Reed Electronics Group

MICROPROCESSOR REPORT

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

THE EDITORIAL VIEW

Augmented Reality—and Larrabee

By Tom R. Halfhill {12/28/09-02}

Virtual reality is so...1990s. Sure, artificial environments are beguiling, whether they are created for videogames (*World of Warcraft*), virtual worlds (*Second Life*), Hollywood blockbusters (*Avatar*), or professional training (flight simulators). But now, virtual reality is looking like a stepping stone toward a grander concept: augmented reality.

"Augmented reality" is a term reportedly coined at Boeing in 1990 by Thomas Caudell, now a professor at the University of Arizona. Augmented reality combines some features of virtual reality with actual reality. It can overlay a live view of the real world with computer-generated graphics or textual information, building an enhanced version of reality that's easier to interpret or navigate. Sometimes, augmented reality fabricates astonishing illusions that are entertaining as well as informative. Eventually, actual reality may come to seem drab, confusing, even dangerous.

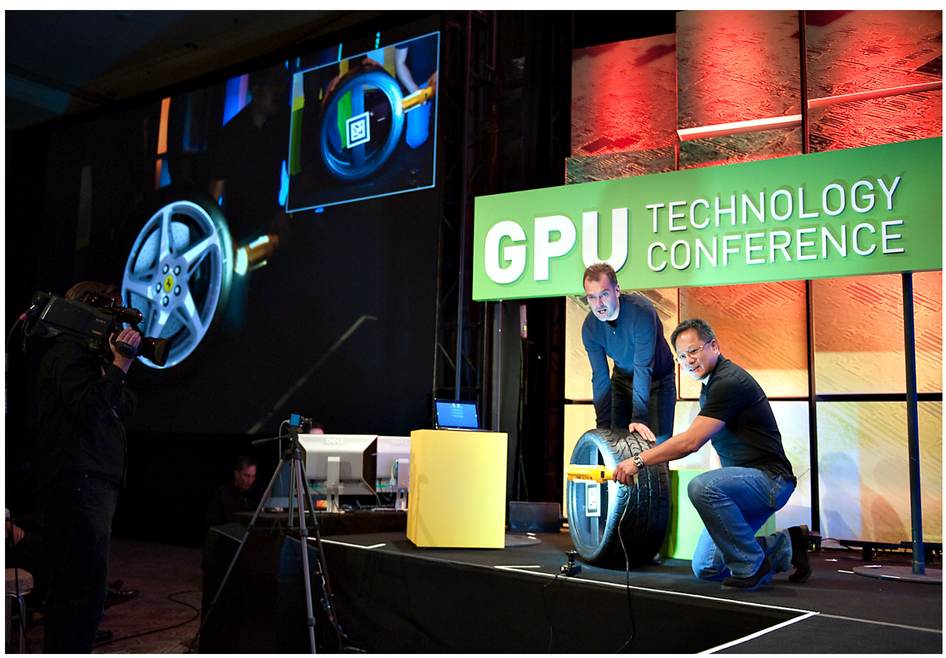

Ludwig Fuchs of RTT AG (left) and Nvidia CEO Jen-Hsun Huang demonstrate augmented reality at the GPU Technology Conference. A live video image of the tire—with a barcode mounted in the center—is visible in a window at the top of the large video screen at left. Further left is an augmented-reality version of the same image with a Ferrari wheel rendered into the tire. The rendered image responds in real time to changes in lighting and tire orientation. Nvidia photo by Olivier Giroux

Is It Wheel or Fake?

I've seen some mighty impressive product demos in my time. Highlights include the first Macintosh, Andy Warhol creating impromptu artwork on a Commodore Amiga, the first CD-ROM drives, the Mosaic web browser, and cross-platform programming in Java. But an augmented reality demo at Nvidia's GPU Technology Conference in September was the most eye-popping thing I've seen in years.

The demo was prepared by Real Time Technology (RTT) AG, which is designing point-of-sale software for automaker Ferrari. RTT's demo began with an excellent but conventional example of virtual reality. Prospective car buyers can view 3D rendered images of Ferrari's sports cars. With a few mouse clicks, shoppers can customize a car with any available colors, accessories, and options. RTT's software displays a fully navigable view of the car's interior and exterior. Nvidia's parallel-processing GPUs provide the horsepower for this real-time rendering, producing images almost indistinguishable from photographs.

It's all good. However, the second part of RTT's demo had the standing-room-only crowd at the conference gasping in disbelief. RTT cofounder Ludwig Fuchs rolled an automobile tire on stage. The tire wasn't mounted on a wheel. Instead, there was only a white placard with a large barcode fixed in the center. Offstage, a cameraman pointed a computer-tethered video camera at the tire. The live image appeared on large video screens flanking the stage.

But, whoa! On the screens, the tire appeared to be mounted on a fancy Ferrari wheel that wasn't really there. Nvidia CEO Jen-Hsun Huang aimed a handheld shop light at the tire while tilting it from side to side. On the video screens, the shiny metal "wheel" moved in perfect synchronization with the tire and reflected the light perfectly. RTT's software and Nvidia's GPUs rendered the ray-traced image in real time. The stunned audience looked back and forth from the reality on stage to the augmented reality on screen. The wheel wasn't real, yet there it was.

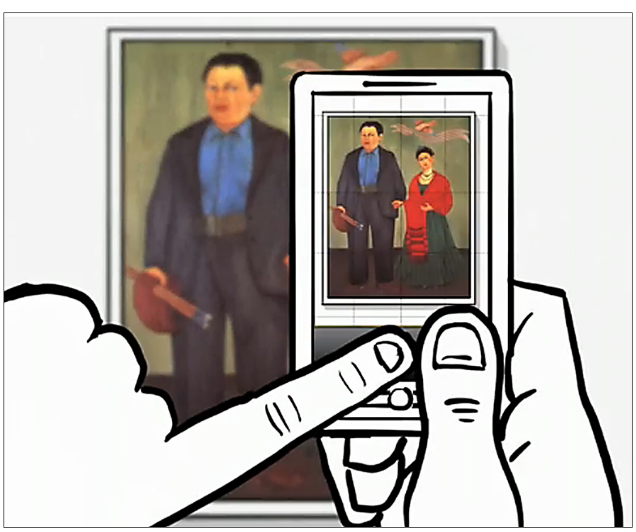

Google Goggles can identify books and paintings photographed with an Android smartphone camera.
(Composite image by Google)

With a few mouse clicks, RTT's software could render any of Ferrari's optional wheels into the tire, instantly. The reflections, lighting effects, and specular highlights looked absolutely authentic. It was an enhanced version of reality—augmented reality.

Google Goggles for Android

The RTT demo showed what's possible when an element of virtual reality—the computer-rendered wheel—is combined with a video image of actual reality. The result was a composite image that simulated a view of the real world so accurately that the augmentation was undetectable. If we hadn't seen the unmounted tire on stage, we wouldn't have known that the wheel was a rendered image. (The barcode placard helped the computer draw the graphics in the proper orientation.)

Another application of augmented reality is to add information to a view of the real world. This type of augmented reality doesn't attempt to fool the viewer into seeing something that isn't there. Instead, it's like a photograph with superimposed captions, except it's created in real time on a live video image of a real scene.

On December 7, Google entered this field by introducing Google Goggles. It's an app for the Android mobile operating system, which runs on smartphones like Motorola's Verizon Droid. Google Goggles combines search-engine information with live video or still photos captured with a cellphone camera.

The still-photo feature is impressive, but it's more like a visual search engine than augmented reality. In one example, a Google product manager photographs a painting with an Android cellphone and initiates a search. The Google search engine identifies the painting and the artist and finds websites about them. This program can also identify products from closeup pictures of their universal product barcodes—a cute trick that some existing smartphone apps can do.

More impressive is the real-time combination of search results with live video images from the cellphone camera. In this example, a Google engineer points the Android phone at a building. Using Global Positioning System (GPS) data and a built-in electronic compass, the phone knows where it's pointing and finds information about the building. Google Goggles identifies the building as a restaurant and displays a link to the business's website. (To see a YouTube demo, visit *www.youtube.com/watch?v=Hhgfz0zPmH4*.)

All this happens instantly. Standing on a sidewalk, you can view a live video image of a street scene on the phone's LCD. As you point the phone in different directions, Google Goggles superimposes text labels on the locations it recognizes. Some labels are hyperlinks to websites with information about those places. You're still looking at the real world, but reality is augmented with data that helps you interpret the scene and navigate it. (Just remember to look both ways with your real eyes before crossing the street!)

Applications for Augmented Reality

Augmented reality has obvious applications for tourism and for anyone trying to find their way to a destination in an unfamiliar part of town. It will be valuable for people with physical impairments. At museums and historical sites, it will be a visual supplement to the rented headphones that play prerecorded guided tours. The U.S. military is using similar software to identify locations of previous ambushes, bombings, and other incidents important to know for soldiers on patrol. Like the Internet, augmented reality is one of those technologies with endless possibilities.

Google is working to make Google Goggles a navigational aid for the natural world, too. Point your cellphone camera at a leaf to identify the plant...point it at an animal to identify the species. Last spring, I photographed a strange bird in southern Virginia that I couldn't identify. I tried looking it up on the Internet, but I couldn't find a website that allows a layman to search a zoological database by specifying visual features. Months passed before I identified the bird using an old-tech solution: I sent the photo to a former colleague who is an avid birder. (The bird was a yellow-crowned night heron.)

Of course, any new technology has downsides. Will people become so dependent on augmented reality that they will fear to venture outside without it? Today, many people won't drive anywhere without a cellphone, in case of a mechanical breakdown or other emergency. Some people have discarded their paper maps and can't find their way to an unfamiliar destination without a GPS unit that tells them exactly where to turn. Many folks can't imagine a vacation or even a weekend without e-mail and text messaging. But let's face it, humans have been depending on technology since the invention of fire.

One thing I worry about is the temptation to merge augmented reality with a heads-up display in automobiles. Yes, it would be useful to have real-time GPS tagging and other information projected in front of my windshield. However, all that data will compete with the road signage and my limited attention.

If this technology is carried too far, someday a heads-up display will tell us the name, age, sex, height, weight, occupation, and hobbies of the pedestrian we're about to run over.

Intel Stumbles With Larrabee

Much has been written about Intel's decision not to introduce Larrabee in 2010 as a discrete graphics processor to compete with GPUs from AMD and Nvidia. Intel will distribute small numbers of preproduction chips for use as development platforms but is postponing the introduction of discrete GPUs. The delay may also affect Intel's plans for future CPU/GPU integrated chips.

Microprocessor Report and In-Stat fielded several press inquiries after Intel's December 7 announcement. We had several reactions: the setback was good news for AMD and Nvidia, obviously; GPUs are complex microprocessors that are very difficult to design, even for the world's biggest semiconductor company; designing a competitive GPU on the general-purpose x86 architecture is even more difficult than designing a GPU with a special-purpose architecture; and the stumble won't stop Intel from pursuing its long-term strategy, because graphics processing and high-performance computing on GPUs are too important to ignore.

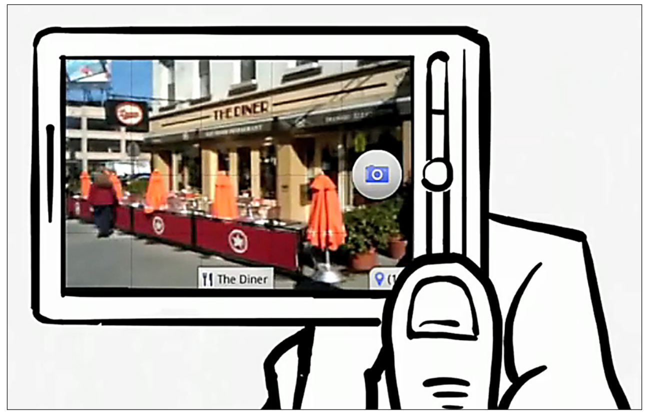

Google Goggles can identify places of interest in live video images captured with an Android smartphone.
(Composite image by Google)

Keep in mind that Larrabee isn't solely a GPU. It's also designed to be a manycore chip for general-purpose parallel processing. Both AMD and Nvidia are promoting their GPUs as alternatives to CPUs for high-performance computing. GPUs are enabling a new class of "desktop supercomputers" and are finding their way into real supercomputers as well. Nvidia has been particularly aggressive with this strategy, heavily promoting its CUDA platform and new Fermi "computational GPU" architecture. (See *MPR 10/5/09-01*, "Looking Beyond Graphics," and *MPR 1/28/08-01*, "Parallel Processing With CUDA.")

Intel remains committed to discrete graphics and high-performance computing. However, Intel's virtual admission that the early Larrabee silicon isn't competitive for discrete graphics certainly detracts from its FUD factor. Developers who were waiting for Larrabee to shake up the field will be more open to adopting the proven technology from AMD and Nvidia.

Feeling the Heat

More than a year ago, *MPR* published a detailed analysis of the Larrabee architecture. We wondered if an x86-based GPU could really beat the specialized architectures from AMD and Nvidia: "At times in this article, we have expressed doubts that Larrabee will exceed or even match the power/performance benchmarks of GPUs from ATI and Nvidia. Of course, this is pure speculation on our part, because Larrabee silicon isn't yet available for independent testing. One reason for our skepticism is that an x86-based GPU with minimal hardware acceleration seems poorly matched against highly specialized GPUs from vendors with more experience in the field."

We identified other potential obstacles for Larrabee: power consumption, if the x86 cores aren't as efficient as special-purpose graphics cores; manufacturing congestion, if Intel tries to compensate for a handicap by moving production to the latest fabrication processes on which Intel's CPUs depend; the difficulty of integrating a GPU with a CPU, for chips that will compete with AMD's future Fusion processors; the learning curve of developing software for a new GPU, even if it's based on the x86; and the challenge of optimizing the graphics driver software, which must translate DirectX or OpenGL calls into native code for Larrabee's unconventional graphics pipeline. (See *MPR 9/29/08-01*, "Intel's Larrabee Redefines GPUs.")

Which obstacle tripped up Larrabee? Our educated guess is that the early silicon isn't competitive with AMD and Nvidia in graphics performance, power consumption, and manufacturing cost. Fixing those fundamental hardware problems will likely require an extensive redesign. Although Intel still hopes to sell the first Larrabee implementation for high-performance computing and software development, those chips will probably be pricey, without the consumer graphics market to drive large volumes.

Despite Larrabee's setback, we'll be surprised if Intel veers from its strategy of using the x86 architecture for these processors. Even if that strategy isn't technically the best route, we suspect that Intel is too devoted to the x86 and too far down the road to change directions now. ◇

Tom R. Halfhill