issue 02 // summer 07

Transformations

A quarterly publication from NVIDIA.

FEATURES

A WORD FROM MIKE

At our Analyst Day in June, Jen-Hsun and the executive team delivered a clear message: "The GPU is extending into new markets and devices." It was fitting that Analyst Day was the venue for the launch of Tesla, our new product line for high-performance computing. Tesla is a disruptive technology. It places an unprecedented amount of computing power into the hands of scientists and engineers, transforming today's workstations into 'personal supercomputers.' We believe Tesla products will be used to solve some of the most difficult challenges that exist today, in important areas like financial risk analysis, molecular biology, and medical diagnostics. With 128 parallel processors delivering up to 518 gigaflops of parallel computation, the GPU has truly become one of the world's most powerful processing technologies.

Michael W. Hara
VICE PRESIDENT
INVESTOR RELATIONS AND COMMUNICATIONS

Q&A
With Jen-Hsun Huang

PRESIDENT AND CEO, NVIDIA CORPORATION

Fiscal 2007 was a record year for NVIDIA. Why are GPUs doing so well in the market?

JH: The GPU (graphics processing unit) is becoming one of the most important processors in the computer. There are several reasons why this is happening:

First, computer graphics is becoming increasingly central to the computing experience. Information expressed in a visual way is more delightful, exciting, and captivating. Computer graphics technology enables application developers to realize their imagination and make us part of it.

Second, we have invested billions of dollars over the years in GPU R&D, driving a rapid increase in processing power, flexibility, and programmability. This has enhanced the versatility of the GPU and made it possible for all kinds of exciting applications to utilize its power – from enhancing photos and adding effects to home videos, to innovative computer vision applications that you will see in the future. In fact, many of the compelling new applications today are enabled by the GPU – from the rich user interfaces of the Mac to 3D maps from Google Earth.

The GPU is one of the most important technology choices for millions of computer users – from the gamer to the creative professional, and now, due to our initiatives in GPU Computing, the scientist endeavoring in fundamental research. Because the GPU is so central to today's computing experience, an increasing portion of the system budget is shifting to the GPU. I expect billions of GPUs to be sold over the next decade. We will serve the markets with our industry-leading brands: GeForce® for the enjoyment of digital experiences, NVIDIA Quadro® for the creation of digital content, and now Tesla for exploration and discovery.

> "We don't build things that our competitors are building; we build things we believe the world will need."

What are the implications of Tesla (NVIDIA's newest product line) in the high-performance computing arena?

JH: There has been a looming crisis in supercomputing. Fundamentally (since Cray's invention of vector processing) no new computing technology has been created to specifically address the challenges of processing a massive amount of data. And since the structure of the modern computer industry makes it unlikely that a computer company will invest billions in R&D to create a processor specifically for high-performance computing, processing speeds have fallen significantly behind the demands from scientists. Yet there are important discoveries that are just beyond our reach that can meaningfully advance humanity. The challenge is to create a processor architecture and programming model that will allow data parallel applications to achieve high throughput, near linear scalability, and ease of programming.

Learning from nearly five years of R&D and partnerships with researchers around the world, we've created a new GPU architecture that can make massive computational resources, which are otherwise used for graphics processing, available to high-performance computing applications. We call this new GPU "Tesla." With up to a several hundred times increase in performance, Tesla can enable scientists and engineers to utilize simulation, rather than cumbersome experimental techniques, to dramatically improve their productivity. Scientists working with Tesla have already reduced simulation times from weeks to hours. I believe we've created the modern high-performance computer architecture – a CPU, a GPU for data-parallel processing, and a GPU for graphics processing. We are entering the era of the personal supercomputing, where complex simulations will be performed and the results visualized in real time.

You often refer to NVIDIA's unique culture. How is NVIDIA's culture tied to the success of the company?

JH: One of our greatest assets is our culture of innovation and execution. NVIDIA drives one of the most relentless paces in the technology industry, so execution is key to our strategy and leadership. At the same time we are highly regarded for our imagination and spirit of innovation. At the core of our culture is a capacity to imagine, to take risks, and to have the willingness to make mistakes. After all, not all ideas are good ideas. Many ideas lead to dead ends, some become embarrassing mistakes, and many cost time and money. But that's the innovation process. We know that trial and error is an integral part of discovery. I really love this about our company: we're willing to take risks and put ourselves on the line, to try out an idea and see if it works. Tesla is a great example of this spirit of innovation and the imagination to create a new market. At NVIDIA, we don't build things that our competitors are building; we build things we believe the world will need. We build things that will change the world. I absolutely believe that this passion for innovation will drive our long-term success and our relevance to the world.

TRANSFORMATIONS // NVIDIA CORPORATION // WWW.NVIDIA.COM

NVIDIA Tesla

A New Architecture for the High-Performance Computing Market

High-performance computing (HPC) in fields like the geosciences, molecular biology, and medical diagnostics enables discoveries that transform millions of lives every day. However, universities, research institutions, and companies in these and other fields face a daunting challenge: as the complexity of their simulation models grows exponentially, so does their need for vast computational resources.

NVIDIA took a giant step in meeting this challenge with the announcement this quarter of a new class of processors based on a revolutionary new GPU. Under the Tesla brand, NVIDIA now offers a family of GPU Computing products that will place the power previously available only from supercomputers into the hands of every scientist and engineer. Today's workstations will be transformed into "personal supercomputers."

Applications that will benefit from the new Tesla solutions include geophysical analysis (with application speedups up to 100X), mobile phone antenna simulations (speedups up to 45X), molecular dynamics (speedups up to 240X), neuron simulation (speedups up to 100X), MRI processing (speedups up to 415X), and atmospheric cloud simulations (speedups up to 50X).

For example, Houston-based Headwave Solutions Inc. specializes in geophysical data analysis in the area of oil and gas exploration, which is becoming increasingly more difficult. Large reservoirs are now found at greater depths and in sediments that are much harder to analyze, such as the recent Jack Field discovery in the Gulf of Mexico, where oil was found at more than 20,000 feet under the sea floor. To discover and interpret these reservoirs, it is necessary to acquire and process huge amounts of seismic data. Due to the complexity of the sediment layers, better resolution is needed in the images, which means acquiring even more data.

Headwave, utilizing NVIDIA GPU Computing technology, allows geophysicists to apply advanced filters to their data and instantly see results even on multi-terabyte datasets. In addition, geophysicists can analyze the original acquired seismic ("pre-stack") data in multiple dimensions as part of their daily workflow. The processing of these terabyte datasets traditionally required months of manual labor and compute time for number crunching. With NVIDIA, Headwave is able to increase compute rates by up to 100X, while reducing time spent in manual operations.

Working with terabyte data sets in real time would not have been possible without the recent advances in GPU Computing. Because many workstations used in the oil and gas industry are already outfitted with GPUs, much of the hardware infrastructure to exploit this technology is already in place. As a result, companies are poised to start taking advantage of this new technology immediately. As Steve Briggs, VP of Systems Integration at Headwave, stated: "NVIDIA Tesla is going to make discovery of huge oil reserves possible through faster and more accurate interpretation of geophysical data."

NVIDIA CUDA

Computing on NVIDIA Tesla is available to any software developer through NVIDIA CUDA, the world's only C-language development environment for the GPU. CUDA is a complete software development solution that includes a C-compiler for the GPU, debugger/profiler, dedicated driver, and standard libraries. CUDA simplifies parallel computing on the GPU by using the standard C language to create programs that process large quantities of data in parallel. Programs written with CUDA and run on Tesla are able to process thousands of threads simultaneously, providing high computational throughput to enable the GPU to quickly solve complex, computational problems. The NVIDIA CUDA development environment is currently supported on the Linux and Microsoft Windows XP operating systems.

The new NVIDIA Tesla family includes

NVIDIA Tesla GPU Computing Processor, a dedicated computing board that scales to multiple Tesla GPUs inside a single PC or workstation. The Tesla GPU features 128 parallel processors and delivers up to 518 gigaflops of parallel computation. The GPU Computing processor can be used in existing systems partnered with high-performance CPUs.

NVIDIA Tesla Deskside Supercomputer, a scalable computing system that includes two Tesla GPUs and attaches to a PC or workstation through an industry-standard PCI-Express connection. With multiple deskside systems, a standard PC or workstation is transformed into a personal supercomputer, delivering up to eight teraflops of compute power to the desktop.

NVIDIA Tesla GPU Computing Server, a 1U server housing up to eight Tesla GPUs, containing more than 1000 parallel processors that add teraflops of parallel processing to clusters. The Tesla GPU server is the first server system of its kind to bring GPU Computing to the datacenter.

For more NVIDIA Tesla product information, please visit www.nvidia.com/tesla

TRANSFORMATIONS // NVIDIA CORPORATION // WWW.NVIDIA.COM

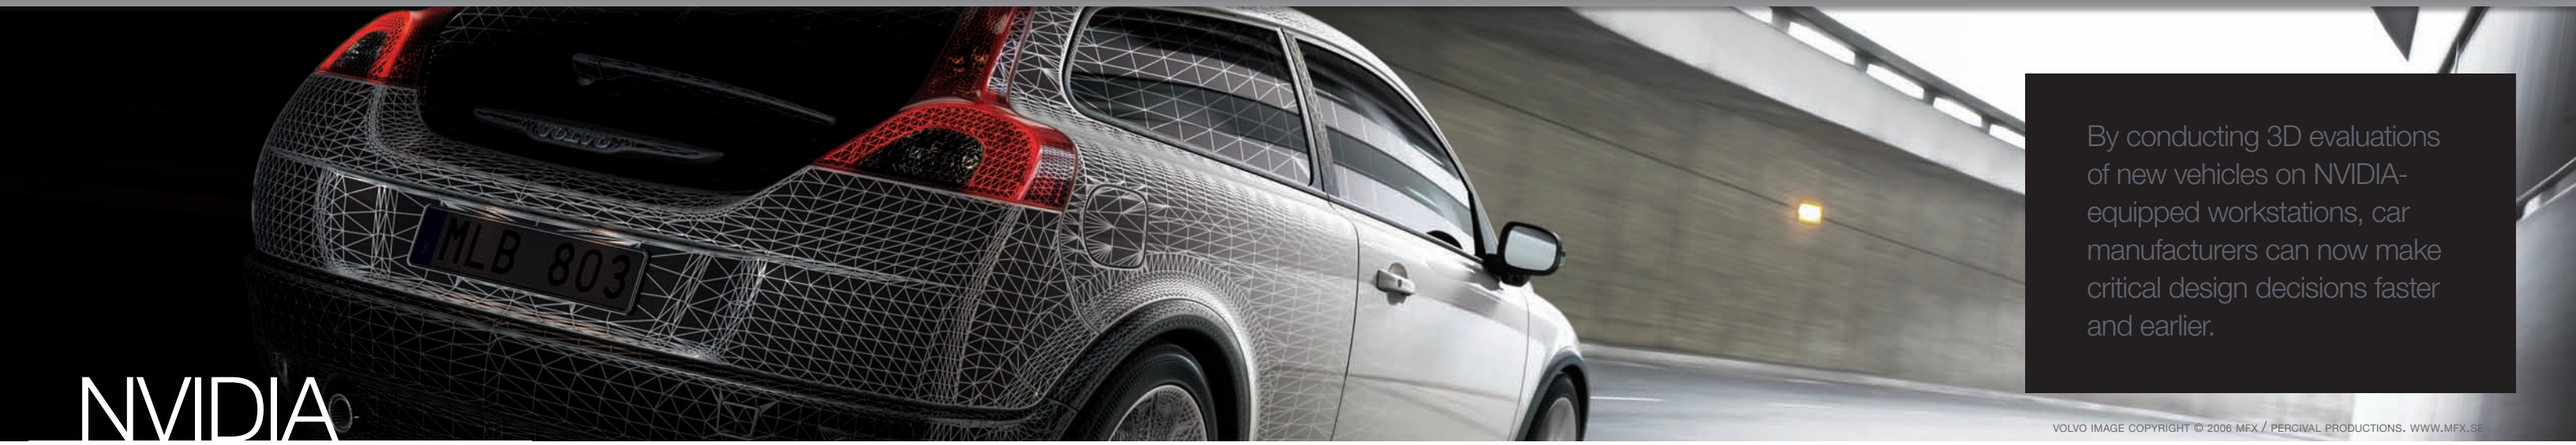

By conducting 3D evaluations of new vehicles on NVIDIA-equipped workstations, car manufacturers can now make critical design decisions faster and earlier.

NVIDIA Transforms the Automotive Industry

By conducting 3D evaluations of new vehicles on NVIDIA-equipped workstations, car manufacturers can now make critical design decisions faster and earlier.

The auto industry has long used Computer-Aided Design (CAD) tools running on NVIDIA GPUs as an integral part of the design process. When the first commercial CAD packages came to market, automotive companies quickly moved away from drawing boards and clay models into the world of 3D computer graphics, saving time, energy, and money. Now with new capabilities in visualization, simulation, and photorealistic rendering, NVIDIA GPUs are again transforming the auto engineering and marketing pipeline, aiding in everything from design visualization and crash simulations to production of 3D images for commercials and marketing materials—reducing costs, and speeding time to market for new models. Currently, it takes auto manufacturers about 48 months to design and build a new car—from the initial sketch to the first car rolling off the production line. NVIDIA technologies are reducing this time to market by tackling time-consuming tasks at all stages of the vehicle development process.

Visualization

One stage to which NVIDIA adds value is the early visualization process. By conducting 3D evaluations of new vehicles on NVIDIA-equipped workstations, car manufacturers can now make critical design decisions faster and earlier. With GPUs, life-sized car models can be displayed on a 3D wall, where designers and executives are able to immerse themselves in the model of the car and even get “inside” it. Designers are able to experiment with elements such as color, lighting, body style, and paint reflection on the fly. An interactive 3D design model can save many hours of work.

Simulation

Another important step in the creation of a new automobile is testing how all the materials that make up the car would perform in a crash. In the past, crash simulations would be conducted with a physical mockup of a car, which took weeks and cost several hundreds of thousands of dollars. Now, with NVIDIA-enabled workstations, crashes can be simulated in hours and the results are immediately available. With Mechanical Computer-Aided Engineering (MCAE), the properties of prospective materials can be modeled and tested in a variety of different crash simulations quickly, accurately, and less expensively, helping to create the safest car for consumers.

Data Processing

In the area of data processing, car manufacturers can benefit from GPU Computing. GPU Computing makes the task of computing large data sets faster and easier, bringing control back into the hands of the engineers. No longer do engineers have to send out data to be computed on a mainframe computer by another company. Now, when engineers need to process large data sets, they have the power at their fingertips within the timeframe that works best for them and the company.

Marketing and Advertising

Finally, when the car is almost ready to hit the market, the 3D CAD model is used for photorealistic rendering for television advertising and other marketing materials. In the past, if a car company planned to shoot video footage of a car racing through the French Alps on a winding mountain road, they would have to place the actual car and a large crew of producers and photographers at the physical location, a process that often took months and could cost many thousands of dollars. Now with the photorealistic rendering capabilities of NVIDIA GPUs, the 3D model can be inserted into existing imagery of the French Alps. The cost is much lower, but the impact of the advertising is just as compelling. The time it takes to produce the ad is reduced to about a month and the ad production doesn’t depend on the availability of physical models of the cars.

From design to production to marketing, the NVIDIA GPU has found its place in the automotive industry by adding value to real-time, specific problems in innovative ways.

RIGHT IMAGE COURTESY EDENFX

TRANSFORMATIONS // NVIDIA CORPORATION // WWW.NVIDIA.COM

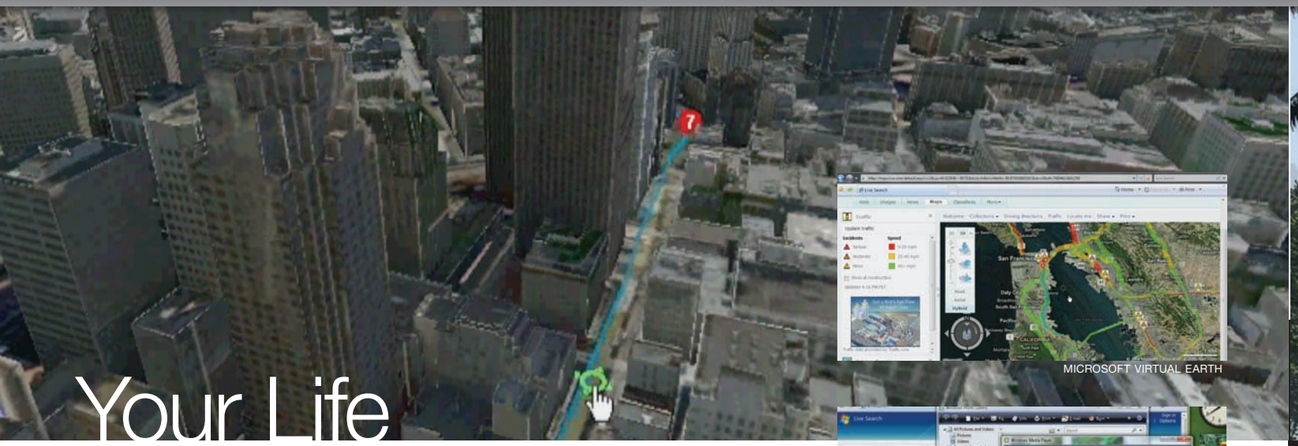

Your Life

Enhanced with 3D Graphics

3D imagery permeates our lives, whether it's for work or entertainment. These images awaken our senses, excite our emotions, and help us make decisions based on increasingly visual information.

The list of popular 3D-based applications that take advantage of an NVIDIA GPU to improve performance is rapidly growing. For example, online map applications such as Google Earth and Microsoft Virtual Earth enable you to view street images and buildings in 3D, so it's easier to visualize the surrounding environment. Google Picasa offers 3D effects when you scroll through your photos. And with Zinio, you can open an online magazine and actually "turn" the pages as you would with a physical document. Other applications and devices that benefit from the GPU include Adobe Acrobat and Acrobat Reader 8, Apple iTunes, and the Blu-ray and HD DVD players. Additionally, the new Microsoft Windows Vista operating system requires a GPU for functionality related to the user interface. While Vista has many important security and file management features, what grabs your attention is the compelling Aero 3D look and feel.

3D is changing the way we use our computers and consumer electronics devices, making work more efficient and play more enjoyable. As a result, NVIDIA GPUs are becoming more and more indispensable to the everyday computer user.

MICROSOFT VIRTUAL EARTH

MICROSOFT VISTA

ADOBE ACROBAT

GOOGLE PICASA

NVIDIA

Leadership in DirectX 10 Games

Microsoft's release of DirectX 10, which enables a completely new 3D rendering paradigm, represents the most significant step forward in 3D graphics application programming interfaces (APIs) since the invention of programmable shaders.

Games utilizing DirectX 10 deliver a stunning new level of graphical realism through vivid environments, detailed facial animation, and special particle and lighting effects. The NVIDIA GeForce 8 Series architecture was designed specifically for DirectX 10 and represents more than three years of intensive R&D and intimate collaboration with Microsoft. The GeForce 8800 GTX, which launched in November 2006, was the world's first DirectX 10–compliant, volume-shipping GPU. Developers were able to create and test their DirectX 10 content on the GeForce 8800 GTX and these cutting-edge games are hitting the market now. Since the launch of the GeForce 8 Series in November, NVIDIA has shipped millions of DirectX 10 parts, including a line up of GPUs in the sub-$100 range.

Here is a list of DirectX 10 games and expected availability dates:

Game	Available
Company of Heroes	Now available
Lost Planet	Now available
BioShock	August 21st
Age of Conan	Late Summer
Hellgate London	Q3 2007
Flight Sim X	Fall 2007
Company of Heroes: Opposing Fronts	Q4 2007
Crysis	Late 2007
World in Conflict	Holiday 2007
PT Boats	TBD

TOP CRYSIS, IMAGE COURTESY OF CRYTEK
RIGHT WORLD IN CONFLICT, IMAGE COURTESY OF MASSIVE ENTERTAINMENT

Improving Education
in South Africa

"Educators worldwide are looking for a cost-effective and sustainable way to bring computers into the classroom; Userful and NVIDIA make this a reality."

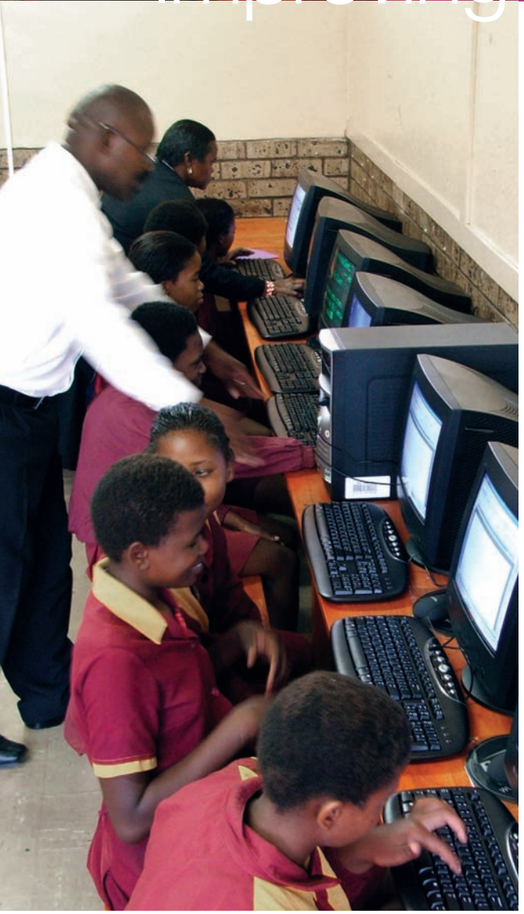

Providing computers and Internet access to South Africa's 28,000 schools and 12 million students is a challenge. In addition to the familiar issue of limited educational budgets, schools in South Africa have additional infrastructure challenges, including a lack of skilled technical personnel, the need for increased protection from theft, and classroom electrical upgrades to support networks of computers.

With hardware from NVIDIA and Desktop Multiplier® software from Userful, a software firm headquartered in Calgary, the government of South Africa has found a viable classroom computing solution. Through use of NVIDIA graphics cards and Userful software, schools in South Africa are providing up to seven students with a full desktop experience (monitor, mouse, and keyboard) for each computer. Each computer includes four NVIDIA graphics cards which, along with Userful's Desktop Multiplier software, allows multiple students to simultaneously use the same computer. The schools have fewer systems to buy and maintain, reduced power requirements, and the systems can be securely stored in a server room.

So far the Userful platform has been deployed to 105 schools in South Africa's North West Province, allowing up to 2,205 students to have computers and Internet access using only 315 systems. With this new access to technology, students will learn more and expand their technical skills, narrowing the "digital divide" and helping diversify the economy of South Africa.

With NVIDIA GPUs, Userful is able to provide high-powered, high-performance, high-speed computing and Internet connectivity—all the functionality of a high-performance computer—at an affordable price.

IMAGES COURTESY USERFUL

TIMOTHY GRIFFIN,
PRESIDENT OF USERFUL

"NVIDIA's leadership in price, performance, and worldwide availability has made them an ideal hardware platform for our software. NVIDIA technology allows Userful to offer our high-performance approach to affordable desktops to system builders on a global basis. Educators worldwide are looking for a cost-effective and sustainable way to bring computers into the classroom; Userful and NVIDIA make this a reality."

For more information, visit www.userful.com

TRANSFORMATIONS // NVIDIA CORPORATION // WWW.NVIDIA.COM

Transformations is a quarterly publication
of NVIDIA's investor relations and communications group.

Please send feedback to nvtransformations@nvidia.com.

NVIDIA Corporation | 2701 San Tomas Expressway | Santa Clara, CA 95050 | **www.nvidia.com**

NVIDIA