



What's next in computing?

Computing has reached an inflection point fueled by a notable shift in consumer culture.

A dramatic increase in the availability and quality of digital content fuels appetite for rich visual experiences on PCs, tablets and other mobile devices. Users expect applications to be more engaging, more portable and more intuitive. This shift in user preference requires a fundamental improvement in processing efficiency.

Continuing constraints on power and scalability in multi-core CPU development have led software and system designers increasingly to leverage GPUs for general purpose computing. Using multiple processors, each specialized for certain kinds of workloads, is referred to as “heterogeneous computing”, and in a power conscious computing environment, has been accepted by the industry as the way of the future.

For example, an APU is a heterogeneous system that incorporates a GPU and a CPU. The GPU's vector processing capabilities enable it to perform parallel operations on large sets of data, mathematically intensive computations, in addition to the traditional task of graphics processing. That means the APU is ideally suited for:

- streaming movies in HD,
- 3D gaming,
- enabling flawless HD videoconferencing, virtual meetings and bi-directional HD video chats,
- helping to improve productivity, especially in light of the emergence of mainstream operating system support for advanced multitasking on multiple monitors and seamless application switching between graphics-intensive content like rich PowerPoint presentations, web browsing in HTML 5, product demos and simulations.

Why is heterogeneous computing important?

Heterogeneous computing enables some incredible innovations. More media in brilliant HD, improved productivity, all with greater power efficiency resulting in all-day battery life – as much as 10+ hours of continuous use (as measured with the Windows Idle test).

By designing an APU, AMD leverages the GPU's parallel processing capabilities which perform parallelizable tasks at much lower power consumption relative to the serial processing of similar data sets on the CPU. By balancing workloads on the GPU and the CPU, the APU creates a balanced and efficient computing model that improves performance and battery life. By way of an APU, heterogeneous systems leverage the processing and power efficiency of a GPU to leap ahead of the CPU-only processing curve and offer teraflops of processing power to usher in a new era of innovation in computing. It's personal supercomputing with all-day battery life!



What's next? The next level of computing – how do you want to interact with your device?

The human computer interface can be visual (web cameras, depth cameras, etc.), auditory (microphones, audio input processors, etc.), environmental (GPS, accelerometers), physical (multi-touch, haptic, etc.) and biometric (security sensors, medical sensors, etc.) Using these interfaces, new modalities of interaction with computers that free us from menu driven interaction as it exists today, can be developed using middleware that translates human behavior into actions such as voice recognition, facial recognition and gesture recognition – all of which require new levels of processing throughput and will be enhanced with GPU technology and industry standard compute APIs, such as OpenCL. Bringing these capabilities to mainstream users requires that the applications target multiple platforms from multiple hardware vendors. Therefore, using an industry standard API is critical as it greatly reduces the adoption barrier with standardization so it's worthwhile for software developers to write complex code for mass consumption. OpenCL is well on its way to making itself the dominant industry standard API for GPU compute. But today's OpenCL driver model is just the first step in opening up the true potential of the APU.

Enter AMD - The opportunity we are seizing with FSA and FSAIL

It's now evident to the whole industry that in order to take advantage of GPU capabilities that could bring new and exciting experiences to mainstream users, the GPU has to be as accessible for programmers as the CPU is today.

The AMD Fusion System Architecture (FSA) roadmap is the key to opening up the APU for easier programming, optimization, load balancing, and ever higher performance at lower power. Most importantly, FSA will be an open architecture with published specifications.

The state of affairs before FSA is to access GPU compute through the OpenCL or Direct Compute APIs and languages, which are supported by run times on top of our drivers. Post FSA, you can access GPU compute in a number of ways. You can program in a high level language such as C++ AMP (Accelerated Massive Parallelism), directly to our hardware. You may co-opt some of the FSA domain libraries which we will make available, to be called from C++AMP or other high level languages. Or you may continue to use OpenCL, which will evolve to run better on FSA, taking advantage of the efficiencies of that architecture, such as low latency dispatch of work to GPU, improved memory models, shared pointers and avoidance of data copies.

So briefly, FSA makes OpenCL run better on it, and it supports C++AMP and other higher level languages so they can run better on AMD hardware.

No matter which ISAs a platform includes (x86 CPU, GPU, ARM), FSAIL (FSA Intermediate Layer) acts as a virtual ISA, translating code to apply to the hardware. At [AMD's Fusion Developer Summit](#) (AFDS), AMD announced that the FSAIL spec will be published, which will provide an opportunity for vendors with other chip architectures to improve their performance with increased programmability by co-opting this approach to heterogeneous computing.



(AFDS hosted developers, academics and emerging innovators to learn about heterogeneous computing, Accelerated Processing Unit (APU) technology, GPU compute, parallel processing and programming tools. Keynote presentations were delivered by industry thought leaders from AMD, ARM, Corel and Microsoft. To view archived videos of keynote presentations, read keynote abstracts or watch the highlights, please visit: <http://developer.amd.com/afds/pages/keynote.aspx>)

Industry players such as ARM are completely committed to heterogeneous computing and open standards. Microsoft's announcement around the C++ AMP compiler also points in that direction. Today, a high percentage of C++ Windows apps are developed with Visual Studio and Microsoft has indicated that the next version of Visual Studio will include C++ AMP, and will directly use the GPU for enhanced user experiences.

Ecosystem adoption

Less fragmentation is good for software and application development.

Imagine being able to access your favorite applications from any of your personal devices, and that those applications deliver a seamless and scalable experience whether you're using your PC or tablet or smartphone. Let's go crazy and imagine that same experience being agnostic to the different operating systems.

An open ecosystem is compelling because it facilitates programming across different hardware platforms and operating systems. The potential result is a combined (therefore much larger) ecosystem of users for any application written to an FSA architecture.

With APUs and FSA, AMD has laid out the path to heterogeneous computing, and opened up the capabilities of this incredibly efficient processing engine to the entire computing ecosystem.

**AMD defines "all-day battery life" as 8+ hours of continuous use as measured with the Windows Idle test.*

