802.11n Wireless Connectivity Supports Seamless Industrial Networks

The 802.11 WLAN standard, and in particular 802.11n, helps provide seamless connectivity and a means of achieving universal IP-based networking in industrial environments.

N.VENKATESH, REDPINE SIGNALS

Industrial environments have long been extensively networked. The objective has primarily been automation—the need to control and monitor complex manufacturing processes that bring together machines, actuators, controllers, data acquisition systems, supervisory systems and communication infrastructure, among others. These individual entities communicate among themselves, following somewhat rigid hierarchies of communication links. Networks have traditionally been designed through a layered approach, with the ISO seven-layer protocol being a well-known and standardized hierarchy, albeit rarely followed in toto. Various segments of an industrial network use their own methods of communicating, often resulting in a complex structure of data links.

There are almost as many types of networks as there are applications that use them. The attributes of the network are split into the attributes of each of the layers that comprise them—and they derive from the requirements of the application under consideration. For example, some devices on a network may require service at rigidly periodic intervals, demanding a transport mechanism that provides guaranteed timely delivery of information. Other devices may communicate in bursts, sending out unsolicited data at irregular intervals, and without needing a deterministic latency bound. Some devices may communicate to a multitude of other devices on the network, while others may only communicate with one other. These requirements translate to specifications on the physical layer, data link layer, network layer and the application layer, usually skipping the other layers in between.

Figure 1 illustrates this diversity. The physical layer in these connections could be RS-232, RS-485, Ethernet or Wireless, among others. Popular protocols include Fieldbus variants, Controller Area Networks or CAN, ModBus and DeviceNet, among a host of others.

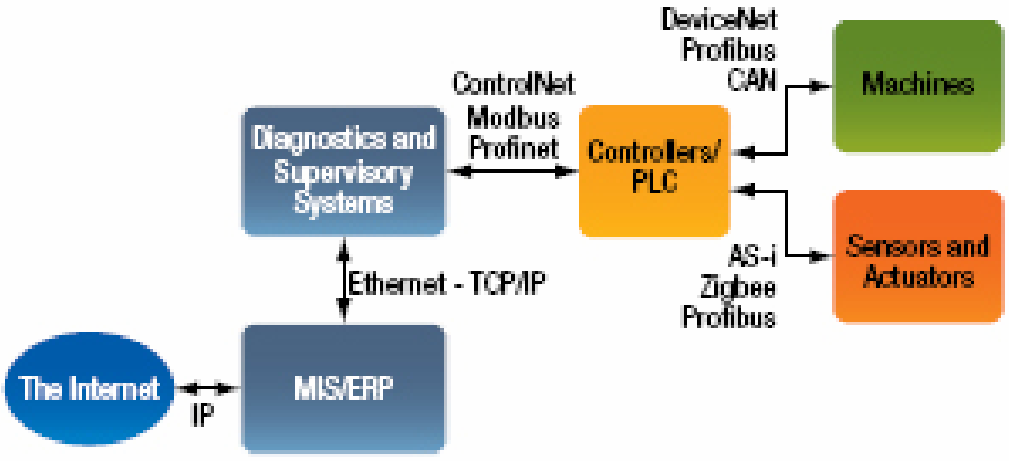

FIGURE 1

Communication between various entities of an industrial environment is often via disparate protocols, each handling a segment of the overall network, and resulting in an obfuscation of the network itself.

IP – the solution for LAN as well as WAN

If there is one network type that can be called near-universal, it is the IP-based network. From within a home or an office to the whole world via the Internet, the TCP/IP protocol suite is omnipresent. The connectionless, packet-switched transport mechanism it provides is ideal for exchange of data; and it also helps that computing platforms everywhere have shed their diversity and adopted one of a few common architectures. Commonality helps reduce cost, and reduced cost in turn enables faster adoption of a technology—but apart from this, the TCP/IP suite also provides for the robustness required by industrial networks. The use of IP by itself does not make a network universal, for there can still be differences in the physical link and the applications that make use of the transport layer. The RS-232 serial link can carry IP packets, and so can Ethernet and optical links, among others (Figure 2). The focus here is on the use of a common wireless link in these networks.

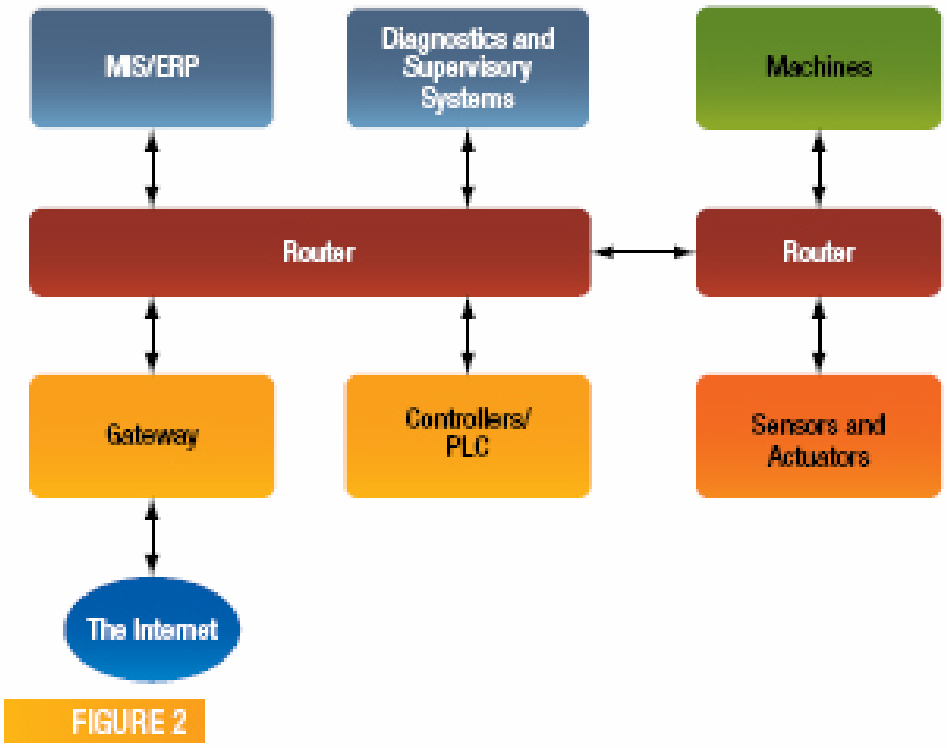

FIGURE 2

The Industrial Network based on IP transport protocol. Note that the physical links may still be different.

The Wireless Network

While networks may still be universally connected without having a common physical layer, there is sound motivation for using wireless as a standard physical medium in industrial or enterprise environments. The number of devices that are connected is increasing rapidly—in fact, there is today an explosion in the deployment of "connected" devices in a whole variety of environments, giving rise to the term "M2M" or machine-to-machine communication, where a plethora of equipment, sensors, actuators, monitors and other devices are connected to each other—something like an Internet of things. The merging of sensors—by definition small, inexpensive, easily deployable and with miniscule energy consumption—into the industrial network has accelerated the adoption of wireless transport mechanisms.

In planning to connect a large number of devices within the premises of an organization, it is evident that extensive cabling would have to be introduced. Routing increasing numbers of cables is expensive and difficult—and the setting up of a new production line or reorganization of a laboratory would see considerable time and resources spent in the installation of cables. In many cases, the cost of the cabling can equal or exceed the cost of the device being networked, and may even prove to be a deterrent to increased automation in the enterprise.

The nature of the devices themselves may dictate the use of wireless as the connecting medium. For example, wireless sensor networks require a few important features to be successfully deployed, and these include the ability to be deployed flexibly in large numbers, work in a "mesh," offer a long battery life, and the ability to transport data in a standard format. These requirements are met only with a wireless transport mechanism, as illustrated in Figure 3.

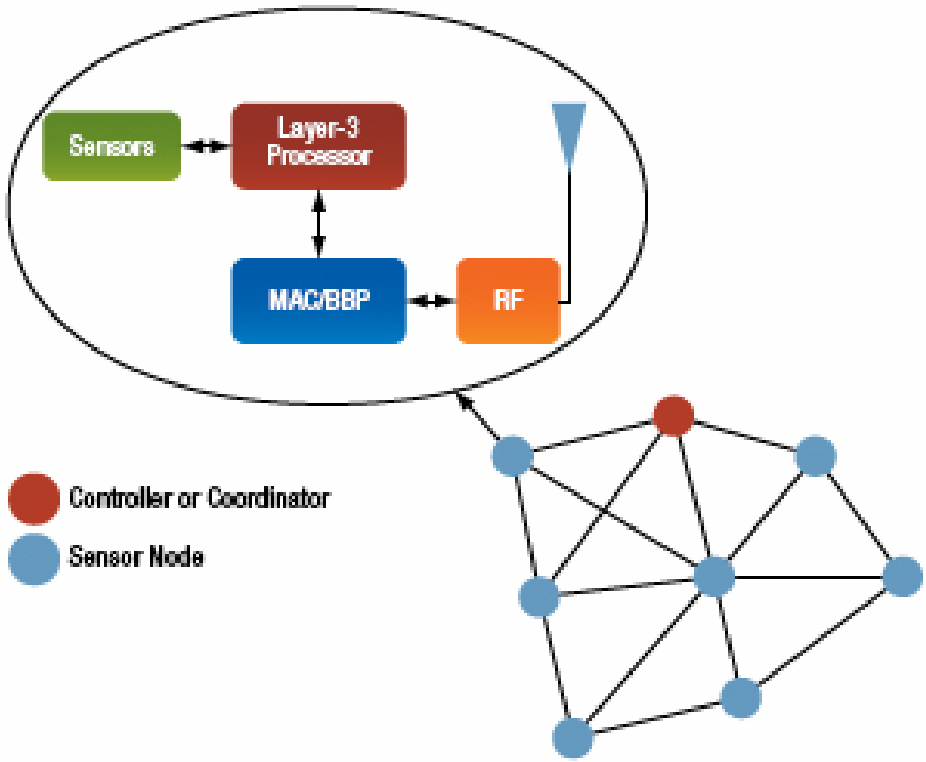

FIGURE 3

A Wireless Sensor Network shown connected in a mesh formation, with the inset showing constituents of a wireless sensor node

The Advantages of WLAN

Three of the most popular wireless standards in this area are the IEEE 802.15.4 "'ZigBee," Bluetooth and IEEE 802.11 Wireless LAN (or WLAN). The focus today is increasingly on WLAN due to its many distinct advantages.

The IEEE 802.11 family of wireless LAN standards was created as descriptions of a system that would be an equivalent of the wired

Ethernet infrastructure. Data throughputs, however, were originally very low—1 Mbit/s to start with—but over the years the standard has been evolving to eventually include, as of early 2009, data rates up to 600 Mbits/s. The 802.11n addition to the standard (presently still in draft form) was proposed as a means of bringing in high user level throughput to WLANs—and therefore has addressed both the PHY and the MAC layers. Apart from addressing high throughput, the WLAN standard has also evolved to support a variety of applications with different needs—,for instance those that require high reliability and timeliness of data transfer, and those that require mobility and roaming.

At the physical layer, 802.11 defines operation in both 2.4 and 5 GHz bands. The 802.11b standard specifies digital sequence spread spectrum (DSSS) and complementary code keying (CCK) modulation schemes with data rates up to 11 Mbits/s. 802.11a, 11g and 11n specify the Orthogonal Frequency Division Modulation (OFDM) modulation scheme along with error correction coding, which has since become the dominant modulation method for high-speed wireless data transport. OFDM is inherently multipath tolerant and provides for relatively simple and inexpensive implementations.

At the MAC layer, 802.11 addresses medium-sharing through a CSMA/CA (Carrier Sense Multiple Access with Collision Avoidance) mechanism. Collisions can and do occur, and cannot be detected as such. A transmitter deduces that a collision has occurred or a packet has not reached the receiver without error by noting the absence of a return acknowledgement. Before attempting to transmit again, the source node would wait for a random back-off interval (Figure 4). This mechanism does not totally eliminate the possibility of lost packets, since the transmitter would only attempt a finite number of retries before abandoning the packet. However, under good link conditions, the probability of lost packets can be very small. The random back-off mechanism creates another problem: packet transfer latency cannot be guaranteed. This affects time-sensitive traffic like control packets for critical process control and voice calls. The IEEE 802.11e standard addresses this issue and provides for prioritizing certain categories of traffic over others.

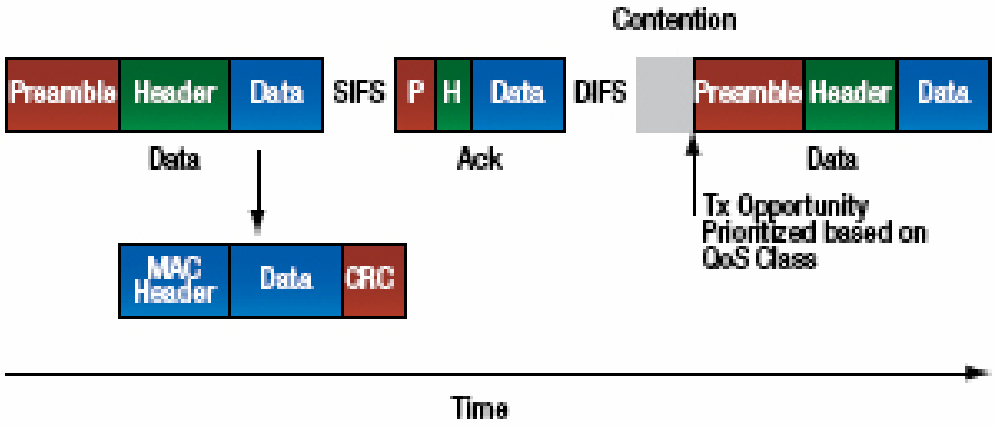

FIGURE 4

802.11 frames are acknowledged by the receiver, following which the medium is thrown open for further traffic. While contending for the medium, devices use a random back-off count with a provision for varying priority based on the QoS class of the data to be transmitted.

One of the main reasons why 802.11 finds itself as the wireless of choice in emerging networks is its ability to scale up and cater to increasing densities in wireless node deployment. The commonly used 802.11g defines “on-air” data rates of up to 54 Mbits/s. It can therefore cater to dozens of nodes each communicating at a few tens or hundreds of Kbits/s, as is common in industrial scenarios. At the same time, a wireless node that has only a limited quantity of data to send at infrequent intervals can put itself into an 802.11 power-save or sleep mode and achieve significant savings in battery drain. However, it must be noted that transmitting small packets is wasteful of bandwidth, since the 802.11 overhead is constant.

WLAN nodes can connect to each other in a standards-defined ad hoc mode, but more commonly, the network would be configured in an infrastructure mode, with all data transfers being routed through an access point (AP). Now, APs are routinely deployed for handling data traffic in most enterprise and factory settings, and this greatly eases the deployment of WLAN-enabled equipment and sensors in those environments. The planning of the network—involving decisions on frequency reuse, coverage of cells and security settings among others—would have been done, paving the way for quick and flexible installation and commissioning of equipment and devices. This is a second significant reason behind the adoption of 802.11 as the choice of wireless technology in industrial networks.

802.11n

The IEEE 802.11n standard primarily provides for high throughput, high efficiency and long-range data connectivity, and includes the use of multiple antennas and transmit-receive chains. Higher throughput is achieved through both PHY and MAC level enhancements. At the PHY level, the standard enables the use of 52 OFDM subcarriers for the carriage of data in every symbol, against the 48 that were used according to 802.11g or 802.11a. A new coding rate of 5/6 is also defined, up ¾ from the rate in the previous versions. The Guard Interval—the period of silence between two symbols—is also halved from 0.8 us to 0.4 us. A symbol is 3.2 us in duration, excluding the guard interval. Occupied bandwidth is increased from 20 MHz to 40 MHz, doubling data rates, with the caveat that the 40 MHz bandwidth mode be used only when an adjacent band is free of traffic. The revolutionary concept of MIMO, or multiple-input multiple-output, is introduced, which makes use of multiple antennas at the transmitter and at the receiver to enable the separation of two or more streams of data sent on the same channel at the same time.

At the MAC level, the 802.11n standard improves efficiency by defining modes where longer packets can be transmitted with a single set of header bytes, or where a return acknowledgement packet contains reception information on a burst of packets previously sent. It also defines other methods like reduced inter-frame spacing (RIFS).

The 802.11n standard primarily addresses the needs of high-traffic nodes, but it also includes a single-stream mode that is intended to provide the benefits of 11n to low-power small form factor devices including sensor nodes. The use of single-stream 802.11n WLAN in these client devices provides the following benefits, while retaining the size, cost and power consumption benefits of legacy devices:

- Higher throughput and lower transmit times—achieved through better efficiency in PHY and MAC.
- Longer range—through use of multiple antennas at the access point.
- Preservation of 802.11n network capacity—the presence of legacy 802.11a/b/g clients forces the 11n nodes to use protection mechanisms and results in overall drop in network capacity. 802.11n helps avoid this.

The Integrated Industrial Network

Devices based on the 802.11 or 802.11n standard can be engineered to provide for the requirements of industrial networks (Figure 5).

WLAN devices can offer the wireless replacement of a serial cable, or construct a flexible and configurable sensor network, or help integrate every piece of equipment into the overall enterprise network. Controllers and supervisory systems built upon standard computing platforms can easily be provided WLAN connectivity through devices that connect via a standard host bus like SDIO, PCI or USB.

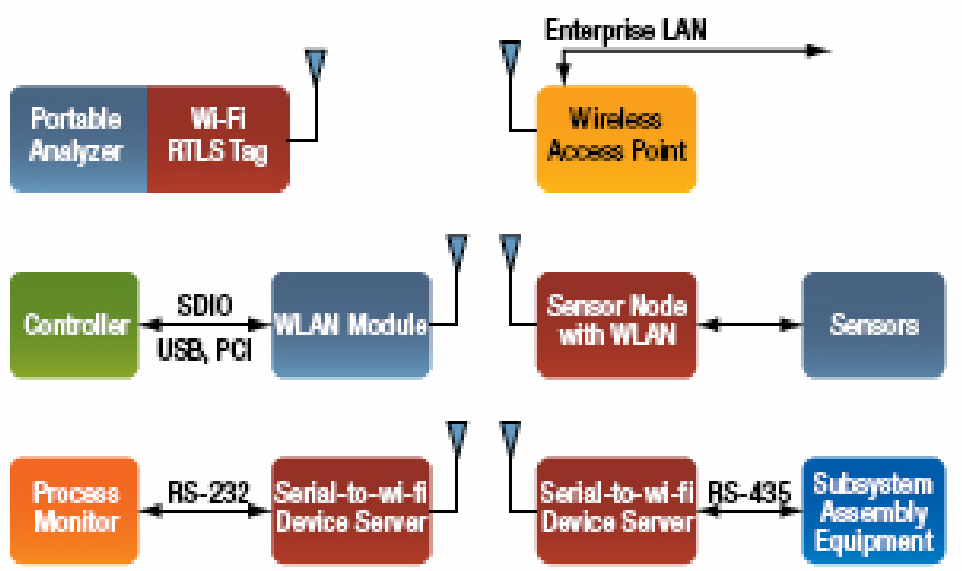

FIGURE 5

A WLAN network in an industrial environment, supported by available 802.11 products.

Monitoring equipment that previously used the serial interface can be provided with external serial-to-Wi-Fi device servers that bring in wireless connectivity without actually requiring the equipment to be redesigned or upgraded. Sensors are available with different forms of wireless connectivity, but those that provide it via WLAN can be chosen for fresh installations. Even RFID tags, which help track assets, can be WLAN-enabled. With the availability of appropriate wireless devices and systems, industrial environments can easily be universally networked based on wireless transport.

Redpine Signals

San Jose, CA.

(408) 748-3385.

[www.redpinesignals.com].