

# Forbes

January 7, 2008

## NVIDIA COMPANY OF THE YEAR



**Jen-Hsun Huang**  
NVIDIA'S CHIEF EXECUTIVE  
INSIDE THE CRYSIS GAME

To stay ahead in a merciless business, computer graphics champ Nvidia is setting its own pace | By Brian Caulfield

SHOOT TO **KILL**

# NVIDIA COMPANY OF THE YEAR

**They accost him at airports.** They pump him for technical information at conferences. Even his son is on his case. When you sell a product that tens of millions of videogame fanatics rely on, there is no escape. Nvidia Chief Executive Jen-Hsun Huang came home one evening in 2003 to find the latest copy of *Maximum PC* magazine waiting for him. Inside was a harsh review of Nvidia's new graphics card, the brawny set of microchips that makes the aliens and bullets come to three-dimensional life on a personal computer.

"Dad," Huang's then 13-year-old son Spencer, an avid gamer, wrote on a sticky note stuck to the page, "I think you need to kick it up a notch."

As if he needed the extra encouragement. In the past four years Huang has taken Nvidia from a trailing position to dominance in the breakneck, highly lucrative graphics-chip business. Intel gets all the credit for being the brains inside most computers, but without Nvidia's graphics cards there would be no high-definition Web video, no visual kick from battling aliens in the lush digital forests of the *Crisis* game.

The two-dimensional world of the Web is rapidly going 3-D with immersive applications such as Google Earth and World of Warcraft. That \$10 graphics processor used in cheap PCs won't cut it anymore, unless all you're doing is looking at low-quality YouTube videos. This is the age of laptops and desktops that can deftly handle hi-def video playback and video editing, with help from graphics cards that start at \$200 and go up from there. This is where Nvidia rules.

"What if Google Earth had a 3-D representation of every single building, every single square meter on Earth, and you could access it in the blink of an eye?" Huang says. "Take the Forbidden City in China. I know the model exists. We've seen the models. It is photorealistic. It is cinematic. It is breathtaking. All of those experiences are possible already."

As bandwidth becomes cheap and ubiquitous, people will

want 3-D power as much online as they do now in desktop games, and the demand for graphics processors will grow. "The more content there is, the more visual interest there can be, the more processing horsepower people need," says Huang, 44.

Nvidia's graphics cards are now the most expensive part of high-performance PCs. Gaming fanatics will frequently supercharge a \$1,200 desktop from the likes of Alienware, a Dell subsidiary, with another \$1,700 worth of Nvidia graphics chips. The Santa Clara, Calif. firm has 62% of the market for desktop PC graphics cards, up from 57% a year ago, according to Mercury Research. Its archrival ATI stumbled badly last year in the two-company horse race for technical superiority and was bought by chipmaker AMD for \$5.4 billion. In a PC industry racked by deflation, Nvidia has managed to increase its gross margin from 29% in 2004 to a current 46%. In the fiscal year ending in a few weeks Nvidia will gross \$4 billion, up 33%, and net \$900 million, up 50%. Since Huang took the company public in 1999, Nvidia's shares have risen 21-fold, edging out even the mighty Apple over the same time period. Such accomplishments, over the last 12 months and the past five years, earn Nvidia the title of FORBES' Company of the Year.

**Huang is now turning his focus** beyond gamers to a host of new customers that will need number-crunching power: oil companies doing deep-sea seismic analysis, Wall Street banks modeling portfolio risk and biologists visualizing molecular structures to find drug target sites. Pricing 150,000 equity options in a second or assessing how much crude oil sits in a pocket 6 miles below the earth's surface are massively parallel mathematical computing problems for which graphics processors are peculiarly equipped.

A typical Intel chip such as the Core Duo is called a central processing unit, or CPU. Like a sprinter (or, more appropriately, a pair of them), this CPU is both fast and extremely versatile, but it can do only so much computing in the space of a microsecond. It is ideal for tackling tasks that require lots of interactivity, for example, calling up data from a hard drive in response to a command from a user. To speed those tasks along, CPUs are built with onboard memory to store the next string of commands to be carried out.

By contrast an Nvidia graphics processing unit, or GPU, res-

## NEVER ENOUGH

Gamers are always hungry for new eye candy, and Nvidia keeps them satiated with a faster graphics card every six months or so. Its latest chips can render hundreds of millions of polygons per second, up from 1 million per second ten years ago.

**1995**  
**1 MILLION TRANSISTORS**  
Nvidia NV1  
Source: Nvidia.

**1999**  
**22 MILLION TRANSISTORS**  
Nvidia GeForce 256


BLOOMBERG NEWS (3)

**1993** Nvidia founded on Feb. 17, Chief Executive Jen-Hsun Huang's 30th birthday. Id Software releases *Doom*, which pits a futuristic space marine against a hellish array of creatures.

**1995** Computer-generated image from the game *Virtua Fighter 1* by Sega. Created with Nvidia's NV1 GPU, which had 1 million transistors.

**1999** Nvidia goes public on the Nasdaq stock exchange.

**2000** Acquires the assets of 3dfx Interactive, once a key rival.



ders graphics by simultaneously hashing out the same problem for multiple pixels: textures, lighting and smooth rendering of shapes flying across the screen. That makes a GPU more like a brigade marching in serried ranks. A GPU is less versatile but for repetitive tasks much more powerful than a CPU. The graphics chip has less need to store upcoming operations in memory, so it can cram in more processing cores than a CPU can. Intel has up to 4 processing cores on a single chip; Nvidia's new GeForce 8800 GTS has 128.

Parallel computing has been the holy grail of supercomputer design almost since the dawn of the silicon age.

Parallelism—simultaneous operations on sets of numbers—is suited to many engineering tasks, such as bomb design, weather forecasting and image processing. Seymour Cray, the supercomputer genius of the 1970s, had a form of parallelism with his vector-processing computers. The same kind of approach lay behind a generation of “array processors” used in computerized X-ray machines and seismic analyzers. But parallel processing proved to be much easier to describe than to execute on a commercial scale. The technology humbled some once promising outfits, like Danny Hillis’ Thinking Machines and Larry Ellison’s NCube.

How curious that parallelism is finally proving itself in, of all things, the manufacture of toys. The Nvidia GPUs are in, for example, the original Xbox and the newer Sony PlayStation 3.

**2006**  
681 MILLION TRANSISTORS  
Nvidia GeForce 8

**2007**  
754 MILLION TRANSISTORS  
Nvidia GeForce 8800 GTS

**2005**  
302 MILLION TRANSISTORS  
Nvidia GeForce 7

What's next? Using Nvidia to tackle classic engineering challenges of the sort that once preoccupied Cray.

This year Nvidia delivered its Tesla line of muscle-bound computing machines built around graphics processors. Now the Nvidia chips teenagers use to blast each other in Unreal Tournament are basically the same ones used by scientists at Evolved Machines in Palo Alto, Calif. who are simulating neurons used by the brain to see and smell. “You can see essentially a very high performance cluster the size of a desktop, at the cost of a desktop,” says Nvidia board member Steven Chu, a Nobel Prize-winning physicist and director of the Lawrence Berkeley National Laboratory.


Mighty Intel has noticed. In 2008 it plans to demonstrate a multicore chip named Larrabee that may steal some of the market for high-end graphics chips. Intel is also a threat at the low end, where it competes with Nvidia and AMD to sell the cheaper integrated graphics chipsets used in budget desktop computers and many notebooks. AMD, having digested its acquisition of ATI, is coming after Nvidia with a pair of powerful new graphics cards starting at \$220. “Their relative position versus Nvidia has improved compared to six months ago,” says Dean McCarron, president of chip tracker Mercury Research. AMD is also building ATI's graphics capabilities into a new CPU that may see the light of day in the second half of 2009.

Nvidia can never get complacent. Its customers are a fickle bunch, always ready to tear up their PCs to cram in the next hot graphics card from ATI. In 2007 Hewlett-Packard began selling a PC called the Blackbird 002 that allows gamers to use either a pair of Nvidia graphics cards or a pair from ATI,

**2002**  
63 MILLION TRANSISTORS  
Nvidia GeForce4

**2003**  
130 MILLION TRANSISTORS  
Nvidia GeForce FX

**2004**  
222 MILLION TRANSISTORS  
Nvidia GeForce 6



**2001** Microsoft releases the Xbox gaming console in November, along with first-person shoot-'em-up ▲Halo, showcasing Nvidia's graphics prowess.

**2002** ATI releases its Radeon 9700 graphics card in August to raves. Nvidia unveils its GeForce FX card for 2003, which gets panned by tech critics.

**2003** ATI wins the contract to design graphics processors for the ▲XBox 360.

**2004** Nvidia releases the GeForce 6 series, grabs strong marks from reviewers.

**2005** Sony unveils the ▼PlayStation 3, which uses Nvidia's graphics technology.



**2006** AMD agrees to acquire ATI for \$5.4 billion in cash and stock.

**2007** Nvidia releases its Tesla high-performance computing products. Actor ▲Doug Jones, rendered with an Nvidia GeForce 8800 Ultra card. Still from PS3's Rock Band▼.





and change between them as easily as changing socks.

Maybe that's why Huang is again stepping up the pace. A year ago Huang tore down his cubicle and replaced it with a conference table so his colleagues won't hesitate to plop themselves down and let him know what's going on.

The company's next generation of chips will be launched in the spring. In Nvidia's labs engineers hover around a cluster of 16 six-foot-tall computers that mimic the behavior of a future Nvidia product hooked up to an ordinary PC. Nvidia's engineers will be able to play games, watch videos and run ordinary applications on these hulking machines to see exactly how their next product will perform. Brian Kelleher, a senior vice president of engineering at Nvidia, pats another of the boxes with pride: "This box here will turn into a slice of silicon one-quarter of an inch on a side. Only it will run 1,000 times faster."

**Huang's background prepared** him well for the tumult and uncertainty of the chip industry. He was born in Taiwan, and when he was 9 his parents sent him and his older brother to live in Tacoma, Wash. with an aunt and uncle he'd never met. They in turn shipped the boys to what they thought was a prep school in rural Kentucky but turned out to be a Baptist reform school. Huang thrived by helping his classmates with their homework and learning to stand his ground. He bonded with his 17-year-old roommate, who was recovering from a fight that had left him with seven stab wounds. "He was my best friend," says Huang in his American-accented English. That mix of toughness and thoughtfulness has marked him ever since.

The Stanford-trained engineer cofounded Nvidia just as he turned 30 in 1993, after stints at AMD and LSI Logic. Huang and cofounders Christopher Malachowsky (currently Nvidia's vice

**The Emulator  
NVIDIA'S CHIP-TESTING  
MACHINE IN ITS  
SANTA CLARA LABS**

president of information technology) and Curtis Priem (who retired in 2003) looked at the cartoonish PC games of the time, such as Castle Wolfenstein 3D, and saw their future in improving the visual experience. The company capitalized on the "fabless" trend that allowed a chip firm to concentrate on design and avoid billions of dollars of capital spending by outsourcing production to Asian factories. Nvidia is a big customer at wafer fabs run by Taiwan Semiconductor and United Microelectronics.

Nvidia stumbled in the first few years. Its first product tried to do too much and was based on Nvidia's own quirky standard. After a round of layoffs in late 1995 Huang gathered the company's 35 remaining employees in a room and leveled with them: The company would have to find a way to catch up, fast. He gave up Nvidia's software to focus on the DirectX graphics standard that Microsoft was building into millions of PCs.

Huang spent \$1 million, a third of the company's cash, on a technology known as emulation, which allows engineers to play with virtual copies of their graphics chips before they put them into silicon. That allowed Nvidia to speed a new graphics chip to market every six to nine months, a pace the company has sustained ever since. Says Richard Bergman, senior vice president with the graphics unit at competitor AMD: "They raise the bar every time."

Intel is going to make its own graphics product out of its general-purpose processors, and AMD will seek to fold sophisticated graphics capabilities into its CPUs, stealing business that would otherwise have gone Huang's way. The problem is, by doing twice as much, Huang says, the bigger firms won't be able to offer the flashiest graphics capabilities quite as often. And Huang knows that if you can deliver something better, even if for just six months, you'll satisfy that fickle customer. **F**

This 6 month E-Print will expire June 3, 2008.

Posting by Permission of **Forbes** Magazine - January 7, 2008 issue © 2008 **Forbes** LLC.

For information on ordering **Forbes** reprints, please call (212) 620-2399. To subscribe to **Forbes** magazine, please call (800) 888-9896.