WIRELESS STANDARDS DEMYSTIFIED

Author: Niek Van Dierdonck

Introduction

Wireless products and technology for sensing & control applications have become a reality. Analysts, technology providers and product integrators agree that widespread adoption of wireless technology is only a matter of time.

When will this happen? ... is the question everybody is asking. Timing of technology adoption is influenced by many factors. Amongst those factors the ongoing development and acceptance of technology standards is probably the most important one. Product integrators require technology standards because it provides product interoperability, a large body of knowledge and development sources, second sourcing flexibility, etc.

This white paper describes the different standards, their strengths and limitations and investigates the target applications that drive them.

This article provides the high-level overview that technology integrators are looking for to start their search for the technology that best meets their needs.

A very diverse application landscape

The average home user has fifty light switches in his home... a facility manager receives a daily status update of all 10.000 lights in the large office building... an industrial plant operator receives an alarm the very second that that the mains power has failed and that the Uninterruptible Power Supply has commanded all heavy machinery to switch to a safe state.

All three examples are typical sensor and actuator cases: the home switch triggers the home light, the office building luminaries deliver a remote status report, the factory Uninterruptible Power Supply changes the state of the machines. In all three applications devices communicate over a network. In all three applications wireless communication offers a huge cost saving opportunity (see inset “Cabling Cost...Candidly”).

Under the hood however, the three applications are very diverse: In the home application the main drivers are low-cost and low-power for the wireless switches and low-cost for the wireless light. Of course reliability is important too, but an occasional second pressing of the switch to make the light turn on won’t

create a critical situation. The situation is quite different in the office building. The facility manager usually guarantees a minimum service level to the building owner or occupant and relies on his automation systems to perform maintenance tasks. Failing systems causing a hick-up in the maintenance schedule can have huge financial impact. Worse, an intruder switching off all lights of a large office building at 6 PM on a dark winter day can cause panic and result in casualties: nothing less than a terrorist attack. So reliability is essential in commercial building applications. Moreover, low-power is probably even more important than in the home, because devices running on a battery will eventually fail and need costly labor effort to replace. Industrial automation is probably even higher up on the reliability scale: even a slight glitch in a safety application may cause fatalities. On the other hand, industrial automation is usually less cost sensitive than home and commercial building applications.

These are only a few of the parameters that define the diversity of sensor applications. Other parameters include latency of the communication (i.e. the time a message travels from the origin – the sensor – to the destination – the actuator), the number of nodes in a network, the complexity to install, commission and maintain a network, etc. It should come to no surprise that for such a diverse application space a *one-size-fits-all* strategy just does not work. Neither for the technology, nor for the wireless standards that specify how wireless technology works. Standardization organizations have understood that scoping is required to answer the vast diversity in requirements. Some technology providers have usually taken one or two fields of specialization in their quest to be excellent in one or two key areas rather than try to do it all equally well, but not well enough for every individual application. Integrators, OEM companies and users of the wireless technology are further away from the technology and generally are aware of their own key requirements much better than of the key requirements of completely different application. Therefore many people are often bewildered by the number of seemingly competing standards emerging.

The standardization organizations and technology providers have not done a good job in addressing the bewilderment. If anything, many have contributed to the general frustration by refusing to being vague about their application scope.

Hardware, embedded software, stacks, chips... are you still following?

The vital forces behind standardization are: interoperability across brands, second sourcing availability, competition between technology providers to drive prices down, compliancy with global regulations and the opportunity to tap into a large body of knowledge. But there is more. Some technology components are so expensive to develop that they can only generate an economic return through very high volumes. And when volumes need to be big, the presence of a global market is paramount. Standards are an excellent vehicle to generate global awareness and to prepare for such a global market for ramp up.

The basic architecture of a wireless sensor system consists of three layers, as depicted in figure 1. In communication protocol theory as much as seven layers are defined, yet for the scope of this paper it is sufficient to cluster them in just three.

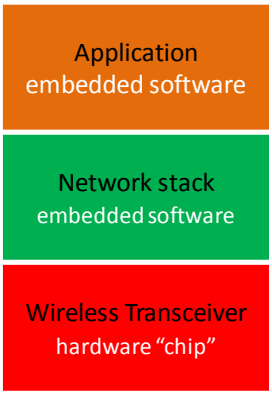

Figure 1: basic architecture of a wireless sensor device

The lowest block in the architecture is the *Wireless Transceiver*. The Transceiver is required to translate digital information (the bits and bytes) into a wireless *electromagnetic* signal that has the right format to be injected in the ether injected through the antenna at the transmitter-side and be reconstructed at the receiver end. The word *transceiver* is a concatenation of *transmitter* and *receiver*. In previous generations of wireless technology, you either had a transmitter for transmission only, or a receiver which was only capable of reception. Nowadays technology has shifted to combined reception and transmission devices as it enables powerful concepts that improve reliability and performance. A straightforward example is the *acknowledgement* principle: the receiver upon successfully reception of a message sends an acknowledgement to the original transmitter in order to confirm correct reception. Without this principle, the transmitter has no way of knowing whether the message ever arrived. Nowadays transmit- or receive-only technology is considered as unreliable and obsolete. The transceiver comes in the form of a chip, as depicted in figure 2.

Figure 2: a chip: containing the transceiver functions to send and receive data wirelessly

The transceiver has by itself no understanding of what a network is. A network is a universe where each device has a unique address, can join and leave the network (sometimes intermittently) and can send and/or receive messages from other devices. These are the responsibilities of the Network Stack. The Network Stack is generally implemented in software that is deeply embedded in a microprocessor. Very recently chip companies implement part of the network stack in hardware. The main motivation for this

migration is to boost the performance in specific areas such as ultra low power operation in application relying on energy harvested from the environment rather than on batteries.

On top of the Network Stack resides the Application. The Application encompasses the algorithms that the user intended in the first place: reading and interpreting temperatures, switching valves, etc. The application is usually implemented as a piece of software embedded in a microcontroller in combination with the sensor and actuator interfacing hardware such as the electronics required to transform the resistance of a PT-100 temperature probe into a meaningful digital signal for the microcontroller.

The wireless Transceiver chip

As established higher up, chips needs high volume sales to generate meaningful return, high volumes require global markets and for a global market to take off technology history has shown that the existence of a standard is essential. This was true for WIFI (wireless internet), technically termed IEEE 802.11 (a/b/g/n/...). Bluetooth chips as well are based on a standard defined in the IEEE 802.15.1 specification. For sensor networks the IEEE 802.15.4 (a/b) standard was set up in 2003. The fact that all three mentioned technologies where standardized under the wings of the same organization, the IEEE, proves that they were conceived for different purposes and not to compete with each other. Indeed, WIFI was conceived as an alternative to wired Ethernet PC communication: high data rate networks with a base station at the center and PC's nearby (i.e. a star-network topology). In order to achieve the application requirements WIFI consumes a fair amount of power – usually sourced from a laptop battery – and data rates degrade quickly when the distance to the base station increases.

Bluetooth was conceived with the mobile phone as the center of the universe: it connects the phone to an earpiece, to a GPS device and to a laptop. The Bluetooth data rate of 1 Mbps is large enough to carry voice, but is at least one order of magnitude smaller than that of WIFI. In return the power consumption is lower, most often sourced from a mobile phone battery. In general, the communication range is also smaller than that of WIFI, which is perfectly compatible with the applications as the phone is usually in the vicinity of the earpiece, the laptop and the GPS device.

Sensor applications have totally different requirements. Power consumption is probably the most apparent difference: sensors often have to work for years on a coin cell battery or on energy harvested from the environment through a solar panel of a vibration harvester. The battery cannot be recharged like a laptop or a phone's battery. Other sensor-specific application requirements are related to automatic network organization, reliability, communication range, the large number of nodes to be supported in a single network, etc. In return a lower data rate is generally acceptable because most sensors generate fairly small amounts of data and not even continuously.

For wireless sensor transceivers the dominant standard and probably only real standard is the IEEE 802.15.4 specification. The first version was ratified in 2003, with an update in 2006. Several vendors offer transceiver chips. Some of them are a minimal implementation of the standard. Others offer add-ons which are useful in some application segments, such as GreenPeak's own GP-2000 transceiver which has a lot of power reducing features destined at coin-cell and battery-less applications.

Table 1 lists the main parameters of the IEEE 802.15.4 and compares these to Bluetooth.

Parameter	**IEEE 802.15.4**	**Bluetooth**
Wireless frequency	2.4GHz / 868MHz / 915MHz	2.4GHz
Data rate	20 kbps up to 250kbps	1000 kbps
Typical average power consumption	1 µA	5000 µA
Network size	Up to 65536	Up to 8 nodes
Range	30-300ft (10-100m)	30-300ft (10m-100m)

There have been efforts to use Bluetooth and WIFI for sensor applications. In all the cases reported Bluetooth and WIFI were used in a non-standard way, in fact weaving the principles of IEEE 802.15.4 in their native implementation. It is nowadays widely accepted that the IEEE 802.15.4 offers the best basis for wireless sensor applications.

Besides the IEEE 802.15.4 standard, a number of technology suppliers have chosen to build a proprietary transceiver. The main motivation seems to be a reduction of the complexity and thus a potential lower cost point. It remains to be seen if a proprietary solution will ever reach sufficient volumes to actually reach that theoretically lower cost point. Additionally, reducing the complexity automatically goes hand in hand with sacrificing performance and thus limiting the applicability.

The Network Stack

In essence the network stack has two responsibilities. To start it is in charge of forming and maintaining the network. An important consideration in wireless network stack design is that is the ability to cope with the constantly varying quality of the wireless links between nodes. For example in a building automation application, the effect of people moving around has a formidable effect on the link quality, because when a person stands in between two nodes the link quality will reduce drastically. So the network stack needs to take into account that links can disappear at any moment, possibly isolating a network node or even a whole branch of the network. In response the network stack needs to re-organize the communication

routes through the network by establishing new links in order to provide uninterrupted connectivity to all parts of the network.

The other responsibility of the network is to ensure that messages can travel from a source node to a destination node in a reliable and efficient way. Efficiency here means that latency requirements - this is the travel time of a message - should be met and that bottlenecks in the routing of messages need to be avoided.

As established in the beginning of this paper, the broad application space has widely varying requirements and thus calls for flexibility in the communication technology. We also noted that the hardware transceiver cannot offer this flexibility. The network stack comes to rescue here, because a large part of it is generally implemented in software. And software as opposed to hardware does not have as high an up-front investment cost, meaning that a software investment can live with lower volumes than hardware and still lead to a healthy return.

The consequence of these economics is that today we see several Network Stacks standardized, some of them progress, others already completed. All the current standards build on top of the IEEE 802.15.4 specification. In other words these standards assume an IEEE 802.15.4 foundation and sit on top of it.

Figure 3: a view on the most prominent sensor Network Stack standards

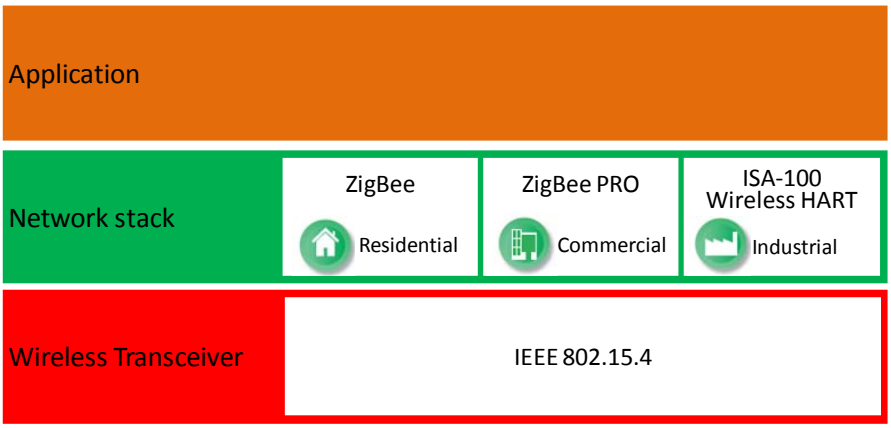

The ZigBee Alliance is an independent standardization organization that has a history stemming from the turn of the millennium. It is driven by a large group of technology providers and OEM companies alike. The most recent milestone of the alliance achieve at the end of 2007 was to finalize the specification of two Network Stacks: the ZigBee Network Stack and the ZigBee PRO Network Stack. In essence ZigBee PRO is a superset of ZigBee, adding functionalities related to the ability to scale up the network size and to better

cope with wireless interference from other technologies. From a usage point of view the ZigBee Network Stack is very suitable for residential “home” applications, where home networks typically contains 10’s or maximum a few 100’s of devices. The ZigBee PRO features make it especially suitable for larger application, very often in commercial building space. The drawback of ZigBee PRO versus ZigBee is that the extra features require a larger program memory size, which automatically translates into higher cost. In the extremely cost-sensitive consumer market every extra cost limits the likelihood of adoption. However, thanks to the ever decreasing cost of silicon, governed by “the law of Moore”, we predict that in the short term the cost difference between ZigBee en ZigBee PRO will be negligible and that most applications will adopt ZigBee PRO.

Although the ZigBee Alliance does not explicitly rule out industrial applications, a number of large industrial automation companies have identified the need for extra features which are not on ZigBee’s top priority list. The two most important “Industrial” features are deterministic latency and deterministic reliability. Latency is the time a message needs to travel from the source to the destination. If the source is a PLC and the destination is a machine than it is easy to see why a tight control over latency is important. That is why the standards that explicitly target industrial automation exploit the IEEE 802.15.4 feature called Guaranteed Time Slots to offer latency determinism. In different words, the IEEE 802.15.4 has a feature that allows to better control when a message will arrive. Guaranteed Time Slots are not exploited by ZigBee.

The second most visible add-on in industrial automation standards is related to reliability. Reliability is related to the availability or absence of a communication path between two wireless devices. The most important enemy of reliability is wireless interference coming from other users of the same wireless frequency band. The most notable interferers for IEEE 802.15.4 based devices that operate in the 2.4GHz frequency band are WIFI transceivers. Most interferers will not fully block out an IEEE 802.15.4 device, but will cause some wireless packets to get lost, regardless of the Network stack operating on top of it. The industrial standards provide a mechanism that allows packet losses to become evenly spread out over time, even if the number of lost packets will not substantially decrease. The result can be called deterministic reliability.

ISA-100 and Wireless HART are the two driving industrial wireless automation standards. ISA-100 is the brainchild of the Instrumentation, Systems, and Automation Society (ISA), a non-profit technical society for focusing on industrial automation. The ISA-100 is expected to deliver a standard specification in the course of 2008-2009.

Wireless HART is not a full industrial sensor protocol but an add-on to the old but very popular HART industrial (wired) bus standard for industrial automation. In essence Wireless HART provides an alternative to the wired message transmission protocol of HART.

As ISA-100 and Wireless HART are fundamentally solving the same problems, they have recently joined hands in an effort to examine whether both standards can be merged into one. In a first version they will most likely not be interoperable and need a *network bridge* – a translator between the two systems – to interface. A follow up version might define a common language.

The advantages of the industrial standards are not totally meaningless in commercial building automation, but probably not essential to it either. At the same time the industrial standard features add substantial cost, which residential and commercial application are not likely to accept as these markets and typically much more cost sensitive than industrial applications.

The table below lists some of the features of the standards discussed.

Feature	**ZigBee**	**ZigBee PRO**	**ISA-100**	**Wireless HART**
Transceiver technology	IEEE 802.15.4	IEEE 802.15.4	IEEE 802.15.4	IEEE 802.15.4
Support for wireless mesh routing	Yes	Yes	Yes	Yes
Ability to cope with very large networks	No	Yes	Yes	Yes
Latency determinism	No	No	Yes	Yes
Reliability determinism	No	No	Yes	Yes
Built in security features	Yes	Yes	Yes	Yes

Proprietary wireless technology

As in all field of technology there are proprietary wireless sensor technologies. We define proprietary as a technology which is dominated by a single company. Proprietary does not mean that the specification is not open, because sometimes it is, but a single company still controls the direction of the technology, effectively leading to a monopoly. Proprietary standards have often been designed around a single or a limited set of applications. In practice a proprietary technology develops much faster than a technology

standard because there is no need to reach consensus among different companies. Quite often the proprietary standard can be superior to standards when used within their limited set of target applications. Conversely it is uncommon that a proprietary technology is able to address the broader space of applications that a standard addresses.

Proprietary technologies are vulnerable, for two reasons: (1) the owner of the technology controls the specification and thus also the price, and (2) the customer depends on the technology owner for upgrades and uninterrupted sourcing. History has shown that a company owning proprietary technology has four options. The first option is that the technology is adopted by an industry leader and many times the technology owner is acquired by this industry leader. Indeed only large players can keep on sustaining a proprietary technology to the market. The second option is that the technology owner relinquishes its exclusive ownership of the technology, upon which the technology can become a *de facto standard* when other technology providers start picking it up and offer a second source of supply to customers. In the third option, which is in reality a further development of the previous option, the technology owner pushes the specification to an independent standardization institute, thereby increasing the chance of widespread market adoption. When none of the first three options produces, then the fourth inevitably will occur: the technology disappears.

The two most notable proprietary technologies in wireless sensor communication are called Z-Wave and Wavenis. Z-Wave has been developed by a company called Zensys. Z-Wave is targeted towards residential automation, such as exemplified by the support of a maximum of 237 nodes. This number is sufficient for homes, but is not suitable for larger *commercial* installations such as hotels and office buildings. As Z-Wave has developed before a standard was in place and because of its fairly limited specification complexity it has been able to generate substantial market traction. We expect that Z-Wave will encounter increasing pressure from the ZigBee standard as more ZigBee products appear in the market and as the higher performance features of ZigBee come almost for free.

Wavenis is another proprietary standard developed by Coronis, acquired by Elster in 2007. Wavenis has generated traction in Automatic Meter Reading applications, and is currently marketed for other applications as well.

Recent evolutions

Even within the boundaries of standards, technology providers discover differentiation opportunities. As an example GreenPeak has provided Transceiver and Network Stack technology compliant to the IEEE 802.15.4 standard and with additional functionalities for ultra low power applications. An ultra-low-power

application is an application that is able to live off a coin cell battery or off energy harvested from the environment through a solar cell, a vibration energy harvester or any other environment energy converter.

Another evolution that is likely to appear soon in standards is low power routing (LPR). In an LPR network battery powered devices are able to receive messages from nearby devices and forward these further down a longer communication chain. Standards offer this functionality only for mains powered devices, because a device is required to be in a continuous listening state, consuming a significant amount of power. LPR adds a time synchronization mechanism to the network, allowing devices to wake up simultaneously to initiate communication and avoiding the need to be always on.

Inset "Cabling Cost…Candidly"

Cables have not fundamentally changed over the last 30 years. The material cost of a cable is generally an insignificant part of the total cost. The installation cost is eating the bulk of the total cost. With the high labor rates in the Western world, the following unit costs are representative:

Total cost (incl. material and installation labor cost) for a stretch of 1 meter

- New home: $ 5
- Retro-fit in existing home: $ 50 (including restoration cost)
- New commercial building: $ 3
- Rewiring of commercial building : $ 50 (including temporary opportunity cost as building space is temporarily unavailability for occupancy)
- Industrial site or plant: $ 5 for non-critical installations up to $ 150 for critical applications in hazardous zones

Niek Van Dierdonck

GreenPeak - VP Strategy & Product Management
www.greenpeak.com <http://www.greenpeak.com/>
Lindestraat 19 - 9240 Zele – Belgium

Ph	+32 52 45 87 20	Direct	+32 52 45 87 28
Fax	+32 52 45 87 29	Mob	+32 475 95 60 62

About GreenPeak

GreenPeak is focused on ultra-low-power wireless sensor communication technology. Our products are designed according to leading standards such as IEEE 802.15.4 and ZigBee and used in devices and applications that need to run for 5 years or more on battery power or in a battery-less setting drawing energy from energy harvested sources. The four low-power techniques introduced in this paper are just a few of the considerations we have taken when designing our technology. They have proven critical to large scale adoption of wireless communication technology in most sensor applications.

Figure 4: the GreenPeak CM-08 ultra-low-power wireless module