sponsored by
ADVANTECH
Enabling an Intelligent Planet

Special Issue
ISSN 1470-5745

# industrial ethernet book

The Journal of Industrial Network Connectivity


IEEE802.11 technology: wireless for the factory

General industrial wireless for process

Industrial wireless mesh networks

Industrial wireless case studies and examples

www.iebmedia.com ■ www.advantech.com

Photo: Geoff Hodgkinson

# GET CONNECTED . . .



www.iebmedia.com ■ www.advantech.com

## Contents



### Industrial Ethernet Book special issue sponsored by Advantech

This special compilation of wireless technology articles, originally published during the last three years in the Industrial Ethernet Book, has been sponsored by Advantech. Both we and our sponsor hope that it will provide an easy introduction into the use of industrial wireless technologies for both Factory and Process Automation.

### About Advantech

Founded in 1983, Advantech Co.,Ltd.is a leader in providing trusted, innovative products, services, and solutions. Advantech offers comprehensive system integration, hardware, software, customer-centric design services, embedded systems, automation products, and global logistics support. We cooperate closely with our partners to help provide complete solutions for a wide array of applications across a diverse range of industries. The company's mission is to enable an intelligent planet with Automation and Embedded Computing products and solutions that empower the development of smarter working and living.

ADVANTECH
*Enabling an Intelligent Planet*

**Editor:** Al Presher, editor@iebmedia.com
**Supplement Editor:** Frank Ogden, frank.ogden@iebmedia.com
**Publisher:** Leopold Ploner, info@iebmedia.com
Tel.: +49-(0)8192-933-7820 · Fax: +49-(0)8192-933-7829

Published by IEB Media, Bahnhofstrasse 12, 86938 Schondorf am Ammersee, Germany

IEB MEDIA

# Industrial wireless mesh networks can use IEEE802.11 standards

**There are many types mesh network technology which, taken together, can bring wireless connectivity to just about every conceivable industrial and process application. However, there is only one – IEEE802.11n – which truly meets the throughput and reliability demands of industrial automation, plus compatibility with enterprise IT standards says Advantech's Alex Tsai and Kunhong Chen**

ARTWORK: FRANK OGDEN

ANY NETWORK which can function without pre-ordered infrastructure may be described as having mesh topology. To meet this requirement, each node must not only capture and disseminate its own data, but also serve as a relay for other nodes. In other words, it must organise its own collaboration to propagate data across the network.

The expression 'mesh topology' generally suggests MANET – as in Mobile Ad hoc Network – or WSN (as in Wireless Sensor Network), but it may also be applied to freely associated nodes operating with standards-based WLAN technologies such as IEEE802.11n. Of course networks constructed with GSM mobile or wireless sensor comms technology have incredibly useful properties. For instance GSM can provide interactive Scada and telemetry in places and over distances where it simply is not possible to get a wire or a signal in any other way. Likewise, WSN provides short range, low data rate connectivity with a power budget applicable to disposable batteries. However, only IEEE802.11n technology presently offers network-friendly IP routing and the sort of data rate applicable to industrial automation applications in a mesh topology format.

Routing technique largely determines the characteristics of a mesh network. The routing function propagates the message by hopping the data from node to node along a path until the destination is reached. To ensure the availability of a path, a routing network must allow for a continuous connection and reconfiguration around broken or blocked paths through the use of self-healing algorithms. Self-healing capabilities enable routing based networks to operate when one node breaks down or a connection goes bad. As a result, the network is considered to be very reliable since there is often more than one path between a source and a destination in the network.

## Mesh network parameters

The three types of mesh with different uses are detailed here, all of which use the standards-based features available in IEEE802.11n hardware. However, for the greatest flexibility a hybrid configuration is probably the most useful as will be shown.

**Client**. Using one type of radio device e.g., an EKI-6340-3, as a client mesh provides peer-to-peer networks among client devices and the client nodes to perform routing and configuration functions.

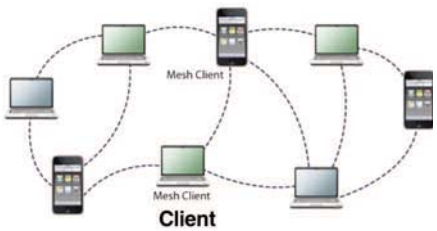

Client

Infrastructure

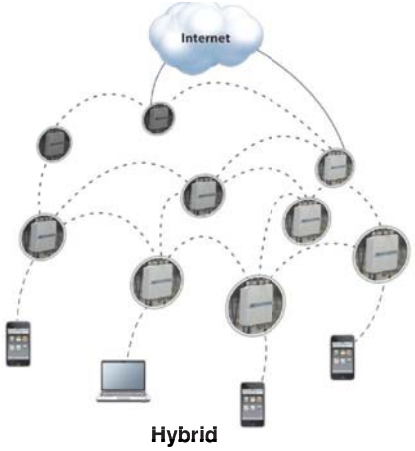

Hybrid

**Infrastructure**. By using routers as the infrastructure for clients that connect to them, the routers form a mesh of self-configuring, self-healing links. With their inherent gateway functionality, mesh routers can also be connected to the Internet, thereby providing a backbone for traditional networks while enabling integration with existing wireless networks.

**Hybrid**. A combination of the previous two architectures, hybrid architecture performs the functions of infrastructure and client and, as such, can access the network through mesh routers as well as communicating with other mesh clients. This flexibility makes it the most suitable configuration for a majority of applications.

In addition to these three designs, there are also three key feature classes that can be applied to each of them.

**Self-healing and Self-forming functions**. Intelligent mesh capabilities with self-healing and route choosing algorithms (self-forming) determine their form by a calculation involving the number of hops and the radio signal quality for each hop. Each wireless connection in a mesh network will have a path score to represent the signal quality between nodes. The path score calculation includes RSSI (Received Signal Strength Indication), noise level and bandwidth flow information in providing the reference for calculation. The number of hops from source to destination is usually a minor consideration in a routing algorithm.

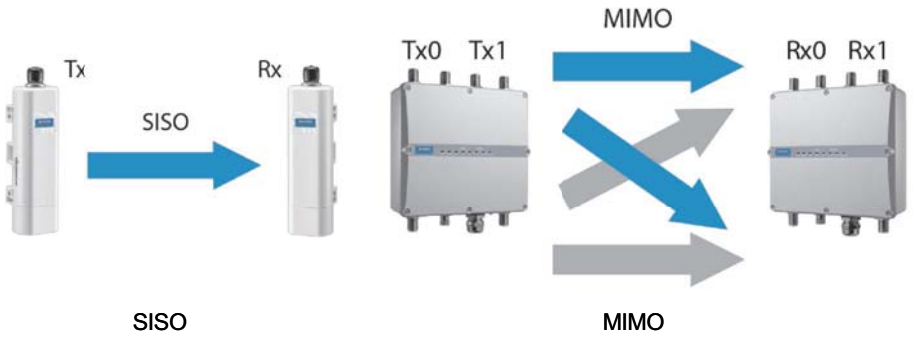


**Multiple hopping functions**. Intelligent wireless mesh systems may have a throughput of higher than 150Mbps at two hops and can still transfer data at 100Mbps after as many as 10 hops. These sorts of throughput rate minimise concern about the carriage of high bandwidth deterministic traffic such as video and motion control data. Furthermore, the self-healing function may reduce maintenance and other network management costs. A typical application might involve monitoring the status of remote oil fields; each derrick would use a node with self-healing and self-forming mesh capability (such as Advantech's EKI-6351), transmitting data to an EKI-6340 AP node set up for direct backhaul to a control centre.

Oil refineries provide an excellent example in the use use of self-healing, self forming and multi hopping algorithms. These sites have many separate areas that need to be continuously monitored and managed, but due to the nature of the business, fires are real possibility and may destroy the hardware monitoring a particular section. Once a link has been broken, routers up or downstream of damaged one might not be able to send their data back to the control centre with some conventional network configurations. Thanks to self-healing and forming and multiple hopping algorithms, the data can still get through.

## IEEE802.11n explained

802.11n includes many enhancements to the earlier a/b/g versions of the protocol. These improvements include an increase in speed, range and reliability.

| IEEE WLAN standard | Over the air estimates | Media Access Control layer, service access point estimates |
|---|---|---|
| IEEE802.11b | 11 Mbps | 5Mbps |
| IEEE802.11g | 54 Mbps | 25 Mbps (when 11b not present) |
| IEEE802.11a | 54Mbps | 25 Mbps |
| IEEE802.11n | 300 Mbps | 150 Mbps |

*Comparison of different 802.11 transfer rates* (source: Intel Labs)

Versions /g and /b operate at 2.4GHz while /a operates at 5GHz. The advantage of 802.11n is that it can operate at both frequencies, being able to support simultaneous 802.11a and 802.11n links at 5GHz, or 802.11b, /g and /n links at 2.4GHz. Where backwards compatibility is not an issue, 802.11n hardware can be configured to run special features. To improve wireless network range, throughput and reliability, 802.11n has three properties over and above older standards: MIMO, packet aggregation and channel bonding. Taken together, these offer a fivefold increase in performance over 802.11a/b/g networks.

**MIMO** (multiple in, multiple out) systems are built using multiple vector antennas at both the transmitter and the receiver, thus providing a mimo system with its desirable qualities. Since it can employ both diversity and multiplexing of simultaneous data streams, it potentially increases system capacity by three or more times. Depending on where mimo signals are processed, a mimo system can be classified into three distinct types: receiver processing only, transmitter processing only, both TX and RX processing systems.

**Receiver processing only**. Receivers employ multiple front ends rather than mimo signal processing. Antennas at the receiver are connected to multiple independent front ends producing separate data streams. These streams are then multiplexed (muxed) into a single data stream providing a much higher data rate than a single antenna system.

Transmitter processing only. In this reverse scenario a single data stream is demuxed and transmitted as multiple substreams. When the signals from different antennas arrive at the receiver, mimo signal processing must be performed using one of three schemes: space-time coding, vertical Bell Lab Layered Space-Time (V-Blast), and maximum likelihood detection (MLD). MLD provides the best performance of the three.

Transmitter and Receiver Processing. Perhaps the best of both worlds, but with better performance? Well partly, but it comes at a price as the hardware is both complicated to configure and administer. The most popular method of performing the two functions is known as singular value decomposition which diagonalizes the mimo channels to form independent channels, to which water filling – data-packing – schemes can be applied to maximise overall system capacity. With enough processing power available within the mesh routers all types of mimo systems can be applied. However, for ease, transmitter-processing-only mimo is applied from mesh routers to mesh clients and receiver-processing-only mimo for links from the routers to the clients.

## Packet aggregation

Packet aggregation increases efficiency by aggregating multiple packets of an application into a single transmission frame, and so enabling them to be sent with a fixed overhead cost of just a single frame. Packet aggregation works best for data applications such as file transfers. For real-time applications such as voice or video transmission, packet aggregation has no effect and it is better to minimise the number of 'packed' packets to reduce latency and eliminate jitter contention.

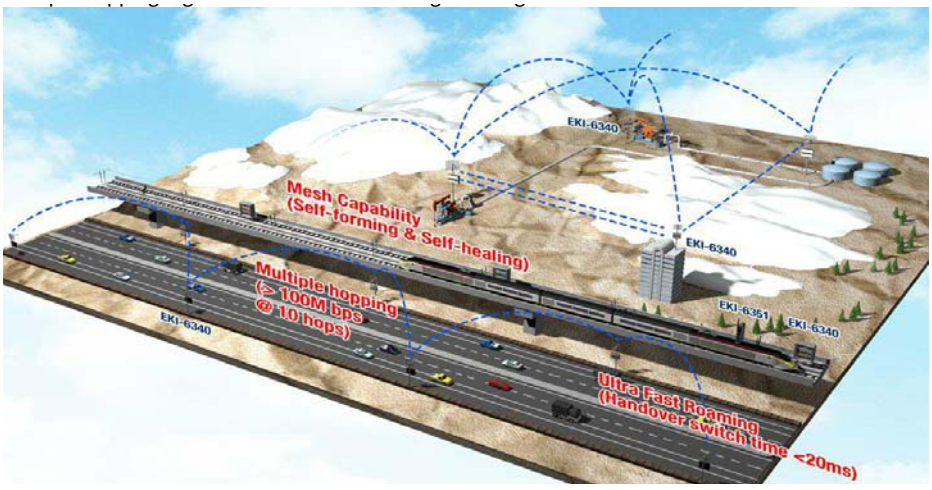


***Fast enough for express handover:*** *This sort of wireless mesh system using an ultra-fast roaming algorithm is suited to handling communication between fast-moving trains and the trackside.*

## Channel bonding

Where 802.11a/g only supports 20MHz spectrum width to carry a maximum of 54Mbps of raw data per channel, 802.11n increases that to 150Mbps per channel. A technique called channel bonding combines two adjacent 20MHz channels into a single 40MHz channel, thereby doubling the throughput to over 300Mbps. Channel bonding works best at 5GHz because there are over 100 channels in the spectrum block, whereas at 2.4GHz only three non-overlapping 20MHz are available for use.

## Fast roaming

A Mesh solution would comprise one mesh gateway (one way connected with switch by Ethernet cable and one way connected with mesh node or mesh AP via radio), a number of entirely wireless-connected mesh nodes (which of course connect with the mesh gateway via radio and and all the others in the mesh). These in turn hop on to mesh APs (and other mesh nodes), hooking one-to-one with other WLAN mesh stations or regular Wi-Fi clients. The mesh nodes are dependent upon the throughput requirements.

Fast roaming is a special feature of a mesh station – it is not a regular Wi-Fi client – and the handover time between two mesh APs can be as fast as 20ms. They achieve this in the following way. The Mesh APs are set to periodically and proactively broadcast information to nearby mesh stations. The mesh stations which are under the coverage of mesh APs can periodically generate a list of path scores. Once a new path score is generated and shown to be better than the current link score, the mesh station will handover the path to another mesh AP without further authentication and association processes. These two steps were performed when the mesh station first joined the mesh network.

This sort of wireless mesh system using an ultra-fast roaming algorithm is suited to handling communication between fast-moving trains and the trackside. By installing an EKI-6351 mesh station inside a train and an EKI-6340 access point along the side of the rail, communication can be maintained without connection loss since the short handover time permits this. Even in an environment where fibre links to the AP cannot be installed, the hardware of this example has three radios, so establishing a wireless backbone through its two spare radios.

***Alex Tsai*** *is Industrial Communication Product Development Manager for Advantech's Industrial Automation Group.*

***Kunhong Chen*** *is Industrial Communication Product Manager for Advantech Industrial Automation Group. He has 10 years working experience with Industrial Wireless and device server technology for Intel, Gemtek and now Advantech.*

# IEEE 802.11n: the next step for industrial wireless LANs

**In less than a decade, wireless LANs have evolved from a niche technology useable only by a few specialised applications to the default media for the last few metres of the network for consumers, enterprise and industry. And WLANs continue to evolve. The latest generation of high speed wireless LAN technology, based on the IEEE 802.11n standard, is now widely used in the world of Office Automation but has yet to make an impact with industrial WLANs. As this primer from Adam Conway suggests, It is just a matter of time.**

THE TECHNOLOGY behind IEEE 802.11n is projected to deliver as much as a six-fold increase in effective bandwidth, as well as increased WLAN reliability compared to existing 802.11g and 802.11a deployments. This promise has led some to consider the wireless LAN as a viable alternative to the wired network.

At a minimum, the advances realised by it will cause many enterprises to reconsider the role of WLANs in their network, as well as the effect of such a deployment on their infrastructure. Before deploying 802.11n, however, organisations will need to understand the answers to some basic questions, including:

- What are the operational differences between 802.11n technologies and existing WLAN elements?
- Is 802.11n backwards-compatible with existing wired and wireless network design?
- What modes can the deployed?

While these questions are simple, the answers to them are not. IEEE 802.11n uses complex technologies more frequently used in the world of radio/broadcast than in networking. Indeed, there is no shortage of material claiming to demystify 802.11n, but which only succeeds in introducing a plethora of new four letter acronyms.

Here, we will look at the basic elements of 802.11n functionality, with an emphasis on how it differs from WLAN technologies presently in use. The primary focus will be on the major methods that 802.11n uses to deliver on the claim of large increases in throughput and reliability.

## Performance and reliability

802.11n touts major improvements in both performance and reliability, yet also purports to have backward compatibility with 802.11a and 802.11b/g equipment. The backward compatibility, higher performance and increased reliability come about through the action and interaction of two key technologies: Multiple In/Multiple Out (MIMO) transmit/receive capabilities and Channel Bonding.

Incremental improvements are also seen by combining a myriad of additional technologies, but for the sake of simplicity, we will consider only the primary changes.

## Multiple In, Multiple Out (MIMO)

MIMO is the biggest innovation that comes with 802.11n. Though there are different kinds of MIMO techniques, we will limit our discussion to the most useful and prevalent form in building enterprise WLANs, often called 'spatial diversity MIMO' or 'multipath MIMO'.

**Multiple In.** When you use only one antenna on the transmitter and one receiver in an indoor environment, you are subject to multipath interference. Multipath interference happens when a number of packets are encoded and sent out over the air. The waveform will interact with anything it encounters on its way from transmitter to receiver. Some of these things, like a metal fire door, will reflect the signal; some things, like a working microwave, will interfere with it; some things, like organic material such as plants and people, will absorb it. The result is that the receiver can end up with multiple copies of the original signal.

This is similar to how a single sound produced in a canyon can result in an echo, sounding to the listener as though the sound is produced many times over, out of phase with the original. This echo effect makes it difficult to sort out the original message, since signals received in different phases can combine or even completely cancel one another.

We have all encountered this effect when listening to the radio in the car. The signal might be just fine until you come to a particular place such as a stop light where, suddenly, the signal seems to disappear. If you move a bit, however, the signal comes back.

What is really happening when the radio station appears to go away is that multipath interference is creating a null – the signals received are offset from each other and, when combined, net a zero signal. When you move, you've shifted what the receiver 'hears' and the signal appears to come back. With a complex signal, it can be virtually impossible to determine where one message ends and another begins.

One way that WLAN providers have worked around multipath is to provide a diverse set of antennas. Antenna diversity, however, is not MIMO. It is simply that only one of the set of antennas is actually transmitting or receiving – the WLAN is just able to select the antenna set with the best signal-to-noise ratio.

Multipath has traditionally been the enemy of WLANs, because the echo-like effect typically serves to detract from the original signal. When using MIMO and its multiple receiving antennas, however, the effects of multipath become additive – that is, multiple messages can be received by multiple antennas, and combined.

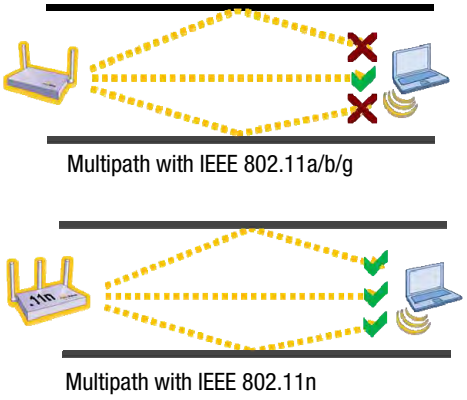


***Fig. 1.*** *Multipath use for 802.11a/b/g vs. 802.11n*

When using MIMO, we still get multipath as always, but this time we can sort out a message more easily and actually use the multipath reflections to our advantage to gain significant signal strength and thus improve reliability (Fig. 1). What does this mean in practice? Reliability translates to a greater coverage area for a given data rate, or to higher data rates for a given coverage area. That translates to more bandwidth per user…

**Multiple Out.** MIMO allows for multiple (from 2 to 4) transmitting and receiving antennas that operate simultaneously. Using advanced signal processing at both the access points and clients, MIMO transmitters can multiplex a message over separate transmitting antennas. The receivers digitally process the signal to identify separate bit streams – commonly known as spatial streams – and reassemble them. This multiplexing dramatically increases the effective bandwidth. Thus the two biggest improvements MIMO brings are:

- The ability to more easily sort out multipath echoes which increases reliability (and as shown, more reliability equates to more bandwidth-per-user)
- The ability to multiplex different data streams across multiple transmitters which increases effective bandwidth.

## MIMO APs with legacy clients

MIMO can assist operational reliability for legacy 802.11b/g and 802.11a clients. This occurs because with 802.11a/b/g hardware, APs cope with multipath interference by scaling back the data rate. This means that clients that could get 54Mbps throughput in an interference-free environment might have to drop to 48 or 36Mbps at a short distance from the AP in the presence of multipath.

Even if MIMO is used only in the access points, the technology still delivers up to 30% performance enhancement over conventional 802.11a/b/g networks simply because MIMO receive antenna technology handles multipath in a much better way. This efficiency means that clients that would normally have to drop from 54Mbps data rates to 48 or 36Mbps at a short distance from the AP can now remain associated at 54Mbps.

## MIMO hardware options

The 802.11n standard allows for several different configurations of transmitters and receivers, from two to four transmitters and from one to four receivers. MIMO systems are described by quoting the number of transmitters 'by' the number of receivers. Thus a '2x1' system has two transmitters and one receiver.

Adding transmitters or receivers to the system will increase performance, but only to a point. For example, it is generally accepted that the benefits are large for each step from 2x1 to 2x2 and from 2x3 to 3x3, but beyond that the value is diminished for the current generation of 802.11n.

Additionally it is often recommended that access points are optimised in a 3x3 configuration whereas clients function best in a 2x3 configuration. The AP can make use of the additional transmitter because it handles multiple clients.

**Channel Bonding** is a technique where two adjacent contiguous 20MHz channels are combined into a wider 40MHz channel. In fact, the bandwidth on both edges of a 20MHz channel are typically not used at 100%, in order to prevent any channel overlap. Channel bonding allows the use of both 20MHz channels as well as this gap between the channels, resulting in slightly more than double the bandwidth (Figures 2 and 3).

For example, the highest data rate for 802.11a or 802.11g is 54Mbps for a single transmitter on a 20MHz channel. In 802.11n, a 20MHz-wide channel was made more efficient using various incremental improvements to increase the maximum data rate of a single channel from 54 to 65Mbps.

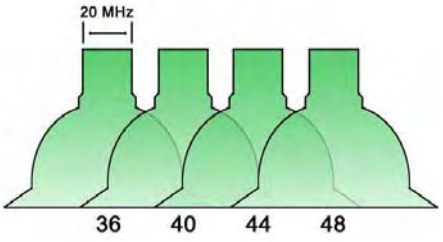


***Fig. 2.*** *Four 5GHz 20MHz channels*

With the addition of channel bonding and better spectral efficiency, a 40MHz bonded channel on a single transmitter gets you slightly more than double the 54Mbps data rate, or 135Mbps.

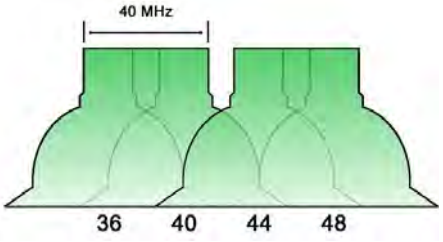


***Fig. 3.*** *Four 5GHz 20MHz channels bonded to form two 40MHz channels*

The drawback to channel bonding, as we will demonstrate, is that it can really only be implemented in the 5GHz band.

**Channel usage.** IEEE 802.11n wireless networks can operate in either the 2.4GHz (802.11b/g) or 5GHz (802.11a) bands. If you have an access point with two 802.11n radios, it can operate in both 2.4GHz and 5GHz bands simultaneously. The access point can be configured to use the same channels as 802.11b/g and 802.11a, and thus remain backward compatible with clients still running 802.11b/g or 802.11a.

When building a wireless LAN for the enterprise, it is important that no two APs operate on the same channel when they are in close proximity. Doing so causes co-channel interference and is similar to having two radio stations transmit on the same frequency, in that the receiver ends up getting mostly garbage.

To avoid this, the APs need to change the channels they use so as to not interfere with each other. That can get tricky if there aren't enough channels to choose from. 802.11 b/g in the 2.4GHz range has a three to one reuse pattern for useable channels. This three to one pattern is the very minimum number that can be used to build a non-interfering network.

So what happens with 802.11n when you introduce things such as increased range and channel bonding? It becomes clear that the 5GHz band is the only choice when using 802.11n with channel bonding, as it easily allows enough 20MHz non interfering channels to get to a seven to one reuse pattern, and a three to one reuse pattern with 40MHz channels. This allows plenty of spectrum for building out a WLAN without co channel interference (Figures 4 and 5).

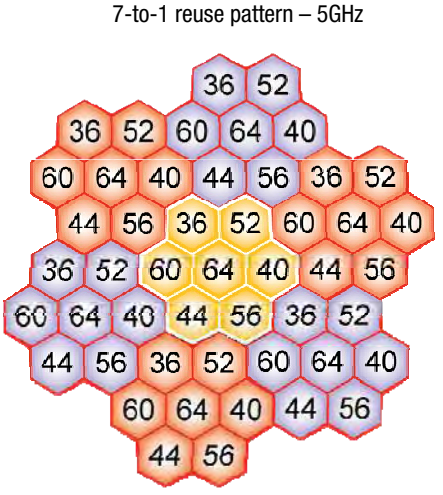


***Fig. 4.*** *Reuse pattern for 5GHz using 20MHz channels*

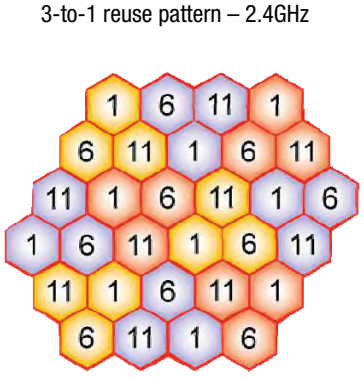


***Fig. 5.*** *Reuse pattern for 2.4GHz using 20MHz channels*

## 802.11n maximum data rate

Several factors determine the maximum performance that can be achieved with 802.11n. Spatial streams and channel bonding that were mentioned earlier provide the biggest benefits, but there are several other items that can also increase performance.

**Guard interval (GI).** A guard interval is a set amount of time between transmissions, designed to ensure that distinct transmissions do not interfere with one another. The purpose of the guard interval is to introduce immunity to propagation delays, echoes and reflections. The shorter the guard interval, the more efficiency there is in the channel usage – but a shorter guard interval also increases the risk of interference.

A (short) guard interval of 400ns will work in most office environments since distances between points of reflection, as well as between clients, are short. Most reflections will be received quickly – usually within 50-100ns. The need for a (long) guard interval of 800ns

## What is Spatial Multiplexing?

SM provides a wireless system an opportunity to increase throughput without the use of additional spectral bandwidth. A simplified view of SM can be thought of as transmitting N unique data streams using highly directional antennas aimed at N different receive antennas. Each receiver detects a unique data stream resulting in an N fold increase in throughput. With sophisticated signal processing techniques it is possible to achieve a similar N-fold improvement using N (or more) omni-directional antennas.

For example, in a standard WLAN 802.11g access point a second data stream can be transmitted from a second antenna. It is possible for a laptop with two receive chains to decipher the two (different) data streams, effectively doubling the throughput. A third antenna (at both the transmitter and receiver) can triple the data rate and so on.

There is an upper bound, however. Analogous to solving an algebra problem of N unknowns from M independent equations (M N), in SM systems the maximum number (N) of data streams is restricted to the number (M) of independent (uncorrelated) signals received. Here, uncorrelated signals are radio signals that took different physical paths (multipath) from a transmit antenna to the receive antennas. In other words, throughput improvement is limited by the number unique signal paths. Furthermore, the multipath conditions are completely dependent on the environment where the WLAN is deployed.

**Winston Sun** Atheros Communications

becomes more important as areas become larger, such as in warehouses and in outdoor environments, as reflections and echoes become more likely to continue after the short guard interval would be over.

The guard interval that was set in IEEE 802.11 specifications prior to 802.11n was longer than needed in many environments. A shorter guard interval was added as an option in the 802.11n specification to allow for higher data rates where the long guard interval is not required.

**Frame Aggregation.** Data over wired and wireless networks are sent as a stream of packets known as data frames. Frame aggregation takes these packets and combines them into fewer, larger packets allowing an increase in overall performance. This was added to the 802.11n specification to allow for an additional increase in performance. Frame aggregation is a feature that only 802.11n clients can take advantage of since legacy clients will not be able to understand the new format of the larger packets.

**Reduced Inter Frame Spacing (RIFS).** The standard spacing between IEEE 802.11 packets is known as the Short Inter Frame Space (SIFS). IEEE 802.11n adds a smaller spacing between the packets when a larger spacing isn't required. This reduces the overhead and slightly increases throughput. This was added to the 802.11n specification to increase performance where possible. RIFS is a feature that only 802.11n clients can take advantage of since legacy clients will not be able to receive packets with the shorter spacing.

The maximum possible data rates when using 802.11n with and without channel bonding, using one through four theoretical spatial streams, with both long and short guard intervals are listed in Table 1.

Of particular note is that today's radio chipsets generally do not support more than two spatial streams, nor do they support a true 'green-field' configuration. Also the data rate describes the PHY-level encoding rate over the air which has significant overhead. The actual wired bandwidth throughput is roughly 50% of the data rate.

One simple conclusion is that we will see future generations of chipsets capable of even higher bandwidths than exist today.

| | 1 Spatial Stream | 2 Spatial Streams | 3 Spatial Streams | 4 Spatial Streams |
|---|---|---|---|---|
| 20 MHz Channel | 65 Mbps | 130 Mbps | 195 Mbps | 260 Mbps |
| 40 MHz Channel | 135 Mbps | 270 Mbps | 405 Mbps | 540 Mbps |

Long (800ns) Guard Interval

| | 1 Spatial Stream | 2 Spatial Streams | 3 Spatial Streams | 4 Spatial Streams |
|---|---|---|---|---|
| 20 MHz Channel | 72 Mbps | 144 Mbps | 217 Mbps | 289 Mbps |
| 40 MHz Channel | 150 Mbps | 300 Mbps | 450 Mbps | 600 Mbps |

Short (400ns) Guard Interval

***Table 1.*** *The effects of channel bonding and Guard Interval set against the maximum attainable data rate using IEEE 802.11n*

## Backwards compatibility

An 802.11n AP is backwards-compatible with legacy IEEE 802.11b/g (2.4GHz) or 802.11a (5GHz) clients. Please note, however, that there is a performance trade-off in this configuration, similar to that observed with an 802.11g AP supporting 802.11b clients.

- Though legacy clients will benefit somewhat from the extended range that an 802.11n AP can offer, they are not capable of the higher data rates.
- An 802.11g client takes longer to send a given amount of data when compared to an 802.11n client, therefore the 802.11g client will consume more 'air time.' This has the impact of limiting the airtime available to 802.11n clients which, in a congested state, will reduce 802.11n performance.

**802.11n compatibility modes.** An 802.11n access point can be configured to operate in three modes; Legacy, Mixed and Green-field Modes.

**Legacy mode.** In this mode, the access point is configured to operate just like an 802.11a or 802.11g device. No benefits of 802.11n such as MIMO or channel bonding are used. This mode could be used when an enterprise buys a new 802.11n access point and, although some clients may have 802.11n capabilities, the company chooses consistency among user experience over maximum possible speed. In Legacy mode, 802.11n capabilities exist, but are not turned on.

**Mixed mode.** This mode will be the most popular of the possible deployments. Here, the access point is configured to operate as an 802.11n AP while also communicating with 802.11 a/b/g stations. When configured for mixed mode, the 802.11n access point must provide 'protection' for the older 802.11 devices, in much the same way that 802.11g access points would communicate with 802.11b clients. Thus the presence of an 802.11a/g client reduces the overall bandwidth capacity of the 802.11n access point, in part because of the lower data rates at which the a/g clients communicate.

**Green-field mode.** This mode is described in the standard and assumes that only 802.11n stations operate on the network, with no protection mechanisms for 802.11 a/b/g necessary. Most current 802.11n chipsets do not support this mode, as the incremental performance benefit is small and it is expected that mixed mode will be prevalent for the near future.

## IEEE 802.11n in the future

802.11n provides significant improvements in WLAN performance and reliability for 802.11n clients, as well as performance and reliability improvements to existing legacy clients. MIMO takes the challenge of multipath interference and uses it to increase performance and reliability of the overall network. The addition of channel bonding can realise significant benefits in performance as well. The combination of these innovative features allows immediate advantages to be seen when migrating to an 802.11n wireless network, even with legacy clients. The benefits only increase as more clients become 802.11n-capable over time. The increase of performance, throughput, and reliability of 802.11n allows the WLAN to become a viable alternative/companion to the wired network for high bandwidth and robust applications.

*From the paper* The Network Inpact of 802.11n *by* ***Adam Conway****, VP Product Management, Aerohive Networks*

www.aerohive.com

First published in the ***industrial ethernet book*** July 2010

# IEEE802.11n WLAN – lessons learned from 1000 installations

**The introduction of the wireless IEEE802.11n standard is an opportunity to provide Ethernet capability without costly and disruptive installation of Cat5 cable. However, 802.11n offers its own challenges – by following the guidelines in this article, taken from a Xirrus white paper, 802.11n networks can achieve similar performance to 100baseT networks, and even Gigabit networks. While the emphasis of the paper is about enterprise usage, the newest LAN standard brings significant benefits to industrial applications. Although quoted wireless client numbers are unlikely to be met in factory and plant use, many of the thoughts offered here may be usefully applied for a significant operational advantage.**

SINCE EARLY 2008, wireless infrastructure company Xirrus has deployed over 1000 IEEE802.11n Wi-Fi networks for universities, schools, enterprises, hospitals, convention centres and others. During the design and implementation of these networks, the company gained much experience and knowledge in 802.11n technology and what it takes to successfully install and operate high performance, resilient 802.11n networks. A white paper outlines key lessons learned from these 802.11n deployments, and this article summarises it.

Proper planning is crucial with any network, but it is especially important with 802.11n as this has more flexibility and configuration options than legacy Wi-Fi. Network designs can take into account end-user needs and environmental issues to optimise performance and robustness. A site survey is important because it lets the network administrator know exactly where equipment needs to be placed for best performance and RF propagation. Real equipment should be used to determine the best placement.

802.11n networks use Multiple Input Multiple Output (MIMO) – traffic is carried on two or three lower powered streams, increasing network throughput while ending reflection nulls through active phase management of the carriers. Because of MIMO, 802.11n RF propagation can be significantly different from 802.11abg networks. A site survey should always be carried out to test the RF characteristics of the environment before deployment. To realise the benefits of 802.11n, the network should operate in the 5GHz band. During the site survey, readings should be taken for both 2.4GHz and 5GHz to ensure that both bands can be seen from all areas to be covered. In most environments, 2.4GHz will propagate further than 5GHz – changing equipment location or adding additional equipment may be necessary to provide full 5GHz band coverage (Fig. 1).

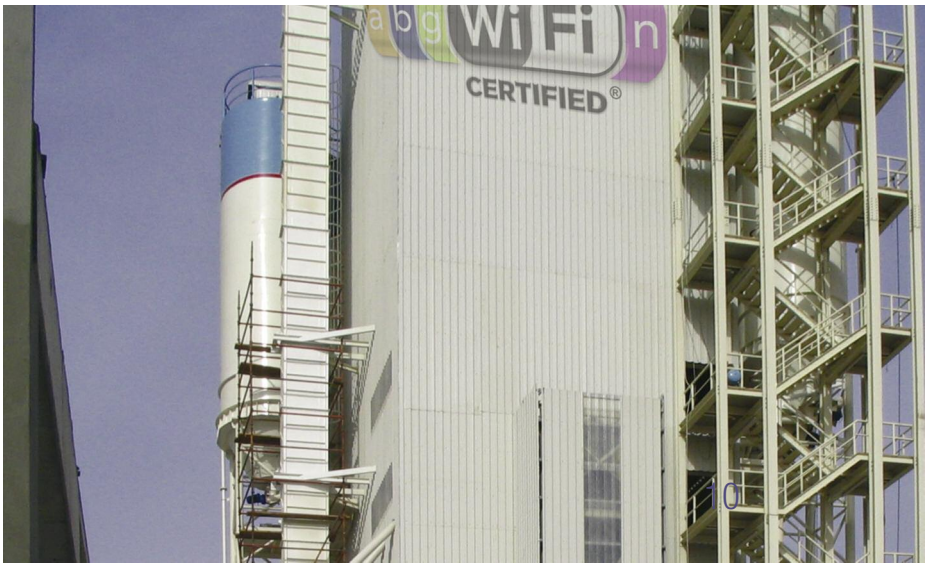


ARTWORK: FRANK OGDEN

The site survey should establish that multiple radios operate at an RSSI level of -72dBm or more from every area to be covered by the network. For a resilient connection, there should be multiple radios from which a station can choose, should one of them be too heavily used or go down.

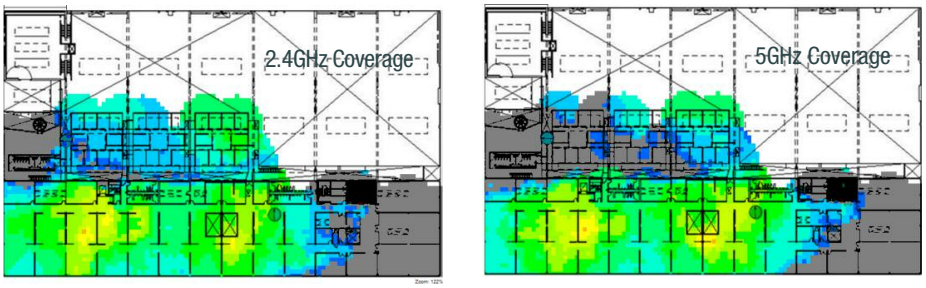


***Fig. 1 . Note that 2.4GHz (left) and 5GHz coverage differs significantly as shown in this example.*** *Lesson Learned: When doing site surveys, look at both 5GHz and 2.4GHz bands. 802.11n can operate in both and to fully realise its benefits, both bands should be supported throughout the entire network.*

## Setting cell size

Setting the appropriate cell size is important when planning a wireless network. Choosing a cell size as large as possible for best equipment utilization may seem to be the way forward. However, if a single radio covers too large an area, overall performance will drop because too many users are forced to use a single radio, and stations at the edge of the cell will have slower connection speeds. One or two slow stations will affect the overall throughput of the network. If the goal is to do wire-switch replacement, care should be given to set the cell size for optimum performance.

Many network administrators assume they can use fewer 802.11n access points (AP) than with 802.11abg. However, if there are legacy 802.11abg stations in the network, APs will need to give adequate coverage for 802.11n and 802.11abg stations. If the existing 802.11abg network was providing adequate coverage, it can be assumed that the new 802.11n APs can be placed in the same locations providing similar coverage – this should be confirmed by the site survey. (Fig. 2).

The recommended security configuration is 802.1x for authentication and WPA2/AES for encryption. Other authentication options are

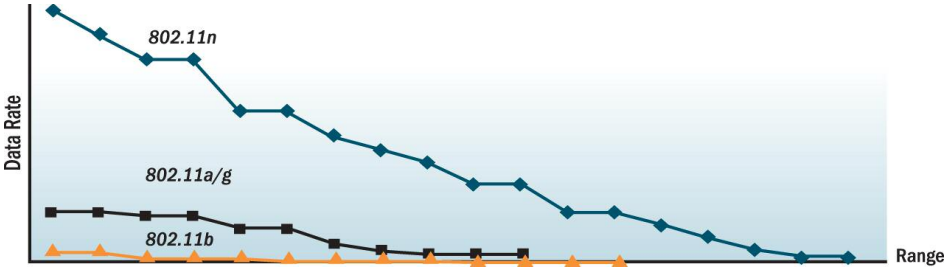


***Fig. 2. 802.11 relative rate and range comparison:*** *Many network administrators assume that they can use fewer IEEE802.11n APs than with 802.11abg. However, most networks will still see many 802.11abg stations present, so similar numbers of 802.11 APs will be needed for adequate 802.11n and 802.11abg stations coverage.*

none, WEP, and TKIP. Using no encryption is not a realistic option for today's networks. WEP and TKIP provide some security, but are not nearly as robust as AES. Besides inferior security, WEP and TKIP limit data rates to 54Mbps, so the increased speed of 802.11n cannot be realised.

## Greater encryption power

While AES is more secure, it does create an increase in required encryption power on the core network. Since 802.11n supports approximately six times the data throughput of 802.11ag, six times as much traffic may need to be encrypted. Many back-end controllers cannot handle this increase in encryption. In fact, some vendors state an 80% drop in throughput capability when using encryption.

A wireless network should have a continuously operating network threat sensor. If the network is migrated to 802.11n, the monitoring tool needs to be migrated as well.

Using 100baseT wired switches in the network is inefficient with 802.11n. Two-radio 802.11n APs can generate up to 250Mbps of data traffic, oversubscribing a 100baseT connection. This effect is more marked when four- and eight-radio arrays are used. Gigabit Ethernet switch ports are required to support the data traffic from 802.11n APs and arrays.

Core switches will also be affected by 802.11n, particularly when using controller-based wireless networks. With a controller, all wireless traffic must go from the APs to the controller for processing and then back to the APs ('tromboning'). To minimise tromboning, traffic should be processed at the edge where possible (Fig. 3).

802.11n equipment typically requires too much power to be used with 802.3af provided over a wired Ethernet connection (PoE), so special power injectors are required. Some vendors have a low power mode that allows their devices to be run with standard 802.3af PoE ports; however many of the features of 802.11n must be turned off and/or performance reduced. It is better to plan for full power and all of the functionality by using higher-powered injectors.

## RF design considerations

An 802.11n network will always perform better than an 802.11abg network, but there are design decisions that can be made to increase the performance even more.

802.11n is the first Wi-Fi standard that can operate in both the 2.4GHz and the 5GHz bands. 5GHz has many advantages over 2.4GHz and every effort should be made to move the network to primarily 5GHz. Most enterprise class equipment can support 5GHz today – however there are some handheld devices (e.g. Blackberrys and iPhones) and netbooks that are primarily 2.4GHz, and so it is prudent to leave some 2.4GHz APs in place to support handheld Wi-Fi devices. In buying netbooks, care should be taken to ensure that both 5GHz and 2.4GHz operation is supported (Fig. 4 *over page*).

Some APs are fixed with one radio operating at 5GHz and one radio at 2.4GHz. As the wireless network transitions to predominately 5GHz with this type of equipment, the 2.4GHz radios will go unused. Better alternatives are APs and arrays that have software-selectable frequencies which may be used with both 5GHz and 2.4GHz, convertible to 5GHz as a majority of the stations become 5GHz-capable.

Network administrators often assume that they can put more users in the higher throughput 802.11n network than they could

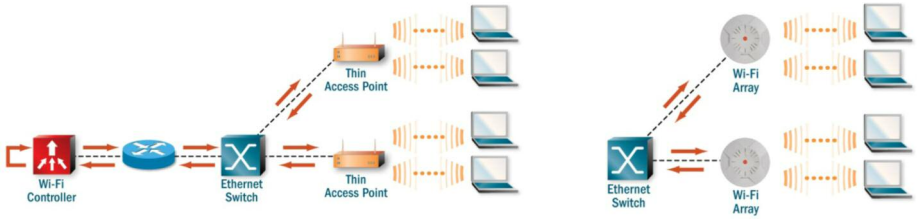


***Fig. 3. Wi-Fi architecture comparison:*** *Minimise traffic load from 802.11n networks on the core by processing as much traffic as possible at the edge. However, the core must be capable of dealing with higher 802.11n network traffic volumes*

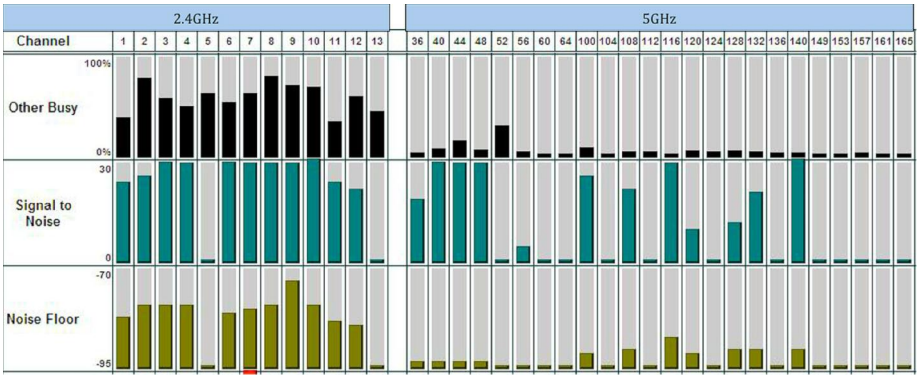


***Fig. 4. Spectrum analysis of 2.4GHz and 5GHz:*** *The 5GHz band contains less RFI noise. Any device using under these conditions will experience a much cleaner signal and provide a more reliable connection compared with 2.4GHz operation.*

with 802.11abg. However, as more users are added to the network, more of the bandwidth is taken up by network overhead and less by user data traffic. In general, there is a negative correlation between number of users and overall throughput per radio (Fig. 5).

Channel bonding takes two 20MHz channels and combines them together to form a single

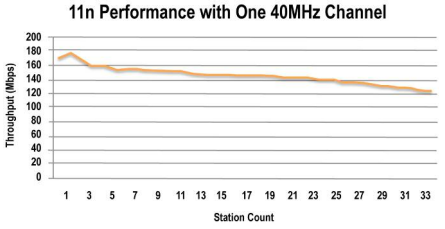


***Fig. 5. Single radio performance with different station counts:*** *As the number of users per radio increases, the total throughput starts to decrease*

40Mhz channel, effectively doubling the channel bandwidth. This effectively cuts the number of available channels to use in half and so is most practically implemented in the 5GHz band where many more channels are available. Bonding is not recommended in the 2.4GHz band as there are only three non-overlapping channels available, so bonding is only possible on one pair of channels

The doubling of capacity on a given radio with channel bonding is evident with low numbers of stations. However, as the number of stations operating in the network increases, bonding is less beneficial because the reduced number of channels reduces the overall bandwidth available. With channel bonding

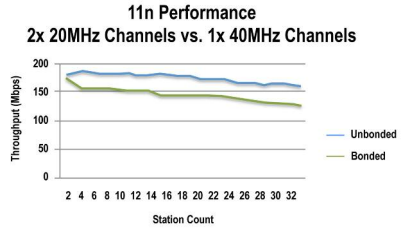


***Fig. 6. Single bonded radio vs two non-bonded radios:*** *If the user density is high on the network and/or 5GHz channel availability is an issue, using non-bonded channels should be considered*

turned off, the bandwidth can be more easily distributed.

If stations in the network do not handle channel bonding well, they will either not connect or will revert to 2.4GHz. In such cases, bonding should not be used in the design (Fig. 6).

### 5GHz Wi-Fi vs 2.4GHz: the benefits

5GHz Wi-Fi outperforms 2.4GHz in a number of ways:

Congestion – Many non-Wi-Fi devices operate in the 2.4GHz range (microwaves, Bluetooth), causing frame loss, retransmissions and reducing bandwidth for Wi-Fi devices. The 5GHz range contains the least amount of noise. With less interference, any device operating in 5GHz will have a cleaner signal and a better user experience than one operating in the 2.4GHz range.

More channels – In the 2.4GHz frequency, there are three non-overlapping channels. In the 5GHz band, there are 24 non-overlapping channels. By moving to 5GHz, bandwidth increases eight-fold.

Higher throughput – 5GHz stations typically perform better than 2.4GHz stations. Experience shows that 5GHz stations will typically have a 3:2 performance advantage over 2.4GHz stations in both legacy Wi-Fi and in 802.11n.

802.11n increases the overall network performance, but the slowest connected station will always limit a wireless network. When a legacy station transmits, no other stations can transmit, and management traffic must be sent at slower rates so that the legacy station can process it. The slower the connection rate, the slower the station can send its traffic and the longer it takes for the wireless network to become available for any other stations to send traffic. 802.11b APs cause the highest drop in performance – up to 75% – while 802.11g APs can reduce performance by almost 50%. Rather than taking a performance hit by supporting 802.11b, 802.11n networks may be designed to disallow 802.11b stations from connecting (Fig. 7).

## Station considerations

To support 802.11n, compatible network interface controllers (NIC) are needed. Fortunately many laptops today already ship with 802.11n NICs, and 802.11n adapters are available for upgrading older laptops. Before buying new 802.11n equipment, it is worthwhile testing to see which stations or network adapters give the best performance. If 802.11n stations already exist, taking the time to qualify their behaviour will help optimise station and network settings for the best 802.11n network possible.

Performance can vary greatly between network adapters. In general, integrated adapters typically perform better than external adapters, but it is always good practise to confirm what works best for a specific network. Encryption can also affect performance. Most network adapters do encryption at the chip level so performance is not affected, but with some adapters the encryption is carried out in software and can be much slower than when using open networks.

Some stations will not associate to particular channels in the 5GHz band. If a Wi-Fi network is using these particular channels, the station may not associate at all or will associate to another array using different channels (even if that array is further away). Support for the mid-band channels in 5GHz (channels 100-140) is inconsistent between different NIC brands, but some more commonly used channels can cause issues as well with certain network interface controllers (NIC).

Many NICs favour the 2.4GHz band over 5GHz, connecting to 2.4GHz radios even if 5GHz radios are available with the same or stronger signal strength. With many adapters, upgrading to the latest drivers will fix this, or the station could be set to only use 5GHz, so improving performance for those stations and freeing bandwidth for stations that are not capable of running at 5GHz. Another option to maximise 5GHz utilisation is to turn down the power on some 2.4GHz radios (assuming sufficient 5GHz coverage, which a well-planned 802.11n network will have) – stations will be more likely to associate to the higher-powered 5GHz radios.

Many stations stick to a particular array, lowering performance as they get further away from the array – even if there is an array nearby with a stronger signal. By increasing the level of roaming aggressiveness in the NIC driver, stations will more readily change to an

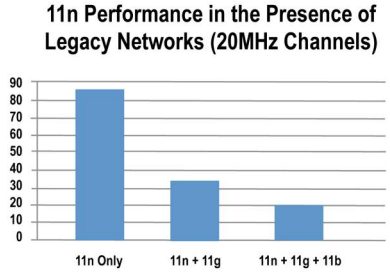


***Fig. 7. Effect of legacy stations on 802.11n performance:*** *This drops by almost 50% if an 802.11g station is causing congestion in the network. The performance is cut by almost 75% if there is an 802.11b station causing congestion*

## Lessons learned from IEEE802.11n Enterprise deployments

- Because of MIMO, 802.11n RF propagation can significantly differ from traditional 802.11abg networks. An active site survey should always be done to test the RF characteristics of the environment prior to deployment.
- Survey for at least two radios at all locations. Be sure at least one 2.4GHz and one 5GHz radio are visible from anywhere, and preferably, two 5GHz radios.
- Use multi-radio arrays/APs so that large numbers of users do not have to share a single radio. Set the cell size so that a minimal number of users are at the fringe of the cell.
- Do not create large cell size for an 802.11n design if you want to achieve the full performance 11n has to offer.
- Don't assume a lower network device count when designing 802.11n networks. Design for high performance and plan support for legacy stations.
- Use WPA2/AES encryption for 802.11n deployments; anything else is a compromise on security and performance.
- Check the encryption performance of your wireless controllers and/or APs; many cannot handle the load and will become oversubscribed.
- Legacy 802.11abg monitoring tools may not catch all of the threats for an 802.11n network. If the network is being migrated to 802.11n, the monitoring tools needs to be migrated as well.
- 802.11n arrays/AP must be plugged into Gigabit Ethernet switch ports to take full advantage of the potental throughput improvement availble with 802.11n.
- Minimize traffic load from 802.11n networks on the core by processing as much traffic as possible at the edge.
- Be sure the core is capable of seeing an increase in data traffic from an 802.11n network.
- A 5GHz Wi Fi network will be more resilient and higher performing than a 2.4GHz Wi Fi network will be. Every effort should be made to transition to 5GHz when deploying 802.11n.
- Buy networking equipment that allows the radio operating frequency to be selected in software rather than fixed at a certain frequency.
- When moving to 802.11n, don't expect to support more users per radio. The goal of deploying 11n should be to get more bandwidth per user. Because wireless is a shared medium, total performance of a radio will drop as more users are added.
- In high density, high performance environments, using non-bonded 20MHz channels may improve overall performance.
- Some clients may not operate properly with bonded channels. The network should be designed to handle these clients by turning bonding off on certain radios, or entirely.
- 802.11b support should be turned off to improve 802.11n network performance if at all possible. Small numbers of 802.11b-only stations can be inexpensively upgraded to 802.11agn.
- Test network adapters before making a purchasing decision. Throughput can vary greatly depending on the adapter and laptop configuration.
- Check that wireless stations can connect to the specific channels being designed for use in the 802.11n network.
- Be sure all network adapters are running the latest drivers for optimal frequency and radio selection.
- If stations prefer 2.4GHz, it may be possible to coax them to 5GHz by doing some RF planning.
- If stations stick to an array/AP for too long, performance may drop. Tweak the Roaming Aggressiveness to optimise roaming behaviour.
- If possible, use the same NIC cards on all stations. At a minimum, use NICs with the same 802.11n feature set.

array with a stronger signal. The level should not be too high, or the station will spend too much time associating to new arrays and not enough time sending traffic.

Similar wireless NIC cards offering 802.11n and WPA2/AES encryption will improve network performance. Increasing the performance of as many stations as possible will have a tremendous affect on network performance.

## Conclusion

To achieve a high performance network capable of replacing wired switches, a multi-radio, 5GHz 802.11n Wi-Fi network will provide the best results. The overall network must be appropriately designed, from the supporting wired network, to wireless device placement, to the RF design. Stations should be capable of operating at 5GHz to take full advantage of 802.11n functionality and achieve maximum performance.

With a properly designed 802.11n Wi-Fi network, IT managers can deploy wireless networks to replace wired networks. Existing wired network budgets can be re-allocated to wireless equipment that, with proper design, will deliver similar end-user experience, but with all the flexibility and mobility benefits that wireless brings.

*From the Xirrus Inc. white paper* ***802.11n: Lessons Learned from the First 1000 Installations***

http://j.mp/hiRtYx

First published in the ***industrial ethernet book*** April 2011

# Car to roadside communication using IEEE 802.11p technology

**Provision of external services to vehicles has – until recently – been limited because of the lack of high-speed communications between them and service providers. The lack of standardised communications interfaces between various car manufacturers hasn't helped. The IEEE 1609 family of standards for Wireless Access in Vehicular Environments (WAVE) addresses these issues. Josef Jiru reports.**

THE EC has launched a Road Safety programme aimed at cutting the number of road deaths by half from 2011-2020. To help achieve this objective, a number of organisations have been carrying out research and development into car-to-roadside communication techniques. Among them is Fraunhofer ESK, the Munich-based institute that undertakes applied research into networked information and communication systems. The institute has been working on networking vehicle and the communication infrastructures (Fig. 1) to provide drivers with a more complete operational perspective of the vehicle, which can add to safety.

Early warning detection of road construction or accidents can also be provided to reduce road congestion. In addition, linking the vehicle and its environment enables optimal use of existing infrastructures.

Ultimately, nationwide, Europe-wide and global networks are needed to enable communications between all vehicles and roadside access points, or other vehicles.

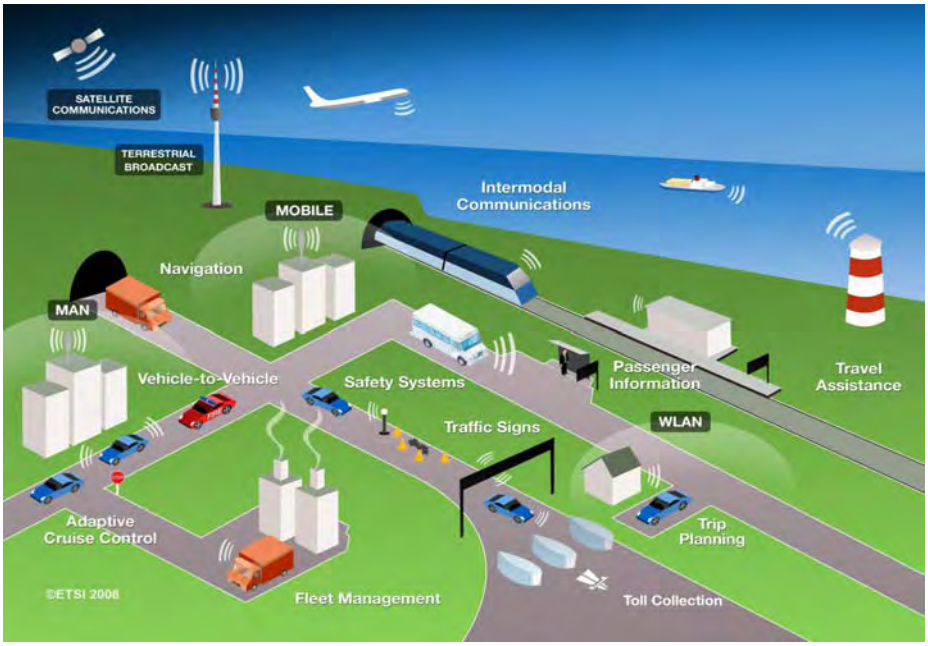


***Fig. 1.*** *Intelligent Transport Systems (ITS) include all types of communications in vehicles, between vehicles and between vehicles and fixed locations. Such systems include the use of information and communication technologies (ICT) for rail, water and air transport (picture ETSI).*

## IEEE 802.11p

Car-to-roadside communication is based on a WLAN (IEEE 802.11p) platform developed especially for vehicles, while IEEE 1609 is a higher layer standard on which IEEE 802.11p is based. IEEE 802.11p is an approved amendment to IEEE 802.11 that adds WAVE. It therefore defines the enhancements to 802.11 required to support Intelligent Transportation Systems (ITS) applications.

WAVE standards define an architecture and a complementary, standardised set of services and interfaces that together enable secure vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) wireless communications. Transportation benefits include greater vehicle safety, better navigation and traffic management, plus automated tolling.

Note that 802.11p will also be the basis for Dedicated Short Range Communications (DSRC). Based on the ISO Communications, Air-interface, Long and Medium range (CALM) architecture standard, this is a US Department of Transportation project that addresses vehicle-based communication networks, especially toll collection, vehicle safety services, and commerce transactions via cars.

## The technology

The car-to-roadside communication that Fraunhofer ESK has been looking at is based on this WLAN IEEE 802.11p platform. It is combined with a satellite positioning system to enable the exchange of vehicle positioning and sensor data with the vehicle's environment. The Universal Mobile Telecommunications System (UMTS) can be used as an alternative or additional technology. For time critical applications, however, UMTS can only be considered conditionally because the transmission of timed messages with short delays cannot yet be guaranteed.

The system will have to operate at vehicle speeds of up to 200km/h and support a transmission range of 1km. These requirements place high demands on the wireless communication network. Scalable transmission power will help avoid collisions on the wireless segment when traffic is highly dense. To enable time-critical, safety-relevant applications, the data will be prioritised and partitioned into different channels.

Because of the potential high vehicle speeds, topological changes are continuously generated that place many demands upon routing functionality. Many routing algorithms are available, but so far none has been found to be adequate on its own. A situation-dependent, hybrid solution is therefore needed.

On Board Units (OBUs) and the Roadside Units (RSUs) (Fig. 2) were designed for the demonstrator. These are based on embedded hardware and real-time Linux. The RSUs are linked in a multi-hop, fault-tolerant meshed network, with the routing protocol being based upon modified Optimised Link State Routing (OLSR) protocol. This is an optimisation of the classical link state algorithm adjusted to the requirements of a mobile wireless LAN. Multipoint relays (MPRs) are selected nodes that forward broadcast messages during flooding. The technique significantly reduces the message overhead compared with classical flooding, where every node retransmits each message when it receives the first message copy.

The current network topology is then displayed in real-time in the form of graphs, and the RSUs offer a variety of services through the meshed network. Requested services or services in the near vicinity of the vehicle can

## IEEE 802.11p MAC protocol for vehicle ad-hoc networks – VANETs

Vehicular Ad-hoc Networks (VANETs) have emerged as a key underlying technology in the realisation of Intelligent Transportation Systems (ITS). Vehicle communications can greatly reduce the incidence of accidents. One application is the Intelligent Transportation System (ITS) that includes automatic control services for improving safety, reducing traffic congestion and increasing passenger comfort.

Although VANETs are particular cases of the general mobile ad-hoc networks, they possess some characteristics that make its nature exclusive – and present challenges that require a set of new protocols. Vehicle speeds can be up to 150km/h, while the network topology changes repeatedly and unpredictably. Hence, the time duration for which a link is active between the vehicles is very short, especially when the vehicles are travelling in opposite directions.

One approach to enlarge the lifetime of links is to increase the transmission power, but increase in a vehicle's transmission range will result in increasing the probability of data collision and cause degradation in the overall system throughput. The other solution requires new protocols exhibiting a very low latency.

The usefulness of the broadcasted messages depends on latency. For example, if a vehicle is stopping or suddenly stops, it should broadcast a message to warn other vehicles of the probable danger. Considering that the driver needs at least three quarters of a second to initiate a response, the warning message should be delivered with virtually no system latency.

Position of nodes changes quickly and unpredictably in VANETs, so that, building an efficient routing table or a list of neighbour nodes will tire out the wireless channel and reduce the network efficiency. Protocols that rely on prior information about location of nodes are likely to have a bad performance. However, the topologies of a VANET can be an advantage because the vehicles are not expected to leave the covered road; therefore, the running direction of vehicles is predictable to some level. Privacy, safety and security fundamentally affect public acceptance of the technology. With VANETs, each node corresponds and reports a specific person and location.

### Implementation

VANETs involve both vehicle-to-vehicle (V2V) and vehicle-to-infrastructure/roadside (V2I or V2R) communications that rely on short-to-medium-range communication techniques.

The IEEE-1609 set of standards for Wireless Access in Vehicular Environments (WAVE) specifies an architecture that includes new standards for vehicle communication aimed at supporting ITS applications. IEEE-802.11p forms the bottom layers of the WAVE protocol stack and contains MAC and PHY layers derived from IEEE-802.11a. This makes IEEE-802.11p more suitable for high speed vehicles. The WAVE protocol forms a LAN to facilitate ITS applications; this LAN defines a WAVE Basic Service Set (WBSS) comprising vehicle on-board units plus roadside units.

The IEEE-802.11p MAC protocol is derived from the IEEE-802.11 distributed coordination function (DCF), and also uses the IEEE802.11e EDCA based quality-of-service (QoS) amendments. 802.11 MAC uses CSMA/CA that is specified in almost all variants of IEEE-802.11 (802.11a, 802.11b, 802.11g and 802.11p). RTS (Request-To-Send) and CTS (Clear-To-Send) mechanisms are used to resolve hidden and exposed node problems (such as found in vehicle movements).

In IEEE-802.11p, both the MAC (medium access control) and the PHY (physical) layers belonging to the DSRC/WAVE protocols are enhanced. Except for slight parameter changes to enable high user mobility, the IEEE-802.11p physical layer is identical to IEEE-802.11a. Moreover, the transmission power may be higher (up to +44dBm) in 802.11p compared to that in 802.11a. IEEE-802.11p MAC layer is derived from the basic IEEE-802.11 DCF.

The operating frequency for IEEE-802.11p is the 5.85-5.925GHz range in the licensed 5.9GHz ITS band. For safety applications requiring higher priority, one channel is dedicated to control. Other channels are service channels that can serve safety and non-safety services.

The function of the MAC is to coordinate the use of the communication medium. MAC layer protocols decide which node will access the shared medium at any time. As safety critical applications are designed to alert drivers about immediate danger, therefore requiring tight delay bounds, a MAC protocol has to take into consideration these strict application requirements.

### A study

The focus of a paper[1] on which this box is based is a performance analysis of the IEEE-802.11p MAC protocol. In the paper, the IEEE-802.11p MAC method was evaluated by developing a VANET model using MATLAB simulations. The scenario included time-critical packets periodically broadcast in V2V communications. Channel access delay and the probability of channel access were evaluated.

It was found that when the number of vehicles increased from 100 to 200, channel access delay increased by around 20ms, and the probability of channel access decreased by about 5%. More vehicles meant a greater chance of collision, as vehicles have less chance to access the channel. Moreover, increasing the VANET sensing range from 500 to 1000m reduced channel access probability by 35%.

The work showed that VANET performs well in small ranges because it becomes easy for vehicles to communicate with each other. Physical interference and the probability of hidden terminals also decreased, because in small ranges vehicle location is more likely to be known to other vehicles. Finally, when simulation times were varied between 50 to 80ms, the probability of vehicles accessing channels increased by 15%.

1. *Performance Evaluation of IEEE 802.11p MAC Protocol for VANETs*; Shahzad A. Malik, Madad Ali Shah, Shahid A.Khan, M.Jahanzeb, Umar Farooq and Adnan Khan. *Australian Journal of Basic and Applied Sciences*, 4(8): 4089-4098, 2010.

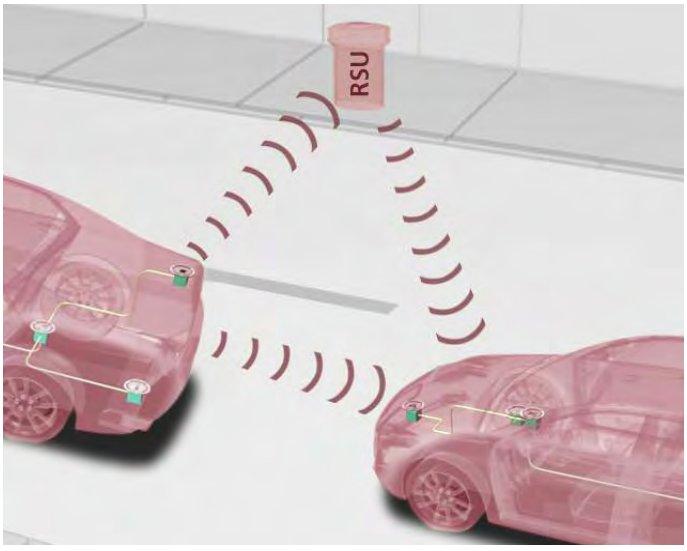


***Fig. 2.*** *On Board Units (OBUs) and the Roadside Units (RSUs) were designed for the demonstrator. These are based on embedded hardware and real-time Linux. The vehicle OBU communicates with the traffic signal, so drivers know when the next signal change occurs, allowing time to adjust driving and route accordingly. This can reduce traffic congestion.*

be displayed on the OBU. The OBUs and RSUs transmit beacons at regular intervals. To avoid placing an unnecessary load on the network, connections are established only on demand.

### Security

The different scenarios call for a reliable and robust communication platform to ensure that the entire system will be unaffected by local disruptions and interference. A security concept must also be developed to exclude attacks and manipulation. At the forefront are data security, authentication and data integrity, issues that can be addressed through encryption, digital signatures and certificates.

The deployment of a public key infrastructure is costly and, as a result, cannot be efficiently implemented on low-cost hardware. Even so, the system must guarantee the anonymity of the driver. This can be achieved by using temporary and revocable pseudonyms.

### The advantages

In contrast to car-to-car communication, car-to-roadside communication delivers advantages as soon as the first vehicle is equipped with the technology. No minimum penetration rate is called for. The benefits lie mainly in extra safety, better traffic management and good access to information and entertainment.

**Improved traffic flow** – Because the vehicle OBU communicates with the traffic signal, the driver knows when the next signal change occurs, allowing him/her to adjust driving and/or route accordingly. This can reduce traffic congestion.

**Increased safety** – Vehicle sensors register accidents and their severity. The vehicle transmits the accident information to the next RSU and a high-priority emergency call is disseminated throughout the mesh network. This warns other drivers in the vicinity of the accident, or forwards the accident report to emergency services control centres, such as police, fire and rescue services.

**More convenience** – Service stations supply the RSUs with up-to-date fuel prices and hours of operation and forward the data to RSUs in their vicinity. The vehicle can query the data for all service stations located along the planned route and display the service station with the best prices on the OBU.

Fraunhofer ESK offers concepts based upon fault-tolerant wireless

multi-hop networking of RSUs and the creation of a corresponding service platform. The organisation is looking to jointly develop the storage of up-to-date information in the RSU network (such as service station data), analysis and display of the processed data on the OBU (aggregation and filtering) and global networking (Internet).

To ensure reliable implementation of such concepts, an end-to-end simulation environment comprising traffic and network simulators is required. These will allow simulation of processes such as the intelligent aggregation and filtering of the data, or the routing algorithms in mesh networks with mixed architectures (stationary and ad hoc) and their impact on network loads. The information forwarding process and the driving strategies can then be analysed to determine their effectiveness. Such general frameworks are currently being developed.

Network continuity, an increasingly important issue, is also being improved. This applies not only to the continuity of multimodal navigation with up-to-date data, but also office-to-car continuity, in which the workplace shifts from the office to the car and is then adapted to the vehicle environment.

## Participation

To ensure interoperability between automobile manufacturers, all of the communication protocols and security concepts must be standardised. Various committees, including IEEE, C2C CC, ETSI and ISO, are working together to develop such standards. As member of the Car to Car Communication Consortium (C2CCC) Fraunhofer ESK is participating in this process.

Other current topics are new network technologies, software methodology, networks embedded systems, car-to-environment networks, electro mobility, wireless communication and sensor networks as well as mobile expert systems.

***Josef Jiru*** *is a research fellow at Fraunhofer ESK in Munich, Germany.*

First published in the ***industrial ethernet book*** June 2011

# WLAN: the future for railway communications networks?

**Technologies underlying train communications systems have advanced little since the first 'modern' rail systems were established 200 years ago. However, the latest wireless data technologies now enable the creation of advanced train communications systems. WLANs offer an optimal combination of bandwidth and cost-effectiveness for such operations, and Gigabit bandwidth enables real-time performance for passenger comfort and security. The future of railway communications is coming, and operators need to be ready, says Paul Hsu.**

DRIVEN BY CHANGING market expectations, bandwidth, response time and advancing technologies, train communications technology is in the midst of a highly significant transition. Train communications systems now must do more than ever before. The railway applications of today, tomorrow, and beyond demand more bandwidth, a faster real-time response time and more reliability from their communications networks, be they intra-train, train-to-ground, or trackside networks.

Typical train networks include the Ethernet train backbone, ground-to-train communications and onboard IP video surveillance.

## Response time

Traditionally, train control relied on human operators being given directions through some combination of radio, visual signals and track circuits, but the response time is slow, so for safety reasons, tracks were divided into long segments ('blocks'), with only one train allowed on a block at a time to prevent collisions – not that dissimilar to Industrial Ethernet in fact.

The introduction of Communication-based Train Control (CBTC) technology improved the efficiency of train operations by allowing operators to reduce the length of the blocks without compromising safety. However, the

| | Satellite | Cellular | WLAN |
|---|---|---|---|
| **Max Data Rate** | 20Mbps down, 384Kbps up | 7.2Mbps down, 384Kbps up | 54Mbps down, 300Mbps up |
| **Throughput** | Fair | Poor | Very good |
| **Train installation cost** | High | Low | Very low |
| **Infrastructure install cost** | Very high | High – covered by carrier | Low |
| **Service charges** | Yes | Yes | No |
| **Total cost** | Very high | Very high | Low |
| **Roaming** | None needed, but satellite occlusion blocks coverage in some areas | ISP-dependent | 100ms or less with fast roaming |
| **Mobility** | 300kph | About 150kph | About 150kph |

***Table 1****: Wireless technology solutions compared*


***Fig. 1.*** *Suitable WLAN communications equipment includes Ethernet switches and IP cameras with EN50155 and EN50121-1/2 certification to confirm their resilience in harsh railway environments, plus hardened wireless devices.*

efficacy of a CBTC system is highly contingent on the communications response time. With real-time response, the CBTC can safely and efficiently maximise the number of trains on the track.

In particular, communications must be sufficiently resilient to overcome the unique hazards of rolling stock operations - weather, shock, vibration and electromagnetic interference. The EN50155 and EN50121-1/2 standards are useful benchmarks for confirming that the communications devices are sufficiently robust for onboard and trackside applications.

## WLAN provides an answer

WLAN (Fig. 1) offers an optimal combination of bandwidth and cost-effectiveness for railway operations. Wireless technology frees operators from the limitations and complications of cabling a communications system, which is a particularly arduous task in an application with as many moving parts as a train system. Of all currently available solutions, WLAN stands out as the solution with the best balance of capabilities and cost – see Table 1.

As Table 1 shows, WLAN is clearly the superior choice, as it offers the most bandwidth and the lowest total cost. Monthly service charges are a significant continuing expense of satellite and cellular communications. With WLAN, not only are the installation costs low, but there is also no need to pay a satellite or cellular provider for data service. In terms of mobility, the development of optimised roaming technology has made WLAN mobile enough to support train to ground communications, even at relatively high cruising speeds.

## Trackside communications

WLAN eliminates cabling headaches for trackside networks, such as theft, maintenance and damage. Trackside networks consist of numerous wayside cabinets that share data up and down a length of track. Replacing the cables with WLAN units in each wayside cabinet eliminates such vulnerabilities.

## Intra-train communications

Compared with couplers, WLAN has a higher throughput and lower maintenance costs. Train operators are calling upon intra-train networks to support more applications, including passenger information, public announcement, video surveillance, intercom, HVAC, and data-driven train control systems.

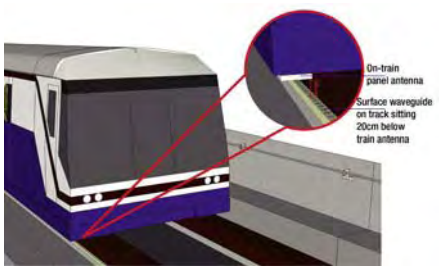


***Fig. 3.*** *The antenna remains within 20cm of the waveguide strip on the ground. This helps provide an even more secure link between client and AP*

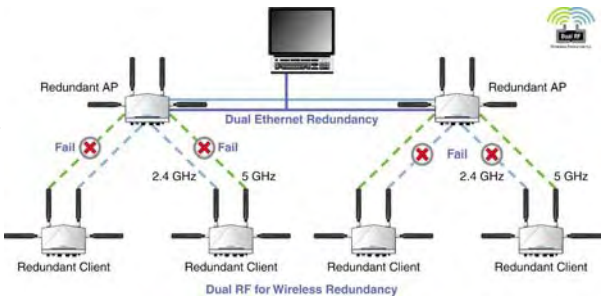


***Fig. 2.*** *Dual RF for wireless redundancy – concurrent redundant wireless connections. Hardened outdoor wireless APs with fast roaming and dual RF redundancy are ideal for train to ground communications, especially when deployed with a high AP density*

Conventional wired intra-train communications relied on couplers between carriages to send data down the line, but these have a fixed bandwidth and limited data rate, which places severe constraints on the upgradeability of an intra-train network. WLAN is a natural fit for intra-train networks that can reduce maintenance costs while increasing throughput to support more applications today and in the future. WLAN can comfortably support advanced CBTC and passenger infotainment systems.

## Train to ground communications

As the key link between the trackside network infrastructure and the intra-train network, train-to-ground communications is the lynchpin that enables innovative, next-generation railway applications such as rich passenger infotainment systems and Automatic Train Operations (ATO) through CBTC.

These systems simply would not be possible with conventional train-to-ground communications systems. ATO coordinates trains to maximise track utilisation and increase the service efficiency and frequency beyond that which is possible without central control. In order to safely do this, the control centre must receive and send a dizzying amount of data, including train status, passenger status, video data from cameras, and emergency controls. In addition, the Passenger Infotainment System must transfer real-time video, ads, news content, and more. 10 Mbps (or greater) is a reasonable estimate of how much throughput is needed to sustain these next-generation applications.

The sheer magnitude of throughput required is far beyond radio capabilities. This level of demand even strains the capacity of modern satellite and cellular data technology such as GSM and HSDPA. However, the IEEE 802.11 WLAN standard allows transfer up to 300 Mbits of data to comfortably enable all the applications envisioned today, with plenty of throughput left for future applications.

## Train-to-ground communications

A reliable and capable train-to-ground communications link is the foundation of many valuable next-generation train systems, but creating such a link can seem like a daunting task. How is it possible to maintain a consistent, uninterrupted link between fixed trackside access points and a quickly moving train that traverses many different operating environments? Luckily, WLAN technology has a large toolbox of solutions. Train tunnels are examples of clear environments having limited interference. Antennas are a cost-effective coverage solution for this kind of environment.

There is far more interference above ground, especially in busy urban areas. Still, WLANs can use a number of strategies to create networks that remain reliable in these conditions. Hardened outdoor wireless access points (APs) with fast roaming and dual RF redundancy are ideal for such a scenario, especially when deployed with high AP density. (Fig. 2).

For truly exceptional operating environments, waveguides (Fig. 3) or leaky (leakage) coaxial cables (LCX) provide an even more secure link between client and AP, albeit at increased cost. A track lined with waveguides or LCX cables offers stable, interference-proof access.

Suitable communications equipment includes industrial embedded computers, industrial Ethernet switches and IP cameras with EN50155 and EN50121-1/2 certification to confirm their resilience in harsh railway environments, and the rugged wireless devices to bring it all together. Gigabit bandwidth enables real-time performance for passenger comfort and security. Outdoors wireless AP/bridge/client devices should feature fast roaming technology, as well as the dependability provided by dual independent RF modules, power redundancy, and a weatherproof, dustproof, wide operating temperature design.

With two such units at each wayside cabinet, a railway can replace trackside cabling and its attendant maintenance headaches with dual redundant RF links. Units can connect to a solar-powered POE switch that supports, for example, an IP camera and a third wireless device that acts a local wireless AP.

Such a system can provide railway staff out in the field with convenient wireless access to system maintenance tools. Such equipment can also provide the foundation for a future train communications network with the addition of even more advanced wireless technology, complete with fast train-to-ground WLAN communications.

***Paul Hsu*** *is Product Manager for Moxa's Industrial Wireless Division*

First published in the ***industrial ethernet book*** June 2011

# Applying wireless to EtherNet/IP Industrial Automation systems

**While EtherNet/IP has many advantages, cable installation is often expensive, and communications to remote sites or moving platforms may not be reliable or cost-effective. Wireless Ethernet technologies have emerged that can now reliably reduce network costs while improving plant production. Applying these technologies is not a simple matter as industrial Ethernet systems vary greatly in terms of bandwidth requirements, response times and data transmission characteristics. Gary Enstad and Jim Ralston**

ETHERNET INDUSTRIAL Protocol (EtherNet/IP) is a network protocol defined by the ODVA. An important part of the EtherNet/IP standard is definition of Common Industrial Protocol (CIP) messaging. CIP defines the information packet with recognition that the message attributes will vary as applications do. Thus CIP message definition takes into account a wide range of applications including programming, diagnostics, data collection, PLC data exchange and I/O communications.

CIP defines two different types of connections. The first type is Explicit CIP which uses TCP/IP for its communications protocol. Explicit messages are unscheduled and use a request/response communications procedure or client/server connection. Examples of Explicit: executing a MSG statement between PLCs, HMIs, device diagnostics and program uploads/downloads.

The second type of CIP is Implicit; Implicit uses UDP/IP for its communication mechanism. Implicit connections are time critical, are scheduled and use a requested packet interval (RPI) parameter to specify the rate at which data updates.

Implicit connections use UDP packets for produce/consume data over an EtherNet/IP network. The UDP packets are usually multicast if there is more than one consumer of the data. This multicast address is assigned by the EtherNet/IP interface and is unique for each produced tag. Multicast IP addresses are used

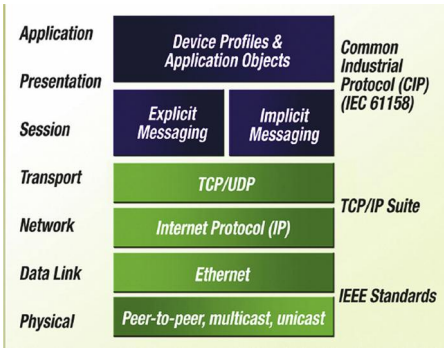


*EtherNet/IP uses the standard 7 layer OSI model for protocol definition. UDP over wireless can be problematic since successful frame reception does not form part of the UDP implicit messaging protocol*

PHOTO: FORD MOTORS

to make the network more efficient. A producer of data can produce data for multiple consumers. By using multicast packets, many devices can receive or consume this packet without the producer having to send it to each individual consumer.

EtherNet/IP I/O blocks may support two major implicit connection types: direct and rack optimised. A direct connection is a real-time, data transfer link between the controller and a single I/O module. Rack optimisation is a connection option where multiple discrete I/O modules in a chassis can be consolidated to use a single connection. Analogue modules typically can not be rack optimised and each analogue channel uses a separate CIP connection.

Proper network design is critical for implicit networking systems in order to achieve predictable deterministic I/O performance and to ensure that I/O traffic does not 'leak' outside of the automation network causing network degradation. ODVA recommends specific design strategies to ensure optimised network performance such as segmentation (isolating sub-networks), the use of managed layer three switches (IGMP snooping and multicast packet filtering) and high speed network infrastructure (100Base-T or faster). Wireless design is particularly critical as the wireless media is by nature slower than wired networks.

CIP safety is an extension of standard CIP. CIP Safety simply extends the application layer by adding a CIP safety layer to it. CIP Safety has generally been used where reliable communications is a must or stop on a failure is required. It has many triggers in place to detect critical and non critical errors and to close the connections in order to assure a safety condition.

New specification enhancements have been added to allow safety applications to have longer fault tolerances and the ability to assist in maintaining operations over wireless networks. Some of these enhancements include extending the RPI multiplier and the ability to configure the packet time expectation. These parameters are especially helpful in the wireless world where latency tends to be higher. This setting also could allow for re-transmission of RF packets if required to assure the safety packets get through. These changes lend themselves to make CIP safety well suited for wireless communications.

## EtherNet/IP wireless architecture

Industrial wireless applications can be divided into two broad categories: Those requiring high speed, low latency performance, and those permitting slower speed with longer packet latency. Wireless technologies are available to accommodate both.

A commonly made error is to assume that faster technologies are better. If the application can handle slower speeds, then using relatively slower frequency hopping technology may be the best approach. Frequency hopping is the most robust especially regarding communications in high RF noise areas, and easier to implement. As applications demand higher speeds, then more considerations and engineering challenges are typically encountered.

One of the most popular uses of wireless is in sharing I/O information between PLCs. As previously discussed, Explicit Messaging uses

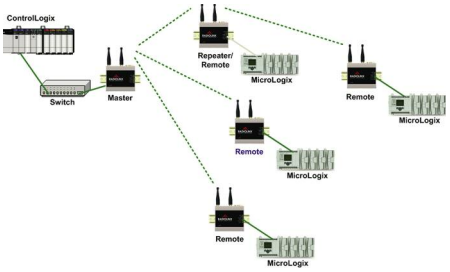


***Explicit Messaging application:*** *Wireless TCP/IP system used for data exchange*

TCP/IP based communications. Because these messages are unscheduled at the protocol layer, slower wireless Ethernet technologies may be used. MSG blocks in PLC ladder code may be programmed to accept long delays in transmission. If the application (process) is not time-critical, then a slower (but robust) frequency hopping technology may be the best choice.

There are many factors influencing how fast an explicit MSG may be executed. Generally, applications requiring 200ms response time or slower make a good candidate for FHSS. Faster response times may require faster technologies such as OFDM available in IEEE 802.11a and 802.11g standards.

## Antenna placement

Antenna placement is always important. Spread spectrum systems perform best with clear, unobstructed line-of-sight (LOS) between antennas. If there are obstructions (such as metal structures, concrete walls/floors, trees, etc.), then communications will be impeded. Obtaining clear LOS may not be practical or possible. Fortunately, the lower frequency 900MHz band offers relatively good reflectivity and penetration characteristics. When used with frequency hopping techniques, this band has the best chance to provide reliable data transmission in applications without clear line-of-sight although with relatively slower speeds. Applications without line-of-sight should always be thoroughly tested before implemented.

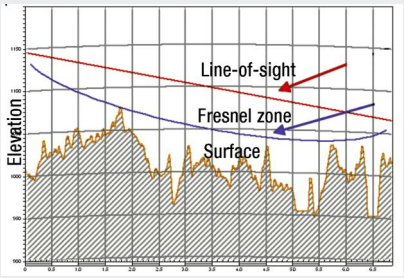


For longer range, outdoor systems clear line-of-sight between antennas is even more critical. Additionally, RF transmission theory dictates that the earth can reflect the signal in such a way to improve or impede it. A buffer zone in between the earth and the LOS is also needed for maximum signal levels. This area is called the Fresnel Zone, and is important when engineering the antenna system, particularly antenna height.

## Wireless for HMI networks

Another popular application for wireless is connecting HMIs to plant networks or machines. HMIs use TCP/IP communications and are not time critical other than to meet the needs of the process and, most importantly, the operator. While HMI screens may look complex and data intensive, the actual data being transmitted (updated) is usually minimal. If programmed efficiently, slower wireless technologies such as frequency hopping may be used. The key consideration is update times and the amount of information actually being transmitted.

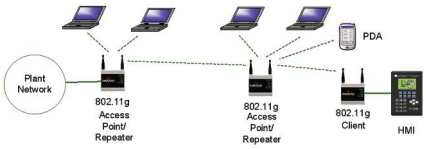


***Wireless TCP/IP*** *system for HMI connectivity*

FHSS may be the best choice if appropriate. However, if portable computers or PDAs are used, then the industrial wireless network ▶

must support the standard built into these portable devices. 802.11b and 802.11g (Wi-Fi) are of course the most common. This technology supports mobile operators while providing high speed, low latency communications.

## Remote SCADA Systems

Utilities extensively use remote SCADA networks where remote pump stations, tanks, sub-stations and pipelines are controlled and monitored from a central site. Wireless is a very good alternative to leasing phones lines. The main challenge with remote SCADA is distance and terrain between the sites. Wireless technologies that support EtherNet/IP speeds require unobstructed line-of-sight (LOS) and adherence to the Fresnel Zone for optimal performance. Analyzing the terrain and LOS obstructions is vital in determining the feasibility of such links. Fortunately, repeater sites are often available to achieve LOS and many FHSS radios have store-and-forward repeating capabilities.

Most EtherNet/IP SCADA systems do not require fast communications as Explicit Messaging is used to communicate from the remote PLCs back to the central plant. Frequency hopping is often the best choice here because FHSS offers the longest range due to excellent receiver sensitivity and support of the 90 MHz frequency band.

Bandwidth is limited in FHSS systems, so careful examination of network traffic is prudent. There is sometimes a temptation to add in other types of communication (such as VoIP, surveillance video, Internet connectivity, etc.) which will quickly exceed the capabilities

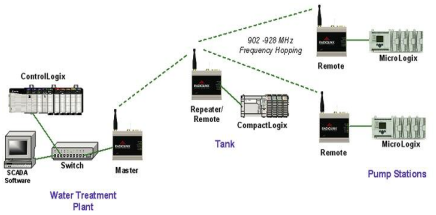


*Remote Wireless SCADA system using EtherNet/IP TCP/IP*

of an FHSS wireless system. IEEE 802.11-based systems offer the highest speeds for multiple use remote communications, but are limited in distance as their receiver sensitivity is lower and they do not support the 900MHz band.

## Wireless for Implicit Messaging

An emerging application for wireless is communications to distributed I/O blocks using EtherNet/IP. Wireless offers many advantages in these applications including elimination of mechanical coupling methods used in moving systems (e.g. rails, slip rings) and general cost savings due to reduction of Ethernet infrastructure. Communication to EtherNet/IP I/O blocks can also reduce automation costs compared to using remote PLCs. Programming is simpler using I/O instead of remote PLCs because MSG blocks are not required in the main controller's ladder program. But remember that Implicit Messaging is based on UDP/IP, not TCP/IP. Wireless networks must be carefully designed, and the plant RF environment more closely managed to ensure reliable communications.

Several factors should be carefully considered before choosing this architecture, including:

- Lack of remote PLC control (intelligence) in case of communication failure;
- Amount of I/O and required scan times (network traffic);
- Packet handling ability of wireless technology;
- Efficiency of the RF technology with multicast UDP packets;
- 802.11 clear channel availability.

However if circumstances are right, wireless EtherNet/IP I/O can be a significant cost saver and actually improve system reliability when correctly implemented, especially in moving systems.

## Practical implementation

Our company has done much work with EtherNet/IP and has many successful customer installations. The following information is provided only as a guideline towards predicting wireless performance and should not be relied upon for any other purpose. We always recommend field testing to confirm wireless performance and reliability.

**Predicting wireless I/O performance**. Because EtherNet/IP I/O messages are scheduled, it is possible to predict scan time performance over industrial wireless systems if the following conditions are met:

- Packets per second performance of wireless technology;
- Wireless behaviour in multipoint systems (handling of UDP Multicasts);
- Number of CIP connections.

The first step in designing a wireless EtherNet/IP system is to calculate the number of packets per second which determines minimum wireless bandwidth requirements. Start by counting the number of CIP connections. To calculate how many connections are in an I/O system, sum up all direct connections and rack optimised connections. To determine how many packets per second the system will be using, multiply each connection by two. *Two...?* This is because each CIP connection is bi-directional meaning that during every Requested Packet Interval (RPI), a produced packet is sent by each end of the connection.

For example, if there are five direct connections and two rack-optimised connections (with six digital modules in each) this equals seven total CIP Connections, the total number of packets is then calculated:

## The importance of modulation type and channel spacing

Direct Sequence and Orthogonal Frequency offer the fastest spread spectrum data rates as the wide channel permits transmission of complex modulation schemes. OFDM uses a complex modulation technique and is capable of high data rates and low latency (the transmission time a packet takes from one end to the other). OFDM is also significantly more immune to multipath fading, a problem due to RF reflections that high data rate systems frequently exhibit. (Note that slower speed frequency hopping systems are relatively immune to multipath).

Direct Sequence and OFDM are the methods used by all popular open Wi-Fi standards today including IEEE 802.11b, 802.11g (both transmitting in the 2.4GHz band) and 802.11a (transmitting in the 5GHz band). IEEE 802.11n is nearing ratification at the time this paper was written, and will also incorporate these techniques. While the wide band modulation offers high speed, it unfortunately is more prone to noise problems when multiple systems are operating in close proximity. For example, IEEE802.11b/g has thirteen channels available (eleven channels in North America), but only three channels don't overlap.


*Non-overlapping 802.11b/g Channels (1, 6 & 11). Due to overlapping channels and the popularity of Wi-Fi systems in plants, band crowding and RF saturation can lead to poor wireless performance*

## Frequency hopping sequencing

Frequency hopping has a high RF energy per bit ratio, and by incorporating error correction techniques, frequency hopping offers the best chance for successful data transmission as the transmitter will send the packet over and over again using different channels until an acknowledgement is received. However the process is slower than Direct Sequence/OFDM and has longer data latency. Most Frequency Hopping systems are limited to 1Mbps or lower data rate. But if the data rate is fast enough for the application, the reliability of frequency hopping is tough to beat, especially in high noise environments.

IEEE 802.15.1 (Bluetooth) is one of the few open standards incorporating frequency hopping. Because of the distance limitation of IEEE 802.15.1 Bluetooth devices are seldom applied to wireless Ethernet systems. Most industrial frequency hopping modems are proprietary; vendor X will usually not communicate with vendor Y. While this potentially presents a disadvantage for commercial systems, it can be desirable for industrial systems for two reasons: Security and isolation from the wireless IT system.

Because industrial frequency hopping technologies are not typically based upon an open standard, the manufacturer can use unique authentication processes and sophisticated encryption techniques to ensure high levels of security. While security has significantly improved in Wi-Fi systems with WPA and WPA2 standards, hackers will continue to look for holes. Fortunately many industrial Wi-Fi manufacturers now include an option to hide the access point by not transmitting its SSID beacon, effectively hiding the access point from potential hackers. Other security techniques include cryptographic encryption, key management and rogue access point detection providing a high degree of security just as frequency hopping systems provide.

Frequency hopping also offers plant managers the ability to operate their own wireless network separate from the IT department. Because of the popularity of 802.11 technologies for wireless network access, proprietary frequency hopping systems may offer the best choice for industrial systems – and keep the peace between department managers.

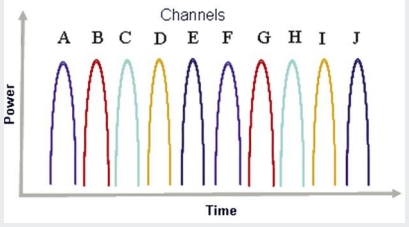


*Frequency hopping channel sequencing*

7 CIP Connections x 2 = 14 Packets

Note that the six rack-optimised modules in each rack only count as one connection. Rack optimisation (if available in the I/O hardware) can significantly reduce wireless traffic.

Next multiply the packets by the scan time (derived from RPI setting) to calculate packets per second (pps). Let's assume that in the above example, the required RPI time is 20ms (actual RPI time is application dependant), we know that there are 50pps at an RPI time of 20ms (1/0.02). We then multiply the 14 connections by the 50pps to get the overall packets per second rate:

14 x 50 = 700 packets per second

The overall packets per second rate for 802.11 a/g radios can be in the thousands. However it is best practice to not operate the radio network at maximum capacity. Rockwell Automation suggests reserving 10% of each adaptor's bandwidth so it is possible to use its RSLogix 5000 software for remote programming.

It is also suggested that 30% of the radio's packet rate be reserved for RF overhead. In a congested RF environment a radio contending for the RF medium will use up valuable time if the radio determines the channel is busy by ▶

| RPI Setting (ms) | CIP Connections | | | | | | |
|---|---|---|---|---|---|---|---|
| | 5 | 10 | 20 | 30 | 40 | 50 | 60 |
| 5 | 1000 | 2000 | 4000 | 6000 | 8000 | 10000 | 12000 |
| 10 | 500 | 1000 | 2000 | 3000 | 4000 | 5000 | 6000 |
| 20 | 250 | 500 | 1000 | 1500 | 2000 | 2500 | 3000 |
| 30 | 167 | 333 | 667 | 1000 | 1333 | 1667 | 2000 |
| 40 | 125 | 250 | 500 | 750 | 1000 | 1250 | 1500 |
| 50 | 100 | 200 | 400 | 600 | 800 | 1000 | 1200 |
| 60 | 83 | 167 | 333 | 500 | 667 | 833 | 1000 |
| 75 | 67 | 133 | 267 | 400 | 533 | 667 | 800 |
| 100 | 50 | 100 | 200 | 300 | 400 | 500 | 600 |
| 200 | 25 | 50 | 100 | 150 | 200 | 250 | 300 |

***Table 1.*** *Calculating packets-per-second where:* CIP connections x (1/RPI) = pps

using its carrier sense mechanism. If the radio determines the medium is busy it ceases packet transmission while it runs its back off algorithm and then re-accesses the channel. All this takes time.

A radio network operating in a highly congested RF environment can easily use 10% of its packets doing RF retries. RF retries can occur if a packet is lost or corrupt due to poor signal to noise ratio, antenna placement or multipath fading problems. Point-to-multipoint systems consume higher amounts of bandwidth Selecting a clear channel is good practice in wireless EtherNet/IP I/O networks.

The next step is to determine a reliable packet rate that the wireless technology will reliably support while keeping in mind that at least 40% should be reserved for other applications and RF overhead.

## Impact of multicast UDP packets

As previously discussed, produced CIP packets are multicast over the Ethernet network to accommodate multiple consumers. In wired systems, managed switches with IGMP querying are recommended to direct multicast UDP traffic only to the segments that need them. This ensures that high speed I/O data does not reduce performance of the plant Ethernet network, a major concern of IT managers.

Similarly, 802.11-based access points will re-broadcast multicast UDP packets to all active wireless clients. This represents a major problem because the unnecessary broadcast of high speed UDP packets will quickly clog an 802.11 channel – significantly reducing performance and even dropping UDP packets which will cause system errors.

One way to correct this problem and optimise wireless performance is to invoke IGMP Snooping and multicast packet filtering at the RF layer. By determining which devices are actually consuming the packets, the radio can build a consumption table and eliminate needless re-broadcasts. In point-to-multipoint systems, this feature can improve throughput by as much as 30% while significantly reducing dropped packets.

By applying multicast filtering to the 802.11 standards, it is possible to predict packet-per-second performance even in multipoint systems. For example, ProSoft Technology's 802.11a/b/g Industrial Hotspot has been determined to support 1800 packets per second. This equates to a little over 1000 packets per second available for EtherNet/IP CIP packets after subtracting 40% for RF overhead and other applications.

Table 1 shows how to calculate packets-per-second and highlights in green where an 802.11a or 802.11g wireless system may be used.

## Wireless I/O application example

A manufacturer had a problem with moving carriages in the production plant. Each carriage was controlled by a ControlLogix PLC to Flex I/O and drives onboard the carriage. EtherNet/IP was used over flex Ethernet cable. The carriages travel over a 140m track at moderate to fast speeds.

Ethernet cable breakage would occur due to the frequent motion of the carriage. When the cable failed, the Flex I/O system would create E-Stop condition shutting down the carriage. This resulted in the carriage abruptly stopping without warning to the operator riding onboard. Production on the line could not continue until the cable was repaired.

The plant engineers researched the possibility of replacing the flex cable with a wireless Ethernet system. They calculated that they needed a wireless technology capable of supporting an RPI time of 32 and of supporting 12 side by side lines with no interference. After consulting with us, they selected the 802.11a/b/g Industrial Hotspot. It was selected because it would support the RPI time, has 12 non-overlapping channels, excellent vibration specifications and diagnostics.

The plant engineers installed a pair on a test line to prove to management (and themselves) that the wireless technology would be more reliable than the flex cable. This line performed without failure for over 30 days, and they have since converted three other lines to wireless.

## Newest wireless technologies

While this paper focused on widely available FHSS and IEEE standards such as 802.11a and 802.11g [*at the time of writing* – Ed], there are several wireless standards on the horizon that promise higher performance and connectivity options for EtherNet/IP networks.

**IEEE 802.11n** promises several features that are attractive for EtherNet/IP communications including dual band (2.4 and 5GHz) support, significantly faster packet transfer rates with a reported throughput up to 300Mbps and RF propagation that actually takes advantage of reflected signals (quite common in industrial plants with lots of metal) using multi-input, multi-output (MIMO) antenna systems.

**IEEE 802.16** (WiMax). While popularly known as an emerging cellular technology, WiMax technology will soon be available in the spread spectrum (license-free) bands including 2.4GHz. WiMax offers high speed (up to 70Mbps) at potentially long range. WiMax technologies may dramatically improve data rates to remote industrial sites and SCADA systems.

**ISA100.11a**. The ISA is working on the ISA100.11a standard for wireless enabled devices, such as sensors. Operating in the 2.4GHz band, the technology will 'sense' existing 802.11b/g systems and work around them. While designed primarily for embedded devices, EtherNet/IP adapters and gateways will likely be supported.

*Gary Enstad has a BS in Electrical Engineering and has been involved in wireless design and technical support for over nine years. His current role is Wireless Application Development Engineer for ProSoft Technology*

*Jim Ralston has been involved with the design and support of industrial wireless systems for over 12 years. He is currently the Northeast Regional Sales Manager for ProSoft Technology*

*This is an edited version of the Paper* '**Applying Wireless to EtherNet/IP Automation Systems**' *presented at the ODVA Conference, 2009*

First published in the ***industrial ethernet book*** May 2009

# Simulcast RF wireless networks aid data transmission integrity

**Radio interference has been and always will be a major concern for wireless applications. In a wireless environment, data transmission is over air and due to the characteristics of this medium, a basic appreciation of RF engineering knowledge is required to ensure a reliable wireless connection. Because interference normally occurs at a particular frequency, if two or more different frequencies are used to communicate at the same time, then data transmission can continue, even if there is interference on one of the frequencies.**

THE STANDARD architecture of wireless infrastructures includes access points (AP) that connect many clients to an Ethernet network. Since the APs and clients are connected by a

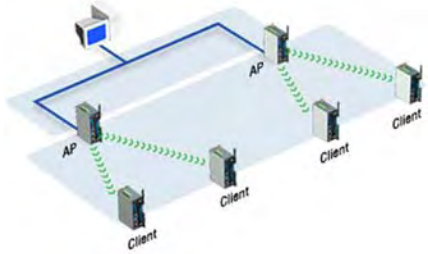


*Traditional Single-RF Wireless Architecture*

single-RF connection, if the RF connection fails, the system and network behind the client will be disconnected. With dual-RF wireless architecture two independent RF modules are used to form independent wireless connections using different frequencies to avoid interruptions in transmission. To achieve network redundancy without needing to change existing wireless LAN architecture, APs and clients support dual RF channels – usually operating at 2.4 and 5GHz simultaneously.

## Dual-RF redundancy

The concept has been applied by incorporating two RF modules in a single wireless LAN device to enable two independent wireless connections. The hardware uses this IEEE 802.11a/b/g-

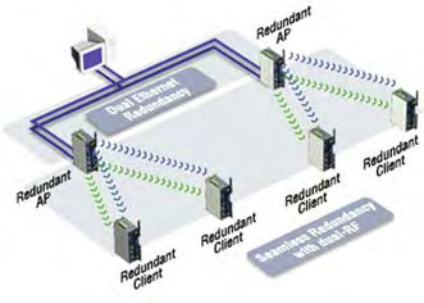


*Dual-RF architecture*

compliant wireless redundancy scheme. To achieve reliability beyond wireless redundancy, Ethernet redundancy is also added to the implementation. The external manifestation of additional wired redundancy takes the from of

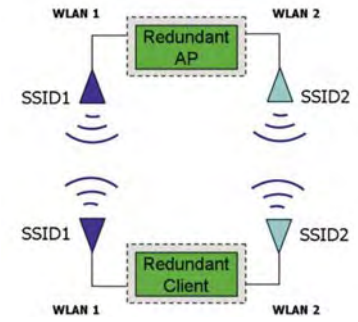


*Dual-RF—Wireless Redundancy Mode*

two Ethernet ports capable of running both RSTP and Moxa's own proprietary ring protocol over the wired side of the network.

In practical terms the devices would implement two independent wireless connections between the redundant AP and redundant client devices. One path uses 2.4GHz while the other operates at 5GHz to prevent interference. If one of the two wireless connections fails, the other connection will continue providing service between the redundant AP and redundant client devices.

Configuration involves setting up a redundant AP on the AP side, and a redundant client on the client side, each RF path using a different SSID. The figure at the top of this column shows the system layout in which WLAN1 is set to SSID1 and WLAN2 is set to SSID2. In addition to connecting up its redundant clients, the redundant AP can also serve one or more traditional single RF clients using standard single path service.

## Wireless bridge mode

The hardware implementation leverages the dual RF internals to provide what it calls a wireless bridge mode in which WLAN1 is configured as the master AP and WLAN2 as the slave client. The mode will not reduce the bandwidth but will extend the wireless range. More importantly, this is designed to optimize the WDS (Wireless Distribution System) mode in light of its throughput performance. WDS mode's normal throughput is

$$25\text{Mbps}/(n-1)$$

in which n is the number of WDS nodes. For example, if there are 4 mesh nodes, the throughput is around 8Mbps. The AWK hardware's wireless bridge mode can increase the throughput from 10Mbps to 15Mbps says the company. This way, the performance of each bridge connection will remain the same.

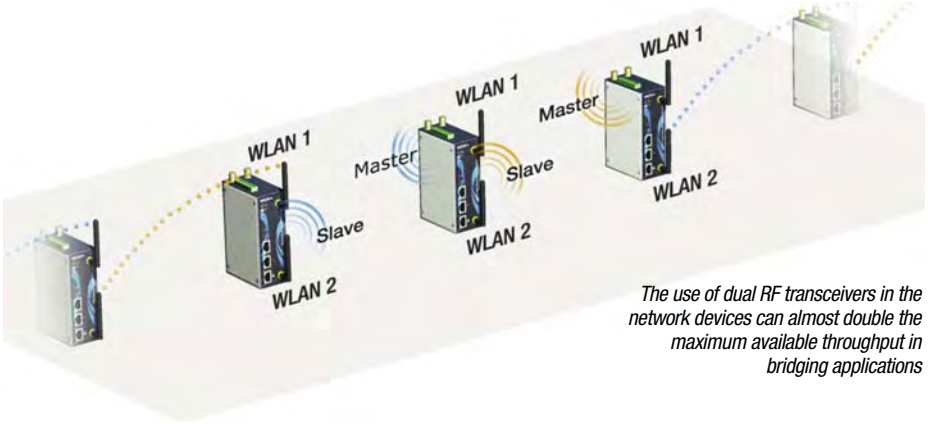


*The use of dual RF transceivers in the network devices can almost double the maximum available throughput in bridging applications*

## AP-client connection mode

Most WLAN applications use infrastructure mode. In AP-client mode, a wireless AP is required to set up a basic infrastructure service set (BSS) for wireless connectivity. The AP can be used by itself to set up a WLAN, or can be used to connect the WLAN to a wired network. The hardware supports AP-client connections, which can be used to provide Internet access in areas where cabling would be too expensive or otherwise impractical to install.

*From a Moxa technical application note*

First published in the ***industrial ethernet book*** November 2009

# Controllers to centralise the management of big WLANs

**The use of Wireless LAN has gained increasing acceptance in automation. Hardware manufacturers have accordingly designed products for use in line with industrial requirements. This radio technology mostly finds employment within application islands; additional requirements have to be fulfilled if it is to be used over a wider area. Achieving coordinated wireless operation requires a degree of centralised management. A new design take on the WLAN controller enables the creation of extended networks comprising hundreds or thousands of access points says Hirschmann's Olaf Schilperoort**

IN WHAT SEEMS like an age ago back in 2002, a few entrepreneurs linked to American universities set out to turn the WLAN world on its head. Before then, radio networks consisted of one or more access points connected to a wired network which enabled mobile clients such as laptops and PDAs to gain network access. These access points, which were used predominantly in the office sector, had just the basic functions necessary for a Layer 2 infrastructure. Only few of them included any sort of routing or firewall mechanisms.

The start-up companies initially considered adding extensive Layer 3 functions. They wanted to take WLAN beyond local applications in order to construct and centrally administer campus and company-wide networks. These networks could extend to several thousand access points. There was also a question of the software needed to manage this number of devices. But the promise was to administer large networks with a server-like device.

## The original idea

This device was initially known as a WLAN switch, because it was meant to control all the access point functions centrally. The access points were given the commercially attractive name *Thin AP*. Their task was limited to that of providing an interface between cable and radio link. Every packet sent and received – including all control and management information – was initially sent to the central WLAN switch, before the latter relayed the user data to the local area network (LAN). In order to make this possible, the switch and access points were inter-connected by a Layer 2 tunnel.

This central approach was actually advantageous for large installations since WLAN clients that move around an extended site had to remain within the same IP sub-network at all times. The reason for this was to avoid new IP addresses having to be assigned when switching to other sub-networks. The inevitable interruption of the network connection would otherwise last several seconds, causing VoIP for example to suffer service interruptions. WLAN switches were therefore regarded as a solution to cut down roaming handover time.

PHOTO: BELDEN

Roaming in this context is the term given to the change by a client from one access point or sub-network to another.

However WLAN switches turned out to be a bottleneck as each data packet first had to be transported from the receiving access point across the entire network to the central device which only then dispatched it to the assigned destination in the LAN (Fig. 1). This led to the switch's bandwidth quickly becoming exhausted. This problem was further exacerbated with the advent of WLAN Standard IEEE 802.11n, which hiked the data rate possible with each access point. With bandwidth continuing to be an important sales argument in the office sector, the development of WLAN switches looked as though it was going to reach an impasse.

## Decentralised improvement

For such reasons of bandwidth management, it became necessary to work around some aspects of the centralised approach. WLAN switches became WLAN controllers, which now only exchange control data with the access points. Many functions, which had previously been centralised, migrated back to the access points. In addition, the CAPWAP Standard (Control and Provisioning of Wireless Access Points) had since been defined, which describes how the Layer 2 tunnel between controller and access point should be structured; this standard was later also included in IEEE 802.11w. This subsequently led to a symbiosis of the centralised and decentralised approaches. In other words, it combined the flexibility of the local, standalone access point with broadband network connection and the benefits of a high-performance controller solution.

A WLAN controller locates all available access points when the network is first configured, provided a cable connection exists to the

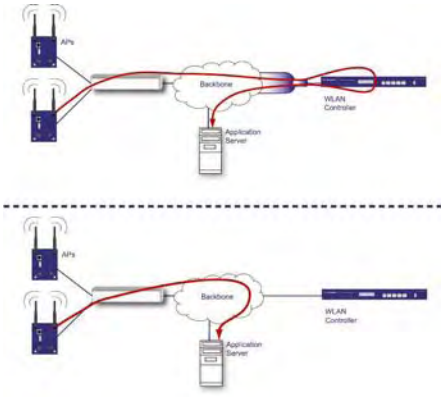


*Fig. 1. WLAN switches were regarded as a solution to cut down roaming handover time. However WLAN switches turned out to be a bottleneck as each data packet first had to be transported from the receiving access point across the entire network to the central device which only then dispatched it to the assigned destination in the LAN*

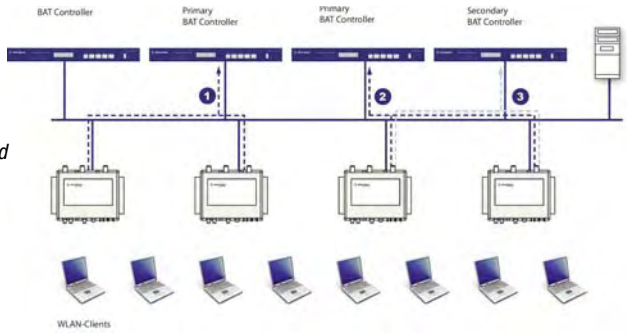


*Fig. 2. WLAN controllers enable defective access points to be replaced rapidly and also support and replace each other*

controller's network. The controller then checks whether the access points have the right firmware and also sends out configuration data assigning a logical network position. This saves the administrator time and effort in putting a functional network together. During operation, the controller provides all the necessary information via the network. For instance, if an access point fails, the controller immediately recognises the alternative device as soon as it is replaced and connects it to the network. The management software therefore no longer has to call up the data from each individual access point. This reduces network load and information is available more quickly.

In addition, further functions can be provided which would not be possible without a complete overview of the network. For example, controllers can be used to set up the WLAN network in such a way that channel distribution results in the fewest possible overlaps and problems between the access points. As the controller represents the central access to the WLAN network, it also becomes the gate keeper between the wired and wireless part of the network. The controller then provides a firewall and security function. The devices can also serve as VPN gateways that can interconnect several WLAN networks. Moreover, controllers can also be set to provide redundancy if one of them fails (Fig. 2).

## Controller-compatible access points

WLAN controllers are not economic for small network use. In this case, there is still a need for stand-alone access points designed for industrial applications. But if 50 or more of these devices are used, this investment will certainly be worthwhile. Industrial access points that can be managed by WLAN controllers are an important consideration since enhanced devices provide a ready made upgrade path should controller-based operation prove necessary.

In the office sector, controller-supported networks are currently used mainly for VoIP where roaming across several sub-networks is the main benefit. In automation, a WLAN network rarely has to handle several applications simultaneously for all the reasons that one might expect: diagnostics are too complex and the possible error sources too great. The main purpose behind the use of a controller for industrial applications would be to ease network administration.

In addition, client localisation is also desirable for process automation. This enables the administrator to see at any time where personnel with a portable WLAN client are physically located, whether they are completing their work in the right place. Controller-based networks are especially suitable for such tasks, as they provide the necessary data centrally.

## And finally...

There is an increasingly wide range of WLAN applications for automation within the context of networks with a growing complexity. With the WLAN controllers, it becomes possible to manage networks with hundreds or thousands of access points from a central location.

***Olaf Schilperoort*** *is with Belden, Inc*

First published in the *industrial ethernet book* April 2010

# Tracking assets: RFID meets industrial Wi-Fi networks

**Globalisation and mobility trends have profoundly affected industrial plant operations and entire business models, with assets that make up the supply chain being constantly in motion. Efficiently managing people, products, equipment and raw materials is essential to keep global manufacturers competitive. One fast growing technique that enables this is a combination of wireless networks with RFID technology.**

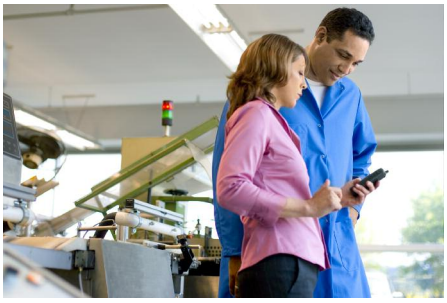

PHOTO: CISCO

IN THE PAST, a lack of visibility into the location of valuable assets has hindered efficiency and resulted in operational problems such as misplaced mobile tools, jigs, machinery, parts, or work-in-process (WIP) inventory. This results in high equipment leasing and/or replacement costs to offset losses.

Worker safety and security is another issue. Accidents and security incidents can make it life-critical to quickly locate employees and building exit routes. Context-aware conditions help minimise injury and losses. As manufacturers seek to continuously improve business agility, quality, reliability and safety, they are taking advantage of increasing intelligence in devices and networks.

Increasingly 'smart' devices, which include radio frequency identification (RFID) tags and sensors that have advanced diagnostics, are contributing to the billions of devices now connected to IP networks. This proliferation of smart devices is referred to by some as the 'Internet of Things', and it is projected to grow to trillions of devices that will be connected using the emerging IPv6 protocol[1]. For manufacturers, a growing number of connected smart devices promises to revolutionise portability, mobility, context-aware condition and use of critical assets.

However, devices must be continuously connected – including in tough manufacturing environments – and be integrated to deliver context-rich information, alongside with data. The combination of smart devices with wireless networks is now able to deliver powerful smart services, such as context-aware, real-time asset management and location services.

## RFID + wireless

Technology analyst firm IDC included smart services using RFID as one of its ten Manufacturing Industry Predictions for 2011[2]. Indeed, the past year has seen an explosion of RFID solutions used to track almost everything imaginable.

Research firm Gartner suggests that the explosion of mobile smart devices is leading to an equivalent explosion of mobile applications, which will create new wireless infrastructure requirements[3]. When RFID solutions are teamed with wireless networks, manufacturers can suddenly increase operational visibility almost anywhere. Today, RFID and wireless networks are able to deliver immediate, unprecedented visibility into assets, WIP and people's locations.

Manufacturers have been piloting and deploying RFID in their organisations for some time, and several trends are developing.

## Passive or active?

There are both 'passive' and 'active' RFID tag types. Passive RFID uses low cost tags that are generally disposable as consumables (as used in shops to prevent theft etc), while active RFID uses more sophisticated tags that cost a little more at typically Euro 17 – 68 ($25-$100). There is also a third type – semi-passive tags. These are similar to active tags in that they use a battery to run the microchip's circuitry, but they cannot communicate with RFID readers. To conserve battery life, some semi-passive tags remain in a 'sleep' mode until contacted by a reader.

Active RFID tags possess greater functionality and tend to be associated with supply chain or manufacturing process workflows. They are reusable and have batteries that last five or more years. Active RFID tags can also function as low-cost remote sensors that broadcast telemetry back to a base station. Active-RFID deployment is a sector that is seeing good payback and return on investment for manufacturers.

Looking at wireless transmission, work with RFID and the Real Time Location System (RTLS) has led to the development and support of multiple RFID technologies. The first is Received Signal Strength Indicator (RSSI). This is a measurement of the power present in a received radio signal. RSSI is generic radio receiver technology metric, which is usually invisible to the user of the device containing the receiver, but is directly known to users of wireless networking of the IEEE 802.11 protocol family.

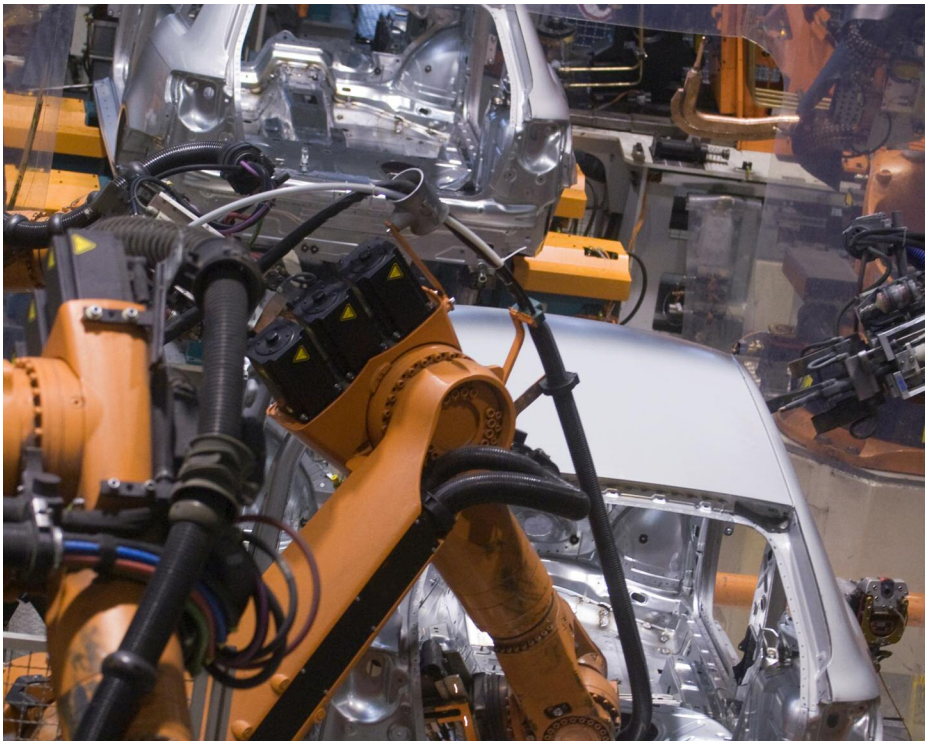

PHOTO: CISCO

## RFID technology

Active RFID tags have an onboard battery to power the microchip's circuitry and transmit signals to readers. Typically, active tags can be read from distances of around 30m. They also allow more applications by offering two-way communications, sensor integration, independent system intelligence, and constant viability. Active RFID tags work using Received Signal Strength Indication (RSSI) or Time Difference of Arrival (TDOA) technologies.

RSSI: Some tags associate with APs to provide a regular beacon or 'chirp' (a 416-bit 802.11 frame). This reduces tag/AP interaction to a simple unidirectional packet – there is no state for the AP to maintain and no IP address required. Also, less radio time means a longer battery life. This is common amongst IEEE 802.11 wireless protocol families, whether a, b, g or the newer n flavour. Other tags send out a more constant signal, become associated with the AP, providing a more constant awareness with stronger integration. However, battery life is shorter.

In an 802.11 Wi-Fi network, the APs are arranged so that they receive signals from the RFID tag and triangulate its position based on the received signal strength at each AP. The stronger the signal, the closer the tag is to the AP. Each AP sends it's strength metric to the wireless control system, which interprets the signal strength via an appliance and middleware. This provides coordinates via APIs to user interface software, which – in turn – specifies location. The results can be sent as coordinates or shown on a visual map.

For Lightweight Access Point Protocol (LWAPP) environments, location information is sent to the Location Appliance and Wireless Control System (WCS) rather than an alternative separate non-LWAPP engine, which can also calculate tag positions. Accuracy can be 2 – 3m, depending on environment, location and AP density.

TDOA: This is similar to RSSI in that the network APs triangulate tag locations based upon the signal, but here it is the time it takes for the signal packet(s) to arrive at the AP that is important. The shorter the time taken for the signals to be detected, the closer the tag is to the AP, and a position is estimated using one or more locating algorithms in the location appliance and WCS. This method is more resilient for high ceilings and outdoor applications, and is less disturbed by building structures.

In another application there are chokepoints ('exciters'). This works using strategically placed readers or triggering devices – the exciters. When a tag passes within range of a 125 kHz exciter, the tag emits a short burst of Wi-Fi data that includes that exciter's unique identification. With an exciter, immediate tag recognition occurs as the tag is 'excited' and emits a response. Conventionally, it can take several minutes before tags chirp and the location is updated. Exciters, therefore, provide a more real-time experience for tracking moving items. Additionally, tags can be set to turn themselves on or off as they go through chokepoints.

In summary, RSSI is good for most indoor tracking applications and uses existing standards-based Wi-Fi networks. TDOA uses a more proprietary approach, but can be backhauled via Wi-Fi networks and is good for outdoor and difficult environments. Choke-points are good for tracking movement of assets through doorways and designated areas, so helping to prevent asset loss or tracking mislaid items.

Organisations that define standards and regulate the use of RFID include ISO, IEC, ASTM International, DASH7 Alliance and EPCglobal.

## Open standards platforms

With different protocols and standards abounding, a key technology development has been the introduction of open standards based platforms.

The pace of adoption is starting to pick up. The customer stories below tell the business benefits side, but there are also IT and process benefits. With an open-standards WCS4, the application can integrate with warehouse, ERP and other business systems. Tracking WIP can become part of the overall methods of manufacture, not only for the supply chain, but also plant-floor activities, such as MES and WIP.

Challenges remain. Firstly, manufacturers want reliable and available systems. A system is now available that uses a wireless interference monitoring technology. This comprises special hardware built into certain APs, plus advanced interference identification algorithms that automatically detect radio interference and can map its source, rather like an integral spectrum analyser. Microwave and blue-tooth devices, for example, can cause interference. The system identifies the interference and detects where it's coming from and can mitigate it, providing a more reliable network.

Secondly, 802.11n adoption continues to grow. These are higher performance than the older networks, but manufacturers' IT departments may still have legacy deployments, so technologies are still needed to cope with both old and new. There are ways of improving the network performance of legacy 802.11a and g devices (older Wi-Fi networks), whilst still enjoying the benefits of a higher performance 802.11n Wi-Fi network.

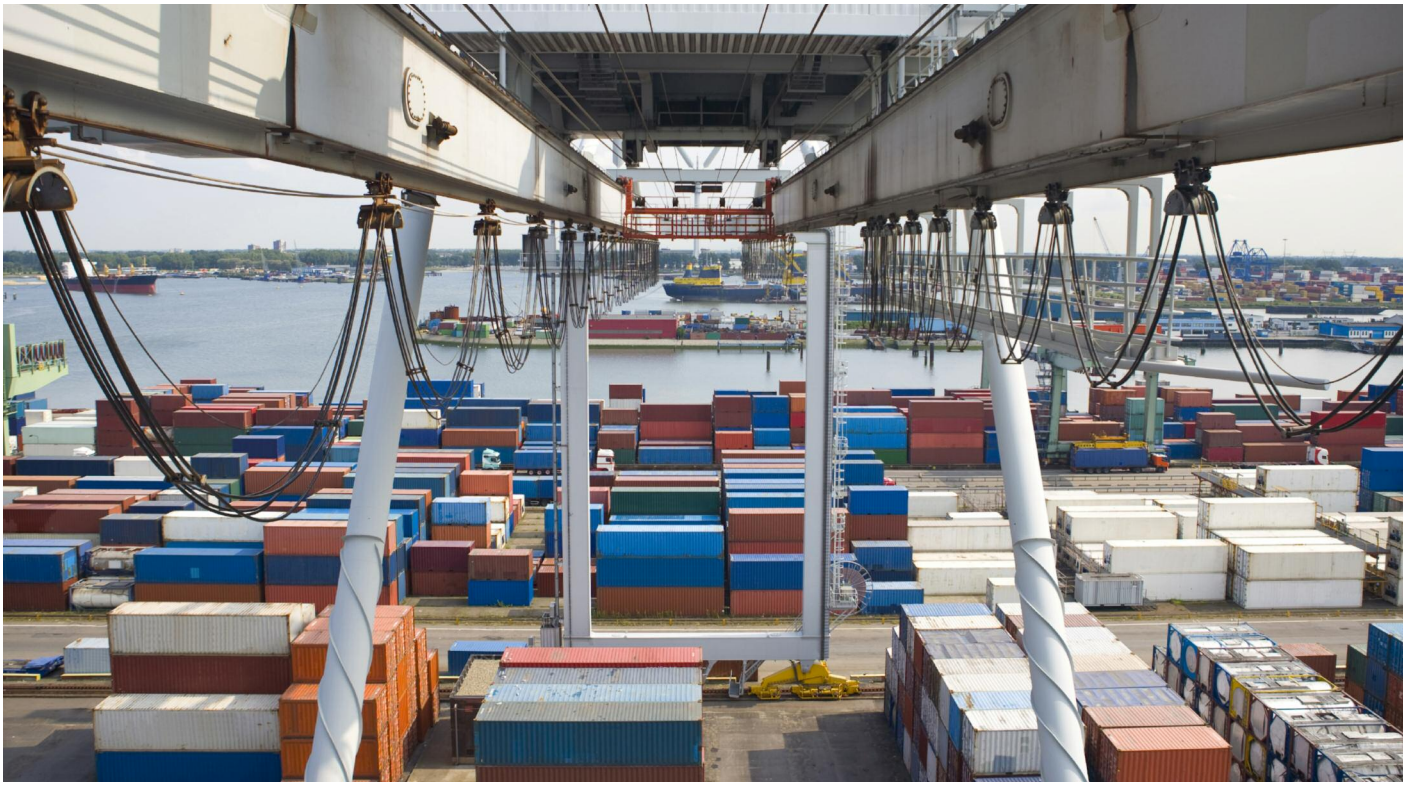

PHOTO: CISCO

Thirdly, videos of faulty machines, pictures of prototypes and digital signage are becoming increasingly common and add up to more, potentially disruptive, load on the network. So technologies will need to help manage loads, prioritise network activities and give a better quality of service. Moreover, as wireless network and fixed LAN convergence increases, manufacturers and their IT departments can use one system for wired and wireless management.

## Case studies

**Aircraft manufacturer efficiency gain** – RFID and wireless solutions can help manufacturers overcome space and time barriers in keeping track of parts and tools. In Boeing's vast jetliner production environment, essential components and tools can easily be misplaced – and must be replaced to avoid production delays. Additional purchases and lost productivity quickly become very costly.

Boeing used its wireless network, RFID tags and location tracking software to help reduce losses. RFID tags were fitted to 1700 critical parts, tools and factory machines. Any part can be located instantly. Because each part has a tag, it can be quickly identified to resolve service issues, saving troubleshooting time and misdiagnosis. This improved productivity by enabling employees to quickly locate parts and tools, which also helped reduce production delays and the potential for government fines.

**Big scrap reduction** – Leading architectural glass fabricator Viracon uses an RFID and wireless asset tracking solution with which RFID tags are attached to all of its 5800 work-in-process glass carriers at three manufacturing campuses. Paired with a wireless network, this eliminated the need for a dedicated network of RFID tag readers. When operators need to find a particular glass carrier, they search for it using shop floor computers. The real-time location of the required carrier is presented on a site map, allowing carrier location within seconds and providing complete visibility. This system increased success rates in precisely locating carriers to 99%, reduced reproduction and scrap of lost carriers by 65%, and ensured that the correct glass is transported on time to the right destinations.

**Boosting tyre production** – Continental Tyre of the Americas' Illinois plant produces over 1000 different tyre stock keeping units (SKUs). Demand was increasing beyond capacity, so the company set out to remove bottlenecks and delays in its production process to increase throughput. Inefficient management of its thousands of carriers was causing unnecessary idle time and setups.

The company implemented a wireless network with RFID tags as part of a solution that tracks and manages tyre assembly and material carriers. This system tracks Work-In-Process (WIP) functions without needing a proprietary reader/sensor network. The system tells operators where the freshest rubber and tyre components are, how to get to them using their fork-lift-truck computer's visual map, where they're needed, and even the status and location of empty component and rubber carriers. The latter is important, since components can 'go off' during the curing process; lost WIP can cause wastage. As a result, there are reduced delays and production stoppages, minimised machine idle times, and significantly increased daily throughput. Scrap has been eliminated. Payback is around six months, and component tyre losses decreased by around 20%.

## Tracking the Future

It is clear that RFID technology and wireless networks create a winning solution for a range of asset management, cost reduction, location tracking, and safety initiatives. The beauty of these solutions are their flexibility. Now a manufacturer can identify and manage almost anything to gain better operational visibility and intelligence.

RFID systems can also communicate directly with PLCs via Profinet, ModBus and Ethernet/IP etc., so are ideal for industrial applications that require the exchange of RFID data with ERP systems, as well as with devices such as PLCs.

### References


1. Internet Protocol version 6 (IPv6) is a version of the Internet Protocol (IP) designed to succeed IPv4 as it allows for vastly more addresses.
2. IDC Insights: 10 Manufacturing Industry Predictions for 2011
3. Top Technology Trends You Can't Afford To Ignore, Oct. 5, 2010
4. Web Map Service (WMS) is a standard protocol for serving georeferenced map images over the Internet generated by a map server using data from a geographic information system (GIS) database.



*From the publication* ***Cisco Tracks RFID with Active RFID and Wireless LANs***

First published in the ***industrial ethernet book*** July 2011

# Extended WLAN is a single network

**Administering remote installations centrally using thin access point and wireless controller technologies provides added value for the company and its service provider customers says Wolfgang Bölderl-Ermel**

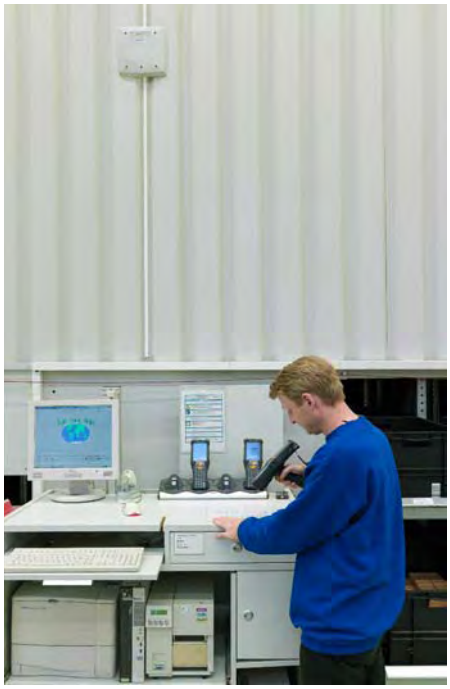

WITH WIRELESS networks becoming more important, the need for integrated wireless network management solutions linking office and production areas is important.

However, wireless components such as wireless access points (WAPs) and clients must be both industry-proof and and usable in the office. They should also be capable of being jointly operated and managed in the same wireless network.

## Computer centre example

The Siemens Erlangen Frauenaurach computer centre is one of two redundant computer centres in a redundant dark fibre backbone that connects several of the company's sites. The centres in this Metropolitan Area Network (MAN) provide technical departments with individual network services. Intranet and Internet access services use a standardised, transparent firewall, information security and server operation processes all in a single package.

Examples include APs and Clients for industrial wireless LAN applications indoors, outdoors or in Zone 2 hazardous areas. Included are up to three wireless modules, internal or external antennas and RJ45 and fibre optic cable connections.

The 'thin access point' Scalance W786 2 HPW operates with, or on, the company's HiPath Wireless Controller. Thin access point solutions place very little 802.11 intelligence on the AP – instead, they pass the 802.11 packets to a centralised controller or switch. The wireless controller acts as a dynamic router that combines and coordinates all connected APs and manages its clients such that the wireless LAN works as an individual, centrally administered IP sub network. This reduces on-site installation costs and the need for distributed AP servicing, and gives expanded central control functions.

This type of WLAN architecture simplifies network topology bridging between offices and the production environment. This is because it allows a universal, integrated wireless network to be used for all APs and clients across company divisions, such as Internet, operating data, automation equipment, VoIP and WLAN capable telephone integration.

Comments Helmut Mack, project manager at the Erlanger computer centre: 'The thin access points can be conveniently administered from a central point. After the initial installation, nobody from our company has to be on site to put extra or replacement equipment into operation, to install software updates or to make changes to the configuration.'

Users can connect the hardware without recourse to trained personnel. The devices then automatically and securely log onto an available controller using a secure, encrypted tunnel connection via a fibre optic backbone. The device is configured and all tasks can be handled controller based from the central office.

Up to 200 controller-capable APs for as many as 4096 simultaneous users can be centrally managed and coordinated. Using the HiPath Wireless Manager, several controllers and hundreds of access points are administratively grouped together. This wireless manager presents the WLAN as several subnetworks, but it also manages connections such that participants can move around within the network. It also keeps an eye on security rules and the defined quality for different devices.

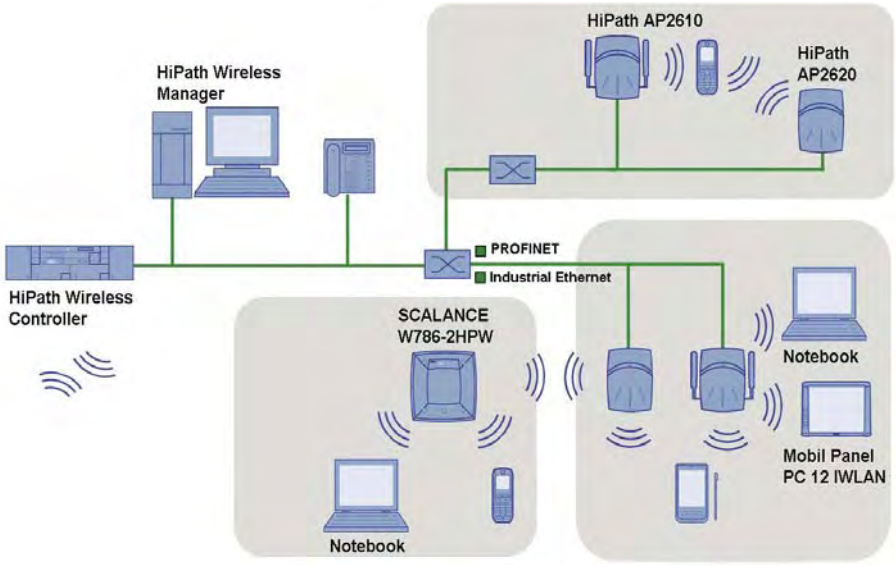


*The wireless controller implements a single network*

## In operation

One of the beneficiaries is Siemens Gerätewerk Erlangen (GWE), which draws its IT services from the computer centre at the same site. Take the reliable picking in GWE's picking warehouse. Employees accept their orders as bar codes with a scanner that wirelessly communicates with higher level order management system.

The industrial wireless LAN is centrally managed from the Erlangen and Karlsruhe computer centres where redundant wireless controllers are installed. The redundant MAN, which uses Coarse Wavelength Division Multiplex (CWDM) laser technology for the large sites, and which allows eight simultaneous Gigabit connections, connects the Nuremberg, Regensburg, Amberg and Chemnitz sites via the 'dark fibre' fibre optic ring. On site technical divisions centrally draw upon tailored network services and don't need their own special expertise or hardware.

The services are based on standardised, cost effective and proven processes and protocols. Currently, around 360 HiPath AP3610 and AP3620 series APs are operating in the MAN, together with over 80 hardened Scalance W thin access points.

***Wolfgang Bölderl-Ermel*** *works for Siemens Industrial Communications*

First published in the ***industrial ethernet book*** November 2010

# Sensor networks: wireless mesh or wireless backbone?

**Automation plant operators should carefully consider their current and future needs before choosing an industrial wireless system. Some applications are well suited to a field device meshing network, while others are better served by an infrastructure meshing network. To gain the maximum benefit that meshing can offer, the selected system should support both topologies simultaneously and seamlessly in a single network. In the following article, Soroush Amidi explains the salient characteristics of each topology to help end users decide which one best serves their needs.**

THE PACE OF ADOPTION for industrial wireless by the very conservative automation and manufacturing industry testifies to the strength of wireless meshing technology. Every year, thousands of plants opt to use wireless devices, such as mobile handhelds running apps for maintenance and process monitoring, video collaboration cameras, asset location tags and wireless field instruments. This improves organisations' operational and capital expenditure performances. The attractiveness of industrial wireless lies in its mobility, flexibility and lower cost.

The use of wireless technology in the process industries is not new. Automation professionals began using wireless transmitters more than a decade ago to collect data from remote areas or equipments where the use of wired transmitters were not feasible, either because of physical or financial constraints.

What is new, however, is the development of wireless meshing technology, which offers the reliability and robustness that was lacking in point-to-point wireless products used at the beginning of the 21st century.

Today, a growing number of industrial end users are implementing wireless devices, including wireless transmitters. In fact, wireless transmitters are even being installed for process monitoring in areas where wired transmitters can be used. Why? Again, it is because of the reliability offered by the wireless meshing technology.

When it comes to implementing a wireless meshing network, automation professionals are faced with choosing from several different topologies. They can implement a field device meshing topology where field devices, typically battery powered wireless field instruments or wired field instruments with wireless adapter, form a peer-to-peer meshing network. Alternatively, they can implement an infrastructure meshing topology where infrastructure nodes, i.e. line-powered industrial access points, form a peer-to-peer meshing network connecting wireless field devices, field instruments and/or Wi-Fi devices.

## Field device meshing topology

Field device meshing enables a wireless peer-to-peer network to form among wireless field instruments. This approach does not require any lined powered wireless infrastructure to be present. Communication packets can hop between transmitters to reach the final destination. Transmitters auto-discover neighbouring transmitters and establish a communication path with each other, thus forming a mesh network.

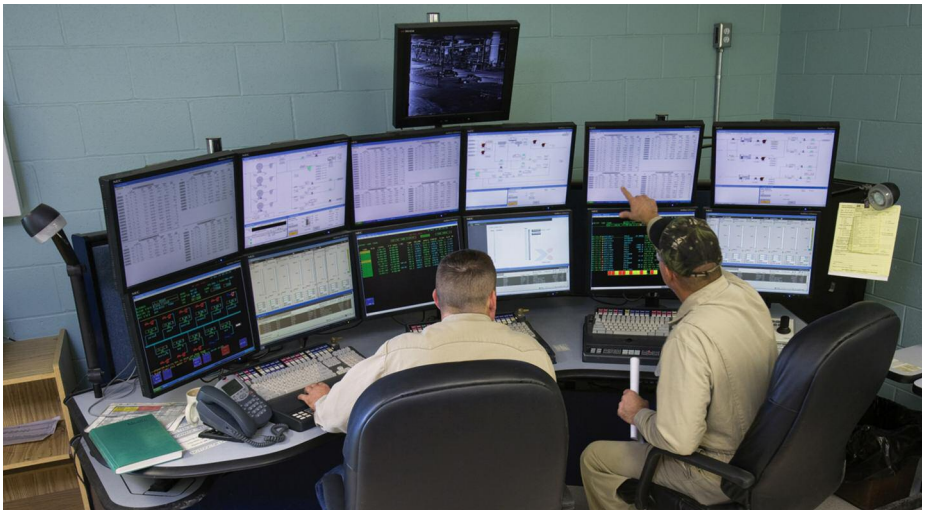

PHOTO: HONEYWELL

***Wireless sensor meshing or wireless sensors with high-speed backbone?*** *Automation plant operators should carefully consider their current and future needs before choosing an industrial wireless system.*

PHOTO: HONEYWELL

***Wireless sender for ISA100.11a:*** *It can operate with an update rate of 1s, but at the expense of battery life.*

Each transmitter acts as an independent router (i.e., the transmitter can send its own data as well as route data received from other transmitters). This allows for continuous connections and reconfiguration around broken or blocked paths by 'hopping' from transmitter to transmitter until the packet reaches the wireless sensor gateway.

## Infrastructure meshing topology

In an infrastructure meshing network, field instruments do not act as routers. Instead, line-powered infrastructure nodes route the data. These infrastructure nodes auto-discover each other and establish a peer-to-peer network. This allows for continuous connections and reconfiguration around broken or blocked paths, again by 'hopping' from node to node.

A Field Device Access Point hosts an ISA100.11a backbone router board (refer to ISA100.11a standards for backbone router definition) in an industrial enclosure. Data can be routed through other ISA100 field devices using the ISA100.11a standard or via a high-speed backbone router such as an IEEE 802.11 WLAN via the Ethernet port hosted on each node.

Multinodes route data through other Multinodes using the IEEE 802.11 standard. They also host a backbone router board connected to a meshing access point board in the same industrial enclosure. They can route ISA100.11a data as well as Wi-Fi data.

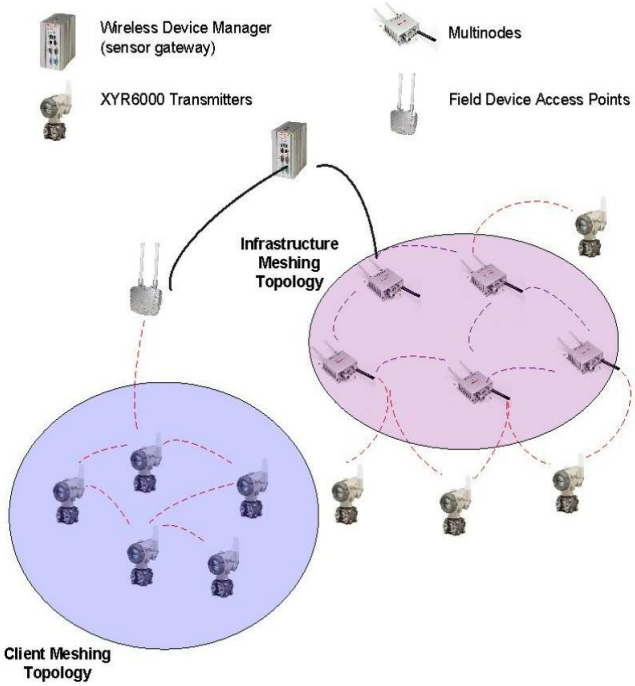


***Fig. 1. Field device meshing and infrastructure meshing:*** *The OneWireless system supports field device meshing and infrastructure meshing topologies in either standalone mode, or in combination.*

## The ISA100.11a standard

So as to realise the maximum benefits offered by meshing, both field device and infrastructure meshing topologies need to be supported simultaneously and seamlessly in a single network. This is one of the core requirements of the ISA100.11a standard for wireless field instruments, which was developed by ISA100's 500-plus members representing the plant automation industry. This technology was designed for use in remote applications with low frequency update rates, as well as in control and critical monitoring applications with high frequency update rates.

ISA100 members ensured that ISA100.11a transmitters could operate with both field device and infrastructure meshing topologies. Figure 1 illustrates how ISA100.11a field instruments operate in either field device meshing or infrastructure meshing mode.

## A comparison

Table 1 (*over page*) shows the characteristics of a field device meshing topology versus that of an infrastructure meshing approach.

Latency. The ISA100.11a standard is intended to optimise the battery life of wireless transmitters. When in field device meshing/routing mode, ISA100 field instruments wake up on a periodic basis to listen to other transmitters and route data based on reporting rates. This capability allows field instruments to save power and maximise their battery life, which effectively translates to lower latency.

In a star meshing topology, routing infrastructure nodes listen and route data in real-time. This is possible since the devices are line-powered. Field instruments simply need to transmit their data to the routers. Network timeslot allocation is much easier in an infrastructure meshing topology.

As just described, there is a significant difference in data latency between a field device meshing network and an infrastructure

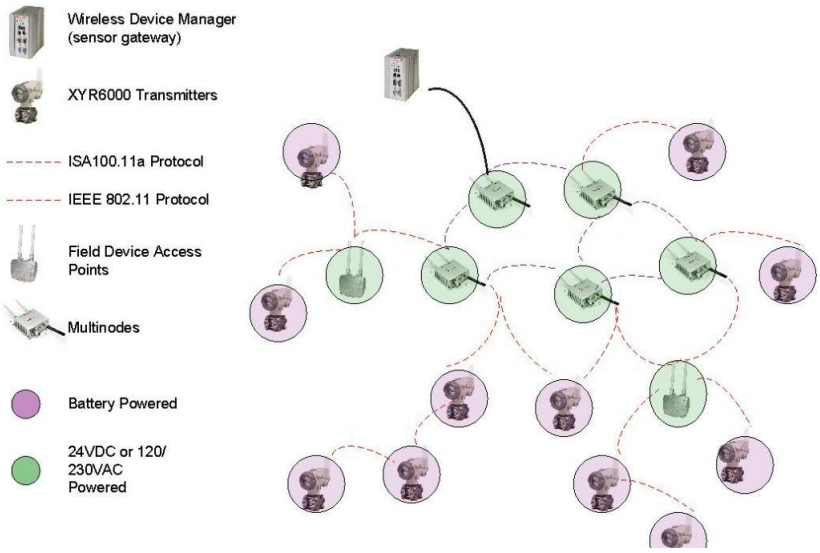


***Fig. 2. An ISA100.11a-compatible flexible network:*** *The infrastructure meshing network will cover most of the site's area and the field device meshing network will connect remote devices to the high-speed infrastructure meshing network via one or two battery-powered routing field devices.*

meshing network. That is why Honeywell recommends field device meshing topologies for applications requiring slow reporting rates from the field, and infrastructure meshing topologies for applications requiring fast reporting rates from the field.

**Battery life**. In a field device meshing topology, routing transmitters have to actively listen and send data to neighbouring transmitters. The battery life of a transmitter is a function of the number of transmitters meshed with the device. It decreases as the number of meshed transmitters increases. In an infrastructure topology, transmitters send data to routing infrastructure devices. This translates into a significant difference between the battery life of a transmitter acting as a routing node in a field device meshing topology, and a transmitter in an infrastructure meshing topology.

Honeywell recommends field device meshing topologies for applications requiring slow reporting rates, and for plants not planning to scale up to tens of transmitters. Infrastructure meshing topologies are suitable for applications requiring fast reporting rates and millisecond latency, and for plants planning to scale up to hundreds of transmitters.

**Reporting rate**. In order to avoid frequent replacement of batteries, routing transmitters are only used for applications requiring slow reporting rates. A non-routing transmitter in an infrastructure meshing topology can be configured at the highest reporting rate. Some transmitters can report as fast as a 1-second update rate.

Honeywell recommends field device meshing topologies with ISA100.11a field instruments for users with applications requiring slow update rates (30 seconds or more). For applications requiring fast update rates, it recommends an infrastructure meshing topology with either Multinodes or Field Device Access Points (see box).

**Scalability**. Routing transmitters can route a limited number of field instruments. This limitation is primarily because of the 802.15.4 radio. This has a 250 kps throughput, which is just sufficient to route data from a handful of transmitters. The other limitation is because of the components selected for use with wireless field instruments. Power consumption and management are key requirements driven by customers' need to have wireless field instruments with multi-year (10 years) battery life.

The components are selected based on their power consumption and power management capability. This typically means a small memory and computing footprint, which differs from line-powered routing infrastructure devices where power consumption and power management are not an issue. Wireless field instruments can route up to four transmitters efficiently, while routing infrastructure nodes can route up to 80 transmitters.

**Flexibility**. Field device meshing topology limits users to an ISA100.11a field instruments-based application. An infrastructure meshing topology having industrial meshing access points designed to extend the process control network into the field, can support ISA100.11a field instruments and Wi-Fi devices, and therefore enable a multitude of applications. Such devices (Multinodes) can simultaneously route data between Wi-Fi (IEEE 802.11 b/g) clients, ISA100.11a field instruments and Ethernet/IP devices and host applications.

It is possible to start small with a field device meshing network and scale up to an infrastructure meshing network. Most installations will be a combination of infrastructure meshing and field device meshing networks. The infrastructure meshing network will cover most of the site's area and the field device meshing network will be used to connect remote devices to the high-speed infrastructure meshing network via one or two battery-powered routing field devices. Figure 2 shows the Honeywell implementation of an ISA100.11a-compatible network.

## And finally...

Automation professionals should carefully consider their current and future needs prior to selecting an industrial wireless system. Certain applications are ideal for infrastructure meshing networks, while others are more suited to field device meshing networks. To gain the maximum benefit that meshing can offer, the selected system should support both topologies simultaneously and seamlessly in a single network.

Current market trends indicate that most installed wireless systems will use infrastructure meshing complemented by field device meshing when infrastructure meshing is installed at the core of the plant where power is readily available. Field device meshing is used in remote areas of the plant where power is not readily available.

| | Field Device Meshing | Infrastructure Meshing |
|---|---|---|
| **Latency** | 100ms/hop | 20ms/hop |
| **Fastest recommended reporting rate for field instruments** | 30 seconds | 1 second |
| **Routing capability** | Four transmitters at 30 seconds | 20 transmitters at 1-second or 80 transmitters at 5 seconds |
| **Field instrument battery life** | <1 year at 1s update rate when routing data from other transmitters | 3 years at 1s update rate, 10 years at 10s update rate |
| **Scalability** | 100 transmitters | 4000 transmitters |
| **Flexibility** | ISA100.11a transmitters only | ISA100.11a transmitters, Wi-Fi devices and IP/Ethernet-based devices |

***Table 1:*** *Characteristics of a field device meshing topology versus an infrastructure meshing topology*

***Soroush Amidi*** *works for Honeywell Process Solutions.*

First published in the *industrial ethernet book* November 2011

# Internet Protocol for wireless connected Smart Objects

**The emerging application space for smart objects requires scalable and interoperable communication mechanisms that support future innovation as the application space grows. IP has proven itself a long lived, stable, and highly scalable communication technology that supports both a wide range of applications, devices, and underlying communication technologies. The IP stack is lightweight and runs on tiny, battery operated embedded devices. IP therefore has all the qualities to make 'The Internet of Things' a reality, connecting billions of communicating devices. Internet Protocol for Smart Objects (IPSO) Alliance**

SMART OBJECTS are small computers with a sensor or actuator and a communication device, embedded in objects such as thermometers, car engines, light switches, and industry machinery. Smart objects enable a wide range of applications in areas such as home automation, building automation, factory monitoring, smart cities, health management systems, smart grid and energy management, and transportation.

Until recently, smart objects were realised with limited communication capabilities, such as RFID tags, but the new generation of devices has bidirectional wireless communication and sensors that provide real time data such as temperature, pressure, vibrations, and energy measurement.

Smart objects can be battery operated, but not always, and typically have three components: a CPU (8 , 16 or 32 bit microcontroller), memory (a few tens of kilobytes) and a low power wireless communication device (from a few kilobits/s to a few hundreds of kilobits/s). The size is small and the price is low: a few square millimetres and few dollars. The technical development in low cost sensors and actuators combined with low power communication technologies such as IEEE 802.15.4, low power Wi-Fi, and power line communication has been rapid. Nevertheless, the emergence of smart object applications has not been as fast because the large number of proprietary or semi closed systems has lead to partial and non-interoperable solutions.

The current situation for smart objects is similar to what computer networks looked like about two decades ago: islands of computers communicating with their own protocol, for example SNA, IPX, and Vines, interconnected by complex multi-protocol gateways. Subsequently, these architectures evolved to IP based tunnelling mechanisms such as DLSw or XOT. Today, these networks operate on fully end-to-end IP based architectures.

Many of today's non IP based sensor architectures are evolving toward a protocol translation gateway model, similar to the path computer networks went through before quickly moving to fully IP based architectures. Have we not learnt from the past? Protocol gateways are inherently complex to design, manage, and deploy. The network fragmentation leads to non efficient networks because of inconsistent routing, QoS, transport and network recovery techniques. end-to-end IP architectures are widely accepted as the only alternative to design scalable and efficient networks of large numbers of communicating devices.

## The Internet of Things

To support the large number of emerging applications for smart objects, the underlying networking technology must be inherently scalable, interoperable, and have a solid standardisation base to support future innovation as the application space grows.

IP has proven itself a long lived, stable, and highly scalable communication technology that supports both a wide range of application, a wide range of devices, and a wide range of underlying communication technologies. The layered architecture of IP provides a high level of flexibility and innovation. IP already supports a plethora of applications, such as email, the World Wide Web, Internet telephony, video streaming, and collaborative tools. Over the past 20 years, IP has evolved to support new mechanisms for high availability, enhanced security, support of Quality of Service (QoS), real time transport, and Virtual Private Networks (VPNs).

IP has a long history as a communication mechanism for general purpose PC computers and network servers. It was therefore long believed that IP was too heavy weight to run on highly constrained devices. Several recent lightweight IP stacks have demonstrated that they can be designed to meet the requirements of light footprint devices with a few kilobytes of RAM and ROM, limited processing power and severe energy constraints.

IP provides standardised, lightweight, and platform independent network access to smart objects and other embedded networked devices. The use of IP makes devices accessible from anywhere and from anything; general purpose PC computers, cell phones, PDAs as well as database servers and other automated equipment such as a temperature sensor or a light bulb.

IP runs over virtually any underlying communication technology, ranging from high speed wired Ethernet links to low power 802.15.4 radios and 802.11 (Wi-Fi) equipment. For long haul communication, IP data is readily transported through encrypted channels over the global Internet.

Memory efficient implementations of the IP stack show that IP can successfully work in as little as a few kilobytes of RAM, and require less than 10KB of ROM. Figure 1 shows the memory footprint of five embedded TCP/IP stacks: the open source µIP stack from the Contiki operating system, one commercially available TinyOS based IPv6 stack, the commercially available NanoStack, and the open source lwIP stack. Their footprint is around 10KB, except for lwIP that is around 20KB.

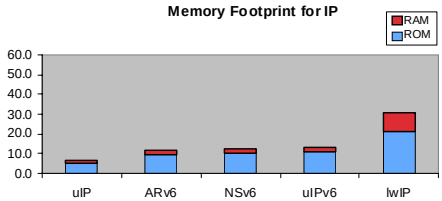


***Fig. 1.*** *Memory footprint for five embedded TCP/IP stacks*

For power constrained devices, recent standardisation work has made IP power efficient enough to run over sub-milliwatt radio links such as 802.15.4. Such low power operation enables years of lifetime on typical AA batteries, even for multi-hop routing nodes.

*IP is Scalable.* With the global Internet, IP has proven itself to be inherently scalable. No other networking technology has ever been deployed and tested at such an immense scale and with such a large number of devices. As smart objects will connect an even larger

number of devices than that of the existing Internet, scalability is a primary concern.

The next generation Internet protocol, IPv6, expands the address space of IP to $2^{128}$. Such a large address space has been said to be enough to provide every grain of sand on the planet with an IP address.

*IP is End-to-End*. IP provides end-to-end communication between devices, without intermediate protocol translation gateways. Protocol gateways are inherently complex to design, manage, and deploy. The objective of a gateway is to translate or map between two or more protocols. Such translation, however, typically requires significant semantic and functional translation for the protocols to work together. Mechanisms on both sides usually differ significantly, thus requiring the adoption of a least common denominator approach that leads to inefficient networks because of inconsistent routing, QoS, transport and network recovery techniques. With the end-to-end architecture of IP, there are no protocol translation gateways involved.

With the IP end-to-end architecture, there is no single point of failure. Intermediate routers may fail, but the end-to-end communication will chose alternate paths through the network. In contrast, if a protocol translation gateway fails, the entire network fails.

In the IP architecture, protocols can change without affecting the underlying network. Routers operate independently of the protocols running over them. In contrast, a protocol translation gateway needs to be updated every time a protocol changes, no matter how small the change.

With the success of today's global Internet, the end-to-end architecture of IP has proven itself scalable, stable, and efficient. For the future Internet of Things, scalability, stability, and efficiency is even more important than ever. IP therefore is the future proof choice for the Internet of Things.

Smart objects enable a wide range of applications that will improve our lives in many areas such as energy management, healthcare, and safety. The recent progress in low cost embedded devices is about to make the Internet of Things a reality. For this to come true, we must learn from the lessons of the past and adopt a flexible, scalable, efficient and open based networking technology. IP has proven itself to fulfil these requirements and it is now a fact that IP can meet the strict requirements of highly constrained smart object networks.

***Adam Dunkels Ph.D*** *is a senior scientist, Swedish Institute of Computer Science*

***JP Vasseur*** *is an engineer with Cisco Systems*

*From the IPSO Alliance paper* ***IP for Smart Objects***

# Lightweight IPv6 Stacks for Smart Objects

**Historically, smart objects have used a plethora of communication technologies, both at the physical and medium access control layers, and at upper layers.**

The upper layers of the communication stack remain either proprietary or specified by exclusive alliances. This plethora of solutions renders interoperability between different sensor networks difficult. It also makes the seamless integration of sensor networks with existing IP networks impossible. IP is an ideal solution to this end-to-end interoperability issue. However, the adoption of IP as the Layer 3 protocol to connect wireless or wired smart objects has been impaired by the common belief that IP is not well suited for the memory and energy constraints of such devices.

## Efficient IPv6 stacks

Three independent implementations of a lightweight IPv6 stack for smart objects have already been developed. The data structures necessary to implement a minimal IPv6 stack are an interface address list, a neighbour cache, a prefix list, a routing table, and a default router list as defined in RFC4861. These structures are updated whenever packets are sent or received. Moreover, timers are used to manage entries that have a lifetime associated with them.

The IPv6 specification requires the link to support a 1280 byte MTU, imposing significant challenges in supporting IPv6 on memory constrained devices. A minimal implementation can use a single 1280 byte buffer for the basic data plane. However, the current Neighbour Discovery (ND) specification (RFC4861) requires a packet buffer for each neighbour on the link and is prohibitively expensive for smart objects.

Wireless mesh networks (e.g. IEEE 802.15.4) often span multiple hops and require some or all devices to forward datagrams (and thus act as routers) to provide access availability throughout the network. As a result, the store and forward technique requires nodes to buffer additional packets that are being forwarded.

One way to implement the forwarding queue is to allocate full packet buffers for each entry in the queue. More sophisticated approaches support variable sized buffers that allow more effective use of the memory especially when datagrams do not consume the full 1280 byte MTU.

## Layer 2 dependencies

The layered IP protocol architecture allows the network layer to remain agnostic to the specifics of the link layer below. IPv6 stacks implemented for smart objects do not have to give up the layered protocol architecture in order to fit within the severe resource constraints typical to smart objects.

In wired networks, the network layer only expects the link layer to provide services necessary for delivering packets. Networks operating on constrained (e.g. low throughput) links, however, can benefit from visibility between layers. On wireless links, for example, knowledge of when packets are acknowledged and some indication of signal strength can provide invaluable information to link estimators that are used by IP based routing protocols.

Similarly, routing protocols could indicate which neighbours are most important to optimise link layer transmissions by providing visibility into the routing table. Furthermore, the link layer could deliver packets more efficiently if it had visibility into the forwarding queue. For example, the link could optimise transmission of multiple frames to the same destination and the link could appropriately schedule transmissions to minimise average queuing delay. By providing this extra visibility in a link independent manner, we can significantly enhance the performance of an IP based communication architecture for smart objects while maintaining the layering.

## Optimisation for constrained links

Link technologies used in Smart Objects are often highly constrained in terms of throughput, channel capacity, and MTU. Energy constraints can also limit how much information can be communicated. We describe a number of optimisations to enable more efficient operation over constrained links. While the optimisations described below were originally designed for low power wireless networks, they equally apply to low speed wired links as well.

**Compressing IPv6 headers.** IPv6 has a base header of 40 bytes, which is fairly large especially when compared to the small link layer MTUs and application payloads typical to smart objects. As a result, compressing IPv6 headers can help save memory in buffering those datagrams as well as reduce energy costs for delivering datagrams. With RFC 4944, the IETF has defined the 6LoWPAN adaptation layer that includes a compression mechanism. This mechanism is stateless which means that it creates no binding state between the compressor-decompressor pair.

Stateless compression gives nodes the necessary flexibility to communicate with any neighbour in compact form at all times. While the current specification only supports link local communication, the IETF is currently working on mechanisms to support stateful compression for global communication. The stateful compression mechanisms rely on shared context throughout the entire network, allowing state synchronisation to remain simple and nodes to communicate with any of their neighbours in compressed form.

Many link technologies designed for smart objects do not support the full 1280 byte MTU.

For several reasons, including cost reduction, guaranteeing an acceptable Packet Error Rate (PER), and microcontroller buffering constraints, IEEE 802.15.4 only supports a 127 byte MTU. To support the IPv6 1280 byte MTU, datagrams must thus be fragmented before they can be delivered to the link layer. With RFC4944, the IETF has defined the 6LoWPAN adaptation layer to support fragmentation.

Fragmented datagrams also provide additional opportunity for more effective buffering techniques when forwarding datagrams. For example, fragments can be delivered before a node reassembles the entire datagram, allowing nodes with severe memory constraints to forward complete datagrams.

## Extending Neighbour Discovery

IPv6 ND is currently defined for operation only on links that support a single broadcast domain. As such, many of the ND primitives rely on multicast to discover and communicate with neighbouring nodes on the same link. For this reason, IPv6 Neighbour Discovery (ND) does not map well to smart objects connected by a wireless network that spans multiple hops.

One option involves emulating a single broadcast domain over the entire wireless network, allowing support for IPv6 ND as specified today. However, for networks that reach any reasonable scale, a simple multicast is prohibitively expensive as the message must be delivered to all nodes within the network.

An alternative solution is to concentrate ND operations at routers that serve as egress points for the network. By relying on egress routers, ND can now use unicast communication which is much more efficient than multicast in wireless networks. The use of edge routers to support ND is currently being specified within the IETF.

An IPv6 Ready protocol stack can be implemented in approximately 11.5KB of ROM and 1.8KB of RAM. To support the 1280 byte minimum MTU, a single packet buffer requires 1280 bytes of RAM.

Of the remaining data structures, the neighbour cache requires 35 bytes per neighbour, the prefix list requires 23 bytes per prefix, the router list requires seven bytes per router, and the interface address list requires approximately 100 bytes.

Table 1 provides a breakdown of the typical memory requirements for individual components. (Fragmentation and per neighbour buffering are not included). Neighbour Discovery consumes the largest portion of the complexity.

Together with a complete run time (timers, scheduler, etc.) as well as RFC compliant UDP and TCP protocols above, an OS that provides a complete IPv6 network stack can be implemented within 35KB of ROM and 3KB of RAM. Complete IPv6 based applications fit comfortably within a microcontroller providing only 64KB of ROM and 4 to 8KB of RAM.

| Function | ROM | RAM |
|---|---|---|
| ND Input/Output | 4800 | 20 |
| ND structures | 2128 | 238 |
| Network interface management | 1348 | 118 |
| Stateless address autoconfiguration | 372 | 16 |
| IPv6 (headers processing, etc) | 1434 | 44 |
| Packet buffer (Ethernet case) | 0 | 1296 |
| ICMPv6 | 1406 | 16 |
| Total | 11488 | 1748 |

***Table 1.*** *Typical code and memory requirements for a complete IPv6 protocol stack*

## Energy requirements

Battery powered smart objects are constrained not only in memory, but in energy as well. With appropriate link layer mechanisms and experimental modifications to IPv6 (e.g. compression and ND optimisations), it is possible to implement an IPv6 network stack that consumes very little energy. In production deployments, an IPv6/802.15.4 network can operate at <1% duty cycles (including forwarding nodes), with very low per hop latency and high reliability under realistic work loads.

Table 2 presents network statistics from an IPv6-based home monitoring application. The application consists of 15 nodes deployed in refrigerators, solar power inverters, outdoors, and indoors in or near the intended sense point. Nodes periodically report environmental data (e.g., temperature and humidity) as well as a variety of sensors attached to each node. Application traffic was about one datagram per minute per node, each datagram nearly filling a full 802.15.4 frame. The routing topology consisted of 7 nodes within 1 hop of the egress router, the remaining half being 2 to 3 hops away.

| | |
|---|---|
| Application Data Rate | 1 UDP datagram/minute |
| Average Radio Duty Cycle | 0.65% |
| Expected Per-Hop Latency | 0.125 seconds |
| Datagram Delivery Rates | 99.98% |

***Table 2.*** *Network statistics for a production home-monitoring application*

This work demonstrates that even the most memory constrained devices can be IPv6 Ready. We have shown that it is also possible to implement an IPv6 based network on energy constrained devices and are currently specifying optimisations (e.g. IPv6 header compression and IPv6 ND modifications) to reduce the energy requirements of supporting a complete IPv6 network stack. These efforts provide a framework that will enable further advancement of IPv6 support for smart objects. In addition to routing, support for end-to-end IPv6 based security mechanisms (e.g. IPsec, SSL, or TLS) still require some evaluation.

***Julien Abeillé**, software engineer, Cisco Systems*

***Mathilde Durvy Ph.D**, software Engineer, Cisco Systems*

***Stephen Dawson Haggerty**, Berkeley, University of California,*

***Jonathan Hui Ph.D**, Arch Rock Corporation*

*From the IPSO paper* ***Lightweight IPv6 Stacks for Smart Objects***

# 6LoWPAN: Incorporating IEEE 802.15.4 into IP architecture

**IP for Smart Objects seeks to extend the use of IP networking into resource constrained devices over a wide range of low power link technologies – IEEE 802.15.4 represents one such link.**

Extending IP to low power, wireless personal area networks (LoWPANs) was once considered impractical because these networks are highly constrained and must operate unattended for multiyear lifetimes on modest batteries. Many vendors embraced proprietary protocols, assuming that IP was too resource intensive to be scaled down to operate on the microcontrollers and low power wireless links used in LoWPAN settings. However, 6LoWPAN radically alters the calculation by introducing an adaptation layer that enables efficient IPv6 communication over IEEE 802.15.4 LoWPAN links.

Several leading radio manufacturers have implemented IEEE 802.15.4, which specifies a wireless link for low power personal area networks (LoWPANs). 802.15.4 is widely used in embedded applications, such as environmental monitoring to improve agricultural yields, structural monitoring to track building and bridge integrity, industrial control to provide more sense points and control points at lower cost. These applications generally require numerous low cost nodes communicating over multiple hops to cover a large geographical area, and they must operate unattended for years on modest batteries. Such requirements target a very different set of applications than do WPAN technologies such as Bluetooth, which eliminate wiring for headsets, game controllers, and personal devices. Accordingly, 802.15.4's capabilities are more limited than other WPANs and WLANs – they have small frame sizes, low bandwidth, and low transmit power.

Additionally, the microcontrollers typically coupled with LoWPAN radios have limited memory and processing power. These constraints led many LoWPAN vendors to embrace proprietary protocols and link only solutions (such as ZigBee), presuming that IP was too memory and bandwidth intensive for them to scale it down as necessary.

While not following the IP standard, many of these technologies still have not proven their effectiveness in constrained environments. 6LoWPAN radically alters the landscape by introducing an adaptation layer between the IP stack's link and network layers to enable efficient transmission of IPv6 datagrams over 802.15.4 links, dramatically reducing IP overhead.

The adaptation layer is an IETF proposed standard and provides header compression to reduce transmission overhead, fragmentation to support the IPv6 minimum MTU requirement,

and support for layer two forwarding to deliver an IPv6 datagram over multiple radio hops.

6LoWPAN achieves low overhead by apply ing cross-layer optimisations; it uses information in the link and adaptation layers to compress network and transport layer headers. Drawing on IPv6 extension headers, it employs the header stacking principle to separate the orthogonal concepts and keep the header small and easy to parse.

## IPv6 over IEEE 802.15.4

The IPv6 protocol is designed as the successor to IPv4 and enables the Internet to scale for decades to come. To overcome dwindling unallocated address space – and in anticipation that networked appliances and instruments will vastly outnumber conventional computer hosts – IPv6 expands the IP address space from 32 to 128 bits.

Recognising the growth in link bandwidth, IPv6 increases the minimum MTU requirement from 576 to 1280 bytes. To simplify routers and increase performance, IPv6 implements fragmentation at the endpoints, rather than in intermediate routers. To increase protocol efficiency and eliminate the need for ad hoc link level services to bootstrap a subnet, IPv6 includes scoped multicast as an integral part of its architecture.

Core IPv6 components, such as Neighbour Discovery (ND), use link local scoped multicast for address resolution, duplicate address detection (DAD), and router discovery. Stateless address auto-configuration (SAA) simplifies configuration and management of IPv6 devices by enabling nodes to assign themselves meaningful addresses. IPv6 also reflects the advances in link technologies the Internet uses. Ethernet has prevailed as the dominant link, and its throughput has increased at an extraordinary rate.

Current WLAN technologies such as IEEE 802.11a/b/g mirror Ethernet capabilities by supporting similarly sized MTUs and high link rates. Both links operate in the context of ample power and highly capable devices. WPAN technologies, on the other hand, operate with lower power. IEEE 802.15.4 was designed specifically for long lived application domains that require numerous low cost nodes, and these constraints limit the capability of LoWPAN links and the microcontrollers to which they're attached. Throughput is limited to 250 kbps. The frame length is limited to 128 bytes to ensure reasonably low packet error rates when bit error rates are non negligible and reflects microcontrollers' limited buffering capabilities.

802.15.4 defines short 16-bit link addresses, in addition to IEEE EUI 64 addresses, to reduce header overhead and memory requirements. Communication range is short (tens of metres) because transmission power increases polynomially with range. Unlike most typical WPAN and WLAN installations, LoWPANs communicate over multiple hops. Finally, the associated microcontrollers typically have about 8KB of data RAM and 64KB of program ROM.

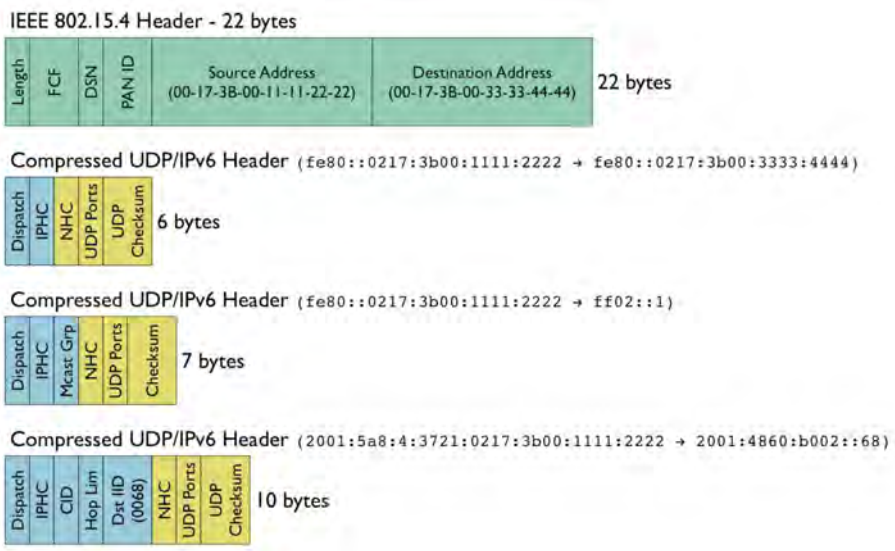


***Fig. 1.*** *6LoWPAN Improved Header Compression examples*

Due to these resource constraints and LoWPAN's multihop nature, supporting IPv6 over LoWPAN networks presents several challenges. First, IPv6 datagrams aren't a natural fit for LoWPANs. Low throughput, limited buffering, and frames that are one tenth the size of the IPv6 minimum MTU requirement make datagram fragmentation and compression a necessity for efficient operation.

For example, link headers can limit effective link payload to 81 bytes, making the IPv6 (40 bytes), UDP (eight bytes), and TCP (20 bytes) headers seem exceedingly large.

Second, because 802.15.4 is both low power and low throughput, it's more prone to spurious interference, link failures, dynamic link qualities, and asymmetric links. Such characteristics require the network layer to be responsive and adaptive while remaining energy efficient, and they affect all aspects of networking, including fragmentation, compression, forwarding, and routing.

Third, a LoWPAN's expected topology is a mesh of short range connections. This negates the assumption that the link is a single broadcast domain on which a core of IP architectural components – such as IPv6 ND and SAA – relies. The IETF 6LoWPAN working group addressed these issues with RFC 4944.

## 6LoWPAN Adaptation Layer

The 6LoWPAN format defines how IPv6 communication is carried in 802.15.4 frames and specifies the adaptation layer's key elements. 6LoWPAN has three primary elements:

- Header compression. IPv6 header fields are compressed by assuming usage of common values. Header fields are elided from a packet when the adaptation layer can derive them from link level information carried in the 802.15.4 frame or based on simple assumptions of shared context.
- Fragmentation. IPv6 packets are fragmented into multiple link level frames to accommodate the IPv6 minimum MTU requirement.
- Layer 2 forwarding. To support layer two forwarding of IPv6 datagrams, the adaptation layer can carry link level addresses for the ends of an IP hop. Alternatively, the IP stack might accomplish intra PAN routing via Layer 3 forwarding, in which each 802.15.4 radio hop is an IP hop.

The key concept applied throughout the 6LoWPAN adaptation layer is the use of stateless or shared context compression to elide adaptation, network and transport layer header fields – compressing all three layers down to a few bytes, combined.

We can see the possibility of compressing header fields to a few bits when we observe that they often carry common values, reserving an escape value for when less common ones appear. Common values occur due to frequent use of a subset of IPv6 functionality (such as UDP, TCP, and ICMPv6 as Next Header values) and simple assumptions of shared context (for example, a common network prefix assigned to the entire LoWPAN).

6LoWPAN also removes redundant header information across protocol layers (for instance, UDP and IPv6 length fields and IPv6 addresses are derived from lower layer headers).

Traditional IP header compression techniques are stateful and generally focus on optimising individual flows over a highly constrained link. These methods assume that the compressor and decompressor are in direct and exclusive communication and compress both network and transport layer headers together. They optimise for long-lived flows by exploiting redundancies across packets within a flow over time, requiring the endpoints to initially send packets uncompressed.

Flow-based compression techniques are poorly suited for LoWPANs. Traffic in many LoWPAN applications is driven by infrequent readings or notifications, rather than long lived flows. Communication over multiple hops requires hop by hop compression and decompression and per-flow state at each intermediate node. Many LoWPAN routing protocols obtain receiver diversity via rerouting, which would require state migration and reduce compression effectiveness. By contrast, stateless and shared context compression in 6LoWPAN doesn't require any per flow state and lets routing protocols dynamically choose routes without affecting compression efficiency.

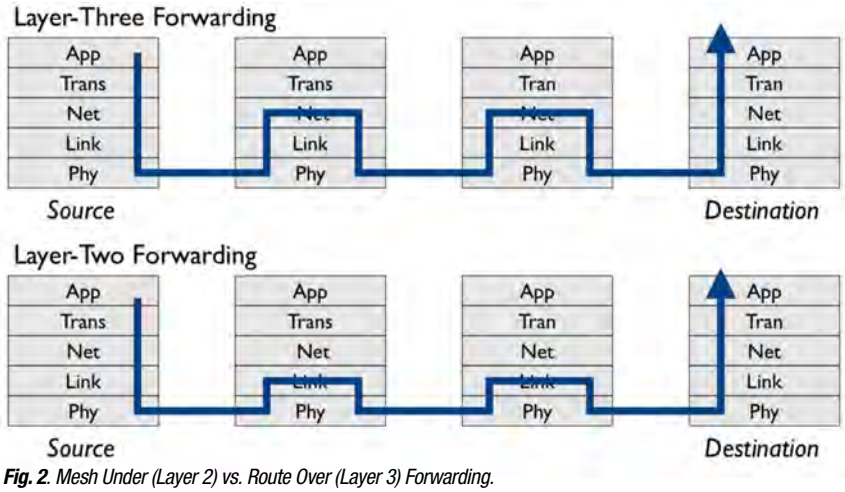


***Fig. 2.*** *Mesh Under (Layer 2) vs. Route Over (Layer 3) Forwarding.*

## UDP/IPv6 header compression

Typical header configurations using IPHC and NHC are shown in Fig. 1. As with RFC 4944, the best case compression efficiency occurs with link local unicast communication – IPHC and NHC can compress a UDP/IPv6 header down to six bytes. The Version, Traffic Class, Flow Label, Payload Length, Next Header, Hop Limit, and link local prefixes for the IPv6 Source and Destination addresses are all removed. The suffix for both IPv6 source and destination addresses are derived from the IEEE 802.15.4 header.

The improvements made in Improved Header Compression (IPHC) become obvious with multicast communication and global communication. As shown in Fig. 1, IPHC can compress communication to well known multicast addresses down to seven bytes (vs. 23 bytes with HC1). Well known multicast addresses have limited their group IDs to only the bottom few bytes, and IPHC takes advantage of this property. When communicating with global addresses, IPHC can compress a UDP/IPv6 header down to 9 or 10 bytes (vs. 31 bytes with HC1). Using context based compression, the prefix of both addresses can be compressed. The Source Address IID may be compressed if it is derivable from the IEEE 802.15.4 header. Finally, unlike RFC 4944, IPHC need only carry two bytes of the IID inline if the upper six bytes are all zeros.

## IPv6/6LoWPAN architecture

The 6LoWPAN format specification defines how fragmentation, compression, and Layer 2 forwarding are represented in an 802.15.4 frame. However, the implementation of those capabilities is out of that document's scope. 6LoWPAN's dependencies on the specific operations defined in the 802.15.4 MAC are minimal, supporting essentially any MAC protocol that provides the 802.15.4 frame format.

Similarly, the 6LoWPAN format doesn't specify how IPv6 capabilities, such as ND and SAA, are orchestrated to configure the LoWPAN to be consistent with the adaptation layer.

## IEEE 802.15.4 in practice

802.15.4 presents several pragmatic issues that have significant architectural impact beyond the 6LoWPAN adaptation layer. Whereas in conventional WPAN settings, the user typically adjusts device and host placement so that the link between them is adequate, in typical LoWPAN settings a network of many devices is embedded in a physical environment at particular, meaningful locations. Network protocols must deal with the many exigencies that arise. Multi hop routing extends range and helps avoid obstacles. Thus, a LoWPAN network isn't typically a single broadcast domain.

Moreover, the link quality between any node pair is often complex and time varying due to environmental factors. Hop by hop retransmission schemes help make lossy 802.15.4 links viable for multihop communication, but alone they aren't sufficient.

Links that are reasonably good on average – say, with 90% packet reception reliability – will often experience bursts of loss due to changes in the noise floor and spurious interference. Routing can overcome such bursts when forwarding datagrams by selecting an alternate path. In effect, routing can exploit receiver diversity by dynamically selecting from multiple next hop candidates. To deal with these link challenges, the network layer requires extra visibility into detailed link behaviour to build and maintain effective routing structures.

Many LoWPAN applications have significant device mobility within the LoWPAN, giving rise to time varying connectivity relationships, in addition to variations induced by changing environmental factors. For example, package tracking might involve numerous devices moving among a set of stationary ones. This isn't IP mobility in the traditional sense because nodes might remain in close physical proximity and be connected within the LoWPAN. However, such variations require that the routing topology adapt to connectivity changes.

The 802.15.4 specification defines only a limited set of power management mechanisms for edge devices and no power management for forwarding devices. Consequently, most commercial implementations and industrial standards built on 802.15.4 forego the defined power management mechanisms when defining routing protocols. To conserve energy, nodes must duty cycle the radio, but doing so requires both transmitter and receiver to coordinate when and how to communicate. Common mechanisms for this involve sampled listening techniques, in which the receiver periodically listens for lengthened transmissions, or scheduling techniques, which involve time synchronisation between nodes. 6LoWPAN, so far, avoids requiring particular MAC features.

## Mesh Under vs Route Over

Two important architectural issues for IPv6 over LoWPAN are concerned with the method by which link level factors inform routing, and at what layer datagram forwarding occurs within the LoWPAN.

Traditionally, IP routing occurs at the network layer in a manner largely independent from the underlying links that implement the individual hops. 6LoWPAN, in its role as an adaptation between the link (layer two) and the network (layer three), can support routing at either layer. Figure 2 shows the difference in packet processing between the two approaches.

## Related work in low power WPANs

Support for IP in resource-constrained environments has a long history, including over telephone modems that gave rise to PPP, DHCP for auto configuration, and header compression.

6LoWPAN differs in the way that it exploits shared context, frequently occurring simple cases, and cross-layer redundancy to reduce vastly header size when communicating over a dynamic, multihop topology.

It builds on prior work with stateless IP header compression. Many efforts have addressed links in which multihop forwarding is required, including frame relay and Asynchronous Transfer Mode. 6LoWPAN is

unique in that it also addresses severe resource constraints.

IEEE 802.15.1 (Bluetooth) is another wireless link technology that falls under the WPAN classification. Intended to serve as a cable replacement technology, Bluetooth supports relatively high throughput for a limited number of nodes within a small range. IEEE 802.15.3 pushes WPAN capabilities further, with greater throughput and support for more nodes. Although both are intended for battery operation, they only target lifetimes of several days to several weeks.

In contrast, 802.15.4 is intended for low data rate applications in which numerous nodes must be low cost and have multi-year lifetime on modest batteries.

The 802.15.4 standard supports up to 64,000 nodes within a PAN compared to a small handful with other WPAN links. 802.15.4 has also reduced complexity, intended to function with eight bit microcontrollers providing 8 Kbytes of RAM or less. Although IP over Bluetooth using the Bluetooth Network Encapsulation Protocol has been around for several years, it is typically used to provide a point-to-point connection over a single radio hop.

Researchers have developed numerous mesh network layers over 802.15.4, as open source projects (such as TinyOS, industrial forums (ZigBee and WirelessHART), or proprietary offerings (Dust Networks, Sensicast and Millenial Net). Each has defined its own set of incompatible packet formats tied to particular MAC features, routing algorithms, and addressing. Many address only the individual 802.15.4 subnet, leaving all further communication protocols to be defined via *ad hoc* gateways. 6LoWPAN potentially offers unification of this disparate activity and enable embedded 802.15.4 devices to be incorporated into Ethernet, Wi-Fi, General Packet Radio Service, and other environments within a uniform IP framework.

Many embedded TCP/IP stacks provide IP host functionality and are widely used in wired and powered settings. However, few embedded IP stacks directly address the issues related to supporting IP over low power mesh topologies in LoWPANs.

Within the IETF, the mobile *ad hoc* networks (MANET) working group and related research activities, tremendous effort has been devoted to reactive and proactive routing protocols for mobile devices. This work has assumed capable, high bandwidth links and powerful hosts with high, random mobility. As such, it used conventional IP datagrams and frame formats and hasn't attend to the impact of resource constraints.

Work in the IETF Auto configuration (AUTOCONF) working group is devoted to developing solutions for stateless address auto configuration and Neighbour Discovery in settings in which IP connectivity is naturally viewed as a collection of overlapping partial broadcast domains.

The Routing Over Low power and Lossy links (ROLL) working group was recently chartered to address routing in LoWPANs (independently of the link layer).

Until recently, extending IP out to wireless industrial networks was thought to be impractical, if not impossible. Vendors embraced proprietary protocols because they presumed that IP, which is memory and bandwidth intensive, couldn't be scaled down to operate on the microcontrollers and low power wireless links used in these environments. Recent efforts within the IETF make IP over low power communication links now feasible, including IEEE 802.15.4. These developments make IP attractive for use in low power devices, everything from handhelds to instruments.

***Jonathan Hui Ph.D**, Arch Rock Corporation*
***David Culler Ph.D**, University of California, Berkeley*
***Samita Chakrabarti**, IP Infusion*
*From the IPSO paper **6LoWPAN: Incorporating IEEE 802.15.4 into the IP architecture***

## 6LoWPAN Neighbour Discovery: a high level overview

**The 6LoWPAN Neighbour Discovery protocol is an extension to the standard IPv6 Neighbour Discovery protocol. 6LoWPAN ND mostly uses the same standard ND messages with some additional extension of options. A few new messages have been added to consider ND support for mobility, fault tolerance, bootstrap, header compression etc. The goal of this protocol is to support both IP-routed and link-layer mesh networks.**

IPv6 provides 128-bit address space which can support huge number of unique IP addresses. The core of IPv6 protocol functionality is the Neighbour Discovery Protocol which is mainly used for address resolution, address auto configuration, router discovery and neighbour reachability.

The IETF standardisation effort for 6LoWPAN Neighbour Discovery has taken the following issues related to low power wireless links into the design considerations:

- A continuously changing wireless link;
- Assigned short-addressing;
- The flat network model of 6LoWPAN;
- The use of both link layer mesh and IP layer mesh routing techniques;
- The lack of native multicast;
- Deep sleep cycles of 6LoWPAN nodes;
- The mobile nature of wireless nodes
- The limited energy and computational capacity of nodes.

The new 6LoWPAN ND specification takes these 6LoWPAN issues into account, providing optimised ND mechanisms. In order to do so certain assumptions are taken into account. For simplicity, IPv6 addresses have a direct mapping to link layer MAC addresses, and because of the nature of low power wireless networks a subnet is extended to cover whole 6LoWPAN islands. The use of multicast is avoided, and the operation of hosts is kept as simple as possible.

### Protocol overview and operations

6LoWPAN ND defines new terminology which is useful for understanding how it works. The nodes in 6LoWPAN networks are 6LoWPAN Hosts, 6LoWPAN Routers and Edge Routers. Hosts send and receive packets but do not route; 6LoWPAN routers perform IP routing within the LoWPAN, and Edge Routers route between 6LoWPANs and other IP networks.

A 6LoWPAN IPv6 subnet includes all the nodes which share the same IPv6 prefix covering a whole LoWPAN or Extended LoWPAN. Figure 1 shows the 6LoWPAN architecture with different kinds of IPv6 LoWPANs and network elements.

A LoWPAN is defined by the nodes sharing the same IPv6 prefix, usually with a single Edge Router as shown in the figure. A LoWPAN may be connected to other IP networks over a backhaul link, e.g. xDSL or GPRS.

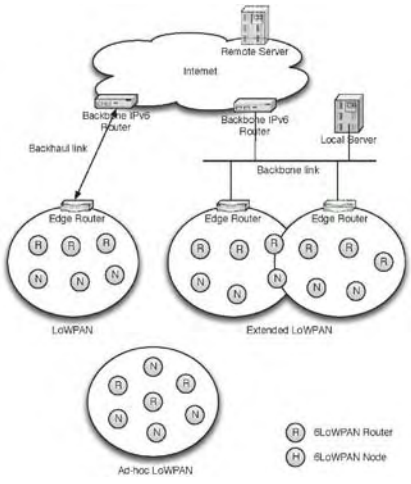


***Fig. 1.** The Model for 6LoWPAN architecture*

The Extended LoWPAN in the figure is made up of two Edge Routers and their nodes. Extended LoWPANs use a backbone link, e.g. Ethernet, to coordinate information about the LoWPAN, which is also part of the LoWPAN IPv6 subnet. Finally, an Ad Hoc LoWPAN is also shown which can operate without a backbone infrastructure. In Fig. 1, IP routing is used as an example within the LoWPANs with both Host and Router nodes.

After the underlying wireless link technology has been initialised, 6LoWPAN-ND provides basic mechanisms for a new node to bootstrap itself onto a 6LoWPAN using IPv6 stateless auto configuration.

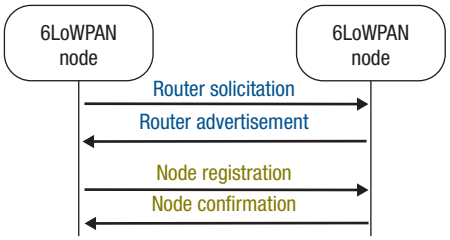


***Fig. 2**. 6LoWPAN ND message exchanges*

All nodes in a 6LoWPAN register with Edge Routers (ERs), which are on the border to other IP networks. It is assumed that the nodes are reachable by the Edge Router either by direct radio link or via 6LoWPAN routers. Nodes that cannot reach any Edge Router, may be configured to form an ad-hoc LoWPAN by configuring one to act as an Edge Router. A typical registration exchange between a host and router is shown in Fig. 2.

These Edge Routers keep track of the nodes in their network, and are able to perform duplicate address detection, address resolution and short address generation on behalf of nodes across the entire 6LoWPAN network island.

In addition, advanced features are provided for building large subnets including multiple Edge Routers handling large numbers of nodes. This technique is not only limited to infrastructure based networks, and can also be applied in isolated ad-hoc networks. All nodes in a LoWPAN have a unique IPv6 address (or possibly several), which does not change while the node is attached to the same or extended LoWPAN. This allows for efficient mobility within a LoWPAN and good manageability of nodes from outside the LoWPAN.

Finally, 6LoWPAN ND also includes fault detection techniques such as the use of secondary Edge Routers when the primary Edge Router fails or becomes unreachable. Moreover, mobile 6LoWPAN nodes can move between Edge Routers in the 6LoWPAN network, when the serving Edge Router becomes unreachable.

## Network example

Here follows a Building Automation example of how 6LoWPAN can be used to build different kinds of practical networks. The same principles could be applied widely to other processes.

Building Automation is an important application for wireless embedded networking, especially in energy efficiency and improved facility management. Applications requiring wireless networking include energy monitoring, lighting control, HVAC control, asset management, and door access control. 6LoWPAN is well suited to building automation as it provides a horizontal, reusable network for all the applications above along with the required security, scalability and mobility support.

Most building automation backend systems are IP- or Web-based today, leveraging the seamless Internet integration and manageability of IPv6.

Figure 3 shows a simplified example of a 6LoWPAN network for building automation. Such a network would leverage the existing Ethernet infrastructure, interconnecting all the Edge Routers of a facility which make up an Extended LoWPAN. Both remote and local servers can access nodes using IPv6 or IPv4 using a IPv4/IPv6 transition mechanism such as 6-to-4 tunnelling. All nodes have a unique IPv6 address, which stays the same while nodes move freely between Edge Routers in the facility. This is a crucial feature for asset management using 6LoWPAN active RF tags or IPv6 addressed sensors.

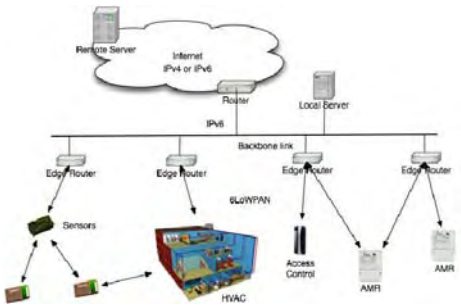


***Fig. 3**. Building automation network example*

## Router behaviour

**Edge Router:** An Edge Router has both a 6LoWPAN interface and a regular IPv6 interface and it is located at the junction of backbone and 6LoWPAN network. The Edge Router is responsible for Router advertisement to the 6LoWPAN network, address resolution of a node, maintaining a whiteboard of IPv6 addresses, duplicate address detection for the addresses it defends, forwarding packets from one 6LoWPAN to another 6LoWPAN or data packets between the 6LoWPAN network to the backbone network. When 6LoWPAN packets are forwarded to the backbone network, the 6LoWPAN adaptation layer is stripped off, the header is uncompressed and it makes sure that global IPv6 source address is used for the outgoing packets.

For incoming packets from backbone to the 6LoWPAN network, it adds 6LoWPAN specific adaptation layer and possibly 6LoWPAN IPv6 header compression mechanism and then forwards them to the 6LoWPAN network. Edge Router's 6LoWPAN interface also joins the 6LoWPAN_ER anycast address and it listens for Router Solicitation and Node Registration messages. The Edge router can also generate IPv6 addresses using IEEE 802.15.4 short addresses on behalf of the registering nodes. Note that the Edge Router follows standard IPv6 Neighbour Discovery procedures on the backbone IPv6 link.

**Backbone Router:** A regular legacy IPv6 router which communicates with the backbone side of the Edge router. A backbone router is not required to recognise the Owner Address Identifier when received with NS/NA Duplicate Address Detection messages.

The IPv6 backbone router may also be responsible for translating or tunnelling the IPv6 packets into IPv4 networks. IPv6 backbone router may use prefix delegation to Edge routers for globally unique IPv6 addresses in 6LoWPAN network. In future, more interactions between backbone routers and 6LoWPAN Edge Routers may be defined for the configuration and control of Edge Routers.

**6LoWPAN Router:** This type of router is present in the multihop mesh topology of 6LoWPAN networks. 6LoWPAN routers are essential components of 6LoWPAN links and they forward data packets across links and relay the control packets between the Edge router and the 6LoWPAN nodes. They also respond to the Router Solicitation messages from the nodes on the same link with Router Advertisements. It should run a 6LoWPAN compatible routing protocol, such as that being defined at IETF removes working group. 6LoWPAN routers are aware of the 6LoWPAN Edge Router address and may be able to forward a packet addressed to the 6LoWPAN_ER anycast address.

## 6LoWPAN node behaviour

6LoWPAN nodes are required to do very little ND signalling, and do not need to perform Duplicate Address Detection, Address Resolution, or Neighbour Solicitation. After performing boot-strapping as described in the previous section, the node needs to periodically renew its registration with one or more Edge Routers. A 6LoWPAN node forms an optimistic IPv6 link local address from the EUI64 bit unique global MAC address. It should then send a Router Solicitation to its own link router (either 6LoWPAN router or the on-link Edge Router). The 6LoWPAN node is recommended not to use multicast or broadcast address to send neighbour discovery messages in IEEE 802.15.4 networks. A 6LoWPAN node can request multiple addresses for registration and receives confirmation from the Edge router(s). Since the IPv6 address contains the EUI64 style MAC address in its lower 64 bit address, Neighbour Solicitation to the Edge Router for address resolution is not required. After successful Registration, a 6LoWPAN node sends the data packet directly to a 6LoWPAN neighbour's address via the 6LoWPAN routers that act as default routers. The 6LoWPAN node should configure a globally unique IPv6 address through auto configuration if it wants to communicate with nodes outside the 6LoWPAN networks.

***Samita Chakrabarti**, IP Infusion*
***Zach Shelby**, Sensinode*

*From the IPSO Alliance paper*
***6LoWPAN Neighbor Discovery: A Highlevel Overview***
www.ipso-alliance.org
First published in the ***industrial ethernet book*** July 2010

# Will 'The Internet of Things' change Industrial Wireless?

**The development of an IP-based stack specification Using ZigBee IP created a great deal of interest within the power and utility industries. But what does all this activity in networked wireless development mean for the Process and Factory Automation sector? Long before governments around the world felt a need to pour stimulus money into data infrastructure as an investment strategy in its own right, the Process and Factory Automation businesses had been making quiet progress with dedicated wireless technology development. Is this all about to change with the development of low cost, low power, mass market wireless stacks capable of making direct connections to the Internet asks Frank Ogden**

LET US STATE from the outset that WirelessHART is a good, secure and robust method of exchanging relatively low rate data around physically distributed applications without wires. It also runs the proven PHY and MAC layers of IEEE802.15.4 which provide an excellent platform for the higher layers of wireless mesh networking and other features essential to wireless in Process Automation.

So why are we getting so excited about the appearance of new wireless stacks which fit onto 802.15.4 and are not WirelessHART?

The answer is simple. Enter the words 'smart grid' qualified by the word 'stimulus' into Google and then amaze at the countless big money initiatives which the search engine throws up. Then look more closely at the names involved… Cisco, PG&E, GE, Landis+Gyr, etc, not to mention a second order of both the eminently sensible organisations as well as complete eco-barmies. This is serious stuff and the money which surrounds it will prime the development pumps in a way that wireless data network programmes have not seen before. The money now promised by governments right

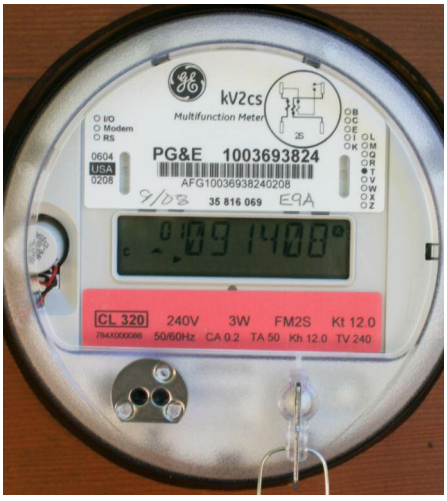


around the world (not just in the USA) dwarfs Process and Factory Automation developments by orders of magnitude.

In practical terms, this money will (mostly) mean the development of public utility meters which can implement SCADA over wireless enabling the exchange of billing information, interactive adjustment of tariffs and control of both consuming and microgeneration plant. Of course this might be achieved by throwing some of those development $billions at new wireless protocols designed specifically for smart grid and nothing else – a kind of WHART for electricity meters. But it is not going to happen like that for the same reason that the HCF wireless device protocol is unlikely to achieve long term dominance in the Process industry: lack of Internet-compatible direct connectivity.

## Gateways can be good

Actually, you can make two good arguments for addressing HCF and proprietary wireless stacks through network gateways rather than the IP-based routers – the equivalent device in IPv6 – for that is what we are really talking about. Gateways offer security through natural segmentation, always a good argument when talking about Process Automation.

The second argument, perhaps cynical, is even more compelling. If an application does not conform to Internet standards, then it becomes a marketing lock provided that you are large enough to enforce it.

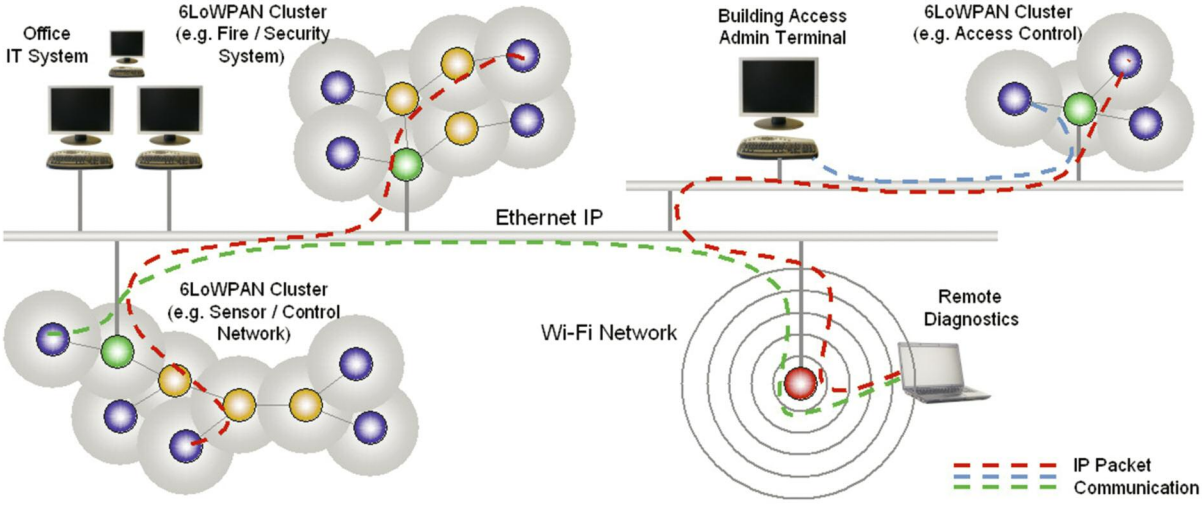


***British fabless semiconductor company Jennic** is just one of many companies around the world attempting to leverage 6LoWPAN. The company's protocol stack provides a wireless connectivity based on the IEEE802.15.4 standard at 2.4GHz allowing embedded devices to communicate wirelessly using Internet Protocol (IP). Designed to work with Jennic wireless microcontroller, the IP connectivity solution provides a low power single chip implementation for the development of wireless networking products that offer multi-year battery life, which can communicate with other IP devices in an existing network. Point-to-point and star connectivity are supported as standard in IEEE802.15.4, but the device can also connect over a proprietary networking stack, providing a self-healing cluster tree with IPv6 direct addressing to individual nodes.*

*If it will work for smart meters, it could also work for remote sensors in a way presently beyond today's wireless Process Automation systems*

This happened in the early days of wired fieldbus where technology and marketing were synonymous. It is worth speculating how the early fieldbus markets might have developed if today's Internet had been around at the time…

Well, the Internet is here and the virtually inexhaustible IPv6 addressing, presently without immediate value in Process or Factory automation, promises the 'Internet of Things', a connected web of smart objects. And it is a reasonable bet that some of those Things will find their way into both Factory and Plant.

Actually, the gateway vs router argument loses some of its force when you start talking about wireless mesh. Given that both WHART and the new IPv6 stacks have highly robust security layers, Internet gateways merely create a potential system vulnerability accessible from a single IP address; dumb(ish) wireless routers and IPv6 have both good stacks and sheer numbers to ward off attacks. A malware malcontent might just be able to sidestep firewalls and conventional security to find an individual device to target, they will also have to find the right number address and there are an awful lot to choose from! 192.168.2.235 probably won't hack it…

## The problem clearly stated

While we are talking very early days in IPv6 wireless technology albeit backed with big money, it seems important not to get too carried away. The *Problem Statement* written by Bormann and Sturek of the IETF 6LoWPAN Working Group is essential reading for anyone of serious disposition towards wireless development. The introduction to the document provides quite a good summary:

'The 6LoWPAN and ROLL [Routing Over Low power and Lossy networks] WGs are laying the groundwork to make the Wireless Embedded Internet a reality, but what application protocols will we use with these networks?

6LoWPAN's lowpower area networks (LoWPANs) and, more generally, ROLL's lowpower lossy networks (LLNs) exhibit severe constraints on the bit rates achievable and the packet sizes that can be efficiently sent. Many (but not always all) nodes in these networks also are constrained in the energy they can expend for computing and communication, the memory available, and the code size that can be accommodated in each node.

More generally speaking, there are many factors that play together to limit the per-node capabilities compared to today's typical Internet nodes. (Moreover in most cases links in LoWPANs and LLNs are "lossy" by nature, thus experience high bit error ratios, making the communication sometimes challenging. The existing transport protocols used for reliability such as TCP may be too expensive to implement for LoWPAN nodes and may not perform well in lossy environments.)

The established application protocols for Internet applications may be a poor fit for LoWPANs and LLNs. For instance, request response protocols like HTTP may require battery-operated, mostly sleeping nodes to be listening for requests much more frequently than their application processing requirements alone would need them to wake up. In addition, some of the applications may require optimizations in terms of bandwidth usage; for example, the usual data formats (both headers and body) are perceived to be too chatty for the 50 to 60 byte payloads possible in LoWPANs. Their interpretation and generation may require too much code for the 8-bit and 16-bit processors dominating LoWPAN nodes.

Still, it would be a mistake to start a new silo of application protocols that do not benefit from existing application area Internet experience. A number of concepts well-established in the IETF application protocols, such as identifying resources by URIs, may transfer very well to LoWPANs…'

## Reality kicks in

It is the answering of such questions which appears to be tying up ISA100 with inextricable knots in the committee's pursuit of a wireless mesh Internet Protocol standard(s) for everything Industrial and Process. Meanwhile, the real world needs a more practical approach to wireless application. IEEE802.11a/b/g/n is fine and fast for cable network replacement where the nodes can draw hand-warming electrical power.

IEEE802.15.1 Bluetooth is much slower, less power-hungry but more robust: an altogether excellent all-round candidate for high reliability sloth-like industrial wireless devices which do not require direct network connection or self-forming networks. Can work off batteries.

ZigBee. Great for opening garage doors and other jobs around the house. It recently declared an IP stack version for smart metering which could (eventually) find industrial and Process applications.

WirelessHART and other IEEE802.15.4 base platforms. Great for SCADA-ing up oil refineries and slothful, but physically spread applications in need of low power wireless mesh that do not require direct Internet device control.

ISA100 and variants. I'll keep watching because this is what I am paid to do, but the signals are rather confusing.

6LoWPAN (IEEE802.15.4 PHY and datalink). Far too early to make any serious statements about future relationships with Process and Factory Automation. However, the IPv6 network mesh, a moniker which includes *The Internet of Things* and software stacks with a size inverse to those of the development dollars make this the one to watch.

First published in the ***industrial ethernet book*** September 2009

# Untangle the Mesh: Comparing mesh networking technologies

**Over the past few years, mesh networks have become increasingly popular with the appearance of more wireless devices. As with other technology trends, there has emerged a plethora of different mesh networking technologies and architectures. This article is based on a paper entitled *What a Mesh!* by Joel K Young given at Embedded Systems. The author discusses mesh networks, including wireless network basics, the criteria for evaluating different wireless mesh networking technologies and an evaluation of some offerings**

NETWORKS TOPOLOGIES are not always what they seem to be. In the wired world, they generally follow the path of the wires – very simple. If devices are wired in a ring, then the network topology is a ring. Wireless paths and access methods are not always obvious. For example, is a Wi-Fi access point a star topology or a bus?

A mesh network employs some level of more complete interconnection among nodes. This means that paths are not defined by a specific architectural pattern, but rather by the connections themselves. In the full mesh topology, each node (workstation or other device) is connected directly to each of the others. In the partial mesh topology, some nodes are connected to all participants, but others are connected only to those other nodes with which they exchange the most data.

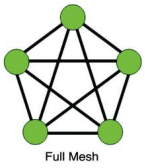


***Fig. 1.*** *Full mesh example where all nodes interconnect*

Figure 1 illustrates a full mesh, where each of the five nodes is connected to all the others. Another important thing to note about a mesh is that some or all nodes may be routers and some or all nodes may be end-points since full interconnection is not achieved unless the network is very small. Full interconnection gets complex very quickly. Figure 2 illustrates three different instantiations of mesh networks. The green nodes are end devices, the yellow are routers (which may also be end devices) and the purple is the network coordinator – responsible for allowing joining and departing from the mesh (more on this later). Note that one instantiation of a mesh can be a star – a mesh with one router and the rest end points. The Cluster Tree network is a combination of near full connectivity among routers and end points hanging off individual routers. The Peer to Peer mesh generally gives equal rights to all nodes, including routing and end point functionality

Wireless – particularly mesh – networks present problems not found in wired networks such as for instance, accessing the medium. With wireless in an open space, listening is more important than talking. If all nodes talk at once, listening is difficult. So radios must be good listeners if they are going to have a chance to get a word in edge wise, so to speak.

Determining paths in a wireless mesh network is also difficult because the environment is dynamic. In this case, there are two choices: planning a route in advance, or hooking up one step at a time. Occasionally doing both is best – this usually involves retracing ones steps and repeating well travelled routes. Also, changes in the working environment will cause some paths to nodes to disappear only to re-appear later. This is due to changing signal conditions or traffic conditions.

Sleeping and waking is a power issue. Wireless might be able to eliminate the traditional power cable, but using batteries requires effective power management. The most common way of handling power management is putting the nodes to sleep when they are not being used. This sounds fine but then how do you wake them up when you need them?

## Which technology is best?

Essentially this requires comparisons made on the basis of security, reliability, power management, scalability, data movement, and cost.

Security is as much about the perception of threat as the actual threat itself. Nonetheless, security is easily evaluated by well understood traditional factors. The first is encryption – protecting the information itself. Modern encryption wants at least AES128 as an algorithm (128 bit key). The next is authentication, which is validating that the users (or nodes) are who they say they are. This is typically handled by key exchange or authenticated certificate. Last is authorisation, which should be thought of as granting permission associated with having the right key or certificate. Beyond these, there are other factors which are associated with the ease of distributing and configuring the authorisation and authentication mechanisms.

Reliability can be defined as the likelihood that a message will arrive at its intended destination on time, and if not on time it must arrive eventually.

Power Management – When discussing wireless sensor networks, how long will batteries last? In the context of network architecture, power management is analysed in terms of end nodes, router nodes and network coordinators. It is most important to have power-efficient end nodes because they are likely to be at the greatest distance from power sources. The routers are second: Battery powered devices that sleep extend the flexibility of the architecture. Finally, the coordinator is usually always powered.

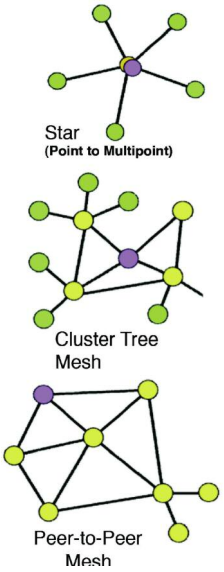


***Fig. 2.*** *Three different variations onmesh networks. The green nodes are end devices, the yellow are routers (which may also be end devices) and the purple is the network coordinator*

In the context of nodes that can sleep, average power consumption is key. This is best assessed by looking at the combination of how they wake up, how frequently they wake up, total transmitting time and total listening time. Since the most power is consumed when radios transmit, it is important to keep transmit time to a minimum.

Scalability is about the largest practical size to which a network can be extended before unreliability sets in. All the networks have large physical limits in the 10s of thousands, but a practical design is always much smaller. This is because scalability is related both to reliability mechanisms and nature of the application. If a network never experiences problems which cause rerouting, then the routing tables will never change, meaning cached routes will always work and there will be few retransmissions or reroutes because of failure. This describes a stable network that can be very large.

Scalability is also dependent on the type and

## Glossary

- **DSSS** – Direct Sequence Spread Spectrum. This is a method of encoding a signal which distributes information over a wide path of spectrum using a pseudo random code. Because of the wide spreading, the signal appears to be noise for those without the spreading code.
- **FHSS** – Frequency Hopping Spread Spectrum. Similar to DSSS, the big difference is that is uses a more constrained spreading algorithm and changes channels as a function of time, theoretically making the transmission more immune to interference.
- **TSMP** – Time Synchronized Mesh Protocol. This is a mesh protocol that uses time slots to allocate spectrum for communication between two nodes. Because time slots differ over pairs, interference is minimized because access to the channel is controlled by timeslot.
- **Cluster Tree** – Region based mesh network routing algorithm. In this algorithm, routes are formed and maintained between clusters of nodes. Route discovery is completed and maintained between the clusters – providing access to the children of each cluster.
- **PAN ID** – Personal Area Network Identifier. This is the term for the network name assigned to particular personal area network.
- **CSMA** – Carrier Sense Multiple Access. This protocol defines the channel access technique deployed by Ethernet, Wi-Fi and bus oriented networks. It provides a method for detecting collisions and retransmitting as a method to acquire a communications channel.
- **TDMA** – Time Division Multiple Access. The protocol defines the channel access technique used by TSMP and GSM networks in which a communications channel is divided into time slots. Each node is allocated a specific time slot for communication.

volume of data. Data flow can be placed in three categories: dribble data, burst data and streaming data. Dribble data is periodic, infrequent and slow, while streaming data is constant, etc. A network can be large if the traffic is dribble data because the flow follows consistent patterns, with plenty of bandwidth. Sleeping networks do well with dribble data, but scale poorly with streaming data.

**Data Movement** is about raw carrying capacity and the classic trade off: Does the application require lots of data with low latency or does it require dribble data with long, non deterministic latency? This can be resolved to five variables, namely data rate, latency, packet size, fragmentation, and range.

**Cost** is measured by the unit cost per node as well as the cost to maintain the network. It becomes more complicated in the case of battery powered sleeping nodes. As one might expect the combined unit and maintenance cost of wireless mesh networks tends to be evaluated relative to the cost of the devices served.

# Network architectures

## Point to multipoint

**Key Characteristics.** Mostly a simple star, it is often confused with a mesh network. These networks, of which Bluetooth is the prime example, tend to use either FHSS or DSSS (802.15.4) air interfaces. They require static configuration for PAN ID, routes, and security. All nodes can see all other nodes and need to be told which nodes to talk to. Security tends to be pair-wise for both the encryption and key. End points may go to sleep or stay awake, but the central router is always awake.

**Network Architecture.** All nodes are on the same channel (or hop to the same channel). Bandwidth/throughput limited by simultaneous data at the concentration point. Figure 3 illustrates a typical topology. Collisions happen with lots of traffic or lots of nodes.

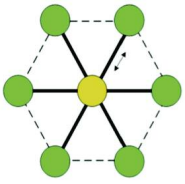

***Fig. 3.*** *Point to multipoint network. Mostly a simple star, it is often confused with a mesh network. Bluetooth is a prime example of the network type*

**Strengths.** The beauty of the basic non-mesh point to multipoint network is simplicity. Communication, unless traffic is very heavy, is relatively deterministic since there are no hops and minimal, or managed, collisions. It also allows for maximum throughput because there is no added routing or route discovery. Finally it is easy to understand and easy to manage. Because of the simplicity, it also tends to drive the lowest cost for its specific size and function.

**Limitations.** Unfortunately, the simplicity provokes limitations. Networks tend to be small. Large networks only work if polled from central point. This requires very specific message management. There are also single points of failure and no ways to route around changing conditions. The network follows the belief that if it worked the first time, it will work forever – which it will if RF environmental conditions are maintained.

## ZigBee 2007

**Key Characteristics.** ZigBee is built on top of 802.15.4 using DSSS at 2.4GHz. End points sleep, routers don't sleep and a coordinator is needed to start the network and to allow points to join the network. ▶

## Waiting in the wings – ISA-100.11a

This article was written for last year's Embedded Systems conference and before ISA ratified its Wireless for Process Automation specification in September of this year. We feel that Joel Young's paper would not be complete without a suitable update.

With Honeywell as its highest profile backer, the ISA-100.11a standard 'Wireless Systems for Industrial Automation: Process Control and Related Applications' is a mesh network built on an IEEE802.15.4 MAC and physical layer, TDMA mesh protocol (as championed by Dust Networks), uses the WirelessHART-style DSSS frequency hopping for interference mapping and avoidance. It also of course implements mesh routing.

ISA describes the protocol as providing 'reliable and secure wireless operation for non-critical monitoring, alerting, supervisory control, open loop control, and closed loop control applications. The standard defines the protocol suite, system management, gateway, and security specifications for low-data-rate wireless connectivity with fixed, portable, and moving devices supporting very limited power consumption requirements.'

The application focus is on applications such as monitoring and process control where latency in the order of 100ms can be tolerated, with optional behaviour for shorter latency. The declared aim is to provide reliable and secure wireless operation for non-critical monitoring, alerting, supervisory control, open loop control, and closed-loop control applications.

While the first version requires network connection through a gateway and does not directly support IP-based address space at device level, other ISA working groups plan coexistence with other industrial wireless devices based on IEEE 802.11x, IEEE 802.15x, IEEE 802.16x, and other relevant standards. This plan includes ISA-100 proposals capable of direct network addressing using IPv6.

**Frank Ogden**

The ZigBee standard has evolved in three different versions: 2004, 2006 and 2007. ZigBee 2004 is no longer used and ZigBee 2006 had significant limitations. ZigBee 2007 includes key features for frequency agility, message fragmentation, and enhanced security associated with key management. The routing of messages follows Cluster Tree methodology where routes to all points are maintained at each cluster. This allows a very short routing time, but requires many routes. Discovery of routes uses the AODV algorithm where paths are explored between clusters. A ZigBee release using IPv6 addressing is presently under development.

**Network Architecture.** The network consists of three specific device types. A ZigBee Coordinator (ZC) is required for each network to initiate network formation. The coordinator may act as a router once the network is formed. The ZigBee Router (ZR) is actually an optional network component, although a network without routers becomes a point to multipoint network. The router participates in multi-hop routing of messages. Finally, the ZigBee End Device (ZED) does not allow association and does not participate in routing. Figure 4 illustrates an example network.

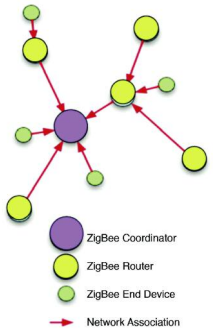


*Fig. 4. ZigBee network example. End points sleep, routers don't sleep and a coordinator is needed to start the network and to allow points to join the network.*

**Strengths.** End devices consume very low power. Cluster Tree routing provides quick knowledge of routes and thereby efficient routing. With ZigBee 2007, frequency agility skips interference-prone channels automatically. Long messages are allowed with message fragmentation support and security is flexible with support of separated keys. Finally, the network can scale to very large.

**Limitations.** Routers must be powered at all times: no sleep mode. Cluster Tree routing creates heavy route discovery traffic on network changes. Heavy traffic volume produces collisions and potential message loss. A coordinator is needed to start and manage the network: if the coordinator goes down, the network can't start.

## WirelessHART

**Key Characteristics.** WirelessHART uses the Time Synchronised Mesh Protocol created by Dust Networks. Unlike other networks, the time-based system uses TDMA as access method. The network is optimised for low power and all nodes can be sleeping routers and every node is a router. A Gateway is required to implement the critical time synchronisation of sleeping and waking functions. Like ZigBee, it uses 802.15.4 DSSS, with the addition of a more advanced frequency hopping algorithm. Security includes encryption and authentication.

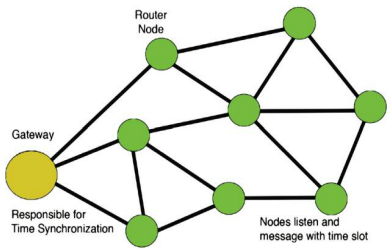


*Fig. 5. WirelessHART network example. A Gateway is required to implement the critical time synchronisation of sleeping and waking functions.*

**Network Architecture.** Note that all the nodes are routers. Figure 5 illustrates a typical network topology. The illustrated routes change dynamically based on visibility within specific time slots as it hops through the different DSSS channels. The relationship between any two nodes is negotiated to be in a specific time slot, thereby minimising the probability of any collisions. When sleeping, nodes awaken during their time slot and listen to see if there is any traffic. Clocks are kept synchronised by the gateway.

**Strengths.** Every node provides a routing function with very low power consumption. Since transmissions occur only within the allocated time slot, retransmissions are minimised. Communications are reliable with every message acknowledged. Networks are able to scale to moderate level – around 1000 nodes. Frequency hopping minimises the probability of interference. Security includes encryption and appropriate authentication.

**Limitations.** Because of the time slot approach, latency is long and non-deterministic. It takes a network some time to form, and for the participating nodes to negotiate their individual time slots. Because communications is slotted, the available 802.15.4 bandwidth is split up, meaning that throughput is minimised for burst traffic. A powered gateway (coordinator) is required to keep the network functioning – opening up a single point of failure if the gateway becomes unavailable for an extended period of time. Finally, the radios are expensive compared to the other available solutions.

## 6LoWPAN

**Key Characteristics.** 6LoWPAN is an acronym for IPv6 over low power wireless personal area networks. Presently it is a proposed standard based on the IETF RCF 4944. It is designed to be used over 802.15.4 chips and radios. Unlike traditional IPv6, 6LoWPAN deals with packet size incompatibilities in message transport (128 bytes vs MTU of 1280 bytes in IPv6) and it is designed for a small memory footprint systems. Today it is a point to multipoint architecture but with proposed augmentation to a mesh routing scheme.

Figure 6 provides an example of 6LoWPAN network topology. Note that for now it offers only point to multipoint. Unlike the other networks discussed in this paper, the figure shows an end to end IP-based link from a host computer to an end device. In this case it is illustrated by an electricity meter. The end device is directly addressable by the host computer on the far end of the network. The interworking function provided in the pictured box provides a transport change and repacketing at the data link level.

**Network Architecture.** Typical ad hoc network topology. Unlike the Cluster Tree

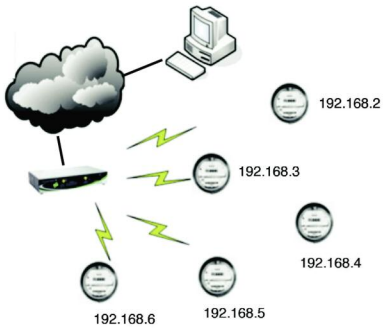


***Fig. 6.*** *IPv6 LoWPAN network example. Unlike the Cluster Tree method described in ZigBee, routes are only determined on an as-needed basis. This means that routes that are never used never get routing table entries and routes that are used frequently are continuously updated, thus optimising efficiency.*

method described in ZigBee, routes are only determined on an as needed basis. This means that routes that are never used never get routing table entries and routes that are used frequently are continuously updated, thus optimising efficiency.

**Strengths.** Every node is a router with very low power consumption. Further, because every message is acknowledged and routes are determined on an as needed basis, the network is not overwhelmed with discovery traffic – important with sleeping battery powered routers. Efficient route discovery and routing means that the network only learns routes that actually get used (AODV). Frequency agility is supported and security involves both encryption and authentication. Reliability is projected at 99.99%. Finally, the system supports larger payloads with support for message fragmentation.

**Limitations.** Efficient power management means long latency and non-determinism. Even though throughput is not limited by time slots, it is still limited depending on loading and discoveries. The network can scale to a moderate size of around 500+ nodes and can be very large if traffic is light and message flow doesn't change much.

## System comparisons

Using the criteria defined at the beginning of this document, they all do very well in security in that they have well defined encryption, authentication and authorisation schemes. ZigBee and 6LoWPAN have a slight advantage in that their key systems should be easier to implement and thus be more flexible.

With respect to reliability, point to multipoint takes the biggest hit because it inherently has a single point of failure. Some schemes may have frequency agility options while others do not. Prior to the 2007 standard, ZigBee had a weakness in the frequency agility area which has now been fixed along with added support for message fragmentation. The others are similar – WirelessHART is designed to never lose a message while 6LoWPAN does well on the assumption that the existing TCP/IP protocol suite has class of service built in. While our own proprietary DigiMesh has a similar approach to WirelessHART, it is still somewhat unproven in large deployments.

Power management will no doubt be hotly debated. WirelessHART defines systems where all nodes in the network, including routers, can sleep. Even though sleeping ZigBee end devices are most efficient when it comes to power, the fact that routers can't sleep bumped the rating down. Until 6LoWPAN publishes a mesh and power management strategy, the rating will remain unknown.

The scalability rating follows directly from the question of how big can the network get and still function. This is where the ZigBee 2007 Pro stack shines. The Cluster Tree architecture creates a hierarchy which enables scalability. WirelessHART scales well; particularly if most communication is kept local – however, the networks can tend to get very slow when they become too large. Finally, point to multipoint architecture has an obvious limitation in the number of nodes that can be attached to one central point.

The best data mover is no doubt the simplest system – namely point to multipoint. The simple network design means that focus can be made on short, deterministic latency and high data throughput. There is a direct trade off here with power. WirelessHART rates lower because it is focused on minimising power and maximising reliability – this naturally leads to less deterministic latency and lower throughput. Of course, as a network gets bigger, these two networks will actually do better. However, this is represented in the high scalability ratings for these networks. ZigBee fits in the middle because the backbone of powered routers can move data very efficiently – but can get stuck if too many route discoveries are needed.

Cost may create the most debate. The ratings here were based primarily on the cost of available chip set solutions under the assumption that the right architecture is chosen for the right job. If not, then the cost ratings become meaningless. For example, attempting to deploy a ZigBee solution where battery powered routers are desired means infrastructure costs will skyrocket. So, given this caveat, point to multipoint, ZigBee and our DigiMesh protocol have common costs because they all use similar chipsets.

Each wireless mesh architecture has its respective benefits and there is no single approach that fulfils every wish list. Hence, it is important to match the needs of the application to the capabilities of the network.

***Joel K. Young*** *is senior VP and Chief Technology Officer, Digi International Inc.*

First published in the ***industrial ethernet book*** November 2009

# New routing improves wireless mesh network performance

**The main goal of a wireless sensor network is to enable reception of data from the sensor field without the need for physical connection or access. But there should also exist a data path which can run in the opposite direction: most serious applications also demand control over the networked sensors – true SCADA – which requires the use of bi-directional networks. Most wireless mesh implementations use process-intensive algorithms to create a self-forming, self-healing single path through the network. Diversity Path Mesh offers an attractive alternative by routing everything everywhere. Marius Gafen describes a new mesh paradigm**

THE MAJOR APPLICATION driver for most wireless mesh systems presently relates to smart electrical metering. These wireless networks can be used to build Advanced Metering Infrastructure (AMI) to continuously monitor power consumption, control the functionality of the meters, as well as limit the end-user consumption or report about tampering with the metres. They are expected to handle all these functions in something close to real time. Such capabilities enable electric utilities to reduce operational expenses, implement flexible management systems based on real time energy consumption, increase the system's reliability and ultimately help to reduce both cost and energy usage.

Flexible mesh topology has shown itself to offer the best choice for a large range of wireless sensor network applications; more often than not it proves to be the only feasible topology for modern projects. Critical factors influencing performance of a wireless mesh network vary from application to application and usually require most of the following capabilities:

- Range and coverage. These are probably the most obvious requirements for a wireless network, starting from the node-to-node range at a given transmission power vs antenna gain vs data rate compromise. Range is affected by the quality of the physical layer and the efficiency with which the data propagates through the network. In mesh topology, the technology also necessitates subtler and often more important range-related detail such as the influence of multipath conditions and changes, or the way in which the requirements are affected when the nodes operate in mesh topology rather than just between two separate nodes. Coverage requirements are closely related to range and mainly involve the elimination of dead spots in the network. The extent of the coverage area relates to multiples of the basic node-to-node range.
- Robustness to network changes and RF interference. Some changes can be internal and direct such as adding, removing or moving nodes. Other propagation path factors may derive from incidental events such as the addition (or removal) of walls or buildings or, more trivially, the replacement of a wooden door in the RF path with a metallic one. Such influences lead to subtle, but important issues concerning how the network reacts to changes. For example, could the network go down as a result of these events and, if so, for how long?
- Scalability. This means the ability to form network cells comprising just a few nodes, to cells with thousands or even tens of thousands of nodes. While it seems unlikely that industrial applications would require such scaling, the ability to increase the size of an existing network by orders of magnitude is a definite requirement in the Utility world. Spatial scalability is another aspect of the problem. If network cells which measure a few metres can be expanded in physical size to kilometres without recourse to other technologies such as cellular GSM, this can be a most desirable property. The number of hops involved in the geographical deployment of the nodes is a related concern.
- Power consumption. For many if not most applications, the longest period of time that a node can operate from a battery power source becomes a vital factor, as is the ability to mix between nodes with various power supply types and capacity.
- Ease of integration. The way in which a particular type of network physical layer can connect up to and coexist with other network types is an absolutely critical factor. For example, in some organisations a need to train the deployment staff in specialist network installation may preclude a move to wireless; in others, the network management resources may reduce or even annul the operational cost benefits of the network. Networks that enable an immediate move from wire to wireless without need for new management software or thorough modifications constitute a powerful force for network integration, for example, just by cutting the existing RS485 or Modbus wires and inserting the wireless network nodes.

PHOTO: EDF ENERGY

Despite the fact that mesh topology often represents the best option for numerous applications, the vast majority of existing wireless network types can be configured to all the relevant installation topologies, thus addressing their lowest common denominator. Does this present a paragon of flexibility or pointless complexity? We think that such a strategy can be justified only if the application requires several topologies to be applied concomitantly, or at least in the same project. However, when it is clear that the only possible wireless application topology is mesh, it makes sense to use a technology targeted specifically towards mesh topology.

## Diversity Path Mesh

This technology is designed to maximise the most important performance factors of wireless mesh networks against their trade-off characteristics while operating in the crowded ISM frequency bands. Diversity Path Mesh is a multi-hop, bidirectional communication technology, developed for wireless sensor networks using mesh topology and operating in the unlicensed frequency bands.

The basic technique behind DPM is the use of synchronised flooding, where a node transmits the very same message to all the nodes in the network cell. The theoretical concept of flooding in networks as an alternative to routing is well known. In a flooding-based network, the message data is sent to all the network neighbours, thus eliminating the need to route messages, and enabling the network to use multiple propagation paths (actually all the available propagation paths). Such networks benefit from space diversity, improving overall network robustness tremendously.

Routing calculation is a most demanding network management task, and its elimination by use of message flooding greatly reduces the demand for processing resources in the node in terms of CPU power and memory. It also decreases the amount of data carried in the network since messages contain only payload data; there is no overhead in the carriage of

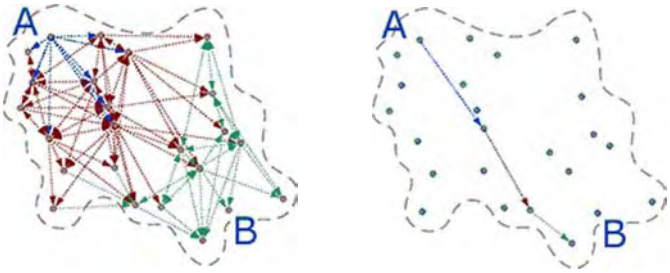


***With Diversity Path Mesh*** *(left) all data packets propagate across the mesh, and every node in the mesh, in a single wave. A strict synchronistic mechanism prevents transmitted packets from echoing around the nodes as unwanted reflections and loops. The simple propagation concept requires little processing power for routing algorithms when compared to that required in isolating a standard single routing path.*

*[One might imagine that every node carrying every data packet would end up using more battery power overall, but the authors are adamant that this is not the case – **Ed**]*

routing information. Furthermore, there is no set up time and any number of nodes can be inserted or removed. As long as the added or remaining nodes are within the reception range, the network simply continues to operate – without healing time and network downtime.

Until recently, the use of flooding in standard mesh network architectures has been avoided for several reasons, commonly known as the 'broadcast storm problem'. Simply put, nodes within reception range retransmit the message whenever received. This results in an uncontrollable series of collisions which degrade network performance and increase energy consumption to unacceptable levels.

To channel the flooding technique into a practical and useful solution, Diversity Path Mesh makes use of high level synchronisation. Messages at each node are relayed to the surrounding nodes with precise timing, thus forming multiple concomitant transmission paths on the way to the destination. The retransmission of messages through the network is synchronised to sub-bit level through using TDMA as the master construction framework. Instead of interfering which each other, the multiple and identical, ±noise, transmissions received by the node receivers can then be summed together in a demodulator function block. This summing action increases the strength (so decreasing the error rate) of the received signal.

Overall, the scheme returns increased reliability since there is no single point (node) of failure, and increased propagation robustness due to the inherent spatial diversity of the propagation through the multiple relay paths. The same summing mechanism which is at the heart of the proposition translates to greater range. It also virtually eliminates dead spots in steady state conditions, multipath occurrence, the effects of RF interference and the effects of changes in propagation paths.

There is a significant reduction in the probability of a message failing to reach its destination. With sequential propagation from node to node in a standard mesh network, the time delay adds sequentially with each propagation step. With simultaneous propagation across the mesh, the time delay to get from one side to the other is greatly reduced over a sequential propagation mode. Relaying the messages also enables DPM networks to extend as far as needed, with the cell overall range and robustness rising with an increasing number of nodes. The maximum number of nodes which could be deployed exceeds any practical requirement for such applications.

In a nutshell, the behaviour of a DPM network has these characteristics:

- The nodes are the network and there are no routers. Once the nodes are connected and powered, the network is up and running. Also there is no need for network management and consequently no need to develop software to manage the network. Compared to router-based networks, extended software development and field tests are eliminated.
- The operation of adding or removing nodes is immediate and effortless. There is no such notion as reprogramming or resetting, hence there are no corresponding delays and no downtime.
- Messages propagate in simultaneous, parallel paths, thus increasing the range between nodes and improving the resiliency to external conditions and RF interference.
- The network traffic comprises practically pure data with no management overhead, resulting in high data throughput and lower power consumption as a consequence.
- The number of times the data may be retransmitted by relay (the number of legs or hops) is practically unlimited, with actual numbers far exceeding the requirements of real applications. For example, with the basic range between nodes being some 2km (with a high power module) a coverage area spanning 60km can be easily achieved within 30 hops within a DPM network cell.
- Increasing the number of nodes either increases the size of the network cell, or the robustness of the network, or both.

## Conclusion

Flooding is arguably the most appropriate technique to be considered for creating a network mesh as it addresses critical performance factors of wireless mesh networks, as well as the accompanying system tradeoffs. However, the flooding technique carries several grave drawbacks which have precluded its use in the past.

Diversity Path Mesh is probably the first mesh wireless network to use the flooding technique instead of standard routing; synchronised flooding produces a surprising performance benefit over and above the positive benefits of flooding propagation.

Products based on this technology and deployed by its originator, Virtual Extension, have proven the capability of Diversity Path Mesh, now successfully deployed in projects where other technologies have been unable to deliver. These applications include smart metering – electricity water and gas, smart and emergency lighting, building and industrial automation, vending machines, agriculture and security.

***Marius Gafen*** *is involved with Israeli start-up company Virtual Extension Inc*

www.virtual-extension.com

First published in the ***industrial ethernet book*** May 2010

# Wireless makes inroads across the process automation sector

**Wireless technology is commonplace and WirelessHART, using the IEEE 802.15.4 radio operating at 2.4GHz, is providing tangible benefits in the process industries. The network range has been extended and redundant communication routes provide the crucial 99.9% reliability, reports David Walker who covers integration, device diagnostics, security and standardisation on HART IP as part of his working brief.**

WIRELESS NETWORKS have been around for about 10 years in the process industries. w-HART is a widely recognised standard that provides a process industry application that, arguably, is as easy to use as a Bluetooth connection to a mobile phone.

WirelessHART (w-HART) was developed specifically to meet the needs of process sector users, who demanded coexistence, reliability, security, multi-vendor interoperability and long battery life. Crucially, power consumption is limited such that process devices can be battery-powered for up to 10 years. The HART Communication Foundation (HCF) worked with process equipment vendors and experts on radio frequency and wireless communications, and came up with a scheme that satisfied all parties. A typical wireless net is shown in Fig. 1.

The network uses the IEEE 802.15.4 radio operating at 2.4GHz. Radios use direct-sequence spread spectrum (DSSS) technology and channel hopping for communication security and reliability, plus time division multiple access (TDMA) to ensure latency-controlled communications between devices on the network. Each device in the mesh network can serve as a router for messages from other devices. This extends the network range and provides redundant communication routes with 99.9% reliability.

Complementary Wi-Fi solutions find applications in video over wireless, field data backhaul and control network bridging. Such Wireless Plant Network (WPN) solutions are all IEEE 802.11 – 2007 based.

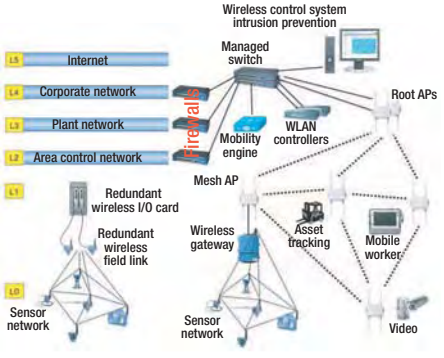


***Fig. 1: A typical wireless plant network architecture.*** *This illustrates the various operational levels and field devices based on a w-HART setup*

PHOTO: EMERSON PROCESS

***Remote pumps communicate wirelessly to a WirelessHART gateway.*** *These Middle-Eastern pumps need performance monitoring so that engineers can make informed maintenance decisions. The gateway connects to the main control system over a WAN so that the engineers gain essential data at their desks. Installing wireless sensors on remote wellheads and other similar applications is much less expensive than hard-wiring.*

IMS Research reports that worldwide shipments of industrial products that are wireless-enabled are set to grow from an estimated 1.2 million in 2009 to more than three million in 2015, an average annual growth rate of 18%. In addition, Emerson itself announced in October 2011 that its customers had achieved 580 million total hours of wireless operation across 6100 networks, all based on w-HART.

In total, there are now around 17 vendors making w-HART products, including most of the major process control equipment manufacturers, with more arriving each month. So wireless has certainly gone mainstream.

Having a wireless infrastructure in a plant allows users to cover many applications, such as rotary kilns, safety relief valves and steam traps. Acoustic transmitters allow monitoring of steam traps and pressure relief valves, and they can easily be added to the wireless network. In the US, it has been estimated that there is a ⊠7940 (US$10,424) annual fuel loss per leaking steam trap. With wireless, it's easy to monitor critical steam traps for leaks, and start saving money immediately.

Wireless cameras can be used for monitoring people, hazardous locations or site security, asset tracking and safety mustering. Personnel tracking applications require a plant to have Wi-Fi coverage for tracking RFID tags, often accomplished by aerial mounting of wireless transmitters and receivers. Finally, wireless can be used for backhaul networks.

## Integration

Wireless field instrumentation devices integrate with the host control system through the w-HART gateway using native DeltaV node (version 10.3); OPC server connection; Modbus

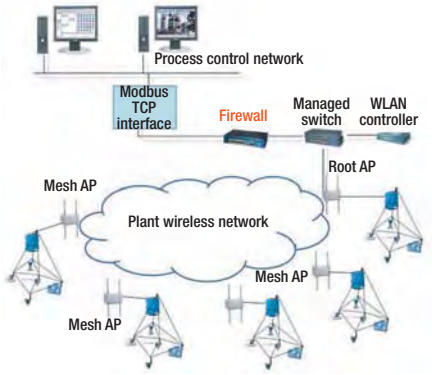


***Fig. 2: Wireless integration of the field network through Modbus TCP/IP interface.*** *For all host systems supporting Modbus TCP/IP or OPC.*

## Two case studies

Wireless solves problems that were too difficult or too costly to approach previously using hard-wired solutions. For example, a Middle-Eastern customer's remote pumps needed performance monitoring so that engineers could make informed maintenance decisions.

Previously, this customer had used a portable vibration analyser. The maintenance people would go to the pumps, collect data and make a decision. However, these pumps are sited in the middle of the desert with an eight-hour drive for access. Often, trips would be made and a pump would not be running, so data couldn't be collected. It was too expensive to install an online vibration monitoring system, so each pump ran until failure.

To resolve this problem, accelerometers fitted to the pump casing with the transmitter located close by communicate wirelessly to the w-HART gateway, which connects to the main control system over a WAN. This vibration data is communicated through the customer's automation system to the enterprise historian, which makes the information available to the reliability engineer at his/her desk through the corporate Intranet.

If customer's reliability engineer – 350km away – sees a pump problem on the historian, he or she, could phone the operations team, who could go out and correct the problem. Such difficult to traverse distances lend themselves well to wireless applications.

In another application at a Middle East wellhead, a customer has about 2500 wireless devices across many gateways collecting well data, something that was difficult and slow to deliver with a hard-wired system. Temperature devices show the customer whether product from each well is flowing. Previously, they had teams driving around the field in trucks and putting their hands on pipes to determine if the well was producing. This customer found that the payback period was nearly instantaneous, given the tremendous savings in maintenance and operating costs. Water wells are a similar application in the Middle East, and in many parts of Africa.

TCP/IP connection; AMS HART TCP/IP; HART port; and Modbus serial connection.

The wireless field network comprises several w-HART devices communicating in a self-organising mesh network to a w-HART gateway. For host systems not supporting w-HART native integration, Modbus Serial, Modbus TCP/IP and OPC DA can all be used to directly connect to the w-HART gateway. Figure 2 shows a wireless integration of the w-HART field network via a Modbus TCP/IP interface.

## Working with wireless

In brief, the advantages of wireless over hardwiring are:

- Less expensive;
- Faster to install;
- Easier to maintain;
- Less susceptible to wear and damage;
- Works in almost any terrain;
- Better for mobile assets (Fig. 5).

Instrument engineers know how difficult it can be to hard-wire a field instrument such as a flowmeter or a level transmitter, back to an automation or information system. One reason that w-HART has been adopted quickly by the process industries is because it's so easy to use.

With conventional hard wiring, precautions have first to be taken to ensure that it will be safe and will meet all hazardous-area requirements. Next, hard-wired cables have to be routed through cable trays and I/O concentrators, and through marshalling cabinets or junction boxes at great expense in labour and equipment. Such wiring can be difficult to maintain. Wireless obviates such difficulties, and specific advantages include savings in wiring and related components, fewer cabinets and conduit systems, less need for ancillary equipment such as space heaters and cabinet cooling systems, savings in time and money, as well as a smaller footprint and less weight.

Being a mesh network, w-HART wireless devices talk to each other at a range up to 100m, and autonomously form wireless communications paths back to a local gateway (Fig. 3). This makes for very easy and fast installation and commissioning.

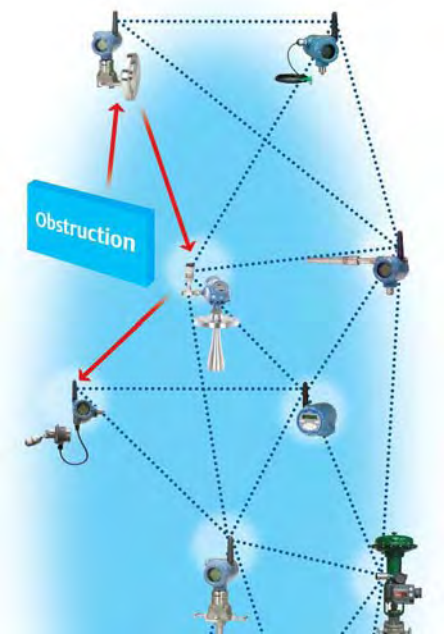


***Fig. 3: WirelessHART system architecture diagram.*** *Modern wireless mesh networks like WirelessHART allows users to link wireless sensors and instruments from multiple vendors to an automation system and/or information system.*

To commission an installation of wireless sensors, a user starts with devices closest to the gateway. As these power up, they become aware of their neighbours and make automatic connections. As more devices are added, they self-organise into mesh networks. A user doesn't have to tell the devices which other devices or gateways they should look for as these tasks are handled by the gateway, which may contain a network manager. Once the mesh has formed, any changes – such as by a permanent or temporary obstacle – will be managed by the devices themselves finding an alternative path to the gateway.

Overall, a wireless system is much simpler to modify than its wired counterpart. If a device changes or moves, it's not necessary to make wiring changes. If more devices are needed, another gateway can be added which will handle 100 more devices.

The redundant data paths inherent to a mesh network eliminate single points of failure. Real installations consistently demonstrate greater than 99% data reliability. w-HART uses a full mesh-topology as many as seven hops deep, so a costly infrastructure of multiple backbone routers installed throughout the plant within range of every wireless device is not needed. Running expensive hazardous area power supplies to backbone routers is also not required.

Available tools help plan and monitor the networks, analogous to Foundation Fieldbus and other hard-wired sensor and instrument systems. As with these hard-wired networks, the required tasks are similar – follow best practices, plan the network and it should work without problems.

## Built-in security

With no physical barrier surrounding wireless plant networks, wireless security is crucial to the successful deployment of field instrument networks and plant application solutions. Attack vectors include rogue APs, ad-hoc wireless bridges, man in the middle (Evil Twin, Honeypot AP, etc), denial of service (DoS), jamming (also considered DoS), plus reconnaissance and cracking.

However, w-HART communications use a security protocol with 128-bit encryption and 24 hex join codes, making hacking difficult. Another important security feature is that there is no IP address in the wireless devices, making it very hard for would-be hackers to make initial penetration. Other security measures include encryption, authentication, verification, key rotation and sequence number. Moreover, this security cannot be turned off, ensuring these measures are constantly active.

There are three main parts to a wireless defence-in-depth model:

PHOTO: EMERSON PROCESS

*Fig. 5. Personnel and tracking applications need a plantwide Wi-Fi coverage for tracking RFID tags, often accomplished by external aerial mounting of wireless transmitters and receivers.*

**Protecting networks** – Each mesh AP's digital certificate authenticates it to the wireless controller and allows it to participate in the secure network, so rogue or perhaps unauthorised APs cannot mimic genuine APs.

All lawful wireless network communications are encrypted to prevent eavesdropping or packet manipulation. Rogue APs cannot insert themselves in the middle of the wireless infrastructure or otherwise compromise the network. Ideally, wireless user access should be deployed with a wireless intrusion prevention system. Also, scan for rogue clients/APs.

**Control access** – Every user/device must authenticate with a centralised network domain authority. One way is to use an Authentication, Authorisation, and Accounting (AAA) server with the RADIUS authentication protocol coordinating access to the wireless network resources with the existing IT security infrastructure.

**Ensure client integrity** – Antivirus software must be installed to prevent any primary infection of the device. Good security practices should be in force. Ideally, control devices (wired or wireless) should have no email or Internet access.

WirelessHART field network inherent security features include:

- AES-128 encryption (NIST/IEEE compliant) for all communications within the device mesh network and the gateway;
- Individual device session keys to ensure end-to-end message authenticity, data integrity, receipt validation, and secrecy through data encryption;
- Hop-by-hop CRC and MIC calculations to ensure message authentication and verification as to communications source/receiver;
- Devices must have a pre-configured 'join key';
- White listing (ACL). If individual join keys are used, devices are explicitly given permission to join the network through the gateway or network manager via an ACL entry (also includes their globally unique HART address).

## Internal firewall

The connectivity from the w-HART gateway to the host system is secured by an easily configured internal firewall that allows only the protocols and ports required for the field solution to be enabled for communication. Ethernet-based protocols (Modbus, OPC, AMS, HART Port, https) all support SSL-protected communications, while the gateway's internal bi-directional firewall should default to 'reject all'. Note that the firewall needs no active management.

The above security features provide an easy to manage yet very robust communications system. Figure 4 shows a table of field wireless attacks against mitigating defences.

Some wireless solutions use an 802.11-2007-based Wi-Fi Mesh technology. Note, however, that 802.11 can be a security risk because wireless signals can be received by any commercially available 802.11 compliant device. By authenticating users before allowing them to access the wireless network, most attackers can be deterred, but it is recommended that all wireless data transmitted within the Wi-Fi mesh network, and between it and all client devices, should be encrypted.

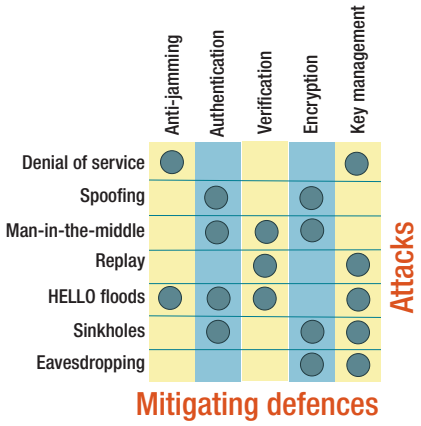


***Fig. 4: Plant wireless attacks against mitigating defences.*** *WirelessHART field network inherent security features include AES-128 encryption, individual device session keys, hop-by-hop CRC and MIC calculations, a pre-configured 'join key' for devices, and white listing (ACL).*

## Device diagnostics

The diagnostics information available from a wireless sensor or instrument is similar to that from a hard-wired fieldbus device or a conventional HART-based instrument. It is known that many users installed HART-based instruments mainly because of the extensive diagnostic information available, but estimates are that only 10% of installations actually use this information to the full extent. Many companies instead limit their use to handheld devices employed to manually calibrate and check field instruments during commissioning and calibration.

Part of the problem with wired HART devices is they have to communicate over a relatively slow 4-20mA connection, and special software is involved. However, with w-HART, the full range of HART diagnostics is available via the high speed wireless connection, and asset management software can extract the HART information. Yet w-HART is still familiar to users, who don't need to buy new tools or undertake more training programs.

Older, wired instruments can be added using a w-HART adapter, opening up a new area for maintenance and diagnostics. For example, many industrial and process plants have valves that have to be pulled from the line for an overhaul on a regular basis.

By adding a wireless adapter to a HART-based valve actuator, users can gain access to all the details needed to make an informed and proactive maintenance decision, instead of running to failure or performing unneeded maintenance. The same applies to flowmeters and related instruments.

## On the horizon

The HCF continues to develop new technologies for HART and w-HART. Coming in the near future is standardisation on HART IP, an abbreviation for Highway Addressable Remote Transducer over IP.

HART IP can use both TCP and UDP as the transport protocol. Typically, most wireless gateways are connected to automation and information systems using Modbus, but the available bandwidth is low. HCF envisages HART IP over an Ethernet physical layer, as a better way of getting data into automation and information systems.

Being able to obtain information on the performance and operational state of the mesh network is also desirable. HCF is responding to feedback from NAMUR testing to see if this is possible.

***David Walker*** *is Sales Director for Wireless Solutions, Middle East & Africa at Emerson Process Management.*

First published in the ***industrial ethernet book*** May 2012

# 60GHz Industrial Wireless: perfect for point-to-point

**Radio systems operating in the licence-free 60GHz band – yes there is one – have characteristics that make them significantly different from the more familiar hardware operating at the traditional 2.4 and 5GHz ISM frequencies. These altered qualities give 60GHz millimetric wavelength radios operational advantages not found in other wireless systems. For instance it is possible to achieve orders of magnitude higher link budgets than those possible with IEEE 802.11n and Ultra Wideband (UWB) systems – which translates into reliable and affordable wireless connections. Frank Ogden reports**

THERE ARE OCCASIONS when network cabling

PHOTO: HUBER & SUHNER

just can't do the job and the practical alternative is wireless Ethernet. Most industrial applications can be adequately served by IEEE802.11a/b/g/n WLAN hardware operating at 2.4 or 5GHz, certainly where the link distance is no more than a couple of hundred metres. Beyond this span, office-style wireless installation gets a bit more problematic with both performance and security concerns. If the link requirement is fixed point-to-point, then 60GHz operation might be an alternative.

With its own licence-free operational frequency allocation just like existing ISM, it is possible to set up and operate 60G fixed links to implement a plug and play extension to a network without wires. Of course a millimetric operational frequency brings with it specific media characteristics: the necessary patch antennas are both physically small and highly directional while oxygen spectral absorption restricts path length to about 1.6km. And of course the signal path must be entirely line of sight for the link to operate. Against this, the highly directional antennas make the signal – a beam – almost impossible to intercept and virtually immune to interference, either from all the usual industrial sources of EMI or from other inhabitants of the crowded lower frequency ISM bands. Other 60GHz operations are unlikely to interfere since narrow beam transmissions only 'illuminate' what they point at and, in any case, oxygen signal attenuation places a limit on interference from distant sources.

When it comes to possible data bandwidth, the RF baseband is 7GHz wide – getting on for 100 times wider than the channel bandwidth available under 5GHz ISM band operation. There is already plug and play 60GHz COTS hardware on the market which delivers true 100Mbps Fast Ethernet full duplex performance but the potential data bandwidth achievable is much greater, limited only by the confines of Shannon's Law.

Possible industrial applications include all those where one might presently use multimode optical fibre media but are precluded because a cable might get in the way, or otherwise be difficult or expensive to install. For instance, a 60G fixed link would be perfect for connecting up buildings separated by public roads, implementing low installation effort temporary links around tank farms and utilities and similar. Away from industrial plant, 60G is receiving serious interest for 'last mile' telecoms distribution, cellular backhaul and data centre connections, in fact anywhere which needs flexible, non-invasive point-to-point network segments.

PHOTO: HUBER & SUHNER

***When you can't dig up the road,*** *this 60GHz wireless bridge from Huber & Suhner implements a useful Ethernet-based line-of-sight wireless solution…*

## Friis' Law is good!

The high antenna gain typical for 60GHz systems comes in part from the Friis Equation for Path Loss. This equation, also known as Friis' Law, holds that as the frequency of operation changes, the effective area of any particular antenna varies proportionately to the square of the operational frequency. Put another way, this means that the gain possible from an antenna of any given size increases by the operational frequency squared. An antenna with an area of of just over 6cm$^2$ – a square an inch on a side – will have a gain of approximately 25dBi at 60GHz versus a gain of approximately 3dBi at 5GHz. This characteristic provides the mechanism by which millimetre wave systems deliver better link budget performance than systems operating at lower frequencies.

First published in the ***industrial ethernet book*** February 2011

## Airbus explores 60GHz for wireless IFE, WAIC

The Fraunhofer-Heinrich Hertz Institute in Berlin has carried out several projects with Airbus in the use of 60GHz technology for wireless in-flight entertainment as well as low rate WAIC (Wireless Avionics Intra-Communications) at lower frequencies for sensor and crew communication applications.

The basic idea of wireless onboard communication is to replace wired infrastructure by wireless links. The applications range from traditional in-cabin connections such as crew communication through new passenger services, and individual in-flight entertainment to use in sensor networks with hundreds of sensor nodes inside the whole aircraft structure. The motivation for these developments include the saving of weight and wiring harness maintenance, the higher robustness of a self organised wireless networks compared to wires and connectors, higher flexibility in cabin configurations, and the implementation of new functions such as sensors at moving or non-accessible parts.

The 60GHz frequency band is a most promising candidate. It provides several gigahertz of bandwidth for license-free worldwide use, and the high free space attenuation minimises interferences between different aircrafts and allows a reuse of frequencies, even within the same cabin.

Similar work is presently being done at Boeing, mainly in the use of WiFi, but 60GHz is also under investigation by the US plane maker.

**Dr.-Ing. Wilhelm Keusgen**

# Leaky feeder cables provide non-contact WLAN operation

**The Austrian plant manufacturer Berndorf Band produces endless transport and process belts made from stainless steel. The company has equipped its biggest production hall to date with the latest automation technology. The production plant makes 3.5mm thick steel bands with a length of up to 270m. Providing WLAN connectivity for mobile production cells along the length of the hall required a signal field with special properties. Peter Hallas explains.**

PHOTO: SIEMENS A&D

***Long job:*** *The steel belts are almost 300m long and have to be worked on with movable stations along their length*

THE PROCESS and transport belts produced at Berndorf Band are mainly used in the chemical industry, food & beverage production and by producers of timber and laminates. Each band is produced exactly to customer requirements from the raw materials which include stainless steel, carbon steel or titanium. Production takes place on two lines which simplifies handling of the band rolls and the individual processing stations. This makes production both faster and more flexible.

## Communication by radio

In building a new production hall of enormous dimensions, the company encountered a challenge in getting power to the individual workstation cells. These need to move about on rails with the actual position of individual stations determined by the length of the belt and the processing steps such as welding, grinding and polishing to be carried out. The power supply to these stations is provided from above by loop conductors so that the hall floor can be kept free of cables and wires, an essential advantage over the old solution with cable drag chains.

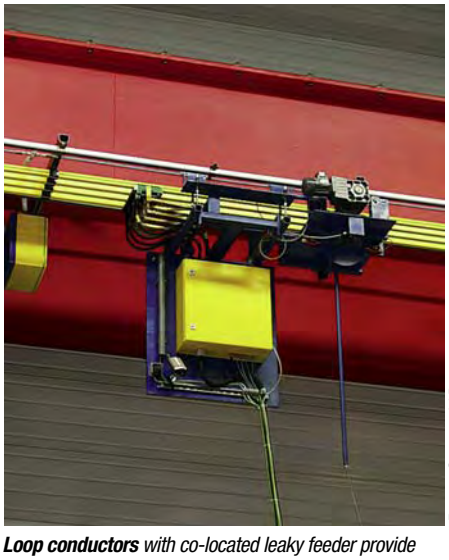

PHOTO: SIEMENS A&D

***Loop conductors*** *with co-located leaky feeder provide data and power to the production line from above*

Communication in the plants takes place via Profinet and Profibus. This also includes the entire safety-oriented communication sent using Profisafe. This data, which includes the Profisafe layer, has to be sent reliably to the work stations. To achieve this over the dimensions of the production required the installation of a leaky feeder cable, which is installed under the hall roof together with the loop conductors for the power supply. This special transmission line cable has deliberate discontinuities along its entire length. Each discontinuity acts as an antenna for the Scalance W access points. The leaky feeder cable provides a well-defined radio field along its length and can provide standard two-way WLAN communication at either 2.4 or 5GHz with reliability. Using co-located power cables and feeder, the actual radio path runs to just a few centimetres into the WLAN interface components mounted on the movable work station.

The non-contact data transmission takes place without wear and is thus maintenance-free, a major advantage over a wiper contact application.

The Industrial Ethernet system connects all the Simatic S7 standard and failsafe CPUs used in the system within the distributed automation architecture, as well as the distributed ET 200 I/O units, each with others, and with the master WinCC visualization system and the control panels. Profibus is used locally within the workstations. There are no differences between this radio variant and a fixed wired Profinet communication in the programming with Step 7 and the connection of the wireless components. An AS-i cable and a link to the Profibus via a DP/AS-i link reduce the wiring effort in the band coiling machine. The system has been running stably since commissioning.

***Peter Hallas*** *is with Siemens A&D, Austria*

First published in the *industrial ethernet book* September 2009

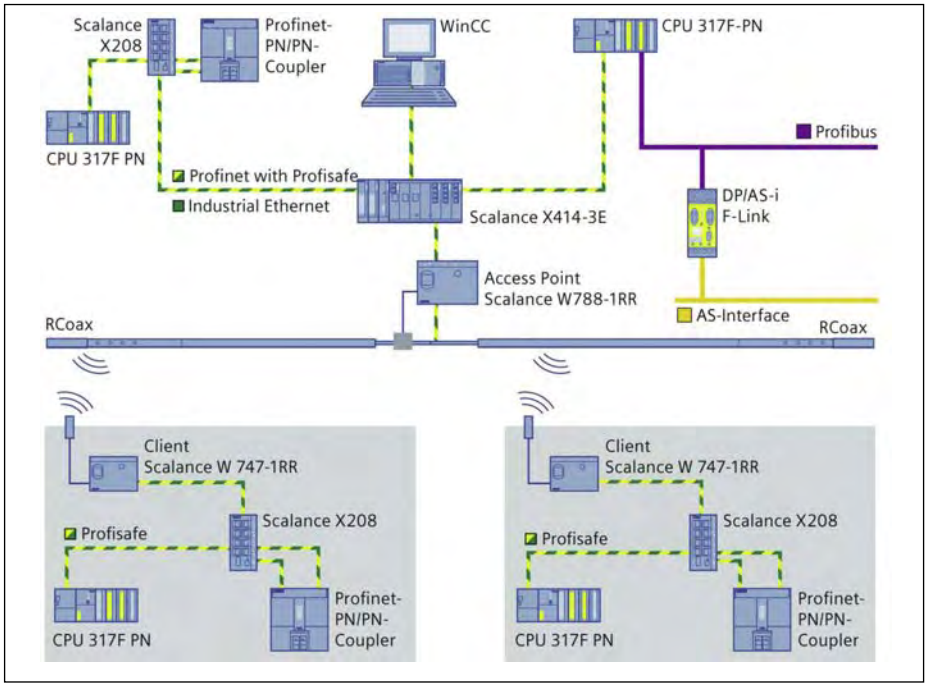

# Wireless avoids cable trouble on electroplating line automation

**Easily damaged and maintenance prone trailing cables on the suspended monorail of a plastic part electroplating line have been eliminated through the combined use of an industrial wireless LAN, Profisafe and an extension of IEEE 802.11 through iPCF. Christof Kreienmeier explains.**

ELECTROPLATING SPECIALIST, BIA Kunststoff- und Galvanotechnik GmbH of Solingen, Germany has, together with Aucos Elektronische Geräte GmbH of Aachen, planned, built and commissioned a new electroplating production line. Unusually, the suspended monorail tracks above the electroplating tanks are not powered via trailing or armoured wire cables, but communicate with the control system using Profisafe via wireless LAN (WLAN).

This approach was chosen because cables collect contamination and dust particles that end up in the tanks, and because the ever-increasing speeds of the trolleys (typically up to 120m/min) result in cable breaks. The resulting production losses can be very expensive. BIA suggested replacing the trailing cables with the only viable alternative – a radio-based industrial solution that had to be capable of providing cyclic and time-critical data traffic at predictable transfer rates. This would ensure reliable emergency stop circuit functioning.

## Controlled wireless field

Siemens supplied the wireless automation technology, which included rapid-roaming-enabled IP65-protected industrial access points (APs), plus leaky feeder cables (Fig. 1) to provide a controlled radio field along the tracks.

Aucos chose a failsafe controller that also operates the Profinet system connection to Aucos' control system. The controller connects to a managed switch having eight electrical ports that communicates with four industrial WLAN rapid roaming APs (IEEE 802.11-compliant) distributed in the field. These do not set up a typical omni-directional wireless network in free space, but use the special leaky feeder cables to create a uniform field strength along the four tracks that can be reliably monitored.

Their counterpart in each suspended monorail trolley is the directional transmitting and receiving antenna situated alongside the leaky feeder cable to allow optimal communication. In the trolley vehicle switch box, each antenna is connected to a WLAN device used as a gateway between WLAN and Profibus, forwarding the relevant Profibus and safety signals. Failsafe signal modules, motor starters and frequency converters for the suspended monorail trolleys completed the Siemens supplied equipment.

Despite the advanced automation technology, plastics electroplating still needs much manual intervention, so there are no barrier fences. Ensuring operator safety at all times depends on the reliability of the Profisafe communication components, as well as a clear prioritisation of safety signals.

In a fixed installation, the rapid roaming APs support the free movement of several nodes in the wireless network and can transfer them seamlessly from one AP to another. This mobile installation uses rapid roaming to ensure high signal transmission speed to and from the trolleys using the extended iPCF protocol (see box). It therefore ensures that the safety-relevant data is transmitted to the controller with a cycle time of only 16ms, allowing predictable response times well below 100ms – more than sufficient for time-critical emergency-stop decisions. In addition, process times can be controlled very precisely.

PHOTO: SIEMENS

***Fig. 1.*** *With their IP65 enclosures the access points are installed freely in the field and connected directly to the RCoax leaky feeder cable (blue).*

## Interference-free

The leaky feeder cable provides safety in a very different way. Aucos installed it parallel to the suspended monorails in strands of around 50m, ensuring that the field strength along the line provides reliability. Because the radio field is directional, it offers optimum reception, even with significantly reduced field strength. As a result, electromagnetic exposure is reduced – and only a few meters away, the radio network is practically too weak for external users to receive it. Advanced encryption algorithms mean that there is effective protection against unauthorised access.

The decision to use the 5GHz band was made to avoid interference caused by wireless sensors, Bluetooth devices and other devices that operate in the more common 2.4GHz band. The wireless components use both frequencies as standard.

BIA is satisfied and says that in future, new systems will always be equipped with this WLAN solution, being found to be far superior to the maintenance-prone trailing cable. It is quieter, a better system overview is now possible and the transport trolleys have virtually unlimited route options. One unplanned downtime can easily cost several thousand euros, but this system has eliminated cable breakages.

***Christof Kreienmeier*** *works for Siemens AG's Industry Sector – Industry Automation in Düsseldorf, Germany.*

First published in the ***industrial ethernet book*** July 2011

## Industrial Point Coordinated Function – speed equates to safety

In safety technology, the primary issue is transferring safety signals as quickly as possible to the control to head off critical situations. It is therefore necessary to assign the available bandwidth clearly, and to afford safety signals the highest priority. Siemens implements this by extending the IEEE 802.11 standard with the Industrial Point Coordinated Function (iPCF) a transmission protocol that – in contrast to the DCF method – governs access by multiple participants to the wireless network deterministically.

IPCF is a proprietary process that supports roaming in a WLAN infrastructure at times of 50ms. Both AP and client must support rapid roaming. In automation technology, it is often necessary to transfer process data and alarms in real time. Existing DCF and PCF roaming procedures don't allow such real-time data transfer without interruption.

iPCF is based on the centralised PCF procedure, already described in the first WLAN standard IEEE 802.11 in 1999. With PCF, the AP polls the connected stations and assigns time slices for data transmission. In contrast, iPCF works with a proprietary modified data frame protocol structure, so that iPCF clients can log on to an AP on which iPCF is enabled. The AP therefore becomes the ultimate authority and assigns transmission times to clients as time slices. This produces stable and predictable response times (typically 16ms with five to eight clients) with intensive use of the available bandwidth capacity in the radio network. For this reason, iPCF is very suitable for typical production fieldbuses that achieve their real-time behaviour through short cycle times.

Compiled and published by IEB Media, Bahnhofstrasse 12, 86938 Schondorf am Ammersee, Germany. 