

MIKE VILLAS'S WORLD



The augmented-reality wonderland of Pyramid Hill and Fairmont High School is taking shape today
 BY HARRY GOLDSTEIN

ABOUT THE BIG PICTURE, there isn't much doubt. Sensing, monitoring, networking, and computing technologies of incredible variety and profusion will converge over the next 10 to 20 years to give us—and those who would keep tabs on us—incredible powers of observation. But exactly how it will change our lives, we can only imagine.

Some fear a simple eruption of technology-based Orwellian repression. Others anticipate the emergence of a hyperengaging form of existence based on really cool toys that (caveat emptor) spy on us every now and then. We'll drift casually in and out of augmented reality and have dizzying access to an unceasing torrent of information.

In the cool toys category, some of the most compelling and detailed scenarios have come from Vernor Vinge, a science fiction author and former computer science professor at San Diego State University. In the preceding story, "Synthetic Serendipity," which Vinge adapted for *IEEE Spectrum* from his upcoming novel, *Rainbows End*, he introduces us to Mike Villas and


his friends. Through their eyes, we see how we might integrate the coming technologies into our lives.

Vinge's sensor planet of 2020 teems with billions of wireless ultra-wideband communications nodes connected to countless pinhead-size cameras, microphones, motion detectors, and biometric and other sensors to form a fine-grained mesh of networks that cover every square millimeter of the globe. Equipped with full-color, see-through displays that cover each pupil like a contact lens and clothing that senses muscle twitches, people will exploit an immensely sophisticated successor of today's Internet. They'll be able to immerse themselves in gripping gaming environments, silently communicate with friends just by tensing their muscles, and hunt down information about other people.

In the next 30 years, Vinge believes, we will reach a point where the combination of powerful processors, limitless data-storage capacity, ubiquitous sensor networks, and deeply embedded user interfaces will create a bond between human and machine "so intimate that users



ALL ILLUSTRATIONS: PETER BOELENBERG



their position in the network by communicating with other nodes located in a 10- to 20-meter radius in every direction.

In a network of thousands of nodes, each individual node talks only to another dozen or so nodes in its local cluster. To fix its position in space, an individual node measures both the time it takes to transmit a train of pulses to a neighboring node and how long it takes to receive an answering pulse from that node. Like old friends trading gossip, one node will tell another node about other neighbors it has contact with, so that every node in a cluster knows its position relative to all the others.

Mated to the story's game server in Vinge's Pyramid Hill Amusement Park, these localizers keep track of players to within a fraction of a millimeter, to overlay full-color graphics on top of players' real-world views.

But here at the turn of the century, we're still fumbling with the Global Positioning System and location accuracies measured in meters or, at best, centimeters. Nevertheless, an early approximation of Pyramid Hill exists on the campus of the University of South Australia in Adelaide. Researchers lugging backpacks stuffed with notebook computers, cables, batteries, and GPS receivers sport bulky electronic halos—head-mounted displays and head-tracking devices. As they stalk around campus, they wield haptic guns that vibrate when fired at monsters from a modified version of *Quake*, the popular desktop shooting-gallery game. Bruce Thomas, Wayne Piekarski, and their colleagues at the university's Wearable Computer Laboratory took *Quake*'s open source code and adapted the game to the campus environment. They kept the guns and monsters but removed the texturing for the ground, structures, and sky so that the real world shines through.

The game, which they call *ARQuake*, works pretty well when a player is more than 10 meters away from, say, an approaching monster, which appears to the gamers amid real objects. A combination of magnetometers, gyroscopes, accelerometers, and GPS readers track each player and match the position of physical objects like trees and buildings to the blank spaces in the computer game where the

graphical images of these objects have been erased.

But the closer the virtual character is in relation to the player, the more the GPS tracking error erodes the gaming experience and the virtual monsters do things they shouldn't, like pass through walls. When a player is within a meter or two of a three-story building, which blocks the GPS signal, the game can no longer track the player at all, and the virtual world stops moving with him or her.

"The trick is designing the games so they force the user to stay within the desirable operating area," says Piekarski. "So if the GPS doesn't work well next to the walls, encourage the user to stay out in the open by making all the interesting stuff happen there."

To get the kind of seamlessness that Vinge envisions will require an area bristling with localizers. It's also going to demand some cleverness in matching virtual props and characters to real ones—for example, a virtual dinosaur to the robotic mechanism built solely to be the framework for the graphics-created beast.

"The robot's going to have to be pretty much the same size and shape as the dinosaur—otherwise it's not going to bite players at the right time," says Blair McIntyre, an assistant professor at Georgia Tech's Augmented Environments Lab in Atlanta. "When you hide the physical world behind these virtual overlays but then expect people to be able to touch things in the physical world, that means that the virtual world must precisely correspond to the physical."

The stumbling block for immersive games so far is a vicious little problem called simulator sickness, according to Wright, cofounder of the computer game company Maxis Software Inc., in Walnut Creek, Calif., and creator of such blockbuster games as *SimCity* and *The Sims*. "It's hard to get a really accurate read on which way you're pointed and to have it update fast enough," says Wright. "With today's systems, you turn your head, and all of a sudden reality moves, but the virtual image lags behind." For most people, nausea ensues.

Overcoming the problem comes down to minimizing latency—the lag that occurs when radio signals travel among players and the computers that compute the frame-by-frame updates for each player. In "Synthetic Serendipity," localizers spread all over Pyramid Hill deliver this information at a rate fast enough to keep the experience smooth and immersive. But

may reasonably be considered super-humanly intelligent."

Vinge (pronounced VIN-jee) closely tracks emerging technologies by staying in touch with such influential computer scientists as Robert Fleming and Cherie Kushner and with research engineers like Georgia Institute of Technology's Thad Starner. *Spectrum* talked about Vinge's story with his friends, as well as with such techno-gurus as Jaron Lanier, a pioneer of virtual reality, and Will Wright, creator of the hit computer game *The Sims*. Though they squabble about how the technologies described in "Synthetic Serendipity" will come together in the end, they all agree that the result will make fact out of fiction.



"YEARS AGO, GAMES AND MOVIES

were for indoors.... Now they were on the outside. They were the world."

Vinge's brand of immersive reality in "Synthetic Serendipity" can be thought of as an electronically created, shared hallucination. It combines virtual reality (where your ears, eyes, and skin are fed computer-generated sounds, images, and sensations) and augmented reality (where computer-generated images are laid over your view of the physical terrain). In San Diego, circa 2020, you can see what you choose to see: suburban homes can become castles; your friends, velociraptors. The key challenge to mapping the virtual onto the real is access to streams of sufficiently specific and precise location information. In Vinge's conception, the task is performed by localizer nodes, which are transceivers that determine

what about players who are not located on the Hill, such as the mysterious stranger who projects Big Lizard onto that robot? For players in remote locations to share a virtual reality experience, the basic strategy is to pick a prime location for the computers that run the software that predicts, for each player, what frame needs to appear next in his or her display.

Jaron Lanier, the dreadlocked hipster credited with coining the term virtual reality, wants to embed VR displays into enclosed environments such as offices, classrooms, and laboratories so that people across the globe or across the street can share the same virtual environment—a VR teleconference of sorts. For the room-sized tele-immersive displays he and others have been developing, real-world latencies are in the 40- or 50-millisecond range, more than enough to turn the experience into a herky-jerky nightmare. To achieve the bone-crunching realism Vinge describes, engineers will have to get laten-

cies down to 10 ms or less, says Lanier, who is now a visiting scientist at Silicon Graphics Inc., in Mountain View, Calif.

When it comes to chipping away at latencies, different virtual interactions call for different strategies. For the “haptic carnage” Vinge writes about, a local cache of simulation data related to dinosaur teeth and jaws stored on Fred’s wearable could feed actuators in his shirt to create the sensation of being munched. But for real-time interactions between people on opposite sides of a continent, where there is an unpredictable two-way flow of data, engineers will have to put the computers midway between them in what Lanier calls “virtual-world prediction megacenters.”

In 2000, as chief scientist of the National Tele-immersion Initiative, a coalition of research universities studying advanced applications for the next generation of the Internet, Lanier helped demonstrate that placing computation in the middle of a network dramatically improved the quality of prototypical immersive applications. His team used the Pittsburgh Supercomputing Center to link researchers at the University of North Carolina, Chapel Hill, with staff from Advanced Network & Services in Armonk, N.Y. With virtual laser pointers, the participants moved computer-generated furniture around in a three-dimensional space projected in real time onto large screens at each venue. Though more like a low-fi holodeck from *Star Trek* than Pyramid Hill, the demo showed that we’re on our way to sharing virtual experiences in real spaces.



“THE TWINS LOOKED AT EACH OTHER. Mike could tell they were silent messaging.”

The idea of communicating without uttering a word or typing on a keypad seems like so much parapsychological mumbo jumbo. But there’s solid technology behind it. In Vinge’s story, Mike and his friends, the Radner twins, play games, surf the Web, and “silent message” each other by shrugging and twitching to control the electronics embedded in their clothing.

And such gesturing isn’t the only alternative to speaking or typing. Researchers today are already investigating several alternatives, including eye blinks and actions known as subvocal utterances.

Vinge’s twitching communication scheme seems the most fantastical, but it could capitalize on medical technology that has been around for years. In a typical setup

for monitoring muscle activity, surface-electromyography electrodes placed in contact with the skin detect the minute electrical signals associated with the contraction of a muscle. In the wearable application Vinge imagines, analog signals picked up by such electrodes would be amplified and then sent along conductive threads to chips that would digitize the signals and route them to a processor running gesture-recognition software. The program would first determine whether the series of muscle contractions is intentional and, if so, match them to a library of gesture cues and associated meanings. Finally, the textualized message would be wirelessly transmitted to the intended party.

“In theory, we could train ourselves to twitch different muscles so we could effectively type,” says Georgia Tech’s McIntyre. “All that you need to learn how to control these muscles is feedback. And if the input device reacted predictably and consistently, then we could learn to control things that were peripheral to our bodies just by having the computer monitor signals to and from the brain.”


For now, researchers are working on different silent messaging techniques. Elsewhere at Georgia Tech, Thad Starnier’s “blinkprint” technique measures eye blinks to identify blinkers and allow them a measure of rudimentary control.

But for many, using vocal cords is still the most natural way of communicating—even if you don’t make a sound. Researchers at NASA’s Ames Research Center, in Moffett Field, Calif., led by Chuck Jorgensen, recently proved that the tremors in nerves controlling the vocal cords can be detected when someone is speaking very quietly or even just reading silently. His team is experimenting with button-size surface-contact electrodes, placed beneath and on either side of the larynx, which detect the nerve signals that would otherwise become speech. The sensors relay those subvocal nerve signals to a digital signal processor and then to a software package trained to recognize certain signals as simple words, such as “stop” and “go” and the digits “0” through “9.” Researchers used the system for hands-free Web browsing.



“MIKE GAVE A SHRUG and a twitch just so. That was enough cue for his Epiphany wearable.”

In “Synthetic Serendipity,” clothing is the interface for everything from personal communications to gaming. Today, the



primitive forerunners of these garments are being tested by e-textile pioneers like Sundaresan Jayaraman [see “Ready to Ware,” *Spectrum*, October 2003]. His prototype machine-washable SmartShirt contains sensors, actuators, processors, and communications circuitry all embedded directly in the fabric and connected by a flexible data bus.

The SmartShirt controller, now the size of a pager and running on a 3-volt battery, processes the signals from sensors woven into the fabric to compute vital signs like heart rate. It wirelessly transmits the data to a display device, such as a wristwatch, a PDA, or a PC. By 2020 such a controller will be smaller than a dime and powered by microfuel cells or solar cells woven into the fabric, like those demonstrated recently by photovoltaics maker Konarka Technologies Inc., in Lowell, Mass.

The garments of Vinge’s imagination, which can faultlessly sense even subtle messages amid random twitches and shrugs, are years away. Nevertheless, the raw computational power and storage capacities in today’s experimental wearables are already impressive. As he strolls the campus of the College of Computing at Georgia Tech, where he is an assistant professor, Starner views e-mail documents and surfs the Web by peering through a head-up liquid-crystal display clipped onto his eyeglasses. He’s also “typing” single-handedly on a palm-size keyboard-and-mouse combo called a Twiddler2, from Handkey Corp. in Denver. Obviously, he’s never far from his computer, a shoulder-bag unit from Charmed Technology in Los Angeles. It’s based on a low-power Transmeta Crusoe processor, 256 megabytes of RAM, and an 80-gigabyte hard disk.

Starner predicts that within five years, the shoulder-bag unit will shrink to fit in his pocket and carry 1 terabyte of disk space. That’s more than enough to store a year’s worth of your e-mail, music, video, medical records, and every word you utter.

Tiny, light, and bright displays are clearly central to Vinge’s future world; few people would gladly wear the geeky clip-ons that Starner habitually sports. In Vinge’s story, the display of choice is a retinal-scanning system embedded in contact lenses. The basic technology is already here, albeit in much bulkier form, in displays from Microvision Inc., in Bothell, Wash. The company’s scanned-beam display uses extremely small semiconductor lasers to scan images directly onto the retina [see “In the Eye of the Beholder,” *Spectrum*, May].

The display allows augmented-reality software to superimpose graphs and text over your view of real objects—schematics of the underground utility infrastructure of Pyramid Hill, say, or navigation arrows that guide you to the right classroom. When green diode lasers become available to combine with the red and blue ones already here, full-color, 3-D displays will give gamers gut-wrenchingly vivid scenes that will make the images Starner now sees on his display look like grainy old snapshots.

The scanned-beam display, including the lasers and the 2.5-mm-diameter microelectromechanical scanner that paints the light onto the retina, needs to shrink to fit comfortably and unobtrusively on a contact lens. John R. Lewis, a research fellow at Microvision, insists he could build such a prototype today for US \$5 million to \$10 million.

Of course, we might skip the contact lens altogether and pump images directly into the brain’s visual cortex. Research scientist Lanier, who counts himself an expert in user interfaces, among other things, believes that prosthetic sensory implants are almost inevitable; it’s just a matter of how far into the nervous system the embedding takes place.

“The machinery of the retina and its connectivity to the visual cortex and the visual cortex’s integration with the brain are all so exquisitely good that probably the best engineering decision will be to implant a display inside the eye but on the outside of the retina,” Lanier says. “But engineering culture has to articulate a more spiritual and more beautiful vision for the future in order for any of these things to be accepted.”



“WITHOUT A COMPLETE localizer mesh, nodes could not know precisely where they and their neighbors were. High-rate laser comm could not be established....”

In Mike Villas’s world, every object is instrumented and networked: its description, current states, capabilities, and relationship to its context are known to anyone who gains access to the network.

Vinge bases his fictional sensor networks on the localizer nodes being developed by Robert Fleming and Cherie Kushner, cofounders of Aether Wire & Location Inc., in Nicasio, Calif. Over the last decade, under contracts from the U.S. Defense Advanced Research Projects Agency and other government sponsors,



the company has been making a device, currently the size of a pager, capable of providing location information accurate to within a centimeter.

As described earlier, localizer nodes are transceivers that determine their position in the network by communicating with other localizers placed in a 10- to 20-meter radius in every direction. By forming a mesh network, the devices talk to each other instead of a central base station, sending packets for one another along routes that are recalculated every few milliseconds to find the fastest path.

Unlike the sensor networks based on various flavors of IEEE 802.11 (Wi-Fi) and IEEE 802.15.4 (ZigBee) that are now becoming ubiquitous in developed countries, Aether Wire’s networks rely on low-frequency ultrawideband. While traditional narrowband radio communications transmit data by modulating sinusoidal waves and emitting a great deal of power in a narrow band of frequencies, ultrawideband transmitters blast low-power streams of pulses at a rate of 40 million to a billion times per second over a broad swath—at least 500 megahertz—of spectrum. Information is impressed onto the pulse train by varying the amplitude, spacing, polarity, or duration of the individual pulses in the train.

In Vinge’s world, localizers serve a variety of functions. They precisely steer lasers that send and receive data to and from people and devices at tremendously high rates over the Internet. Localizer nodes can also be integrated with just about any kind of sensor imaginable, including the hundreds of cameras embedded at Fairmont High School, which Mike and his friends tap into to peer over a classmate’s shoulder or to locate one another across campus. ■