

Enhancing Products through Connectivity

ember

Technical White Paper

Reliable Wireless Networks for Industrial Systems

Mesh Network Topology Developed at MIT's Media Lab
Exhibits Superior Reliability Over Traditional Wireless Systems

By Robert Poor and Brent Hodges

Ember Corporation

Boston, Massachusetts, USA



Wireless systems for industry have mostly used “cellular phone” style radio links, using Point-to-Point or Point-to-Multipoint transmission. Research at MIT’s Media Lab in Cambridge, Massachusetts conclusively proved that traditional wireless formats have substantial liabilities in industrial applications: rigid structure, meticulous planning requirements, and dropped signals.

In contrast, wireless mesh networks are multi-hop systems in which devices assist each other in transmitting packets through the network, especially in adverse conditions. They can be dropped in place ad hoc with minimal preparation and provide a reliable, flexible system that can be extended to thousands of devices.

◎◎◎ The Potential of Wireless

Factory communication wiring can easily have an installed cost of \$5-10 per foot. What if expensive communication wiring could be replaced with reliable wireless links?

Productivity programs demand more and more information from smart devices. What if industrial gear could gain more local intelligence by sharing information with nearby sensors?

More and more maintenance systems require remote data acquisition. What if you could continually monitor the condition of all the equipment on your factory floor, not locating failures after the fact, but predicting them before they happen?

Eliminate Cable

An obvious problem that can be addressed with wireless solutions is simple wire replacement, where the radio frequency (RF) communication link emulates wire in an existing system. No changes are made to the system architecture. Rather, wireless links are used to transmit the same data that the physical wire once carried.

Consider an instrument connected by a serial cable to a control panel using Modbus as a communication protocol. Wireless RF links can replace the serial cable as the physical layer to carry Modbus packets back and forth, requiring no physical changes to the instrument, the control panel or the underlying software architecture.

The serial cable is taken away, and a wireless transceiver is physically connected to the serial port at both the instrument and at the control panel. Neither the control panel nor the instrument can tell that it is not using a cable.

The labor required to run this cable and conduit is not cheap. Installation cost is a growing concern for designers and facility managers, and labor rates continue to rise in most parts of the world.

So the benefits of wireless are fairly obvious: Expensive cabling has been eliminated. If these cables were in a hazardous environment such as a chemical processing plant, they would have to be isolated from potential contact with chemicals and placed inside conduit run through concrete walls to reach the instrumentation deployed throughout the plant.

Speed Up Installation, Reconfiguration and Expansion

Another benefit of wireless is the speed of deployment. Wired systems can take days or weeks to be properly installed, isolated and commissioned. Wireless networks require only the end points to be installed, saving hours or days for each instrument installed. Other instruments can be added as required without the need for expensive, disruptive cabling and labor.

A further benefit of the wireless system is the ease of reconfiguration and expansion. If there is need for a plant expansion, or relocation of instruments, there is no expensive conduit to be moved or added. If the instruments to be connected to the control panel need to be placed on mobile equipment, such as the mobile batch containers found in bio-tech, pharmaceutical and other specialty chemical installations, wireless offers an attractive solution.

“What if Radios and Microprocessors Were Free?”

This was the question that spurred Robert Poor to launch the MIT wireless mesh research project in 1997. There is a larger implication of simple, affordable wireless technology: The fact that billions of smart devices in businesses, factories, homes and vehicles can be transformed by the addition of wireless communications. We have already seen the staggering impact of the Internet since the mid 1990's, and there are far more embedded devices than people or PC's. As more and more devices are equipped with wireless links, an important revolution is underway.



But Has Wireless Really Been That Easy?

So far, using wireless systems to replace common links like serial cables seems easy. If only it were so straightforward! Users who have tried previously available wireless systems have a common list of complaints:

- 1 RF links are not as dependable as wires. Anyone who has used a cell phone, portable radio or CB knows first-hand about RF links. The signal is constantly changing as conditions change between the two points.
- 2 Expanding or moving an RF point is not always as easy as claimed, because a new position on the network may be out of range of the control point for the wireless network. This control point is commonly placed at the control panel in an industrial application.
- 3 Wireless installers sometimes offer assistance of professional technicians who perform RF site surveys to determine control points for the wireless network based on planned coverage areas. While useful, this adds highly skilled labor back into the installation cost and doesn't address ease of reconfiguration or expansion.
- 4 Some installations require additional wireless control points (sometimes referred to as wireless access points) in addition to the control panel.

These problems are evident when you try to talk on a cellular phone:

- Signals appear and disappear simply because you move your phone six inches
- You can hear the other person, but they can't hear you
- Calls get "Dropped" and require re-connection
- Problems are compounded by reflections inside of buildings (multipath)
- Interference from other RF sources garbles reception
- There has to be a cell tower nearby – cell phones can't directly call each other
- If too many phones are in use, the system is "busy" and calls can't get through

This is a very serious set of problems to face when reliability is of prime importance. So despite the increasing popularity of IEEE 802.11 Wireless LAN (Local Area Network) systems and the promise of Bluetooth systems, wireless communication has yet to be widely adopted in industrial applications. And while wireless systems seem like an obvious solution for industrial applications, in reality the cure can be worse than the disease. Solutions based on these standards were not designed with the industrial environment in mind. Industrial users need a network architecture that takes the unique challenges of the industrial environment into account.

The ARBORNET research team at MIT Media Lab examined embedded wireless solutions for smart devices from 1997 to 2001, led by Robert Poor. The resulting thesis, *Embedded Networks: Pervasive, Low-Power, Wireless Connectivity* concluded that traditional "cell phone" style wireless systems were simply inadequate for industrial applications, and could never gain widespread acceptance in their current form, but that a fresh approach was needed and, ultimately, was identified.

This new approach had to have the following characteristics:

- 1 The network should not require sophisticated planning or site mapping to achieve reliable communications. That adds specialized, expensive labor to the installation.
- 2 Human intervention should not be necessary for the network to move a packet from one end to the other. The network should figure this out by itself.
- 3 All devices must be able to transmit from where they are right now, and not have to be moved.
- 4 Nobody uses wireless for "fun," only for pragmatic reasons – namely lower cost and ease of installation. If it's not easier and less expensive than copper, then the promise of wireless will never materialize.
- 5 The network error rate should be below acceptable levels, as defined by the customer.

Control & Sensing Networks versus Data Networks

To understand why standard wireless networks do not work well for industrial applications, it is important to distinguish between control and sensing networks versus data networks.

Wireless data networks are primarily designed to link together computers, PDAs, printers, Internet access points, etc. where large amounts of data are sent both directions.

In data networks, the emphasis is on speed: faster is better. The design and evolution of 802.11 networks is a good example. 802.11b, one of the first wireless LAN standards to be widely adopted, connects devices at up to 11 Mbps. One of the latest updates to this standard, 802.11a, will allow for data speeds up to 54 Mbps, enabling more rapid downloads of music and video files by end-users.

But wireless networks for industrial control and sensing, above all, must be reliable, adaptable, and scalable. Because industrial sensors send only a few of bits of data per second or minute, providing information like temperature, pressure and flow, data rates of 11 Mbps or even 54 Mbps are rarely needed. Although speed is often the focus for data networks, the primary design objectives for industrial control and sensing networks are reliability, adaptability and scalability.



Essential Wireless Requirements of Industrial Environments

Requirement #1: Reliability

For most industrial applications, reliability is crucial: wireless systems must be just as reliable as traditional copper wire. Depending on the application, garbled or dropped data can result in anything from a disruptive glitch to a devastating failure.

Three factors determine the signal reliability between a radio transmitter and receiver:

- 1 Path loss
- 2 RF interference
- 3 Transmit power

Consider a conversation between two people. Path loss corresponds to how muted one person's voice becomes, due to distance or obstacles between them. The listener will have a hard time understanding the speaker if they are too far away or talking through a closed door.

RF interference corresponds to ambient noise: it will be difficult for the listener to understand the speaker in a noisy environment. Many other factors—including receiver sensitivity and data encoding technique—affect the reliability of a link. However, between a given radio transmitter and receiver, the path loss, interference, and transmit power determine the bit error rate.

The problems of path loss or interference can be overcome by moving closer to the listener or by shouting loud enough to be heard. In the wireless world, this corresponds to repositioning radios or by transmitting with a higher power.

Unfortunately, neither of these are generally viable options. Increasing the transmit power creates a situation similar to people shouting over loud music at a party. A sophisticated antenna design that directs the RF signal towards the receiving radio might help. But this is much like using a megaphone to shout at the listener. It does improve the path loss situation, but may fail if the listener or the megaphone moves.

Requirement #2: Adaptability

The network should adapt to the existing environment. The environment should not have to be altered to make the system “wireless ready.”

If you need a wireless link between a tank level sensor and a data logger, it is not practical to relocate the tank or the data logger just to create a reliable connection. In fact, a wireless link may be unsuitable for connecting tank level sensors and data collection points in pre-existing structures, as these are often immovable objects. If cables were already being used for this, more wire could always be run, though at a prohibitive cost.

In the wireless world, the network should integrate seamlessly with the environment. A key attribute of a good wireless network is that daily work activities and the facility layout are not a concern.

You never want to ask someone to move in order to hear them speak more clearly. Likewise, repositioning radios and equipment in order to increase communication reliability is not always a realistic option.

Requirement #3: Scalability

Any network, wired or wireless, should scale gracefully as the number of endpoints increases. Scalability is one of the attractions of fieldbuses over hard-wired “home run” systems: once the trunk line is in place, adding new devices is relatively easy. In many multi-drop networks, adding a new device is as simple as wiring the device directly into the network cable or a termination block at one end of the network. Eliminating the need to “home run” wire the new device back to the control panel has reduced wiring.

In a wireless system, all devices on the network share the airwaves. Simply transmitting with more power can increase the reliability of a single transmit / receive pair, but as soon as multiple devices share the airwaves, this approach may actually decrease overall reliability.

It's not unlike being in a large, noisy restaurant where people are speaking loudly to be heard and no one can understand anyone else. Similarly, transmitting with more power in order to increase reliability is not consistent with building scalable networks that support numerous endpoints.

◉ ◎ ◎ Wireless Formats

Point-to-Point Links

Sometimes referred to as a “wireless bridge”, a point-to-point link serves as a replacement for a single communication cable. A point-to-point link might be used to connect a Programmable Logic Controller (PLC) to a remote monitoring station as shown in Figure 1 below.



FIGURE 1: POINT-TO-POINT LINK

Point-to-point links can communicate reliably as long as the two endpoints are located sufficiently close to one another to escape the effects of RF interference and path loss. If a reliable connection is not initially achieved, it is sometimes possible to relocate the radios or boost the transmit power to achieve the desired reliability.

Point-to-Multipoint Links

Point-to-multipoint wireless systems, such as those based on IEEE 802.11 or Bluetooth, have one base station or access point which controls communication with all of the other wireless nodes in the network. Also referred to as a “hub and spoke” or “star” topology, this architecture has similarities to wired “home run” systems, in which all the signals converge on a single terminal block. A point-to-multipoint example is shown in Figure 2.




FIGURE 2: POINT-TO-MULTIPOINT NETWORK

Signals in point-to-multipoint networks converge at a single access point. The reliability of these networks is set by the quality of the RF link between the central access point and each endpoint.

In industrial settings, it can be hard to find a location for an access point that provides dependable communication with each endpoint. Moving an access point to improve communication with one endpoint will often degrade communication with other endpoints.

While it may be possible to wire together multiple access points in order to improve reliability, the cost of additional wiring can defeat the original reasons for choosing a wireless solution.



Wireless Mesh Networks

This wireless mesh network topology for industrial control and sensing, developed by the MIT Media Lab and produced by the Ember Corporation (www.ember.com) is a “point-to-point-to-point” or “peer-to-peer” system called an ad hoc, multi-hop network. A patent was granted for this technology, United States Patent number 6,028,857.

A node can send and receive messages, but in a mesh network, a node also functions as a router and can relay messages for its neighbors. Through this relaying process, a packet of wireless data will find its way to its ultimate destination, passing through intermediate nodes with reliable communication links. An example of a mesh network appears in Figure 3.




FIGURE 3: WIRELESS MESH NETWORK

There are some important things to notice in this figure:

- 1 The resemblance to a map of the Internet is not entirely coincidental. Like the Internet and other router-based communication networks, a mesh network offers multiple redundant communication paths throughout the network.
- 2 If a single node fails for any reason (including the introduction of strong RF interference), messages will automatically be routed through alternate paths.
- 3 In a mesh network, the distance between wireless nodes can be shortened, which dramatically increases the link quality between nodes. If you reduce distance by a factor of two, the resulting signal is at least four times more powerful at the receiver. This makes links more reliable without increasing transmitter power in the individual nodes.

More Nodes = Greater Reliability

Notice the addition of freestanding “repeater” nodes in the middle of the network. In a mesh network, it is possible to extend distance, add redundancy, and improve the general reliability of the network simply by adding repeater nodes.

Mesh = Self-Configuring

A network should not need a person to tell it how to get a message to its destination. A mesh network is self-organizing and does not require manual configuration. Because it is both self-configuring and self-healing, adding new gear or relocating existing gear is as simple as plugging in a wireless node and turning it on. The network discovers the new node and automatically incorporates it into the network without the need for a system administrator.

A mesh is not only inherently reliable, it is also highly adaptable. If your tank level sensor and data logger are placed too far apart for a solid RF communication link, you just lay down one or more repeater nodes to fill gaps in the network.

Mesh = Self-Healing

On the Internet, if one router goes down, messages are sent through an alternate path by other routers. Similarly, if a device in a mesh network fails, messages are sent around it via other devices. Loss of one or more nodes does not necessarily affect its operation. A mesh network is self-healing because human intervention is not necessary for re-routing of messages.

Mesh = Redundant

The actual meaning of “redundancy” in the real-world is a matter of degree and must be carefully specified. In a mesh network, the degree of redundancy is essentially a function of node density. A mesh network can be deliberately “over-designed” simply by adding extra nodes, so that each device has two or more paths for sending data. This is a much simpler way of obtaining redundancy than is possible in most other types of systems.

Mesh = Scalable to Thousands of Nodes

A mesh is also scalable, and can handle hundreds or thousands of nodes. Since the operation of the network does not depend upon a central control point, adding multiple data collection points or gateways is convenient.

It is clear that reliability, adaptability and scalability are the most important attributes of a wireless network for industrial control and sensing applications. Point-to-point networks can provide reliability, but don't scale to handle more than one pair of endpoints. Point-to-multipoint networks can handle more endpoints, but the reliability is determined by the placement of the access point and the endpoints.

If environmental conditions result in poor reliability, it is difficult or impossible to adapt a point-to-multipoint network to increase the reliability. By contrast, mesh networks are inherently reliable, adapt easily to environmental or architectural constraints, and can scale to handle thousands of endpoints. These attributes are summarized in Table 1 below:

Topology	Reliability	Adaptability	Scalability
Point-to-Point	High	Low	None (2 endpoints)
Point-to-Multipoint	Low	Low	Moderate (7-30 endpoints)
Mesh Networks	High	High	Yes (1000s of endpoints)

TABLE 1: SUITABILITY IN INDUSTRIAL APPLICATIONS



Real Industrial Applications of Wireless Mesh Networks

Wire replacement

At the beginning of this paper, an example of wireless serial link replacement was given. This is most commonly done with point-to-point or point-to-multipoint technology, but mesh networks still provide complete transparency. The network does not know that copper has been replaced with an RF link, but the mesh network is inherently more reliable, more adaptable, and scalable.

Distributed Control

A specific opportunity for wireless, multi-hop, mesh networks is in distributed control systems. There has been a trend in recent years to place more intelligence throughout the control system. The IEEE 1451 standard Smart Transducer Interface for Sensors and Actuators is evidence of this. Distributed intelligence is naturally served better by wireless multi-hop mesh networks, which do not require a central control topology.

The control of the wireless system is distributed throughout the network, allowing intelligent peers to communicate directly to other points on the network without having to be routed through some central point.

Modular distributed control systems are easier to install and maintain. Since more of the system logic is at the instrument or sub-system level, clusters of instruments can interact and make local decisions. This is often done with small PLCs, which gather information from nearby instruments or sensors and then provide processing power and decision making for this local instrument cluster. These clusters can then be connected as a group back into the main control system. The result is a less complex installation because individual instruments and points do not have to be directly connected to the main control panel.

Is Distributed Control Cheaper to Maintain?

Proponents insist that modular control systems are easier and less costly to maintain. The rationale is that highly modular control systems enable localized decision making, which results in faster isolation of problems within the system. These problems can usually be diagnosed back to a single instrument cluster, allowing engineers and maintenance staff to focus their attention on one area of the system.

Fast problem solving means less downtime when something goes wrong. Likewise, when the system is operational, local decision making by intelligent instruments and small PLCs identifies problems before they impact the entire system and cause bigger problems.

Finally, these modular sub-systems can be replaced or upgraded without affecting the entire system. These many factors make systems much cheaper and easier to operate and maintain.

Matching multi-hop, wireless mesh communication with distributed control facilitates a whole new dimension of interactions between sensors or sensor clusters.

Sensors can now communicate directly to other devices on the network. This topology allows a tank level sensor to communicate directly with nearby valves, alerting them to open or close to prevent an overflow situation. Monitoring equipment could take readings from sensors without having to directly access the sensor with wired connections. This is useful in calibration and troubleshooting.

Diagnostic Monitoring

A third area of application for wireless, multi-hop, mesh networks is in the diagnostic monitoring of devices. This monitoring can occur outside the normal control loop and wireless communication can be sent to notify the system user of any abnormal operation of the device. Take for instance the schematic of a sensor control loop shown in figure 4 below:




FIGURE 4: SCHEMATIC OF TYPICAL SENSOR CONTROL LOOP

In this control loop, an additional signal is extracted and analyzed during the course of normal operation of the sensor. As the sensor operates, the signal is monitored for abnormalities without affecting the sensor's operation. If an abnormal signal or trend is observed, an alert is triggered.

The beauty of using a wireless link for on-board monitoring and alert is that the monitoring link remains independent of the control loop. By using a wireless, multi-hop mesh network, data can be routed dynamically to similar wireless devices. Surrounding devices can respond to the alert from the failing device, even as the alert is being sent to maintenance personnel.

Another benefit of wireless is that maintenance personnel can directly connect to the diagnostic output of the sensor, without running wires. This can eliminate a huge task in the case of a tank level sensor in a large storage tank, or a temperature probe at the top of a tower stack at a chemical refinery.

In a wireless, multi-hop mesh network, a user can get that data via any wireless node on the network.

By using a diagnostic device with additional processing power (such as a laptop computer, handheld computer, handheld diagnostic device), maintenance personnel can check on configuration and other information about any node on the network. This information is a valuable tool for checking and verifying sensor operation when questionable data is received from a sensor through its primary control loop.

◎ ◎ ◎ Case Study: Water Treatment

To validate wireless mesh networks in challenging industrial environments, Ember Corporation deployed a system in a water treatment plant. The environment was typical of such facilities, with significant wireless environment hurdles such as thick reinforced concrete walls segmenting giant tanks of water with large numbers of metal pipes running between tanks (Figure 5).



FIGURE 5: TYPICAL WATER TREATMENT FACILITY

The goal was to connect the instruments in the pipe gallery back to the control panel located in the control room on the third floor of the water filtration plant. A look below at figure 6 provides a geographical representation of the instrumentation topology.

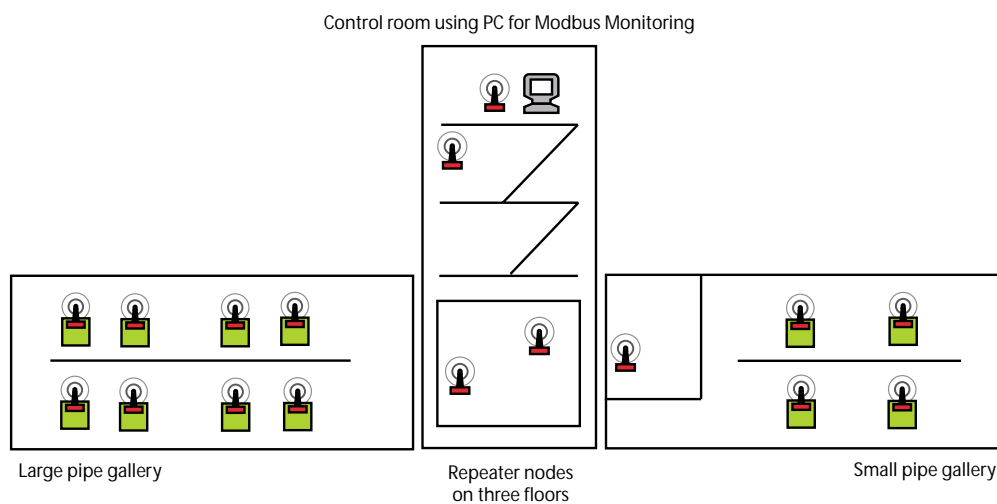


FIGURE 6: SCHEMATIC OF INSTRUMENT LOCATIONS

Figure 6 shows the approximate locations of eight instruments in the large pipe gallery along with four instruments in the small pipe gallery. The control room was located on the third floor of an attached concrete building.

Before wireless communication, a data collection PC in the control room communicated with process instruments over a RS-485 serial bus. The first step in converting this system to wireless networking was to replace this computer's bus connection with a wireless networking card connected to its serial port.

Each process instrument also had bus connections replaced with wireless networking cards, which self-configured on power-up and began attempting to send data to the control room. After all 12 instruments had wireless cards installed, it was possible to analyze the RF network traffic and determine where link reliability was below standards.

These areas included spots where RF signals had to pass through reinforced concrete walls and where a single link spanned two flights of metal stairs. Improving these RF links was a simple matter of dropping down additional RF relay points. This step was made possible by the network's lack of a central wireless control point and each node's ability to cooperatively relay packets on behalf of its neighbors.

After these "repeater" nodes were placed, the network was complete.

Installed and Operating in Two Hours

Time for complete installation was under two hours - compared to approximately twenty hours when each instrument had to be wired back to the control panel. The software on the PC did not discern any difference in the wireless communication network versus the wired serial cable network.

The wireless network exhibited less than 0.1% packet loss before any attempt was made to re-send lost packets through the network. This was accomplished via the mesh networking algorithms used by the wireless network. Neighboring nodes cooperatively relay packets over the best RF link.



Conclusion

Daily experience with some of the challenges of wireless consumer products, university research, and the commercial industry's slow adoption of wireless for use in enterprise applications, are indicators that products based on point-to-point and point-to-multipoint topologies are not well suited for use industrial enterprise communication.

Multi-hop mesh technology, however, is inherently reliable, redundant, and can be extended to include thousands of devices. Real-world examples cited in this paper demonstrate that mesh networks can be installed in hours instead of days or weeks and that these networks are highly dependable.

Wireless mesh networks now meet the objectives originally defined by the MIT ARBORNET team:

- 1 The network does not require sophisticated planning and site mapping to achieve reliable communications. There is no need for specialized, expensive labor to complete the installation.
- 2 The network is self-configuring, and does not require the assistance of a network specialist just to send a packet from one end to the other.
- 3 All devices are able to transmit from where they are originally, and do not have to be moved. A weak signal or "dead zone" can be fixed simply by dropping a repeater node in place.
- 4 Compared to the cost of specialized knowledge that's required to install traditional wireless systems, or the cost of point-to-point copper wiring and conduit, mesh is must less expensive. We are now closer to realizing the real potential of wireless smart devices.
- 5 The network error rate is very low (under 0.1% in the example cited above) and can be further reduced if occasional re-transmits are allowed.

Industrial systems can now benefit from a wireless format that satisfies the multiple conflicting demands of redundancy, distributed communication, flexibility and reliability.

Furthermore, flexible, self-configuring, self-healing networks are inherently less expensive to install and maintain and the ideal of “free radios and microprocessors” is closer than ever before. A significant barrier to low-cost connectivity has been removed.

○○○ Further Steps

Ember Corporation is installing beta systems at a select number of industrial sites, and will also be licensing wireless mesh technology to a select group of Original Equipment Manufacturers in specific vertical markets. For consideration for installation or licensing, please contact Brent Hodges at 617.951.0200.

A Conversation with Dr. Robert Poor,

Founder and Chief Technology Officer, Ember Corporation

“I was two weeks away from pursuing an audio digital signal processing project for my Master’s thesis at MIT when I tore up the proposal and threw it away,” says Robert Poor, whose extensive work in acoustics already included design of the legendary sound system used by the Grateful Dead. “I was gripped by the deep implications of a very simple question: What would happen if radios cost nothing and microprocessor power cost nothing? Right then and there I decided I had to pursue wireless networking.”

“I used to talk about networking simple devices like light switches and thermostats. People would look at me like I had two heads. That was a really radical thing to be talking about in the mid-’90s. Now when I say that, people shrug their shoulders and say “so what?”

Networking small amounts of data – sometimes as little as a few bits per hour – is an entirely different proposition than cramming megabytes through a computer data network. “We had to completely throw away all our assumptions about networking, shatter the mold and start afresh,” he says.

“It used to be that if you wanted a reliable network, you’d hire an expert RF contractor with signal generators and spectrum analyzers. He’d take measurements over time, charge you a boatload of money, give you a list of permissible access points and node locations. Heaven help you if you ever need to move something. What if you build a wall or move a filing cabinet? You have a “brittle” network now. That isn’t what you want. What you want to do is decrease cost of connectivity. In the 802.11 world of wireless voice and data, it’s done expressly for mobility. But in the industrial world, the mandate is simple: Get rid of wire.”

All the benefits of wireless are lost if: 1) it’s not dependable 2) it costs more than laying wire, or 3) if it costs money to make minor change. “All of these things must be true. And when you say “reliable,” I mean unimpeachably reliable. Something that your own guys can deploy without experts.”

He insists that reliability must be built in and cannot be tweaked by a human after the fact. “What if your wireless transceiver is installed on a 2 ton tank, and it doesn’t work? Are you gonna move the tank? Of course not. In a star topology, out of range equals out of luck. We had to have something fundamentally different: Nodes wherever dictated by application. If there’s not enough reception, just plunk down some repeaters. More repeaters makes a better mesh. Plunk is the operative word.”

Poor explains that a mesh works because it circumvents all the common obstacles to good radio transmission. “You need a good Signal to Noise Ratio, and environmental noise from motors and arc welders is a constant enemy. Another enemy is multipath, which is the RF equivalent of echo. If you can’t get closer to the source, your only recourse is more power. But power doesn’t solve multipath. Worse yet, imagine an oil field with 3000 valves. You don’t want everyone shouting!”

“The solution is a dense network, where nodes can relay data at low power, in a multi-hop fashion, instead of turning up the power. When you go to a multi-hop mesh, you can have thousands of devices, use small amounts of power, and eliminate multipath and interference.”

The project began in computer simulation, but then caught the interest of DARPA for military applications through their SensIT program – networks for battlefield awareness. It also received interest from the MIT Media Lab for its commercial potential. The MIT Media Lab is funded by a consortium of over 200 companies who come through on a regular basis to see what projects have potential for their designs. In academia, the mandate is “publish or perish” but in the Media Lab, it’s “demo or die.”

Andy Wheeler joined Robert Poor and played a significant part in developing working hardware. “Wheeler built a robust mesh network for a microclimatology experiment in Hawaii, for example. This whole project was never pie-in-the-sky, it was always grounded in reality. So by the time we started Ember, we’d already been through 2 hardware revisions. The company that grew out of this, Ember Corporation, has always had a thick streak of pragmatism right along side the vision.”

Dr. Poor emphasizes practical results over starry-eyed idealism, and this is evident in the success of the company’s installations as well as the demeanor of its staff. “Companies on outside say how impressed they are with everyone who works here, with their professionalism & enthusiasm.”

Even his approach to work itself has a guiding philosophy. “Life is short. We spend most of our adult days going to work, at work, or recovering from work. With that in mind, I think it’s a sin to squander a life on anything less than amazing. What a company does for you had better be amazing. And it’s a funny thing, because that ingredient is hard to quantify. But it has really put us in good stead, especially in the tough economy. Ember is really fortunate to have a solid product, great people and a team atmosphere. It’s infectious.”

The logo for Ember Corporation, featuring the word "ember" in a lowercase, bold, sans-serif font. The letter 'b' is stylized with a red dot above it.