

A survey of several low power Link layers for

IP Smart Objects

Internet Protocol for Smart Objects (IPSO) Alliance White paper # 6

> JP Vasseur, Distinguished Engineer, Cisco Paul Bertrand, Founder Watteco, Bernard Aboussouan, VP Marketing, Gainspan Eric Gnoske, Atmel Kris Pister, Chief Technologist, Dust Networks David Culler, Chief Technology Officer, Arch Rock Roland Acra, CEO, Arch Rock Allen Huotari, Technical Leader, Cisco

> > June 2010

1. Introduction and architectural consideration

One of the fundamental properties of IP is to be based on a layered architecture, which means that each layer can evolve independently without compromising the architectural model. In other words there is no layer dependency or so-called layer "violation". As far as the Physical media (PHY) and Medium Access Control (MAC) layers are concerned, also referred to as the Layer 1 and 2 (or PHY/MAC), IP networks are sometimes said to be "link layer agnostic" in the sense that new link layers (PHY/MAC) can be supported in the network without requiring any change of the IP architecture. The exact same protocol (IPv4) specified in 1984 that was initially made of low speed links (64 Kbits/s) now supports 40 GBits/s optical links and the number of media that are used in the Internet and IP networks in general is impressive: serial links, SONET, SDH, ATM, Frame Relay, Optical fibers, Ethernet, Wifi are only a few examples.

It is worth mentioning that the lack of layer dependency does not mean that IP simply "ignores" the link layer characteristics thus leading to sub-optimal exploitation of the link layer! Knowing the characteristics of the media is critical for proper IP operation. For example, the routing engine needs to know these characteristics in order to make appropriate routing decisions and calculate the "best" path according to some objective function and routing metrics characterizing the link in some manner. Several routing protocols support more than one metric and the network administrator can then decide to use a metric to reflect the bandwidth, another metric for the link quality, or even some administrative values. A second example is related to the failure notification so as to quickly compute an alternate path at the IP layer in the presence of a link layer failure. Modern routers can make use of a fast keepalive mechanism to detect a failure, or may have local PHY/MAC drivers capable of sending a quick notification to the IP layer in order to trigger a rerouting event.

With no doubt, one of the main advantages of the IP architecture lies in its layered model allowing for a great deal of flexibility. Without such flexibility along with other principles such as the end-to-end principle, the Internet would not be what it is today. IP Smart object networks are no different and it is indeed specifically important to allow for a wide range of a number of low power link layer technologies such as IEEE 802.15.4, IEEE 802.11 and other low power PLC link layers to mention a few, without the need for complex, hard to manage multi-protocol translation gateways that poorly scale.

The aim of this paper is to provide a technical overview of several of these low power wired and wireless technologies: IEEE 802.15.4, IEEE 802.15.4e, Low Power Wi-Fi, a low power Powerline Communication (PLC) technology known as WPC and the newly specified Homeplug green PHY. Several other low power link layer technologies are being defined by several standardization bodies thus augmenting the spectrum of choices for these networks (e.g. low power PLC technologies used for smart metering, ...) that will be covered in further revision of this paper. Note that some technical aspects described in this document are related to a specific implementation: still it was worth providing some details to show how low energy could be achievable on specific media. Furthermore, some technologies are still proprietary (this is explicitly indicated in this case) and on the path to standardization.

It is worth mentioning that there is sometimes a confusion between 6LoWPAN networks that refers to IPv6 networks over IEEE 802.15.4 links and IP for smart object networks (a more generic terms referring to IP-enabled networks made of a number of low power link layers), sometimes also referred to as LLNs (Low power and Lossy Networks).

In some cases low-power link layers were not designed with the intention of carrying IP traffic. Often this requires an adaptation layer between the MAC and the IP layer above it. For example, a PHY layer that is only designed to carry relatively small packets such as IEEE 802.15.4 may need a fragmentation/reassembly layer in order to support potentially larger IPv6 packets. The IETF 6LoWPAN specifies an adaptation layer for carrying IPv6 packet over IEEE 802.15.4 links. This adaptation layer is responsible for handling fragmentation and reassembly, IPv6 header compression, IPv6 mesh headers (see [6lowpan]). In order to connect this new class of low power link layers, the IETF ROLL Working was formed to specify a new routing protocol for LLNs called RPL, optimized for low power and lossy networks; in particular a new set of routing metrics characterizing low power and lossy links used by RPL has also been specified. See [IP-smart-object] for a detailed reference of these IP-protocols in IP smart object networks.

2. IEEE 802.15.4

The IEEE 802.15.4 standard ([IEEE-802-15-4]) describes a LR WPAN (Low Rate Wireless Personal Area Network). In addition to low rate the standard also attempts to achieve several goals simultaneously: extremely low cost, short range operation with a reasonable battery life. Finally, the networks should be simple to install and offer reliable data transfer.

The two major parts of the standard are the PHY and the MAC. These two layers are the common foundation layers of the OSI model and are found in almost all other communication protocols.

The Physical Layer: PHY

The *PHY* layer describes the modulation, operating frequency, over the air data rates, channels and other important aspects of radio operation such as receiver sensitivity and transmission power.

Frequency range

There are four frequency ranges that the standard defines (IEEE 802.15.4 C defines the Chinese band). The ranges are:

- * China: 779 to 787 MHz
- * **Europe**: 863 to 870 MHz
- * North America: 902 to 928 MHz
- * Worldwide: 2400 to 2483.5 MHz

Channels

The Chinese band allows for 4 channels with channel spacing of 2 MHz and center frequencies at 780, 782, 784 and 786 MHz. One channel is available for the European band at 868.3 MHz. Ten channels are available in the North American ISM band with 2 MHz channel spacing and center frequencies at 906, 908, 910, 912, 914, 916, 918, 920, 922 and 924 MHz. Finally, 16 channels are available in the worldwide band with 5 MHz channel spacing and center frequencies at 2405, 2410, 2415, 2420, 2425, 2430, 2435, 2440, 2445, 2450, 2455, 2460, 2465, 2470, 2475 and 2480 MHz.

Modulation

There are two DSSS (Direct Sequence Spread Spectrum) modulation modes and one PSSS (Parallel Sequence Spread Spectrum) modulation mode. These are:

<u>DSSS</u>

* BPSK (Binary Phase Shift Keying)

* O-QPSK (Offset Quadrature Phase Shift Keying)

<u>PSSS</u>

* ASK (Amplitude Shift Keying)

Spread spectrum means that the energy of the signal is spread out over the channel bandwidth. One of the benefits of this is that the signal is less susceptible to narrow band interferers. Phase shift keying is a modulation mechanism that changes the phase reference of the signal or carrier wave. Amplitude shift keying is similar to phase shift keying except that the amplitude of the signal is varied.

Bit rates

There are various bit rates within the channels and modulation modes. These are summarized as follows:

PHY Frequency Band	Channel(s)	Modulatio n	Bit Rate (kb/s)
868 MHz	0	BPSK	20
902 - 928 MHz	1 - 10	BPSK	40
868 MHz (optional mode)	0	ASK	250
902 - 928 MHz (optional mode)	1 - 10	ASK	250

PHY Frequency Band	Channel(s)	Modulatio n	Bit Rate (kb/s)
868 MHz (optional mode)	0	O-QPSK	100
902 - 928 MHz (optional mode)	1 - 10	O-QPSK	250
2400 - 2480 MHz	11 - 26	O-QPSK	250

Figure 1- Modulation and Bit Rates

Transmission power

Maximum transmission power is regulated by government agencies such as the FCC in the United States and ETSI in Europe. Generally, in 802.15.4 systems, a node must be capable of transmitting at least -3 dBm.

Clear channel assessment

The PHY needs to be able to detect whether or not another radio is transmitting and employ a method to avoid interference. The mechanism used is CSMA-CA (Carrier Sense Multiple Access with Collision Avoidance). In this algorithm the radio first listens for energy or modulated data on the air. If any is sensed the algorithm provides for random wait times (backoffs) to retry the transmissions.

The MAC layer

The Media Access Control (MAC) layer provides the network and higher layers an interface to the radio (PHY) layer. Its primary function is to limit when each node transmits on the shared media (the wireless channel) so that transmissions occur one at a time. Like most links, IEEE 802.15.4 supports a Carrier Sense Multiple Access with Collision Avoidance (CSMA-CA) mechanism. In this algorithm the MAC first listens for energy or modulated data on the air. If none is detected, it can transmit immediately. If the channel is not clear the algorithm provides for random wait times (backoffs) before retrying the transmissions. Variants of this approach are used in Ethernet, Wi-Fi, and many other networks because it is simple, completely distributed, and requires no global coordination. Alternative MAC algorithms, such as Time Division Multiple Access (TDMA) used in Bluetooth, portable phones, and WiMAX, arrange in advance each node to have transmission time slots to avoid collisions. Before transmitting, a node waits for its timeslot to arrive.

In low power wireless links the communication duty cycle is typically very low so potential collisions are rare and easily avoided. However, the MAC typically implements

additional power management functionality to allow devices to turn off their radios to save power. The MAC, or more appropriately the Media Management Control (MMC), instructs the PHY layer to turn on the radio when it might receive a transmission [JHui].

Each IEEE 802.15.4 device has a hardwired 64 bit unique address called a MAC address, also referred to as the long address. This is similar to the MAC address in an Ethernet or Wi-Fi interface card in a PC. IEEE 802.15.4 also permits devices to use a 16-bit 'short address', which need be unique over the PAN.

Much as with Wi-Fi WLANs, distinct 15.4 PANs may share the same airspace. With Wi-Fi these networks each have an SSID and, typically, security credentials; with 15.4 they have distinct PANIDs with associated security keys. Associating with a 15.4 network involves obtaining a PANID and presenting valid security credentials.

The MAC layer may provide certain management functions, such as to instruct a node to become a PAN Coordinator and allow other nodes to join with the PAN Coordinator to form a star or peer-to-peer network. In such a configuration, the PAN Coordinator is the central node within the network that manages *network* activities such as association. The PAN Coordinator can issue short, 16 bit, addresses to nodes that join with it.

Either the short and long address can be used to construct addresses that can be used for connectivity to IPv6 based networks. Specifically, there is a mechanism called stateless address auto-configuration, which combines the IPv6 network prefix and the MAC address to form a unique network address. Since the MAC and IPv6 addresses are related to each other there is a mechanism (6LoWPAN) that allows the compression of the IPv6 address that allows the IPv6 packet to be transmitted via an IEEE 802.15.4 frame (described below). DHCPv6 provides a means of dynamically allocating unique IP addresses. As an alternative to a PAN coordinator, these can be assigned so that unique short address is easily extracted from the IP address, allowing further compression.

There are two logical constructs that describe activities that the MAC can perform. These are MCPS (MAC common part sub-layer) and MLME (MAC sub-layer management entity). To broadly generalize the parts, the MCPS is concerned with transporting data and the MLME is concerned with managing the radio and network attributes. For instance, the MLME interface is used to scan channels, associate nodes and generate beacons while the MCPS is responsible for preparing data frames for transmission.

Frame

A unit or number of data bytes transmitted at the PHY or MAC layer is called a frame. A frame is very similar to the packet transmitted at the network layer. The frame can be a maximum of 127 bytes. Note that the IPv6 minimum required MTU is 1280 bytes: as previously discussed, the 6LoWPAN Working Group at the IETF specified an adaptation layer capable of performing layer 2 fragmentation as explained in the IPSO White Paper #2. It is generally composed of a *frame control field* which defines the type of frame it is and the addressing used. A *sequence* byte is next and is a numerically increasing count of

frames transmitted by the node (until the count reaches 255 which causes it to reset to 0 and begin the count again). The *addressing fields* are next and usually contain the source and destination addresses. The heart of the frame is next, the *payload* and finally, the *frame check sequence* ends the frame. An example (2003 IEEE standard) frame is illustrated below:

Bytes: 2	1	0 or 2	0 or 2 or 8	0 or 2	0 or 2 or 8	variable	2
Frame Sequence control number	Destination PAN ID	Destination address	Source PAN ID	Source address	Frame	FCS	
	number	Addressing fields				рауюао	(CRC)
		Ν	IHR			MAC payload	MFR

Figure 2 - General Frame Composition

The *frame control field* (bytes 0 -2) above is constructed to describe salient information about what kind of frame it is. Some of the information may be what kind of addressing is used (none, 16 bit or 64 bit), whether or not acknowledgments are required and if the frame source and destination addresses are within the same PAN. The figure 3 below shows the general structure of the FCF:

Bits: 0 - 2	3	4	5	6	7 - 9	10 - 11	12 - 13	14 - 15
Frame type	Security enabled	Frame pending	Ack. request	Intra- PAN	Reserved	Dest. addressing mode	Reserved	Source addressing mode

Figure 3 - Frame Control Field Composition

Beacon

A network can be constructed to implement beacons. Beacons are essentially time markers that are sent out on a periodic basis. They can be used to synchronize network activity and to provide a mechanism for contention free operation within the beacon interval. For example, certain time periods, or slots, can be allocated to a single node for reception or transmission. Beacons can also be generated when requested such as during an active scan when a node is searching for a PAN to join. In most network implementations beacons are not used with the exception of a beacon request during scan.



Figure 4 – Beacon Interval

The image above is a graphical representation of beacon frames occurring periodically over time.

Primitives

Both the PHY and MAC employ a notion called primitives. Primitives are services that a layer offers or uses. In some cases a primitive can be considered nothing more than a function call. These primitives provide higher layers with an API to the MAC layer. For instance, transmitting a data frame, in MAC terminology, is called *MCPS-DATA.request* while receiving a frame is called *MCPS-Data.indication*. These two examples are the real worker primitives in any network. There are many more primitives, especially on the MLME side, that allow for network management. For instance, there is *MLME-Associate.request*, which is called by a node in order to attempt to join (or associate) with a PAN Coordinator.

Summary

As mentioned at the beginning of this document, an IEEE 802.15.4 network is supposed to deliver reasonable battery life along with short-range reliable data transfer. Let's have a look at an example of range and battery life.

Range

Range is dependent on a multitude of variables. Among them transmission power and receiver sensitivity are primary, and their ratio is the link budget, often represented in decibels. A typical link budget might be 100dB, meaning that the receiver only needs

one part in 10 billion of the transmitted power to correctly receive a signal. In the simplest case, with no obstructions in or near the radio propagation path, the path loss increases as the square of the distance. For typical radios, this means that the range across an open canyon can be a kilometer or more! Under these conditions, an increase in link budget (either by increased transmission power or improved sensitivity) of +3dB can increase the range by 1.4 times while an increase of the link budget by +6dB, will double the range.



Figure 5 - Power vs. Range

Unfortunately, in most applications there are many obstructions in or near the RF propagation path. Obstructions in the path attenuate the RF signal, and obstructions near the path cause multi-path interference. This means that the actual received signal is much weaker than what is predicted by the simple quadratic path loss model. Empirical models typically model loss as distance to the third or fourth power. Still, this means that even in cluttered indoor applications, the RF signal can often travel tens of meters.

Battery Life

Most IEEE 802.15.4 nodes consume 10 mA to 30 mA in active transmit or receive mode. Battery life is heavily dependent upon cycle time between active modes and sleep or quiescent states. The simple graph below shows a typical cycling between sleep and active states of a sensor node. This cycling reduces the current consumption to an average that is tiny compared to the active modes.



Figure 6 - Current Consumption over Time

For instance, consider the scenario of node powered by a 3V battery that has a capacity of 650 mAh. If the average current consumption of the node is 50 μ A, then the battery life will be: 650 mAh/50 μ A = 13,000 hours/24 hours = 542 days/365 = 1.48 years. This requires that the active duty cycle be no more than 0.2%. Duty cycling for transmission is simple because the node turns on its radio when it needs to transmit. Duty cycling for reception is much more challenging because it requires that the device turns its receiver on just when some other device is going transmit to it. Several Media Management Control techniques for IEEE 802.15.4 have demonstrated the ability to minimize idle listening and achieve low duty cycles..

3. IEEE 802.15.4E

In 2007 the IEEE created the 4E working group to provide enhancements to the MAC layer of the 15.4 standard to better support industrial markets [IEEE 802.15.4E]. This has emphasized three major elements: media management to minimize listening costs, improved security mechanisms, and increased link level reliability through the use of multiple channels, especially in the narrow, lower frequency bands. Now, with the 4E standard approaching ratification, IP networks will be able to enjoy the performance seen in prior proprietary protocols. As with most IEEE standards, there are several variations and substantial configurability in 4E.

3.1 Media Power Management

Four low power MMC schemes are represented in the current 15.4e draft:

- DSME (Distributed Synchronous Multi-channel Extension) is an extension to the 15.4 beacon-enabled mode to support multiple channels. It divides up beacon periods into active and inactive periods. Radios are turned off during inactive periods.
- TSCH (Time Synchronized Channel Hopping): is an alternative to beacons that provides globally synchronized channel hopping. Radios are only on for assigned listen slots. Slot assignments are to be addressed by the Next Higher Layer (NHL), such as an IP adaptation layer.
- Low-Latency (LL): uses special short beacons and messages for star topology factory automation networks to optimize latency. Radios only need to be on in assigned slots. It is not intended to support multiple hops.
- Low-Energy (LE): minimizes radio listening without relying on time synchronization by associating a wake-up action with infrequent transmissions. Two approaches are supported. Coordinated Sampled Listening (CSL): has potential receivers perform very short listening probes to check for potential transmission and extends transmission to cover the probe interval. Receiver Initiated Transmission (RIT) has the transmitter listen for a ready-to-receive probe. In either approach, scheduling information can be exchanged to reduce transmission overhead.

With traditional CSMA MACs not augmented by low-energy mechanisms, leaf nodes can be very low power because they wake up infrequently to send their packets, but routing nodes or nodes that need to handle interactive requests need to leave their radio receivers on continuously in order to be ready to receive a packet other nodes at any time. This means that a battery operated routing node would have a lifetime measured in days or weeks, not years.

In contrast, in an 802.15.4E, even routing nodes can have radio duty cycles of less than 0.1%. This means that using 802.15.4E, IP routing nodes will be able to run on batteries for years.

For instance, in an 802.15.4E TSCH network using time synchronization, all of the nodes in the network know when their neighbors will be awake. A receiving node can turn on its radio for a fraction of a millisecond to see if a neighbor has something to say, and only leave it on if there is a packet to hear. In addition to the power advantages, many applications benefit from having an accurate time reference between nodes, for example to provide synchronous sampling of events (like acoustic localization), or to reconstruct the temporal nature of a distributed event (like a failure in a manufacturing plant).

Similarly, in an 802.15.4E LE network a receiving node turns its radio on for a fraction of a millisecond. Knowledge of when neighbors are likely to be awake is maintained locally and transmission need not wait for a scheduled slot, but they must be prepared to initiate a wake up. Similar, very low duty cycles can be achieved with LE as with TSCH.

802.15.4 System-on-chip parameter	range	
Transmit current, 0dBm output	8-30mA	
Receiver current	6-30mA	
Sleep current, RAM retained, timers on	0.4-2uA	
15.4E routing node	10-100uA	
15.4E leaf node	1-10uA	

Figure 7 Current consumption for typical 802.15.4 SOCs and 802.15.4E running on those SOCs

Typical current consumption numbers for 802.15.4 PHY and 802.15.4E MAC are given above. While there is a fairly wide range of performance between the different system-on-chip vendors [CC2530, DN2510, EM35x, MC1322x], the current consumption of all of these chips is dramatically lower than that for traditional wireless routing nodes.

3.2 Channel Hopping

Channel hopping allows for longer range and more stable links than single-channel protocols [Watteyne]. In a 4E network, neighbor nodes know the simple pseudo-random

hopping pattern of their neighbors. Each time a node transmits a packet to a neighbor, it uses the time of transmission, and the hopping pattern to calculate the channel on which the neighbor will be listening. The next transmission will go on a different channel. In this way, if there is external interference from another radio, or multi-path interference due to the environment, and the packet is not acknowledged, then the next transmission (or re-transmission) will not be subject to the same interference, and will most likely get through.

For environments with known interference, any combination of channels can be "blacklisted" so that they are not used in the hopping pattern. Although 15.4E radios are on so infrequently that they almost never cause interference to another radio networking, this blacklisting feature can provide confidence to RF system administrators who are nervous about adding more radios to their RF space.

The benefits of channel hopping depend on the specifics of the environment, but in highmultipath indoor environments the improvement in effective link margin can be 20dB or more [Werb].

The improvement in link stability has important implications for network traffic due to routing protocols. Since the loss of a single link can often cause a flurry of routing traffic in a network, the improved link stability of a multi-channel MAC can have a substantial positive impact on the power consumption of all of the nodes in the network.

3.3 Overhead Reduction

As shown in Figure 2 above, there are quite a few bytes in the 802.15.4 header. In many applications, these bytes remain unchanged in every packet that a node sends or receives. By improving the coding of the 15.4 headers, many of these bytes can be removed, saving energy for longer lifetime.

3.4 Increased Security

Security was built into 802.15.4 from the very beginning, with industry-standard algorithms (AES128, CCM*) for encrypting and authenticating every packet. In the original standard, acknowledgements to packets were not authenticated, allowing an attacker to falsely acknowledge a packet that had not been delivered. This potential weakness has been fixed in 4E.

4. Low Power Wi-Fi[™] (IEEE 802.11)

Until recently, Wi-Fi was not considered viable for battery powered sensor network applications. Wi-Fi silicon had been targeted at laptops and cell phones where the battery can be recharged after several hours of operations or at line powered devices such as access points. With the growing market for smart objects and wireless sensors, several companies have developed application specific integrated circuits that are optimized for sensing applications. These products achieve a similar power profile to other low power wireless link layers architectures while leveraging the huge installed base of over 2 billions Wi-Fi certified devices; a vibrant standard and industry alliance of close to 300 members; well proven encryption, authentication and end to end network security; mature network management systems; making it ideal for residential, commercial, industrial applications, Real Time Location Services (RTLS) and others.

A Power Efficient Protocol

The perception of Wi-Fi as a power hungry protocol arises from the manner in which Wi-Fi systems are designed and used in conventional applications today, and not from any intrinsic inefficiency in the IEEE 802.11 protocol [Wifi]. With energy consumption of 1-17 Joules per Mbyte transmitted for conventional high power Wi-Fi, depending on the protocols or the devices, even the least efficient device running on a AA battery, could transmit 1 Mbyte per day for 4 *years*. The most efficient device is potentially capable of operating for decades on a battery, however current battery technology is limited to 10-20 years even with no intentional current drain.

How is Low Power Wi-Fi Different?

While conventional high power Wi-Fi chips are optimized for fast response, low latency and high data rates, low power Wi-Fi chips are optimized for low power consumption, particularly when the device is in standby mode. For example, in conventional applications the device may actively listen to the channel even when no data is being transmitted to provide good response and low latency. Low power Wi-Fi minimizes power consumption when data is not being transferred. A representative operating scheme for a typical low-power application is summarized in Figure 7. After an initial set of tasks associated with startup, a low-power device spends the great majority of its time doing nothing. The device must wake up periodically to support various applicationrelated or network-related tasks. In the example shown in the figure below, the device sends a Simple Network Management Protocol (SNMP) trap message every minute to reassure the Access Point that it is active, preventing disassociation. Every 2.5 minutes, it awakens to send sensor data. Twice a day, the device sends a Configuration Trap to the SNMP server, to check for pending configuration changes (such as a new sensor time interval). Between each of these very brief operations, the device is in a low-power Standby state. Even during the periods in which it is awake, the device is actually sending or receiving data for only a small portion of the time.



Figure 8: Typical low-power Wi-Fi operation scheme

In order to minimize the power consumed during the vast majority of the time, in which no data is being transferred, the following changes to conventional design approach must be implemented:

- The device must be highly integrated to shorten connections, minimize capacitances and inductances and reduce overall energy consumption. All major system functions, including application programming, task management and network functions, radio management, encryption, MAC and baseband processing, and the radio transceiver itself, should ideally be incorporated on a single die.
- The device must be capable of flexible and rapid power management, including both fast-response states with reduced power consumption, and very-low-power Standby or Idle states employed when no activity is required from the device.
- It must be possible to awaken the device from its low-power states to a fullyoperational condition in a short time, either on a pre-arranged schedule or by an external input.
- Network operations must be arranged so that connection maintenance and remote device management are accomplished with minimal drain on the energy resources of the device.

With an integrated architecture, processors, other functional blocks and clocks can be independently gated off very rapidly to reduce power consumption when a specific function is known to be idle, with full function restored in one clock cycle when processing is ready to continue.

Compared to conventional Wi-Fi, new power saving states are introduced: The whole chip can be rapidly put into "Deep Sleep" state without loss of data, with the high-speed clock oscillator shut down; the system can recover from this "Deep Sleep" state back to full operation in a few milliseconds. An ultralow-power Standby state is also available to minimize energy consumption during long periods when it is known that the system will be idle. In this mode only the RTC block is active and power consumption is few microwatts. The RTC block timer permits the system to awaken at programmable intervals. Alarm inputs are also provided to permit unscheduled wakeup. Volatile data is lost in Standby, so some overhead is encountered in storing state and configuration information, and restoring the same when the system awakens, but the transitions are still much faster and less energy-intensive than the comparable functions when managed by a host CPU through a conventional serial, USB, or Cardbus interface.

These flexible low-power states allow system designers to take full advantage of such features as 802.11 power save operation, since the system can accurately awaken at expected beacon times, rapidly respond to beacon data and request delivery of buffered packets, and return to low power consumption as soon as the requisite functions are complete. In this fashion, low-power nodes can communicate with each other and with networked devices, with the power-hungry always-on packet buffer function taken care of by the powered Access Point. A device can also be awakened asynchronously using one or more input pins, to deal with alarms or unpredictable events that require prompt response.

A highly-integrated architecture thus permits the user to provide high performance and

low latency when data is actually being transmitted and received, while very efficiently minimizing the time spent consuming power when no data is being transferred. Rapid and efficient power management allows small packet transfers to be completed with only a modest increase in per-bit overhead compared to conventional systems.

The use of such an integrated system can provide substantial performance improvements in those parameters important for low-power operation, without compromising Wi-Fi compatibility. As shown in Figure 8, low-power Wi-Fi systems have much lower power idle states available, and can transition between active and idle states much more rapidly, than conventional Wi-Fi systems. They also achieve less remarkable but still substantial improvements in power consumption during data transfer operations, part of it related to the lower transmit power used for lower data sensor applications. As sensors are mostly in standby mode, the average power consumed is generally dominated by power consumption in this mode, and consumption in transmit and receive modes have little effect on the overall power consumption.

Parameter		Conventiona I Wi-Fi	Low-Power Wi-Fi	units
Power consumption	Standby / Idle	NA*	<4	μW
	Processor + clock sleep	13	0.2	mW
	Data processing	115	56	mW
Receive sensitivity, 1 Mbps		-91	-91	dBm
Time to wake from Standby		NA*	10	ms
Time to wake from p	rocessor+clock sleep	75	5	ms

*Not applicable: comparable state does not exist.

Figure 9: Comparison of conventional and low-power Wi-Fi typical performance.

Low Power Wi-Fi= Standard Wi-Fi, IP

Although, as described earlier, low power Wi-Fi chips are optimized for low power consumption typical of sensors applications, low power Wi-Fi conforms to the IEEE 802.11 standards and benefit from the standards' evolution in areas such as security (802.11i), meshing (802.11s) and QOS (802.11e). Low power Wi-Fi also takes advantage of the benefits conferred by the well established IP and Wi-Fi protocols, such as:

- Since Wi-Fi sensors use IP-over-Ethernet networking environment, there is no requirement for an expensive internetworking gateway to handle functions like network address translation or custom provisioning. Sensors are able to get unique IP addresses either static or through DHCP queries, and are able to support ARP for address conflict resolution.
- Sensor nodes can be managed and configured remotely using SNMP a well supported network management protocol. The node may have an SNMP agent that can respond to the SNMP manager's get and set commands, and send SNMP configuration traps to the manager.
- Support of well proven Wi-Fi link-layer encryption and authentication and related Wi-Fi Protected Access (WPA/ WPA2).

- In applications where other Wi-Fi stations constitute an important source of traffic, Wi-Fi sensors can benefit from 802.11's provision for collision avoidance. Every Wi-Fi packet contains a Network Allocation Vector (NAV), informing all stations that hear it that the sending Station wishes to reserve the medium for a time interval long enough to complete the current transmission. A low-power Wi-Fi device can use the NAV value received to reduce power consumption during the requested interval, and avoid attempting a transmission which is likely to collide with that of another Station. Devices that cannot interpret Wi-Fi packets may attempt a transmission during the period reserved by the NAV, and suffer a collision. The transmitted packet is likely to fail to reach its destination, and is either lost, or requires retransmission, in either case wasting valuable node energy.
- Wi-Fi systems benefit from a large installed base and consequent broad-based familiarity with configuration, use, and troubleshooting at the physical and link layers.

5. HomePlug Green PHY

The HomePlug Green PHY is a Low Power, Legacy Interoperable, Highly Reliable PLC from HomePlug.

Although power consumption and cost effectiveness are highly important features for any technology targeting smart energy, home control, home automation, etc. applications, they are not the only important features. Equally important to power and cost are:

- 1) reliability/robustness (which includes range/coverage)
- 2) coexistence and interoperability with existing (and standardized) technologies
- 3) the ability to scale and support many low rate devices on a shared medium (which translates to high, bursted, bit rates) with application throughput comparable to wireless technologies (e.g., 250 kbps)
- 4) time to market (which includes availability from multiple suppliers).

To address all of these criteria, the Home Plug Powerline Alliance is developing a new specification, tentatively called "Green PHY", which is a greatly simplified derivative of the existing and widely deployed HomePlug AV technology.

Overview of HomePlug AV

Before the details of the HomePlug Green PHY can be described, a very brief background of the HomePlug AV (HPAV) system is appropriate.

Figure 10 shows an architectural diagram of the HPAV system. The Higher Layer Entities (HLEs) above the H1 (Host) Interface may be any entity that can provide services to clients below the H1 Interface. The Data Service Access Point (SAP) accepts Ethernet format packets, so all IP based protocols are easily handled.

The Architecture defines two planes. The data plane provides the traditional layered approach with the M1 interface between the Convergence Layer (CL) and the MAC, and the PHY interface between the MAC and the PHY. In Figure 10, the MAC portion of the

control plane is labeled as the Connection Manager (CM) In the control plane, the MAC is a monolith without conventional layering. Although part of the control plane is in all stations, the Central Coordinator (CCo) entity will be active in one and only one station in a single HPAV network.



Figure 10 HPAV Architecture

The HomePlug AV Physical Layer (PHY) operates in the frequency range of approximately 2-30 MHz and provides a 200 Mbps PHY channel rate (with a 150 Mbps information rate). It uses windowed OFDM and a powerful Turbo Convolutional Code (TCC), which provides robust performance within 0.5 dB of Shannon capacity. Windowed OFDM provides flexible spectrum notching capability where the notches can exceed 30 dB in depth without losing significant useful spectrum outside of the notch. Long OFDM symbols with 917 usable carriers (tones) are used in conjunction with a flexible guard interval. Modulation densities from BPSK (which carries 1 bit of information per carrier per symbol) to 1024 QAM (which carries 10 bits of information per carrier per symbol) are independently applied to each usable carrier based on the channel characteristics between the transmitter and the receiver. Since the channel characteristics between each transmitter and receiver are unique, a "tone map" is maintained for each transmitter/receiver pair.

Therefore, the HPAV PHY provides for the implementation of flexible spectrum policy mechanisms to allow for adaptation in varying geographic, network and regulatory environments. Frequency notches can be applied easily and dynamically with region-specific keep-out regions settable under software control. The ability to make soft changes to alter the device's tone mask (enabled tones) allows for implementations that can dynamically adapt their keep-out regions.

The HPAV Media Access Control (MAC) provides a connectionless, prioritized Contention based service to support best-effort applications and applications that rely on prioritized QoS. It also can provide a connection-oriented Contention Free (CF) service to support the QoS requirements (reserved bandwidth, latency and jitter requirements) of demanding AV and IP applications.

The connectionless, prioritized service is based on contending for the powerline medium using Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA) technology in tandem with a brief Priority Resolution phase at the beginning of the contention window. This combination permits pending traffic at the highest priority level to access the

powerline medium ahead of pending traffic with lower priority levels (which is subsequently deferred). It is assumed that the traffic is marked using DSCP or 802.1D

The Contention Free service is based on periodic Time Division Multiple Access (TDMA) allocations of adequate duration to support the QoS requirements of a connection.

To efficiently provide both kinds of communication service, HPAV implements a flexible, centrally-managed architecture. The central manager is called a Central Coordinator (CCo). The CCo establishes a Beacon Period and a schedule which accommodates both the Contention Free allocations and the time allotted for Contention-based traffic.

As shown in Figure 10, the Beacon Period is synchronized to the AC line cycle and is divided into 3 regions:

- Beacon Region
- CSMA Region
- Contention-Free Region

The CCo broadcasts a beacon at the beginning of each Beacon Period and uses the beacon to communicate scheduling within the Beacon Period. By synchronizing to the line cycle, the beacons are extremely robust and reliable and provide stability of the periodic allocations.

The beacon provides announcements of where the beacon will occur over the next few beacon periods—i.e., beacon persistence—to enable continued communications by stations that occasionally miss a beacon.

Furthermore, the schedules advertised in the Beacon are also persistent—i.e., the CCo promises not to change the schedule for a number of Beacon Periods—and the schedule persistence is also advertised in the beacon so that a connection oriented transmitting station can confidently transmit during its persistent allocation(s) even if it has missed several beacons.

The CSMA periods are also persistent so that stations wishing to send CSMA traffic can do so even if they also miss a few beacons.

The MAC also maintains a clock at each station that is tightly synchronized to the CCo's clock (the CCo includes a timestamp in the beacon). This means that the entire HPAV network shares a common network clock for use by HLEs that may have tight timing constraints (e.g., to synchronize surround sound speakers).

Finally, synchronizing to the line cycle provides better channel adaptation to interference on the powerline (interference is typically synchronous relative to the line cycle), thus resulting in improved throughput.



Figure 11 Example of Beacon Period Structure

Overview of HomePlug Green PHY

To reiterate, the architectural objectives for the HomePlug "Green" PHY are to:

- 1) Maintain high reliability/robustness (which includes range/coverage) despite reduced complexity
- 2) Coexist/interoperate with existing, deployed, and standardized technologies
- 3) Scale to support for many low rate devices on a shared medium (which translates to high, bursted, bit rates) with application throughput comparable to wireless technologies (e.g., 250 kbps)
- 4) Provide greatly reduced power consumption
- 5) Be cost effective
- 6) Address time to market concerns (which includes availability from multiple suppliers)

Basing the Green PHY upon HomePlug AV was a reasonable choice considering its field-proven characteristics, large and worldwide installed base, well established Compliance & Interoperability (C & I) program, and multi-vendor ecosystem. In effect, HomePlug AV jumpstarts the Green PHY for time to market.

HomePlug AV employs three robust modes of communication, referred to as ROBO (ROBust OFDM). All ROBO Modes use QPSK modulation, along with a ¹/₂ rate Turbo

Convolutional Code. The Green PHY uses only these modes. ROBO is a form of redundancy coding and is used for low rate / high reliability transmission (including beaconing, data broadcast/multicast communication, session setup, and exchange of Management Messages). Although the use of redundancy coding introduces some inefficiency, the corresponding mitigation of and protection from impairments (e.g., interference and frequency selective fading) in helping prevent packet errors is a significant benefit.

Besides the robust nature of these modes, the further advantages of using only the ROBO modes for the Green PHY are as follows:

- 1) minimum impact on HPAV throughput
- 2) interoperable with current HPAV and future IEEE P1901 solutions baselined on HPAV (the Green PHY will be certified by HomePlug as a profile of P1901)
- 3) able to support multiple PHY Rates (3.8 Mbps, 4.9 Mbps, 9.8 Mbps)
- 4) reuses the entire 2 30 MHz band
- 5) reduces complexity of PHY (e.g., FEC, Analog Front End, Digital Front End)
- 6) low power modes via reduced duty cycle (i.e., "awake" time) can be achieved and still meet application throughput requirements

Since a high speed bursted mode is used, the target throughput (per device/application) of approximately 250 kbps can met by using <7% of the available medium access time (e.g., 2.8 msec within a 40 msec beacon period - which is two 50 Hz powerline cycles) at the most reliable (minimum) PHY rate (3.8 Mbps). Furthermore, since the Green device is only "on" or "awake" for this short interval, the estimated average power consumption is approximately 7% that of HPAV. Additionally, by further exploiting the duty cycle, 10 kbps could be provided using <1% duty cycle which provides opportunity for very low power modes. Alternately, higher throughput rates can be supported using more of the beacon period (but with a consequent increase in power consumption).



HP Green PHY Device Will Be in Sleep Mode >90% of Beacon Period

Figure 12 Example of Beacon Period Structure with Green PHY

The advantages of reusing the entire 2 - 30 MHz band is that optimal performance can still be maintained (i.e., maximizing the number of potentially usable carriers) and tone masking and amplitude (TX power) maps for regulatory compliance and coexistence with HAM bands are still supported.

But it is not only in the PHY where advantages are obtained. Since the MAC is also reused, not only can the Green PHY device receive HPAV beacons, it can also transmit HPAV beacons as a CCo. This means that a single, standalone Green PHY device can be a member of an AV network and can be the CCo in the AV network.

However, in order to reduce MAC complexity, a Green PHY device is only required to be a CCo in support of CSMA/CA only. In HPAV terms, the Green PHY device only needs to be a Level-0 CCo. This means that connection oriented communication cannot be controlled by a Green PHY devices, although it can participate in an AV network that permits connection oriented communications between HPAV devices.

A further simplification for the MAC occurs directly from the PHY simplification. Since ROBO does not rely on estimating channel conditions prior to packet transmission, it implies that there is no need to record/manage/maintain tone maps for each link in the network. The use of ROBO exclusively is a simplification that helps reduce memory size and code space.

Despite simplification of the MAC, reuse of both the multiple priority level QoS mode and Priority Resolution mechanism allows for straightforward low latency support of Demand Response / Load Shedding.

In closing, there is a noteworthy pair of MAC layer features which were added for Green PHY. In order to protect the high throughput of HomePlug AV devices, a new bandwidth "sharing" method was devised. Also, a power saving scheme was developed. Both of these features are described in greater detail below.

As mentioned above, since the required packet layer throughput for the applications targeted by HomePlug GP is <250 kbps, it is appropriate for these devices to be restricted to approximately 10 Mbps peak PHY rate. Nevertheless, in a powerline network supporting heterogeneous applications with extremes of throughput requirements (e.g., high definition multimedia content distribution and smart energy), it would be possible for lower speed applications to block higher speed applications from accessing the powerline medium. For example, a single packet being transmitted at 10 Mbps occupies the medium 20 times longer than that same packet being transmitted at 200 Mbps. While this alone is potentially harmful, consider the case when many lower speed devices are transmitting (perhaps a collection of smart energy devices) and in "competition" with a few higher speed devices. The quality of experience for the users of the higher speed devices could be adversely affected. Therefore, if HomePlug Green PHY devices were allowed to access the medium in an unconstrained manner, it is quite likely that HPAV devices could be adversely affected. To prevent this, a Distributed Bandwidth Control (DBC) function is included in the HomePlug GP specification. If pending high priority traffic is detected, then DBC will limit the aggregate HomePlug GP channel access time. Although this limits the effective MAC throughput rate it still provides ample capacity for smart energy applications. Furthermore, if a pending transmission from a Green PHY device will cause aggregate medium access to exceed 7%, then that device must defer transmission until a subsequent channel access opportunity. In the majority of scenarios, the local powerline medium will never be completely occupied and HomePlug GP devices may exploit unused access time without restriction simply by contending for channel access via the lowest priority level.

Reduced power consumption is a critical requirement for the applications targeting use of the HomePlug Green PHY. Significant power savings for HomePlug devices will be achieved via a Power Save mechanism developed specifically for Green PHY. It is helpful to describe the Power Save Mode by introducing the following terms:

- 1) Awake Window: The interval during which a Green PHY device is capable of transmitting and receiving frames. The Awake Window has a range from 1.5 msec to 2.1 seconds.
- 2) Sleep Window: The interval during which a Green PHY device is incapable of receiving or transmitting frames.
- 3) Power Save Period (PSP): The PSP is the summary interval of the Awake Window and Sleep Window. The PSP possesses a value of 2ⁿ (where n = 0 to 10) beacon periods and may be different for each device. The Awake State always occurs at the beginning of the PSP.
- 4) Power Save Schedule (PSS): The PSS conveys the PSP and the duration of the Awake Window.

It is important to emphasize that the duration of all PSPs are binary multiples of a beacon period. Therefore, a longer PSP will be a binary multiple of a shorter PSP. This allows the CCo to "align" the various PSPs and thus maximize the "overlap" or "coincidence" of the various Awake Windows. The wide range of PSP values enables the use of a long PSP for aggressive power savings or a short PSP for lower latency and response times. For example, if a latency of 30 seconds to 40 seconds is acceptable, a HomePlug Green PHY device can reduce average power consumption by 97% when compared to a device that is always in the Awake State.

6. A Low Power Powerline Communication by Watteco

Communication over power lines or *Power Line Communication* (PLC) has been a real challenge during the last 20 years. Although impressive progresses have been made on modulation and data rate, indeed we can find today 200 Mb/s modem supporting real time video, none of these technologies have kept the level of consumption at a reasonable level.

6.1 Why low power is important

Powerline technology consumption was usually not considered in the past as an issue due to the permanent connection to the mains.

These old times are now over: on one side, we live in the era of **Smart Grid and Smart Metering** where low power is mandatory for every system even those including Powerline technology. On the other side, the "less than ½ Watt" new regulation on standby power in EU for all devices will enforce new paradigms for consumer electronic products including Powerline technologies.

Some numbers first: the peak consumption of the best broadband Powerline modem is today around 6 Watts, while it is close to 3 W for a narrow band one.

In a home environment, 3 or 6 Watts are insignificant compared to a 2kW air conditioning unit. But embedding such modem in a meter or to replace a switch could be problematic due to power supply size and heat dissipation that are particularly challenging in the context of the form factor required by a meter or a switch.

As a consequence some companies are now working on efficient management of **sleeping modes and standby states**. Radio technologies have taken the same path in the past and, at first sight, it appears to be a reasonable solution (see the IEEE 802.15.4 section of this survey). The difference for Powerline technology is that sleeping modes unfortunately don't solve the power supply size problem, as injecting high power carriers is still required. It is a common "marketing mistake" to mix up **Power Consumption and Energy**. A low Energy system, expressed in kilowatt-hour, could still require a very large power supply in Watt, and accordingly would then suppose very long sleep mode periods and an extremely high latency.

Low emission level, efficient in radio thanks to relaying and efficient routing mechanisms, is unfortunately impossible on Powerline due to tough channels and low Analog Front End efficiency. **Improvements of** a**nalog parts** are probably the next challenge for Powerline leading companies.

It is in this context that Watteco developed WPC[™]: a Powerline technology using an ultra low power coupler, allowing Low Rate Wide Band Services, LRWBS, with **less than 20mW**.

6.2 Main features of WPC™ Technology

The WPC[™], for Watt Pulse Communication, technology includes adapted analog parts that excite resonance frequencies of the network producing pulses when connected to the mains. The pulses propagate over the power lines at a long distance (>1 km has been measured in a street lighting environment), keeping a very high signal to noise ratio.

This technology takes advantage of a physical natural resonance phenomenon: the ignition overshoot signal produced when loads connect to an electric network (the Pulse).

A pulse is a very short (a few nanoseconds) spike of energy produced by the networks in response to a load ignition or extinction producing local impedance variation. The emission of pulses can be triggered according to a controlled time schedule in order to communicate between two points of a given low voltage electric network.



Figure 13 Example of pulses datagram at different scales of time

6.3 The physics behind WPC

WPC[™] is based on the transient response of any electrical networks with voltage. By exciting network resonances frequencies, it is possible to create ultra short high level but low energy pulses (compliant with EMC regulations).

As a result the pulse's magnitude can be significantly higher than noise even after propagation and ensure a robust communication signal. The coupling device is very simple and the network reacts with its own resonances frequencies ensuring always optimal pulse emission and propagation whatever the impedance conditions are.



Figure 14 schematic of the pulse emission.

At selected instants defined by a signal command (2) a load (3) is connected to the network. As a result a transient voltage can be observed, caused by the disruption of charge on the established circuit;

Using this principle, pulse modulation frames can be created. The emission level can reach 15V with a typical duration of 40ns.

The next graph shows an example of pulse near emission point (C2). The curve C1 represents the command signal.



Figure 15: C2 Example of pulses close by the emission point, the curve C1 represents the command signal.

During propagation the pulses are transformed, filtered and energy is absorbed depending on the channel response. Figure 15 shows real pulses observed in a real home.



Figure 16: Real pulses observed after few meters of propagation.

The last version of WPC concentrates the energy in the 2-4 MHZ band known as LRWBS for Low Rate Wide Band Services.

The frequency shaping is obtained by sending, instead of one pulse for a bit, a multipulses frame presenting an inter-pulse average spacing of 333 ns (1/3MHz). See figure 18 for an example of multi-pulse bit representation.



Figure 17: A symbol is represented by a multi-pulses signal showing 11 unevenly spaced elementary pulses.

Uneven space create a frequency spreading in the LRWBS Band ensuring frequency diversity, better detection due to higher SNR.

The multi-pulses signals are organized in sequences (Frame) of 26 Bytes represented in figure 18. The absence of bits in the middle corresponds to the zero-crossing of the 50/ 60 Hz cycle.



Figure 18: A complete frame of 26 Bytes.

Main advantages of the Pulse emission mode:

- 1. **The low power consumption** (about 10 mW in emission mode) induced by the method of coupling and the simplicity of the reception. This point is crucial for Smart Grid and energy efficiency applications.
- 2. **The size** (approximately 5 cm²) of a full modem, induced by the network coupling method. This element is crucial to the widespread use of communication points in socket outlets, in meters and all small appliances.
- 3. Pulse modulation limits electronic components and signal processing. The **reduced Bill Of Material of a complete modem** ensures a very low cost.

6.4 Applications

Watteco's technology is critical for applications requiring small size and low power communication. It could be used in smart meters, as well as in-home gateway types of products to enable deployment of Command and Control solutions for load shedding of all equipment within the house. In complement to some higher rate and power consumption systems like HP AV or Green PHY it could be used to add Wake-On-Powerline services.

6.5 IEEE 1901 ISP coexistence.

Powerline is a media naturally shared between users. Until IEEE 1901's coexistence mechanism there was no standard describing the sharing rules between different technologies in the [2-30] MHz band. This lack of standard resulted in interferences and performances drops when two different technologies were sharing the same band and network.

IEEE 1901 defines the ISP (Inter System Protocol) technology, a mandatory mechanism for coexistence of up to 4 different technologies sharing the same media. ISP recognizes the LRWBS band and providing to these services a maximum of 2/8 time slots in the 2-4 MHz band.

Future version of WPC will comply with IEEE 1901 ISP protocol.

7. Conclusion

This white paper provides a short overview of five low power wireless and PLC technologies that can be used in IP Smart Object networks. Thanks to the layering architecture of IP, these technologies can be used in various areas of the network thus forming a seamless end-to-end IP network comprising a variety of link layers. With no doubt, a new revision of this white will soon be required that will show an increasing number of such low power technologies.

References

[6lowpan] IPSO White Paper #3 – "*6LoWPAN: Incorporating IEEE 802.15.4 into the IP architecture* " by Jonathan Hui, PhD, Arch Rock Corporation, David Culler, PhD, University of California, Berkeley and Samita Chakrabarti, IP Infusion.

[CC2530] http://TI.com/cc2530

[DN2510] http://www.dustnetworks.com/products/SmartMesh_IA-510_I/DN2510

[EM35x] http://www.ember.com/pdf/ember-EM300.pdf

[MC1322x] http://www.freescale.com/files/wireless_comm/doc/fact_sheet/MC1322XZIGBEEFS.pdf

[IEEE-802.15.4] IEEE Std 802.15.4[™] -2003, IEEE Standard for Information technology —Telecommunications and information exchange between systems— Local and metropolitan area networks Specific requirements—Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (LR-WPANs).

[IEEE-802.14.4E] IEEE 802.15 WPAN[™] Task Group 4e (TG4e) http://www.ieee802.org/15/pub/TG4e.html.

[IP-smart-object] *"Interconnecting Smart Object with IP: The next Internet"* (The Morgan Kaufmann Series in Networking) by JP Vasseur and Adam Dunkels, <u>www.thenextinternet.org</u>, June 2010.

[JHui] **IP is Dead, Long Live IP for Wireless Sensor Networks** Jonathan W. Hui and David E. Culler. In Proceedings of the 6th International Conference on Embedded Networked Sensor Systems (Raleigh, North Carolina, USA, November 05 - 07, 2008). SenSys '08. ACM, New York, NY.

[Watteyne] **"Reliability Through Frequency Diversity: Why Channel Hopping Makes Sense",** by T. Watteyne, A. Mehta, K. Pister, 6th ACM Intl. Symp. On Performance Evaluation of Wireless Ad Hoc, Sensor and Ubiquitous Networks (PE-WASUN), October 2009.

[WirelessHART] HART Protocol Specification 7.2, HART Communication Foundation, www.HartComm.org.

[Werb] **"Improved Quality of Service in IEEE 802.15.4 Mesh Networks",** by J. Werb, et al., Intl. Workshop on Wireless and Industrial Automation, San Francisco, CA March 2005.

[Wifi] "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications ", IEEE 802.11-2007, June 12, 2007.