

Wireless Sensor Network Challenges and Solutions

Lance Doherty, Systems Engineer, Dust Networks Product Group

Jonathan Simon, Systems Engineering Director, Dust Networks Product Group

Thomas Watteyne, Systems Engineer, Dust Networks Product Group



We live in a world filled with sensors. The buildings that we work in have sensors monitoring temperature, occupancy, smoke and fire, and security. Our cars contain dozens if not hundreds of sensors, monitoring engine performance, braking and passenger safety equipment, to name a few. Manufacturing environments need sensors because you cannot control what you cannot measure. Making products while meeting safety, quality and efficiency targets requires a lot of sensors.

Sensors have become much smaller, less expensive and lower power in the last few decades, driven in part by both Moore's law and the MEMS revolution. Unfortunately, the cost of installing sensors has not kept pace. The cost of running wires to carry power and data typically dwarfs the cost of the sensor itself. Take for example the closest light switch: the wiring for a \$1 switch can cost \$50, mostly labor, even in new construction. If you want to move that switch to the adjacent wall, the price of retrofit is much higher. In industrial process automation, the accepted rule of thumb is \$10,000 to install a sensor—even a simple switch. In this cost environment, many sensors only report data to a local controller—there can be little or no “big picture” when hundreds or thousands of sensors are installed. What is needed is an inexpensive, reliable way to network sensors.

Almost since the time of Marconi, people have used wireless to communicate data from sensors, with mixed results. Traditionally these links have been line-powered and point-to-point, often with time-varying reliability due to environmental conditions. This is fine for some applications, but too restrictive for most.

Markets

Markets for wireless sensor networks (WSNs) include building automation, industrial control, home automation, smart grid and au-

tomated metering infrastructure (AMI), industrial process automation, environmental monitoring, parking and transit infrastructure, energy monitoring and inventory control.

In most cases, these are bidirectional asymmetric data collection applications—large numbers of sense points forward data to a central host that may respond with a process set point or other configuration changes.

Technology Choices

Customers ideally want a technology that is low cost, allows unrestricted sensor placement, receives periodic data reliably with low latency, and runs for the device lifetime with no battery changes. Recent technological advances have enabled us to deliver those features in many markets.

There are several technologies competing to fill this role, including satellite, cellular, Wi-Fi, and a host of solutions based on IEEE 802.15.4 radios. These technologies allow users to form WSNs for collecting sensor data.

Satellite and cellular work well for many applications, but have the highest energy cost per packet. Data plan charges can also be prohibitive, although this is likely to change as carriers develop billing models appropriate for relatively sparse data flows. Coverage can also be an issue. Clearly it can be difficult for a satellite or cell phone signal to make its way out of a heavily obstructed struc-

ture, and the sensors generally do not have the capability of moving from side to side and asking, “Can you hear me now?” For an application sending at a very low data rate (e.g., one data packet per day) with good connectivity, however, satellite or cellular can make a lot of sense.

Wi-Fi (IEEE 802.11b, g) sensors are now widely available. The energy cost for a Wi-Fi packet is much lower than cellular, and there are no recurring fees for data. Connectivity and coverage remain important concerns, as the density of access points necessary for reliable communication with a fixed sensor is typically higher than that necessary for mobile humans with gadgets.

Because of interference and multipath fading, the key to building a reliable wireless system is to exploit channel and path diversity.

With reference to the OSI layer model, the 802.15.4 standard defines a physical layer (PHY) and medium access control (MAC) layer for short range, low power operation that is well suited for wireless sensor networks. The radio is relatively low data rate (up to 250kbps); the packets are short (< 128 bytes) and low energy. For example, sending a few bytes of sensor data, with routing, cryptography, and other headers takes less than 1ms, and burns less than 30μJ of energy (see Figure 1), including receiving a secure link-layer acknowledgement. Sensors can forward radio packets from peers, extending the range of the network far beyond the range of a single radio, and providing the network with immunity to any single radio link failure.

wp001fa

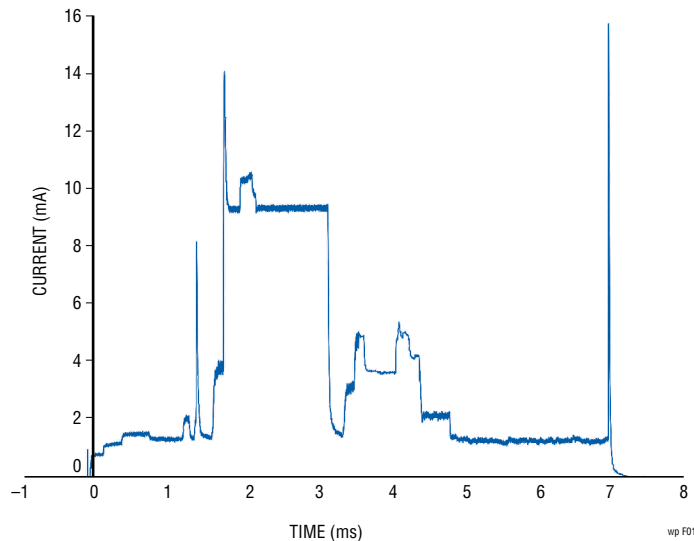


Figure 1. Energy to Transmit a Short 802.15.4 Packet and Receive an Acknowledgement

Performance Metrics

Evaluation of various WSN solutions is based on two questions, “Can I get all my data fast enough?” and “How much is it going to cost?” WSNs must be designed to work in environments with link-layer packet delivery ratios (PDR) down to about 50%.

When developing a wireless data collection system, there are a few performance targets that must be met. First, the system must meet a minimum reliability goal. For industrial applications, the target is typically to receive at least 99.9% of the generated data, as missing data can trigger expensive alarm conditions. Second, the system must support a certain throughput, a number of sensor data packets per second. Third, these data packets are only useful if received within a maximum latency period. Many processes rely on fresh data updates—for control, stale data may have no utility. Fourth, many systems must operate in challenging environments that include wide temperature ranges and intrinsic safety restrictions. Only solutions that meet all four of these requirements are considered suitable for further evaluation.

When considering various solutions that meet the requirements, the key selection criteria become cost of ownership and flexibility. The cost of ownership encompasses several areas: product development, installation, hardware and providing power over the lifetime of the installation. Wireless technologies have reduced installation costs dramatically compared to wired solutions, but battery-powered wireless devices may require bat-

tery changes over the lifetime of the network. There is also a trade-off between building a network with a small number of high power devices to reduce the hardware cost versus using a larger number of low power devices. For devices powered by energy harvesting cells (e.g., solar, thermoelectric), capacitor size may determine a significant portion of the cost. Solutions with deterministic scheduling, such as time-division multiple access (TDMA), can help separate high current events as much as possible to reduce the capacitor size requirement.

Because final deployment conditions are unpredictable, networks must be designed for flexibility. Networks must scale from small

to large numbers of sensors and from low to high density. To be robust across diverse wireless environments, resource provisioning should ensure that devices reliably communicate with moderate interference and that the networks survive the loss of individual devices. Additional resources, including more wireless links, more neighbors for each device, or more signal amplification, improve reliability and latency. All these additions come at increased power costs, which can be minimized with dynamic allocation.

Solutions based on standards provide immunity to the supply chain vagaries of a single vendor component and the assurance that the community has agreed on the governing principles of operation, e.g., security architecture.

Challenges

The wireless channel is unreliable in nature, and a number of phenomena can prevent a transmitted packet from reaching a receiver. One such phenomenon is interference. If two independent transmitters transmit on the same channel such that their signals overlap, they may corrupt each other’s signal at a receiver’s radio. This requires the transmitter to retransmit, at the cost of additional time and energy.

Interference can come from the same network if the underlying medium access technology does not schedule contention-free communications. This is particularly problematic if the two transmitters can hear the receiver, but not hear each other—this is known as the

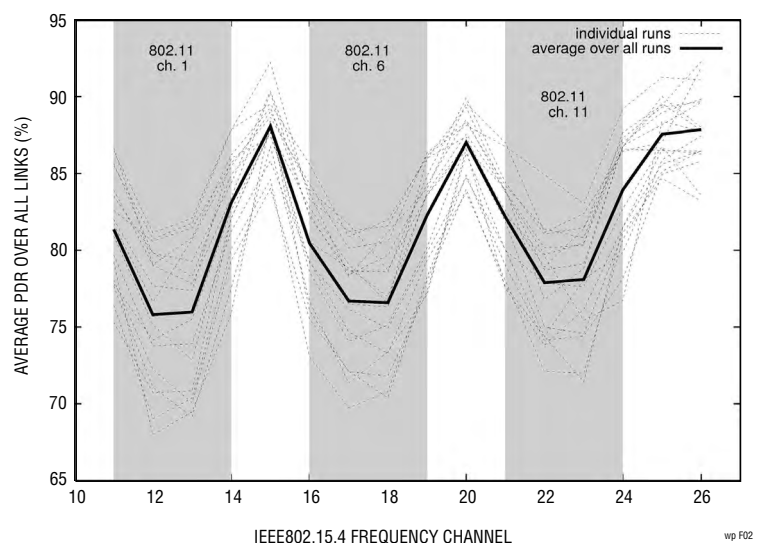


Figure 2. Interference Between Wi-Fi and 802.15.4 in the 2.400GHz to 2.485GHz Frequency Band

“hidden terminal problem,” and it requires backoff and acknowledgement mechanisms to resolve collisions.

Interference can also come from another network operating in the same radio space, or from a different radio technology using the same frequency band. The latter, known as “external” interference, is especially present in unlicensed bands such as the 2.400GHz to 2.485GHz instrumentation, scientific and medical (ISM) band, crowded with Wi-Fi, Bluetooth and 802.15.4.

Figure 2 was obtained by deploying forty-five 802.15.4 nodes in an office environment, and having them exchange 12 million packets, equally distributed over sixteen 802.15.4 channels. It plots the average packet delivery ratio of those packets as a function of the channel they are transmitted on; on channels overlapping Wi-Fi channels, this delivery ratio is lower.

A second phenomenon, multipath fading, shown in Figure 3, can prevent a transmitted packet from reaching a receiver and is both more destructive and harder to quantify. Often described as “self-interference,” this occurs when the recipient receives both the signal traveling over the line-of-sight path from the transmitter, as well as “echoes” of the same signal that have bounced off objects in the environment (floors, ceilings, doors, people, etc.). Since those copies travel different distances, they reach the receiver at different times, potentially interfering destructively. Fades of 20dB to 30dB are not uncommon.

Figure 3 was obtained by having a transmitter transmit 1000 packets to a receiver 5m away, and repeating this with the receiver positioned at each point in a 35cm by 20cm grid. The z-axis represents the packet delivery ratio over that link. While the link is good at most positions, at some positions no packets are received successfully because of multipath fading.

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 Figure 3 changes with the frequency. That is, if a packet is not received because of the multipath fading, retransmitting on a different frequency has a high probability of succeeding.

Because objects in the environment are not static, e.g., cars drive by and doors are opened

and closed, the effect of multipath changes over time. Figure 4 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 sixteen 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 bad, 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 with different channel performance, and there is never any one channel that is good everywhere in the network.¹

Because of interference and multipath fading, the key to building a reliable wireless system is to exploit channel and path diversity.

¹ L. Doherty, W. Lindsay, J. Simon, K. Pister, “Channel-Specific Wireless Sensor Network Path Analysis,” Proc. ICCCN '07, Honolulu, HI, 2007

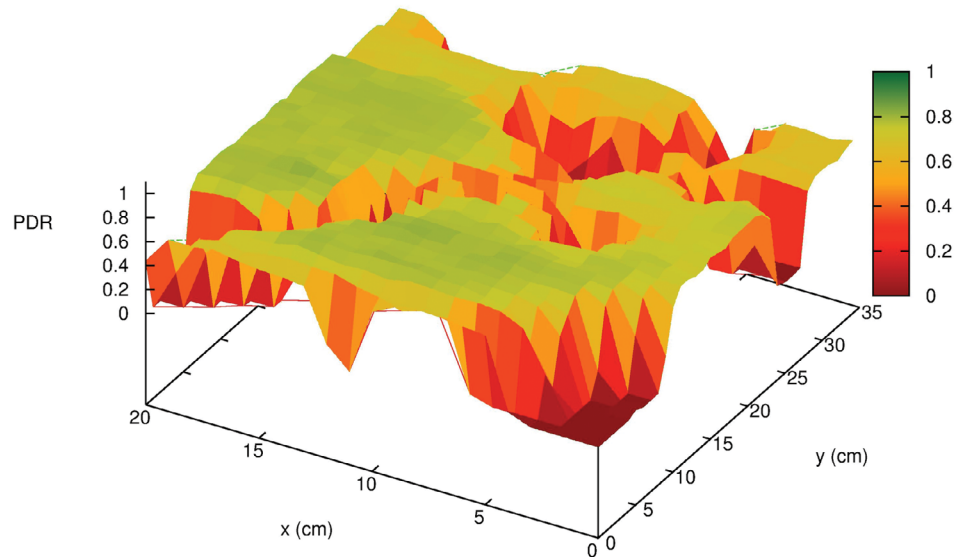


Figure 3. Multipath Fading Causes the Quality of a Link to Vary Dramatically, Even When Moving the Receiver by Only a Couple of Centimeters.

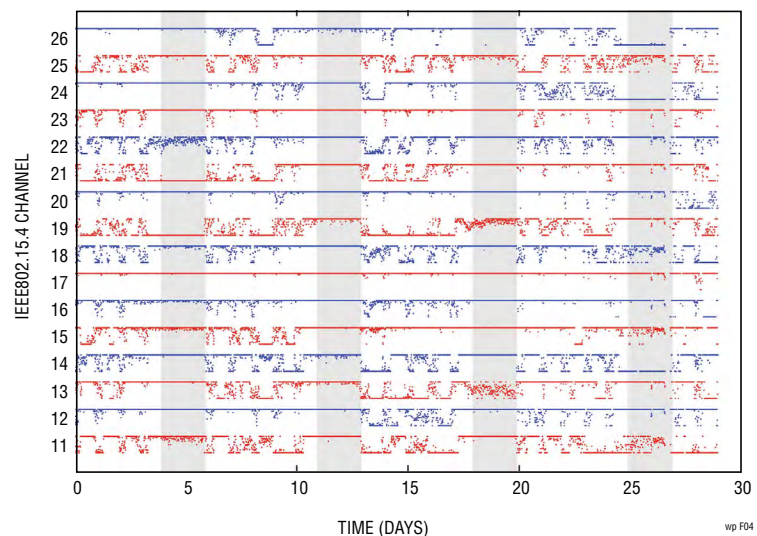


Figure 4. The Packet Delivery Ratio of a Wireless Link Varies Over Time

Solutions

As stated previously, one technology well suited for solving the WSN problem is IEEE 802.15.4—802.15.4 radios, which offer low power, low data rate PHYs in several unlicensed frequency bands, including the 915MHz band, available in North America, and the 2.4GHz ISM band, which is available worldwide. The 2.4GHz band spread spectrum PHYs provide immunity to noise—a particularly important feature for a low energy device designed to operate in a potentially crowded, unlicensed band. The standard also defines a reliable, acknowledged, packet (or frame) based MAC layer with optional encryption and authentication. This flexible solution forms the basis of several proprietary and standards-based protocols including the ZigBee protocol, which uses it to form unsynchronized single-channel networks, and the WirelessHART protocol,² which uses it to form time-synchronized multichannel networks.

The WirelessHART protocol, which Dust Networks® helped develop, has an 802.15.4 2.4GHz PHY and an 802.15.4 based link layer, which adds synchronization, channel

hopping, priority and time-based authentication to the standard 802.15.4 MAC. It has a network layer that provides routing and end-to-end security, and a thin unreliable/reliable mesh transport layer. The WirelessHART standard specifies time slot timing, how devices maintain synchronization, and how devices schedule time/channel communications opportunities by dividing time into slotted communications opportunities (time slots) on repeating superframes. The protocol was designed to allow seamless integration of wireless devices to existing wired HART installations, widely used in industrial process monitoring and control applications. WirelessHART extends the HART application layer command set, adding commands for managing wireless resources and monitoring network health. WirelessHART networks are highly reliable meshes—even with devices that do not have line-of-sight and at tens to hundreds of meters spacing, each device has multiple neighbors to which it can send data—providing the path diversity needed for reliability. WirelessHART networks are centrally managed, with most network intelligence residing in a manager. Field devices

(wireless sensors) report status information that the Manager uses to groom and optimize the network, and sensor data is reported to an application proxy called a gateway.

Recently in 2012, a new 802.15.4e amendment was released, which, among other things, formalized time-slotted channel hopping features like those found in WirelessHART at the 802.15.4 MAC layer. The standard defines the mechanisms for advertising synchronization information to allow devices to synchronize to a network, provides for time-based security, and defines slotted communications and hop sequences. It makes extensive use of data encapsulation in “information elements”—this allows for custom extensions of the MAC without having to wait for the standard to be updated. It is intended to ease development of a multilayer protocol, and was specifically designed to couple to a 6LoWPAN-compressed IPv6 network layer as defined in IETF RFCs 4944 and 6282.³

³ <https://datatracker.ietf.org/doc/>

² http://www.hartcomm.org/hct/documents/documents_spec_list.html

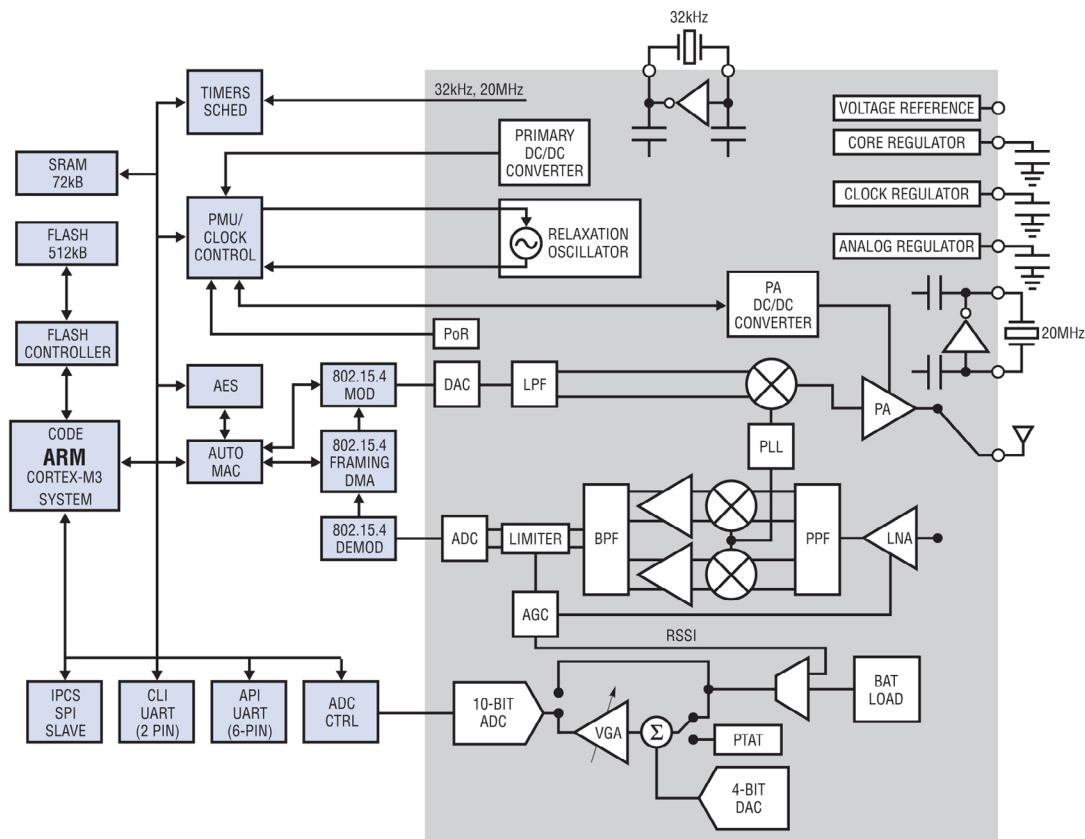


Figure 5. Block Diagram of the LTC5800 Dust Eterna Motes

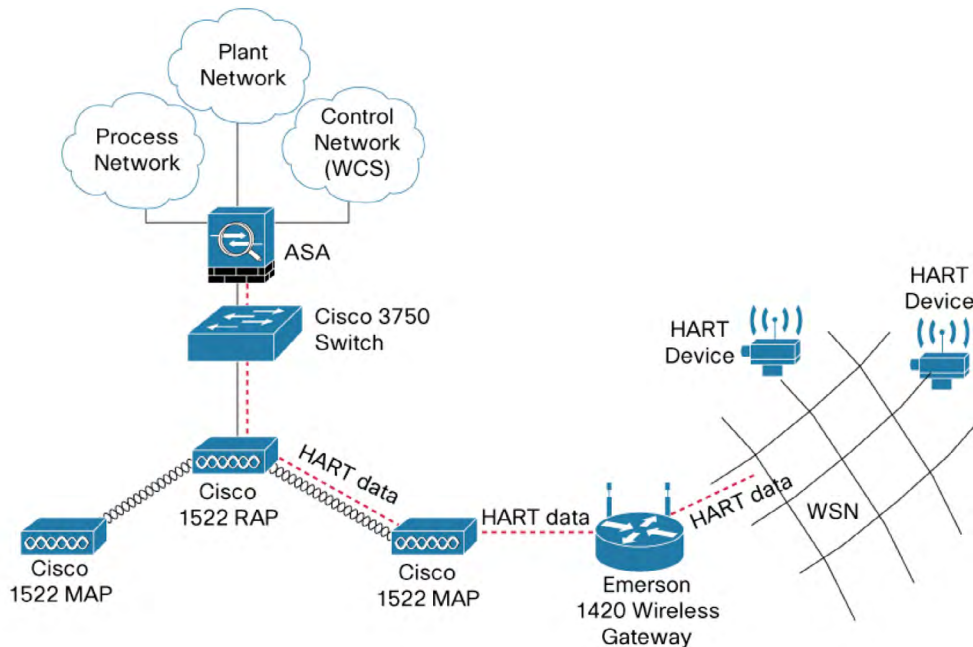


Figure 6. Network Architecture for Refinery Process Control

Applications

Linear's Dust Networks SmartMesh™ product line contains both WirelessHART and 6LoWPAN-compliant IPv6 product offerings that leverage 802.15.4 to provide the most reliable, lowest power WSN solutions on the market. Dust Networks Eterna™ motes (LTC®5800 family) are single-chip devices that couple a Cortex-M3 microprocessor, memory, and peripherals to the lowest power 802.15.4 radio available today (Figure 5). Designers embed a mote in their sensor package, and can rely on the network to form, optimize and carry their sensor data to their application. Dust Networks managers allow for graceful scaling from tens to thousands of devices, providing data and configuration interfaces for the network. Both product families build highly reliable, multihop mesh networks capable of per-node configurable data rates. They are suitable for solving a wide array of WSN problems. Some examples of applications using Dust™ motes and managers include:

Parking: Streetline⁴ is a smart parking provider that monitors the real-time availability of urban parking spaces. Vehicle detectors are installed underneath parking spaces, inside the pavement and flush with the roadway. This brings challenges, as the antenna for the sensor device is located underground,

and then covered by a metal vehicle when the space is occupied. Wireless path diversity is essential as different vehicle positions change the path quality between device pairs. Streetline installs elevated repeater devices on nearby street lamps to obtain line-of-sight to the stall sensors. These repeaters form a multihop mesh to collect all the occupancy data to the local network manager, where it gets aggregated into a citywide database available to customers and enforcement agencies. Wireless technology is critical for this application because it is intractable to wire sensors to each space, and low power wireless decreases the frequency of battery changes.

Refinery Process Control: Chevron uses wireless networks to monitor oil extraction and refining facilities. These networks are often deployed in harsh environments (due to hazardous temperatures, chemicals or risk of explosion) where it is impossible to run conduit for wired sensors. Additionally, wireless enables monitoring of rotating structures and mobile operators. For one deployment (Figure 6), wireless networks were installed in various locations around a large refining facility. To gather the data to a centralized control center, a Cisco IEEE 802.11a wireless mesh was used as the backhaul connection for each IEEE 802.15.4 network manager. This allowed the low power sensor devices to report to their local manager where data was aggregated and reliably shuttled along. This

deployment represents a powerful fusion between the two standards.

Energy Monitoring: Vigilent⁵ provides intelligent energy management systems for indoor environments such as data centers where environmental control is critical. As increased temperature at any location in the data center can cause equipment failures, air conditioning is often run continuously at full power, wasting energy. Facilities managers are reluctant to jeopardize their internal networks, so Vigilent deploys wireless devices that do not interfere with regular operation. The facilities are also sensitive to security, so the wireless protocol is required to have end-to-end encryption of all packets and additional security at the network manager. Data center sense points are typically dense, and Vigilent has had success in deploying multiple overlapping networks to achieve the required number of sensors.

Conclusion

Multichannel time-synchronized mesh networks based on 802.15.4 radios address many of the challenges involved in building flexible, reliable, low power wireless sensor networks.

⁵ <http://www.vigilent.com/>

LT, LTC, LTM, Linear Technology, the Linear logo and Dust Networks are registered trademarks and Dust, Eterna and SmartMesh are trademarks of Linear Technology Corporation. All other trademarks are the property of their respective owners.

⁴ <http://www.streetline.com/>

