Microwave Journal

CoverFeature

MESH NETWORK PROTOCOLS FOR THE INDUSTRIAL INTERNET OF THINGS

Ross Yu, Linear Technology Corp. (Dust Networks Group), Milpitas, Calif.

One of the biggest promises of the Industrial Internet of Things is to leverage real-world data gathered through wireless sensor networks (WSN) to drive higher efficiencies and streamline business practices. The demands on WSNs are diverse, with sensors placed throughout buildings, city streets, industrial plants, tunnels and bridges, moving vehicles or in remote locations such as along pipelines and weather stations. A common requirement across such applications for the Industrial Internet of Things is for WSNs to deliver both low power and wire-like reliability and to do so across a broad spectrum of network shapes, sizes and data rates.

Wireless mesh networks have become increasingly well accepted due to their ability to cover large areas using relatively low power radios that relay messages from node to node and to maintain high reliability by using alternate pathways and channels to overcome interference. One technique in particular, Time Synchronized Channel Hopping (TSCH) mesh networking, pioneered by Linear Technology's Dust Networks and incorporated into the WirelessHART industrial standard, is field proven to deliver the performance needed by the Industrial Internet of Things. TSCH networks typically experiencing >99.999 percent data reliability and all wireless nodes, even rout-

SIDEBAR

The Effects of Multipath Fading on Wireless Communications

Multipath fading depends on the position and nature of every object in the environment and is unpredictable in any practical setup. One good property is that the topography depicted in ***Sidebar 1*** changes with the frequency. If a packet is not received because of multipath fading, retransmitting on a different frequency has a high probability of succeeding.

Because objects in the environment are not static, the effect of multipath changes over time. For example, cars drive by and doors are opened and closed, ***Sidebar 2*** shows the packet delivery ratio on a single wireless path between two industrial sensors over the course of 26 days and for each of the 16 channels used by the system. There are weekly cycles where workdays and weekends are clearly visible. At any given time, some channels are good (high delivery), others poor, and still others highly varying. Channel 17, while generally good, has at least one period of zero delivery. Each path in the network shows qualitatively similar behavior but unique channel performance. There is never any one channel that is viable everywhere in the network. The key to building a reliable wireless system is to exploit channel and path diversity to mitigate interference and multipath fading.

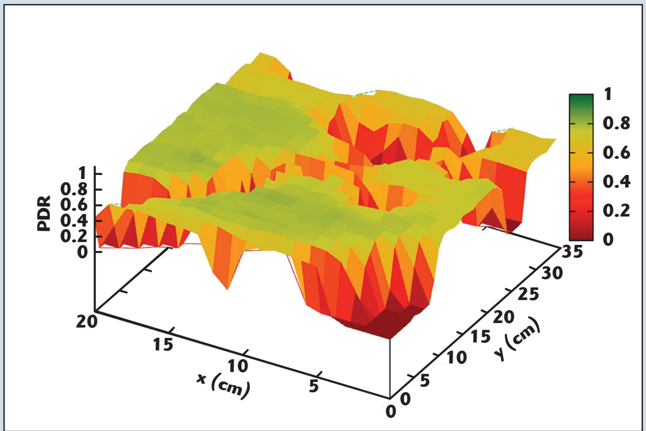

▲ *Sidebar 1 Multipath fading causes the quality of a link to vary dramatically, even when moving the receiver by only a couple of centimeters.*

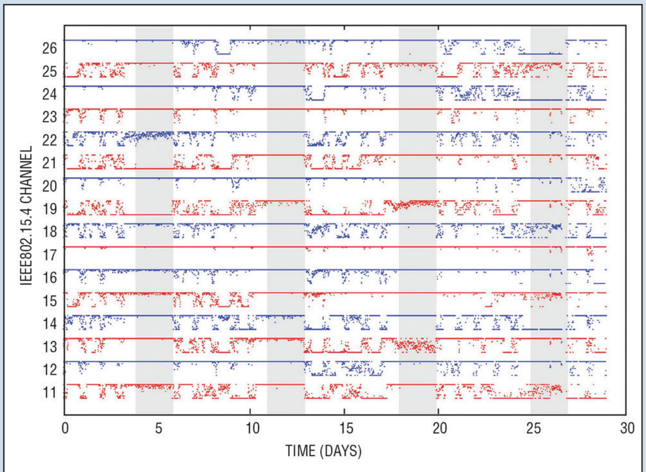

▲ *Sidebar 2 The packet delivery ratio of a wireless link varies over time.*

ing ones, have multi-year battery life on small lithium batteries. However, a variety of mesh networks use similar sounding networking techniques (e.g., "frequency agility" vs. "channel hopping," "sleepy" vs. "time synchronized"), yet yield drastically different performance levels. These wireless networking details determine how such protocol level choices greatly impact a WSN's performance and the network's overall suitability for an application.

WIRELESS SENSOR NETWORK CHALLENGES

Since wireless networks can be unreliable, it is important to understand the sources of unreliability and account for them in a communications system. Unlike wired communications, where the signal is shielded from the outside world by cabling, RF propagates in the open air and interacts with the surrounding environment. There is the possibility other RF transmission sources will cause active interference.

However, much more common is the effect of multipath fading, where the RF message may be attenuated by its own signal reflected off surrounding surfaces and arriving out of phase (see ***Figure 3***). Mobile phone users experience multipath fading every day when their phones seemingly have poor signal strength in one spot, but it improves by moving just a few centimeters. The effects of multipath change over time, as nearby reflective surfaces (e.g., people, cars, doors) typically move. The net result is that any one RF channel will experience significant variation in signal quality over time.

Further adding to the challenge is that multipath fading is unpredictable. By definition, a network must be actively transmitting on a channel to experience (and therefore measure) the channel's performance in the face of multipath fading. Therefore, while the notion of using a simple passive signal strength measurement (RSSI) of an unused channel may be helpful to detect active interferers, it cannot predict that channel's suitability in the face of multipath fading. Fortunately, since multipath fading affects each RF channel differently and changes over time, using channel hopping for frequency diversity minimizes the negative effects of multipath fading. The challenge for WSN protocols is the ability to use channel hopping over large networks with multiple hops.

COMMON APPROACHES

To understand how different WSNs perform in the face of these constraints, let us examine techniques often used in some wireless mesh networks to address frequency diversity and to deliver low power.

Single Channel WSNs and Channel Agility – A common approach in simple implementations of wireless mesh networks is to have all nodes operate on a single channel. Since only one RF channel is used, only one device can be transmitting at a time. Network stack developers often still choose single channel operation due to the relative simplicity of implementation, and in doing so provide a WSN with virtually no frequency diversity.

In order to respond to the presence of active RF interference in channel, some single channel WSNs have a mechanism called channel agility, where a network can broadcast a message to all nodes to change the operating channel. But even in channel agile networks, at any point in time the network is still operating on a single channel. The use of channel agility assumes that there is a single channel that is good for the entire network. However, with the effects of multipath fading, real world data shows that any RF channel will experience severe path degradation during the lifetime of the network, causing nodes to drop out for periods spanning minutes or hours. (See ***sidebar:*** "The Effects of Multipath Fading on Wireless Communications.") While a network with channel agility can change the channel away from an active interferer, the network is still susceptible to the devastating effects of multipath fading.

Duty Cycling by Network-Wide Sleeping - For low power operation, wireless sensor networks perform some form of duty cycling to minimize the percentage of time spent in active operation (e.g., transmit or receive, which typically draws milliamps of power) and maximize the percentage of time spent in a low power sleep mode (typically 1 mA or lower). Some wireless sensor networks incorporate a network-wide sleeping scheme (sometimes called a "sleepy" mesh), in which all the nodes in the network are put into a low power sleep state for an extended period and wake up at approximately the same time to send/receive/forward network traffic. In such sleep schemes, the network is completely unavailable for communications during the inactive period. For example, if a WSN only wakes once an hour for communications, then the network is unable to send an alarm message during that hour, nor can it receive a message from a controller to light up an attached warning indicator. It is also important to consider how the use of network-wide sleeping affects the WSN's ability to cope with real-world operating conditions. During the extended sleep periods, the surrounding RF environment remains dynamic and changing. Any signal pathway that became unusable during network sleep can only be repaired when the network awakens. Even more troublesome is the fact that sleepy networks tend to be single channel networks, placing further stress on the network during its active period and adding to the risk of communication instability.

Another repercussion of using network-wide sleep is that a network-wide sleeping approach forces a user to settle for a slower data rate (and therefore less data) than could be called for by the application. This is an unfortunate trade-off, since the main purpose of a WSN is data reliability and to use that information to enable deeper insight into the user's systems, showing operational trends and inefficiencies such as degrading performance in aging motors, or increased cyclic power draw of old refrigeration equipment in a retail store. When the data delivered by the WSN is sparse due to network limitations, the utility and insight derived from the WSN

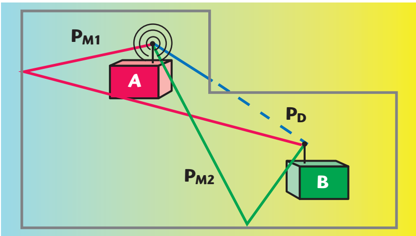

Fig. 3 A radio signal's strength at the receiver (B) is affected not only by the direct path (P_D), but also by reflections (P_{M1} & P_{M2}), which may arrive out of phase and cause significant fading.

becomes limited and runs the risk of reducing the overall value proposition of the monitoring/control system.

TIME SYNCHRONIZED CHANNEL HOPPING MESH NETWORKS

Time Synchronized Channel Hopping (TSCH) mesh networks use tight time synchronization across a multi-hop network to closely coordinate communications and frequency channel usage. In a TSCH network, each node shares a common sense of time that is accurate across the network to within a few tens of microseconds. The nodes exchange timing offset information with neighboring nodes to maintain time synchronization. Network communication is organized into time slots, in which individual packet transmit/receive opportunities are scheduled. Each time slot is long enough (e.g., 7.5 ms) for a transmitting node to wake up, transmit a packet, and receive its link-layer acknowledgment from the receiving node. Data traffic in a TSCH network can be dynamically scheduled, which enables pair-wise channel hopping, full path and frequency diversity, low power packet exchange and high-availability duty cycling.

Pair-Wise Channel Hopping - Time synchronization enables channel hopping on every transmitter-receiver pair for frequency diversity. In a TSCH network, every packet exchange channel hops to avoid RF interference and fading. In addition, multiple transmissions between different device pairs can occur simultaneously on different channels, increasing network bandwidth. For example, there are fifteen usable channels available in the IEEE 802.15.4 2.4 GHz radio specification, which is a popular choice for WSN implementations due to the global availability of this ISM band. This represents up to 15 times the available bandwidth for a TSCH network, compared to that of a single-channel 802.15.4 WSN.

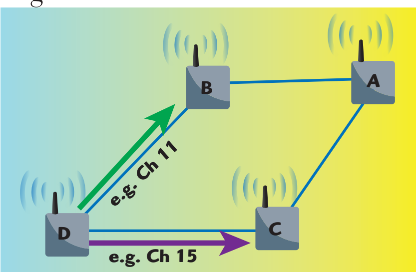

▲ *Fig. 4 If communication fails on the "green" arrow, node D retries on the "purple" arrow using another channel and pathway.*

Full Path and Frequency Diversity - Each device has redundant paths to overcome communications interruptions due to interference, physical obstruction or multipath fading. If a packet transmission fails on one path, a mote will automatically retry on the next available path and a different RF channel (see ***Figure 4***). By exercising path diversity and frequency diversity on each retry (time diversity), the probability of success on each retry is higher compared to a single-channel system.

Low-Power Packet Exchange - The use of TSCH allows nodes to sleep at ultralow power between scheduled communications. Each device is only active if it is sending a packet or listening for a potential packet from a neighbor device. More importantly, since each node knows when it is scheduled to wake up, each node is always available to relay information from its neighbors. Therefore, TSCH networks often reach duty cycles of <1% while keeping the network completely available. Furthermore, since each packet transaction is scheduled, there are no in-network packet collisions in a TSCH network. Networks may be dense and scale without creating debilitating RF self-interference.

High Availability Duty Cycling – Unlike in a network-wide duty cycled network, in a TSCH network individual nodes wake up their transceivers only when they need to transmit a packet or listen for a packet to be received. By scheduling network traffic to the granularity of individual transmitter-receiver exchanges, a TSCH network can easily accommodate heterogeneous data traffic across the network. For example, if there is a tank level sensor that only needs to transmit once an hour, and elsewhere, pressure/flow sensors that report every few seconds, then a TSCH network will wake up nodes (and their parents) only as frequently as needed to reliably support its level of data traffic.

THE POWER OF COMBINING TSCH WITH LOW POWER HARDWARE

The operating currents for 802.15.4 transceivers for general operations, such as transmit, receive and sleep, have steadily reduced over the past decade. For example, the LTC5800-IPM from Linear Technology draws 9.5 mA for a +8 dBm transmit power and 4.5 mA for receive, which is three to five times lower than prior generation 802.15.4 transceivers. Reducing peak current is a good start, but the energy required to send a packet is a function of the amount of charge drawn over a period of time. If current is measured on an oscilloscope and plotted over time (see ***Figure 5***), then the energy required to send a packet is shown as the area under the curve and affected not only by peak currents, but also by how long each operation is active. Products such as this deliver precisely optimized packet exchanges with a successful packet transmission/acknowledgment for a mere 54.5 µC charge at 3.6 V supply voltage (or 196.2 µJ of energy).

A SYSTEM APPROACH TO LOW POWER

By taking a more holistic view of how energy is spent in a wireless sensor network, low power consumption can be thought of as a function of data traffic as well as the energy required to send a packet and the number of retries needed to successfully send a

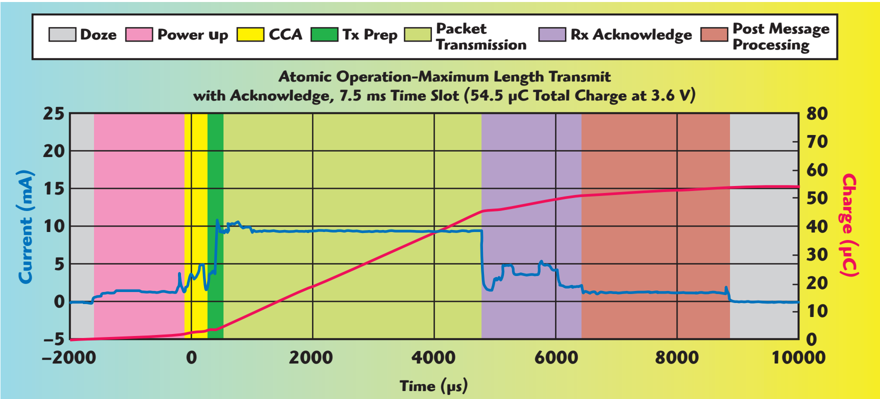

▲ *Fig. 5 The current during packet transmission and receipt of the link-layer acknowledgement. With TSCH-optimized hardware, individual transactions can reach as low as 54.5 µC.*

packet from one node to the other:

$$\text{Average Energy} = f\left(\left(\frac{\text{Num Packets}}{\text{Period of Time}}\right)\times\left(\begin{matrix}\text{Energy per}\\\text{Packet}\end{matrix}\right)\times\left(\begin{matrix}\text{Num Retries to}\\\text{Successfully Send}\\\text{A Packet}\end{matrix}\right)\right)$$

By focusing on energy per packet and using a networking protocol that exercises time, path, and frequency diversity on every retry (thereby reducing the average number of retries required to send a packet), low current consumption can be attained by improving efficiency throughout the system rather than making sacrifices on the application layer. The communication schedules in a TSCH network are highly configurable, with communications timeslots automatically allocated based on application needs. A TSCH network can be configured for slow data rates to minimize power and potentially enable the use of energy harvesting. That same TSCH network can be configured to support heterogeneous report rates, as is commonly done in industrial plants that have slow-changing variables (e.g., tank level) and faster changing variables (e.g., flow in a pipe). A TSCH network will automatically allocate the required timeslots to the portions of a network that need it. Instead of forcing users to tailor their applications to meet the needs of the network, a TSCH network can be tailored to meet the needs of a wide variety of applications.

ENABLING THE INDUSTRIAL INTERNET OF THINGS

TSCH is already a foundational building block of existing industrial wireless standards, such as WirelessHART (IEC62591), and is an enabling piece of emerging Internet Protocol-based WSN standards, including IEEE802.15.4e.[1] Work is underway to standardize a TSCH link layer within the IETF 6TiSCH group as well.[2] The adoption of TSCH into relevant standards will continue to encourage far-reaching adoption. TSCH networks have already proven to deliver multiyear battery life and >99.999 percent data reliability in such diverse and demanding applications as industrial process monitoring,[3] fence line perimeter security,[4] data center energy efficiency,[5] and metropolitan-scale smart-parking solutions.[6] By delivering highly reliable, low-power wireless networks that are highly configurable, TSCH networks are ideally suited for the Industrial Internet of Things. ■

References

1. Reliable, Low Power WSNs for the Internet of Things, www.linear.com/docs/44202.
2. https://datatracjer.ietf.org/wg/6tisch/charter/.
3. Emerson Process, www.linear.com/docs/41383.
4. Integrated Security Corp., www.isc-hydra.com.
5. Vigilent, www.linear.com/docs/41384.
6. Streetline, www.linear.com/docs/41387.