

# White Paper | OpenCL™: The Future of Accelerated Application Performance Is Now

## Table of Contents

- INTRODUCTION** ..... 2
  - What Is OpenCL™? ..... 2
  - Changing the Game ..... 2
  - GPUs: Not Just for Graphics Anymore..... 2
- THE “OPEN” ROAD** ..... 3
  - The Preferred Path to Performance ..... 3
  - Specific Potential Benefits of an Open Approach..... 3
  - OpenCL Has the Momentum ..... 4
- HOW WE GOT HERE** ..... 5
  - The Story Until Now ..... 5
  - OpenCL 1.1: The Turning Point..... 5
- EXPLORING THE POSSIBILITIES** ..... 5
  - CAD/CAM/AEC and Science Applications ..... 5
  - DCC/M&E Applications ..... 7
- AMD: LEADING THE WAY** ..... 8
  - Solid Support for OpenCL and Open Computing ..... 8
  - AMD FUSION APUs: Fulfilling the Ultimate Promise of OpenCL..... 9
- HOW TO LEARN MORE**..... 9



## INTRODUCTION

### What Is OpenCL™?

OpenCL™, or Open Computing Language, is the first open programming standard for writing software for general-purpose computing across both CPUs and GPUs.

That may not sound that important on the face of it. But consider this: applications accelerated with OpenCL can access the combined processing power of a computer's GPU and CPU cores under a single unified platform – a game-changing development for users of professional graphics and computational design software.

### Changing the Game

The ability to perform compute-intensive tasks in parallel, using virtually any multicore processor, opens the door for performance and functionality improvements on an unprecedented scale. If you're working with processing-intensive applications in CAD, CAM, CAE, science, DCC or M&E, get ready for OpenCL to change the way you work.

Just think. Interactive 3D raytraced rendering in near real-time. Simulations that take seconds or minutes to solve instead of hours. Multi-stream video and effects processed in real time. Whatever you do in your job, applications built using OpenCL are designed to help get it done faster and better – much faster and better.

### GPUs: Not Just for Graphics Anymore

OpenCL provides an open, uniform programming environment that enables the processing technique known as GPGPU, or general purpose computing on graphics processing units. GPGPU makes it possible to offload certain processing-intensive tasks to the GPU, rather than relying on the CPU to do them.

Why? Because GPUs are more effective than CPUs when it comes to calculating highly parallelized, "floating point" arithmetic – the type of math calculations required for physics-based modeling, animation and simulation for 3D and CAD/CAM engineering applications; real-time 3D rendering; or image processing for video creation and editing applications. And now, thanks to OpenCL, those GPUs aren't just for powering your graphics application's viewport/GUI anymore.

## THE “OPEN” ROAD

### The Preferred Path to Performance

OpenCL is not the only path to accelerated application performance, but it is the first and only path paved by open, non-proprietary technology.

And that’s critical to the impact it can have on you and how you do your job.

The main reason is that this platform-neutral approach enables ubiquitous cross-platform development. Software developers can create just one set of code, and know that it can be easily “recompiled” to run on devices using a wide range of processors, whether in PCs, servers or mobile devices – even systems powering cloud-based, remote computing applications.

On the other hand, software developed with proprietary solutions only works with the proprietary technologies with which it was designed to work. So if a developer writes software using proprietary GPGPU-enabling technology, it’s not necessarily going to work with a variety of GPUs from multiple vendors. And if you’re someone who wants to use that software, but you don’t have a compatible GPU, you’re going to have a problem.

Proprietary solutions not only limit choices (and limit markets for software vendors), they lock people and companies into technologies, making it difficult for them to respond nimbly to change. For example, if a software vendor releases features in an application that are tied to proprietary hardware, how much time and money are you willing to spare to re-equip in order to leverage those new features? And if your company is part of a larger supply chain, you’re all affected.

With the quickly expanding number of OpenCL-based applications entering the market, you’ll be able to more freely make decisions about applications based on what’s going to be best for you, your business and the companies you work with – instead of having those decisions dictated by a limited proprietary approach.

Software developer Matt Dunbar, manager of platform infrastructure with Simulia, a Dassault Systèmes company, commented to AMD on the proprietary CUDA parallel computing architecture, “It’s pretty clear that proprietary solutions are not a great answer ultimately, because they really limit the freedom of choice with our end users.”

### Specific Potential Benefits of an Open Approach

OpenCL offers a number of specific potential benefits to users of applications that are developed using OpenCL rather than a proprietary solution such as CUDA. And with a community of thousands of programmers and companies making contributions to the OpenCL platform, software developers have a large and growing body of tools, resources and shared “best practices” they can draw from in their efforts to create powerful applications for their end users.

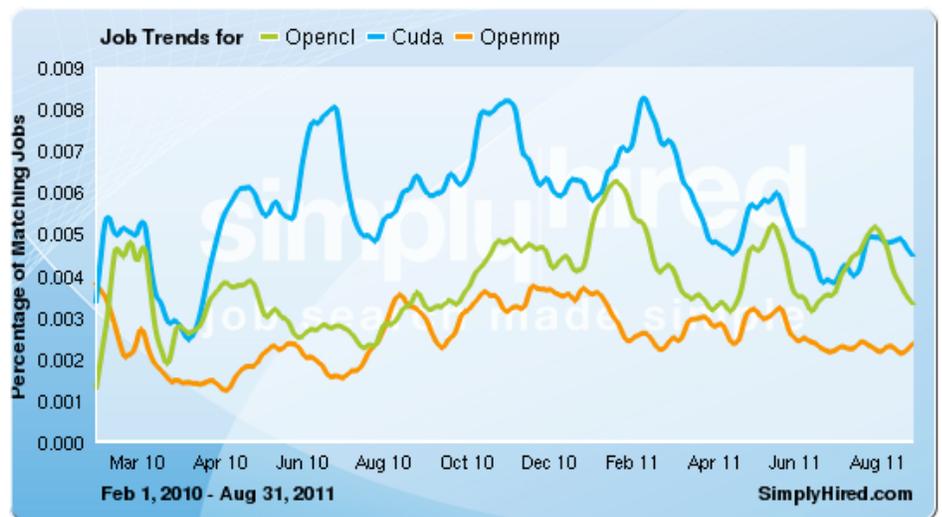
**Performance Optimization for Better Results** – Companies that make processor hardware, whether GPUs or CPUs, can optimize their device driver software for OpenCL – giving them the ability to tune their products to provide the best possible performance for particular applications. This opportunity has always been available to software developers working with OpenCL or CUDA, but, of course, hardware companies have not had the same privilege, since applications developed with CUDA only work with one type of proprietary hardware.

**Broad Choice of Technology Platforms** – A diverse group of technology companies (including AMD, Apple, IBM, Intel, Nvidia and many others) support OpenCL for easy code portability across computing platforms. Software developers who use OpenCL can develop applications that run on a range of supported devices and platforms, giving more technology choices to end users and the companies they work for.

**The Strength of a Robust Technology Ecosystem** – A broad consortium of technology companies and individuals has contributed – and continues to contribute – to the OpenCL specification and related development tools. This has resulted in a more robust technology ecosystem than would ever be possible with a proprietary solution. And users can be assured of continued widespread support for the standard over the long term.

### OpenCL Has the Momentum

OpenCL adoption has been picking up steam since the current version of the standard was adopted in 2010 (see the next section for a brief OpenCL history). Following improvements introduced with OpenCL 1.1, there are now more comprehensive pre-built libraries for developing OpenCL-enabled applications; more courses in OpenCL being offered at colleges and universities (where the added benefit is that there's no need for proprietary hardware); and a tremendous increase in demand for OpenCL developers, as the graph shows. Between March 2010 and June 2011, OpenCL job openings have increased 156% since 2010, CUDA job openings have decreased 53% for the same period.



Recent employment trends for OpenCL and CUDA as of 6/11/2011<sup>3</sup>

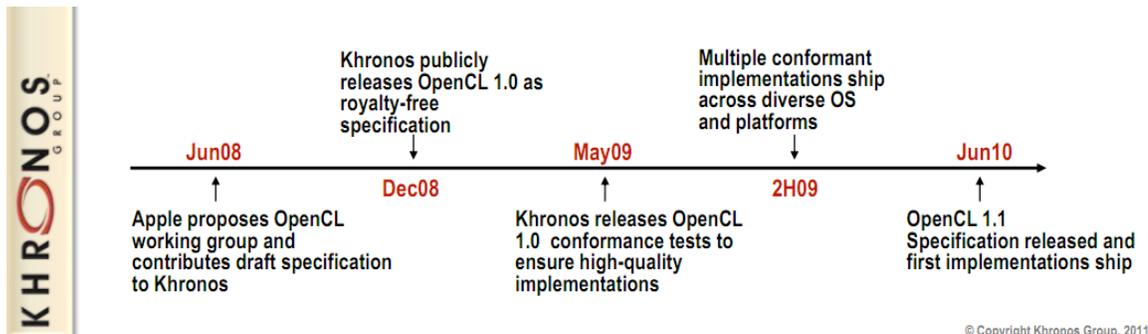
It's no wonder a recent survey<sup>1</sup> by Evans Data Corporation showed OpenCL ranked #2 among the most popular APIs for multi-threaded development in Europe and the Middle East, and ranked #3 in Asia-Pacific countries. The highest ranking achieved by CUDA in the survey, in any region, is #5. In other words, OpenCL is more popular than CUDA in all the high-tech regions of the world.

Increasing demand for open computing among application developers and users today is another factor driving adoption of open standards. OpenCL is part of a larger open-standards ecosystem that includes the industry standards OpenGL™ for 3D graphics and OpenAL™ for computer audio. For development teams who are adopting these standards, this common ecosystem means converging technology roadmaps to simplify technology adoption and future planning. OpenCL is also designed to work with other open and proprietary standards and technologies, giving developers and users maximum flexibility. OpenCL works with Linux® and Macintosh systems, and is not limited to Windows® environments, for example. It also works with other industry-standard APIs like DirectX® and OpenGL, so that developers always have the freedom and flexibility to choose just the right type of acceleration for each individual toolset or feature within any application.

## HOW WE GOT HERE

### The Story Until Now

OpenCL 1.1 is the product of many years of development and refinement by an open consortium of processor and software companies. Originally conceived by Apple, then refined and managed by the Khronos Compute OpenCL Working Group, OpenCL 1.0 was released in late 2008. In keeping with the technology's emergence as an open alternative to proprietary solutions, the Khronos working group reflects broad, diverse industry participation; today, it is made up of industry-leading experts from 33 companies including Apple and AMD. The group has continued to work on the design and function of the standard since its initial release, resulting in the release of OpenCL 1.1 some 18 months after OpenCL 1.0.



OpenCL Timeline

### OpenCL 1.1: The Turning Point

After the release of OpenCL 1.1 in 2010, the standard picked up considerable momentum as an open-standard alternative to proprietary solutions for accelerating application performance. OpenCL 1.1 enhances performance and flexibility with new built-in functionality. And while a number of the added functions focus on ease of application development, ultimately, this works to benefit end users. The ability to handle commands from multiple hosts, for example, provides support for cloud computing environments.

## EXPLORING THE POSSIBILITIES

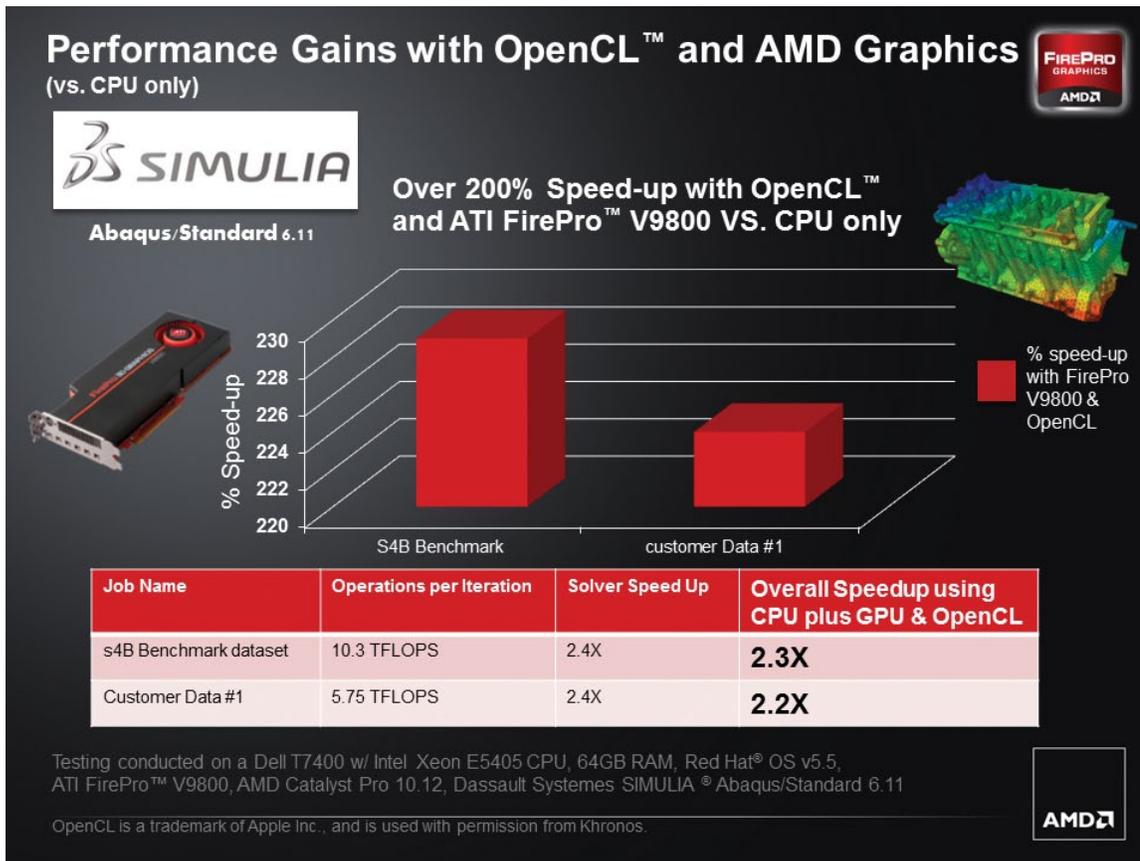
### CAD/CAM/AEC and Science Applications

Applications in computer-aided design and manufacturing; architecture, engineering and construction; and science all benefit in various ways from the speed and flexibility that OpenCL enables. The hardware processing capacity required to accurately and quickly render and simulate designs and physical behavior is critical in all these fields, and requires exceptional capabilities and flexibility in computation. OpenCL is the perfect choice for achieving this capability inside professional applications.

**Numerical Computation** – OpenCL makes it possible to dramatically accelerate various types of computation – including finite element analysis (FEA), discrete element method (DEM) and computational fluid dynamics (CFD) – in industrial design, engineering design and computational manufacturing. For example:

- DEM Solutions makes software for prototyping and optimizing materials handling equipment. By incorporating OpenCL into its EDEM engineering simulation software, the company has dramatically accelerated parts of the compute process to accelerate the processing of simulations so critical in the materials handling and manufacturing process.

- Simulia, a Dassault Systèmes company, delivers simulation solutions including the Abaqus suite of FEA software. When the company tested the performance of its current OpenCL-enabled Abaqus/Standard FEA application, the results showed an overall speedup of more than 200%<sup>2</sup>, as illustrated below.



Abaqus/Standard performance gains with OpenCL

**High-End Visualization** – OpenCL-enabled applications for high-end visualization, running on multiple GPUs, can make it possible for users to easily and flexibly scale their hardware to get the hardware performance required to run extremely accurate lighting simulations in near real-time. In the past, these simulations required hours (or even days) of “software rendering” on CPUs before the visualizations could be viewed. A case in point:

- OPTIS, a leading developer of light simulation software for automotive and aerospace OEMs, recently released its first OpenCL-based program for highly realistic visualization. According to Pete Moorhouse, the company’s vice president of sales and marketing, “We’re delivering physically correct lighting simulations in seconds or minutes where, previously, it was hours or days.”

## DCC/M&E Applications

In digital content creation (DCC), in media and entertainment design pipelines, applications that take advantage of OpenCL can easily break through previous barriers to higher performance and capabilities – just as they do in CAD/CAM/CAE and science environments.

**High-End Visualization** – The same OpenCL performance advantages that automotive and aerospace OEMs rely on in OPTIS' light simulation software can also be brought to bear on DCC and M&E applications. In fact, OPTIS' vice president of sales and marketing believes that the opportunity to expand its market beyond industrial design and into these areas is another benefit of OpenCL for his company, saying "We've shown that by leveraging OpenCL, we can expand our potential user base and develop toolsets for users in areas such as digital content creation and entertainment, without worrying about rearchitecting our software or worrying about end-user hardware."

**Real-Time Video Processing** – Video processing in today's professional content creation environments demands accelerated performance to speed and streamline video creation and editing – and OpenCL is the key to achieving it.

- Sony Creative Software recently enhanced its Sony Vegas Pro 11 collection of applications for content creation by video and broadcast professionals, adding the ability to leverage OpenCL to support faster processing of special effects filters and final rendering.
- Beyond DCC and M&E applications, MotionDSP uses OpenCL to bring accelerated video processing to intelligence and defense analysis. The company has optimized its industry-leading Ikena ISR real-time video reconstruction software to take advantage of OpenCL support for real-time video reconstruction and image enhancement.
- Accelerated and real-time transcoding for archival and distribution/broadcasting of video and media content is another prime use case for OpenCL computing. Companies such as Rovi (formerly Main Concept) have recently released OpenCL-based toolsets designed for television/web broadcasters, studios and other media content providers that can provide massively scalable video processing capacity to reformat and process dozens or even hundreds of high-resolution video streams/files simultaneously – all in real time.

**Physics Engines for Real-Time Simulations** – OpenCL can be used to accelerate the calculation of math algorithms required to "solve" or simulate the effects of real-world phenomena such as gravity, inertia, friction, drag, turbulence, etc. on virtual objects. Software applications can apply these calculations to 3D models to drive lifelike 3D animations quickly and automatically. This minimizes the tedious, frame-by-frame manual manipulation of wireframe objects that is normally required for 3D animation. Real-time physics engines for video games and other interactive programs especially benefit from OpenCL-driven acceleration, because they enable these computations in real time with accuracy that better simulates real-world behavior.

## AMD: LEADING THE WAY

### Solid Support for OpenCL and Open Computing

AMD's commitment to OpenCL is part of a longstanding policy of actively promoting and supporting open computing, rather than creating proprietary solutions. The company is constantly seeking new opportunities to leverage and extend open computing education and collaboration. In support of OpenCL in particular, the company hosts events such as the AMD Fusion Developer Summit, participates in industry forums and round tables, and offers education resources like the AMD OpenCL University Kit academic curriculum.

The greatest demonstration of AMD's commitment to open computing standards is in its products. AMD FirePro™ professional graphics, for example, delivers robust open standards support through:

- Processor architectures and developer resources designed for maximum efficiency and performance in developing and utilizing OpenCL-accelerated computing solutions
- Official AMD Catalyst™ Pro driver releases and technical support synchronized with Khronos OpenCL updates
- Cards and drivers optimized for OpenCL in common applications
- Support of non-proprietary standards and technologies for Stereographic 3D

In addition, AMD maintains strong alliances with leading independent software vendors to ensure that the next generation of applications can exploit the full power and capabilities of OpenCL acceleration. One of the latest examples is the open-source Bullet Physics plug-in for Autodesk® Maya® 2011, delivered as a new toolset for Maya users through close collaboration between AMD and Autodesk. It's designed to give 3D artists easy access (and at zero cost, for Autodesk customers) to an OpenCL-enabled version of the same Bullet Physics library widely used in the development of visual effects for major television shows and blockbuster cinematic film releases. The Bullet Physics plug-in enhances the speed and accuracy of animation workflows for Maya users across a broad range of workstation and PCs configurations, and expands the options that animators can use to create ultra-realistic computer-generated animation and visual effects.

AMD also supports developers with contributions of tools and other zero-cost, freely distributed resources to help speed the creation of the next generation of professional, OpenCL-accelerated applications. AMD performance analysis and debugging tools – along with extensive libraries for out-of-the-box OpenCL functionality – create a foundation for software companies to develop applications that realize the full potential of heterogeneous computing solutions built around the capabilities of OpenCL.

## **AMD Fusion: Fulfilling the Ultimate Promise of OpenCL**

AMD is taking OpenCL-enabled computing to the next level with a new generation of processors that bring GPU and CPU capabilities together in one chip – the APU, or accelerated processing unit. Introduced in 2011, AMD Fusion APUs are already accelerating more than 50 leading applications on millions of PCs around the world.

The shared memory and extremely fast internal data communication between the GPU and CPU processing cores on AMD Fusion APUs enable new levels of computing efficiency for OpenCL-enabled applications that include both professional and consumer use. They allow for advanced capabilities to support multitasking, vivid graphics, lifelike games, lag-free HD video playback, etc. AMD support for developers who want to identify and take advantage of opportunities in these areas includes the annual AMD Fusion Developer Summit.

## **HOW TO LEARN MORE**

Learn more about how OpenCL is enabling unprecedented performance for the applications you count on every day, and how AMD supports OpenCL, by visiting:

- <http://developer.amd.com/zones/opencilzone/pages/default.aspx>
- <http://www.khronos.org/opencil/>

1 EvansData Corporation APAC Development Study, Vol 1, 2011

2 Vs. CPU processing alone

3 Data provided by SimplyHired.com, a search engine for jobs  
<http://www.simplyhired.com/a/jobtrends/trend/q-opencil%2Ccuda%2Copenmp>