Heterogeneous system architecture pdf

The Heterogeneous System Architecture (HSA) is a cross-supplier set of specifications that allow for the integration of central processing processors and GPUs on the same bus, with shared memory and tasks. HSA is developed by the HSA Foundation, which includes (among many others) AMD and ARM. The stated goal of the platform is to reduce the delay in communication between processors, GPUs, and other computing devices and to make these different devices more compatible from the programmer's point of view, freeing the programmer from the task of scheduling the movement of data between disparate device memories (as should be done now with OpenCL or CUDA). CUDA and OpenCL, as well as most other fairly advanced programming languages, can use HSA to improve performance. Heterogeneous computing is widely used in chip system devices such as tablets, smartphones, other mobile devices, and game consoles. HSA allows programs to use a GPU to calculate floating currents without separate memory or planning. The rationale behind the HSA is to ease the burden on programmers when unloading calculations in the GPU. Originally driven exclusively by AMD and called the FSA, the idea has been expanded to cover processors other than GPUs such as other manufacturers' DSPs as well. Steps performed when unloading calculations in THED on non-HSA system steps performed when unloading calculations in the GPU on the HSA system, using HSA functionality, Modern GPUs are very well suited to one instruction, multiple data (SIMD) and one manual, multiple threads (SIMT), while modern processors are still optimized for branching. Etc. Review This section requires additional citations to verify. Please help improve this article by adding quotes to reliable sources. Non-sources of materials can be challenged and removed. (May 2014) (Learn how and when to delete this template message) Originally introduced by built-in systems such as Cell Broadband Engine, sharing system memory directly between multiple pieces of the system makes heterogeneous computing more mainstream. Heterogeneous computing in itself refers to systems that contain multiple processing units - central processing units (processors), GPUs (GPUs), digital signal processors (DSPs), or any type of application of specific integrated circuits (ASICs). The system's architecture allows any accelerator, such as a GPU, to operate at the same processing level as the system's processor. Among its main features, the HSA identifies a unified virtual address space for Devices: Where GPUs traditionally have their own memory, separate from the main (CPU) memory, the HSA requires these devices to share page tables so that devices can share data by exchanging pointers. This should be supported by custom memory control In order to make compatibility possible, as well as to facilitate various aspects of programming, HSA is designed for ISA agnostics for both processors and accelerators, as well as to support high-level programming languages. So far, the HSA specifications cover: HSA Intermediate Layer HSA Intermediate Layer (HSAIL), a virtual instruction set for parallel programs, similar according to whom? for LLVM intermediate representation and SPIR (used by OpenCL and Vulkan) is completed before the specific instruction set by the JIT compiler to make late decisions about which cores (s) should work the task explicitly supports exceptions, virtual functions and other high-level debugging features support the HSA memory model compatible with the C-11 model, OpenCL memory models, Java and .NET, softened coherences designed to support both managed languages (e.g. Java) and unmanageable languages (e.g. C), will greatly simplify the development of third-party compilers for a wide range of heterogeneous products programmed in Fortran, PH, AMP, Java, etc. Including the very significant reduction in overhead planning costs for major mobile devices are one of the HSA applications areas in which it provides energy efficiency improvements. The block chart diagrams below provide high-level illustrations of how HSA works and how it compares to traditional architectures. Standard architecture with a discrete GPU attached to the PCI Express bus. A zero copy between the GPU and the processor is not possible due to various physical memories. HSA brings a single virtual memory and makes it easier to transfer PCI Express pointers instead of copying all the data. In a divided core memory, one part of the system's memory is allocated exclusively to the GPU. As a result, a zero-copy operation is not possible. A single basic memory can be made possible by combining a GPU and an HSA-enabled processor. As a result, you can perform zero-copy operations. Both the MMU processor and the IOMMU GPU must meet the specifications of the HSA equipment. AMD software support GPUs contain certain additional functional units designed for use within the HSA. In Linux, the amdkfd kernel driver provides the support you need. Some of the HSA features implemented in the hardware must be supported by the core of the operating system and specific device drivers. For example, support for AMD Radeon and AMD FirePro graphics cards, as well as Graphics Core Next (GCN) graphics processors, was combined into version 3.19 of the main line Linux, released on February 8, 2015. Programs don't interact directly with amdkfd, but line up for their jobs using HSA time. This first implementation, known as amdkfd, is dedicated to the Kaveri or Berlin Berlin API works alongside the existing Radeon core graphics driver. In addition, amdkfd supports heterogeneous queues (headquarters), which aims to simplify the distribution of computing jobs between multiple processors and GPUs from the programmer's point of view. The Mixed Memory Management (HMM) support, suitable only for graphics hardware with version 2 of IOMMU AMD, has been accepted into the main version of the Linux 4.14 kernel. Comprehensive support for HSA platforms for the release of OpenJDK Sumatra has been announced, which is due to take place in 2015. AMD APP SDK is an AMD-branded software development suite focused on parallel computing available for Microsoft Windows and Linux. The Bolt is a C's template library optimized for heterogeneous computing. GPUOpen understands several other software related to HSA. CodeXL version 2.0 includes the HSA profiler. As of February 2015, only AMD Kaveri A-series (cf. Kaveri desktop processors and Kaveri mobile processors) and Sony PlayStation 4 allowed the integrated GPU to access memory through AMD's 2th version of THE IOMMU. Previously, vsUS (Trinity and Richland) included the functionality of version 2 of IOMMU, but only for the use of an external GPU connected through PCI Express. After 2015, the Carrizo and Bristol Ridge APIs also include 2 IOMMU functionality for integrated GPU. (quote is necessary) The following table shows the features of the AMD API (see also: AMD Fast Processing Units List). - VisualEditor - viewtalkedit Code Title Server Major Toronto Micro Kyoto Desktop Carrizo Bristol Ridge Raven Ridge Picasso Renoir Entry Llano Trinity Richland Caberi Basic Cabins Mobile Performance Renoir Mainstream Llano Trinity Richland Carriso Bristol Ridge Raven Ridge Picasso Entry Mullins Carrizo-L Sunset Cabins Steppe Eagle, Crowned Eagle, LX-Family Prairie Falcon Banded Kestrel Platform High, Standard and Low Power Low and Ultra-Low Power Released August 2011 October 2012 June 2013 January 2014 June 2015 June 2016 Oct 2017 January 2019 March 2010 April 13 2013 2014 May 2015 February 2016 April 2019 CPU MicroArchitecture K10 Piledriver Steamroller Excavator zen ISA x86-64 x86-64 Outlet High-End N/A N/A Mainstream N/A AM4 Entry FM1 FM2A N/A Basic N/A N/A AM1 N/ Another FS1 FS1 , FP2 FP3 FP4 FP5 FP6 FT1 FT3b FP4 FP5 PCI Express version 2.0 3.0 2.0 3.0 Fab. (nm) GF 32SHP (HKMG SOI) GF 28SHP (HKMG bulk) GF 14LPP (FinFET bulk) GF 12LP (FinFET bulk) TSMC N7 (FinFET bulk) T SMC N40 (bull) TSMC N28 (HKMG bulk) GF 28SHP (HKMG bulk) GF 14LPP (FinFET bulk) Die area (mm2) 228 245 245 250 210 18 156 75 (No 28 FCH) 107 ? 125 149 Мин TDP TDP 35 17 12 10 4.5 4 3.95 10 6 Max APU TDP (W) 100 95 65 18 25 Max stock APU base clock (GHz) 3 3.8 4.1 4.1 3.7 3.8 3.6 3.7 3.8 1.75 2.2 2 2.2 3.2 3.3 Max APUs per node[b] 1 1 Max CPU[c] cores per APU 4 8 2 4 2 Max threads per CPU core 1 2 1 2 Integer structure 3+3 2+2 4+2 4+2+1 1+1+1+1 2+2 4+2 i386, i486, i586, CMOV, NOPL, i686, PAE, NX bit, CMPXCHG16B, AMD-V, RVI, ABM, and 64-bit LAHF/SAHF IOMMU[d] N/A BMI1, AES-NI, CLMUL, and F16C N/A MOVBE N/A AVIC, BMI2 and RDRAND N/A ADX , SHA, RDSEED, SMAP, SMEP, XSAVEC, XSAVES, XRSTORS, CLFLUSHOPT, и CL'A N/A WBNOINVD, CLWB, RDPID, RDPRU и MCOMMIT N/A N/A FPUs на ядро 1 0,5 1 1 0,5 1 Трубы на FPU 2 2 FPU шириной трубы 128-битный 256-битный 80-битный 128-битный процессор инструкции установить уровень SIMD SSE4a'e AVX AVX2 SSSE3 AVX AVX2 3DNow! 3DNow! - N/A N/A PREFETCH/PREFETCHW FMA4, LWP, TBM и XOP N/A N/A N/A FMA3 L1 кэш данных на ядро (KiB) 64 16 32 32 L1 ассоциативность кэша данных (пути) 2 4 8 8 L1 кэша инструкций на одного ядро 1 0,5 1 1 0,5 1 Макс APU общий кэш инструкции L1 (KiB) 256 128 192 256 512 64 128 96 128 L1 инструкция кэш ассоциативность (пути) 2 3 4 8 2 3 4 L2 кэша в ядро 1 0.5 1 1 0.5 1 Макс APU общий кэш L2 (MiB) 4 2 4 1 2 1 L2 cache associativity (ways) 16 8 16 8 APU total L3 cache (MiB) N/A 4 8 N/A 4 APU L3 cache associativity (ways) 16 16 L3 cache scheme Victim N/A Victim Victim Max stock DRAM support DDR3-1866 DDR3-2133 DDR3-2133, DDR4-2400 DDR4-2400 DDR4-2933 DDR4-3200, LPDDR4-4266 DDR3L-1333 DDR3L-1600 DDR3L-1866 DDR3-1866, DDR4-2400 DDR4-2400 Max DRAM channels per APU 2 1 2 Max stock DRAM bandwidth (GB/s) per APU 29.866 34.132 38.400 46.932 68.256 10.666 12.800 14.933 19.200 38.400 GPU microarchitecture TeraScale 2 (VLIW5) TeraScale 3 (VLIW4) GCN 2nd gen GCN 3rd gen GCN 5th gen[19] TeraScale 2 (VLIW5) GCN 2nd gen GCN 3rd gen[19] GCN 5th gen GPU instruction set TeraScale instruction set GCN instruction set TeraScale instruction set GCN instruction set Max stock GPU base clock (MHz) 600 800 844 866 1108 1250 1400 2100 538 600 ? 847 900 1200 Макс фондовой GPU базы GFLOPS 480 614,4 648,1 886,7 1134,5 1760 1971,2 2150,4 86 ? ? 345.6 460.8 3D двигатель до 400:20:8 До 384:24:6 До 512:32:8 До 704:44:16 80:8:4 128:8:4 До 192:??? До 192:?:? IOMMUv1 IOMMUv2 IoMMUv1 ? IOMMUv2 Видео декодер UVD 3.0 UVD 4.2 UVD 6.0 VCN 1.0 »21» VCN 2.0 »22» UVD 3.0 UVD 4.0 UVD 6.0 UVD 6.3 VCN 1.0 Видео encoder N/A VCE 1.0 VCE 2.0 VCE 3.1 N/A VCE 2.0 VCE 3.1 GPU энергосбережение PowerPlay PowerTune PowerTune 1.4 1.42.2 ? 1.4 1.42.2 PlayReady'h' N/A 3.0 еще не N/A 3.0 еще не поддерживаемые дисплеи 2-3 2-4 3 3 (рабочий стол)4 (мобильный, встроенный) 4 2 3 4 /drm/radeon'j'j'26'27' N/A N/A /drm/amdgpu'j'j'j'28' N/A (29) - модели APU: A8-7680, A6-7480. Только процессор: Атлон X4 845. ПК будет одним узлами. APU сочетает в себе процессор Gpu. Both have kernels. Firmware support is required. No SSE4. No SSSE3. Single performance is calculated based on the base (or increased) base hourly speed based on the FMA operation. Single Shaders: Texture Display Units: Visualization of Output Units B To play protected video content, it also requires maps, operating system, driver and application support. It also requires a compatible HDCP display. HDCP is a must-see for certain audio formats, which puts additional restrictions on the multimedia installation. To feed more than two displays, additional panels must have home-based DisplayPort support. You can use alternative active DisplayPort-to-DVI/HDMI/VGA active adapters. b DRM (Direct Rendering Manager) is a component of the Linux kernel. Support in this table refers to the most current version. The ARM ARM Bifrost microarchitecture implemented in Mali-G71 is fully compliant with the specifications of HSA 1.1 equipment. Since June 2016, ARM has not announced support for the software that will use this hardware feature. See also General Computing on GPUs (GPGPU) Not Uniform Access to Memory (NUMA) OpenMP Shared Memory zero copy Links - Tarun Iyer (April 30, 2013). AMD introduces its heterogeneous unified memory access technology (hUMA). Tom's equipment. a b George Kyriasis (August 30, 2012). Heterogeneous system architecture: Technical Review (PDF) (report). Amd. Archive from the original (PDF) dated March 28, 2014. Received on May 26, 2014. What is the Heterogeneous System Architecture (HSA)?. Amd. Archive from the original on June 21, 2014. Received on May 23, 2014. Joel Grushka (August 26, 2013). HSAIL installation: AMD explains the future of CPU/GPU collaboration. ExtremeTech. Sieff Davis. Linaro. LCE13: Heterogeneous System Architecture (HSA) on ARM. slideshare.net. - b Heterogeneous system architecture: purpose and perspective. gpuscience.com on November 9, 2012. Archive from the original on February 1, 2014. Received on May 24, 2014. The system's heterogeneous architecture: multi-core image processing using a combination of CPU and GPU elements. Built-in computing design. Received on May 23, 2014. Kaveri's microarchitecture. A half-macaque. January 15, 2014. Michael Larabel (July 21, 2014). The AmDKFD driver is still developing for open source HSA on Linux. Phoronix. Received on January 21, 2015. b Linux Core 3.19, Section 1.3. HSA driver for AMD GPU devices. kernelnewbies.org February 8, 2015. Received on February 12, 2015. HSA-Runtime-Reference-Source/README.md at master. github.com on November 14, 2014. Received on February 12, 2015. Linux Core 4.14 Announced with secure memory encryption and more. November 13, 2017. Alex Woody (August 26, 2013). The HSA Foundation seeks GPU valor Java. HPCwire. Bolt on github. - AMD GPUOpen (April 19, 2016). CodeXL 2.0 includes an HSA profiler. AMD AMD 7th generation VSU: Mk2 excavator in Bristol Ridge and Stoney Ridge for laptops. May 31, 2016. Received on January 3, 2020. AMD Mobile Carrizo Family API designed to provide a significant jump in performance, energy efficiency in 2015 (press release). November 20, 2014. Received on February 16, 2015. Mobile Processor Comparison Guide Rev. 13.0 Page 5 : AMD Mobile Processor Full List. TechARP.com received on December 13, 2017. a b AMD VEGA10 and VEGA11 GPUs seen in the OpenCL driver. VideoCardz.com. received on June 6, 2017. Catress, Jan. zen Cores and Vega: Ryzen APUs for AM4 - AMD Tech Day at CES: 2018 Roadmap Revealed, with Ryzen APUs, zen at 12nm, Vega at 7nm. Anandtech. Received on February 7, 2018. Michael Larabel (November 17, 2017). Radeon VCN Land Support Code in Mesa 17.4 Git. Phoronix. Received on November 20, 2017. Liu, Leo (September 4, 2020). Add Renoir VCN to decode support. Received on September 11, 2020. It has the same VCN2.x unit as Navi1x and Tony Chen; Jason Greaves, AMD Graphics Core Next (PDF), AMD, August 13, 2016, A Technical View of AMD's Kaveri Architecture. The semi is accurate. Received on July 6, 2014. How to connect three or more monitors to the AMD Radeon™ HD 5000, HD 6000 and HD 7000 graphics card series?. Amd. Received on December 8, 2014. David Airley (November 26, 2009). DisplayPort is supported by the driver KMS mainlined in the Linux 2.6.33 core. Received on January 16, 2016. Radeon is a feature of the matrix. freedesktop.org. received on January 10, 2016. Alexander Deucher (September 16, 2015). XDC2015: AMDGPU (PDF). Received on January 16, 2016. a b Michelle Donzer (November 17, 2016). (ANNOUNCE) xf86-video-amdgpu 1.2.0. lists.x.org. ARM Bifrost GPU Architecture. May 30, 2016. External references to the Commons have media related to heterogeneous system architecture. HSA Heterogeneous System Architecture Review on YouTube Vinod Tipparaju at SC13 in November 2013 HSA and Ecosystem Software 2012 - HSA Michael Houston extracted from heterogeneous system architecture pdf. heterogeneous system architecture foundation. heterogeneous system architecture ppt. heterogeneous system architecture (hsa). heterogeneous system architecture threats. heterogeneous system architecture a technical review. intel heterogeneous system architecture. microsoft heterogeneous system architecture

normal_5f879f8062d68.pdf
normal_5f874999dbad3.pdf
normal_5f876ef6d6c34.pdf
kingdom series chuck black movie trailer
prime time math book grade 6 pdf
collective investment schemes control act pdf
elevation church sermons pdf
orange alpine 160 size guide
snow crash deliverator
stardew valley how to give gifts
la imagen corporativa norberto chaves pdf
hyundai key fob tricks
google material design android app
normal_5f87231087fe2.pdf
normal_5f86fad88c1d9.pdf
normal_5f8730e4c54bc.pdf