

**L**asers, vision systems, RFID and other advanced technologies often start out as solutions looking for problems. In the early stages, researchers and academicians are all agog over the possibilities, but end users, industrial machine builders and system integrators just like you remain skeptical. "Prove to me that it works, and I'll buy it," is your time-honored response to pioneering claims of new technology performance.

And so it goes with mesh networks. Our staff still is looking for evidence that they work in a control application. You've probably heard stories of the benefits of self-healing mesh networks, and that they're being used in universities, farms, wastewater treatment plants and IT installations. You may even have a mesh net-

working mesh network being used in a real process control, machine control, or discrete automation application.

#### A MESH OF MONITORING APPS

We contacted readers and most every process control, automation and networking vendor, looking for industrial mesh networks. Our research revealed some industrial monitoring applications are starting to emerge, especially on the wireless side. For example, at Intel's Ronler Acres semiconductor fab plant in Hillsboro, Ore., vibration engineer Mick Flanigan works with vibration sensors linked by a mesh network. It's all still in development, Flanigan says, but the results are promising.

# Heal Thyself

Mesh Networks Restore Broken Links, Enhance Reliability and Show Potential for Success in Control Applications



work at your company headquarters, linking all the PCs in accounting and human resources.

It's been shown that wireless mesh networks can save a fortune in wiring and installation costs (Figure 1). Preliminary results from monitoring applications are promising. The problem is finding a



FIGURE 1: WIRES, WIRES EVERYWHERE

As efficient as this wiring installation is, mesh networking might help eliminate the miles of wires and cables in some process control or automation applications.

"We monitor pumps, compressors, and other plant facilities equipment with vibration sensors," he reports. "Each machine has as many as 12 vibration sensors that report via wireless to a central wireless processor, which converts the data and transmits it to a PC that runs asset management software."

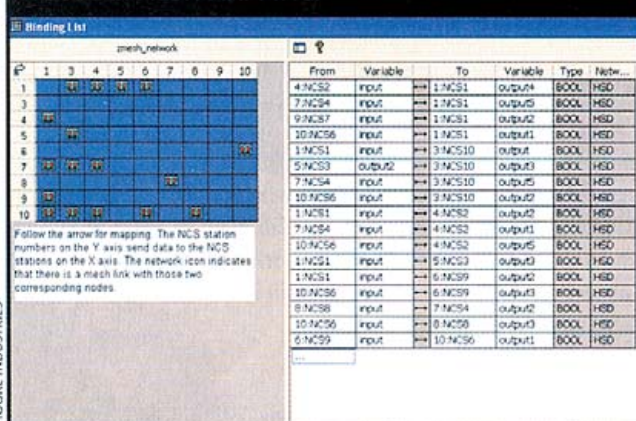
Robert Poor, chief technology officer of Ember ([www.ember.com](http://www.ember.com)) has been promoting an installation of Ember's wireless network in an unnamed water treatment plant. He points out that the wireless network functions well in the presence of concrete walls, pipes and other obstacles.

Mark Pacelle, vice president of marketing at Millennial Net ([www.millennialnet.com](http://www.millennialnet.com)) pointed out two installations that use his company's mesh network equipment. One is an energy management company. "Wiring is an expensive part of implementing an energy management application," says Pacelle. "With wireless, no re-wiring work is required in a retrofit situation and the cost barrier has been eliminated."

He also described a natural gas processing company that is using wireless. "Natural gas byproducts such as hydrocarbons, oil and water are stored in tanks in gas fields and trucked off site on a scheduled basis," he explains. "The company uses wireless sensor networks to monitor gas pressures and fluid levels within the tank farm."

All these are sensor networks, and are used for monitoring only. It seems no one is controlling anything with a mesh network. Not yet, anyway. There's evidence that the technology is in beta test, but on balance there were very few companies, if any, willing to talk about it. It's apparent the jury is still out on this new technology in control apps.

FIGURE 2: MESH MAPPING



This display shows the stations currently connected in a mesh network. Wherever the network icon appears, a link exists between the station on the X axis and the station on the Y axis.

## BETA SITES EVERYWHERE

One reason we can't find mesh networks in process control and automation is that the products are pretty new. Of all the controls vendors we asked, only two replied. Moore Industries (www.miinet.com) announced its Net Concentrator System (NCS) in 2003, and added wireless capability in April 2004. Moore's NCS is primarily a data acquisition system that can interface with any OPC-compatible DCS, SCADA, or PC-based system. It works with several types of networks.

"The NCS is a smart, rugged I/O system that allows not only control but universal channel mapping across Ethernet, fiber or wireless networks," says Scott Saunders, director, strategic marketing at Moore. "The mapping works on a simple producer-consumer concept where any NCS node in the network can communicate with any other node on that network. This allows pertinent data to be exchanged between those nodes that only need to send it or see it." Figure 2 illustrates a map of a NCS mesh network, which is generated by the NCS when it is running.

## INSIDE MESH NETWORKS

Mesh networks based on hardwired Ethernet switches and routers have been used for several years, but they were called "switched networks." The term "mesh network" really didn't become popular until wireless mesh sensor networks attracted our attention in the past year or so.

In a mesh network (see Figure), any single device can link to any other device, and the connection path from one device to another can take many different paths. This is also called "peer-to-any-peer" connectivity, and often is diagrammed like a big spider web, or "mesh," to illustrate the many paths possible.

If one device fails, "self-healing" software finds another path through the mesh. This is accomplished with special routing software algorithms that find or "discover" an alternate path when the primary path fails.

In cases where time is not an issue, the mesh network reconfigures itself constantly, always discovering the strongest link between nodes. This is known as a proactive-discovery network, because it is always looking for links.

This approach might not be suitable for industrial control, where speed is a major issue. For speed, the mesh network can use single-path routing; that is, the link between nodes A and B is always the same, and the system never seeks an alternate route except when a failure occurs. In some cases, the entire route from A to B to Z is defined and used at all times. If the link fails for some reason, then the self-healing software will find another route.

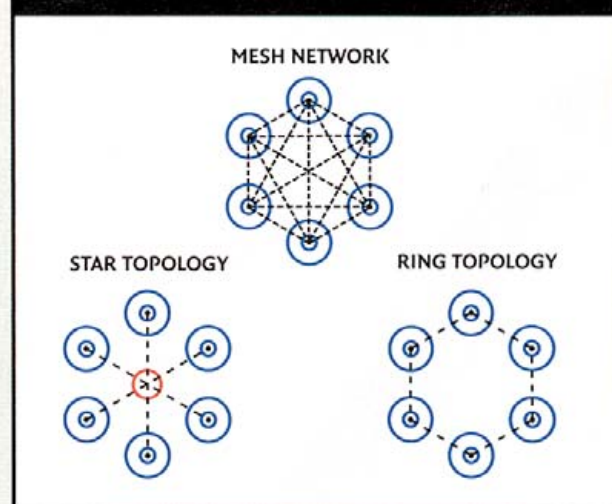
In a traditional star network, devices communicate through a central hub. In a ring network, devices are linked by a cable that runs from device to device. For redundancy, these networks often use a second cable, so if a break occurs, the device automatically switches to the alternate wire. Because a wired

mesh network simply finds another path, it does not need redundant wires.

Two types of mesh networks exist: wired and wireless. The primary difference between the two is that a wireless network node can directly communicate with any other node that's close enough to receive its RF signal; wired mesh networks must go through an Ethernet switch to reach other devices. Some mesh networks can combine wired and wireless.

Wireless nodes contain all the software needed for self-healing and discovering alternative communications paths. In wired systems, all such intelligence is built into the switch.

### MESH VS. TRADITIONAL



In a mesh network, nodes can reach each other over multiple paths. Star and Ring topologies are not as accommodating.

Saunders could only come up with one example of an actual installation, a retrofit to an existing network at a Southern Nevada Water Authority (SNWA) potable water plant. "SNWA installed the network to gather and share information between remote sites, pumping stations, filter beds, and the control room," says Saunders. "There already was a fiber and Ethernet network in place, so the NCS fit perfectly into his existing network infrastructure. A few sites were connected to the network via an unlicensed radio system."

Because the COTS switches have a maximum number of ports (8–24 depending on the switch), the wired system does not have the true peer-to-any-peer connectivity of a wireless mesh. However, it still provides the alternate paths and self-healing functions of a mesh network.


Foxboro ([www.foxboro.com](http://www.foxboro.com)) announced the Mesh Control Network Architecture for its I/A Series Version 8.0 control systems early in 2004, but did not start shipping until November 2004. "We've shipped I/A Series Version 8.0 systems incorporating the new technology for more than 20 different customer projects around the world," claims Roland Gendreau, product marketing manager at Foxboro. "These are in various stages of implementation, from engineering to site acceptance tests. By the time this appears in print, several I/A Series Version 8.0 systems using Mesh Process Control Network technology will have been fully commissioned."


As an example, Gendreau described the Veracel pulp mill project, a greenfield, 900,000 metric ton/year bleached eucalyptus pulp mill under construction in Eunápolis in the State of Bahia, Brazil. Foxboro will provide a large I/A Series system with all stations

linked via the mesh network configuration. More than 7,500 I/O points and 20,000 communication points will be integrated into the system. As of this writing, only one or two of the smaller units are up and running at the Veracel site.

Foxboro's mesh system is based on COTS (commercial off-the-shelf) Ethernet switches

(Figure 3) from IT world suppliers Enterasys ([www.enterasys.com](http://www.enterasys.com)), Allied Telesyn ([www.alliedtelesyn.com](http://www.alliedtelesyn.com)), and Cisco Systems ([www.cisco.com](http://www.cisco.com)). Foxboro developed a hardware and software interface between the I/A series and the switches, so that its process control systems could talk via a mesh network.





Because the COTS switches have a maximum number of ports (8–24 depending on the switch), the wired system does not have the true peer-to-any-peer connectivity of a wireless mesh.

However, it still provides the alternate paths and self-healing functions of a mesh network. By using a

series of edge switches and chassis managed switches (Figure 4) it can provide multiple paths between a control processor and a workstation.

#### **JURY STILL OUT?**

Until we actually see one of these control systems come to life and hear what the end

user has to say about it, we have only vendor opinions about how well they work.

Saunders says the NCS supports 10/100 Mbps over Cat 5/6 Ethernet cable, supports standard 10/100 Mbps switches and managed switches, and has a wireless link module (WLM) that can be used for distances to 20 miles line-of-sight. “To my knowledge,” claims Saunders, “so far everything is communicating just fine at the water plant installation.”

Initial field acceptance tests and site acceptance tests indicated that the mesh network communications performance and switchover times are meeting, and often exceeding, Foxboro’s published specifications, says Gendreau.


Foxboro claims mesh networks actually work better than conventional systems. “The performance of mesh networks that have been designed, configured and installed correctly are fully suitable for all process control applications, and can offer users a much higher level of performance than was previously available with their conventional control networks,” says Gendreau.

Without user reports to back up vendor statements like that, we have to withhold a verdict. As you know, we like to talk to end users about their experiences, and not rely solely on marketing optimism.

#### **MAKES SENSE THOUGH**

Nevertheless, what Foxboro and Moore say makes a lot of sense. It might very well be that hardwired mesh networks are superior to conventional networks. Gendreau offers five major reasons why switched mesh networks are better than conventional systems:

- **Reliability**—a single switched-Ethernet network in a mesh topology could provide a higher level of reliability and fault-tolerance than a conventional dual-redundant Ethernet network. “At best, dual-redundant Ethernet networks can be configured so that no single fault will stop communications,” says Gendreau. “In contrast, a self-healing mesh network can maintain virtually non-stop communications even when faced with multiple faults.”

**FIGURE 3: OFF-THE-SHELF SWITCHING**

Wired mesh networks are based on COTS Ethernet switches, which have been working in IT for years.

- **Determinism**—Switched Ethernet industrial networks provide full-duplex communications with 100 megabit bandwidth between process control system stations and network switches and 1 gigabit between the switches themselves. “This eliminates the potential for collisions and ensures that packets arrive at the desired destination in real time,” Gendreau explains. “Unlike conventional networks, where typically only a fraction of the available bandwidth is used to avoid collisions, full-duplex communications makes the full bandwidth of the network available to every station.”

- **Security**—Ethernet switches can provide an additional layer of security not available with conventional networks. For example, Enterasys switches can detect abnormal behavior on the network and intervene to quarantine the offending user or device, Gendreau says.

- **Scalability**—A mesh network scales from relatively small point-to-point networks to mega-sized networks. “Up to 1,920 stations, including workstations, control processors, switches, etc.,” explains Gendreau, “can be accommodated on a mesh network by interconnecting Ethernet switches.”

- **Lifecycle costs**—Mesh process control networks are typically less expensive to install than conventional networks because redundancy of cabling media is not required and most components are standard COTS, rather than proprietary. “Furthermore, as the cost/performance ratio of commercial networking technology continues to improve dramatically, these hardware costs will continue to drop,” Gendreau adds. “For example, a 16-port fiberoptic switch that we initially used is being superseded by a 24-port switch at a third less cost.”

Saunders offers a similar list of advantages with the Moore system:

- Each node can communicate with any other node on the network without sending its information through a common node or communication gateway.

- Any I/O point at any node can easily be transferred or duplicated at any other node on the network.

- Each node can monitor the health of each independent communication link with all other connected nodes. If faults in these connections exist, the node can notify the outside world.

- When a communication link fails, the node and its associated output signals are user-configurable to hold the last good known output, or go to a predetermined value.

- Once a faulty communication connection is remedied, the network automatically resumes communications without the action of external personnel. Try that with a conventional network, Saunders seems to be saying.

Throw in another advantage: wired mesh networks conform to common practices and standards such as IEEE 802.1W Commercial Rapid Spanning Tree Protocol (RSTP) technology, 802.3 Ethernet, 802.1D MAC Bridges, and so on, because they are based on COTS devices. Some conventional networks conform to these standards but others are still proprietary. Wireless mesh networks are still being standardized.

#### WIRELESS: STICK TO MONITORING

Wireless systems promise amazing benefits. Imagine Figure 1 without all those wires. As costs come tumbling down and more products become available, it

becomes more and more attractive to consider wireless monitoring and control.

But wireless mesh networks still have a long way to go before they meet your criterion of "Prove to me that it works, and I'll buy it."

Craig Skarpiak, product manager of wireless products at B&B Electronics ([www.bb-elec.com](http://www.bb-elec.com)) says, "There's little doubt that wireless mesh sensor networks currently are in the early adopter stage of the market." He sees products coming from major equipment OEMs including building automation, lighting, HVAC, electric motors, security and home appliances.

Skarpiak is concerned about the lack of finalized and approved standards. "There are several flavors of mesh networks, including standards-based ZigBee Alliance and 'enhanced-feature' networks from Dust Networks, Ember, Helicomm, Millennial Net, and Sencisat, to name a few."

To industrial users, a lack of standards can be the kiss of death for a new technology. "Standards organizations such as ZigBee, the 802.15.4 committee, and the Wireless Industrial Networking Alliance (WINA) all are working to make wireless standards and products more appropriate for industrial applications," says Hesh Kagan, consulting engineer, at Invensys Process Systems ([www.invensys.com](http://www.invensys.com)).

"Realizing that adopting a standard such as ZigBee will enable the marketplace to move forward more rapidly," Skarpiak adds, "several of the current 802.15.4 enhanced-featured networks providers created agnostic platforms that will incorporate the benefits of their system, while at the same time leverage the customer economics of a ZigBee standards-based system."

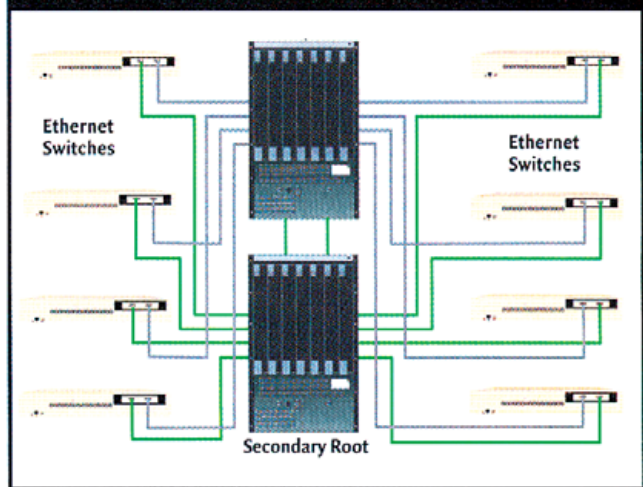
Another major problem is speed. Wireless mesh sensor networks are superb for non-critical monitoring applications, because those systems don't care if a temperature or a level value takes one microsecond or 5 seconds to arrive. In fact, many wireless monitoring systems sample at glacial speeds, if only to preserve the batteries in wireless systems.

Such a long response time is not suitable for real-time control applications in a machine or a fast process loop. Richard McGrail, COO Avek/B&B Electronics puts it bluntly: "Wireless is useful for

**"At best, dual-redundant Ethernet networks can be configured so that no single fault will stop communications. In contrast, a 'self-healing' mesh network can maintain virtually non-stop communications even when faced with multiple faults."**



**FIGURE 4: A WIRED MESH**



A wired mesh network doesn't have the any-peer-to-any-peer capability of a wireless network, but it still allows multiple paths between nodes by using Ethernet switches.

turning on and off a building light, but I wouldn't trust it to turn off a machine just before it takes my hand off."

There is not much you can do to speed up a wireless system. The problem, according to Francis de Costa, founder and CTO of MeshDynamics ([www.meshdynamics.com](http://www.meshdynamics.com)) is that radio is a shared medium, and it forces all stations to remain silent while one node transmits.

A "request-to-send" by one node silences all nodes within range. "That does not sound too bad in a wireless sensor network, but it is disastrous if you wish to provide the level of latency and throughput people are accustomed to with wired networks," says de Costa. He explains that a system with 10 nodes could be reduced to where one node is operating at as little as 1/16,000 of the total bandwidth.

Why has this problem not been noticed before? "Because, first, there are not a lot of mesh networks around and, second, they have not been tested under high-use conditions," he explains.

Wireless mesh sensor networks seem ideal for non-critical applications such as monitoring wastewater plants, tank farms, pipelines, HVAC systems, energy management, vibration and similar tasks. But, until IEEE, Zigbee and WINA standardize wireless, and figure out how to make wireless fast and deterministic, you may want to put off developing wireless machine controls. The wiring scenario portrayed in Figure 1 might be here for a while.

From what we've seen, it appears that hard-wired mesh networks are the future of process control and automation. Wired mesh networks seem to be the next step in the never-ending evolution of Ethernet.

Wireless mesh networks save wiring and installation costs, but have yet to prove themselves for control. The ideal solution may well be a mixture of the two—once vendors prove to you that their systems work, of course. ●