An Introduction to IEEE STD 802.15.4

Jon T. Adams
Freescale Semiconductor, Inc.
2100 E. Elliot Road, MD EL536
Tempe, Arizona 85284 USA
+1 480-413-3439
jta@freescale.com

Abstract—The concept[1][2] of simple sensor nets, devices the size of ping-pong balls, sprinkled liberally on the ground, has been around for a long time. Some of the big challenges have always been cost and complexity, as well as power consumption. While there have been a plurality of proprietary wireless systems developed over the past decade or so for application to this problem, these systems have suffered from an inability to scale well in cost and network complexity. In 2003, the IEEE 802.15.4 standard was ratified, and almost immediately silicon manufacturers began producing compliant single-chip radios. Now, the next generation of transceiver is on the horizon, complete with microcontroller and FLASH memory, as well as the potential for various environmental sensors to be built right into the silicon itself. IEEE STD 802.15.4 specifies the RF, PHY and MAC layers, and there are a variety of custom and industry-standards based networking protocols that can sit atop this IEEE stack. These networking protocols allow the rapid creation of mesh networks that are also self-healing. With energy-saving features designed into the basic IEEE standard, and other possibilities applied by the applications developer, IEEE 802.15.4 radios have the potential to be the cost-effective communications backbone for simple sensory mesh networks that can effectively harvest data with relatively low latency, high accuracy, and the ability to survive for a very long time on small primary batteries or energy-scavenging mechanisms like solar, vibrational, or thermal power. This paper will look closely at the IEEE standard and the features that are natively part of the standard. Some of the various networking protocols that are proposed for or being used on top of this standard will be discussed, including ZigBee networking and IPV6. Practical sensor devices employing the technology will be analyzed and power consumption investigated. In addition, the ongoing updates to the standard taking place now within the IEEE will be discussed in light of their potential to make products developed to this standard even more useful to the sensor community.

TABLE OF CONTENTS

1. INTRODUCTION

In early 2003 the IEEE 802.15.4 standard was ratified after many years of effort. This standard represented a significant break from the "bigger and faster" standards that the IEEE 802 organization continues to develop: instead of higher data rates and more functionality, this standard was to address the simple, low-data volume universe of control and sensor networks, which existed without global standardization through a miasma of proprietary methods and protocols. The lack of a standard approach and protocol was seen as a major impediment to large scale manufacture of inexpensive silicon radios that would serve to drive down the cost per node of these networks.

At its heart the standard defines an RF and PHYsical layer with a Phase-Shift-Key (PSK) transceiver capable of over-the-air speeds of up to 250 kilobits per second (kbps), operating on a subset of 27 available radio channels in specific unlicensed (depending on the geographic region) 800, 900 and 2400 MHz bands. RF channel fading issues are moderated by the use of Direct-Sequence Spread Spectrum (DSSS). There are two channel access methods used: The first is Carrier Squelch Multiple Access (CSMA) with Collision Avoidance (CA), and the second is Time Domain Multiple Access (TDMA) using synchronization beacons and Guaranteed Time Slots (GTS). There are four packet frame types: Data, Acknowledgement, MAC Command, and Beacon. Each frame contains a receiver synchronization sequence, a packet length field, source and destination addresses, various frame control bits, the data payload, and an error-detecting Frame Check Sequence (FCS). Above the radio mechanisms, the Medium Access Control (MAC) layer generates network beacons if the device is a coordinator, synchronizes to others' beacons, supports Personal Area Network (PAN) association and disassociation, manages the channel access, handles and

[1] 0-7803-9546-8/06/$20.00© 2006 IEEE
[2] IEEEAC paper #1055, Version 2, Updated Dec 30 2005

maintains the GTS mechanism, and provides a reliable link between two peer MAC entities.

With the advent of the published standard, silicon manufacturers almost immediately announced the availability of first-generation solutions that demonstrated the simplicity and cost-effectiveness promised by the development of the standard. These first-generation solutions mate an IEEE-compliant RF data modem to a separate 8-bit microcontroller unit (MCU). The total number of passives required is generally under a dozen, mainly capacitors, and one crystal (exclusive of whatever sensor or control functionality is required). A typical IEEE 802.15.4-based, ZigBee-compliant device is shown in Figure 1.

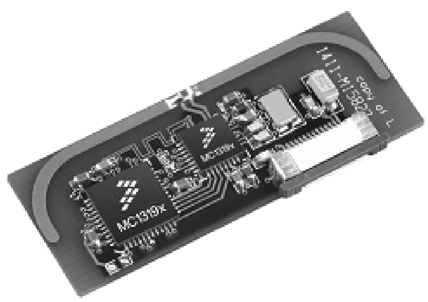

Figure 1. Typical IEEE 802.15.4 Device (including antenna, RF data modem, applications processor, all necessary passives and 16MHz crystal, I/O port flex connector, two-layer circuit board, about 15x40mm) - Courtesy Freescale Semiconductor.

There are currently at least a half dozen manufacturers of IEEE 802.15.4 silicon radios, with more expected to come on line over the next several years as the market matures. In addition, large test and compliance measurement organizations have instituted IEEE 802.15.4 compliance programs, and now provide a method of vendors to demonstrate independently compliance to the specification as an added assurance to the user.

The IEEE standard specifies and controls only the RF, PHYsical and Medium Access Control (MAC) layers. It suggests, but does not describe, networking methods and techniques. This was intentional – the desire by the working group was to develop a standardized protocol and radio, and allow various different approaches to the networking and applications functions. In light of this, several different networking techniques have been developed outside of the standard to take advantage of this radio. Of particular note are the ZigBee Alliance and its ZigBee mesh network, the Internet Engineering Task Force (IETF) "IPv6 over 802.15.4" working group, and, within the IEEE, the mesh networking task group 802.15.5.

Next generation silicon solutions are becoming available now– what was originally two packages (RF Data Modem and MCU) are now being combined into one package, and the power consumption and internal communications speed and efficiency are improved as a result of better optimized interfaces. In the near future, it is expected that silicon sensors and control elements will join the radio and processor on the same silicon, making it practical to have single-chip systems that are optimized for power and cost.

2. IEEE 802.15.4 STANDARD

IEEE 802.15.4 is a member of the IEEE 802 family, but it does not mean that all the features of all the other IEEE 802 standards are included or even desired for this low-rate, low-duty cycle standard. Control of expectations is probably one of the greatest challenges for any standards development organization, and this standard is no exception. The mission for this standard was to empower simple devices with a reliable, robust wireless technology that could run for years on standard primary batteries, was designed to allow a developer who had little ability or interest in the radio technology or communications protocol arts to effectively use and benefit from radios based upon the standard.

While intended to be simple, the standard does employ a number of features that are critical to a reliable, robust wireless link. It brings with it the ability to uniquely identify every radio in a network as well as the method and format of communications between these radios, but does not specify beyond a peer-to-peer communications link a network topology, routing schemes or network growth and repair mechanisms. It is intended for low-duty-cycle communications, which, when combined with the relatively high data rate, means that the period of time required to transfer a small block of data from one device to another is measured in milliseconds, and allows the device, if battery-powered, to spend most of its time sleeping in an ultra-low-power state. While feasible, high duty-cycle applications like voice or low-rate video cannot benefit nearly as much from power consumption savings as much of the power consumption savings is due to the device being able to spend most of its time in a quiescent state. Messages from one node to another can take advantage of receipt acknowledgement to improve transfer reliability. All receivers have some form of channel energy detection to detect potential users of the channel, and there is a Link Quality Indication (LQI) that provides the radios a metric on the signal strength/performance of that link.

Figure 2 shows the overall device architecture, with the RF channel represented as the physical medium, the PHY controlling the RF channel characteristics, and the MAC controlling the PHY. Since it is an 802 protocol, it incorporates the 802.2 Link Layer Control, a standard function of all 802 protocols. Like most other 802 protocols, it does not specify nor describe the upper layers of the OSI/ISO stack (depicted in Figure 2).

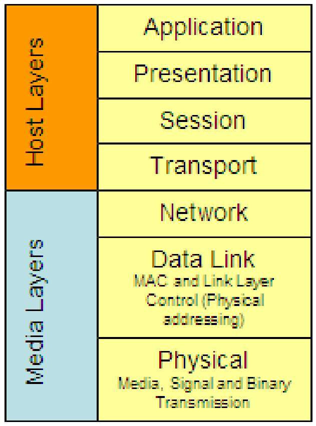

Figure 2. OSI/ISO Stack Model

In fact, comparing Figures 2 and 3, it is apparent that 802.15.4 does not even represent a full set of the Media Layers as specified in the OSI model – the Network layer is somewhat of "an exercise left to the reader".

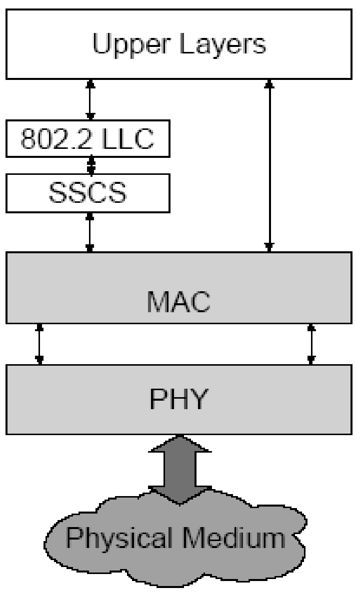

Figure 3. IEEE 802.15.4 Device Architecture

The standard defines two types of physical device. The Full Function Device (FFD) can talk to Reduced Function Device or other FFDs, while an RFD can talk only to an FFD. An RFD is intended for applications that are extremely simple, such as a light switch or an occupancy sensor; they generally communicate infrequently, spending most of their time in a quiescent state. An RFD may only associate with a single FFD at a time. Consequently, the RFD can be implemented using minimal resources and memory capacity.

There are three logical network devices envisioned: the PAN Coordinator, the Router, and the End Device. The Coordinator is built from an FFD, and is fundamental to forming a new network, It may have an overall knowledge of the entire network, whatever that may be, but at all times it is responsible for network address allocation. In a star topology, the "hub" is the Coordinator and it will provide message routing function as well. The Router is built from the same FFD the Coordinator is but, as the name indicates, its primary role in the network is to route packets. The third logical device is the End Device, and while it may be built from either physical device, it handles only communications and data transfer for itself. All logical devices may have non-network functions too, and in fact, it's generally expected that devices like lighting fixtures and other mains-connected devices are also Routers or the Coordinator.

The 802.15.4 standard is not a static animal – even now, there are amendments and additions, clarifications and simplifications being debated that will improve the utility of the standard. Additions being considered right now included increased data rate for the sub-1GHz bands, making optional some of the mandatory paragraphs to simplify what is a "standards-compliant" device, and other useful changes.

RF Link—The IEEE standard specifies the RF link parameters, including modulation type, coding, spreading, symbol/bit rate, and channelization. Currently, the standard identifies 27 channels spread across three different frequency bands, as described in Table 1. The work going on within the standards body right now is adapting the standard to new frequency bands being opened in Asia, as well as making the standard a bit more generic in its ability to be applied to both licensed and unlicensed frequency usage.

	Frequency Band (MHz)		
	868.3	902-928	2400-2483.5
# of Channels	1	10	16
Bandwidth (kHz)	600	2000	5000
Data Rate (kbps)	20	40	250
Symbol Rate (ksps)	20	40	62.5
Unlicensed Geographic Usage	Europe	Americas (approx)	Worldwide
Frequency Stability	40 ppm		

Table 1. IEEE 802.15.4 Frequency Bands, Channelization and RF Parameters

The standard is aligned to specific unlicensed bands available in different geographic regions. The single 868.3 MHz channel for use strictly in the European Union is limited to a 0.1% transmitter duty cycle by regulation. The 900 and 2400 MHz bands have no regulatory duty cycle restrictions.

The nominal transmitter power output specified is 0.5mW (-3dBm), again to limit power consumption but also because the standard is expected to be used in short-range (10-50 meter) applications. However, it is allowed to increase the output power through external amplifiers to whatever the regional regulatory limits are. Interestingly, some of the earliest adopters of the technology have done just this in order to augment or replace the more expensive radios often used in Supervisory Control and Data Acquisition (SCADA) applications.

The nominal receiver sensitivity is specified by Packet Error Rate (PER). The specification requires 1% PER at -85dBm receive power level for the 2400 MHz band and -92dBm for the sub-GHz bands (as measured at the chip's antenna terminals). This represents a receiver with a noise figure substantially worse than 20dB, though most radios in production are between 7 and 9 dB better than this. Thus, manufacture of the radio receiver using low-cost CMOS processes is very feasible. Also, receiver noise figure is a significant fraction of the overall receiver current, so while an extremely sensitive and high performance receiver is not impractical from an engineering point of view, it may act to substantially reduce battery lifetime in a portable application.

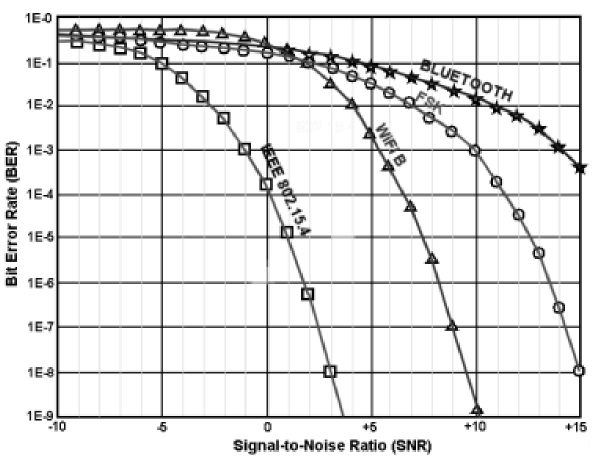

Figure 4. 802.15.4 Performance vs. Other Common Data Radio PHYs

The modulation mode used by 802.15.4 is phase-shift-key (PSK) based, chosen because of its strong ability to be recovered even in very low signal to interference environments. PSK modulation has been employed by NASA for decades in deep space mission telecommunications and it forms the heart of most of the high performance modem standards in existence today – but its use in most simple commercial and consumer wireless applications was traditionally limited due to complexity and cost, with Frequency-Shift Keying (FSK) preferred due to its relative simplicity and the existence of standards like Bell 202 and CCITT V.21. As Figure 2 indicates, the Bit Error Rate performance of PSK with respect to FSK (as well as the modulation techniques of other standards) allows 802.15.4 to enjoy significantly better link margins than other common wireless systems, giving it added robustness in noisy or marginal propagation environments.

The channels below 1 GHz currently use BPSK modulation (Binary PSK), while the 2400 MHz band employs O-QPSK (Offset Quadrature PSK). QPSK is spectrally efficient, but requires a linear transmitter due to the state transitions through zero. So, the developers chose O-QPSK which avoids the zero state and thus allows for a constant envelope transmitter, significantly decreasing transmitter complexity and inefficiency.

Finally, 802.15.4 employs direct-sequence spread spectrum to provide coding gain and added resiliency against multipath. Four bits are packed into a single symbol, then the symbol is coded with a 16-chip sequence per symbol are used for the sub-GHz frequencies, while 32 chips per symbol are used for the 2400MHz band. Interestingly, the spreading code used allows a standard FSK receiver to successfully demodulate the transmitted signal, although at significantly reduced link margin. However, in particularly cost-sensitive applications, it may be advantageous to sacrifice link margin for cost.

PHYsical Layer—The 802.15.4 PHY contains specific primitives that manage the radio channel, and control packet data flow.

The PHY uses Carrier Sense Multiple Access (CSMA) with Collision Avoidance (CA) to access the radio channel. This means that a radio with data to transmit will first listen to the channel and if the channel is clear, then transmit its packet. However, if the channel is busy, either due to another 802.15.4 station transmitting, or due to interference from a non-802.15.4 station (microwave oven, Wi-Fi access point, etc.), the radio will hold off from the channel for a random period of time before again checking the channel for occupancy. In a system where all stations can hear one another, CSMA-CA can provide nearly a 36% channel usage, but in practical environments where all stations cannot hear one another, the channel usage efficiency is as low as the traditional ALOHA mechanism, about 18%. Again, this was understood when the standard was created, and is acceptable given the requirements for system simplicity.

This is a packet radio specification, and the PHY defines 4 different frames that have unique functions: Data, Acknowledgement, Beacon and MAC Command.

As an example, the Data frame is depicted in Figure 5. Its structure is similar to the other 3 frame types. The Synchronization header (SHR) contains a preamble sequence (32 bits, or 4 octets) to allow the receiver to acquire and synchronize to the incoming signal and a start of frame delimiter that signals the end of the preamble. The PHY header (PHR) carries the frame length byte, which indicates the length of the PHY Service Data Unit (PSDU). The SHR, PHR and PSDU make up the PHY Protocol Data Unit (PPDU). The PSDU contains the MAC Header (MHR), which has two frame control octets, a single octet Data Sequence Number, good for reassembling packets received out of sequence, and 4 to 20 octets of address data. The MAC Service Data Unit (MSDU) carries the frame's payload and has a maximum capacity of 104 octets of data. Finally, the MPDU ends with the MAC Footer (MFR), which contains a 16-bit Frame Check Sequence.

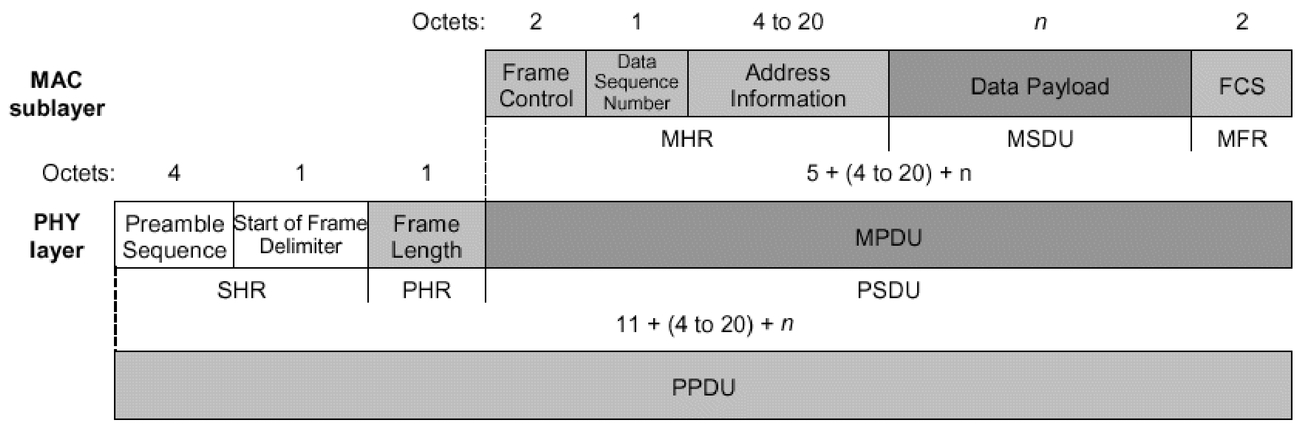

Figure 5. IEEE802.15.4 Data Frame

The Acknowledgement frame is used by a receiving station to "acknowledge" to the transmitting station that a data packet was received without error. The Beacon frame is used by stations that may be implementing significant power saving modes, or by Coordinator and Router devices that are attempting to establish networks. The MAC Command frame provides some unique abilities to send low-level commands from one node to another.

All IEEE 802.15.4 devices have a unique, 64-bit address. This long address similar to the well-known MAC address used in a 802.11 wireless card or 802.3 Ethernet NIC card. However, in complex networks moving small blocks of information, header size is reduced by allowing devices that join an existing network to "trade in" their 64-bit address for a 16-bit local address. This makes in-network communications more efficient and substantially shortens the packet length. The PAN Coordinator is tasked with handing out the short address when a device joins its network.

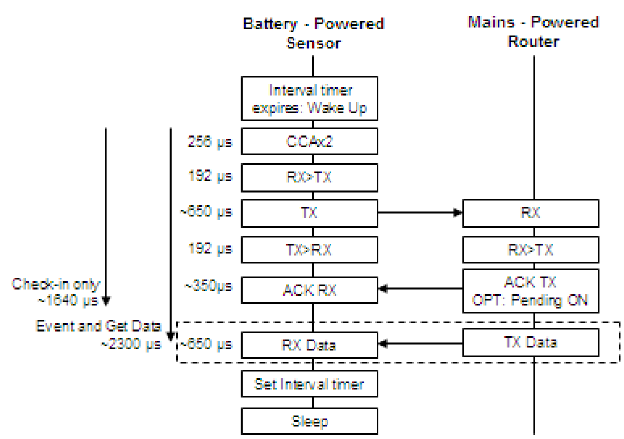

Figure 6. IEEE 802.15.4 Timing and Minimum Latency

The PHY manages all symbol and bit level timing, as well as transmit-receive switching times, intra-packet timings and acknowledgement delays. The PHY can be configured to automatically acknowledge (or not) every packet received successfully, depending on the application.

Figure 6 demonstrates the timing involved in a data transaction between two devices. The sensor wakes up either on an event or at the end of an interval, checks the channel, transmits its message, awaits the acknowledgement, then may go back to sleep or first receive data intended for the node before going back to sleep.

Medium Access Control (MAC) Layer—The 802.15.4 MAC contains over two dozen primitives that allow data transfer, both inbound and outbound, as well as management by higher-level entities of the RF and PHY. In all systems currently on the market, the MAC is implemented in software that runs on some sort of MCU core, but over time it's practical to see that as the standard is proofed out that the MAC could be implemented in a state machine or an embedded core exclusively dedicated to MAC functions. Since system cost and power consumption will remain driving factors for the standard and products based upon the standard, the market will have a strong influence on the systems architecture over time.

The MAC sublayer provides access to the upper layers through two Service Access Points (SAP). Data is managed through the MAC-SAP, while control and monitor functions are accessed through the MAC Layer Management Entity interface, called the MLME-SAP.

The MAC layer is responsible for generating network beacons that allow devices to find an existing network, or in the case of TDMA networks, that provide a timing indication for client devices to access the channel during both contention-based and contention-free periods. Most networks that employ a number of mains-powered (or other source of permanent power) routers probably will use the network beacon for network discovery alone. This beacon may be set in increments from approximately 15.83ms to over 4 minutes, as defined by the equation

$$\text{Beacon Interval} = 15.83\text{ms} * 2^n \quad \text{Where } n = 0 \text{ to } 14$$

The other purpose for the beacon is to signal timing in the operation of TDMA-based networks. Especially in entirely battery-operated networks, it was envisioned that all devices would normally be in a quiescent state, and when an internal timer expired a device would wake up to hear the beacon of its neighbor; the beacon begins an interval called the superframe interval, which provides not only GTS intervals for prearranged traffic but also contention-based interval where any device can "vie" for its neighbor's attention. Like the beacon intervals described by the equation above, the superframe interval is selectable by the network coordinator.

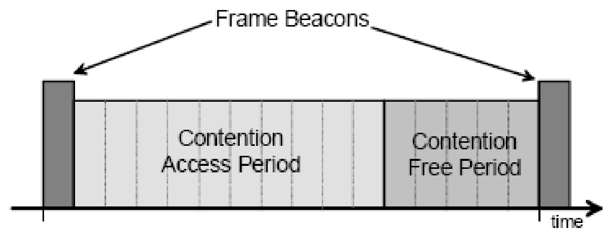

Figure 7: Superframe Structure with GTS Intervals

Figure 7 shows a generalized representation of the superframe interval in-between the network beacons. In the particular case shown in the figure, the superframe interval is equal to the beacon interval. In all cases, there are a total of 16 equal slot times available, their duration dependent on the length of the superframe interval. A device may be allocated one or more GTS intervals in order to transfer network traffic – during that time no other device may use the channel. When a GTS is used, the network does not use the CSMA-CA channel access scheme, hence the phrase "guaranteed time slot". The contention free period always follows the contention access period, where devices not having a prior reservation for a slot time may use the CSMA-CA algorithm to access the channel and pass their traffic. Once the contention free period is completed, the device transmitting the frame beacon may become quiescent, saving energy until the beginning of the next beacon interval.

The MAC layer manages the ability for a device to find a network, to associate to that network and to disassociate as necessary. Upon power up, an upper layer entity commands the transceiver to begin a scan on each channel available in a quest for an existing network. If the device is based upon a FFD, the FFD may attempt to establish its own network, if none can be found. However, assuming the device finds an existing network (discovered by listening on each available channel for an 802.15.4 beacon), the device will attempt to associate to that network. If the network allows the device to associate, then a message is passed up through the MLME-SAP to the network layer above, and that network layer manages the exchange of the device's 64-bit IEEE address with a suitable short address according to the PAN Coordinator's requirements. If a device is required to disassociate itself from a live network, the device will receive a Disassociate command from the PAN Coordinator, with the command initiated above the MAC layer.

3. Networking Techniques

The IEEE 802.15.4 specification provides guidance on possible network types; however, in terms of specification it codifies only tools that are necessary for formation of a network, but of unspecified topology or usage. Figure 8 shows two of the suggested types. The first, a star topology, is common to 802.11 and other host-client networks. All messages from any client device must pass through the hub (PAN Coordinator). The second type, the peer-to-peer, allows each device to communicate directly with peer devices and at its simplest defines direct communications between two devices. However, this method also may be used to create a mesh network if a higher layer entity chooses to do so.

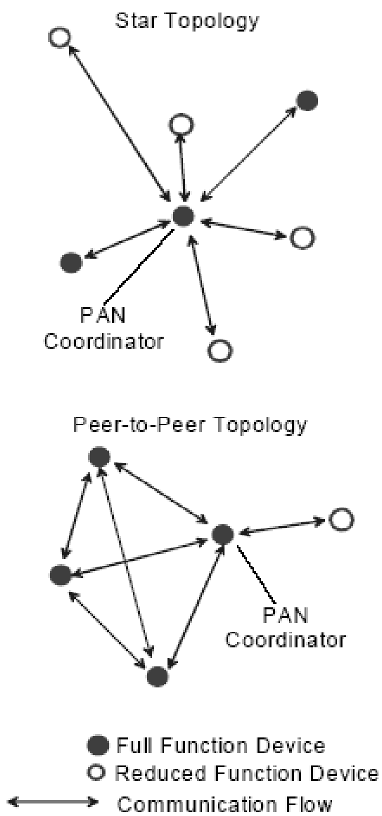

Figure 8. Topology Examples

There are several organized efforts to employ the 802.15.4 radio in larger, organized networks. These groups include the ZigBee Alliance, the IETF, and the IEEE itself. In addition, the functionality and cost-effectiveness of the silicon radios provides those with pre-existing proprietary networking techniques to layer that functionality on top of the IEEE radio, allowing them to reduce cost without redeveloping a network function.

The ZigBee Alliance released their specification to the public in June 2005, and since then the playing field has become much simpler for product designers who want to add wireless to their sensor or control application. An open and growing industry group of more than 180 companies from product/system OEMs to applications developers to semiconductor companies, the Alliance has worked hard to provide a technology that takes best advantage of the robust IEEE STD 802.15.4 short-range wireless protocol, adding flexible mesh networking, strong security tools, well-defined application profiles, and a complete interoperability, compliance and certification program to ensure that end products destined for residential, commercial and industrial spaces work well and network information smoothly. Figure 9 shows the relative organization of the IEEE radio with respect to the ZigBee functionality.

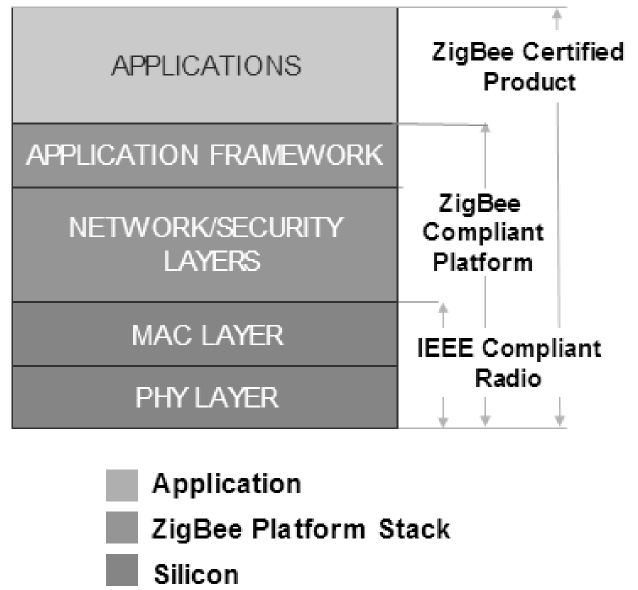

Figure 9. ZigBee Device Construction

ZigBee networking is natively mesh-based. Cost-effective, long-battery-lived radios cannot use high transmit power to ensure successful transfer of data. Instead, the network must be more clever – the most robust route between source and destination may not be the obvious, shortest physical path route, but instead as Figure 10 indicates, a route that requires other radios to "relay" the information.

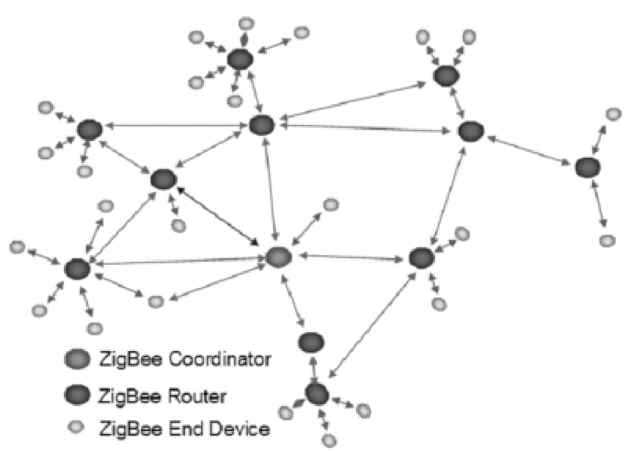

Figure 10. ZigBee Mesh Network and Device Types

The IETF began a task force in 2004 to define a method for using Internet Protocol version 6 (IPv6) over IEEE802.15.4. One of the goals of this group is to define how IPv6 will map to the dissimilar address structure and relatively small and fixed size frame of 802.15.4. Discussions are under way to determine what header compression, if any, should be used, in order that the important details available within the IPv6 header may be preserved when information crosses from the true IPv6 universe to the more constrained 802.15.4 universe.

Another challenge being faced by the IETF task force is the type of network that will be defined. While the ZigBee network could be used, its restriction at 2^{16} active nodes on any one subnet still would require some sort of Network Address Translation (NAT) technique at the interface between the two networks, taking away some of the elegance of the IPv6 addressing ability. Routing methods including some form of Ad-hoc On-demand Distance Vector (AODV) routing may provide a network with routing functionality very similar to the internet. It is also not certain right now what the ultimate size of the network layer that fulfills the needs of this group, but with the "things to things" universe where all devices in the world have an IPv6 address, this effort is very intriguing and bears close watch over the next year or two.

A cousin to the original 802.15.4 task group called 802.15.5 is working now on a mesh networking technique that may be applicable to 802.15.4 radios. This group has been focused to date on larger, more complex systems that are potentially based upon the multimedia PAN 802.15.3 specification, but there are those within the task group that recognize that even small, inexpensive devices can benefit from mesh networking. This group will be one to watch over the next few years as they move toward a specification that may embrace both small, simple devices as well as more sophisticated, data-intensive devices like media servers and client devices.

4. Conclusions

The world of "things to things" communication, of rich sensor networks and simple control systems, was strongly in mind when the IEEE 802.15.4 task group created its specification. Whether ping-pong-ball-sized sensors sprinkled across the ground in a forest or on Mars, or sensor/control networks that allow microcontrol of agricultural environments or commercial buildings, 802.15.4 provides the basic tools that ensure standards-based, reliable, robust communications under most conditions. The existence of the standard has created an already growing silicon solution market, with more vendors joining the fray every month and helping to drive costs and power consumption down, while adding functionality and performance. The basic features of the PHY and MAC layers provide the hooks to upper layer network and applications developers to take advantage of these cost-effective, small radio solutions, allowing those developers to concentrate on their application and the inexpensive delivery of data from a large, amorphous sensor network. The networking techniques already available, like that of the ZigBee Alliance, and the promise of some sort of marriage between IPv6 and 802.15.4, provide a vision of a future world where the simplest machines can interact with their world.

References

The following references are available publicly.

[1] Various authors, *ZigBee Specification*, ZigBee Alliance, 14 December 2004.

[2] McInnis, M. editor-in-chief, *802.15.4 – IEEE Standard for Information Technology*, Institute of Electrical and Electronic Engineers, New York, 1 October 2003.

Biography

Jon T. Adams *is an expert in wireless systems and communications, with his current focus on ZigBee and Ultra Wideband technologies, representing Freescale on the board of directors for the ZigBee Alliance, as president of the UWB Forum, and as a voting member in the IEEE 802.11 (WLAN), 802.15 (WPAN) and 802.16 (WMAN) Working Groups. Additional interests include RFID, Wi-Fi, and Bluetooth, as well as the impact of these on cellular telephony and the consumer experience. Jon presents and is often interviewed or published on these technologies and their impacts on or potentials for the home, commercial and industrial environments. Prior to Freescale, Jon spent 17 years at the NASA/Jet Propulsion Laboratory in Southern California where he was the cognizant engineer or manager for multiple Earth-orbiting or Mars-bound telecommunications or radar payloads. He has BSE and MSEE degrees from UCLA and holds the FCC-issued Amateur Radio license N7UV.*