Graphics Core Next

Graphics Core Next (GCN)[1] is the codename for both a series of microarchitectures as well as for an instruction set. GCN was developed by AMD for their GPUs as the successor to TeraScale microarchitecture/instruction set. The first product featuring GCN was launched in 2011.[2]

GCN is a RISC SIMD (or rather SIMT) microarchitecture contrasting the VLIW SIMD architecture of TeraScale. GCN requires considerably more transistors than TeraScale, but offers advantages for GPGPU computation. It makes the compiler simpler and should also lead to better utilization.

GCN is fabricated in 28 nm and 14 nm graphics chips, available on selected models in the Radeon HD 7000, HD 8000, 200, 300, 400 and 500 series of AMD Radeon graphics cards. GCN is also used in the graphics portion of AMD Accelerated Processing Units (APU), such as in the PlayStation 4 and Xbox One APUs.

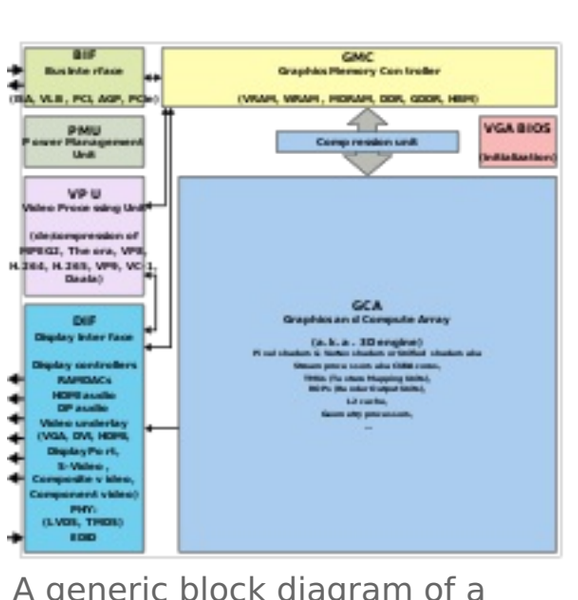

A generic block diagram of a GPU. "Graphics Core Next" shall refer to the entire GPU; hence it is possible that the same version of the GCA (the 3D engine) is combined with different versions of the DIF. AMD refers to the DIF (display interface) as DCE (display controller engine). For example, the Polaris GPUs have the same GCA/GFX as their predecessor. Strictly speaking, GCN originally referred solely to the GCA.

Contents

Instruction set

The GCN instruction set is owned by AMD as well as the X86-64 instruction set. The GCN instruction set has been developed specifically for GPUs (and GPGPU) and, for example, has no micro-operation for division.

Documentation is available:

- for Graphics Core Next instruction set (http://developer.amd.com/wordpress/media/2012/12/AMD_Southern_Islands_Instruction_Set_Architecture.pdf)
- for GCN 2nd generation instruction set (http://developer.amd.com/wordpress/media/2013/07/AMD_Sea_Islands_Instruction_Set_Architecture.pdf)
- for GCN 3rd generation instruction set (http://developer.amd.com/wordpress/media/2013/12/AMD_GCN3_Instruction_Set_Architecture_rev1.1.pdf)
- Documentation for GCN 4th generation instruction set is the same as for the 3rd generation.[3]
- for Vega's instruction set (https://developer.amd.com/wordpress/media/2017/08/Vega_Shader_ISA_28July2017.pdf)

An LLVM code generator (a compiler back-end) is available for the GCN instruction set.[4] It is used by Mesa 3D.

MIAOW (https://raw.githubusercontent.com/wiki/VerticalResearchGroup/miaow/files/MIAOW_Architecture_Whitepaper.pdf) is an open-source RTL implementation of the AMD Southern Islands GPGPU instruction set (aka Graphics Core Next).

In November 2015, AMD announced the "Boltzmann Initiative". The AMD Boltzmann Initiative shall enable the porting of CUDA-based applications to a common C++ programming model.[5]

At the "Super Computing 15" AMD showed their Heterogeneous Compute Compiler (HCC), a headless Linux driver and HSA runtime infrastructure for cluster-class, High Performance Computing (HPC) and the Heterogeneous-compute Interface for Portability (HIP) tool for porting CUDA-based applications to a common C++ programming model.

Microarchitectures

As of July 2017 the family of microarchitectures implementing the identically called instruction set "Graphics Core Next" has seen five iterations. The differences in the instruction set are rather minimal and do not differentiate too much from one another. An exception is the fifth generation GCN architecture, which heavily modified the stream processors to improve performance and support the simultaneous processing of two lower precision numbers in place of a single higher precision number.[6]

Command processing

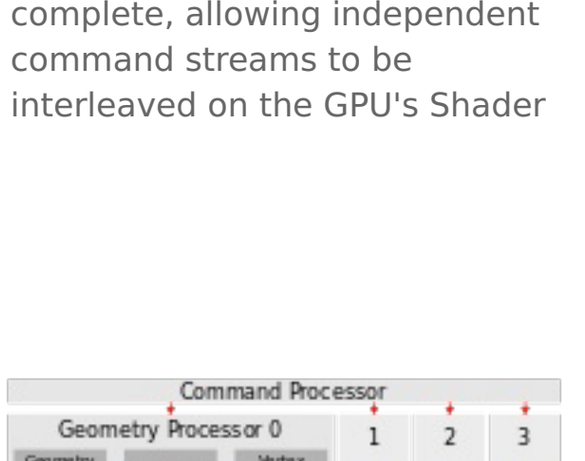

GCN command processing: Each Asynchronous Compute Engines (ACE) can parse incoming commands and dispatch work to the Compute Units (CUs). Each ACE can manage up to 8 independent queues. The ACEs can operate in parallel with the graphics command processor and two DMA engines. The graphics command processor handles graphics queues, the ACEs handle compute queues, and the DMA engines handle copy queues. Each queue can dispatch work items without waiting for other tasks to complete, allowing independent command streams to be interleaved on the GPU's Shader

Graphics Command Processor

The "Graphics Command Processor" (GCP) is a functional unit of the GCN microarchicture. Among other tasks, it is responsible for Asynchronous Shaders. The short video AMD Asynchronous Shaders (http://video.golem.de/games/15945/amd-erklaert-async-compute.html) visualizes the differences between "multi thread", "preemption" and "Asynchronous Shaders (http://amd-dev.wpengine.netdna-cdn.com/wordpress/media/2012/10/Asynchronous-Shaders-White-Paper-FINAL.pdf)[7]".

Asynchronous Compute Engine

The Asynchronous Compute Engine (ACE) is a distinct functional block serving computing purposes. It purpose is similar to that of the Graphics Command Processor.

Scheduler

Since the third iteration of GCN, the hardware contains two schedulers: One to schedule wavefronts during shader execution (CU Scheduler, see below) and a new one to schedule execution of draw and compute queues. The latter helps performance by executing compute operations when the CUs are underutilized because of graphics commands limited by fixed function pipeline speed or bandwidth limited. This functionality is known as Async Compute.

For a given shader, the gpu drivers also need to select a good instruction order, in order to minimize latency. This is done on cpu, and is sometimes referred as "Scheduling".

Geometric processor

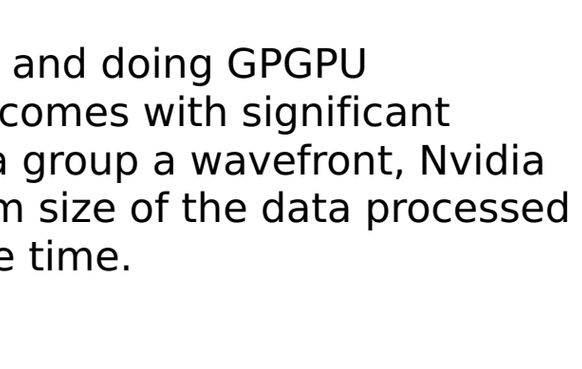

Geometry processor.

The geometry processor contains the Geometry Assembler, the Tesselator and the Vertex Assembler.

The GCN Tesselator of the Geometry processor is capable of doing tessellation in hardware as defined by Direct3D 11 and OpenGL 4.5 (see AMD January 21, 2017) in [8]).

The GCN Tesselator is AMD's most current SIP block, earlier units were ATI TruForm and hardware tessellation in TeraScale.

Compute Units

One compute unit combines 64 shader processors with 4 TMUs.[9][10] The compute unit is separate from, but feed into, the Render output units (ROPs).[10] Each Compute Unit consists of a CU Scheduler, a Branch & Message Unit, 4 SIMD Vector Units (each 16-lane wide), 4 64KiB VGPR files, 1 scalar unit, a 4 KiB GPR file, a local data share of 64 KiB, 4 Texture Filter Units, 16 Texture Fetch Load/Store Units and a 16 KiB L1 Cache. Four Compute units are wired to share an Instruction Cache 16 KiB in size and a scalar data cache 32KiB in size. These are backed by the L2 cache. A SIMD-VU operates on 16 elements at a time (per cycle), while a SU can operate on one a time (one/cycle). In addition the SU handles some other operations like branching.[11]

Every SIMD-VU has some private memory where it stores its registers. There are two types of registers: scalar registers (s0, s1, etc.), which hold 4 bytes number each, and vector registers (v0, v1, etc.), which represent a set of 64 4 bytes numbers each. When you operate on the vector registers, every operation is done in parallel on the 64 numbers. Every time you do some work with them, you actually work with 64 inputs. For example, you work on 64 different pixels at a time (for each of them the inputs are slightly different, and thus you get slightly different color at the end).

Every SIMD-VU has room for 512 scalar registers and 256 vector registers.

CU Scheduler

The CU scheduler is the hardware functional block choosing for the SIMD-VU which wavefronts to execute. It picks one SIMD-VU per cycle for scheduling. This is not to be confused with other schedulers, in hardware or software.

Wavefront

> A 'shader' is a small program written in GLSL which performs graphics processing, and a 'kernel' is a small program written in OpenCL and doing GPGPU processing. These processes don't need that many registers, they need to load data from system or graphics memory. This operation comes with significant latency. AMD and Nvidia chose similar approaches to hide this unavoidable latency: the grouping of multiple threads. AMD calls such a group a wavefront, Nvidia calls it a warp. A group of threads is the most basic unit of scheduling of GPUs implementing this approach to hide latency, is minimum size of the data processed in SIMD fashion, the smallest executable unit of code, the way to processes a single instruction over all of the threads in it at the same time.

In all GCN-GPUs, a “wavefront” consists of 64 threads, and in all Nvidia GPUs a “warp” consists of 32 threads.

AMD's solution is to attribute multiple wavefronts to each SIMD-VU. The hardware distributes the registers to the different wavefronts, and when one wavefront is waiting on some result, which lies in memory, the CU Scheduler decides to make the SIMD-VU work on another wavefront. Wavefronts are attributed per SIMD-VU. SIMD-VUs do not exchange wavefronts. At max 10 wavefronts can be attributed per SIMD-VU (thus 40 per CU).

AMD CodeXL shows tables with the relationship between number of SGPRs and VGPRs to the number of wavefronts, but basically for SGPRS it is min(104, 512/numwavefronts) and VGPRS 256/numwavefronts.

Note that in conjunction with the SSE instructions this concept of most basic level of parallelism is often called a "vector width". The vector width is characterized by the total number of bits in it.

SIMD Vector Unit

Each SIMD Vector Unit has:

- a 16-lane integer and floating point vector Arithmetic Logic Unit (ALU)
- 64 KiB Vector General Purpose Register (VGPR) file
- A 48-bit Program Counter
- Instruction buffer for 10 wavefronts
  - A wavefront is a group of 64 threads: the size of one logical VGPR
- A 64-thread wavefront issues to a 16-lane SIMD Unit over four cycles

Each SIMD-VU has 10 wavefront instruction buffer, and it takes 4 cycles to execute one wavefront.

Audio and video acceleration SIP blocks

Many implementations of GCN are typically accompanied by several of AMD's other ASIC blocks. Including but not limited to the Unified Video Decoder, Video Coding Engine, and AMD TrueAudio.

Video Coding Engine

TrueAudio

Unified virtual memory

In a preview in 2011, AnandTech wrote about the unified virtual memory, supported by Graphics Core Next.[12]

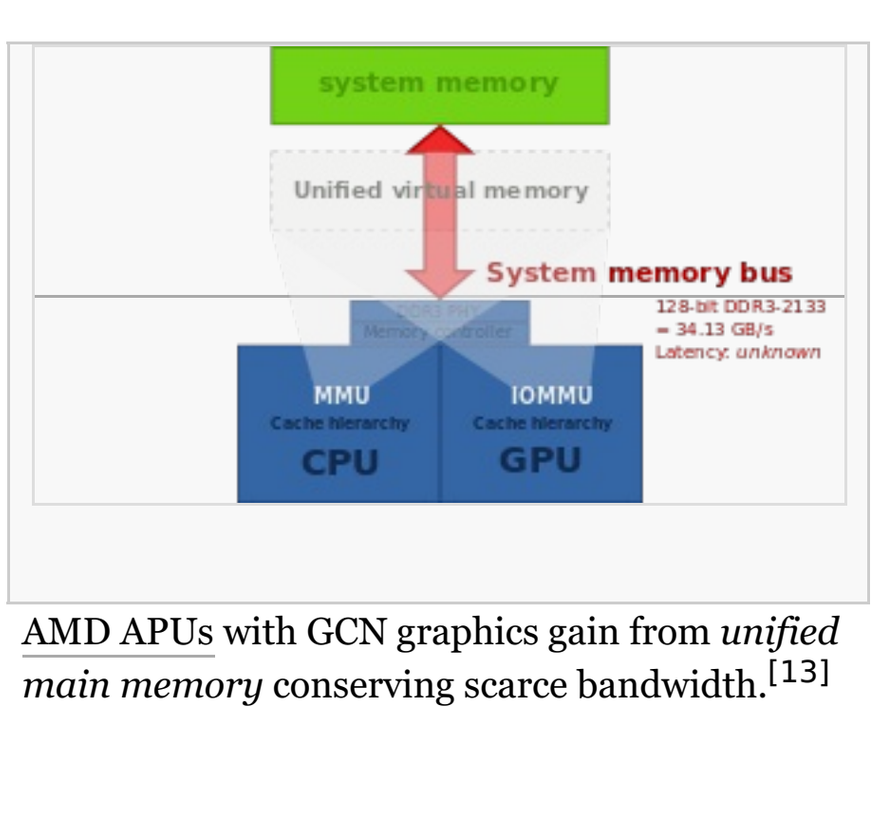

Classical desktop computer architecture with a distinct graphics card over PCI Express. CPU and GPU have their distinct physical memory, with different address spaces. The entire data needs to be copied over the PCIe bus. Note: the diagram shows bandwidths, but not the memory latency.

GCN supports "unified virtual memory", hence enabling zero-copy, instead of the data, only the pointers are copied, "passed". This is a paramount HSA feature.

Integrated graphics-solutions (and AMD APUs with TeraScale graphics) suffer under *partitioned main memory*: a part of the system memory is allocated to the GPU exclusively. Zero-copy is not possible, data has to be copied (over the system memory bus) from one partition to the other.

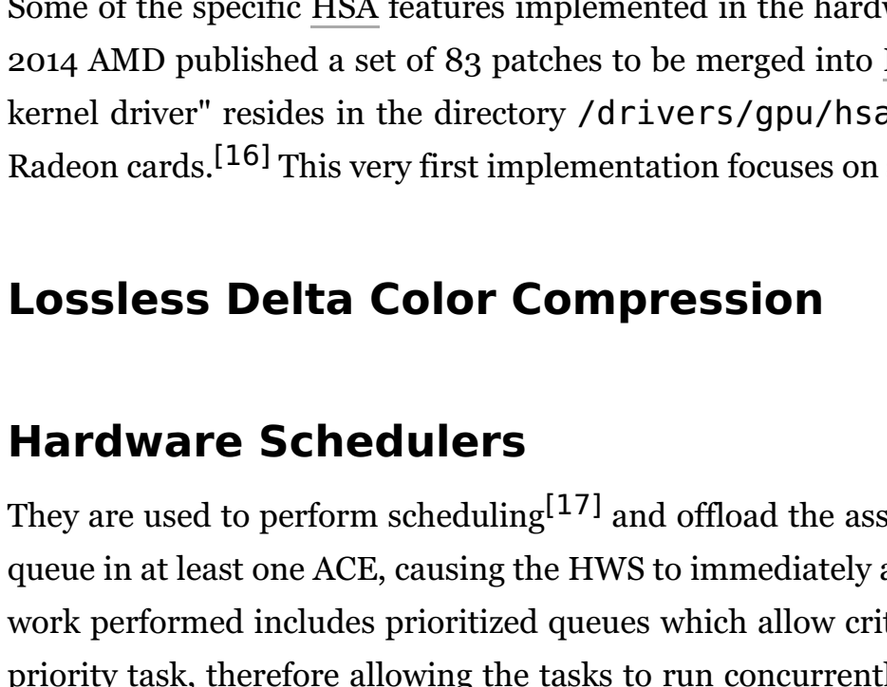

AMD APUs with GCN graphics gain from *unified main memory* conserving scarce bandwidth.[13]

Heterogeneous System Architecture (HSA)

GCN includes special purpose function blocks to be used by HSA. Support for these function blocks is available through amdkfd since Linux kernel 3.19.[14]

Some of the specific HSA features implemented in the hardware need support from the operating system's kernel (its subsystems) and/or from specific device drivers. For example, in July 2014 AMD published a set of 83 patches to be merged into Linux kernel mainline 3.17 for supporting their Graphics Core Next-based Radeon graphics cards. The special driver titled "HSA kernel driver" resides in the directory `/drivers/gpu/hsa` while the DRM-graphics device drivers reside in `/drivers/gpu/drm`[15] and augments the already existent DRM driver for Radeon cards.[16] This very first implementation focuses on a single "Kaveri" APU and works alongside the existing Radeon kernel graphics driver (kgd).

Lossless Delta Color Compression

Hardware Schedulers

They are used to perform scheduling[17] and offload the assignment of compute queues to the ACEs from the driver to hardware by buffering these queues until there is at least one empty queue in at least one ACE, causing the HWS to immediately assign buffered queues to the ACEs until all queues are full or there are no more queues to safely assign.[18] Part of the scheduling work performed includes prioritized queues which allow critical tasks to run at a higher priority than other tasks without requiring the lower priority tasks to be preempted to run the high priority task, therefore allowing the tasks to run concurrently with the high priority tasks scheduled to hog the GPU as much as possible while letting other tasks use the resources that the high priority tasks are not using.[17] These are essentially Asynchronous Compute Engines that lack dispatch controllers.[17] They were first introduced in the fourth generation GCN microarchitecture,[17] but were present in the third generation GCN microarchitecture for internal testing purposes.[19] A driver update has enabled the hardware schedulers in third generation GCN parts for production use.[17]

Primitive Discard Accelerator

This unit discards degenerate triangles before they enter the vertex shader and triangles that do not cover any fragments before they enter the fragment shader.[20] This unit was introduced with the fourth generation GCN microarchitecture.[20]

Iterations

Graphics Core Next (Southern Islands)

- Support for 64-bit addressing (x86-64 address space) with unified address space for CPU and GPU[12]
  - Support for PCI-E 3.0[21]
  - GPU sends interrupt requests to CPU on various events (such as page faults)
- Support for Partially Resident Textures,[22] which enable virtual memory support through DirectX and OpenGL extensions
- AMD PowerTune support, which dynamically adjusts performance to stay within a specific TDP[23]
- Support for Mantle (API)

There are Asynchronous Compute Engines controlling computation and dispatching.[11][24]

ZeroCore Power

ZeroCore Power is a long idle power saving technology, shutting off functional units of the GPU when not in use.[25] AMD ZeroCore Power technology supplements AMD PowerTune.

Chips

Discrete GPUs (Southern Islands family):

- Oland
- Cape Verde
- Pitcairn
- Tahiti

GCN 2nd Generation (Sea Islands)

AMD PowerTune "Bonaire"

GCN 2nd generation was introduced with Radeon HD 7790 and is also found in Radeon HD 8770, R7 260/260X, R9 290/290X, R9 295X2, R7 360, R9 390/390X, as well as Steamroller-based Desktop Kaveri APUs and Mobile Kaveri APUs and in the Puma-based "Beema" and "Mullins" APUs. It has multiple advantages over the original GCN, including FreeSync support, AMD TrueAudio and a revised version of AMD PowerTune technology.

GCN 2nd generation introduced an entity called "Shader Engine" (SE). A Shader Engine comprises one geometry processor, up to 11 CUs (Hawaii chip), rasterizers, ROPs, and L1 cache. Not part of a Shader Engine is the Graphics Command Processor, the 8 ACEs, the L2 cache and memory controllers as well as the audio and video accelerators, the display controllers, the 2 DMA controllers and the PCIe interface.

The A10-7850K "Kaveri" contains 8 CUs (compute units) and 8 Asynchronous Compute Engines for independent scheduling and work item dispatching.[26]

At AMD Developer Summit (APU) in November 2013 Michael Mantor presented the Radeon R9 290X.[27]

Chips

Discrete GPUs (Sea Islands family):

- Bonaire
- Hawaii

Integrated into APUs:

- Temash
- Kabini
- Liverpool (i.e. the APU found of the Playstation 4)
- Durango (i.e. the APU found of the Xbox One and Xbox One S)
- Kaveri
- Godavari
- Mullins
- Beema
- Carrizo-L

GCN 3rd Generation (Volcanic Islands)

GCN 3rd generation[28] was introduced in 2014 with the Radeon R9 285 and R9 M295X, which have the "Tonga" GPU. It features improved tessellation performance, lossless delta color compression in order to reduce memory bandwidth usage, an updated and more efficient instruction set, a new high quality scaler for video, and a new multimedia engine (video encoder/decoder). Delta color compression is supported in Mesa.[29] However, its double precision performance is worse compared to previous generation.[30]

Chips

Discrete GPUs:

- Tonga (Volcanic Islands family), comes with UVD 5.0 (Unified Video Decoder)
- Fiji (Pirate Islands family), comes with UVD 6.0 and High Bandwidth Memory (HBM 1)

Integrated into APUs:

- Carrizo, comes with UVD 6.0
- Bristol Ridge[31]
- Stoney Ridge[31]

GCN 4th Generation (Arctic Islands)

GPUs of the Arctic Islands-family were introduced in Q2 of 2016 with AMD Radeon 400 series branded graphics cards, based upon the Polaris architecture. All Polaris-based chips are produced on the 14 nm FinFET process.[32] The fourth generation GCN instruction set architecture is compatible with the third generation. It is an optimization for 14 nm FinFET process enabling higher GPU clock speeds than with the 3rd GCN generation.[33]Architectural improvements include new hardware schedulers, a new primitive discard accelerator, a new display controller, and an updated UVD that can decode HEVC at 4K resolutions at 60 frames per second with 10 bits per color channel.

Chips

Discrete GPUs:[34]

- Polaris 10 (also codenamed Ellesmere) found on "Radeon RX 470"- and "Radeon RX 480"-branded graphics cards
- Polaris 11 (also codenamed Baffin) found on "Radeon RX 460"-branded graphics card (also Radeon RX 560**D**).
- Polaris 12 found on "Radeon RX 550" and "Radeon RX 540"-branded graphics cards.
- Polaris 20, which is a refreshed(14nm LPP process) Polaris 10 with higher clocks, used for "Radeon RX 570" and "Radeon RX 580"-branded graphic cards.[35]
- Polaris 21, which is a refreshed(14nm LPP process) Polaris 11, used for "Radeon RX 560"-branded graphics card.
- Polaris 22, found on "Radeon RX Vega M GH" and "Radeon RX Vega M GL"-branded graphics.

GCN 5th Generation (Vega)

AMD began releasing details of their next generation of GCN Architecture, termed the 'Next-Generation Compute Unit', in January 2017.[33][36][37] The new design is expected to increase instructions per clock, higher clock speeds, support for HBM2, a larger memory address space. The discrete graphics chipsets also include High Bandwidth Cache Controller, but not when integrated into APUs.[38] Additionally, the new chips are expected to include improvements in the Rasterisation and Render output units. The stream processors are heavily modified from the previous generations to support packed math Rapid Pack Math technology for 8-bit, 16-bit, and 32-bit numbers. With this there is a significant performance advantage when lower precision is acceptable (for example: processing two half-precision numbers at the same rate as a single single precision number).

Nvidia introduced tile-based rasterization and binning with Maxwell,[39] and this was a big reason for Maxwell's efficiency increase. In January, AnandTech assumed that Vega would finally catch up with Nvidia regarding energy efficiency optimizations due to the new "Draw Stream Binning Rasterizer" to be introduced with Vega.[40]

It also added support for a new shader stage - primitive shaders.[41][42] Primitive shaders provide more flexible geometry processing and replace the vertex and geometry shaders in a rendering pipeline.[43]

Chips

Discrete GPUs:

- Vega 10 found on "Radeon RX Vega 64"- and "Radeon RX Vega 56"-branded graphics cards[44]

Integrated into APUs:

- Raven Ridge[45] - Came with VCN 1 which supersedes VCE and UVD and allows full fixed-function VP9 decode

GCN 6th Generation (Navi)

Navi is expected in 2019 and will offer "Next Generation Memory" as well as improved scalability.[46]

See also

- List of AMD graphics processing units

References

1. AMD Developer Central (January 31, 2014). "GS-4106 The AMD GCN Architecture - A Crash Course, by Layla Mah" (https://www.slideshare.net/DevCentralAMD/gs4106-the-amd-gcn-architecture-a-crash-course-by-layla-mah).
2. "AMD Launches World's Fastest Single-GPU Graphics Card – the AMD Radeon HD 7970" (https://www.amd.com/en-us/press-releases/Pages/amd-launches-worlds-fastest-2011dec22.aspx) (Press release). AMD. December 22, 2011. Retrieved January 20, 2015.
3. "AMD community forums" (https://community.amd.com/thread/203070).
4. "LLVM back-end amdgpu" (http://llvm.org/docs/AMDGPUUsage.html).
5. "AMD Boltzmann Initiative – Heterogeneous-compute Interface for Portability (HIP)" (https://www.amd.com/en-us/press-releases/Pages/boltzmann-initiative-2015nov16.aspx). November 16, 2015. Retrieved January 15, 2016.
6. Smith, Ryan (January 5, 2017). "The AMD Vega GPU Architecture Preview" (http://www.anandtech.com/show/11002/the-amd-vega-gpu-architecture-teaser). Anandtech. Retrieved July 11, 2017.
7. Smith, Ryan. "AMD Dives Deep On Asynchronous Shading" (http://www.anandtech.com/show/9124/amd-dives-deep-on-asynchronous-shading).
8. "The Khronos Group" (https://www.khronos.org/conformance/adopters/conformant-products#opengl). October 26, 2017.
9. *Compute Cores Whitepaper* (https://www.amd.com/Documents/Compute_Cores_Whitepaper.pdf) (PDF). AMD. 2014. p. 5.
10. Smith, Ryan (December 21, 2011). "AMD's Graphics Core Next Preview" (http://www.anandtech.com/show/4455/amds-graphics-core-next-preview-amd-architects-for-compute/5). Anandtech. Retrieved April 18, 2017.
11. Mantor, Michael; Houston, Mike (June 15, 2011). "AMD Graphics Core Next" (http://developer.amd.com/wordpress/media/2013/06/2620_final.pdf) (pdf). AMD. p. 40. Retrieved July 15, 2014. "Asynchronous Compute Engine (ACE)"
12. "Not Just A New Architecture, But New Features Too" (http://www.anandtech.com/show/4455/amds-graphics-core-next-preview-amd-architects-for-compute/6). AnandTech. December 21, 2011. Retrieved July 11, 2014.
13. "Kaveri microarchitecture" (http://semiaccurate.com/2014/01/15/technical-look-amds-kaveri-architecture/). *SemiAccurate*. January 15, 2014.
14. Dave Airlie (November 26, 2014). "Merge AMDKFD" (http://cgit.freedesktop.org/~airlied/linux/commit/?h=drm-next&id=33f86ff62c4368c8d6bf3c76dc2fa416e3f90213). freedesktop.org. Retrieved January 21, 2015.
15. "/drivers/gpu/drm" (https://git.kernel.org/cgit/linux/kernel/git/stable/linux-stable.git/tree/drivers/gpu/drm). kernel.org.
16. "[PATCH 00/83] AMD HSA kernel driver" (http://lkml.iu.edu/hypermail/linux/kernel/1407.1/02923.html). LKML. July 10, 2014. Retrieved July 11, 2014.
17. Angelini, Chris (June 29, 2016). "AMD Radeon RX 480 8GB Review" (http://www.tomshardware.com/reviews/amd-radeon-rx-480-polaris-10,4616.html). *Tom's Hardware*. p. 1. Retrieved August 11, 2016.
18. "Dissecting the Polaris Architecture" (http://radeon.wpengine.netdna-cdn.com/wp-content/uploads/2016/08/Polaris-Architecture-Whitepaper-Final-08042016.pdf) (PDF). 2016. Retrieved August 12, 2016.
19. Shrout, Ryan (June 29, 2016). "The AMD Radeon RX 480 Review - The Polaris Promise" (http://www.pcper.com/reviews/Graphics-Cards/AMD-Radeon-RX-480-Review-Polaris-Promise/Polaris-Architecture-4th-Generation-). *PC Perspective*. p. 2. Retrieved August 12, 2016.
20. Smith, Ryan (June 29, 2016). "The AMD Radeon RX 480 Preview: Polaris Makes Its Mainstream Mark" (http://www.anandtech.com/show/10446/the-amd-radeon-rx-480-preview/3). *AnandTech*. p. 3. Retrieved August 11, 2016.
21. "AMD Radeon HD 7000 Series to be PCI-Express 3.0 Compliant" (http://www.techpowerup.com/149225/AMD-Radeon-HD-7000-Series-to-be-PCI-Express-3.0-Compliant.html). TechPowerUp. Retrieved July 21, 2011.
22. "AMD Details Next Gen. GPU Architecture" (http://pixelsmashers.com/?p=10341). Retrieved August 3, 2011.
23. Tony Chen, Jason Greaves, "AMD's Graphics Core Next (GCN) Architecture" (http://meseec.ce.rit.edu/551-projects/fall2014/3-4.pdf) (PDF), *AMD*, retrieved August 13, 2016
24. "AMD's Graphics Core Next Preview: AMD's New GPU, Architected For Compute" (http://www.anandtech.com/show/4455/amds-graphics-core-next-preview-amd-architects-for-compute/5). AnandTech. December 21, 2011. Retrieved July 15, 2014. "AMD's new Asynchronous Compute Engines serve as the command processors for compute operations on GCN. The principal purpose of ACEs will be to accept work and to dispatch it off to the CUs for processing."
25. "Managing Idle Power: Introducing ZeroCore Power" (http://www.anandtech.com/show/5261/amd-radeon-hd-7970-review/11). AnandTech. December 22, 2011. Retrieved April 29, 2015.
26. "AMD's Kaveri A10-7850K tested" (http://www.anandtech.com/show/7677/amd-kaveri-review-a8-7600-a10-7850k/7). *AnandTech*. January 14, 2014. Retrieved July 7, 2014.
27. "AMD Radeon R9-290X" (http://www.slideshare.net/DevCentralAMD/gs4152-michael-mantor). November 21, 2013.
28. http://images.anandtech.com/doci/9319/Slide%2019%20-%20GCN%20Overview.png
29. "Add DCC Support" (http://lists.freedesktop.org/archives/mesa-dev/2015-October/096919.html). *Freedesktop.org*. October 11, 2015.
30. Smith, Ryan (September 10, 2014). "AMD Radeon R9 285 Reveiw" (http://www.anandtech.com/show/8460/amd-radeon-r9-285-review). Anandtech. Retrieved March 13, 2017.
31. Cutress, Ian (June 1, 2016). "AMD Announces 7th Generation APU" (http://www.anandtech.com/show/10362/amd-7th-generation-apu-bristol-ridge-stoney-ridge-for-notebooks). Anandtech.com. Retrieved June 1, 2016.
32. "Radeon Technologies Group - January 2016 - AMD Polaris Architecture" (http://www.guru3d.com/articles-pages/radeon-technologies-group-january-2016-amd-polaris-architecture,1.html).
33. Smith, Ryan (January 5, 2017). "The AMD Vega Architecture Teaser: Higher IPC, Tiling, & More, coming in H1'2017" (http://www.anandtech.com/show/11002/the-amd-vega-gpu-architecture-teaser). Anandtech.com. Retrieved January 10, 2017.
34. WhyCry (March 24, 2016). "AMD confirms Polaris 10 is Ellesmere and Polaris 11 is Baffin" (http://videocardz.com/58634/amd-confirms-polaris-10-is-ellesmere-and-polaris-11-is-baffin). *VideoCardz*. Retrieved April 8, 2016.
35. "Fast vollständige Hardware-Daten zu AMDs Radeon RX 500 Serie geleakt" (https://www.3dcenter.org/news/fast-vollstaendige-hardware-daten-zu-amds-radeon-rx-500-serie-geleakt). *www.3dcenter.org*.
36. Kampman, Jeff (January 5, 2017). "The curtain comes up on AMD's Vega architecture" (http://techreport.com/review/31224/the-curtain-comes-up-on-amd-vega-architecture). TechReport.com. Retrieved January 10, 2017.
37. Shrout, Ryan (January 5, 2017). "AMD Vega GPU Architecture Preview: Redesigned Memory Architecture" (https://www.pcper.com/reviews/Graphics-Cards/AMD-Vega-GPU-Architecture-Preview-Redesigned-Memory-Architecture). PC Perspective. Retrieved January 10, 2017.
38. Kampman, Jeff (October 26, 2017). "AMD's Ryzen 7 2700U and Ryzen 5 2500U APUs revealed" (https://techreport.com/review/32743/amd-ryzen-7-2700u-and-ryzen-5-2500u-apus-revealed). Techreport. Retrieved October 26, 2017.
39. Raevenlord (March 1, 2017). "On NVIDIA's Tile-Based Rendering" (https://www.techpowerup.com/231129/on-nvidias-tile-based-rendering). *techPowerUp*.
40. "Vega Teaser: Draw Stream Binning Rasterizer" (http://www.anandtech.com/show/11002/the-amd-vega-gpu-architecture-teaser/3).
41. "Radeon RX Vega Revealed: AMD promises 4K gaming performance for $499 - Trusted Reviews" (http://www.trustedreviews.com/news/amd-vega-specs-performance-release-date-technology-explained). July 31, 2017.
42. "The curtain comes up on AMD's Vega architecture" (http://techreport.com/review/31224/the-curtain-comes-up-on-amd-vega-architecture/2).
43. "The curtain comes up on AMD's Vega architecture" (http://techreport.com/review/31224/the-curtain-comes-up-on-amd-vega-architecture/2).
44. "The AMD Radeon RX Vega 64 & RX Vega 56 Review: Vega Burning Bright" (https://www.anandtech.com/show/11717/the-amd-radeon-rx-vega-64-and-56-review). Anandtech.com. August 14, 2017. Retrieved November 16, 2017.
45. Ferreira, Bruno (May 16, 2017). "Ryzen Mobile APUs are coming to a laptop near you" (https://techreport.com/news/31918/ryzen-mobile-apus-are-coming-to-a-laptop-near-you). Tech Report. Retrieved May 16, 2017.
46. Usman Pirzada (January 8, 2018). "AMD Radeon Graphics Roadmap Update: 7nm Vega in 2018, Navi in 2019, Followed By "Next-Gen", New Vega Instinct Announced" (https://wccftech.com/amd-radeon-update-7nm-vega-2018-navi-2019/). *WCCFtech*.