



HSA FOUNDATION

Purpose and Outlook

Patrick Moorhead, President & Principal Analyst, Moor Insights & Strategy
November, 2012

TABLE OF CONTENTS

EXECUTIVE SUMMARY.....	2
HSA FOUNDATION VISION.....	2
PARALLEL PROGRAMMING CHALLENGES.....	5
HSA FOUNDATION APPROACH.....	7
HSA BENEFITS.....	9
ALTERNATIVES TO HSA.....	11
RISKS & PRICE OF NOT ADOPTING HSA.....	12
SUMMARY.....	13
REFERENCES.....	15

EXECUTIVE SUMMARY

The HSA (Heterogeneous System Architecture) Foundation, known as the “HSAF”, is an open, industry standard consortium founded to define and deliver open standards and tools for hardware and software to fully take advantage of *high performance* of parallel compute engines, and do so in the *lowest possible power* envelope. This new environment will enable rich new user experiences never been seen before, and done at incredibly low power.



The HSAF aims to simplify and optimize parallel programming of high-performance, low power heterogeneous computing architectures including CPUs, GPUs, APUs and DSP (Digital Signal Processor). The foundation intends to literally change the hardware for participating members to suit the millions of mainstream programmers, different from many other industry approaches that can result in daunting programming challenges and applicability only to high performance computing and workstation workloads. For software, the HSAF will take the simplified hardware and create new software with the ultimate goal of plugging into the most popular, novice programmer development environments. This would then result in opening up heterogeneous computing to tens of millions of programmers.

By aligning the interests of its founding members, AMD, ARM, Imagination Technologies, MediaTek, Qualcomm, Samsung, and Texas Instruments (TI), this new ecosystem will combine scalar processing on the CPU with GPUs and DSPs parallel processing for SoC (system on chip) silicon vendors, SoC IP building block vendors, OEMs (Original Equipment Manufacturer), OSVs (Operating System Vendors), ISVs (Independent Software Vendors) and academia. The HSA will ultimately drive innovation into the future of billions of multi-core devices, spanning from futuristic devices like smart watches and room computers, to today’s innovations in smartphones, tablets, PCs, TVs, to giant cloud server installations.

Alternative approaches do exist to HSAF like NVIDIA’s CUDA, focused on specific and proprietary NVIDIA hardware. HSAF aims to build tools that are simpler and that work across AMD, ARM, Imagination Technologies, MediaTek, Qualcomm, Samsung and Texas Instruments hardware.

HSA FOUNDATION VISION

The HSAF vision is an exciting one. In five to 10 years, the author believes that billions of client devices, networks and data centers will be doing things users may not have imagined *with* maximum battery life *and* minimal data center power draw. Users are given quicker, context-aware, information rich, and more highly interactive experiences. This will require graphics, compute, and fixed function interoperation on the client device, network device, and server. This will be made possible by fully harnessing the parallel performance of the GPU and assigning the right tasks to the right silicon, whether it is CPU, GPU, APU or DSP.

Tens of millions of programmers are expected to carry out this task as tools required to carry out this programming will be simplified over time via HSA Foundation development tools. The HSA revolution has already started with Ninja programmers using OpenCL™ and Open GL ES languages who understand the nuances of the GPU. This wave can then evolve into a simpler process where sophisticated coders will more easily take advantage of HSA with languages like C++, C# and Java. Ultimately, the author believes even the most simplistic of development environments like Apple Xcode and Google Android SDK will need to be accessible to fully tap into the power of the GPU.

Device	Smartphones	Tablets	PC	Servers	Embedded
Max Power	3W	5 W	35-65W	5-250W	50 W
Compute (CPU+ GPU)	15 GF	45 GF	600 GF	5 TF	50 GF
Volume 2015	2 Billion	300 M	400 M	30 M	100 M

Source: Moor Insights & Strategy

Benefits of Parallel Processing

There are three general kinds of processing: serial, parallel and mixed serial/parallel. Generally speaking, the authors believes parallel software workloads perform better on parallel processors like GPUs and serial processing works better on CPUs and the mixed processing tries it best to be efficient on either processor type.


Parallel workloads are booming on modern devices driven in part by video content, 3D modeling and gaming. Computer vision, video and image object recognition, and video "perfection" is taking video content to new possibilities. The author has observed AMD, Intel and NVIDIA showing these processes take enormous amounts of parallelized processing power, which is not currently optimized for scalar processors.

Whether it is in workstation applications or games, 3D modeling is also already highly parallelized. The most complex models are designed, rendered and played on highly parallelized platforms, primarily driven by the parallel GPU and augmented by the parallel CPU. AMD FirePro™ technology and NVIDIA Quadro, Tesla, and Maximus are the most popular and the most highly parallelized examples of parallel computing used in workstations and HPC servers delivering the highest raw performance.

The needs for today's optimizations of video, 3D modeling and gaming have inspired the GPU industry to deliver as many GFLOPS (gigaFLOPS), or billions of floating point operations per

second as possible. This has been a race at the high end by NVIDIA and AMD. Ten years ago, the CPU and GPU were equal processors as measured by GFLOPS. In 2006, programmable shaders really gave GPUs a sizable advantage over CPUs for parallel tasks. According to AMD, this continued even farther and as of 2012, the highest performance GPU delivered 30X times the gigaFLOPS of the fastest CPU.

The challenge is that few operating systems, applications or memory architectures have been optimized yet to take advantage of all this performance.



According to the HSA Foundation, the HSAF aims to seek agreement and create unified hardware and software standards that enable more efficient, shared access to memory and communication by the CPU, GPU, and DSP. By exploiting the power and performance advantages of the GPU, the author believes heterogeneous computing can make possible substantial new features, usage models, and user experiences on tablets, smartphones, and PCs that otherwise have only been available in HPC environments.

Workloads and Usage Models

As explained above, CPUs and GPUs, when architected to work better together, will drive the next generation of client usage models and data center workloads. The following are a sample of application areas with abundant parallel workloads:

- Natural UI and Gestures:** Natural way of interacting as visual gestures like face, hands, and voices are read at greater distances and/or better precision; inches on tablets and smartphones and measured in feet on digital TVs. When multiple streams of user input and sensors are coming in, parallel processing will quickly turn the raw data into usable information.

- **Biometric Recognition:** Secure login authentication as devices recognize the user through two and even three stage methods using face, gesture, voice, mood, fingerprint, location and context.
- **Augmented Reality:** Altered appearance and/or useful information overlaid in foreground, background, and environments.
- **Content Everywhere:** User access to media in at all times with automatic bookmarking. Continuous computing enables users to access all of their content anywhere, anytime, on any device, right where they left off.
- **Beyond HD Experiences:** Media pixel ratios and resolutions are increasing along with the need for greater power efficiencies as parallel workloads rise. Hollywood and the CE (consumer electronics) industry are driving 4K (known as “Ultra HD”) cameras and TV production, creating a 4x increase in density over today’s HD media. The author believes Apple’s Retina Display is just one precursor to 4K and 8K resolutions are the next generation media experiences.

Power

The goal of the HSA Foundation is to maximize application performance but not at the expense of power. This means creating architecture that runs the right workload on the right processor at the right time. It needs to be power efficient enough to run on devices as tiny as a watch to the lowest-power, highest density servers. This is accomplished by quickly assigning the workload to the most efficient compute engine, a CPU, GPU, or DSP. Power is also reduced by giving each compute engine equal access to memory.

PARALLEL PROGRAMMING CHALLENGES

While multi-core processors are designed for greater performance with optimal power consumption, the software development that is needed to maximize the performance potential of multi-core systems is much more complicated than those associated with serial computing and CPUs.

Today, only a handful of "Ninja programmers" are up to the task. AMD estimates that only 100K GPU programmers and over 10M CPU programmers exist. The HSAF aims to implement solutions and methodology to address the issues below through an industry standard framework to implement parallel programming for different software applications.

Ninja Programmers

Today, these parallel programming challenges are being overcome with the most sophisticated programmers. Considering the expertise necessary to overcome the parallel design challenges, the programmers capable of handling GPU/hardware development have been referred to by Adobe at AMD’s last development summit as “Ninja Programmers”. These developers have

significant niche value, creating applications with high performance and low power, but typically requiring high investments in effort and skill and an increased time to market for new features. This greatly reduces the talent pool, as according to AMD, there are approximately 100K GPU coders working with a small number of apps (200+) but amidst growing ecosystem.


The HSA standard can ease the programming difficulty and grow the ecosystem of parallel processing, allowing people to work more effectively and interactively.

Mainstream Programmers

To increase the number of applications that can take advantage of the parallel GPU in harmony with the CPU and DSP, more mainstream tools should be embraced by developers. Mainstream developers consist of a much larger group, greater than 10M compared to the 100K Ninjas, working with 4M+ apps that are well established and enable great user experiences on the CPU.

The author believes even with mainstream approaches coding for parallel content, they are unable to get the performance they need with lower power. The HSAF strives to reduce time to market and lower developer investment than with Ninjas and GPU coding, but also a reduced return in terms of differentiation of features and performance.

The HSA Foundation seeks to create a methodology and framework that reduces developer investment and increases returns in programming applications that take advantage of the power, performance, and energy efficiency of the GPU in parallel.



HSA FOUNDATION APPROACH

Traditionally, when innovation is achieved in hardware, software developers are asked to adjust and rework the software to take advantage of new features. This thinking is shortsighted, outdated, and limiting when compared to the possibilities of a global ecosystem. The HSAF aims to deliver the best parallel hardware that works with existing programming languages and task parallel runtimes that are being coded for multicore CPUs.

Industry Standard

The HSAF was created to drive industry standards. There are two types of standards, “de-facto” and “industry” standards. De-facto standards are those driven by a single company which becomes the standard. Other initiatives are driven by industry-standard, non-profits like the Linux Foundation or Khronos. Industry-standards pull together multitudes of companies who agree on standards which where companies can provide input to and receive the output free of charge.

The HSAF is an industry standard movement that allows multiple companies to work together to bring to market strong development solutions to drive innovative advanced content and applications. Open standards build big ecosystems, and an open standardization body for HSA platforms can champion the development and evangelization of the architecture going forward, desiring to make heterogeneous programming easy and a first-class complement to CPU computing with increased power efficiency, allowing it be a platform of choice from smartphones to the cloud. Creating an industry standard will also foster growth and drive learning and innovation among programming talent through developer training and academic programs.

Bylaws dictate that the HSA Foundation will act as a non-profit corporation for promoting the broad and open industry adoption of the HSA Architecture.

Requirements for Membership:

- **Founder/Promoter/Contributor/Associate/Supporter:** For-profit businesses that engage in or support the research, development, production, manufacture, use, sale, or standardization of heterogeneous system architectures and software.
- **Academic members:** Individuals and Academic entities that engage in or support HSAs and software, provide knowledge of value to the Board, and are invited by Board
- **Supporter:** Businesses with revenue < \$25 Million, primary business of core tool development of compilers, simulators, debuggers, and/or libraries.

	Promoter	Supporter	Contributor	Academic	Associate
Candidate for BOD	Yes	No	No	No	No
Candidate for Working Group Chairperson	Yes	Yes	No	No	No
Working Group Voting Rights	Yes	Yes	Yes	No	No
Working Group Participation	Yes	Yes	Yes	Yes	No

Chart provided by the HSAF

HSAF founders have been working on heterogeneous designs and understand GPU compute. They have all done their own custom modifications or have leveraged Ninja-class tools to GPU-accelerate software. The Founders represent extensive industry knowledge and coverage across a wide swath of the market.

Founders have very impressive market shares in graphics, CPUs, SOCs, and smartphones markets. Via the full value chain of these founding companies, AMD, ARM, Imagination Technologies, MediaTek, Qualcomm, Samsung, and TI, the HSAF will work to bring about broad industry support for heterogeneous computing through open standardization and architecture support and definition for GPU computing and special purpose hardware accelerators. The organization will own and advance the specs and conformance, drive innovative advanced content and applications, and cultivate programming talent through developer training and educational programs.

The Memory Solution

As previously discussed, arbitrating memory between CPU, GPU, and DSP in a heterogeneous computing environment is inefficient and drives complexity in the programming environment. [According to IMS](#), already half of the processor revenue is driven by heterogeneous (hybrid) devices, including at least a CPU, GPU and a non-unified memory controller. According to the HSAF founders, many of these implementations are performing compute on the GPU, but without the unified memory controller, the full potential of performance and low power aren't harnessed and programming remains with the Ninjas.

With system architecture optimized within the HSA framework, the CPU, GPU and DSP are “equal citizens”. All forms of processors can share system memory equally, meaning that system memory is fully accessible per the needs of the application.

Even though the HSAF was formed just a few months ago, there has been a lot of work done already. So far, the HSAF has published the following on the HSAF website which address shared memory:

- HSA Programmer Reference Manual
- HSA Hardware System Architecture Specification
- HSA Software System Architecture Specification
- [AMD IOMMU V2 Architectural Specification](#)

Silicon IP and silicon hardware providers can adopt and then drive these specifications into their products which may then be adopted by device makers.

The Software and Programming Solution

As previously discussed, parallel programming is difficult to program and the goals of the HSAF are to enable a simplified development platform.

The HSAF will provide to tools necessary for programmers to build apps that take advantage of HSA. Specifically, these are called compilers, debuggers, and performance optimizers

HSAF BENEFITS

General HSAF Benefits

Many benefits exist across many of the different constituents:

- **Cost Benefit:** Members can leverage the combined planning, development, test and distribution of the hardware and software specifications and building the software tools that they would have had to do on their own.

- **Time to Market (TTM):** If the alternative of an open standard approach is doing it yourself, members should see an advantage in getting to market quickly.

- **Ecosystem Pull:** With cultured standard bodies that have delivered value for years, comes the power of the combined pull of the thousands of members and their cumulative impact. Like compounding money, ecosystems provide an amplification or ecosystem effect a single company may not provide.

- **Intellectual Property:** Founders can leverage the IP created inside the foundation license-free. The alternative is creating all the IP yourself or licensing it for fee.

- **Silicon Design Simplicity:** According to the HSAF, in today's non-standard parallel programming environment, many silicon architectures need to support every different kind of hardware bus to be able to interoperate correctly between the CPU, GPU and DSP. By standardizing on these interfaces this simplifies the silicon and the required software.

OEM and Developers Benefits

In addition to the general benefits HSAF could provide outlined above, it is beneficial to drill down into specific benefits to device and server OEMs and ISVs could realize. This comes with the assumption that the HSAF and HSA delivers on everything they are committed to deliver.

The following are the potential benefits to ISVs by implementing HSA:

- **Competitive advantage:** Because HSA offers a standard way to leverage the incredible parallel processing and load balancing offered by the GPU, HSA applications will provide rich experiences described in the workload and usage models section above. Richer applications could provide a large business opportunity that is differentiated from ISVs who do not support HSA.
- **Time to Market:** By leveraging standard, non-Ninja tools, ISVs may be able to improve their time to market of their rich applications.
- **Lower Cost:** Because more mainstream programmers should be able to program for HSA applications, more of them will be available and therefore should be lower cost. Ninja programmers are very highly skilled, but also can be very expensive.
- **Investment Protection:** When ISVs support HSA, they will write code that has the supported of multiple hardware platforms from ARM, AMD, Imagination Technologies, MediaTek, Texas Instruments and any other future adopters.

Client device, network and server makers can potentially benefit from HSA as well if they develop HSA-compliant hardware:

- **Access to Rich HSA Apps:** HSA-compliant hardware can fully take advantage of the rich usage models and workloads outlined above.
- **Low Power:** Because HSA optimizes the workload for the appropriate processor, HSA-compliant platforms can be lower power.
- **Lower Development Cost:** Device and server makers have the potential to have varying degrees of common platforms, providing lower hardware and software developments costs.

Initiative Roadmap

The HSA Foundation recognizes that this is just the beginning of the initiative and that hundreds more deliverables need to be completed over the next 10 years to reach the vision described earlier.

There is also the reality that heterogeneous computing solutions *already exist today* and are being utilized by HSAF members and non-members. The main point here is to simplify the programming to open up the power of the architecture to millions, not a few hundred thousand programmers.

The HSAF will strive to bridge the gap between now and the “end vision state” with multiple deliverables, including new hardware memory architecture and software tools like compilers, debuggers, profilers, simulators and libraries. These tools will fill the gap and over time, move simplicity to more and more programmers.

In reality, there will be multiple serial processes going on for multiple programmer toolsets.

ALTERNATIVES TO HSA

A few alternatives for heterogeneous computing exist today, but they exist in the world of Ninja programmers. Some of these programs are full development programs like NVIDIA’s CUDA, other are industry standard languages like OpenCL™ and others are custom, low-level, brute force coding done on many phones today.

To be clear, the author believes there are no alternatives today for a unified memory-based solution set like the HSAF would like to introduce. This does not mean they won’t ever exist, it’s just that today, they don’t.

Here are the following leading non-unified memory architecture solutions:

NVIDIA CUDA

NVIDIA’s CUDA (Compute Unified Device Architecture) is a parallel computing platform and programming model invented by NVIDIA and released in 2006, offering the first solution for general computing on GPUs where no assembly language is required, allowing C, C++, C#, OpenCL and Fortran code to be sent straight to the GPU. CUDA is supported on NVIDIA GPUs and is widely deployed through thousands of applications primarily workstation and HPC. There are over 300 Million CUDA enabled GPUs in notebooks, workstations, compute clusters and supercomputers. CUDA currently has support for 4 numerical analysis tools, 10 language solutions, 5 performance analysis tools, 5 debuggers, and 13 libraries.

OpenCL™

OpenCL, the first open programming standard for writing software across CPUs and GPUs, is now supported by a number of comprehensive pre-built libraries for application development and is being taught at hundreds of colleges and universities. Data from SimplyHired.com indicates that between 2010 and 2011, there was a 156% increase in demand for developers using OpenCL. HSA is not an alternative to OpenCL, but it is optimized platform architecture, and OpenCL will benefit from the avoidance of wasteful copies with low latency dispatch, improved memory model, and pointers shared between CPU and GPU. In addition, HSA exposes a lower level programming interface for those seeking optimized performance and maximized control.

Other Methods

As previously noted, many companies are doing very low level, standard and non-standard programming to take advantage of the compute capabilities inside the GPU. According to the HSAF, similar to OpenCL, some of these standard methods entail using HLSL, OpenGL Shading language (GLSL), Compute Shader, in traditional programming models like DirectX®, OpenGL and OpenGL ES. Many of these implementations are in the phone, tablet and PC market, and are tied to a specific function in a specific application. These are painstaking enhancements that take a lot of time, money, and resources and given any changes in the surrounding code, need to be re-optimized.

PRICE OF NOT ADOPTING AND POTENTIAL RISKS TO HSA

The author believes that by not adopting a standard ecosystem for application development inside this unified memory environment, an enormous opportunity could be missed. The status quo system of requiring a small pool of qualified programmers to build applications that fit parallel processing hardware is maintained as their expertise is locked into the specific hardware of the vendor. Problems still exist in maintenance and debugging, and the time to market can't match the advances in demand for mobile access, processing speeds, and user interactivity with video and massive amounts of data in communication. As Moore's Law advances in more processors on smaller chips, the challenges presented by programming, development costs, and power usage will only be amplified.

Non-Joining Chip Companies

Chip companies that do not join the HSA are relying on their own resources in hardware, tools, libraries, and languages, and possibly most importantly, developers willing to continue programming according to their hardware. If the goals of the HSA Foundation progress as

planned, developers are poised to refine their skills to match the development opportunities of greatest demand. Non-joining companies' already slim pool of qualified "Ninja" programmers with adequate expertise in parallel systems may continue to shrink, which will likely increase those costs. The opportunity of the shared resources among the founding companies and associated members will not be at the disposal of companies that do not join.

Risks to the HSA Foundation

There are always risks associated with anything designed to make major shifts in the ecosystem and those that have a big future pay off.

- **Time to Market (TTM):** Any time multiple companies need to agree on anything that is connected to other companies, it takes time. HSAF standards have the ability to impact hundreds if not thousands of companies in the overall ecosystem. Additionally, operating systems, middleware, and development tools will need to be built or modified to take advantage of the HSA work, and this takes time.

- **Relevant Compromise:** For multiple companies to agree on things there is always the need for compromise. Can a unified memory standard that works well for a phone work well for a cloud server?

- **Degree of Difficulty:** The HSAF members are some of the brightest in the industry. Given the monumental challenges it takes to parallelize code for a single component like a GPU, how difficult will it be to do for a CPU, GPU and DSP at the same time? Yes, the unified memory architecture will make it much easier, but will it make it easier to do it for three different classes of processors? Also, to take full advantage of the unified memory, operating systems will need to be modified, likely requiring support from the major player like Microsoft, Apple, Google or the Linux communities. Middleware and development tools will need to be created or modified to take advantage of the HSAF APIs (application programming interfaces).

- **Other Major Players:** The current HSAF Founders provide a very large footprint in phones, tablets, TVs PCs, and other embedded markets. The addition of NVIDIA or Intel would provide an even greater reach. With greater reach, this increases the probability of broader HSAF success.

SUMMARY

The HSA Foundation champions a long-term, industry-standard approach to solving today's challenges in harnessing the massive raw performance in parallel computing in the lowest power envelope. This open standard will span a wide breadth of devices from smart watches to smart phones to PCs to enterprise servers, and everything in between. The founders recognize both the challenges and the opportunities presented by parallel computing and seek to change the thinking of traditional approaches by "bending" the hardware to fit the software. One of the group's most important tasks are to create an industry standard for unified memory which sees

the CPU, GPU and DSP as equal citizens. This in itself will simplify programming and the development of those programming tools significantly.

Founders AMD, ARM, Imagination Technologies, MediaTek, Qualcomm, Samsung, and Texas Instruments will pool its resources and experience in CPUs, GPUs, DSPs and SoCs to develop a specification for hardware and software upon which they intend to drive into the markets with the support from the HSA ecosystem. The foundation has already opened its doors to promoters, supporters, contributors, academics, and associates to leverage the ecosystem even more.

The HSAF goals are lofty and come with the risks generally associated with most industry standard groups. These risks have been mitigated but not eliminated with goal-aligned participants who have enough market size to drive the standard, the resources to execute the deliverables, and the urgency to do it as quickly as possible. The author believes right now, the industry's alternative comes down to keeping the status quo with complex programming suitable only for Ninja programmers. In that scenario, it would be impossible to achieve the vision the HSAF has laid out.

The HSAF is off to a positive start and was able to leverage the 36 months of development time before the official launch in June 2012 as evidenced by delivery of a draft memory specification. Per the HSAF, the next major milestone will be agreement on the hardware specification. If this can be done over the next 6 months, this would enable HSA-compliant hardware to be in the hands of software developers in 18 months and this would be a very positive sign for the HSAF.

REFERENCES

- AMD. Interview with AMD, Chief HSA Evangelist: [Greg Stoner](http://www.amd.com) www.amd.com
- ARM. Interview with ARM, Media Processing Division: [Jem Davies](http://www.arm.com) www.arm.com
- Borenstein, Greg. "[Adam Harvey Explains Viola-Jones Face Detection](http://www.makemantics.com)" www.makemantics.com
- Chang, Janie. "[Two Extremes of Touch Interaction](http://research.microsoft.com)" <http://research.microsoft.com>
- Christensen, Mikael Hvidtfeldt. "[GPU vs CPU for Pixel Graphics](http://blog.hvidtfeldts.net)" <http://blog.hvidtfeldts.net>
- Cogswell, Jeff. "[Parallel Processing with Many Integrated Core \(MIC\) Architecture for Weather Prediction](http://www.goparallel.sourceforge.net)" www.goparallel.sourceforge.net
- COTS Journal, the Journal of Military Electronics & Computing. "[APU Platform Blends GPU and CPU Advantages for Parallel Processing](http://www.cotsjournalonline.com/)" <http://www.cotsjournalonline.com/>
- DeGrasse, Martha. "[Qualcomm, Broadcom, Texas Instruments Coming on Strong in Semiconductor Market](http://www.rcrwireless.com)" www.rcrwireless.com
- Facebook Developers. "[Open Source](http://developers.facebook.com)" <http://developers.facebook.com>
- Fast Matrix Solver. "[Why GPUs](http://www.fmslib.com)" www.fmslib.com
- Ford, Dale. "[Intel Reasserts Semiconductor Market Leadership in 2011](http://www.isuppli.com)" www.isuppli.com
- GfxS Staff. "[ARM gaining market share in graphics](http://www.gfxspeak.com)" www.gfxspeak.com
- HSA Foundation. "[HSA Foundation Bylaws](http://www.hsafoundation.com)" www.hsafoundation.com
- Hwu, Wen-mei. "[Three Challenges in Parallel Programming](http://www.parallel.illinois.edu)" www.parallel.illinois.edu
- IHS Research. "[Half of \\$111 Billion 2011 Processor Revenue Generated by Hybrids, According to IMS Research](http://www.imsresearch.com)". www.imsresearch.com
- Imagination Technologies. "[PowerVR Graphics](http://www.imgtec.com)" www.imgtec.com
- Imagination Technologies. "[The promise of heterogeneous computing: an insider look at the HSA Foundation](http://www.imgtec.com)" www.imgtec.com
- Imagination Technologies. Interview with Imagination Technologies: [Alexandru Voica](http://www.imgtec.com) www.imgtec.com
- Intel. "[Many Integrated Core \(MIC\) Architecture-Advanced](http://www.intel.com)" www.intel.com
- Jon Peddie Research. "[Graphics Add-in Board Shipments seasonally down just 2% from Last Quarter, Reports Jon Peddie Research](http://www.jonpeddie.com)" www.jonpeddie.com
- Jon Peddie Research. "[Q1 Graphics Shipments Decline .8% Over Last Quarter and Slip 3.38% Over Last Year.... And that's the good news.](http://www.jonpeddie.com)" www.jonpeddie.com
- K, Sarnath. "[At the Epicenter of Heterogeneous Computing](http://ers.hclblogs.com)" <http://ers.hclblogs.com>
- KHRONOS. "[OpenCL-The open standard for parallel programming of heterogeneous systems](http://www.khronos.org)" www.khronos.org
- KHRONOS. "[OpenMAX-The Standard for Media Library Portability](http://www.khronos.org)" www.khronos.org
- MacManus, Christopher. "[Microsoft's 3D computer offers a world for your hands](http://www.news.cbet.com)" www.news.cbet.com
- Makedon, Terry. "[HSA Represents The Evolution of Computing](http://www.blogs.amd.com)" www.blogs.amd.com
- Mattson, Tim. "[How to sound like a Parallel Programming Expert](http://www.software.intel.com)" www.software.intel.com
- McKenney, Paul E, Linux Technology Center, IBM Beaverton. "[Is Parallel Programming Hard, And, If So, What Can You Do About It?](http://www.kernel.org)" www.kernel.org
- MediaTek. "[Wireless Handset Products](http://www.mediatek.com)" www.mediatek.com
- Microsoft Applied Sciences Group. "[Adaptive Hardware. Context Aware Interfaces](http://www.microsoft.com)" www.microsoft.com

- Microsoft Applied Sciences Group. "[Gesture Interaction. A human-computer interaction system](#)" [www.microsoft.com](#)
- Microsoft Applied Sciences Group. "[History](#)" [www.microsoft.com](#)
- Microsoft Applied Sciences Group. "[The Wedge. Seeing Smart Displays Through A New Lens.](#)" [www.microsoft.com](#)
- Microsoft. "[Visual C# 2010 Express](#)" [www.microsoft.com](#)
- NVIDIA. "[Introducing CUDA5](#)" [www.nvidia.com](#)
- NVIDIA. "[CUDA TOOLKIT](#)" <http://developer.nvidia.com>
- NVIDIA. "[CUDA Tools & Ecosystem](#)" <http://developer.nvidia.com>
- Oblong Industries. Interview with Oblong Industries, CEO: [Kwindla Kramer](#)
- Oracle. "[Learn About Java Technology](#)" [www.java.com](#)
- Texas Instruments. "[OMAP Mobile Processors](#)" [www.ti.com](#)
- Texas Instruments. Interview with Texas Instruments, Linux Development Center: [Matthew Locke](#)
- The Linux Foundation. "[About Us](#)" [www.linuxfoundation.org](#)
- Varhol, Peter. "[GPU vs. CPU Computing](#)"

Author

[Patrick Moorhead](#), Principal Analyst

Inquiries

Please contact us at the email address above if you would like to discuss this report and Moor Insights & Strategy will promptly respond.

Licensing

Creative Commons Attribution: Licensees may copy, distribute, display and perform the work and make derivative works based on this paper only if *Patrick Moorhead*, and *Moor Insights & Strategy* are credited.

Disclosures

This paper was commissioned by the AMD.

DISCLAIMER

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

©2012 Moor Insights & Strategy.

Company and product names are used for informational purposes only and may be trademarks of their respective owners.