



Anatomy of AMD's TeraScale Graphics Engine



Mike Houston

SIGGRAPH2008

Beyond Programmable Shading: Fundamentals

Design Goals



SIGGRAPH2008

Focus on Efficiency

f(Perf/Watt, Perf/\$)

Scale up processing power and AA performance

- Target >2x previous generation

Enhance stream computing capability

- Faster and more flexible

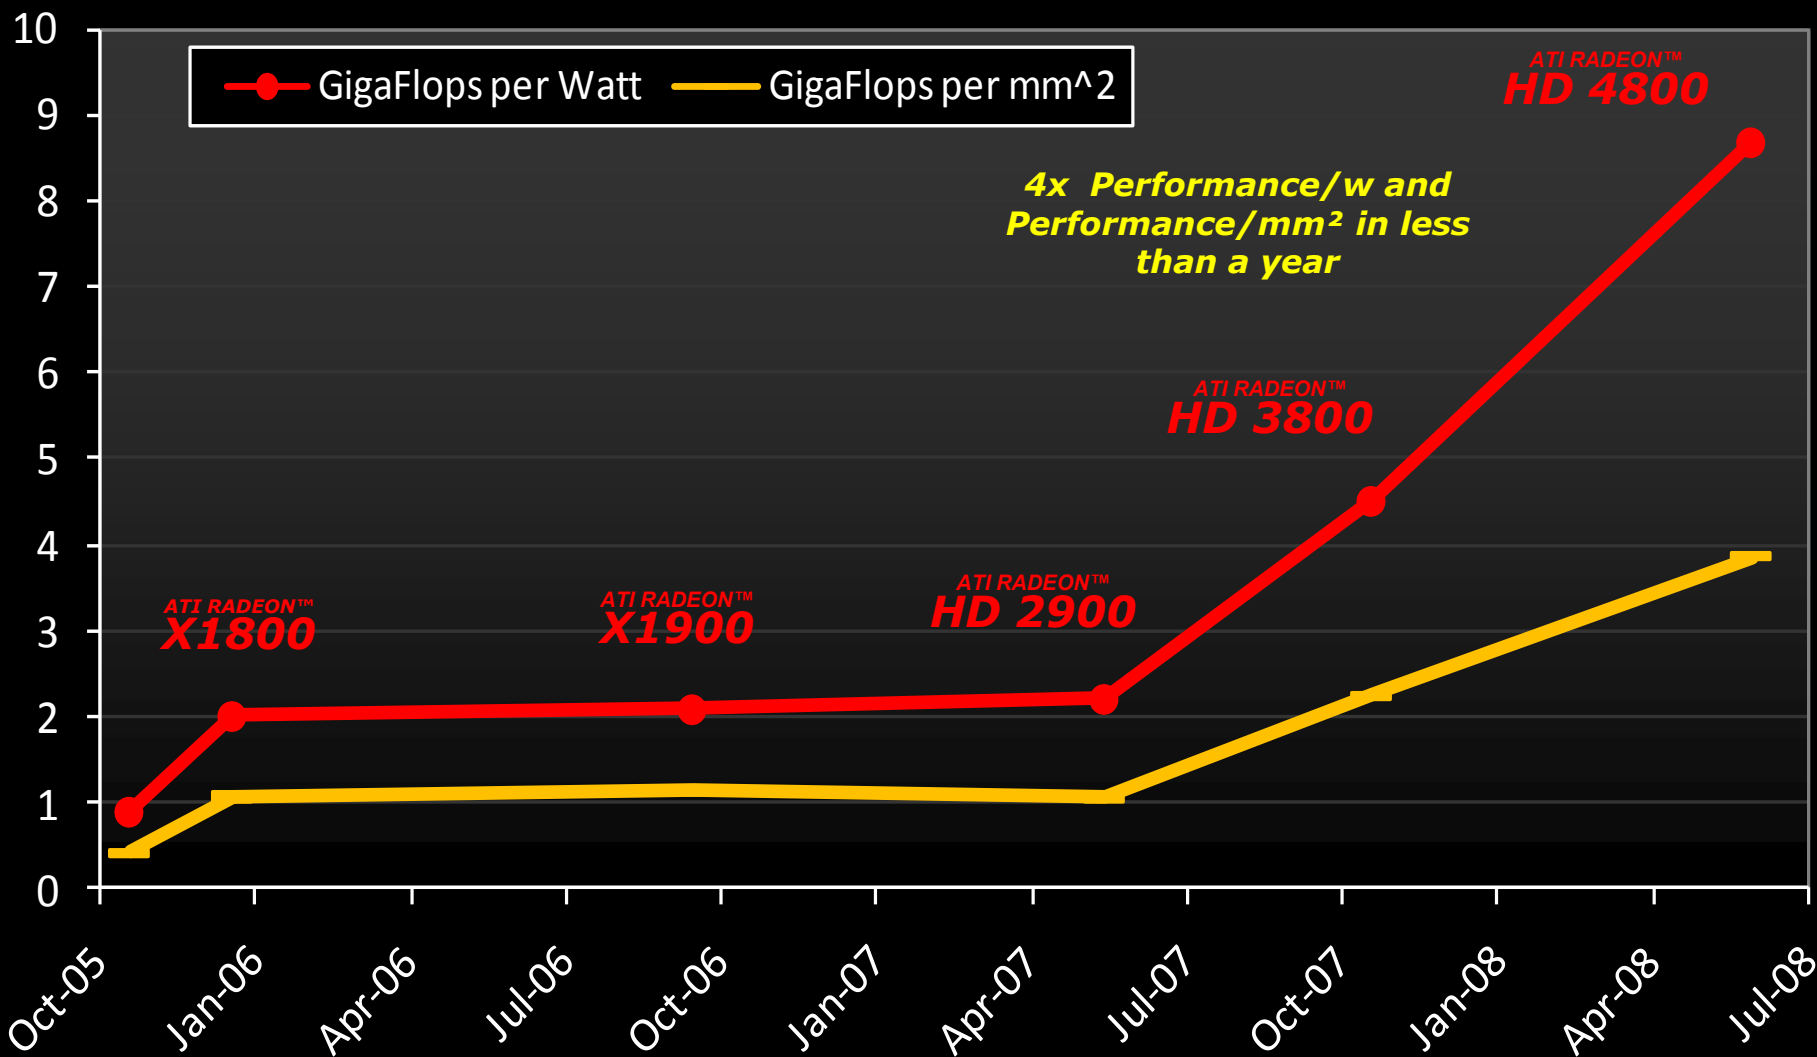
Implement advanced feature set

- DirectX® 10.1, tessellation, UVD2, PCIe® 2.0, and more...

Design Efficiency



SIGGRAPH2008




Terascale Graphics Engine



SIGGRAPH2008

- 800 highly optimized stream processing units
- New SIMD core layout
- Optimized texture units
- New texture cache design
- New memory architecture
- Optimized render back-ends for fast anti-aliasing performance
- Enhanced geometry shader & tessellator performance




ATI Radeon™ HD 4800 Series Architecture



SIGGRAPH2008

- 10 SIMD cores
 - Each with 80 32-bit Stream Processing Units (800 total)
- 40 Texture Units

	ATI Radeon™ HD 3870	ATI Radeon™ HD 4870	Difference
Die Size	190 mm ²	260 mm ²	1.4x
Memory	72 GB/sec	115 GB/sec	1.6x
AA Resolve	32	64	2x
Z/Stencil	32	64	2x
Texture	16	40	2.5x
Shader	320	800	2.5x




SIMD Cores



SIGGRAPH2008

- Each core:
 - Includes 80 scalar stream processing units in total + 16KB Local Data Share
 - Has its own control logic and runs from a shared set of threads
 - Has 4 dedicated texture units + L1 cache
 - Communicates with other SIMD cores via 16KB global data share
- New design allows texture fetch capability to scale with shader power, maintaining 4:1 ALU:TEX ratio


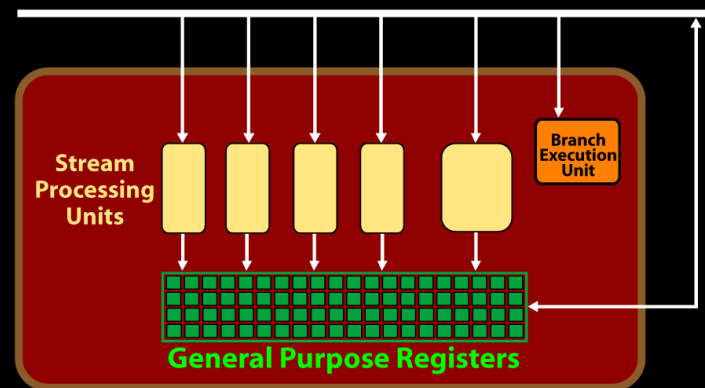


Stream Processing Units



SIGGRAPH2008

- 40% increase in performance per mm²*
- More aggressive clock gating for improved Performance per Watt *
- Fast double precision processing (240 GigaFLOPS)
- Integer bit shift operations for all units (12.5x improvement *)



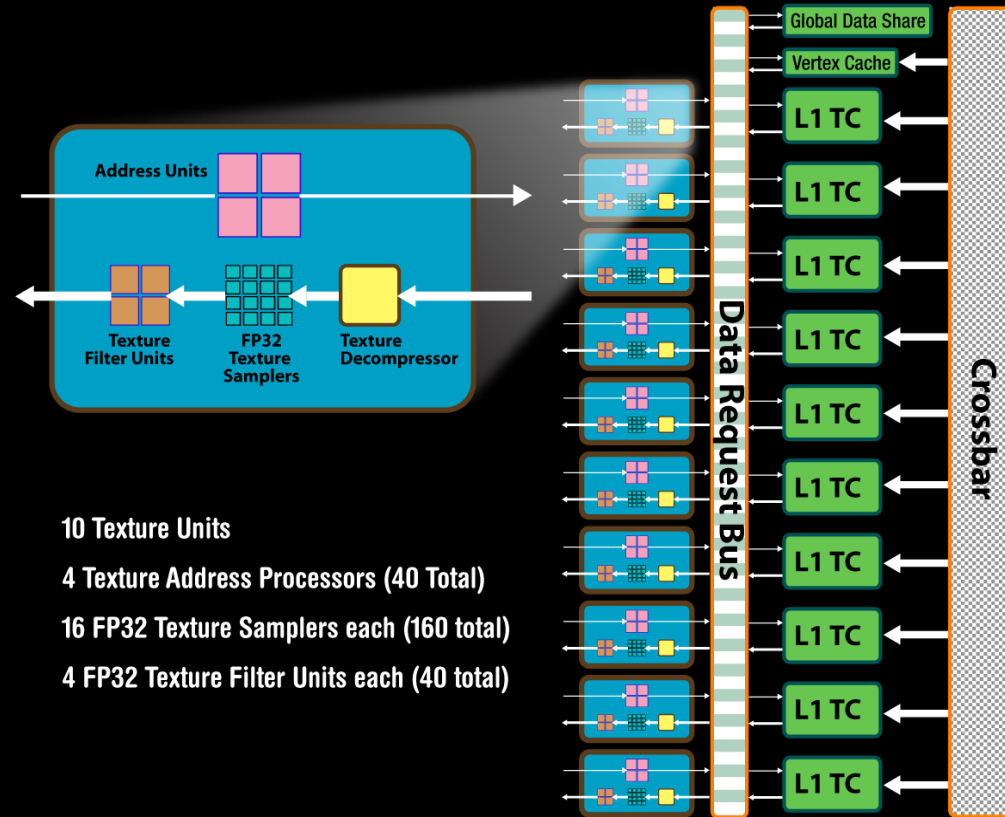
* Internal AMD test results comparing ATI Radeon™ HD 4800 series and ATI Radeon™ HD 3800 series

Texture Units




SIGGRAPH2008

- Streamlined design
 - 70% increase in performance/mm² *
- More performance
 - Double the texture cache bandwidth of the ATI Radeon™ HD 3800 series*
 - 2.5x increase in 32-bit filter rate *
 - 1.25x increase in 64-bit filter rate *
 - Up to 160 fetches per clock *



Peak 32-bit texture fetch rate



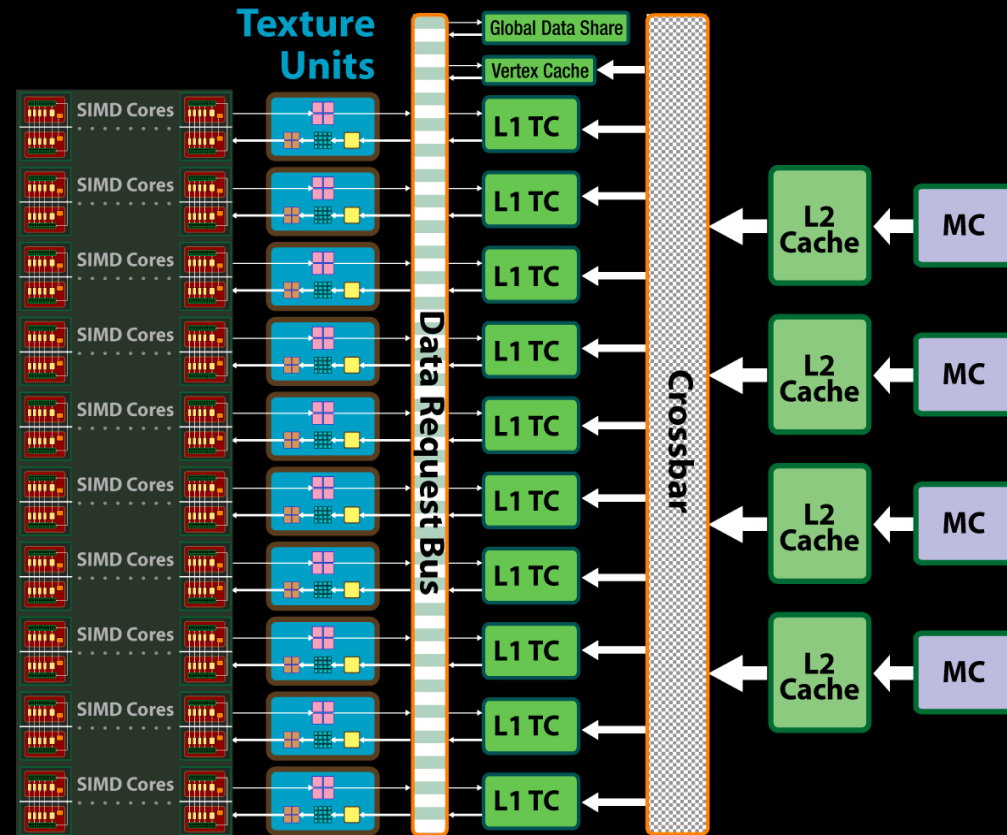
* Internal AMD test results comparing ATI Radeon™ HD 4800 series and ATI Radeon™ HD 3800 series

Texture Units



SIGGRAPH2008

- New cache design
 - L2s aligned with memory channels
 - L1s store unique data per SIMD
 - 2.5x increase aggregate L1*
 - Separate vertex cache
 - Increased bandwidth
 - Up to 480 GB/sec of L1 texture fetch bandwidth
 - Up to 384 GB/sec between L1 & L2




* Comparing ATI Radeon™ HD 4800 series and ATI Radeon™ HD 3800 series

Render Back-Ends

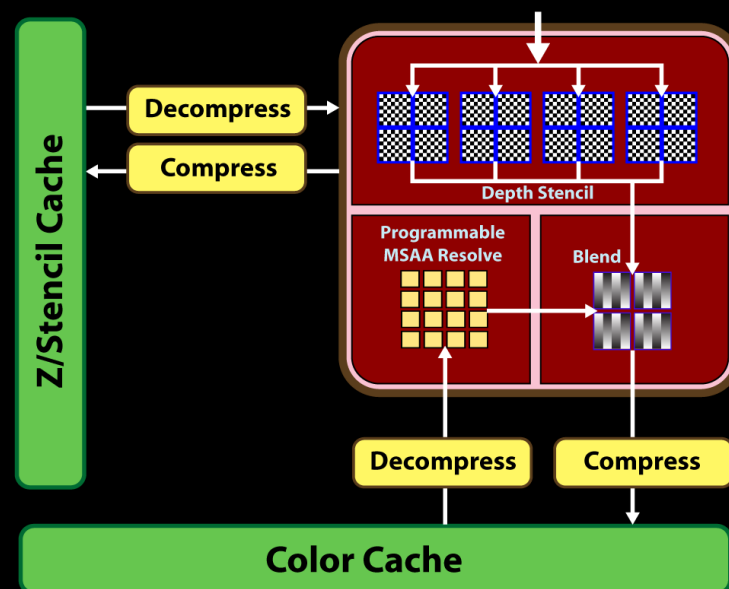


SIGGRAPH2008

- Focus on improving AA performance per mm²
 - Doubled peak rate for depth/stencil ops to 64 per clock
 - Doubled AA peak fill rate for 32-bit & 64-bit color*
 - Doubled non-AA peak fill rate for 64-bit color
- Supports both fixed function (MSAA) and programmable (CFAA) modes



Color		ATI Radeon™ HD 3800 series	ATI Radeon™ HD 4800 series	Difference
No MSAA	32-bit	16 pix/clock	16 pix/clock	1x
2x/4x MSAA		8 pix/clock	16 pix/clock	2x
8x MSAA		4 pix/clock	8 pix/clock	2x
No MSAA	64-bit	8 pix/clock	16 pix/clock	2x
2x/4x MSAA		8 pix/clock	16 pix/clock	2x
8x MSAA		4 pix/clock	8 pix/clock	2x
Depth/stencil only		32 pix/clock	64 pix/clock	2x



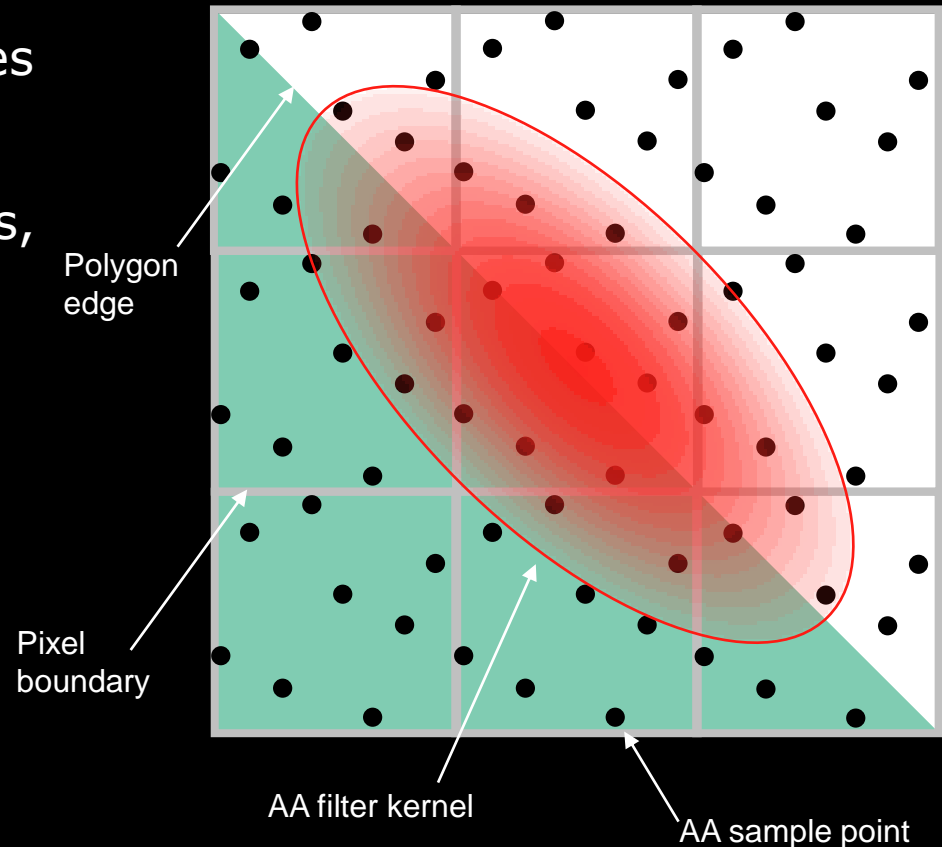
* Comparing ATI Radeon™ HD 4800 series and ATI Radeon™ HD 3800 series

Edge Detect CFAA Filters



SIGGRAPH2008

- Enhanced edge-detect filter delivers 12x & 24x CFAA modes
- Avoids blurring by taking additional samples along edges, not across them
- Same memory footprint as 4x & 8x MSAA
- Works with Adaptive AA



Edge Detect CFAA Filters



SIGGRAPH2008



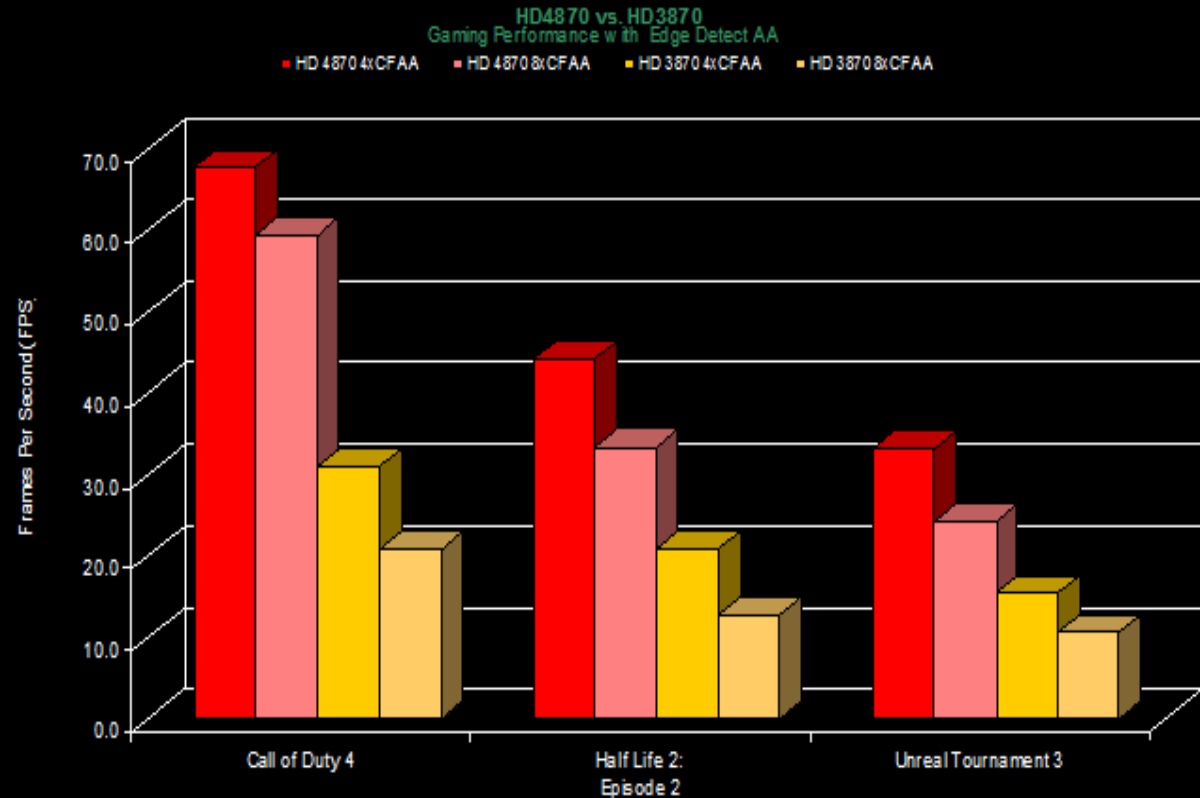
Images captured from Half-Life 2 by Valve Software

Custom Filter Anti-Aliasing Performance



SIGGRAPH2008

- Performance benefits greatly from 2x sample generation rate and 2.5x shader resolve rate
- New fast path between render back-end and shader engine provides further improvements




Test settings were 2560x1600 with 8x AF for Call of Duty 4 and Half Life 2, and 1920x1200 with 8xAF for Unreal Tournament 3. Test platform was an AMD Phenom X4 9850, with Catalyst 8.6 driver

Memory Controller Architecture



SIGGRAPH2008

- New distributed design with hub
- Controllers distributed around periphery of chip, adjacent to primary bandwidth consumers
- Memory tiling & 256-bit interface allows reduced latency, silicon area, and power consumption
- Hub handles relatively low bandwidth traffic
 - PCI Express, CrossFireX interconnect, UVD2, display controllers, intercommunication)

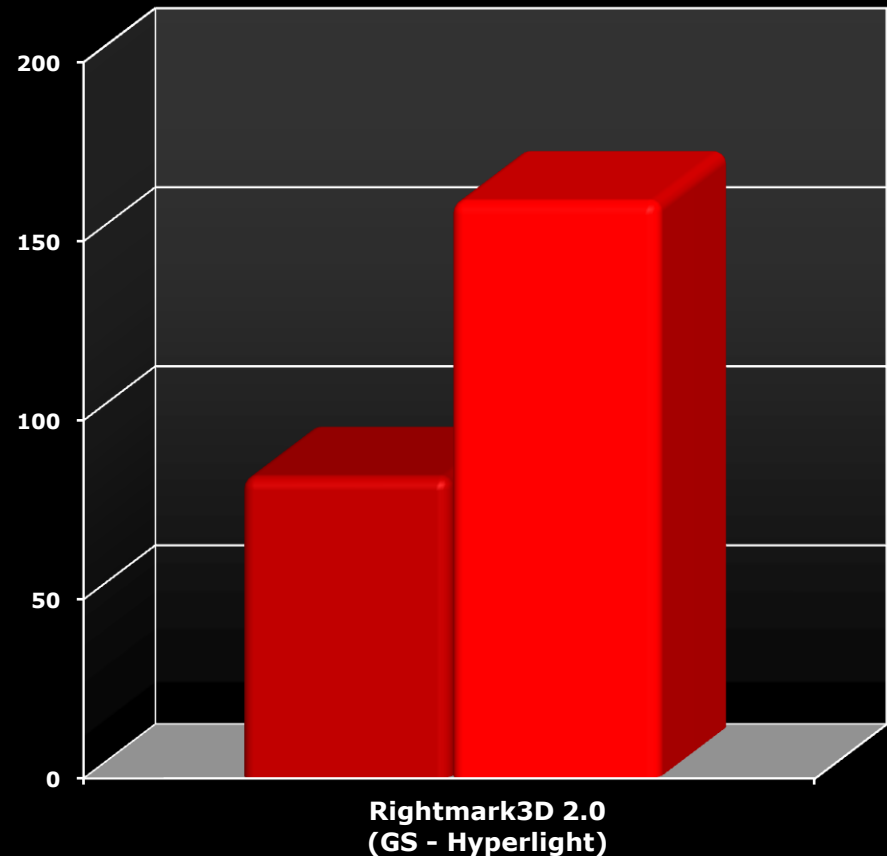


Geometry Shader & Tessellation



SIGGRAPH2008

- Enhanced geometry amplification performance over previous generation
 - Allow more GS-generated data to be kept on-chip
 - 4x more GS threads supported
- Improved tessellation unit
 - Instancing support
 - Compatible with DirectX® 10/10.1



■ ATI Radeon HD 3870

■ ATI Radeon HD 4870

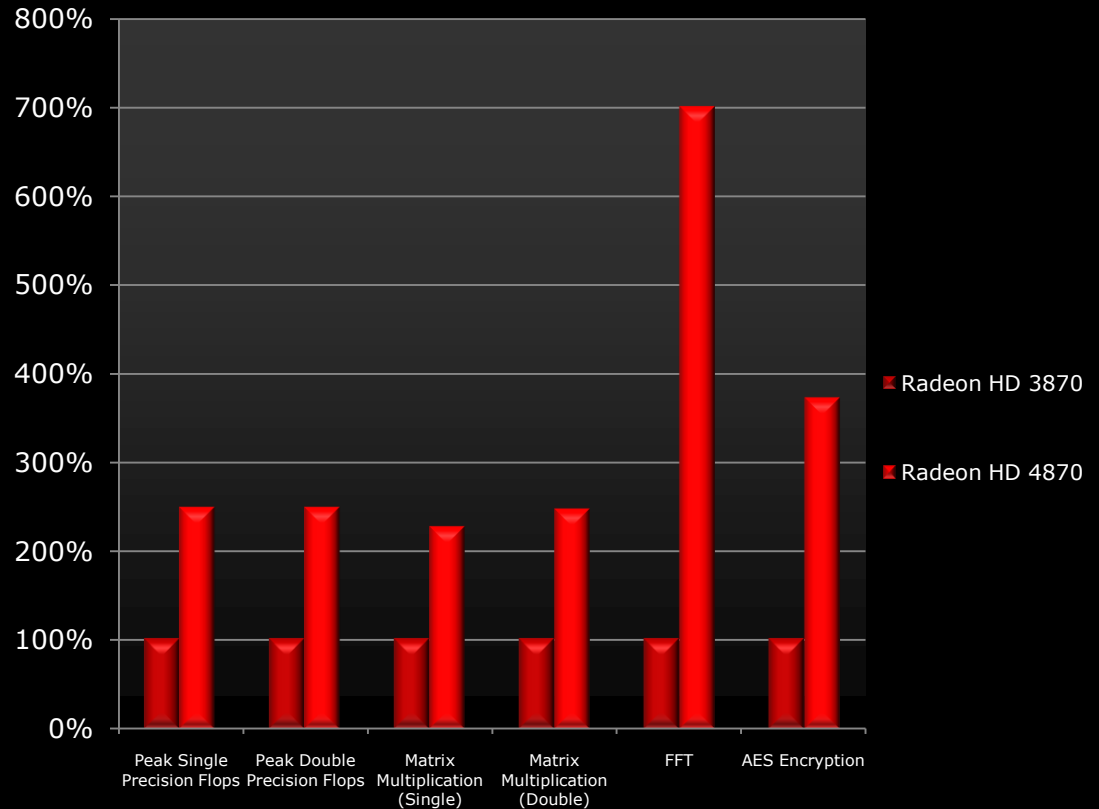
Test settings: High Polycount, Heavy Load settings, 640x480 resolution.
Config: Intel Core2 Extreme X9650 processor, using Catalyst 8.5 driver

ATI Radeon™ HD 4800 Series Stream Architecture



SIGGRAPH2008

- Several enhancements done for stream computing
 - Fast compute vector
 - Local and Global data shares
 - Fast Integer Processing
 - Fast Memexport/Memimport
- Significant increases in performance on many important stream processing workloads



Internal AMD testing, CAL SDK version 1.1, Intel QX6800 CPU, Catalyst version 8.5



ATI Radeon™ HD 4870 Computation Highlights

- >100 GB/s memory bandwidth
 - 256b GDDR5 interface
- Targeted for handling thousands of simultaneous lightweight threads
- 800 (160x5) stream processors
 - 640 (160x4) basic units (FMAC, ADD/SUB, etc.)
 - ~1.2 TFlops theoretical peak
 - 160 enhanced transcendental units (adds COS, LOG, EXP, RSQ, etc.)
 - Support for INT/UINT in all units (ADD/SUB, AND, XOR, NOT, OR, etc.)
 - 64-bit double precision FP support
 - 1/5 single precision rate (~250GFlops theoretical performance)

4 SIMDs -> 10 SIMDs

- 2.5X peak performance increase over ATI Radeon™ 3870
- ~1.2 TFlops FP32 theoretical peak
- ~250 GFlops FP64 theoretical peak

Scratch-pad memories

- 16KB per SIMD (LDS)
- 16KB across SIMDs (GDS)

Synchronization capabilities

Compute Shader

- Launch work without rasterization
- “Linear” scheduling
- Faster thread launch

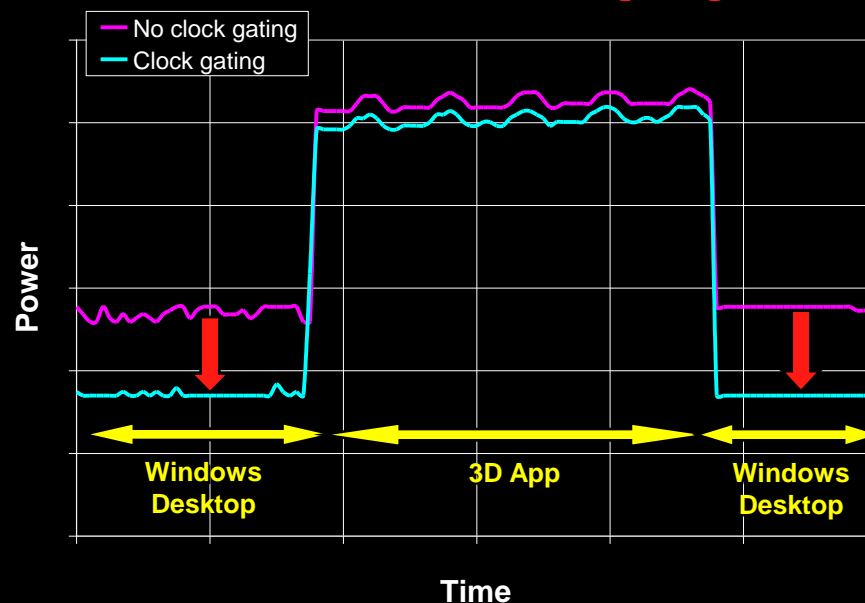
Dynamic Power Management



SIGGRAPH2008

- On-chip microcontroller
 - Constantly monitors thermals sensors and activity of various GPU blocks, PCI Express bus
 - Minimal driver overhead
- Controls clock gating, engine/memory clock speeds, voltages, and fan controller
- Enables 2x Perf/W improvement vs. ATI Radeon™ HD 3800²

Up to 36% avg. power savings from clock gating¹



¹ Internal AMD test results for ATI Radeon™ HD 4800 series

² Internal AMD test results comparing ATI Radeon™ HD 4800 series and ATI Radeon™ HD 3800 series

Terascale Graphics Have Arrived



SIGGRAPH2008

Efficient GPU design

- Major advances in Performance/Watt, Performance/\$

Improved game performance and image quality

- >2x increase in AA frame rates

Massive stream compute power

- Over 1 TeraFLOPS per GPU

Advanced feature set

- DirectX® 10.1, tessellation, CFAA, GDDR5, PCI Express® 2.0

Disclaimer and Attribution



SIGGRAPH2008

DISCLAIMER

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors.

AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION.

AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY DIRECT, INDIRECT, SPECIAL OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

ATTRIBUTION

© 2008 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, ATI, the ATI logo, CrossFireX, PowerPlay and Radeon and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other names are for informational purposes only and may be trademarks of their respective owners.