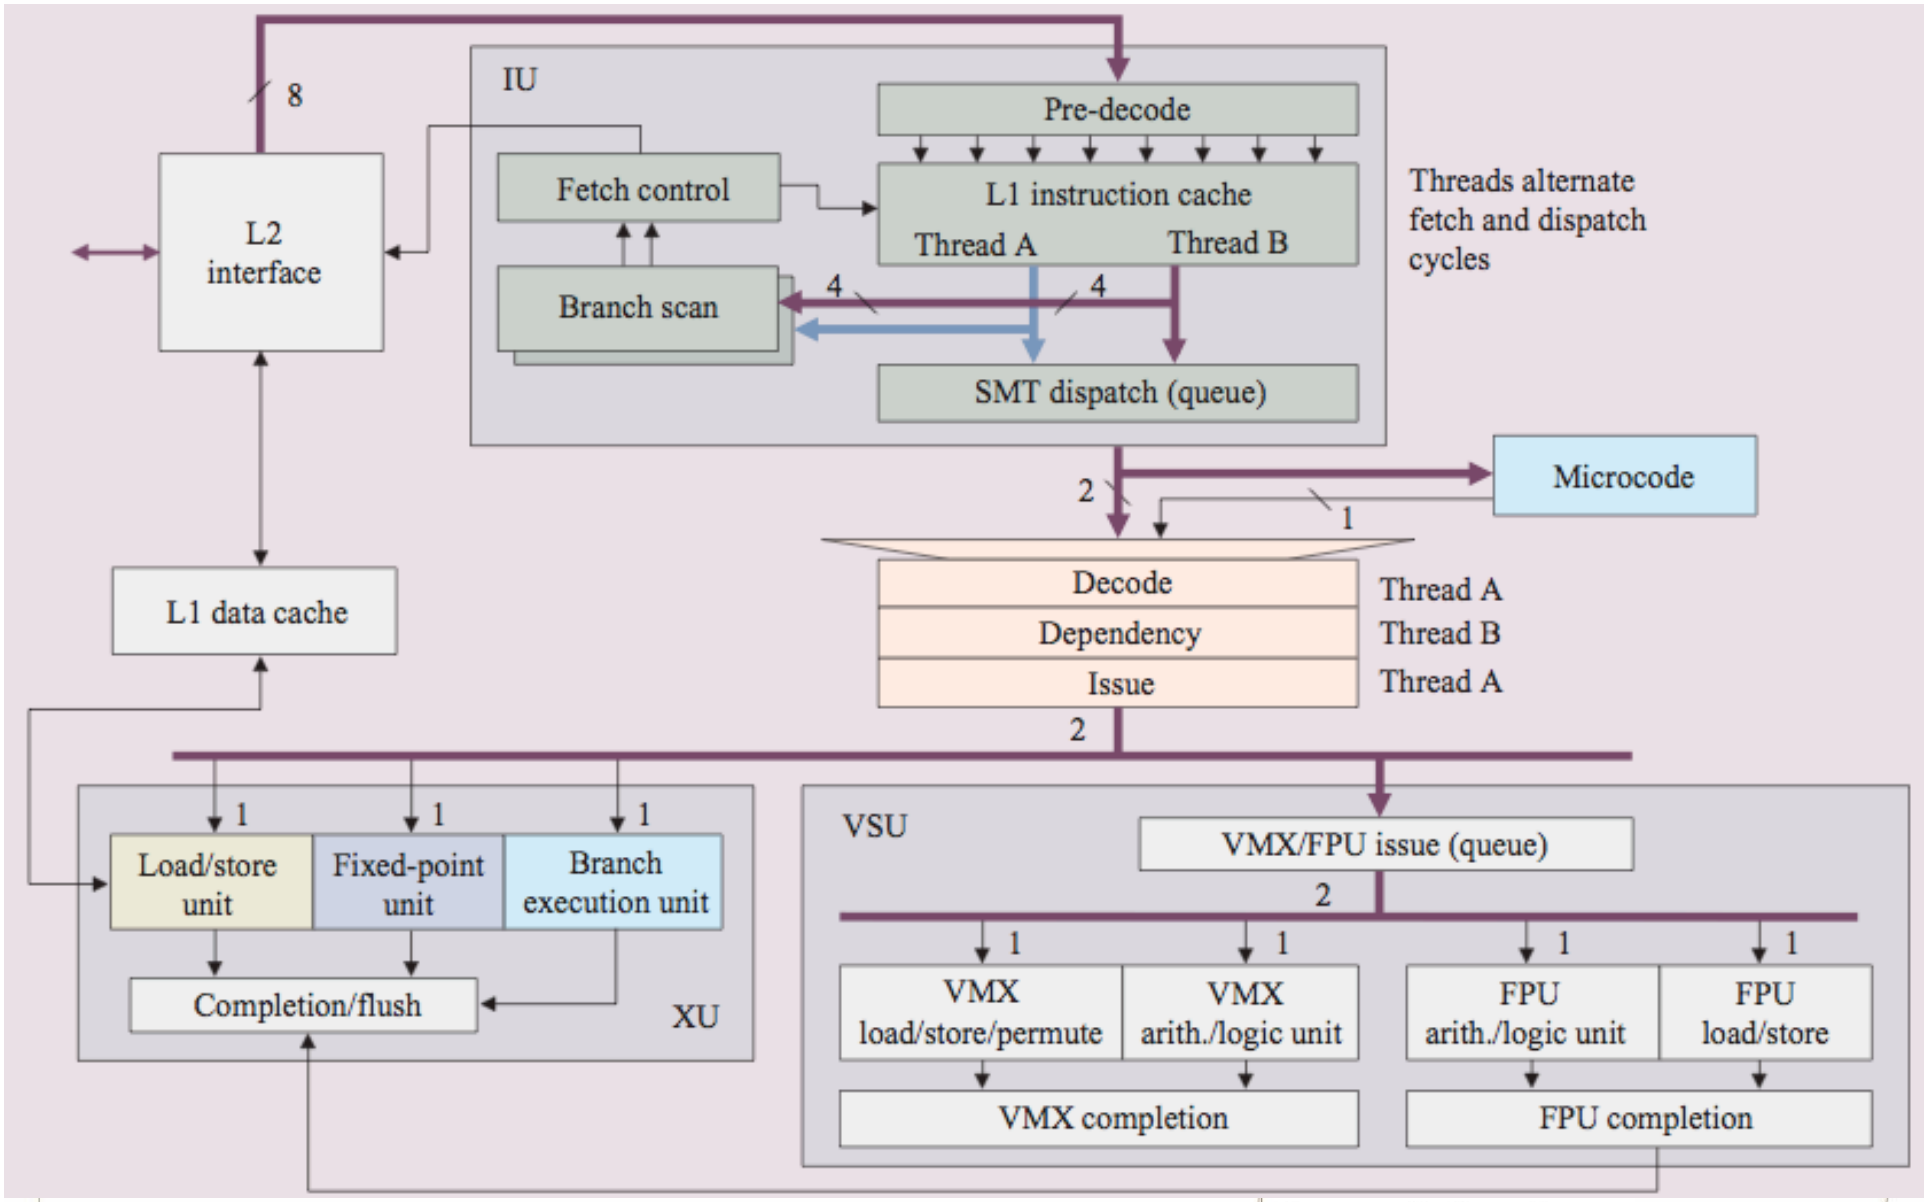


Image from J.A. Kahle et al., "Introduction to the Cell Processor," *IBM J. Res. & Dev.,* Vol. 49, No. 4/5, July/Sept. 2005, pp. 589-604.

# Cell PPE pipeline

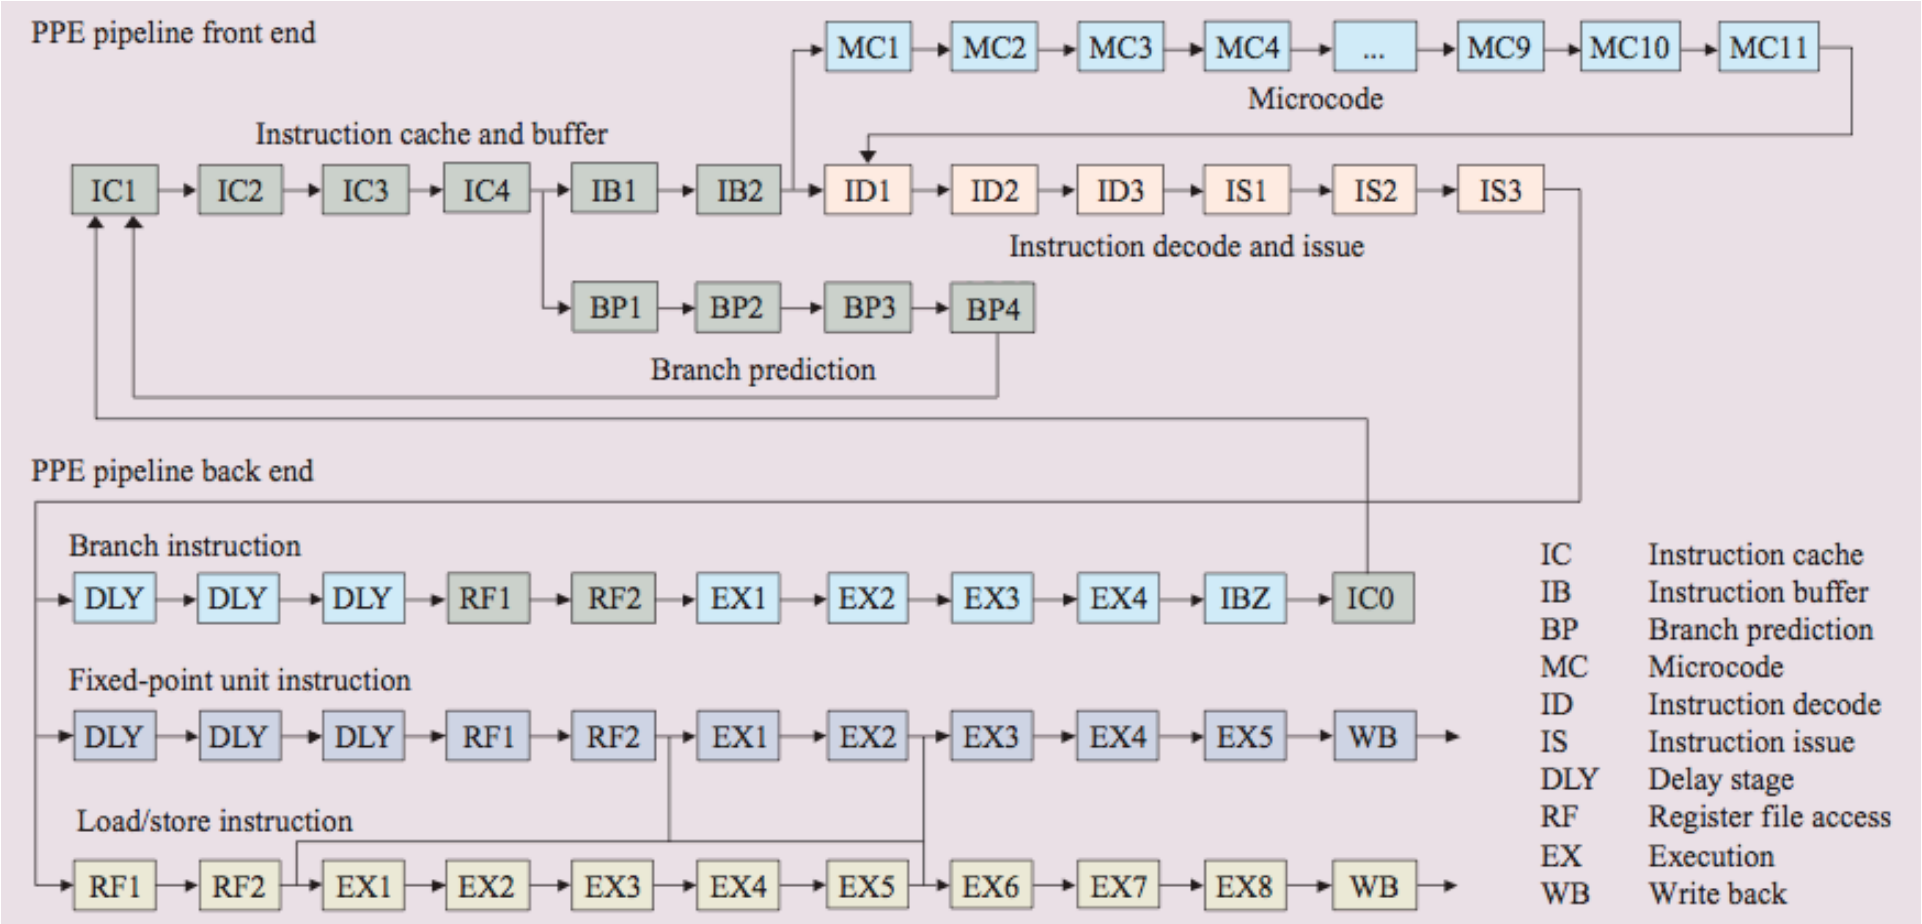


Image from J.A. Kahle et al., "Introduction to the Cell Processor," *IBM J. Res. & Dev.,* Vol. 49, No. 4/5, July/Sept. 2005, pp. 589-604.

Georgia Institute of Technology

# Cell SPE architecture

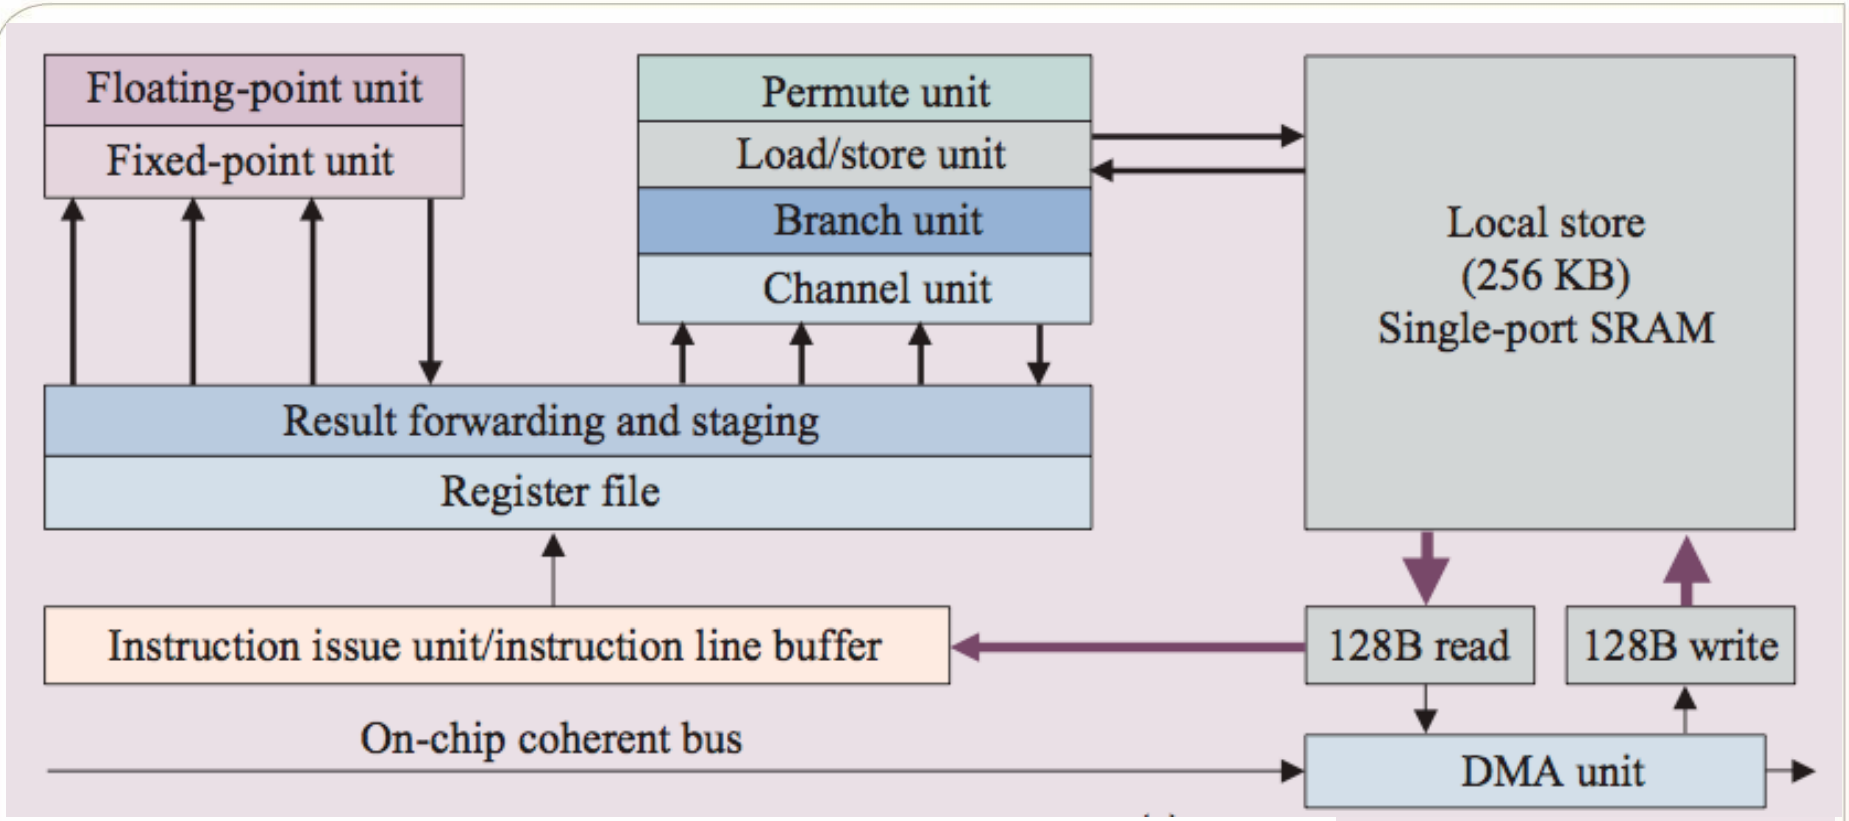


"The SPEs are not coprocessors."

– Mike Acton, Engine Director, Insomniac Games, and keeper of www.cellperformance.com

Image from J.A. Kahle et al., "Introduction to the Cell Processor," *IBM J. Res. & Dev.,* Vol. 49, No. 4/5, July/Sept. 2005, pp. 589-604.

Georgia Institute of Technology

# Cell SPE pipeline

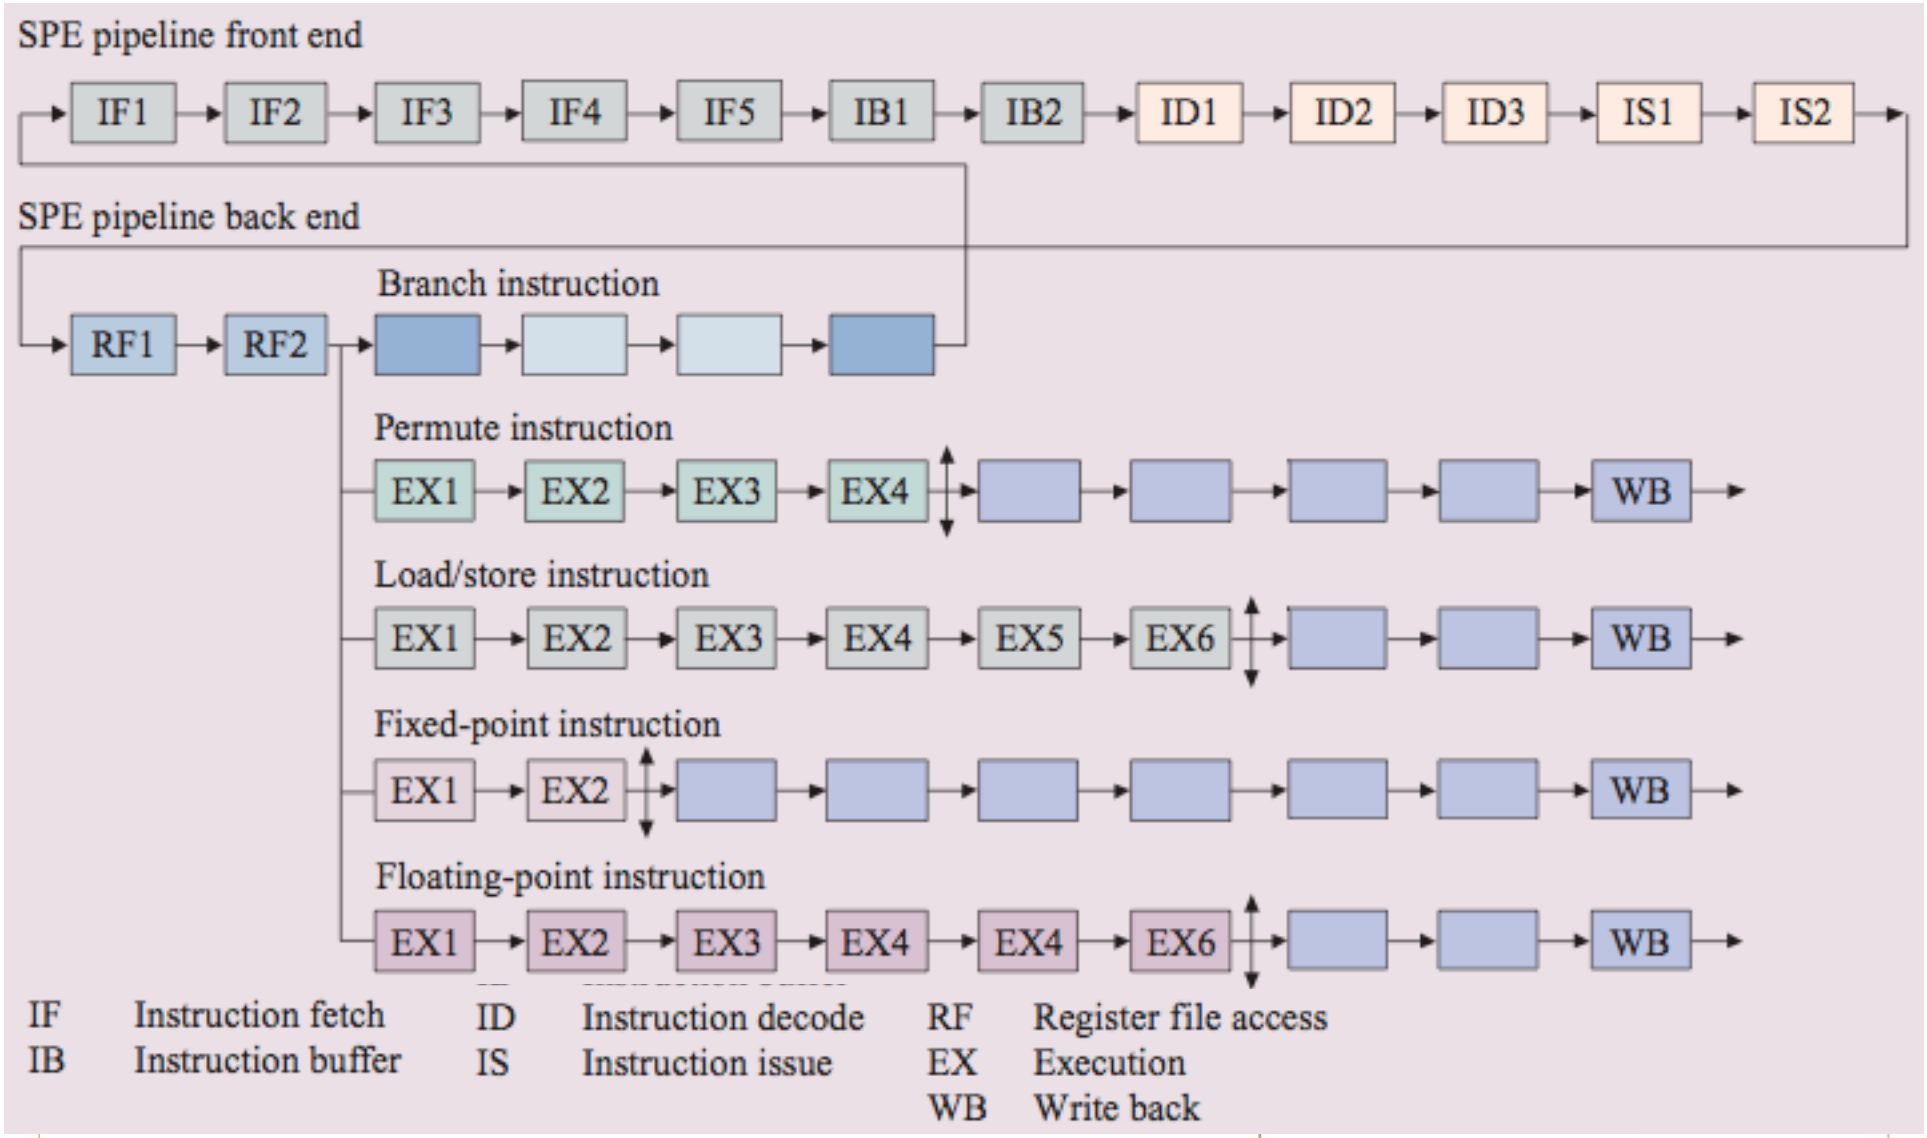


Image from J.A. Kahle et al., "Introduction to the Cell Processor," *IBM J. Res. & Dev.,* Vol. 49, No. 4/5, July/Sept. 2005, pp. 589-604.

# GPUs: Xbox 360 Xenos vs. PS3 RSX

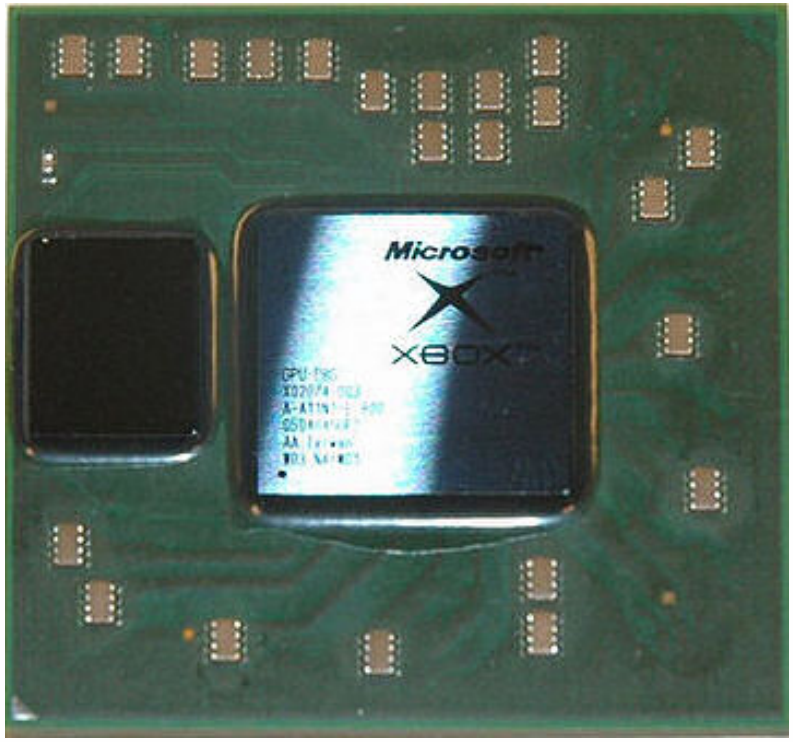



**Images not to scale**

Xenos image from Wikipedia

RSX image from
www.pctuning.cz/index.php?option=com_content&
task=view&id=7787&Itemid=88&limit=1&limitstart=2

Georgia Institute of Technology

# Xbox 360 GPU architecture

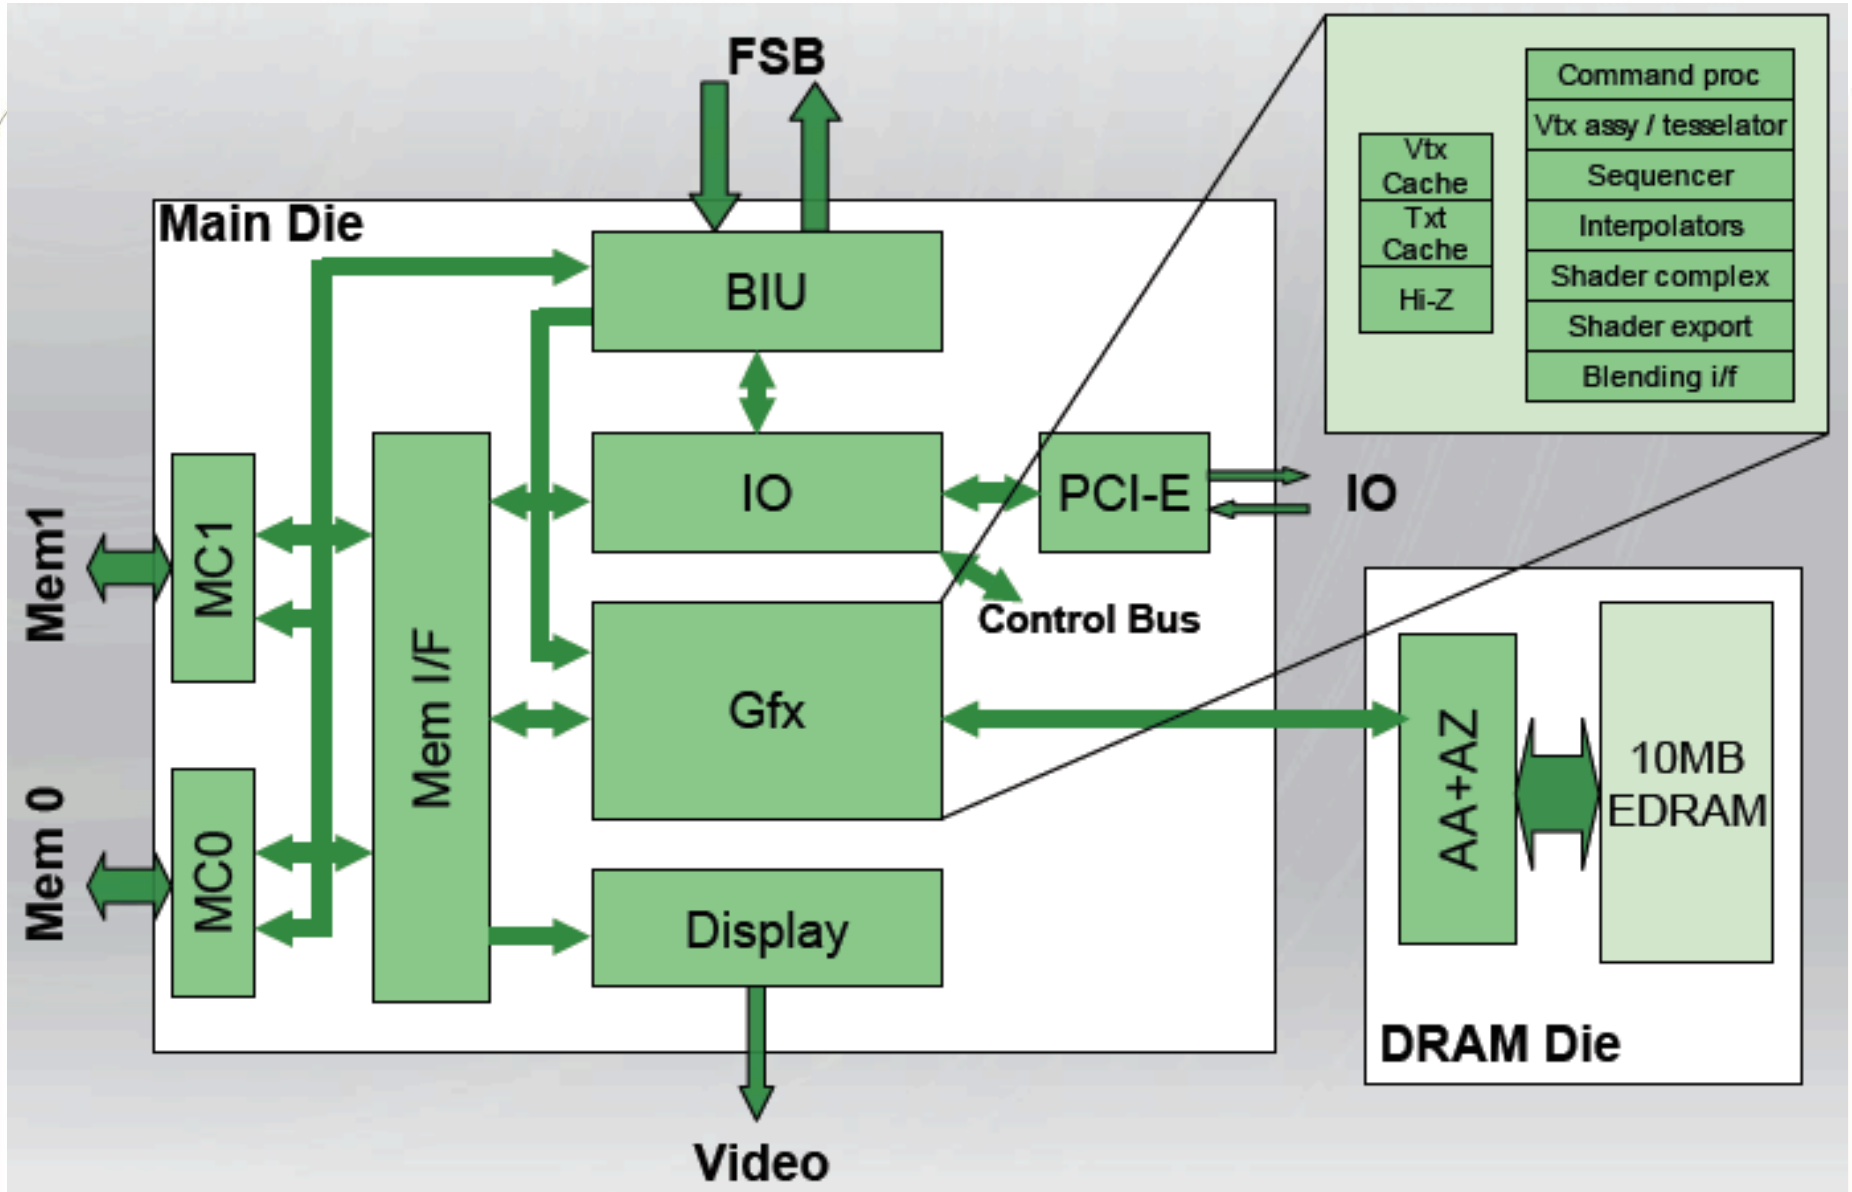


Image from J. Andrews and N. Baker, "Xbox 360 System Architecture," *Hot Chips* Presentation

# Xbox 360 GPU layout

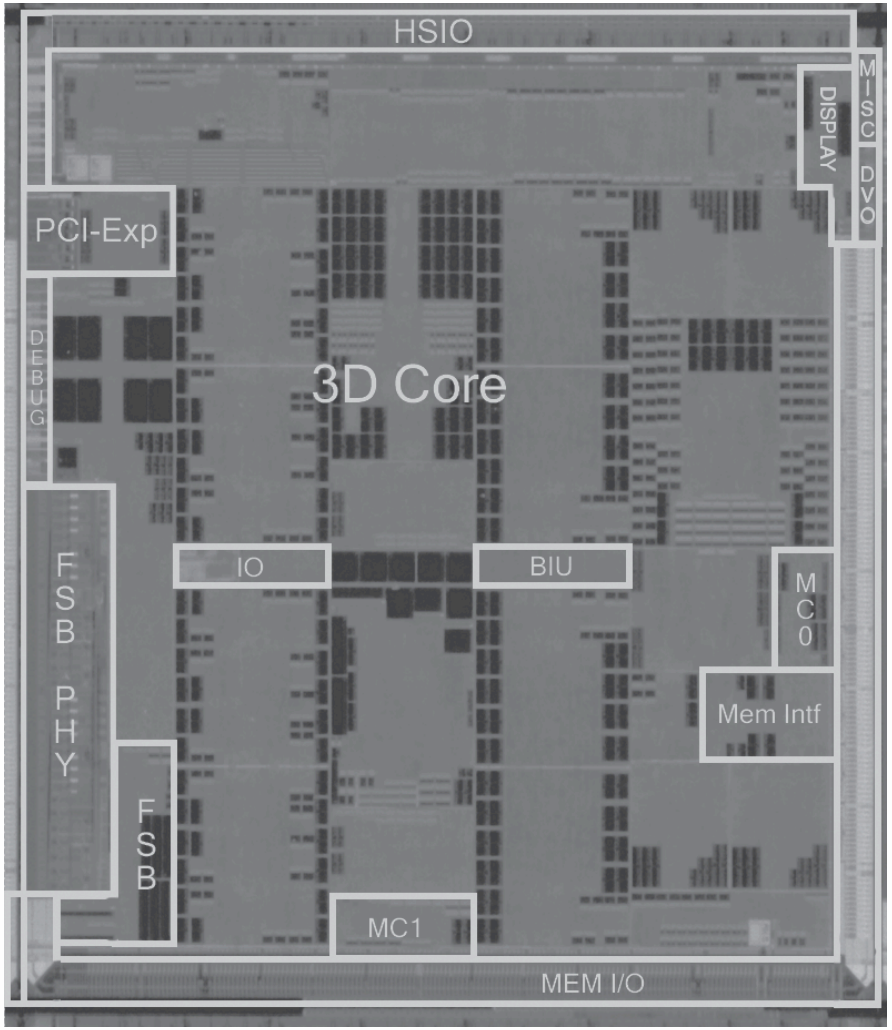


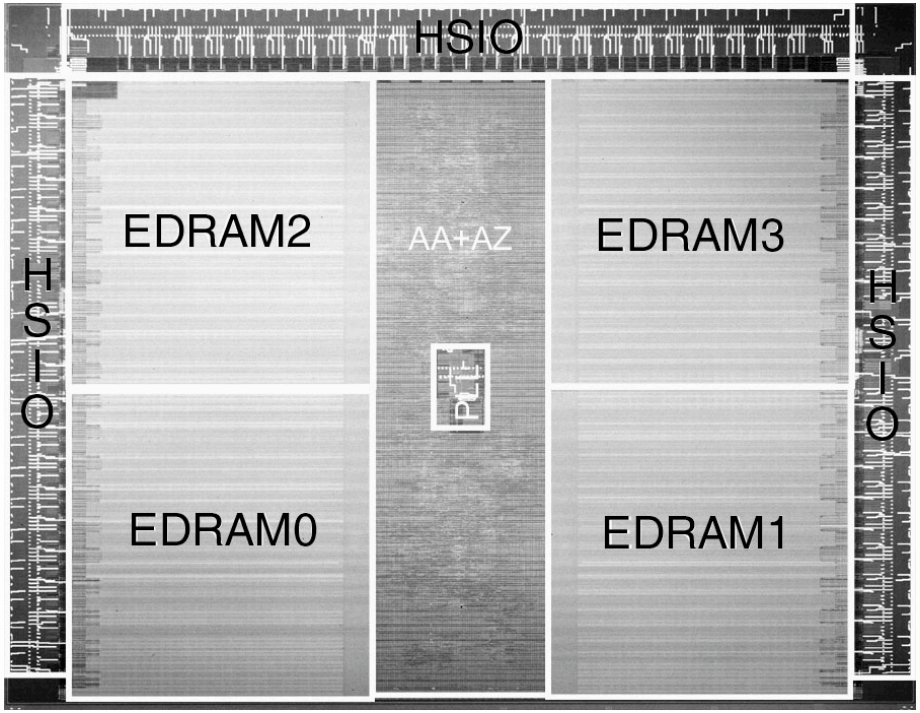


**Images not to scale**

Image from J. Andrews and N. Baker, "Xbox 360 System Architecture," *Hot Chips* Presentation

# GPUs: Xbox 360 Xenos vs. PS3 RSX (1)

- Xbox 360: ATI Xenos
- 500 MHz
- Precursor to Radeon HD 2000 series
- 16 vertex fetch units with built-in tesselation
- 48 unified shaders (can do vertices or pixels)
  - All 48 have to be doing either vertices or pixels in one clock cycle
  - Can change from cycle to cycle
  - Rumored to have more than 48 per chip; gets higher yields
- 16 texture interpolating (filtering) units
- 16 texture fetch (addressing) units
- 8 render output units

- PS3: NVIDIA RSX "Reality Synthesizer"
- 550 MHz
- Somewhat like 7800 (G70)

- 24 pixel shaders
- 8 vertex shaders

- 24 texture filtering units
- 8 texture addressing units
- 8 render output units

# GPUs: Xbox 360 Xenos vs. PS3 RSX (2)

- 10 MB video buffer eRAM die includes some custom logic for color, alpha compositing, Z/stencil buffering, and anti-aliasing
  - Does not include textures
  - 256 GB/sec bandwidth to GPU
  - Currently on separate die on same package
  - Guess will later probably put on same die
  - Buffer in eRAM is copied to main memory for output
- Video buffer part of 256 MB video RAM
- Cell FlexIO bus interface
  - 20 GB/s read to the Cell and XDR memory
  - 15 GB/s write to the Cell and XDR memory

# Xbox 360 CPU/GPU/memory synergy

- GPU can read data directly from CPU's L2 cache through the FSB without going through main memory
- Facilitates XPS (Xbox Procedural Synthesis), in which CPU decompresses 3D data for the GPU
- "For render-to-texture, GPU must first 'flush' appropriate buffer to main memory before using it as a texture"
- Shaders can output directly to main memory instead of frame buffer (good if need to use GPU for physics, etc.)

Information from J. Andrews and N. Baker, "Xbox 360 System Architecture," *IEEE Micro*, March-April 2006, pp. 25-37.

# Xbox 360 data flow example

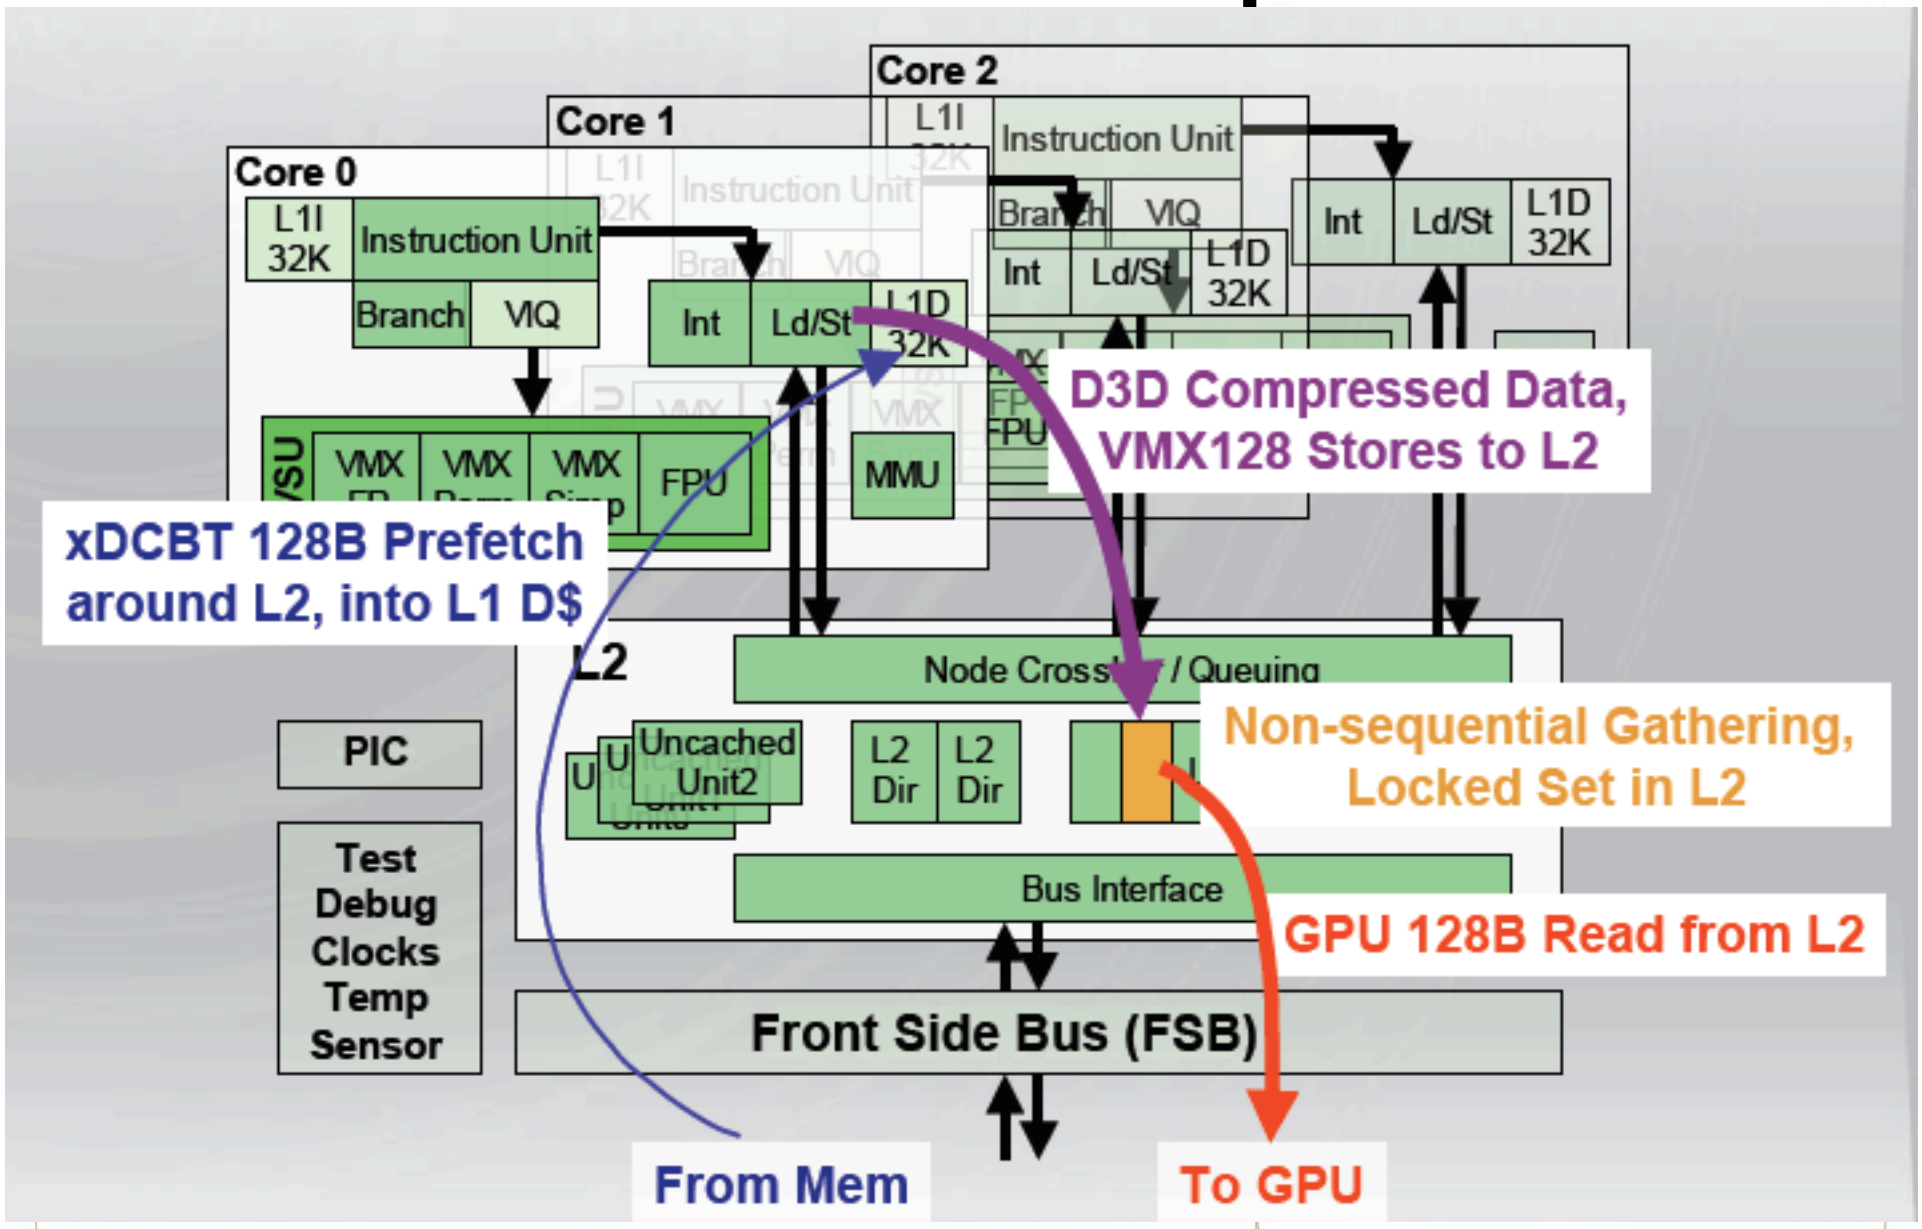


Image from J. Andrews and N. Baker, "Xbox 360 System Architecture," *Hot Chips* Presentation