AMD
Smarter Choice

Radeon™ HD 2900

Michael Doggett

October 3, 2007

Overview

AMD
Smarter Choice

- Starting Point
- Requirements
- Top level
- Pipeline Blocks from ‘top to bottom’
  - Command Processor
  - Shader Setup Engine
  - Threaded Dispatcher
  - Shader Core
  - Texture Filtering
  - Render Backend
- Conclusion

Starting point

AMD
Smarter Choice

- Combine the best of existing technology
  - R5xx series
    - Heavily threaded shader cores
      - Hides latency of memory fetch
    - Vec4+1 Vertex and Vec3+1 Pixel shaders
    - Ringbus memory subsystem
  - XBOX 360 GPU
    - Unified shader architecture
      - Vertex and Pixel
      - Vec4+1
      - Stream Out
    - Unified L1 texture cache
    - Introduced Tessellator

Requirements

- DirectX10 compatible
- Support new driver model
  - Vista driver model
- Scalable family
  - "Number" of shader cores, texture units, render back-ends.
  - Shader scalable in number of pipes, SIMDs.
  - Target specific cost, feature set and performance levels for each part

AMD

RADEON 2900

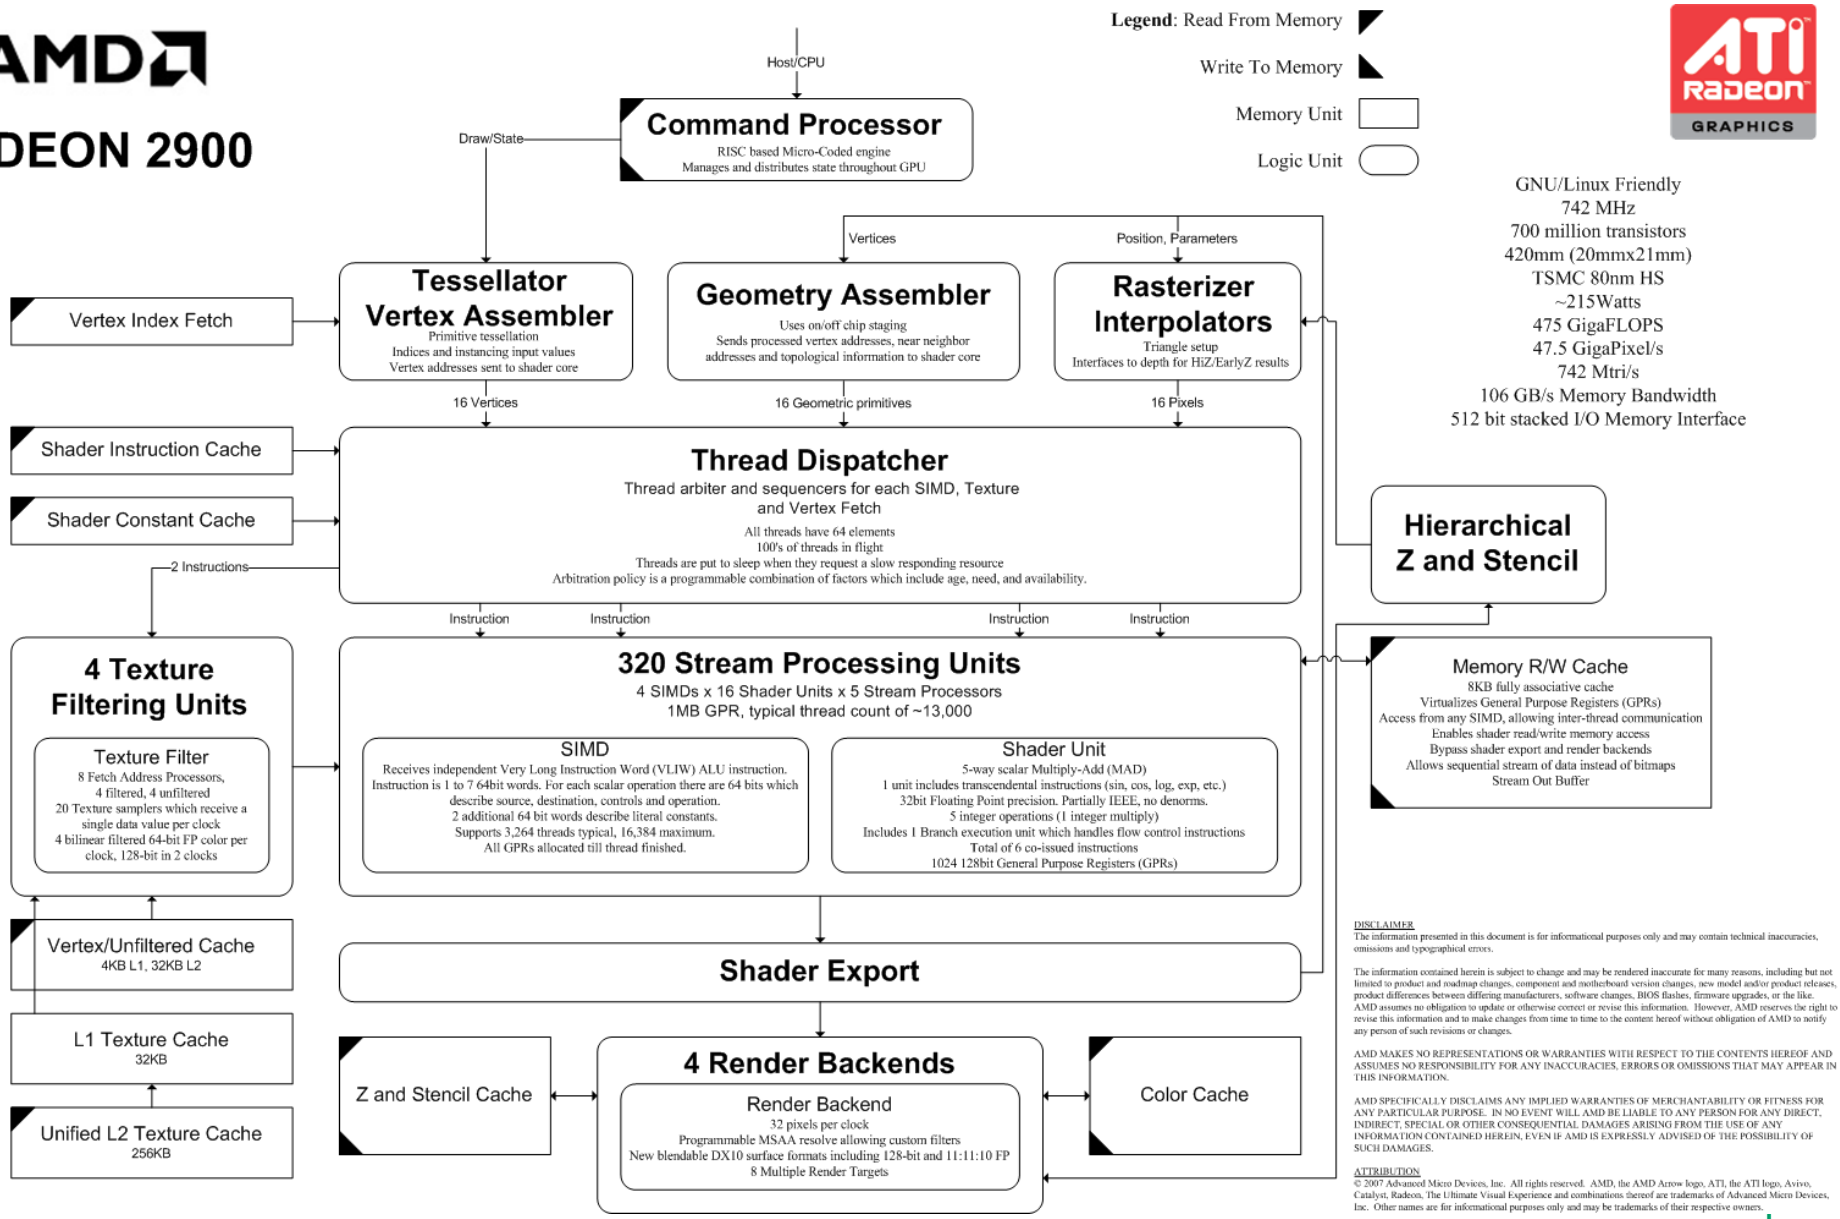

ATI RADEON GRAPHICS

GNU/Linux Friendly
742 MHz
700 million transistors
420mm (20mmx21mm)
TSMC 80nm HS
~215Watts
475 GigaFLOPS
47.5 GigaPixel/s
742 Mtri/s
106 GB/s Memory Bandwidth
512 bit stacked I/O Memory Interface

Radeon HD 2000 Series

AMD Smarter Choice

Radeon	2900	2600	2400
Stream Processors	320	120	40
SIMDs	4	3	2
Shader Units	16	8	4
Texture Units	16	8	4
Render Backends	16	4	4
L2 texture cache (KB)	256	128	0
Technology (nm)	80	65	65
Area (mm2)	420	153	82
Transistors (Millions)	720	390	180
Memory Pins	512	128	64

Where next ?

AMD
Smarter Choice

- Move fixed functions blocks to shader
  - Improve programmability, reduce area, improve reuse, maintain/target performance
- Enhancements for GPGPU
  - Improved precision and compliance
  - New APIs, new functions
- New technologies such as 65, 55, 45, 32…
- Graphics and gaming keeps on evolving
  - DX-next is already being discussed
  - We are well into next generation and next-next generations …

Hardware-Software Interface

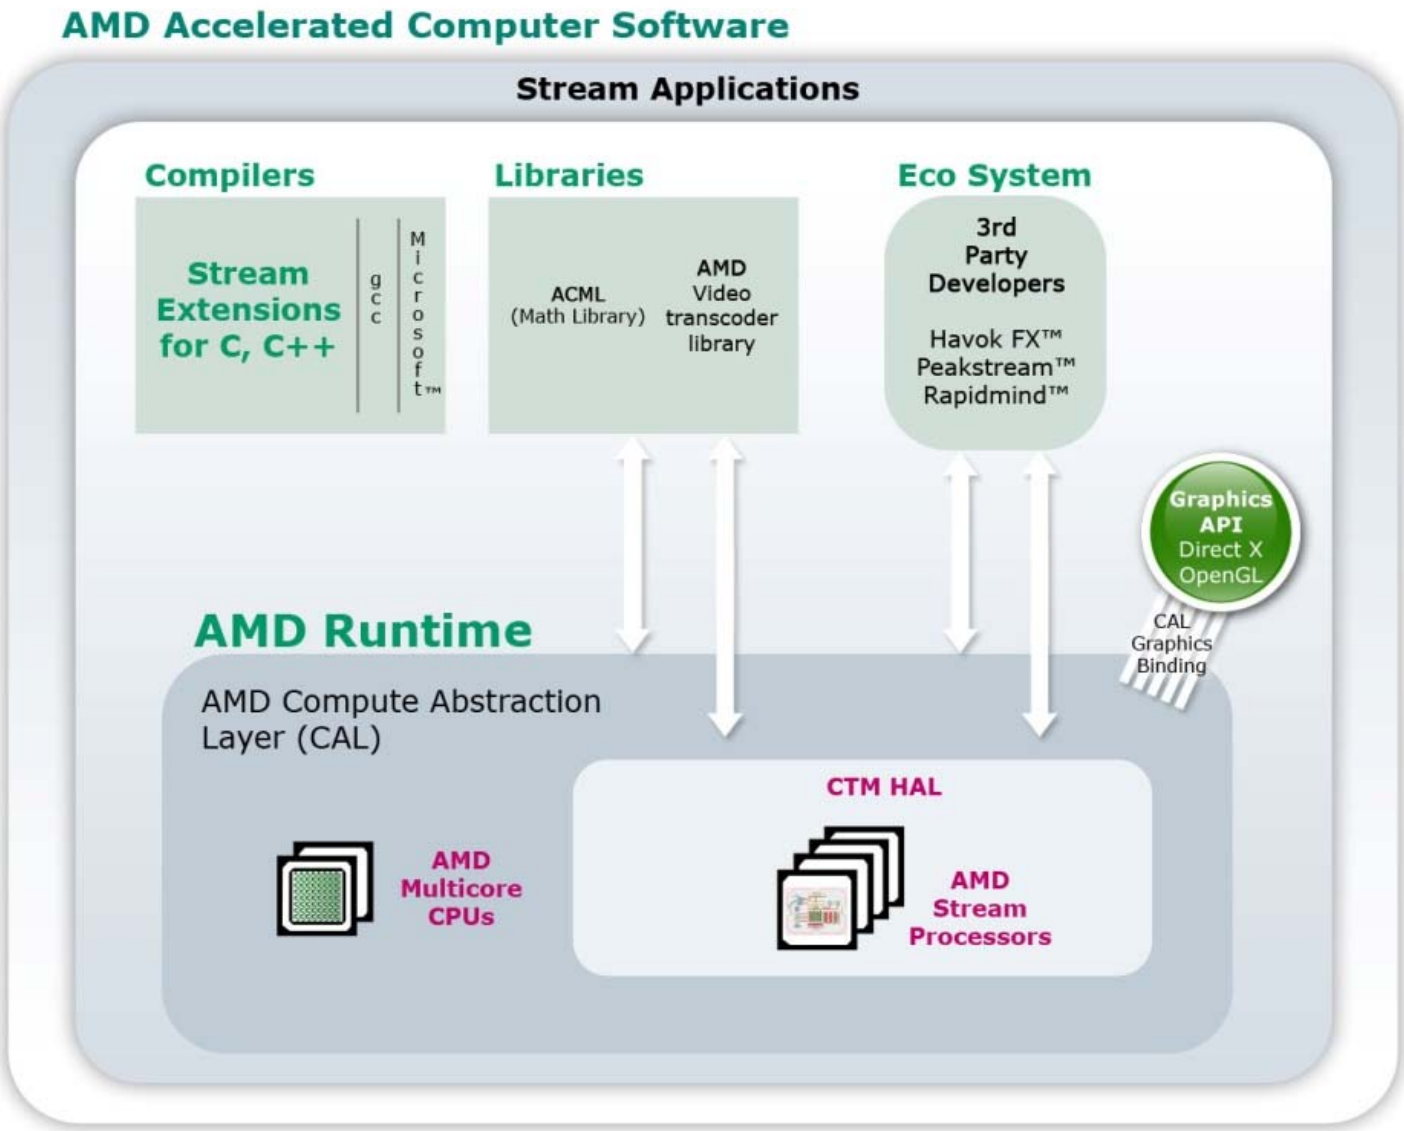

Radeon HD 2900

AMD
Smarter Choice

- Unified shader
  - Vertex, Geometry and Pixel
  - Multiple SIMD
  - VLIW 5-way scalar
- Shader cached memory read/write
- Geometry shader on/off chip storage
- 512 bit stacked I/O Memory Interface
- Full DX10 functionality

Questions and Demo

AMD
Smarter Choice

- Thanks to Eric Demers, Mike Mantor and Norm Rubin

Performance Improvements

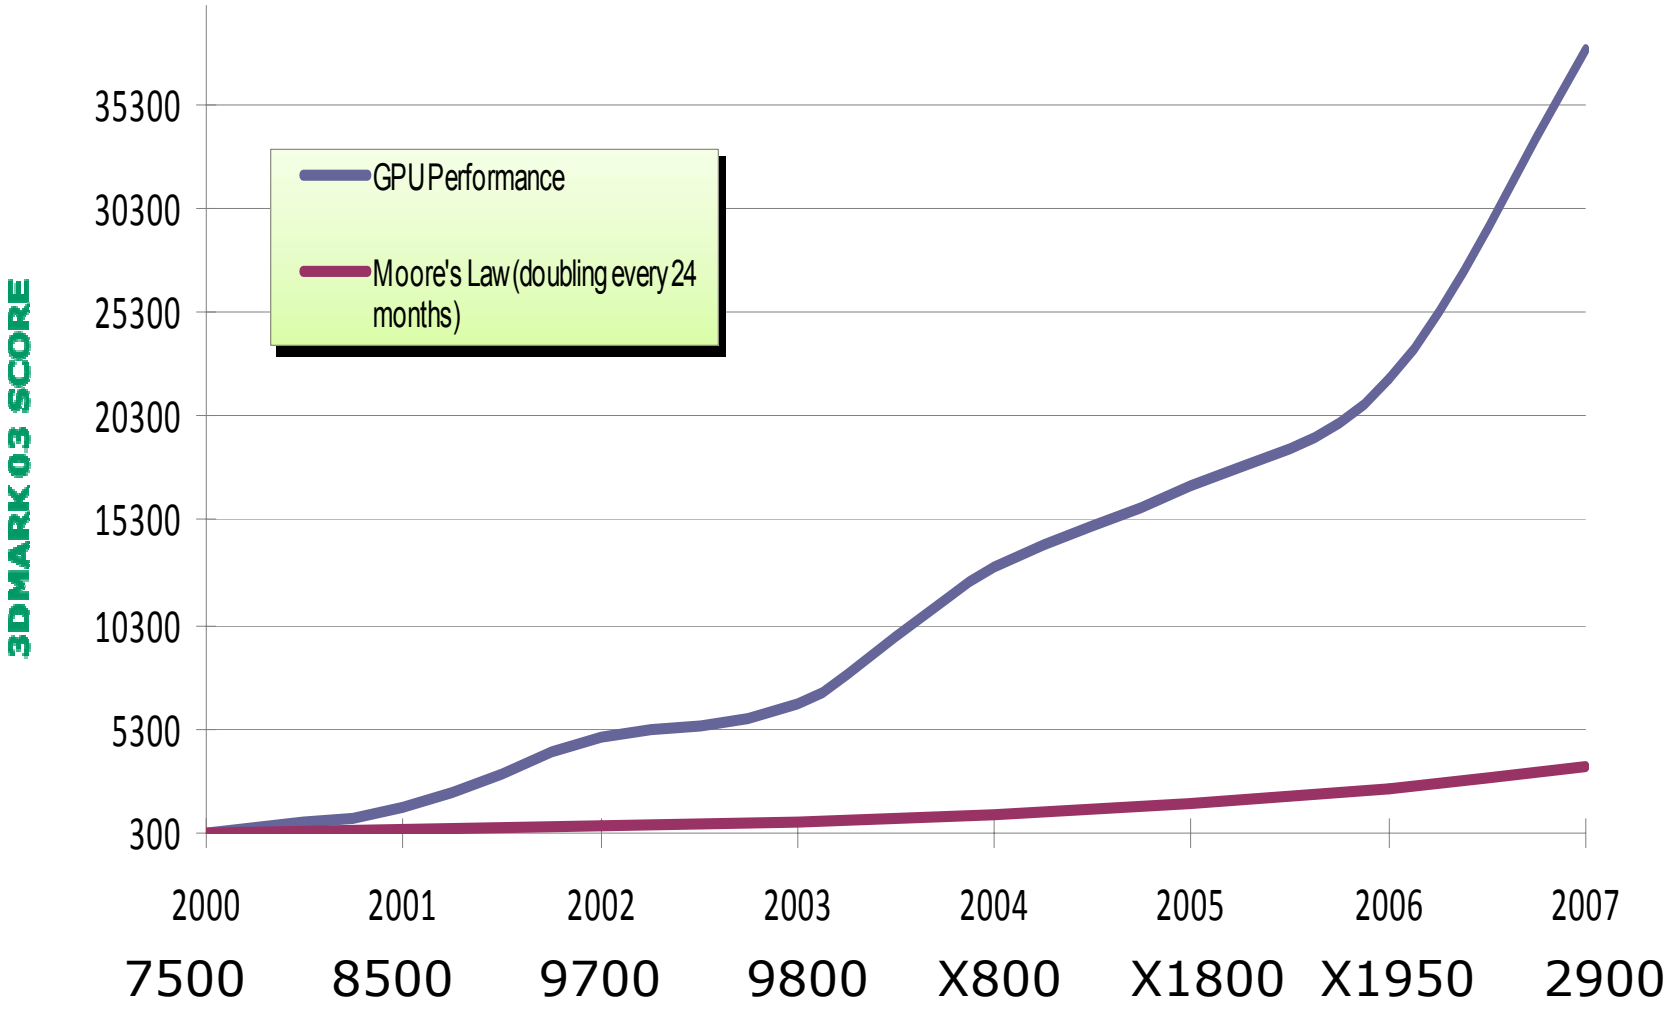

	Rage Pro	Rage 128	Radeon	Radeon 8500	Radeon 9700 Pro	Radeon 9800 XT	Radeon X850 XT Platinum Edition	Radeon X1800 XT	Radeon X1950 XTX	Radeon HD 2900 XT	
Year	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	
Transistor Size	350 nm	250 nm	180 nm	150 nm	150 nm	150 nm	130 nm	90 nm	90nm	80nm	
Transistor Count	5 million	13 million	30 million	60 million	110 million	110 million	160 million	321 million	384 million	700 million	*1.7x / year*
Clock Speed	75 MHz	100 MHz	183 MHz	275 MHz	325 MHz	412 MHz	550 MHz	625 MHz	650 MHz	740 MHz	*1.3x / year*
Rendering Pipelines / Shader Processors	1	2	2	4	8	8	16	16	48	64	*1.6x / year*
Memory Bandwidth (GB/sec)	0.6	1.6	5.9	8.8	19.8	23.4	37.8	44.8	64.0	106.0	*1.8x / year*
Pixel Shading / Fill Rate (Mpixels/sec)	75	200	366	1100	2600	3300	8800	10000	31200	47360	*2.0x / year*
Vertex Processing (Mvertices/sec)	4	8	30	69	325	412	825	1250	1300	11840	*2.4x / year*
System Bus Bandwidth (GB/sec)	0.53	1.06	1.06	1.06	2.11	2.11	4	4	4	4	*1.3x / year*
DirectX Version and Shader Model Support	DirectX 5	DirectX 6	DirectX 7	DirectX 8.1 (SM1.4)	DirectX 9.0 (SM2.0)	DirectX 9.0 (SM2.0)	DirectX 9.0 (SM2.0b)	DirectX 9.0 (SM3.0)	DirectX 9.0 (SM3.0)	DirectX 10.0 (SM4.0)	
Features	Hardware Triangle Setup	128-bit memory interface, AGP 4X	Hardware transform & lighting, DDR memory support	Programmable Shaders, Higher Order Surfaces	Floating Point Processing, 256-bit memory interface, AGP 8X	Unlimited shader instructions, 256MB DDR2 memory support	PCI Express x16, GDDR3 memory support	Dynamic shader flow control, 128-bit precision, HDR output	GDDR4 memory support	Unified Shader Architecture, 512-bit memory interface	

Abstract :

The ATI Radeon HD 2900 developed by AMD is a Graphics Processing Unit (GPU) capable of massively parallel computation for high performance 3D graphics and general purpose algorithms. The unified shader architecture consists of a combination of MIMD and SIMD architectures of 5 way scalar arithmetic units running in parallel. The shader uses multi-threading to hide latency of memory access so that compute units are kept busy. The threads consist of vertex, geometry and pixel threads that represent different programmable stages of a traditional 3D graphics pipeline mapped onto a single scheduled shader unit. Varied distributed and unified caches are used for data, instructions, read only texture reads and vertex data. A ring based memory subsystem allows multiple clients to access multiple memory channels.