AUTOMATIC NETWORKS

Devices that connect themselves could change networking

You've spent the last several weekends hunting for the perfect lamp to brighten up that shadowy corner of the living room. When you finally bring it home and plug it in, the network of motion sensors and light meters in the house immediately senses the torchière and turns it on—but only if it's dark and you're in the room. And if you decide to swap it with the lamp in the bedroom, no problem: the network figures that out as soon as you're done.

In theory, linking together sensors, appliances and other devices so they can communicate and work together could make life easier and more productive. The reality—at least for now—is that setting up such networks is expensive and far from easy, especially if they involve thousands or even millions of components.

Now networks of devices that organize themselves—connecting to one another wirelessly and automatically, without human intervention—are moving out of research labs and into the marketplace. In their first incarnation, they will connect large numbers of sensors in factories and industrial settings, but within a few years they will move into office buildings, homes, even farm fields. Companies like MIT Media Laboratory spinoff Ember, Motorola and San Diego-based Sensoria are moving to create and sell the wireless radios and microchips that will enable devices ranging from temperature sensors to sprinklers to be connected in self-organizing networks [*Technology Review* board member Robert Metcalfe is an Ember investor and board member. Ed.]. "It really is the only form of networking that can work for lots and lots of little objects," says MIT Media Lab researcher Michael Hawley. "The consequences of it really are going to be as magical as anything we've seen in technology."

In a self-organizing network, says Ember cofounder and chief technology officer Rob Poor, "You just plunk these nodes down, and they discover each other and figure out how to get data back to where you want to get it to." In other words, every element automatically recognizes every other element. Without any outside help, the devices must then determine how to get data where it needs to go.

One self-establishing network that is slowly making commercial headway is the Bluetooth wireless system. Originally designed by cell phone maker Ericsson to replace cables running between devices like computers and printers or cell phones and headsets, the system allows up to eight devices to connect to each other. When any device equipped with a low-cost Bluetooth radio comes within about 10 meters of any other Bluetooth-enabled product, the two automatically initiate a connection.

"Bluetooth is pretty well designed for what it was supposed to do, which was allow all the devices you carry around to talk to each other," says Ember cofounder and chief scientist Andy Wheeler. But the eight-device limit and the design of the Bluetooth network both hinder the system's usefulness for applications that require hundreds or thousands of devices spread over a large area.

Most wireless networks, including Bluetooth and the popular 802.11b "Wi-Fi" networks used to connect computers to the Internet, employ a spoke-and-hub organization in which one device acts as a central access point that all the other network members must communicate with directly. This turns out to be impractical in settings like factories, which are filled with machines, thick concrete walls and other radio-hostile things. That's where Ember and similar network designs come in. With Ember's networking protocols, the radio networks look like a mesh: every device recognizes other nearby devices immediately and can talk to all of its neighbors, passing data along. "Every

SOME COMPANIES IN SELF-ORGANIZING NETWORKS

COMPANY	APPLICATION
Sensoria (San Diego, CA)	Home and office building automation; automotive applications
Motorola (Phoenix, AZ)	Wire replacement for factories; home and office building automation
Palo Alto Research Center (Palo Alto, CA)	Military applications and homeland defense
MeshNetworks (Maitland, FL)	Cellular replacement for mobile broadband voice and data services

JOHN CRAIG

The applications of self-organizing networks, says the MIT Media Lab's Michael Hawley, "really are going to be as magical as anything we've seen in technology."

node is a little bit of a router," says Poor. "It sends the message on in the right direction."

The initial application of Ember networks will be the replacement of expensive wiring in areas like factory floors. Other uses for the system will emerge as prices drop. For example, temperature, light and motion sensors could be placed in office buildings and networked to the lighting and ventilation systems. The network would then be able to tell when people were working in different areas of the building and switch on lights, heating or cooling only when needed. The same idea will eventually be applied to home automation, says Ember's acting CEO Adrian Tuck. Tuck cites security systems as another emerging application. For example, networked biological-weapons sensors placed in air-conditioning ducts throughout a building or in water treatment plants could offer early warning of a terrorist attack *(see "Networking the Infrastructure,"* TR *December 2001)*.

Motorola researchers see agricultural applications as well. Moisture sensors distributed throughout a field could be networked to irrigation systems, signaling the giant sprinklers to activate only when a part of the field was dry, instead of at regular intervals, saving water and money. The same scheme might be employed in a backyard sprinkler system.

Self-organizing architectures are also appearing for more complex networks. Engineers at IBM's Almaden Research Center in San Jose, CA, are prototyping a data storage system made up of "collective intelligent bricks": densely packed devices each consisting of a microchip, some memory and several hard-disk drives. Several hundreds of the bricks would be combined to create a single massive storage system. Software allows the bricks to recognize the addition of new bricks and figure out the best way to send data between them for storage. Similarly, if a brick fails, the system finds a way to route around it.

The goal of the brick system is to make storage servers simpler and cheaper to manage. Moidin Mohiuddin, the lab's senior manager of advanced storage systems, estimates that one administrator can currently manage about one terabyte, or one trillion bytes, of data. He hopes a system composed of the bricks could increase that figure a thousandfold. Moidin says the architecture could work as well for other types of servers and ultimately even PCs, making setting up a home or office network as simple as turning on the machines.

As they make their way into more and more systems, self-organizing networks will do no less than transform the way we relate to everything from our computers to our appliances, making them, if not smarter, at least more helpful. "I think [the networks] will turn up in all sorts of creative ways," says MIT's Hawley. "The result is going to be a radical simplification of the way we interact with the stuff around us."
—Erika Jonietz

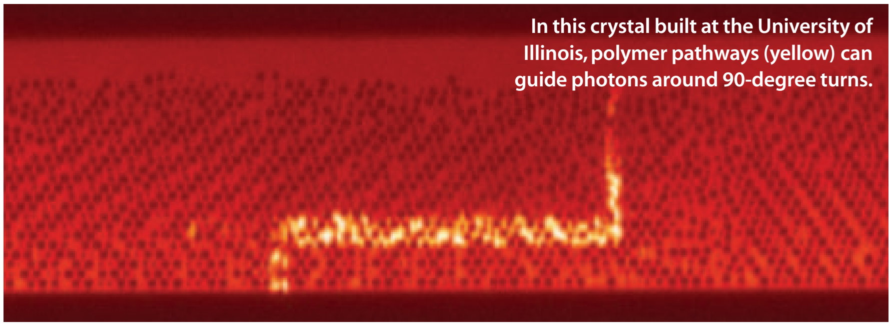

In this crystal built at the University of Illinois, polymer pathways (yellow) can guide photons around 90-degree turns.

GUIDING LIGHT

PHOTONICS | To Paul Braun, the future of optical computing is crystal clear. Braun and his colleagues at the University of Illinois at Urbana-Champaign report that they've found a cheaper and simpler way to construct tiny optical "waveguides" inside photonic crystals. These waveguides have the potential to behave like the microscopic wires on a conventional microchip, except that they would transport photons rather than electrons around tangles of sub-micrometer-scale circuitry. And that could help make photonic crystals the basis for a new generation of far faster telecommunications and computing devices.

Photonic crystals are intricate microscopic structures pocked with regularly spaced holes, like an orderly Swiss cheese. The holes create a barrier against light of certain wavelengths, and in the right arrangement, they can force photons along prescribed paths. Unlike optical fibers, which leak light when bent too far, these waveguides can hurl photons around sharp corners, making them ideal components for optical switches, microlasers, light-emitting diodes, even all-optical integrated circuits.

While companies like Agilent Technologies and a number of academic and government labs are developing photonic crystals, creating pathways that snake through them with the required micrometer-level precision is a major challenge. Several research groups, including one at Sandia National Laboratories in Albuquerque, NM, have built and tested photonic-crystal waveguides on silicon wafers, but their fabrication technique is the same complex, repetitive and expensive lithographic process used to pattern today's microchips. "It's a wonderful technique—if you don't care what it costs," Braun says.

Braun's technique starts with tiny silica spheres that assemble themselves in solution into a three-dimensional, crystal-like structure. Braun's real achievement was finding a way to create precisely shaped pathways through these crystals after they assemble: his group fills the space between the spheres with a photosensitive polymer, then uses a microscope to focus a laser on specific points, causing the polymer to harden. Drain the surrounding, unhardened polymer, and the result is a "defect" in the crystal: a perfectly sculpted pathway through the spheres, built with only one pass of the laser.

"A lot of people have been thinking about how to add defects to these self-assembling materials," says David Norris, a photonics researcher in the Chemical Engineering and Materials Science department at the University of Minnesota. "Paul's group has shown the first example of how that might be done." While Braun says it could take three years or more for self-assembling photonic crystals to find their way into commercial devices, he expects to demonstrate a working prototype—perhaps made from a material such as silicon that transmits light more reliably than the polymer—within the next six months. —*Wade Roush*

SPEEDIER DELIVERY

LOGISTICS | In an effort to reduce futile delivery efforts and the stranding of valuable packages on front porches, researchers at Accenture Technology Labs in Chicago are developing a Web-based software system for package shipping and receiving.

When a package reaches a certain point in shipping—say, the regional hub where trucks are loaded for local routes—it's scanned and identified much as parcels tracked by United Parcel Service and Federal Express are now. The software then sends out an alert to the package's recipient, based on preferences he or she has preselected online: e-mail, an instant text message or a phone call. A customer can even grant the shipper access to an online personal calendar in order to coordinate a preferred delivery time and can send the system new delivery instructions if the original time or address is inconvenient. The software automatically updates the driver's delivery list, saving the precious time it takes to drive to an empty home and write a missed-delivery note.

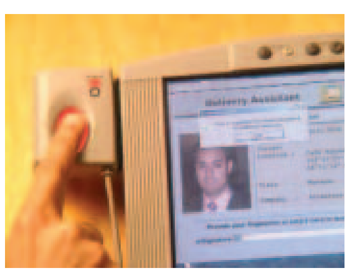

Accenture's package delivery system can confirm a recipient's identity using a portable fingerprint reader.

While Accenture will not divulge which of its corporate clients are interested in the prototype, the researchers foresee widespread use of the software within a few years. And since the system is built using universal Web standards, says Michael Hoch, analyst at the Boston-based Aberdeen Group, "You can expect to see any number of companies develop similar services in the next three to five years." Which could mean the end of those dreaded missed-delivery sticky notes once and for all. —*Kevin Hogan*

COURTESY OF ACCENTURE (SPEEDIER DELIVERY); COURTESY OF PAUL BRAUN AND WONMOK LEE (GUIDING LIGHT)

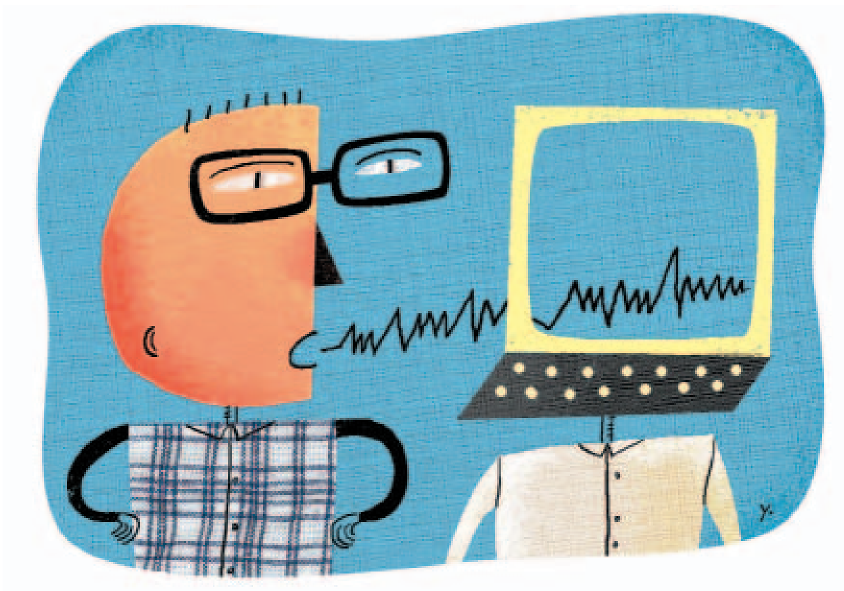

SPEAK EASY

IBM aims to solve speech recognition's nagging problems

SOFTWARE | The idea of computers that accurately understand human speech has both enticed and frustrated engineers. But now, IBM Research in Yorktown Heights, NY, is undertaking a multiyear project to finally solve all the problems that have kept voice recognition systems from comprehending free-form conversations—and becoming mainstream technology.

IBM aims to create a system that understands perhaps 20 languages, including medical and legal terms, with about 98 percent accuracy—a big improvement over the 80 to 85 percent accuracy of IBM's own speech recognition products and those from firms such as Peabody, MA-based ScanSoft. Troubles with accuracy are largely to blame for the limited market for speech recognition, which has so far been relegated mainly to dictation and telephone-based automated-response applications. IBM also hopes to overcome the other limitations of current systems: the need for hours of training, quiet surroundings and steady voice inflections. By making voice recognition more accurate and more broadly applicable, IBM believes it could open markets in real-time transcription for business meetings and new voice interfaces for handheld computers, or for search engines that could retrieve sound bites from audio databases of news broadcasts and speeches.

In current speech recognition technology, algorithms compare the waveform, an electronic representation of a word, to a master waveform database to develop a short list of possible matches, then select the most commonly used word on that list. IBM is exploring ways to make better matches, including new algorithms that make guesses based on the context of the conversation. IBM researchers have also built a lip-reading video system that reduces errors by one-third, says David Nahamoo, group manager of Human Language Technologies for IBM Research. "We're combining audio and visual features together, which we're feeding into our recognition engines," he says. "We're learning how to use one to clean up the other."

Some experts are skeptical. Real-time meeting transcription is still "lab stuff right now," says Steve McClure, vice president at technology market researcher IDC in Framingham, MA. "I've seen IBM demos work fine one time, and another time the damn application wouldn't work at all." Nahamoo concedes the initiative needs years of work and lots of luck to reach its goals. But given the speech recognition industry's history of failing to deliver on its promises, Big Blue's newest push could provide a few words of encouragement to the struggling technology. —*Alan Joch*

HIGH-TECH DIRT PILES

CONSTRUCTION | **Keeping accurate tabs on every element of a construction project, from piles of earth to stacks of steel beams, is complicated and expensive. To ease the burden on builders, engineers at the National Institute of Standards and Technology are developing information networks that could automate the process.**

In one experiment, a construction site is rigged with a Global Positioning System antenna, a computer equipped with a wireless Ethernet connection and a laser-based measuring device. The laser scanner analyzes the size of an object—say a mound of excavated dirt; measurements are sent via wireless Ethernet to databases and file servers that can be accessed by contractors and engineers both on and off the site. Software puts the data into an intelligible form—say, a 3-D model for monitoring job status—and can provide precise measurements for billing purposes. "Right now, many estimates for jobs like ground removal are taken only by how many trucks were used to haul the stuff away," says Geraldine S. Cheok, a research structural engineer at NIST. "This will make the numbers much more exact."

Ultimately, the NIST system would go beyond measuring dirt piles. Researchers plan to use radio frequency identification tags to track every pipe, beam and hammer that enters or leaves a site. While the researchers expect to have the technology ready for field use by 2006, the building industry is notoriously slow in adopting new techniques, says Ken Eickmann, director of the Austin, TX-based Construction Industry Institute, a research organization that looks for better construction practices. "But if it proves to be a cost saver," he says, "you will see it in practice." —*Kevin Hogan*

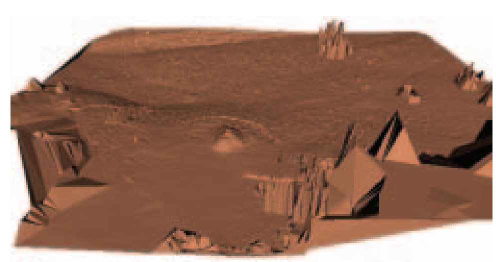

NIST's system generates a 3-D model of a construction site's terrain before excavation.

COURTESY OF THE NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY (DIRT PILES); JAMES YANG (ILLUSTRATION)

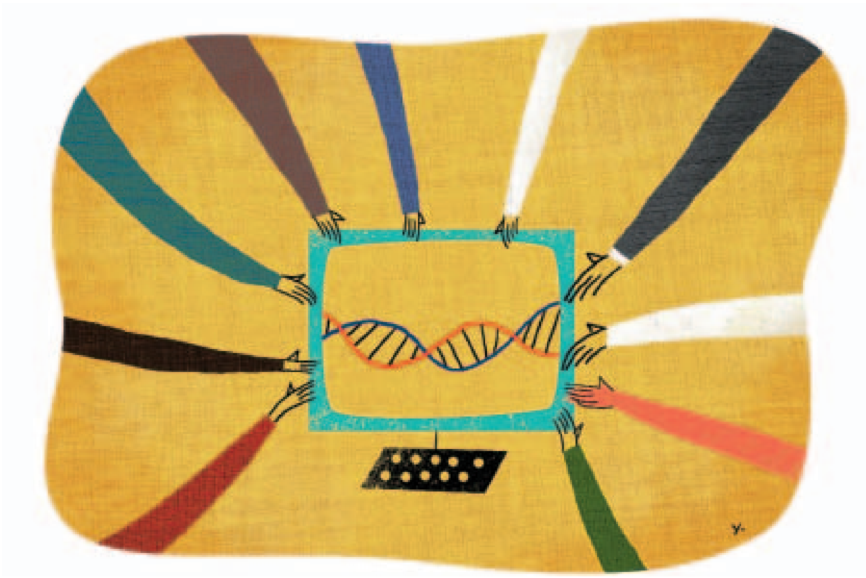

BIOTECH COMPUTING OPENS UP

An African firm turns to open-source programming

SOFTWARE | Open-source programming created a revolution in operating systems, making Linux a popular alternative to Microsoft's Windows. The idea—to make software source code open for modification by anyone—has caught on in large part because more eyes on the software means rapid improvements and fewer bugs; companies like Red Hat have turned the idea into profits by selling easily installed, well-supported versions of Linux. Now Cape Town, South Africa-based Electric Genetics is, for the first time, applying that business model to biomedical software.

In June, Electric Genetics plans to launch its first open-source product, a package of programs that link human DNA sequence data to information about how and when genes are turned on, and about the proteins encoded by those genes. "The science behind genomics is changing weekly," says cofounder Tania Broveak Hide. "A commercial software company, with your typical 12- to 18-month development cycle—I really don't think that works for a fast-moving scientific discipline."

Broveak Hide and her husband Winston Hide—the University of the Western Cape bioinformaticist with whom she cofounded the company in 1997—reached that conclusion after first selling proprietary software—and after their own programmers suggested that open-source development would yield better software faster. "I think it'll make a huge difference for the scientific discoveries," says Broveak Hide. "If we can push out better technology to the pharmaceutical companies, they're going to be able to make their discoveries faster."

Though others have yet to follow Electric Genetics' lead, industry observers say open-source development will be critical to improving the quality of bioinformatics software and, ultimately, biomedical research. *—Erika Jonietz*

A BETTER VIEW FOR ADVANCED SURGERY

MEDICINE | **The advent of "minimally invasive" surgery, performed with slender instruments through tiny incisions, has meant less trauma and faster healing for patients. But the technique requires surgeons to watch a video or ultrasound screen while operating to see what's going on underneath the skin—an awkward proposition for the surgeon. A head-mounted virtual-reality apparatus, developed at the University of North Carolina and now in clinical trials, could offer doctors a more natural view and allow for faster, safer operations.**

The trial—the first conducted with such a device—will involve 24 women undergoing breast tissue biopsies. A surgeon dons headgear incorporating glasses-like displays and two cameras mounted in front of her eyes. A computer merges ultrasound information flowing from a probe held to the patient's skin with video taken by the cameras, showing the operation site. This combination gives the surgeon a view into the body that corresponds with her natural perspective. What's more, the system tracks every twist and turn of the surgeon's head using ceiling-mounted cameras. Software adjusts for these movements to keep the ultrasound and video components of the surgeon's view in sync.

Harvard Medical School's Ferenc Jolesz, director of image-guided therapy at Boston's Brigham and Women's Hospital, says the North Carolina device "may lead to fundamental change in surgical visualization." The technology does, however, have competition. Siemens is developing a similar device and is now seeking an appropriate clinical trial.

The North Carolina device might take a decade to reach operating rooms, says computer scientist Henry Fuchs, who led the team that developed it. So far, a surgeon working with Fuchs has performed four operations; Fuchs hopes that others will be completed this year. After that, he aims to seek an industrial partner to further develop the technology. Eventually, head-mounted devices could be used for more challenging procedures like liver biopsies, which would become less awkward for surgeons—and safer for patients. *—David Talbot*

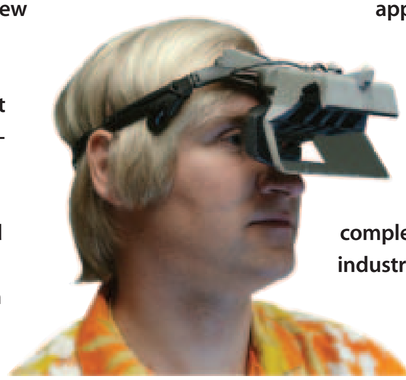

University of North Carolina engineer Kurtis Keller helped build this surgical visualization device.

ANDREI STATE (BETTER VIEW); JAMES YANG (ILLUSTRATION)