A Brief Introduction to CAD, CAE, CPU, GPU, API, SDI, BPP, MHz, GHz, MXM, AFR, SFR, DMU, DCC, ISV, GLSL, AADC, GSX, SPEC and SMPTE.

SIMPLIFYING THE WORLD OF PROFESSIONAL GRAPHICS

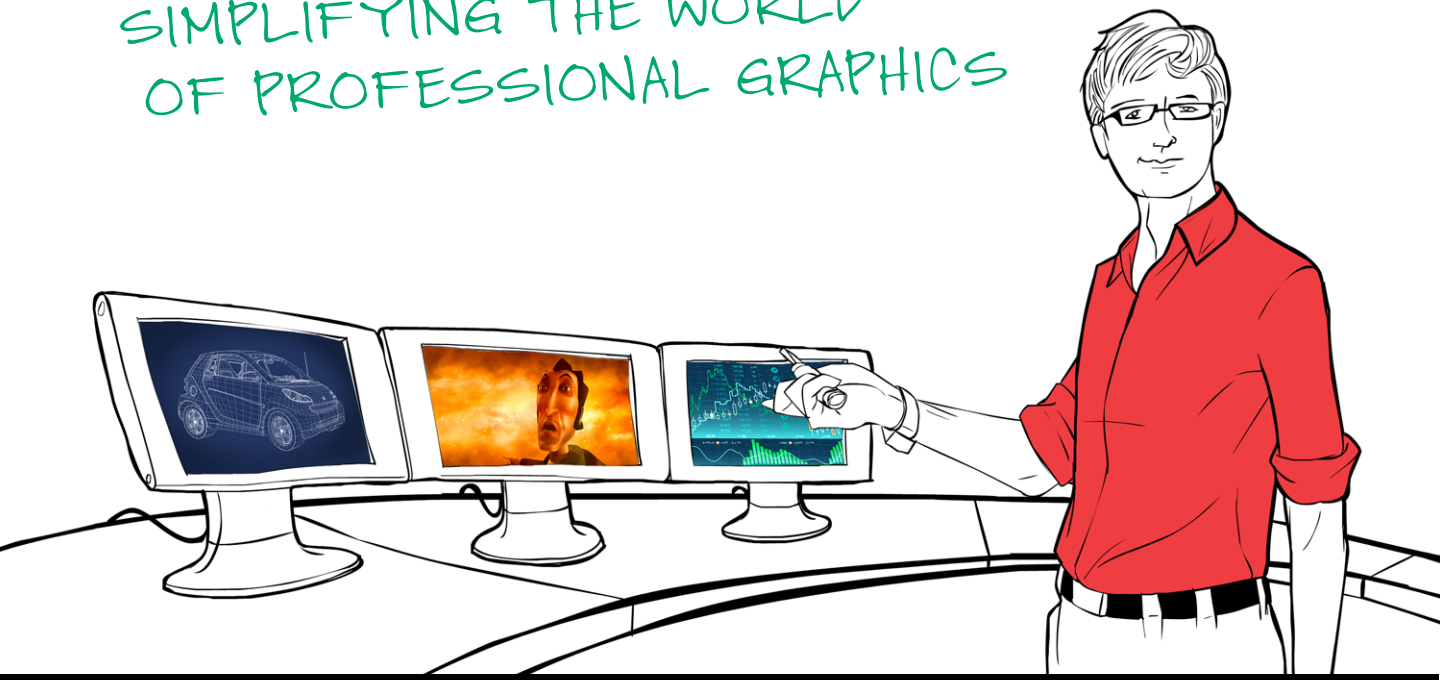

A reference for anyone interested in how graphics work and why they're important, from the experts at AMD.

The information age in hi-res. Technology is helping people and businesses get sharper. It's not enough to have the data — we need it in ever-increasing clarity. It can be the difference between looking at an x-ray and seeing a patient's complete physiology. It's wild-west oil and gas speculation versus targeted exploration and discovery.

It's a new way of seeing things. And it can change everything. You just need to choose the technology that will bring your data to life.

That's where AMD can help. Some of our most exciting advances have been made working with users and industries that demand more than traditional consumer and mainstream graphics. With our customers as our first source of inspiration, we've developed groundbreaking products for game-changing results.

In this guide, you'll discover the impact that AMD FirePro™ professional graphics can have on productivity. We'll show how AMD solutions are working right now for businesses in CAD / CAE, DCC, Scientific Visualization, Medical Imaging, Oil and Gas exploration, and Finance. We'll look at each solution from the client's unique perspective. And we'll introduce you to the latest AMD professional graphics solutions designed to drive very real business success.

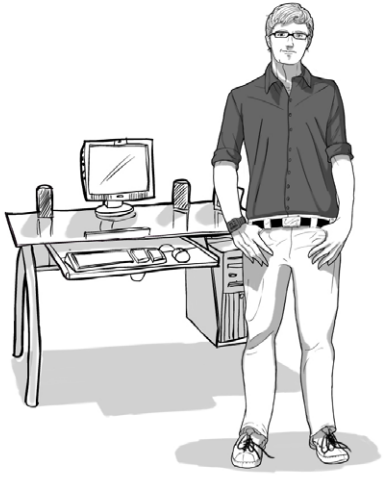

3D 101:
A Primitive Overview

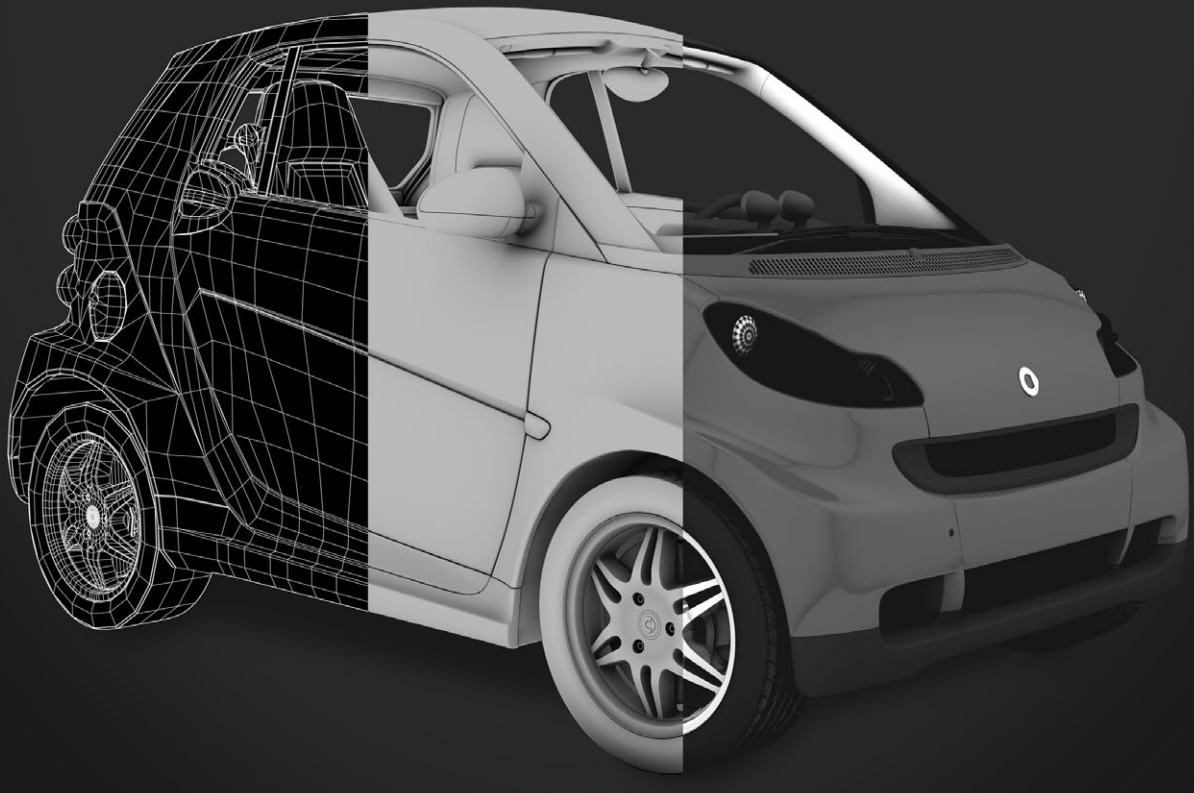

Seeing and being seen. Human beings are used to looking at the real world, which is infinite in size and detail. If you “zoom in” on an image in the real world, you see that image in greater detail.

If you “zoom in” on a digital image, however, you’ll find that its level of detail is less than infinite. Quite the opposite, really — because taking a closer look at a digital image only reveals and magnifies its flaws.

In the same way zooming in on a digital image reveals its limitations, the demand for greater detail from our digital images reinforces the fundamental issue with 3D graphics: No matter how far technology has advanced, computers are still finite.

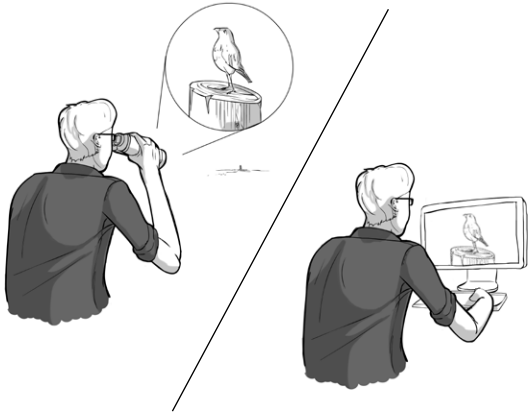

The only way to create, store and view a seemingly infinite 3D version of our “world” is to build and store a huge number of images — each formed from data that allows us to move, modify, or add light and texture to the building blocks we see on our 2D monitors. The problem, therefore, is about maximizing the level of detail in our images — while minimizing the amount of data to be stored and processed.

3D Graphics in Black & White

In this section, we’ll show you how 3D images are built from the ground up.

Breaking it Down

Picture a scene, and then isolate the objects within that scene. Each object has a set of properties that determine how it looks and behaves. That’s because all digital objects are merely sets of vertices. The more vertices you have, and the more attributes you assign to those vertices, the more detailed the object.

Now look closer and you’ll discover that these vertices and objects are formed from standard building blocks known as primitives (see figure 1). It’s an efficient system, since objects can be created using common shapes rather than custom shapes. That means the building blocks (read: the data) are the same no matter what the object looks like, and many of an object’s attributes can be duplicated too.

The most commonly used primitive is the triangle, which is defined by three vertices (the minimum number required to form a polygon). By joining multiple primitives, it’s possible to create virtually any shape. Beyond the primitive, there is the pixel. Pixels (picture elements) are the smallest building blocks of any image. Each pixel consists of a location (co-ordinate) and a color value. Color depth — the maximum number of colors a pixel can be — is expressed in bits per pixel (bpp).

FIGURE 1:
PRIMITIVES

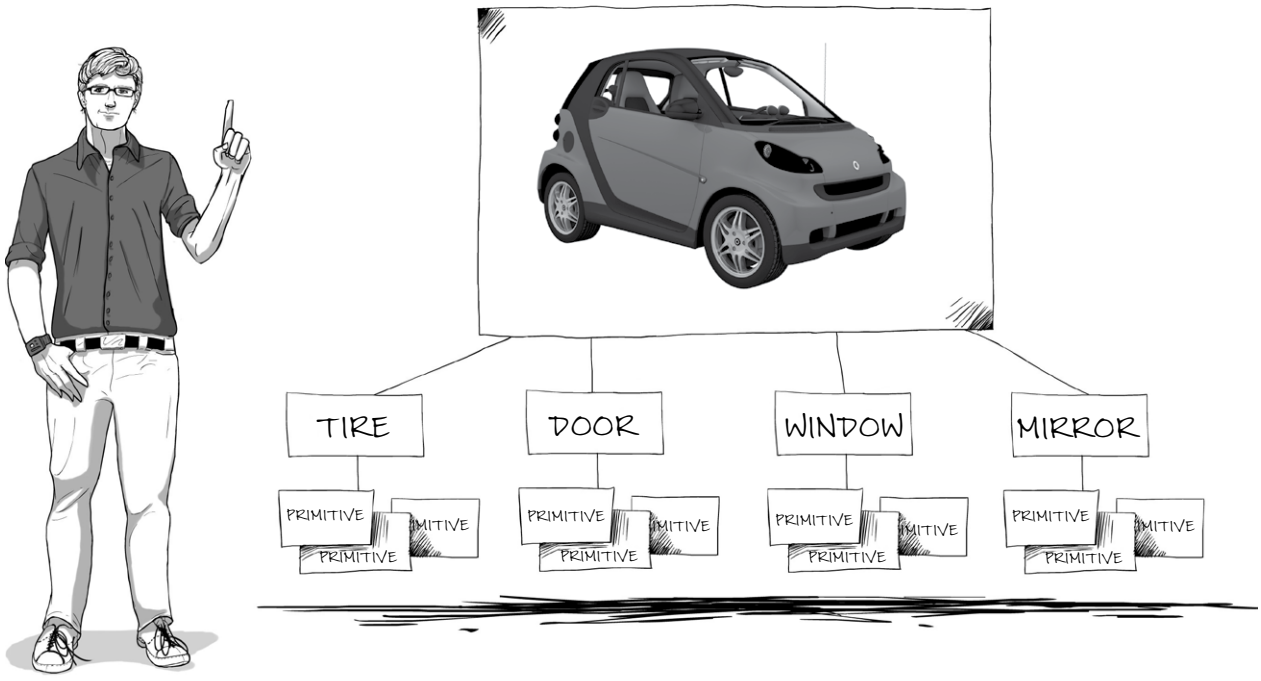

Making the Shade

Pixels and primitives are only part of the 3D picture. Once those elements are created, you have to apply characteristics such as light, shading and texture.

Early 3D graphics tools relied on stored libraries of textures. But the realism of the images was limited by the static nature of the textures. To render more realistic images, you need to use a shader — a small program uploaded to your GPU.

In the past, there were separate shaders for pixels and vertices — each working to process different types of objects. But modern GPUs use a common shader engine based on Shader Model 4.0 to process vertex, pixel, and geometry data simultaneously. This approach eliminates idle time, and allows each shader to consistently operate at full capacity.

The Big Picture

Bottom line, even with a sharp, efficient set of blueprints for building 3D images, all of those pixels, primitives and vertices add up to a lot of data. It takes real processing power to deliver realistic graphics. And that means a serious Graphics Processing Unit (GPU).

Your GPU and You

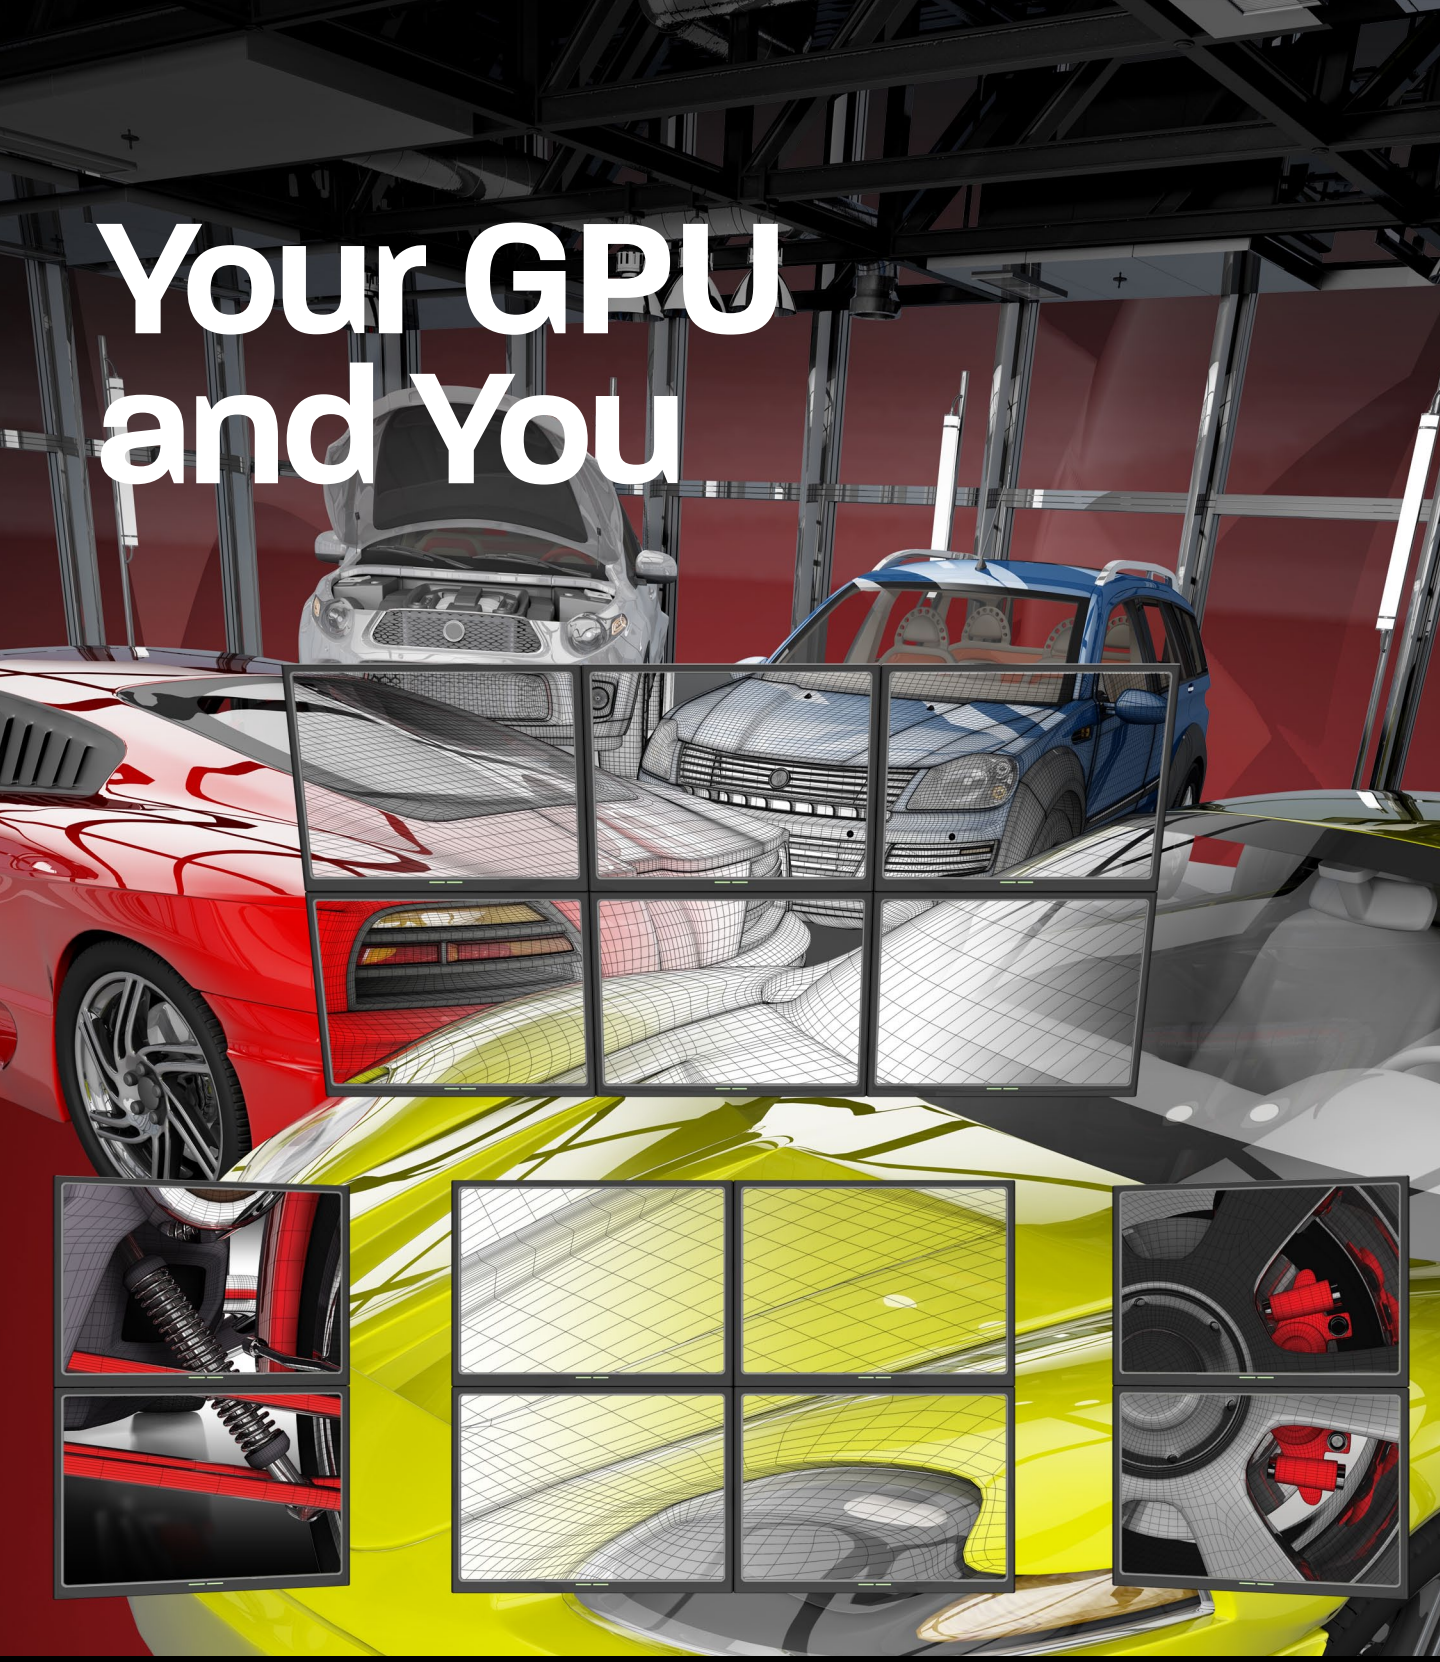

The Graphics Processing Unit (GPU) is the heart, soul, brains and brawn of a professional graphics solution. As shown in Figure 2, a typical GPU is made up of five key components: the 3D engine, the memory controller / memory interface, the bus interface, the 2D engine and the video and display engine.

> 3D Engine: In this guide, we're most concerned with the 3D engine within the GPU. It's a set of processors that handle the mathematical operations required to render 3D graphics. Technically speaking, "rendering" is what happens when a set of data points representing a 3D image is modified by user input or other parameters to create a 3D image on a 2D display.

> Memory Controller/Memory Interface: This is the circuitry that keeps the various processors in the GPU fed with data. It manages the flow of data onto and off of the chip, and it communicates with the local graphics memory.

> Bus Interface: The bus interface enables communication between the GPU, CPU, and system memory. On modern systems, this would happen via the PCI Express® bus, which connects to the north bridge.

> 2D Engine: This component of the GPU handles basic 2D operations such as drawing lines, boxes, and the traditional non-3D that appears on the screen.

> Video and Display Engine: By scaling, color correcting and otherwise processing the data, this engine formats the output image before it reaches the end display. It also performs video decode acceleration and post processing.

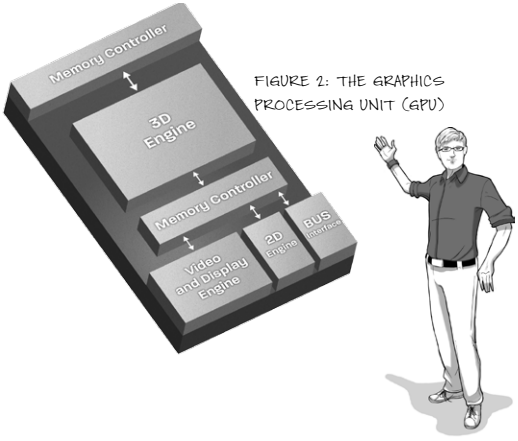

FIGURE 2: THE GRAPHICS PROCESSING UNIT (GPU)

Which API Are You?

While GPUs are capable of rendering complex 3D imagery, each does so in different ways. Plus, over time, new features are added to core rendering engines that allow for additional shading and texture refinements.

Developers typically don't want to write a separate rendering engine for each and every chip on the market. That's where graphics Application Programming Interfaces (APIs) come in.

APIs: Accelerating Applications
APIs make things easier for application and game programmers by allowing them to tell processors how to accelerate the rendering of 3D images and computational tasks using a standard set of parameters. Software included with the graphics card translates the instructions to the processor, and also uses acceleration when possible to improve hardware performance.

There are three dominant APIs: OpenCL™, Microsoft® DirectX® and Open Graphics Library (OpenGL).

> OpenCL™ OpenCL is quickly becoming a de facto standard API for accelerating a broad range of applications across a multitude of compatible CPUs, GPUs, and APUs. A cross-platform, open-computing standard, OpenCL allows software developers to create powerful professional design and engineering applications that can offer end-users exponential increases in performance and drastic improvements to efficiency and productivity. The OpenCL standard is maintained by the non-profit Khronos Group. Applications developed using OpenCL can run on a wide range of modern computing devises including most PCs, tablets/mobile devices, etc. OpenCL also has great potential for web/cloud-based computing as this segment continues to grow quickly.

In June of 2010, OpenCL 1.1 was ratified by Khronos Group and offered vastly advanced function and capabilities over the original OpenCL 1.0 specification. Since then, software developers in the DCC and CAD marketplaces have moved to quickly embrace OpenCL, and as of late 2012, there are numerous professional OpenCL-accelerated applications either available today or currently being developed for release in 2012.

> Microsoft® DirectX® As an API standard that's owned and controlled by Microsoft, DirectX is native to Windows systems only — having been implemented in various versions since Windows 95. DirectX is generally acknowledged as the dominant platform for games programming, and has been implemented directly into the Microsoft Xbox and Microsoft Xbox 360. (See Figure 3.)

> OpenGL is the premier environment for developing portable, interactive 2D and 3D graphics applications. It is the most widely used and supported 2D and 3D graphics API, bringing thousands of applications to a wide variety of computer platforms.

Applications such as games, CAD, virtual reality or animation systems rely on the long list of OpenGL functions to produce complex 3D objects from simple building blocks (See "Primitives," Section II.). The key difference between OpenGL and DirectX is implementation. While DirectX is available only to Windows-based systems, OpenGL has been implemented for Windows, Mac OS X, and various forms of UNIX (including Linux®). With its broad availability and range of features, OpenGL has evolved into the standard API for professional graphics.

OpenGL 4.2 is the latest version to be released by the OpenGL Architectural Review Board (ARB) under the auspices of the Khronos Group.

The feature set and hardware requirements for OpenGL 4.2 are roughly equivalent to those of DirectX 11 (the current version of DirectX). One of the most significant innovations is a high-level shading protocol called the OpenGL Shading Language (GLSL). This intermediate coding gives programmers a more direct conduit into the graphics pipeline so they can render objects without implementing low-level programming routines (or using less sophisticated assembly language) outside of OpenGL.

OpenGL 4.2 is designed to be backward compatible with previous versions, and includes these new features:

> Support for shaders with atomic counters and load/store/atomic read-modify-write operations to a single level of a texture.

> Capturing GPU-tessellated geometry and drawing multiple instances of the result of a transform feedback to enable complex objects to be efficiently repositioned and replicated.

> Support for modifying an arbitrary subset of a compressed texture, without having to re-download the whole texture to the GPU for significant performance improvements.

> Support for packing multiple 8- and 16-bit values into a single 32-bit value for efficient shader processing with significantly reduced memory storage and bandwidth.

For full details, please visit www.opengl.org/documentation/current_version

FIGURE 3: AN APPLICATION PROGRAMMING INTERFACE (API)

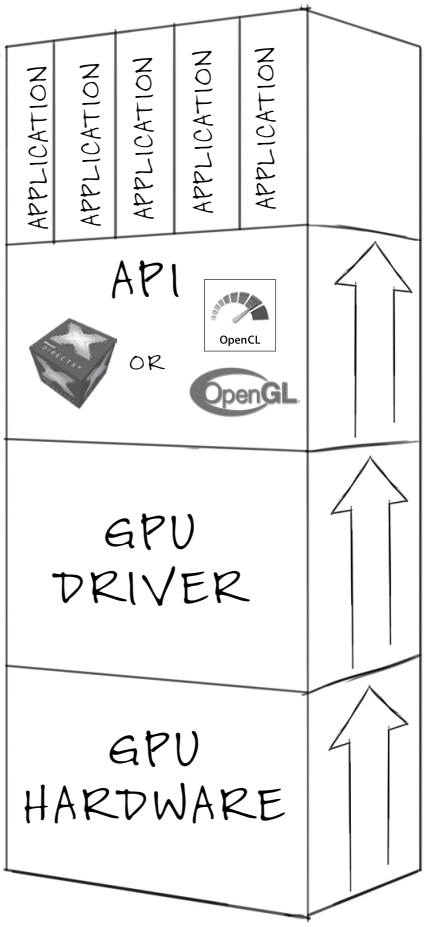

APPLICATIONS: THE TOOLS THAT LET YOU WORK YOUR MAGIC (AND SEE THE RESULTS)

APIS: RENDERS IMAGES USING STANDARDIZED SETS OF PARAMETERS, THEREBY KEEPING PROGRAMMERS FROM RE-INVENTING PICTURES OF WHEELS

GPU DRIVER: SENDS INFORMATION BETWEEN YOUR APPLICATIONS AND THE GPU

GPU HARDWARE: THE BACKBONE (OR, AS IN THIS ILLUSTRATION, THE BOTTOM CUBE) OF THE PROFESSIONAL GRAPHICS SYSTEM

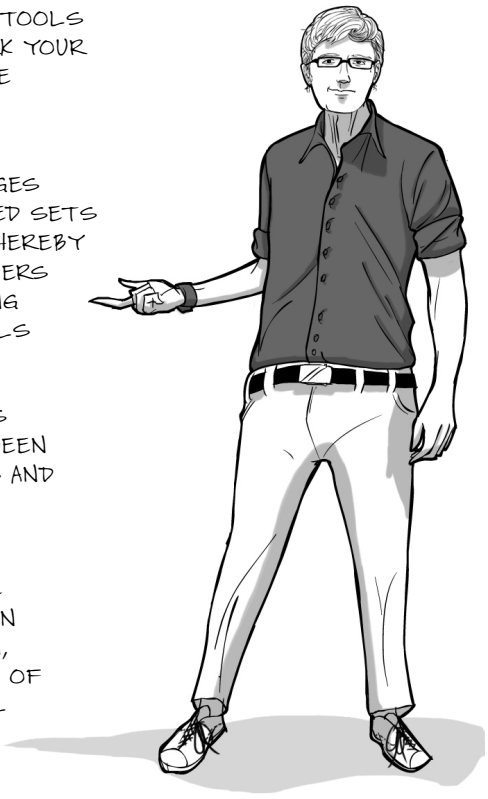

Watch This:

Cool Stuff GPUs Can Do

GPUs are fun. But get them together with a few of their friends, and they get even more interesting. Many of today's special effects, broadcast-quality video and other eye candy are the result of tweaking, tinkering and generally pushing the limits of what GPUs can do. Remember — these are professional graphics. Don't try this at home.

Multiple GPUs
The concept is simple: two brains are better than one. So, by linking GPUs together, users can get higher resolution, faster rendering and overall enhanced performance. There are several ways to make it work:

> Alternate Frame Rendering (AFR)
In this mode, all even frames are rendered on one GPU, and odd frames are rendered on the other. AFR is a good way to boost performance in geometry-limited applications.

> Split Frame Rendering (SFR)
By dividing each frame into two sections — with each section being processed by one GPU — SFR improves performance scaling for VizSim, broadcast, video and other fill-rate limited applications.

> Full Screen Anti-aliasing
Here, two GPUs are used to render the same frame. While both processors have anti-aliasing enabled, each uses a different anti-alias sample pattern. When both versions of the frame are completed, they are merged and sent to the display buffer on their way to the display. The resulting image has effectively twice the number of samples as an ordinary anti-aliased frame.

Applications that require a high level of image/line quality — including digital mock-up (DMU), optical simulations, TV production tools, or medical imaging apps can benefit greatly from full screen anti-aliasing.

Video Synchronization
On a video wall built with multiple small displays, the illusion of one large image is only achieved if the output signals to each monitor are in sync. The same idea applies when computer-generated graphics are combined with video from an external camera. In either case, it takes special processing to synchronize the video signals.

The two methods commonly used to synchronize multiple video signals are genlock (generator lock) and framelock.

> **Genlock** syncs video outputs using an external reference signal generator. This ensures that output sources from multiple devices (such as workstations, cameras, or video recorders) are combined or switched properly. Without genlock, constant switching between sources may cause a picture to jump or to be momentarily lost as the device tries to lock to a new signal.

> **Framelock** synchronizes the frames on multiple video displays. In this system, the hardware internally generates a unified signal and sends it to the monitors. The result is a large virtual canvas comprised of multiple screens. Today's 4K resolution (4096×3702) projectors, for example, require four DVI inputs to create one large 4K image.

Remote Graphics & VDI
Clients like government agencies or businesses with sensitive transactions and intellectual property can benefit from a remote graphics solution. Simply put, it means separating the CPU and GPU from the rest of a system's input / output devices (keyboard, mouse, speakers and displays). On a remote graphics system, critical data is highly secure, because data remains on the server, and only compressed video output is sent for display across the network — and heat and noise levels are significantly reduced.

With this kind of system, user interfaces are connected to a small, central client box which is then tied back to the CPU and GPU. There are three ways the system can then host graphics: integrated graphics and IP compression, MXM form factor graphics for remote blades, and a standard professional graphics card for remote PCIe® slots.

VDI, or "virtual desktop infrastructure," is a new trend in corporate computing that allows a single server unit to run dozens of "virtual" Windows or Linux OS-based desktop computing instances simultaneously. Remote workers connect to their virtual desktop computing session over a network-connected, passive "thin client" device at the user's desk. Using a standard keyboard, mouse, monitor, speakers and other standard peripherals connected to the "thin client" portal device, the VDI user connects remotely to his or her "virtual computer" and enjoys a rich multimedia computing experience almost indistinguishable from one that he or she would get working directly on a PC at their desk.

FIGURE 4: VIDEO SYNCHRONIZATION

Professional Graphics: Who Needs Them?

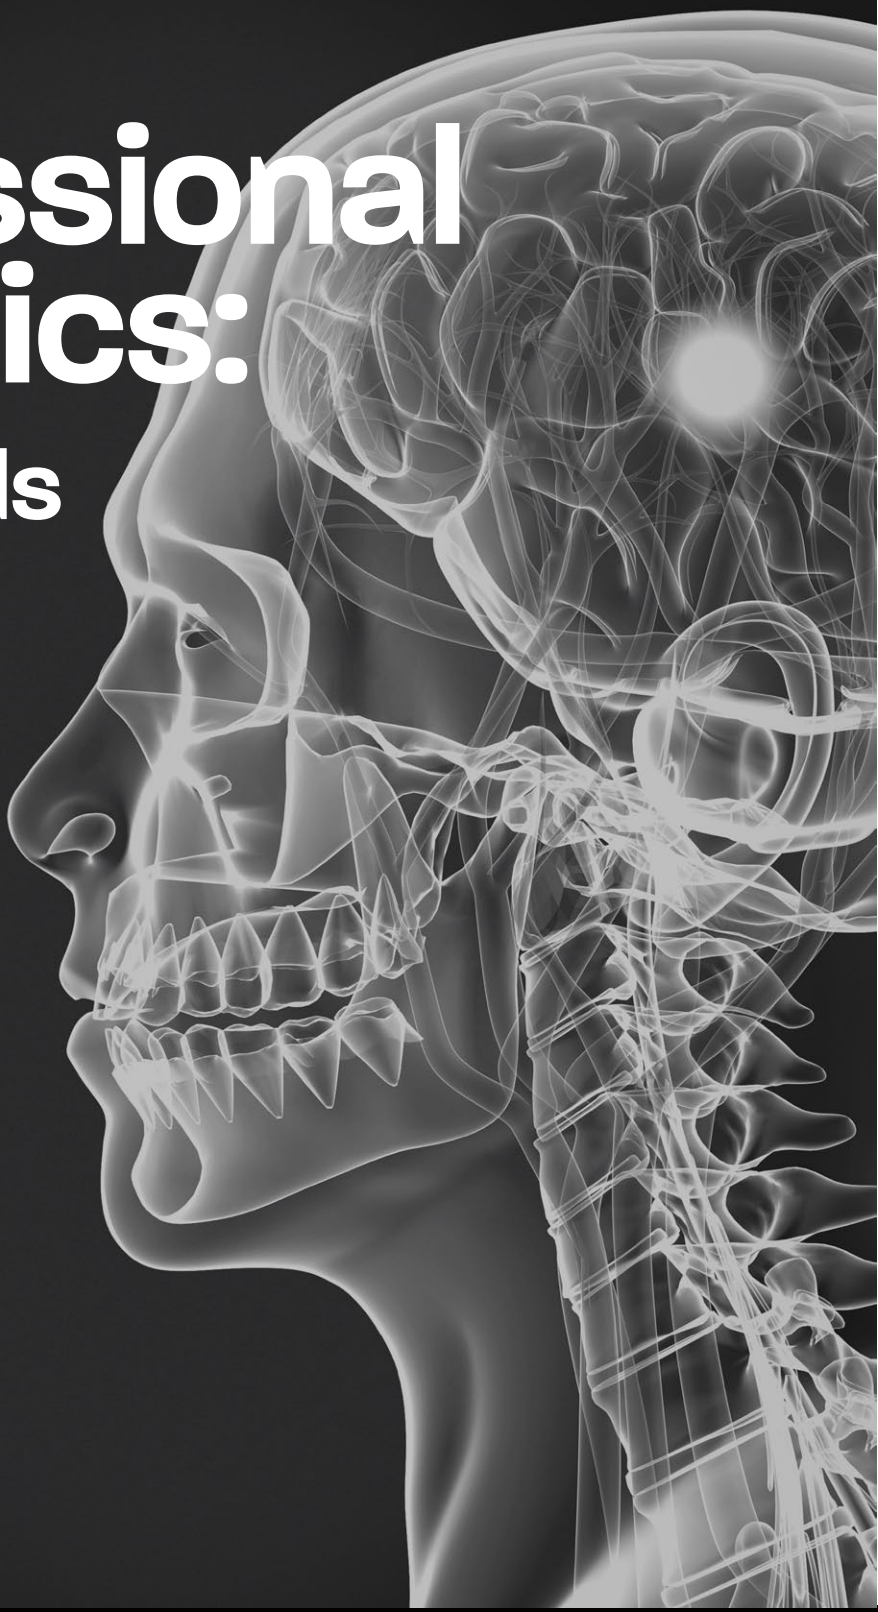

As one might expect, there's a difference between the graphics needs of an animator at a Hollywood movie studio and those of a self-employed accountant who enjoys the occasional PC game.

Naturally, both kinds of users expect performance, and AMD delivers. In either case, we're dealing with speed, clarity and user experience. So, what separates one group from the other? It's more about how AMD technology meets users' demands in the long- and short-term.

Don't Try This at Home

Typically, PC users in the consumer and mainstream business segments are looking to enhance applications — including multimedia players, online video streaming or digital photo tools. At the same time, PC gamers want the fastest performance with the latest game titles. They're always upgrading to the latest and greatest — looking to stay ahead, and to gain a competitive advantage.

For the vast majority of PC users, the AMD Radeon™ family of products are ideal for home use, office applications, and gaming.

Professional Graphics: The Binary Truth

Like mainstream users, those who work with professional graphics demand application performance, realistic imagery and reliability. But the similarities end there. Because in industries like construction, science, healthcare, finance or oil and gas, 3D graphics can mean the difference between success and failure, profit and loss — and even life and death.

The right applications help these users do what they do better, faster and more consistently. So they need the right professional graphics solutions to run those applications better, faster and more consistently, too.

Due to the complexity of 3D graphics systems, professional users also want technology that's designed to stay ahead of the curve — so they don't have to upgrade every time something new hits the market. And since they're making a serious commitment to cards, CPUs, displays and other components, they need reliable, ready support to protect that investment and keep things running smoothly.

Build It First and Make it Last

Professional graphics is a different playing field. And as such, the timing for new products and other innovations is often very different from the usual consumer and mainstream business sales cycle. In professional graphics, it's about building products designed to last — with fewer upgrades and broader compatibility.

At AMD, we don't wait for the traditional PC sales cycle. Instead, we focus on the product life cycle. Professional users need the right technology at the right time with the right certifications. And that's what we deliver.

We Say It With Inflection

Professional graphics users look to AMD as an innovator. And we time our releases with global “technology inflection points” — those places in the ebb and flow of technology where our solutions can achieve the greatest impact for our users, and for technology as a whole. In other words, when we create a game-changing solution, we make it available as soon as possible, we build it to last as long as possible — and we let the industry catch up.

AMD was the first to:

> Supply AGP in quantity

> Lead the transition with native PCIe support

> Provide a notebook 3D graphics accelerator

> Provide discrete graphics products supporting both video and sound solutions through HDMI

> Offer a 90nm and 55nm process GPU package

> Support Windows Vista® WHQL drivers

> Receive VESA DisplayPort certification for PC graphics

> Provide GPUs supporting DirectX 11

> Release a graphics card featuring Graphics Double Data Rate, version 5 (GDDR5) memory

> Release a graphics card (with AMD Eyefinity technology) capable of processing over one trillion floating point operations per second and supporting BIP to Ce displays on a single graphics card

Benchmarking:
Real World Comparisons

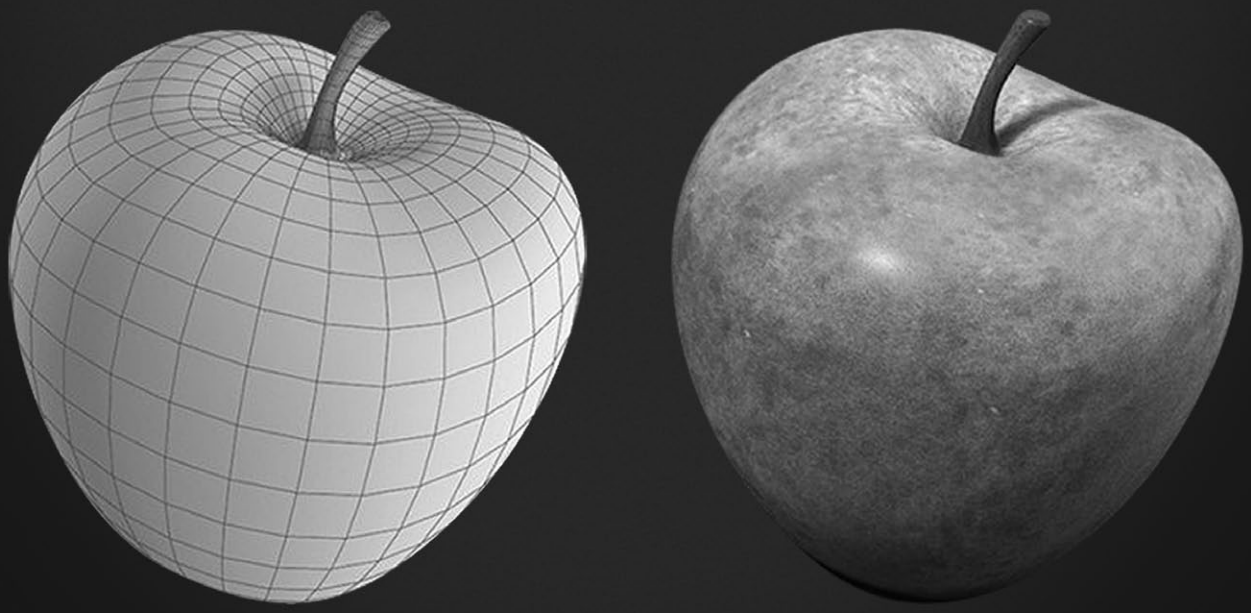

Benchmarking is the process of evaluating a product's relative performance, cost, or quality against others of a similar type.

In the consumer and mainstream business markets, graphics cards are often benchmarked based on frame rates. Professional graphics products, however, are measured using real-world applications in a number of user categories — such as engineering and design, digital content creation (DCC), or life sciences.

The non-profit Standard Performance Evaluation Corporation (SPEC) has established a standardized set of tools and methodologies for measuring the performance of professional graphics technology. Their goal is to help vendors and customers make informed decisions based on unbiased and unambiguous evaluation.

In other words, benchmarking gives users a frame of reference for comparing professional graphics solutions. That way, customers can choose the tools that will perform best for their unique applications.

SPECviewperf®

For graphics subsystems under OpenGL, SPEC designed a benchmarking tool called SPECviewperf®. Through a collection of applications — each with datasets submitted by independent software vendors (ISVs) in the form of "viewsets" — SPECviewperf tests and verifies a system's rendering performance. Each viewset represents the datasets and rendering parameters used in real-world applications.

The following viewsets make up SPECviewperf:

> 3ds Max
> CATIA
> EnSight
> Maya
> Pro/ENGINEER
> SolidWorks
> UGS Teamcenter Visualization Mockup
> UGS NX

Benchmark results from SPECviewperf can be used to compare the performance of professional graphics products. At the same time, SPECviewperf can indicate how well an application will perform under one of the given configurations.

SPECapc

To further evaluate professional graphics system performance, SPEC has also developed the Application Performance Characterization (SPECapc) standard. Like SPECviewperf, SPECapc benchmarks are based on actual software application performance. But SPECviewperf only measures the graphics subsystems, whereas SPECapc tests the overall system — including graphics, I/O, and CPU.

The following SPECapc benchmarks are currently in use:

> SPECapc for 3ds Max
> SPECapc for Maya
> SPECapc for Pro/ENGINEER Wildfire
> SPECapc for Solid Edge
> SPECapc for SolidWorks
> SPECapc for UGS NX

While SPECapc is updated continuously with new test scenarios, it may not be useful in every application scenario. In some cases, ISVs can use their own benchmarking tools.

ISVs: Short for "Certified Performance"

Professional applications are extremely demanding on computer systems — and they're especially tough on graphics subsystems. That's why AMD works with independent software vendors (ISVs) to help ensure applications are performing at their peaks. ISVs expose bottlenecks and suggest optimization updates to graphics drivers, and they use benchmarks like SPECviewperf and SPECapc to make sure each upgrade delivers maximum performance.

ISVs certify that their applications will run effectively on a designated graphics card. The process varies according to the kind of program being tested. With mechanical CAD programs, for example, vendors run a script of several thousand individual tests that check the program's commands. The automated script is designed to exercise all of the software's functions, simulate as many features as possible, test processing speed, and put the graphics card through its paces. The automated test may take as little as an hour or two.

To perform further manual testing, software engineers certify the graphics card driver level, push the system's endurance over a long period of time, perform multiple simultaneous tasks, and dish out whatever idiosyncratic punishment the software vendor feels is important.

Each ISV handles its testing process differently, setting its own requirements for compatibility. This makes the evaluation process different from familiar workstation reviews or independent party analysis — where hardware might be ranked by performance. Instead, ISV certification is about the reliable function of the application itself, whereas benchmarking measures card performance alone. Basically, if you want certification, you need an ISV to run your application through its full functionality.

Beyond Optimized: AMD and Automatic Application Detection and Optimization

In order to support the wide range of professional applications, graphics adapter software drivers must include all of the optimization paths for all of the programs. In the past, users configured their graphics drivers by choosing the appropriate optimization path with the application running. But that meant whenever the user switched to another application (or rebooted), the optimization had to be reconfigured manually.

As of ATI FirePro™ driver version 8.263, AMD introduced a new feature called Automatic Application Detection and Configuration (AADC). This feature lets the graphics driver automatically detect running applications and configure the graphics driver settings — thereby optimizing the application in real time. As the user switches between applications, the correct optimization settings are adjusted for the active, running application. The process is completely automated, with no user intervention required. The table below lists examples of the applications supported by AADC.

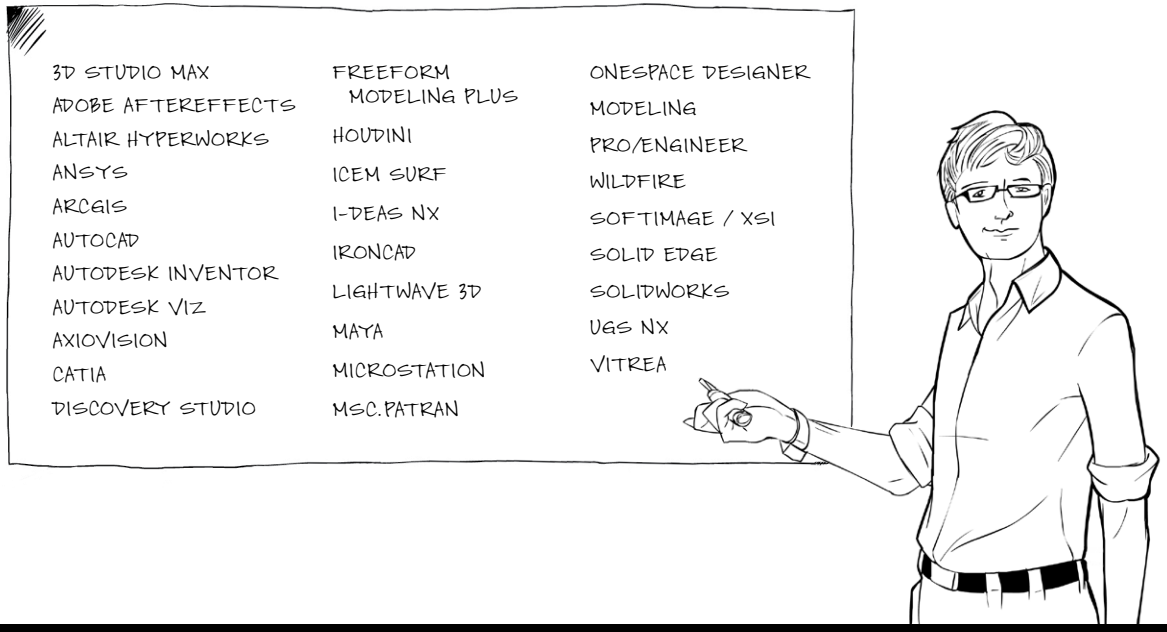

Professional Graphics at Work

Because professional graphics applications are so different — and the requirements so demanding — AMD develops professional graphics products for highly specific end users and market segments.

In this section, we'll summarize the major markets that drive innovation and push AMD to create some of the most advanced 3D professional graphics tools available today. You'll see what each segment needs from their professional graphics — and how AMD responds when performance matters most.

CAD and CAE: Building on Solid Data

Architects and engineers use professional graphics systems to draw, render, model and test their designs. Computer-aided design (CAD) tools are used for drafting and producing other architectural documents, while computer-aided engineering (CAE) software helps

engineers simulate and test everything from building materials to wind and lighting.

CAE, for example, can be used to simulate structural and environmental conditions that might affect the strength of a bridge or a staircase. Programs can validate the choice of materials (i.e. steel, wood or concrete), and engineers can optimize their designs before the project ever breaks ground. The same is true for CAD — where designers can build anticipation and emotion with an interactive, 3D, real-time model of a proposed office space or loft apartment.

In both cases, maximum rendering speed and absolute realism can take a winning design to the next level — or help prevent a catastrophic structural failure. And it goes without saying that testing designs and structures can give architects and engineers a huge boost in productivity — by reducing analysis time and even speeding time to market.

Some notable examples of CAD/CAE applications include:

- Siemens PLM NX
- PTC Creo™
- Dassault Systèmes CATIA
- Dassault Systèmes SolidWorks
- Autodesk Inventor
- Autodesk AutoCAD

DCC: Where the Magic Happens

Digital content creation (DCC) involves image editing and animation. In this industry, it's all about delivering the most realistic, creative or otherwise stunning visual imagery to give viewers a memorable experience.

To create the illusion of movement, animators display an image on a monitor, and then quickly replace it with a new and slightly different image. In 3D animation and graphics, life-like models are built on virtual skeletons — and their limbs, eyes, mouths, and clothes are moved or changed on key frames. The graphics system calculates and renders the differences in appearance between these key frames — giving the appearance of motion.

In the realm of special effects, animators often rely on framelock or genlock functionality. In either technique, computer-generated images are synched with broadcast television or other video input.

No matter what, when you're talking about digital content creation, you're dealing with huge demands on graphics subsystems. The applications require graphics cards with powerful frame buffers and lots of memory bandwidth to handle the massive models, scenes, textures and datasets.

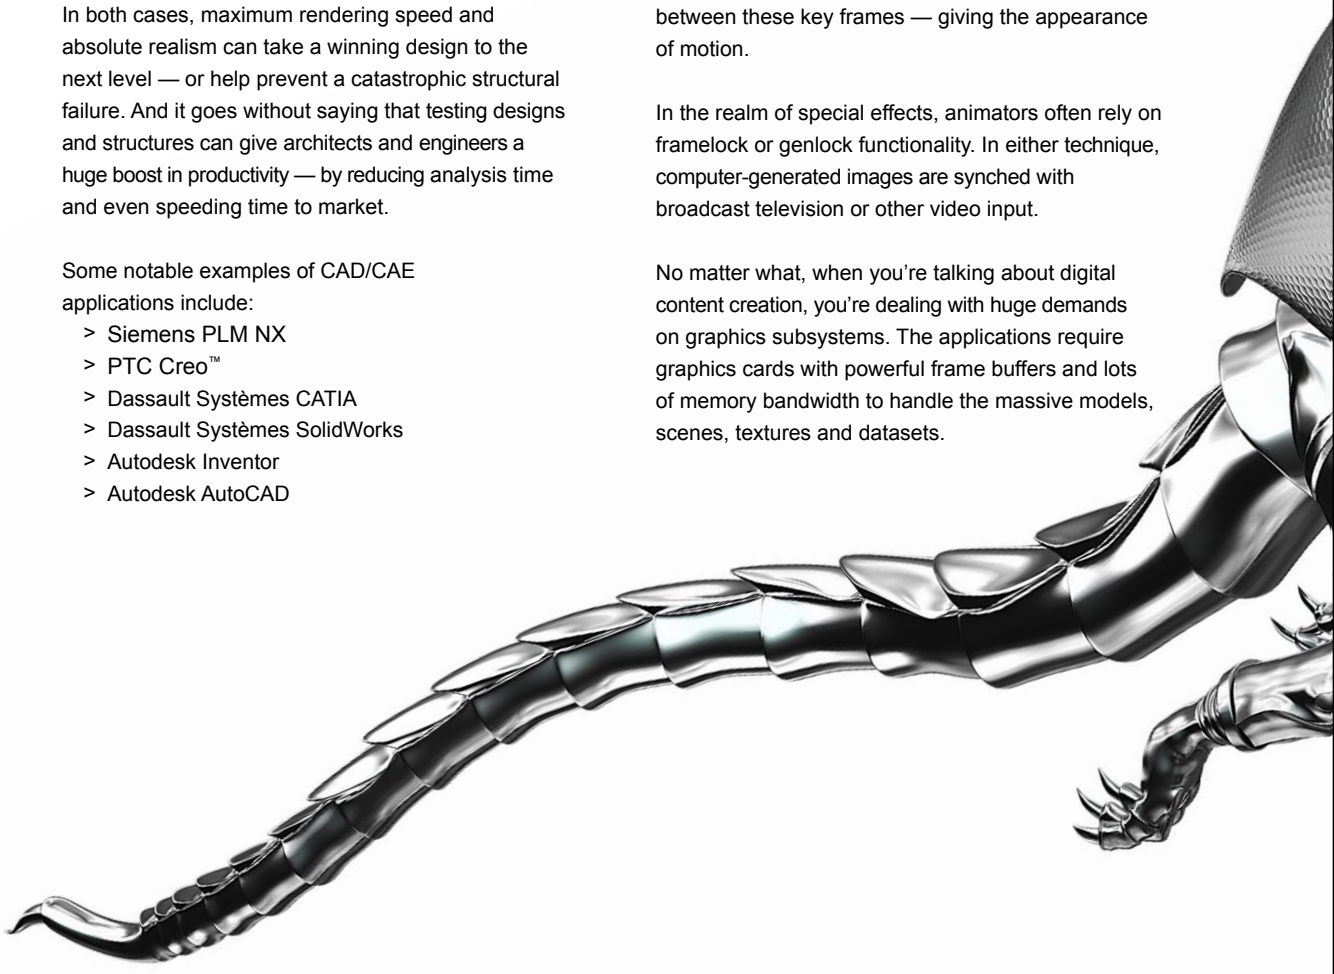

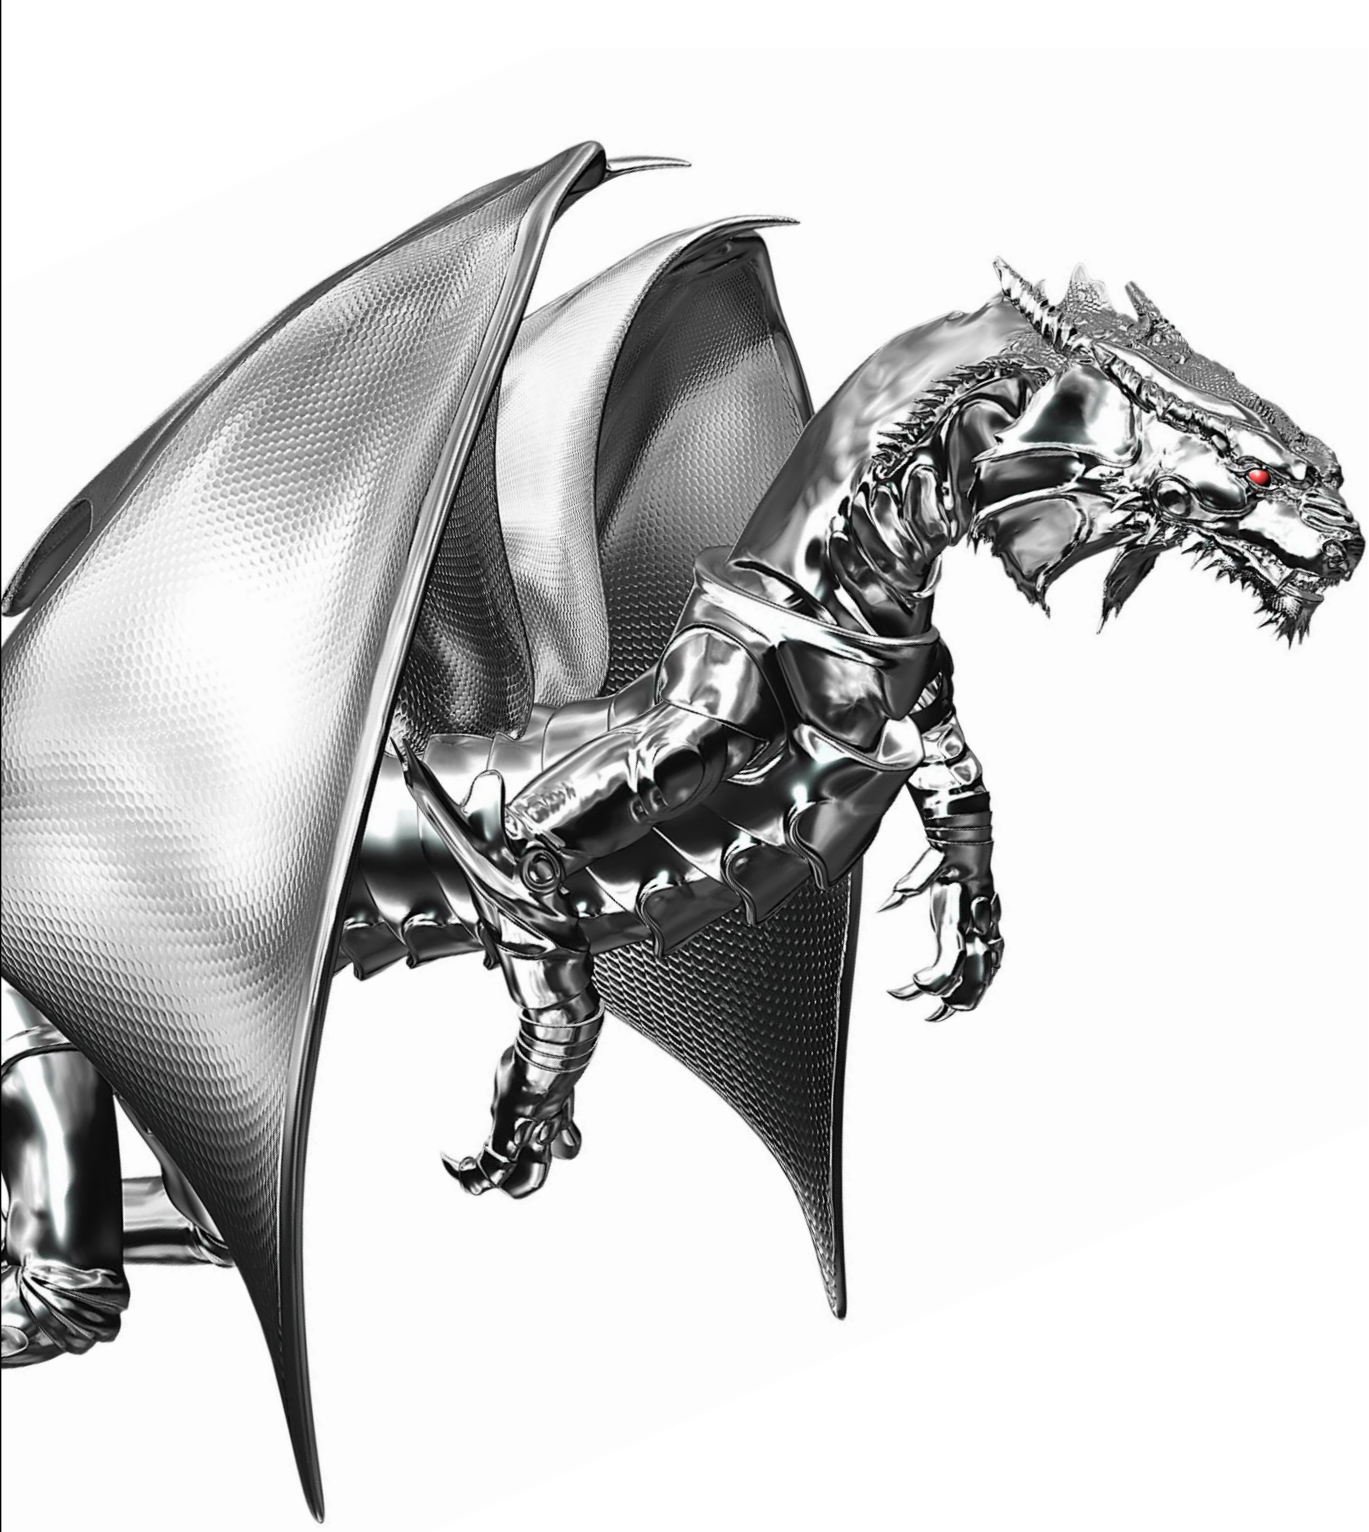

Scientific Visualization: Virtual Pictures for Real Answers

In the field of scientific visualization, experts call on computer graphics to create 3D images of complex, often massive data sets. These images range from simulations of fluid dynamics to detailed MRI scans.

In natural science, for example, visualization tools are used in molecular rendering — which helps scientists judge the potential effectiveness of medicines by forecasting interactions between drug molecules.

Geographers and ecologists use detailed terrain rendering programs to look for patterns, features, relationships, and anomalies in the Earth's surface over time. With this data, they can study water levels, erosion and other phenomena.

These simulations can require anything from programs that run in a few minutes to network-based groups of computers running for days. And as you might expect, it takes graphics cards with large frame buffers and high memory bandwidth to make these massive datasets meaningful.

Medical Imaging: Seeing the Data Helps Us Find the Cure

With medical imaging, it's less about the quantitative datasets (like those involved in rendering a 3D model) and more about the qualitative outcomes. In other words, professional graphics are designed to give healthcare experts the information they need (quantitative data) to make accurate diagnoses (qualitative results).

A radiology team at a hospital, for example, might create dozens of images of a human body for study and clinical purposes. A professional graphics system can take those images and produce a 3D model. By using 3D visualization tools to examine that model, doctors are then able to diagnose and treat any number of pathologies in their individual patients. Such electronic anatomical atlases are available to physicians today — thanks in part to hardware and software that can render the data with crucial accuracy and detail. One of the most important facets of medical imaging, in fact, is the development of 10-bit output — which can reproduce more than 1 billion colors. Often, radiologists work in grayscale rather than color. HD bit output is required for specific modalities such as mammography where the more precise range of grayscale can prove more informative for difficult diagnoses.

Key providers of medical software and hardware include General Electric, Philips, Siemens, Agfa, and Toshiba.

Oil and Gas: Seismic Matters

Petroleum geologists searching for oil are really trying to locate hydrocarbons in the Earth's subsurface. That's because the presence of hydrocarbons indicates the safest and most efficient places to drill wells. And since efficiency is all about saving time and costs, the goal is to find the best tools to help find those hydrocarbons — fast.

Enter seismic interpretative applications that render images of subsurface geological structures from recorded seismic data.

Notable examples include:

- Schlumberger applications
- Paradigm applications
- Landmark GeoProbe

Oil and gas exploration companies that run apps like these depend on graphics cards like AMD FirePro™ to help speed up the imaging and eliminate the guesswork.

Financial: These Charts are So Money

Major financial institutions involved in investment banking, equity trading, and investment management are managing tons of critical transactions at once — and that often means multiple displays. That means applications have to present lots of tables, charts, and graphs — each in their own windows, and all at a glance.

Huge financial institutions generally make smart purchases on graphics systems and are ultra-conservative with costs like electricity. So whatever desktop computers, docking stations, laptops or display configurations the institutions run, they can find graphics cards designed to fit those devices that are highly efficient. That's why AMD FirePro™ Multi-View products offer low-power, low-profile options. They're designed to give users dual- and quaddisplay capabilities in a variety of devices — and to conserve energy in the process.

Where the Pros Go:

AMD FirePro™ Professional Graphics Solutions

→ MachStudio Pro combines with AMD FirePro™ 3D workstation graphics to provide 3D content creators with efficient workflow at the speed of thought.

As outlined in this guide, professional graphics solutions build on familiar 3D graphics concepts and applications — and then take them to an unprecedented level of image quality and performance.

At AMD, our customers drive our innovations. The demands placed on professional graphics tools by a variety of end users inspire us to create unique and purpose-built products that meet or exceed their rigorous standards.

This section introduces the major product categories in the AMD FirePro™ professional graphics lineup — based on the markets for which these solutions are targeted.

For specifications on the latest products, please visit: http://www.amd.com/us/products/workstation/graphics/Pages/workstation-graphics.aspx

AMD FirePro™ 3D Professional Graphics

AMD FirePro 3D professional graphics are designed for professionals who demand the highest quality, reliability and application performance in their profession. AMD FirePro graphics products are fueled

by the combination of leading-edge GPU technology, driver software and comprehensive ISV certifications, and performance across industry segments.

Ideal for CAD, DCC and simulation markets, AMD FirePro 3D professional graphics deliver industry leading features and price performance:

- DisplayPort output for superior 30-bit image quality
- Scalable ultra parallel processing architecture with up to 800 unified shaders
- AutoDetect instinctively optimizes performance for multi-application workflow
- Dual Link DVI output for driving ultra high resolution widescreen monitors
- High Dynamic Range (HDR) rendering with 8-bit, 10-bit, and 16-bit per RGB color component support
- Hardware acceleration of DirectX® 11 and OpenGL 4.2 advanced features
- Optimized and certified for many CAD and DCC applications
- PCI Express® 2.0 compliant
- Maximized Application Interactivity
- OpenCL™ 1.1 support

Thanks to their unique ultra parallel processing architecture, AMD professional graphics maximize throughput by automatically directing graphics horsepower where it's needed. The intelligent management of computational resources means the graphics accelerator is used most efficiently — so users get real-time rendering of complex models and scenes, and increased frame rate performance at the same time.

Seeing More and Doing More with Your Data

AMD FirePro 3D graphics feature a full 30-bit display pipeline, along with support for High Dynamic Range output. The result produces over one billion colors — the most vibrant visual fidelity available. By comparison, the traditional 24-bit display pipeline produces only 16.7 million colors. And with DisplayPort and Dual Link enabled DVI outputs, AMD FirePro 3D graphics cards can generate a multimonitor desktop over 15360 pixels wide.

Interacting with Huge Datasets

AMD FirePro 3D graphics store geometry internally, rather than pushing data back to the CPU to calculate changes. With larger-scale models, this advance can help boost rotations by up to two times over current technology. Plus, AMD FirePro 3D graphics products use AMD's unique AutoDetect technology. As users open 3D applications or move between them, AMD FirePro graphics automatically configures the optimized driver settings for maximum performance in a multitude of workflow scenarios.

AMD Eyefinity Multi-Display Technology: See the Future on Multiple Monitors

The newest revolution in display configurations comes with AMD Eyefinity technology. It's a single GPU system that controls the entire workflow — and gives users new opportunities for handling data at a remarkable cost savings.

AMD Eyefinity technology manages up to six independent and simul-taneous displays per GPU — automatically correcting for bezels, varying monitor proportions, horizontal and vertical orientations and other parameters.[1]

The result: Multiple applications run seamlessly between displays. Visual simulations render with unbelievable speed across giant video walls with up to 6 monitors. And professional graphics users in every industry can easily configure a display setup that fits their unique needs.

Performance, Reliability, and Support

AMD FirePro 3D professional graphics are tested and proven with major CAD and DCC applications. AMD FirePro 3D unified drivers help simplify installation, deployment, and maintenance. And of course, AMD provides direct customer access to a dedicated professional-class technical support team.

AMD FirePro™ Multi-View

AMD FirePro Multi-View 2D graphics acceleration cards are designed exclusively for users in the corporate and financial markets. Here, system compatibility and data stability are critical not only to protecting the customer's technology investment, but to making sure their clients get the most accurate and timely information possible. That's why the AMD FirePro Multi-View family offers the option of dual- or quad-monitor cards based on PCI and PCI Express bus architecture.

Features & Benefits:

- Reliability
- Stability
- Flexibility
- User-friendly software
- Direct user support
- Minimal power consumption
- Space efficient low profile design
- Product longevity
- 3-year limited warranty on all AMD FirePro Multi-View cards
- Maximum Flexibility in a Unique, Compact Design

AMD FirePro Multi-View professional 2D graphics cards feature a unique, low profile dual monitor design — meaning all AMD FirePro Multi-View cards will fit into both low profile and full height systems. This translates to compatibility with a wide range of machines and fewer cards to qualify. The same goes for the new low profile AMD FirePro Multi-View quad monitor cards — their super space-saving architecture lets you choose even the smallest systems (including docking stations) and still achieve simultaneous output to four displays.

The AMD FirePro Multi-View design relies on passive cooling for a higher mean time between failures (MTBF) and exceptional reliability. Plus, its inherently low power consumption helps reduce your overall power load without sacrificing performance — especially for sites with a large installed user base.

Maximum Memory, Stability and Flexibility

AMD FirePro Multi-View cards contain either 32MB, 64MB, 128MB, 256MB or 512MB of dedicated on-board memory for each display for maximum stability and unparalleled performance (total on-board memory ranges from 64MB to 256MB depending on product configuration). For added flexibility, dual and quad AMD FirePro Multi-View cards can be combined to drive even more displays.

Finally, every AMD FirePro Multi-View card, including quad versions, meets the industry's highest standards for resolution:

Maximum monitor resolution for all AMD FireMV™ and AMD FirePro multi-view boards:	**Monitor specification:**
2560 x 1600	Display Port
2048 x 1536	Analog
1600 x 1200	Digital

AMD FireMV and AMD FirePro multi-view customers get three levels of helpful, knowledgeable support:

- Level 1: Direct telephone and email support
- Level 2: In-house problem resolution teams assigned to individual customer issues until they are resolved
- Level 3: In-field and on-site technical support as needed

Bringing Professional Graphics into Focus

We hope this guide has been useful in helping you understand the basics of 3D graphics technology. If you're a part of the world of professional graphics, then that's almost as much as you need to know. After all, whether you're using the tools to make the most of your data — or you're creating the professional graphics solutions that help businesses succeed — the rest of the story is up to you.

It's Your Future. See it with AMD.

At AMD, our customers drive our success. Our solutions are direct responses to real needs in the real world. And nowhere is this more evident than in professional graphics. Because behind the desire for all of this "virtual reality," there's the drive to meet concrete business objectives and to achieve measurable results. It's about saving time, saving money and even helping to save lives. And it doesn't get any more real than that.

And so, as the demand for more clarity, more speed and more performance continues to rise, AMD is there to push beyond the limits and to change the way we manage data and see our world.

www.amd.com/firepro

Images courtesy of University of Hertfordshire.

1 AMD Eyefinity technology can support multiple displays using a single enabled AMD FirePro™ professional graphics card; the number of supported displays varies by card model. Microsoft® Windows® 7, Windows Vista®, or Linux® is required in order to support more than 2 displays. Depending on the card model, native DisplayPort connectors and/or certified DisplayPort active or passive adapters to convert your monitor's native input to your card's DisplayPort or Mini-DisplayPort connector(s) may be required. See www.amd.com/firepro for details. SLS ("Single Large Surface") functionality requires an identical display resolution on all configured displays.