Electronic Design®
Machine Design®
LIBRARY

A compendium of articles from
*Electronic Design* and *Machine Design*

*Sponsored by*
TADIRAN BATTERIES

Copyright © 2024 by Endeavor Business Media. All rights reserved.

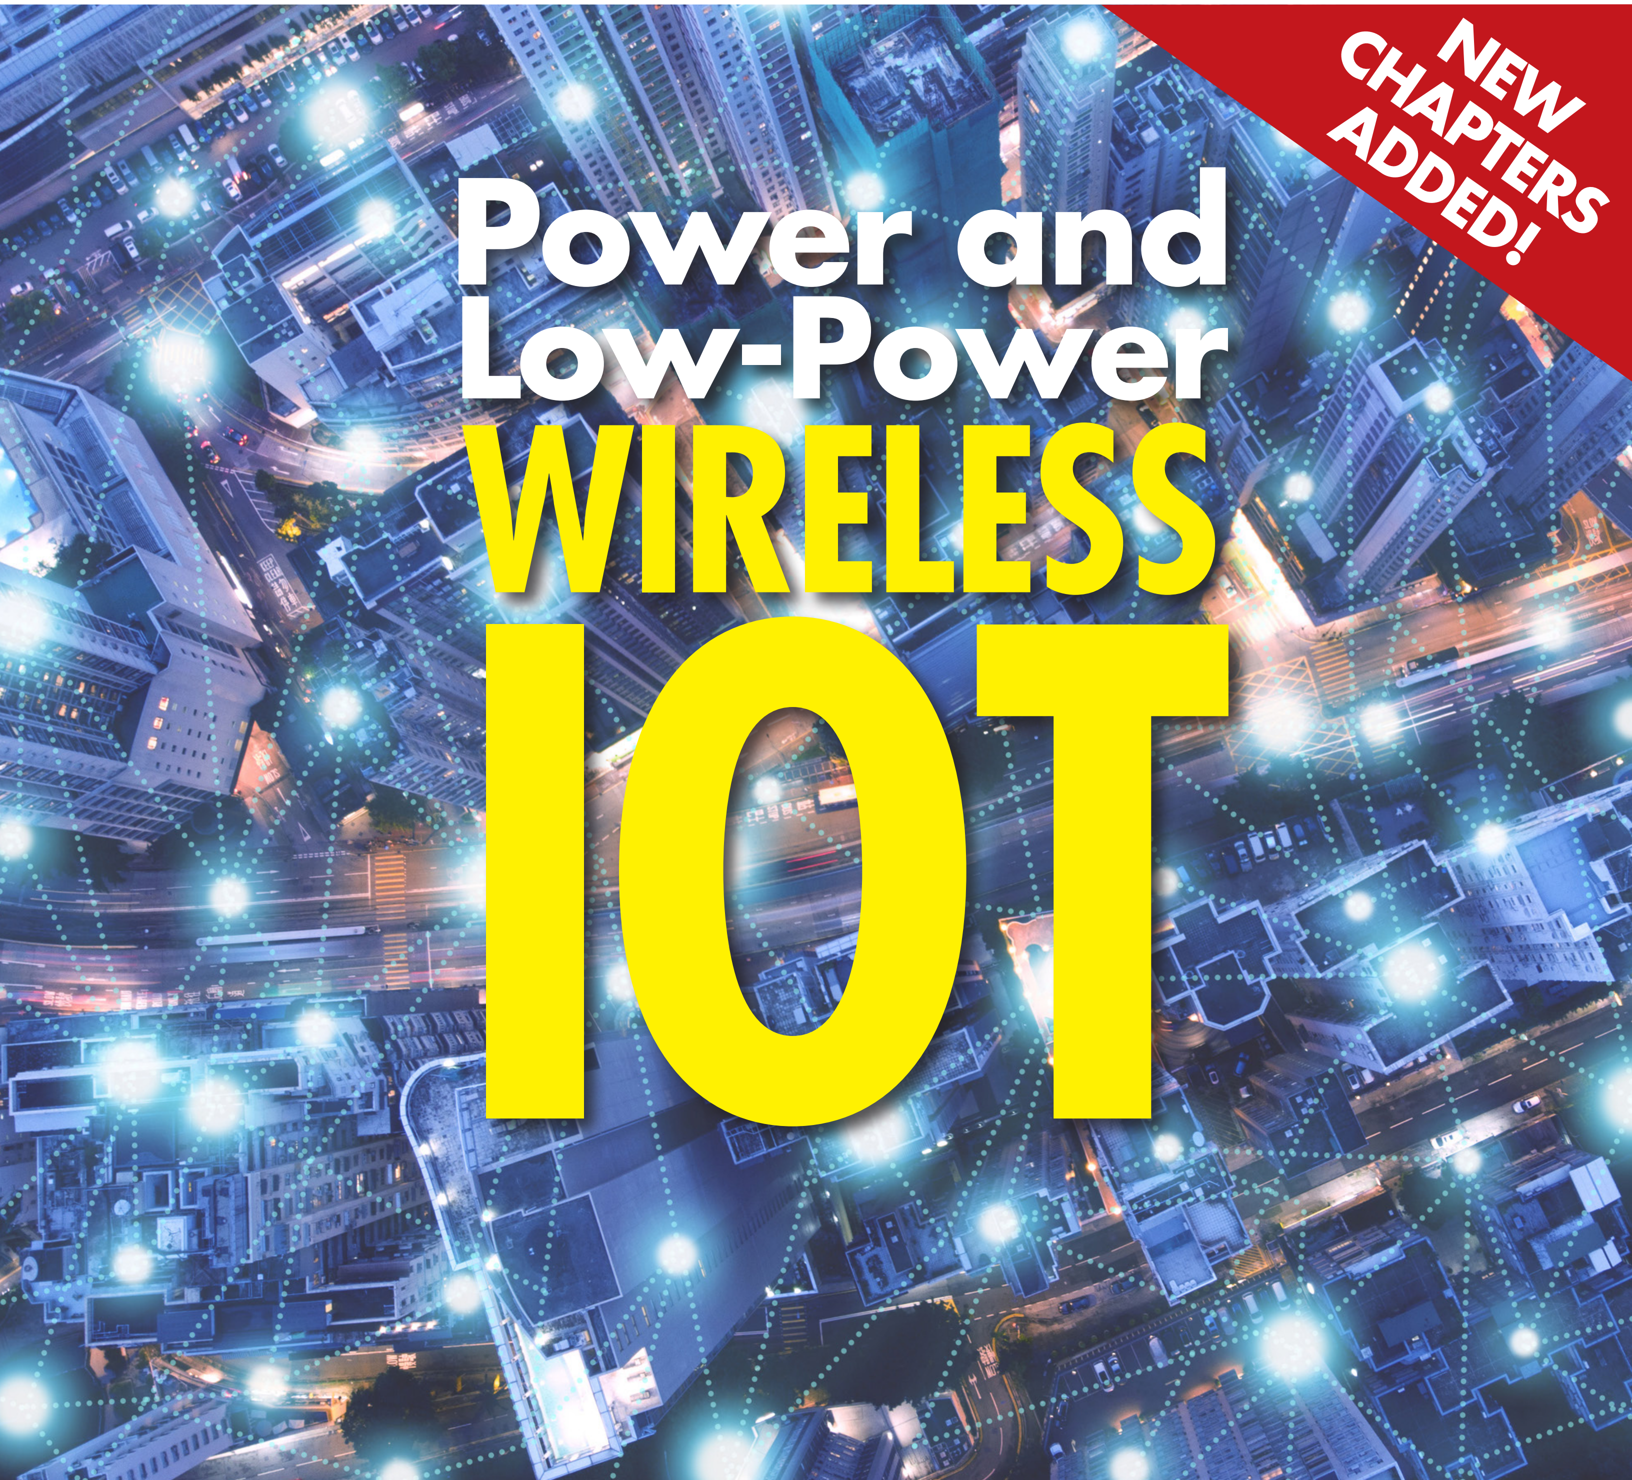

NEW CHAPTERS ADDED!

# Power and Low-Power WIRELESS IOT

# IIoT devices run longer on Tadiran batteries.


PROVEN **40** YEAR **OPERATING** LIFE*

Remote wireless devices connected to the Industrial Internet of Things (IIoT) run on Tadiran bobbin-type $LiSOCl_2$ batteries.

Our batteries offer a winning combination: a patented hybrid layer capacitor (HLC) that delivers the high pulses required for two-way wireless communications; the widest temperature range of all; and the lowest self-discharge rate (0.7% per year), enabling our cells to last up to 4 times longer than the competition.

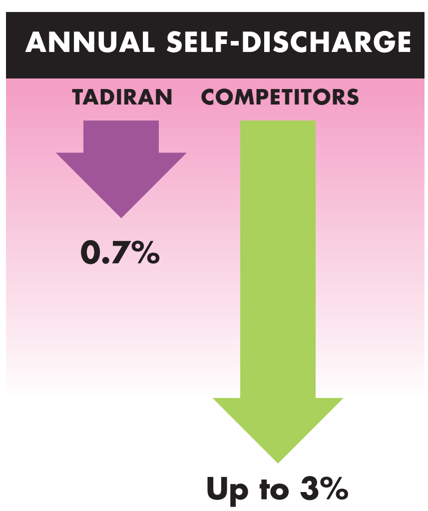


**Looking to have your remote wireless device complete a 40-year marathon? Then team up with Tadiran batteries that last a lifetime.**

TADIRAN BATTERIES

* Tadiran $LiSOCl_2$ batteries feature the lowest annual self-discharge rate of any competitive battery, less than 1% per year, enabling these batteries to operate over 40 years depending on device operating usage. However, this is not an expressed or implied warranty, as each application differs in terms of annual energy consumption and/or operating environment.

**Tadiran Batteries**
**2001 Marcus Ave.**
**Suite 125E**
**Lake Success,**
**NY 11042**
**1-800-537-1368**
**516-621-4980**

**www.tadiranbat.com**


## INTRODUCTION

With low-power wireless spreading across the IIoT, this eBook offers timely insight into highly relevant content regarding this battery life-extending technology. The selected topics range from 11 myths about low-power, using hybrid solar/li-ion battery solutions to deliver long-term power for remote wireless devices, using low-power communications protocols such as LPWAN and Bluetooth to create mesh networks that support two-way wireless communications, and highlighting certain applications that benefit from low-power design.

***Bill Wong***
*Editor,*
*Senior Content Director, Electronic Design & MWRF*

# CONTENTS



*Sponsored by* TADIRAN BATTERIES



CHAPTER 1:

# 11 Myths About Low Power

KEITH ODLAND, Senior Director of Marketing, *Ambiq Micro*

**Though we've heard about low-power design for some time, with power consumption becoming a main figure of merit, myths and misunderstandings still abound.**

We've heard about low-power design over a long period now. In fact, power consumption is acknowledged as one of the main figures of merit for any product. But myths and misunderstandings still abound.

**1. Power consumption is the only thing that matters (part 1).**
Wrong. Even engineers working on power-sensitive products like wearables, smart metering, and IoT devices need ever-increasing processor horsepower so that they can innovate. What really matters is satisfying the need for more processing within defined power and energy budgets.

There are signs that the wearables market in particular is moving into a new phase. Designers have established the parameters of ultra-low-power design and produced the first generation of wearable products. Now they're looking at processor technologies like sub-threshold switching, which allow them to make much more feature-rich products within the same power budget.

**2. Power consumption is the only thing that matters (part 2).**
Instantaneous power consumption is of course important—your product can't ever be allowed to demand more volt-amps than is possible with the energy source. But even more important is energy consumption: The amount of energy required by the product to complete its tasks is a dominant factor in defining battery life. As we know, in consumer markets, time between charges is a key "care-about." It's perhaps even more vital in applications like metering, where a single charge needs to last for years. In areas like this, a few percent on battery life makes a substantial difference to the economics of deployment.

**3. Low power is only required in portable and/or battery-powered applications.**
Wrong. First, every product is now expected to be energy-efficient. Programs such as

Energy Star and increasingly demanding energy performance regulations for buildings have made sure of that.

But the requirement for low-power operation is often driven by more subtle factors. For example, a product like a smart wall switch will commonly derive its power parasitically from the electrical mains. In that case, local regulations will set very strict limits on how much current it can sink out through the neutral wire. Similarly, if you want to add extra sensors in a multi-drop communications system using a protocol like RS-485, you need to be very careful about keeping within the current budget defined in the standard.

**4. The best way to design a low-power MCU is to choose a stripped-down core like a Cortex-M0, which doesn't have much performance.**

Wrong. For starters, your application may simply need the capabilities of the higher-featured processor. As we've already said, the focus in a market like wearables is now switching to the feature set. In any case, it's possible that a higher-spec processor will deliver better energy efficiency, because it will "get the job done" in fewer clock cycles than a less-competent alternative.

One such example is Ambiq's Apollo processor—we chose an ARM Cortex-M4 ahead of the lower-spec M0. A look at its CoreMark benchmarks will help to explain why (**see table**).

ARM CORTEX COREMARK BENCHMARKS

| | Dhrystone (official) | Dhrystone (max opt) | CoreMark* |
|---|---|---|---|
| Cortex-M0 | 0.87 | 1.27 | 2.33 |
| Cortex-M0+ | 0.95 | 1.36 | 2.46 |
| Cortex-M3 | 1.25 | 1.89 | 3.32 |
| Cortex-M4 | 1.25 | 1.95 | 3.40 |
| Cortex-M7 | 2.14 | 3.23 | 5.01 |

*CoreMark data from ARM website and EEMBC.org website

As you can see, the M4 delivers CoreMark figures that are around 46% better than the M0 running at the same clock frequency. So an M4 device will spend a proportionally smaller amount of time processing, and return to a lower-power sleep state more quickly than the M0. Of course, Ambiq's use of subthreshold voltage technology narrows the power-consumption gap between the M4 and M0 even further.

**5. Lower clock rate is better.**

Wrong. As with myth #4, getting the job done quicker means a shorter period of active power consumption; therefore, a faster clock might require less energy overall.

The important thing to remember is that there are two components to active power consumption. The most obvious is what we might call the "dynamic component"—the power consumed when the processor performs a task. The second is what we can consider as a "quiescent current"—power that's dissipated while the processor is in active mode, whether tasks are executing or not. If a task takes a million clock cycles, a 10X increase in clock speed will probably not increase the dynamic component. However, it WILL reduce total execution time (and hence the amount of "quiescent" energy dissipated) by a corresponding factor.

**6. What's the problem? These MCUs spend most of their time in sleep mode. All you need is low sleep mode consumption.**

Wrong. You can't achieve low energy operation without low active power. The fact is that active-mode power consumption is typically three or four orders of magnitude more than the corresponding figure for sleep mode. So the active-passive duty cycle doesn't have to be high for the active mode to dominate the overall performance figure.

Let's take a specific example of the Apollo microcontroller. Its active power consumption is 35 µA/MHz. Therefore, at a typical active clock speed of 24 MHz, it draws a little under 1 mA of current. In contrast, its sleep mode draws just 150 nA. So even if the application is asleep 99.9% of the time, the active-mode contribution to the overall energy budget is over five times that of sleep mode.

**7. A "good" sleep mode is all about how low you can take the power consumption.**

Wrong. You also need to think about how much time and energy you're going to spend waking the processor up. A very deep sleep mode will require a substantial number of cycles to wake up and restore status. Effectively, these are "wasted cycles." Consequently, if the application involves frequent wake-ups, it might be better to strike a compromise and choose a "shallower" sleep mode—one that dissipates more power during the sleep period, but requires less energy expenditure on each wake-up.

**8. Saving power is all about the processor.**

Wrong. Peripherals are increasingly important. You need a low-power peripheral interface, as well as peripherals with a degree of autonomous functionality that can offload tasks from the processor.

In addition, you need to design-in memory structures and architectures that are conducive to low-energy operation. As an example, consider a system in which an application talks to a Bluetooth radio via SPI. To minimize power consumption, we would, of course, want the SPI interface to work without processor intervention. But we would also want to size the incoming message buffer so that the processor only needs to wake up when it has to deal with a complete and valid message. Somewhat surprisingly, you can actually save power by getting the size of a memory buffer right!

**9. Low-power design is about continuous evolutionary improvement.**

Wrong. Technologies are still coming to market that reduce power consumption by orders of magnitude. Ambiq's Subthreshold Power Optimized Technology (SPOT) is a case in point. The company developed a microcontroller that, after testing it to the industry-standard EEMBC ULPBench benchmark, consumes less than half the energy of other comparable MCUs.

**10. If you want to build a low-power chip, you use a low-power process.**

Not necessarily. Most “low-power processes” are actually designed for low leakage. This can often mean that dynamic power is higher, and as we have seen, dynamic power typically dominates total energy budgets.

**11. When you’re working to a power budget, MIPS are always on the “cost” side of the equation.**

Another surprising result, until you start thinking about it. However, the fact is that today, MIPS can be delivered so power-efficiently that they can sometimes be used to save power elsewhere in the system.

As an example, consider a data-logging application, in which data is acquired and written to an external storage device, such as an SSD or memory card. The cost, in terms of power, of the data-acquisition process will typically be microwatts, but writing to the memory may cost milliwatts. So a sensible strategy is to spend some processor cycles on compressing the data and “save big” on the energy cost of reads and writes to and from storage. A 1:10 compression of the data might cost a few extra microwatts in processing. However, if it reduces read/write power consumption by a factor of 10, it’s likely to be more than worth it.

*to view this article online,* ☞ *click here*

☞ BACK TO TABLE OF CONTENTS


Image: Amoklv, Dreamstime

CHAPTER 2:

# Forecast is Bright for Solar/Li-Ion Battery Hybrids Powering Remote Sensors

JEFF CRYSTAL, COO, *Voltaic Systems;* SOL JACOBS, VP & General Manager, *Tadiran*

**Hybrid energy harvesting solutions that combine small photovoltaic (PV) panels with industrial grade rechargeable Lithium-ion (Li-ion) batteries enable long-term power supply solutions for remote wireless devices connected to the IIoT.**

Billions of wireless devices will soon be connected to the Industrial Internet of Things (IIoT). While the vast majority of these will be powered by primary (non-rechargeable) lithium batteries, power management challenges can occur when a wireless device has a relatively high daily power consumption. This will prematurely exhaust the primary battery, which can become highly problematic—especially for devices being deployed in hard-to-access locations and extreme environments.

Of all the energy harvesting technologies currently available, the most popular remote power supply solution involves solar PV/Li-ion battery hybrid technology. These two highly synergistic technologies offer a time-proven and increasingly cost-effective solution for powering industrial grade applications, including:

- GPS sensors
- Asset trackers
- Environmental monitoring systems
- Smart agriculture (i.e. monitoring moisture, temperature, and rainfall)
- Marine buoys
- Machine-to-machine(MZM) communications
- Systems control and data automation

Solar/Li-ion hybrid solutions far outnumber the use of thermoelectric, electromagnetic, and piezoelectric energy harvesting devices that are deployed in highly specialized applications.

## When to Deploy Solar/Li-Ion Hybrid Technology

Generally, solar/Li-ion hybrid systems should be designed to offer enough storage capacity to power a remote wireless device for a minimum of five days on a single battery charge. Otherwise, you may run the risk of experiencing a temporary system shut-down during periods of continuous rain or heavy clouds.

When calculating your panel size, you must plan for the worst months (typically December in the Northern Hemisphere) and factor in regional variances in average daily solar energy capture.

For instance, parts of Washington State receive the equivalent of just one hour of full sun per day during winter months, so the PV panel may need to be enlarged to compensate for low average daily sunlight. Conversely, many sunny desert regions throughout the Southwest can receive up to four times more average daily solar energy during winter months, thus permitting the use of smaller and less-expensive PV panels. Understanding the regional operating environment serves to ensure that your solar/Li-ion hybrid solution will provide reliable long-term wireless power.


**IPS solar-powered parking meters provide municipalities with valuable real-time data and revenue streams. Use of an industrial-grade Li-ion battery provides a highly efficient and economical solution that operates for up to 20 years to minimize long-term maintenance costs.**

## Low-Power Design Can Cut Costs

The daily power consumption of an energy harvesting device can be influenced by numerous variables, including:

- Types of sensors used
- Energy consumed while the device is in "stand-by" mode
- Self-discharge rate of the battery
- Frequency of data collection and transmission
- Average current drawn during "active" mode, including power reserves for actuators, motors, and other advanced functionality.

For example, humidity, gas, motion, and light sensors can draw current in the 1-2 mA range, GPS sensors can draw 10-25mA of

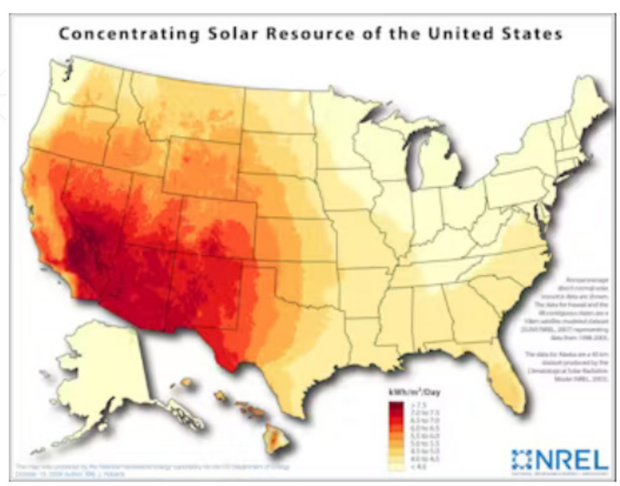

energy, and cellular-based radios may consume up to 1 A of energy while in "active" mode. There are various methods being deployed to reduce energy consumption, including:

- The use of low-power microprocessors and components
- Software solutions that shorten the time required for data sampling, interrogation, and transmission
- Reducing energy leakage while in "stand-by" mode

**COMPARING CONSUMER VS. INDUSTRIAL LI-ION INDUSTRIAL BATTERIES**

| | | TLI-1550 (AA) | Li-Ion |
|---|---|---|---|
| | | Industrial Grade | 18650 |
| Diameter (max) | [cm] | 1.51 | 1.86 |
| Length (max) | [cm] | 5.30 | 6.52 |
| Volume | [cc] | 9.49 | 17.71 |
| Nominal Voltage | [V] | 3.7 | 3.7 |
| Max Discharge Rate | [C] | 15C | 1.6C |
| Max Continuos Discharge Current | [A] | 5 | 5 |
| Capacity | [mAh] | 330 | 3000 |
| Energy Density | [Wh/l] | 129 | 627 |
| Power [RT] | [W/liter] | 1950 | 1045 |
| Power [-20C] | [W/liter] | > 630 | < 170 |
| Operating Temp | deg. C | -40 to +90 | -20 to +60 |
| Charging Temp | deg. C | -40 to +85 | 0 to +45 |
| Self Discharge rate | [%/Year] | <5 | <20 |
| Cycle Life | [100% DOD] | ~5000 | ~300 |
| Cycle Life | [75% DOD] | ~6250 | ~400 |
| Cycle Life | [50% DOD] | ~10000 | ~650 |
| Operating Life | [Years] | >20 | <5 |

Reducing the frequency of data collection and data transmission helps conserve energy, but introduces certain trade-offs that may be avoided by specifying a larger PV panel and/or by using a greater number of Li-ion batteries for energy storage.

Many IIoT sensors and edge devices are being intelligently designed to conserve energy by reducing the clock rates of processors and controllers, adjusting the sensing frequency and broadcast frequency, and by utilizing back-off strategies to reduce communication strength. These devices also utilize low-power communication protocols such as Cat-M1, NB-Iot, LoRa, SigFox, ZigBee, and WirelessHART.

In some situations, average daily energy consumption can be reduced to the point where an energy harvesting device is no longer needed, as the vast majority of industrial-grade remote wireless devices are powered by bobbin-type lithium thionyl chloride (LiSOCl2) batteries that can operate for up to 40 years without battery replacement.

## Extreme Environments Demand Industrial-Grade Solutions

Solar/Li-ion battery hybrids must be ruggedly designed to withstand the "worst-case" environmental conditions. PV panels designed for use in extreme environments should be rated to operate for a minimum of 10 years using high-efficiency cells constructed of the highest-quality materials available. If not, the performance of the PV panels could deteriorate from UV damage, delamination, or corrosion. In addition, consumer-grade Li-ion batteries are not designed to operate or recharge at extreme temperatures, thus necessitating the use of an industrial-grade Li-ion battery.

**CattleWatch hub collars have built-in photovoltaic panels that harvest solar energy. Industrial grade Li-ion rechargeable batteries store this energy to deliver the high pulses required to power remote satellite-based communications between the in-herd mesh network and the rancher.**

The solar/Li-ion hybrid charge circuit needs to be specifically designed to charge well from a solar panel. A panel that is rated at 6 V, 1 A will produce open circuit voltages as high as 7 V. The voltage will drop when connected to a load and continue to fluctuate with solar conditions and temperature. Likewise, current will also vary with solar conditions, temperature and angle of the panel toward the sun. There are a range of design solutions, including maximum power point chargers or linear regulators that can be deployed. Whatever you choose, be sure to capture the smallest amounts of solar energy possible, down to 1 mA, while not restricting your maximum charging rate during peak periods. A common mistake to avoid is setting a voltage target for the panel that is too high. On hot days, the panel voltage will drop and the power output could drop to zero.

Application requirements may dictate certain size and shape constraints for the PV panel, which could impact your ability to utilize an off-the-shelf solution. Custom-designed PV panels may be required based on the recommended voltage, requirements for custom coatings, and specialized mounting requirements.

The total daily estimated voltage output capacity of a PV panel is determined by its surface area and its efficiency. Monocrystalline cells typically operate at 19% efficiency, producing about 0.12 W per square inch (0.019 W per square centimeter). Back contact monocrystalline cells made by SunPower, LG, and others are capable of more than 22%.

They achieve this efficiency gain by moving all or most of the contact points (which appear as thin silver lines on the solar panel) to the back, freeing up more surface area for energy production on the front of the cells.

All PV panels require some form of protective coating: glass, urethane, ETFE, epoxy, or PET. Glass offers high UV resistance, but results in added weight and increased profile, and has a risk of shattering. Urethane is highly UV-resistant and is designed to last up to 10 years. ETFE is lightweight and highly water-resistant, but will have a shorter lifespan than glass or urethane. Epoxy and PET are not recommended for most industrial applications.

PV panels need to be mounted on some form of substrate, including Aluminum-Plastic-Aluminum, PCB plastic, or aluminum. The solar/Li-ion hybrid assembly can be mounted directly to the surface of the wireless device if the PV panel is properly angled to the sun or mounted separately, if necessary.

## Industrial Applications Demand Industrial-Grade Li-Ion Batteries

When an industrial-grade energy harvesting device is being deployed in a remote location or extreme environment, it is essential to specify a rechargeable Li-ion battery with sufficient energy storage capacity. In addition to five days of continuous operation, extended battery life can be critically important in remote, hard-to-access locations. The labor and travel expenses required to replace a battery can far exceed the cost of the battery itself.

**The Spoondrift Spotter operates from a moored or free-drifting position to remotely measure ocean waves, position, and surface currents. Each unit features a waterproof hull, a solar panel array, an industrial grade rechargeable Li-ion battery, and two-way wireless communications.**

Consumer-grade Li-ion batteries are ill-suited for long-term industrial deployments, as they have an expected service life of approximately five years and 500 recharge cycles. Consumer Li-ion cells also suffer from a limited temperature range of 0°C to 60°C and are unable to deliver the periodic high pulses required to power two-way wireless communications.

In contrast, industrial-grade Li-ion batteries are now available that that can deliver up to 20-year battery life with 5,000 full recharge cycles, offering an extended temperature range of −40°C to 85°C, and the ability to operate and recharge at extremely cold temperatures. These industrial grade batteries can also handle 15 A pulses and 5 A continuous discharge current to support two-way wireless communications and remote shut-off capabilities.

## Conclusion

The burgeoning IIoT is fueling the explosive growth of remote wireless connectivity, causing rising demand for solar/Li-ion hybrids that can bring reliable long-term power supply solutions to remote locations and extreme environments.

Achieving a lower total cost of ownership demands long-term thinking, as it makes economic sense to specify a power management solution that supports robust remote wireless connectivity while minimizing future maintenance expenses.

*to view this article online,* ☞ *click here*

☞ **BACK TO TABLE OF CONTENTS**


Image: Wrightstudio, Dreamstime

CHAPTER 3:

# Design Considerations for Bluetooth Mesh Across Industrial, Home Environments


BRIAN BEDROSIAN, Vice President of Marketing,
IoT Compute and Wireless Business Unit, *Cypress Semiconductor*


**BLE Mesh technology expands Bluetooth's reach into an array of applications, but it brings to bear numerous design tradeoffs in both hardware and software. Here's an intro to BLE Mesh and an overview of those tradeoffs framed in a smart-lighting context.**

Bluetooth is a ubiquitous communications protocol with countless applications in consumer electronics, healthcare, industrial automation, and asset tracking. With Bluetooth Low Energy (BLE) Mesh now added as a network layer, there are even greater opportunities for simultaneous control and monitoring of hundreds—even thousands—of devices. These new capabilities come with added complexity for developers, though.

Bluetooth's many advantages have given rise to its now-ubiquitous presence. The Bluetooth standards are maintained and advanced by the Bluetooth Special Interest Group (SIG), which as of this writing has some 33,000 members in 150 countries. The original standards specified point-to-point (1:1) connections, with multipoint (one-to-many or 1:m) and mesh (many-to-many or m:m) added later. Classic Bluetooth supports 1:1 and 1:m communications with both a Basic Rate (BR) and an Enhanced Data Rate (EDR). Bluetooth Low Energy (BLE) is the only mode that supports m:m mesh networking.

The advent of BLE Mesh networking has further expanded the scale and scope of potential applications for Bluetooth. With support for over 30,000 network nodes, BLE Mesh can handle applications that span large buildings, healthcare enterprises, and campuses.

Bluetooth adoption has grown rapidly and is forecasted to continue apace (**Fig. 1**). Such widespread use has made Bluetooth networking ubiquitous across virtually all platforms, including smartphones and tablets, smart watches, laptops, and peripherals, ranging from keyboards and mice to speakers and headsets. Such ubiquity across interfaces and widespread interoperability across brands establishes an installed base and developer/user familiarity that makes Bluetooth an obvious choice for many new applications.

## BLE Mesh Overview

Every device in a BLE Mesh network must meet fundamental requirements identified in the specifications. This section provides an overview of those requirements using terminology adopted by the Bluetooth SIG.

*Mesh-Network Topology*

A mesh-network topology enjoys two significant advantages: virtually unlimited scalability and high resiliency, both of which have contributed to the protocol's popularity among product design engineers. These advantages derive from the many-to-many communications that form multiple paths throughout the network from source to destination (**Fig. 2**).

The m:m connections assure successful communications even when multiple nodes have failed or been taken out of service, whether temporarily or permanently. Putting it another way: BLE Mesh networks can expand far and wide with no single points of failure.

*Mesh-Node Types*

The scalable topology enables BLE Mesh networks to support a theoretical maximum of 32,767 nodes—a number that places no practical limits on real-world applications. The standards define four types of nodes, and any single node can be configured to support multiple types:

- *Relay Nodes* retransmit or relay received messages to propagate them throughout the mesh network. Messages are only relayed when their time-to-live (TTL) value is greater than zero. Except for Low Power Nodes, all BLE Mesh devices should support this capability.
- *Low Power Nodes (LPNs)* are used primarily for battery-powered, low-duty-cycle sensors. To minimize power consumption, LPNs are normally assigned a companion "Friend Node" to serve as an intermediary for messages.
- *Friend Nodes* receive messages on behalf of their assigned LPN(s), storing them in a queue for later delivery. Each LPN periodically "wakes up" and polls its Friend Node to receive any new messages that might be in its queue.
- *Proxy Nodes* relay messages between the connection-oriented General ATTribute (GATT)

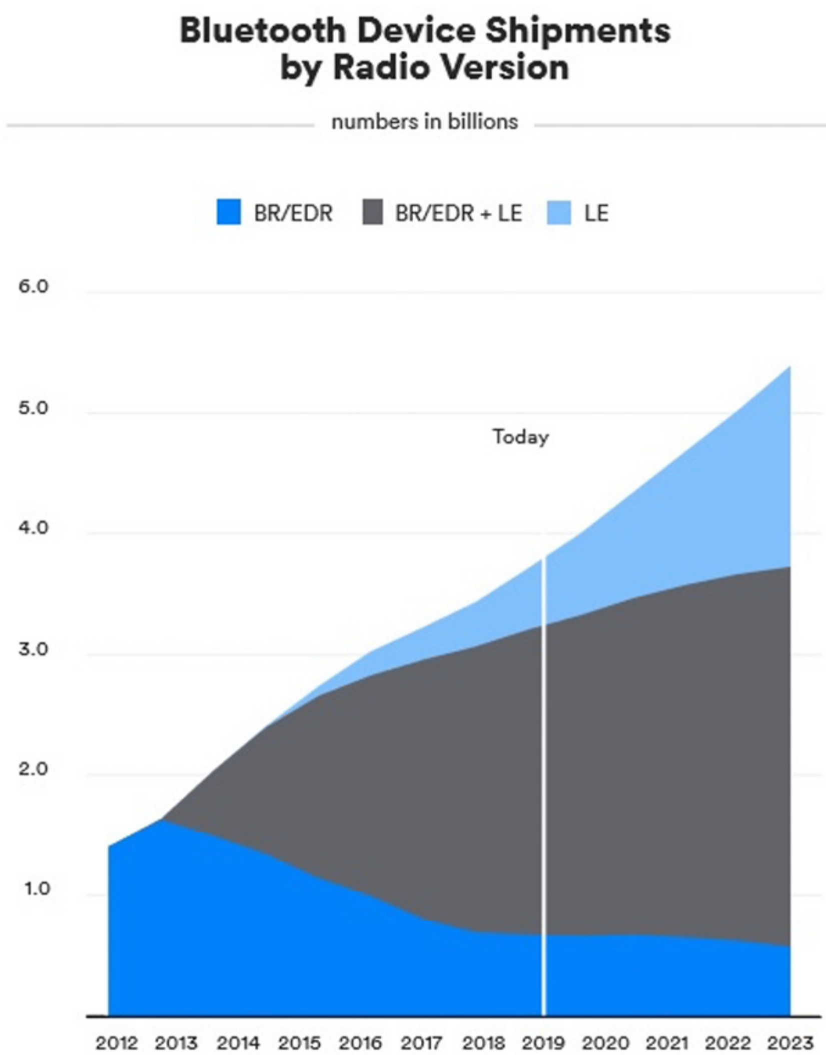


**1. The adoption of Bluetooth continues broadly across end markets, with the Low Energy mode taking an increasingly prominent role over time. (Source: Bluetooth SIG)**

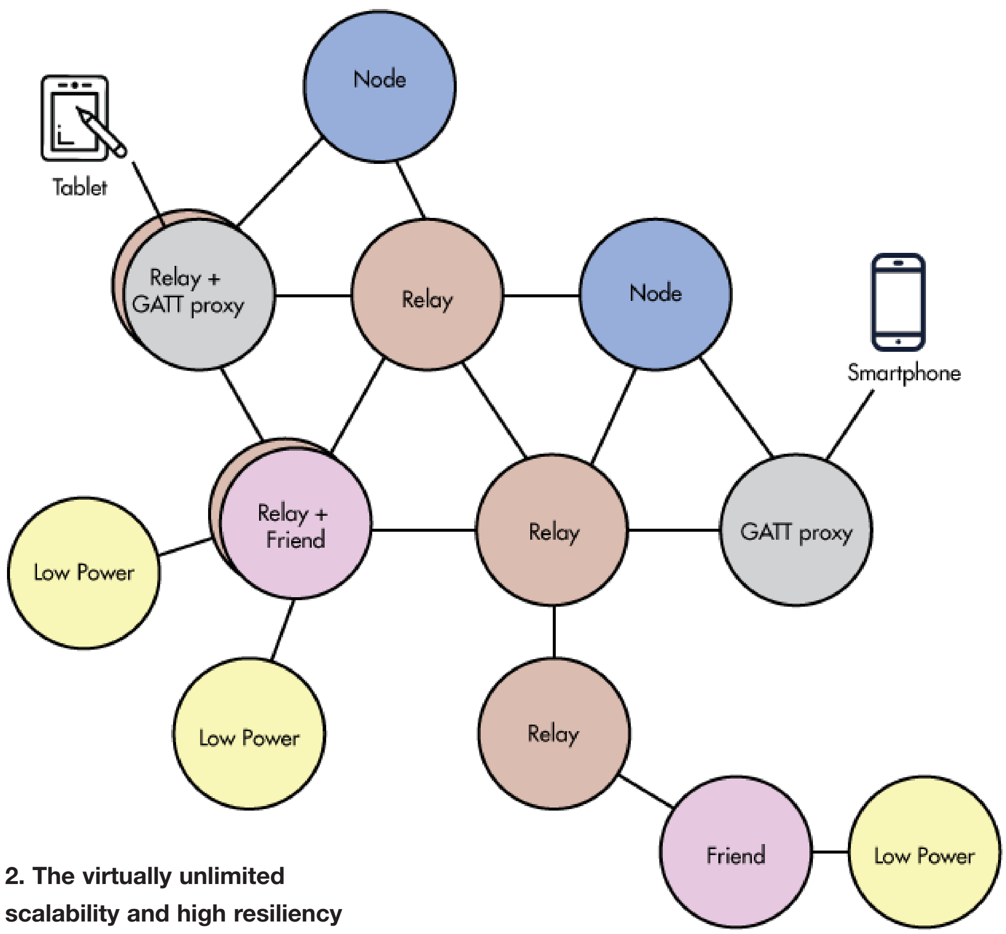


**2. The virtually unlimited scalability and high resiliency of BLE Mesh makes it suitable for a wide variety of new and demanding use cases.**

Bearer and the Advertising Bearer in the BLE Mesh network. This feature enables devices that support BLE (but not the BLE Mesh stack) to communicate with the mesh network without any need for a dedicated gateway or other special provision.

*Mesh-Node Element(s)*

Every node is required to have a primary identification element that defines its basic functionality. They also may optionally have one or more secondary elements to define additional functionality. For example, a switch (the primary element) might have an occupancy sensor as a secondary element, too. Or, an occupancy sensor (the primary element) might also have a light-level sensor as a secondary element.

*Mesh-Network Addresses*

There are four types of addresses in a BLE Mesh network, all of which are assigned during the secure provisioning process. Note that addresses are assigned to elements, which means that a node with multiple elements will have multiple addresses.

- *Unicast addresses* uniquely identify each individual element to enable point-to-point communications.
- *Group addresses* represent multiple elements to enable multicast communications. The Bluetooth SIG has defined four Fixed Group addresses: All-proxies, All-friends, All-relays, and All-nodes.
- *Virtual addresses* create virtual groups of elements or nodes to enable additional, dynamic multicast communications capabilities.
- *Unassigned addresses* identify elements that have yet to be provisioned with their Unicast, Group, and/or Virtual addresses.

*Mesh-Node Models*

BLE Mesh nodes employ one of three different types of models: Client, Server, or Control. These models are determined by a node's basic function or functions, as it's possible to implement more than one model in a single node.

- *Server models* contain and expose the state of an element; for example, a luminaire being on or off or at some intermediate brightness level.
- *Client models* interact with Server models by sending and receiving messages; for

example, when a switch is used to turn off or dim a luminaire.

- *Control models* combine Client and Server models in a single node, and typically include control logic (i.e., rules and behaviors). For example, an outdoor luminaire with an ambient light sensor may be configured to turn on at dusk and off at dawn and to turn on and off an indoor entryway luminaire.

*Messages and Messaging*

There are two categories of messages in a BLE Mesh network: Access messages for implementing an application and Control messages to manage the operation of the mesh network. Access messages are particularly important to product design engineers, as these are the means for requesting, sending, or changing the state values of elements; for example, turning on or off a luminaire.

The three types of Access messages are GET, SET, and STATUS. GET messages are sent to request state values from elements or groups of elements, which send STATUS messages in response. SET messages are sent to change state values in elements or groups of elements, which normally acknowledge the change by sending a STATUS message. SETs can also be unacknowledged, in which case no STATUS message is sent in response to the change. In addition to GET and SET responses, STATUS messages can be initiated by elements to periodically report their state value(s).

Communications occur within a BLE Mesh network as a managed flood of messages. "Flood" conveys how messages flow throughout the entire mesh topology while being "managed" to ensure efficient and effective use of available bandwidth. Key to a managed flood is the publish/subscribe group messaging used. Any node can publish or send a message, and every node is configured to subscribe to or act on only certain messages it receives, with all others being relayed as needed. The combination of these two aspects help give the BLE Mesh network its industry-leading price/performance, scalability, and dependability.

*Device Provisioning*

All devices installed must be provisioned before they can join a BLE Mesh network. The provisioning is normally performed by an application running on a smartphone, tablet, or PC. This is a distinct advantage of BLE Mesh because a Provisioner Application and, optionally, Mesh User Application Code (**Fig. 3**) can be run from a mesh participant. Provisioning is a deterministic and secure process that involves the exchange of keys for mutual authentication.

*Mesh-Network Security*

The BLE Mesh protocols were designed with robust security provisions. Provisioning, authentication, and messaging are subject to strong encryption. Network, application, and device security can all be applied separately using different keys, which provides the means to have multiple entities managing the different elements. These provisions protect against a full spectrum of physical and virtual attacks, including brute force, replay, man-in-

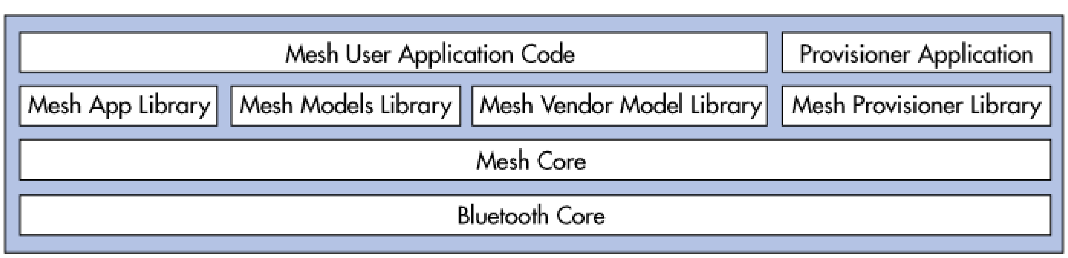


**3. The layered architecture enables software engineers to focus their development efforts exclusively on the Mesh User and Provisioner applications and not the BLE Mesh network.**

the-middle, and trashcan, and provide for user-data privacy.

*Software Architecture*

Figure 3 depicts the layers of software in BLE Mesh nodes. Note how the Mesh User and Provisioner Applications are located at the top of the architecture, above the libraries and core functions, to make them independent of the underlying BLE Mesh network. Note also how the Mesh Vendor Model Library makes it possible to add advanced, proprietary features while maintaining compatibility with the Bluetooth standards.

## Developing BLE Mesh Products

This section highlights the BLE Mesh development process, including the tools used, and explores key design considerations. While BLE Mesh networks are suitable for myriad applications and use cases, smart lighting is used here as an example for two reasons. One is that the standards were established with lighting as a basic use case. The other is the popularity of such "smart building" applications that are forecast to experience a compound annual growth rate (CAGR) of 46% in unit shipments through 2023, according to research conducted for the Bluetooth SIG.

The basic development process begins by defining the device's configuration, which must include, at a minimum, its node type, element(s), model, and all hardware- and/or application-specific states and callback functions. The next step involves implementing the callback functions for both the application and the BLE Mesh network.

Various open-source and product-specific tools are normally used during the development process from beginning to end. A specific and more detailed example of the effort involved is available in the Getting Started with Bluetooth Mesh application note published by Cypress Semiconductor.

For products supporting BLE Mesh networks, design considerations involve, at a minimum, device functions (or elements), mesh-network size, antenna range, memory requirements, power consumption, and cost. Separate design consideration must be given to the provisioning, management, and other software that runs on a smartphone, tablet, or PC.

As is common in all product development efforts, tradeoffs are often needed among the various design considerations. The need for such tradeoffs will be addressed here in the context of three products commonly used in smart-lighting applications: switches, sensors, and luminaires. As will be shown, the design consideration at the center of most tradeoffs is power consumption.

To maximize the versatility and, therefore, the benefit of smart lighting, users will want to be able to deploy switches and sensors virtually anywhere. Because some devices might be deployed in locations that aren't easily accessible, the use of energy harvesting with rechargeable batteries might be a desirable feature. This is especially true for sensors that measure illumination or sense the presence/absence of occupants. Physical switches, by contrast, are by their very nature readily accessible and can, therefore, be designed with replaceable primary batteries when located where ac power isn't available.

For BLE Mesh networks, it's advantageous for battery-powered devices to be designed as Low Power Nodes, which depend on the availability of Friend Node functionality. This requirement can be noted in the product's documentation and/or provided in another product in a family, such as a luminaire, which is assured to have an external power source.

Because a major advantage of BLE Mesh networks is their scalability, they can grow quite large. Though the maximum number of nodes and a high number of hops are rarely limiting factors, products should be designed to work in small-scale deployments that need to span a large area, potentially outdoors, with relatively few nodes. In these use cases, it may be necessary to facilitate increasing a product's antenna range and/or providing Relay Node functionality in a separate (and optionally dedicated) product.

Power consumption is inextricably linked to transmit range, and the BLE Mesh standards give design engineers some powerful (pun intended) capabilities to make the desired tradeoffs that might be needed. One such capability is being able to expand the range without increasing power consumption by decreasing the bandwidth. The converse capability is also possible; that is, boosting the bandwidth by decreasing the range, again while not increasing the power consumption.

Other aspects of a product can elevate the importance of its power consumption. For example, how sensitive does a sensor need to be, and how frequently does it need to poll for status changes? More frequent communication means more power consumption, heightening the need for a larger primary battery or energy harvesting for a rechargeable battery.

**Figure 4** shows an example of an evaluation board that can be used for product prototyping and development. Note the inclusion of three features commonly needed in smart-lighting applications: LEDs for luminaires, switches, and a PIR motion detector for use in an occupancy sensor. The module at the center of the board's bottom edge contains the CPU, memory, and antenna required for the BLE Mesh network, as well as for running the application software.

Given the ubiquity of Bluetooth in smartphones, tablets, and PCs, these systems are normally used for provisioning, configuring, and managing BLE Mesh products, such as the switches, sensors, and luminaires utilized in smart-lighting applications. As noted above, because the Mesh User and Provisioner Applications are layered above the BLE Mesh libraries and core functions, software developers can focus their efforts exclusively on the application and not the network. Here's a sampling of some of the functions Mesh User and

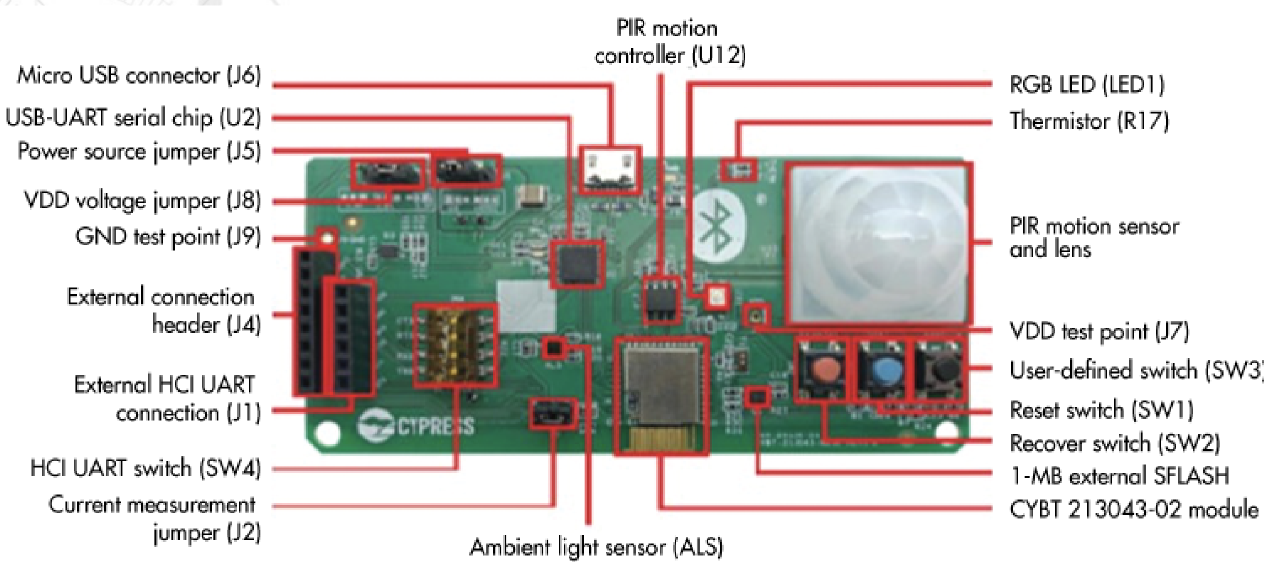


**4. Shown is an evaluation board used for prototyping and developing smart-lighting products. Not shown is the battery compartment on the back. (Source: Cypress Semiconductor)**

Provisioner Applications might need to support:

- Create and delete BLE Mesh networks and groups
- Provision and remove individual nodes
- Configure publications and subscriptions
- Publish GET messages to query the states of elements
- Publish SET messages, which for lighting applications might include On/Off, Level, Lightness and Lightness Hue, Saturation, Lightness Color Temperature, and Delta UV
- Publish SET messages for vendor data and vendor models
- Perform over-the-air (OTA) firmware upgrades

## Choosing BLE Mesh Components

One additional design consideration not covered in the previous section is cost, which is always an important factor in the development of any product. Cost always has two dimensions: designing the product and manufacturing it. Choosing the most cost-effective BLE Mesh components also has two dimensions: the silicon and its software. The silicon is the system-on-chip (SoC) or system-in-package (SiP) modules, and the software is the development tools that accompany them.

With interoperability as the raison d'êtrefor standards, the fundamental requirement when selecting BLE Mesh components is certified compliance with the Bluetooth SIG standards. This applies equally to both the silicon and the software, including the full Bluetooth BR/EDR and BLE Mesh protocol stacks and all pertinent libraries. Using SoCs, SiPs, and other components certified by the Bluetooth SIG eliminates the need for design engineers to conduct rigorous qualification and interoperability tests.

As of this writing, certification is available for version 5.0 of the Bluetooth core specifications and for version 1.0 of the BLE Mesh specifications. Version 2.0 of the BLE Mesh specifications are expected to be published in 2020.

When choosing a BLE Mesh platform, designer engineers should seek a solution that meets most or all of the following criteria:

- A family of modules to accommodate different needs—from the basic battery-powered sensor to the most sophisticated devices that might be needed now and in the foreseeable future.
- Fully integrated modules that minimize the need for external components, accelerate time-to-market, and reduce development and manufacturing costs.
- Ultra-low-power designs with the types of antennae and transmit powers needed to accommodate all anticipated node-to-node distances.
- Adequate CPU, memory (flash and RAM), and I/O for all foreseeable applications and upgrades.

For applications in which the BLE Mesh network may need to communicate with a Wi-Fi network, some form of gateway functionality will be required. For example, a home security system could be used to turn on or off certain lights to simulate people being at home or be deactivated with an authorized code. In these situations, an SoC or SiP combo-module that supports both Wi-Fi and BLE networking simplifies the design effort.

Because the software-development environment, tools, and libraries are just as important as the silicon, endeavor to find a solution that meets most or all of the following criteria:

- An easy-to-use integrated development environment (IDE) that abstracts the complexity of the underlying protocols

- Software development kits (SDKs), sample software, prototyping hardware, and a developer community to assist in software development and testing efforts
- Reference software designs for control applications running on the Android, iOS, Linux, and Windows operating systems

Bluetooth is already ubiquitous in "personal area networking" applications, and the advent of BLE Mesh significantly expands both the scale and scope of potential use cases for this popular protocol. The combination of virtually unlimited scalability and high resiliency afforded by BLE Mesh networking now enables Bluetooth applications to span buildings, campuses, and even entire cities. While the smart-lighting application used here serves as a good example, the potential use cases for BLE Mesh are limited only by the imagination.

Its ubiquity also gives Bluetooth another important advantage: the availability of proven techniques, tools, and software needed to develop new products. Hardware and software engineers interested in learning more about these resources are encouraged to review the documents and links listed in the References section below.

The many resources now available make it easy to get started with BLE Mesh. Proof-of-concept designs can be created quickly and affordably using inexpensive evaluation boards, IDEs, and SDKs. A solid head start may also be available in open-source or vendor-supplied sample application software. As a leading supplier of Bluetooth-certified silicon and software, Cypress Semiconductor provides an array of resources at www.cypress.com/products/ble-bluetooth.

By offering enormous upside potential with minimal downside risk and a short time-to-market, BLE Mesh is destined to become the network of choice for a growing number of residential, commercial, and industrial applications.

**References**

Bluetooth Mesh Networking - An Introduction for Developers
Bluetooth-Mesh-Paving-the-Way-for-Smart-Lighting
BLE Mesh specifications: https://www.bluetooth.com/specifications/mesh-specifications
Getting Started with Bluetooth Mesh (AppNote AN227069)

*to view this article online,* ☞ *click here*

☞ **BACK TO TABLE OF CONTENTS**



Image: Wrightstudio, Dreamstime

CHAPTER 4:

# LPWAN for All: What's the Future of Low-Power, Long-Range Connectivity?

MARC PÉGULU, Vice President, IoT Strategy and Products, *Semtech*

**LPWAN's future looks good amid the shifting IoT technology landscape.**

As 5G connectivity and infrastructure continues to grow, the industry has wondered: Where does this leave low-power wide-area networks, or LPWAN? Will widespread high-bandwidth connectivity render it obsolete?

Here's a spoiler: There's no single network that will render all others obsolete. Let's explore why that is, and how the future of the Internet of Things (IoT) will be built through open networks and collaboration.

## Reliable Long-Range Coverage

Let's start with establishing the benefits of LPWAN and why it's been successful in the first place. The key benefit is in the name: low-power, wide-area network. By nature, LPWAN technology has the ability to transmit information across long distances without using much power. This is especially beneficial for IoT applications, which require small packets of information to be communicated regularly, often in situations that aren't accessible to other networks.

For example, LPWAN is a great option for monitoring operations remotely, whether that's in the middle of the ocean or in a rural agriculture setting. Small amounts of information like machinery status, output, and environmental monitoring can be shared regularly, alerting of any activity out of the norm before there's an issue. In many remote settings, other networks can't offer the same consistent, cost-effective, and reliable coverage.

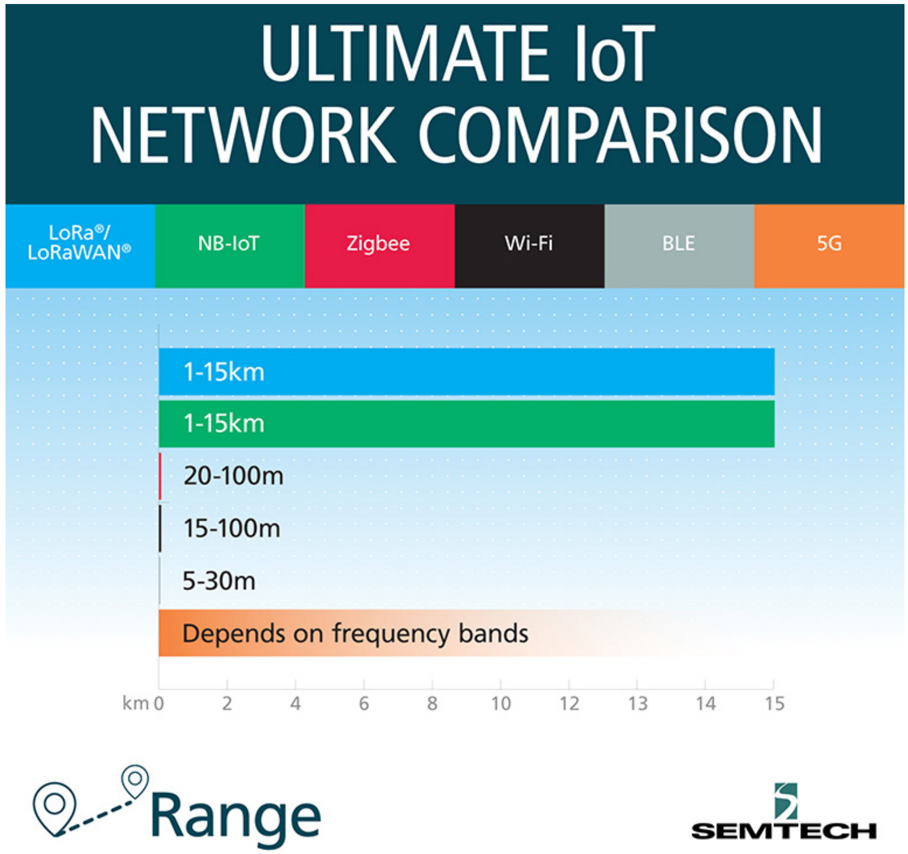


**The chart shows how the range capabilities stack up between six IoT network technologies.**

## Industry Growth

There's also growth in the industry. LoRaWAN, one of the most widespread LPWAN protocols, has a very strong ecosystem and plenty of deployed use cases. It's managed by the nonprofit LoRa Alliance organization. At the time of this article, there are 163 private and public LoRaWAN network operators, a number that continues to climb.

Because it's an open protocol, LoRaWAN gives customers the flexibility they need to develop their IoT applications. In addition, LoRaWAN is a proven standard that's endorsed by the International Telecommunication Union (ITU) and backed by a strong, global ecosystem from chip-to-cloud companies like AWS, Microsoft, Orange, and Semtech.

Devices are ramping up as well: ABI Research predictsthat by 2026, 50% of non-cellular LPWA will feature Semtech's LoRa devices. And Juniper Research reportsthat the overall cellular IoT market (including narrowband IoT and LTE-M) will grow by 95% between 2022 to 2026.

## Interoperability and the Competition

Often the "LPWAN is dead" argument is based in a desire for a complex, low-latency network that can handle massive amounts of data. It's true—LPWAN can't transmit the same amount of data as a network like 5G.

That said, networks aren't competitive, they're complementary. In many IoT use cases, 5G would be the equivalent of using a sledgehammer to crack a nut—an excessive amount of power used where a right-sized technology would be a better fit. As the IoT grows, the capabilities of each network will layer over the next to support the billions of messages regularly transmitting across the IoT.

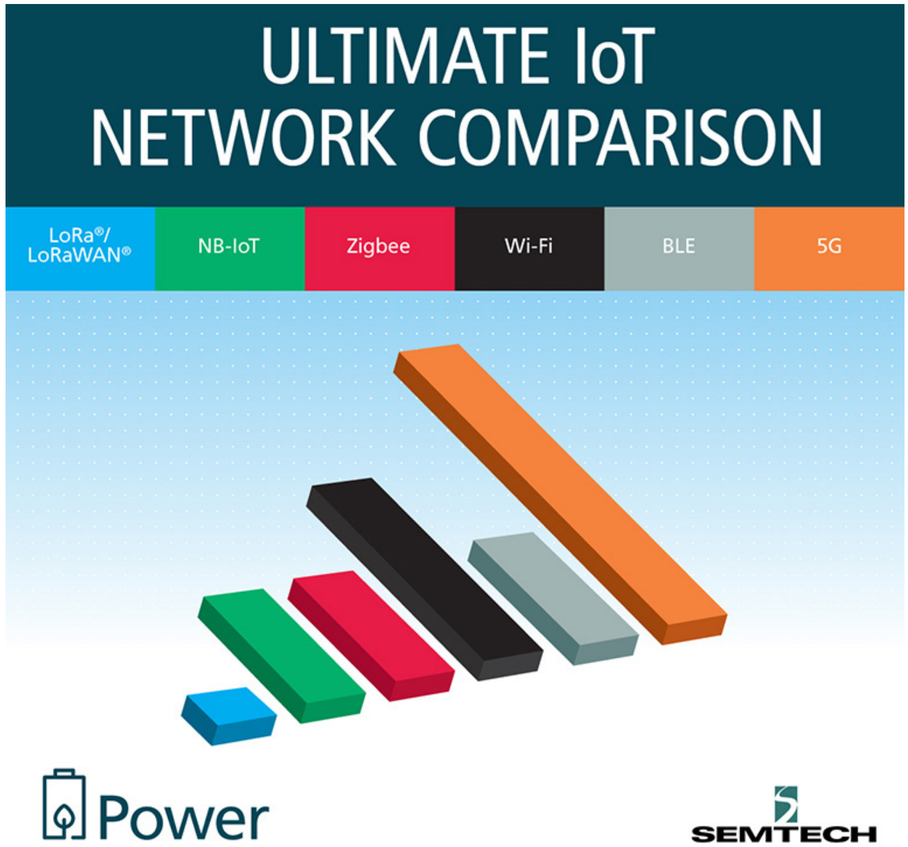


**Comparing the power usage between six IoT network technologies.**

Take, for example, how LPWAN and Bluetooth Low Energy (BLE) can work together. BLE's versatility and ubiquity makes it a go-to choice for short-range communications like mobile payments or consumer IoT. However, it's limited in how far it can transmit data, which just happens to be where LPWAN's strength lies. By pairing the two technologies, LPWAN extends BLE's capabilities and makes deployments more flexible.

That's why interoperability is key for IoT. Everything needs to talk to one another for its true purpose to be fulfilled. As 5G targets low-latency and high-throughput applications, the applications using LoRa and LoRaWAN will make up a larger portion of the massive IoT space, led by mobile operators, unlicensed spectrum operators, and enterprises across private rollouts.

## Final Thoughts

The market is heading toward a multi-radio access network (or Multi-RAN) strategy where each network is complementary rather than competitive. So is there an end in sight for LPWAN? No. As the world of IoT continues to grow with new applications every day, the need for a low-power, long-range network grows as well.

*Semtech, the Semtech logo and LoRa are registered trademarks or service marks of Semtech Corporation or its affiliates.*

*to view this article online, ☞ click here*

☞ **BACK TO TABLE OF CONTENTS**


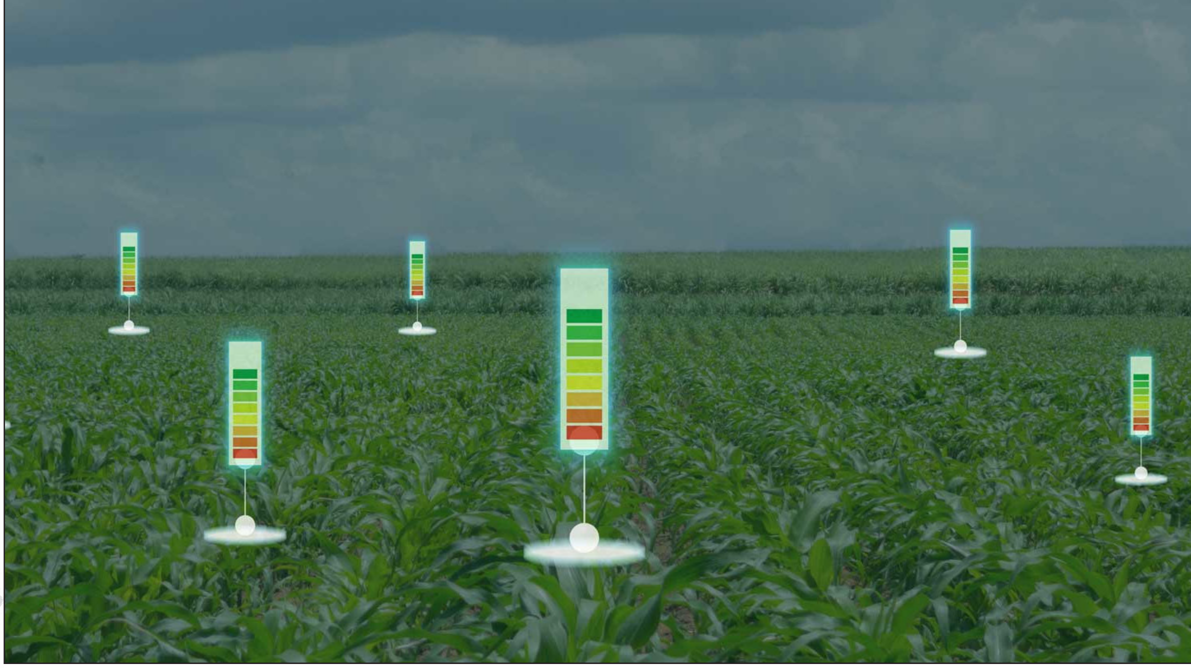

Image: Dreamstime_Ekkasit919_124220683

CHAPTER 5:

# LoRaWAN Brings the IoT Across Longer Distances (Part 1): The Technology

LEE GOLDBERG, Contributing Editor, *Electronic Design*

**LoRaWAN is a low-power, low-cost wireless technology that can enable secure, highly reliable communications to smart buildings, smart cities, smart agriculture, and more.**

Low-power wireless network (LPWN) technologies such as LoRaWAN (long-range wide area network) and Sigfox help extend the IoT into areas and applications that aren't practical for 5G services to support. Though their low data rates (250 bits/s to 50 kb/s) and relatively high latency aren't suitable for some applications, the ability to perform reliably in noisy, congested environments across long distances, while operating for nearly a decade on a single low-voltage battery, make them an excellent solution for many other apps. They include smart agriculture, remote environmental sensing, and monitoring

**1. LoRaWAN is a robust, low-power, long-range wireless protocol that can bring the IoT to nearly any embedded application.**

the structural health of bridges, railways, and other critical elements of the national infrastructure (**Fig. 1**).

Since many of these applications also require low per-node CAPEX/OPEX costs to be feasible, LoRaWAN's open-source specification, use of unlicensed frequency bands, and its ability to leverage existing public networks have helped it gain popularity in many parts

**2. The LoRaWAN transmitter generates chirp signals by varying their frequency over time and keeping phase constant between adjacent symbols (a). A comparison of several popular LPWAN technologies illustrates how LoRaWAN has been optimized to support robust transmission of low-speed data across long distances while delivering extremely long battery life (b). (Credits: Noreen, Bounceur, and Clavier, Ref. 2, and LoRa Alliance)**

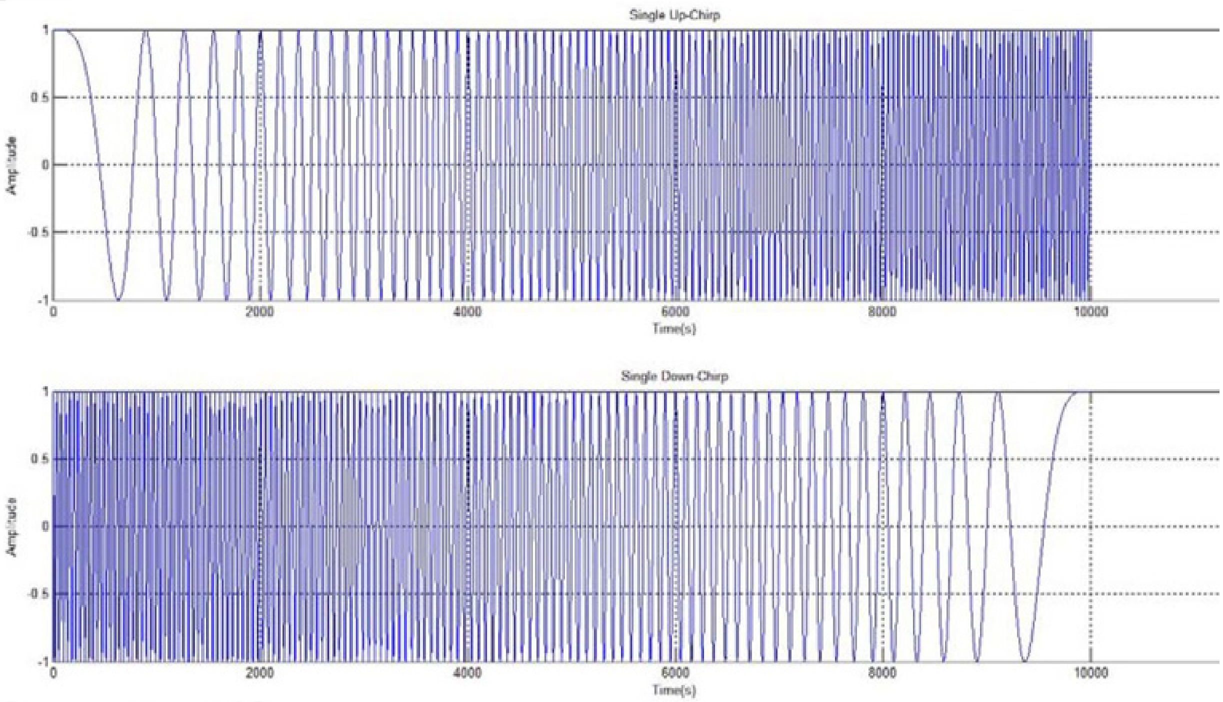


**(a)**

| Feature | LoRaWAN | Narrow-Band | LTE Cat-1 2016 (Rel12) | LTE Cat-M 2018 (Rel13) | NB-LTE 2019(Rel13+) |
|---|---|---|---|---|---|
| Modulation | SS Chirp | UNB / GFSK/BPSK | OFDMA | OFDMA | OFDMA |
| Rx bandwidth | 500 - 125 KHz | 100 Hz | 20 MHz | 20 - 1.4 MHz | 200 KHz |
| Data Rate | 290bps - 50Kbps | 100 bit/sec<br>12 / 8 bytes Max | 10 Mbit/sec | 200kbps – 1Mbps | ~20K bit/sec |
| Max. # Msgs/day | Unlimited | UL: 140 msgs/day | Unlimited | Unlimited | Unlimited |
| Max Output Power | 20 dBm | 20 dBm | 23 - 46 dBm | 23/30 dBm | 20 dBm |
| Link Budget | 154 dB | 151 dB | 130 dB+ | 146 dB | 150 dB |
| Batery lifetime - 2000mAh | 105 months | 90 months | | 18 months | |
| Power Efficiency | Very High | Very High | Low | Medium | Med high |
| Interference immunity | Very high | Low | Medium | Medium | Low |
| Coexistence | Yes | No | Yes | Yes | No |
| Security | Yes | No | Yes | Yes | Yes |
| Mobility / localization | Yes | Limited mobility,<br>No loc | Mobility | Mobility | Limited Mobility<br>No Loc |

**(b)**

of North America, Europe, and elsewhere.

In this article, we'll look at the technologies that underlie the LoRaWAN wireless data standard and a few of the applications it makes possible.

### What is LoRaWAN?

LoRaWAN is a type of low-power wide-area network (LPWAN) that uses open-source technology and transmits over unlicensed frequency bands. It's based on the LoRa protocol,[1] which is designed to wirelessly connect battery-operated "things" to the internet in regional, national, or global networks. LoRaWAN was also developed specifically to support the key functional requirements of the Internet of Things (IoT), such as bidirectional communications, end-to-end security, mobility, and localization services.

The protocol operates on the globally available sub-1-GHz ISM bands, enabling low-power embedded radios to use very-low-power transceivers to cover large distances with true bidirectional communications. Instead of the frequency-shift-key (FSK) modulation techniques used by many wireless networks, LoRaWAN employs chirp-spread-spectrum (CSS) modulation2 (**Fig. 2a**).

CCS modulation maintains the same low-power characteristics as FSK modulation, but it significantly increases the communication range (**Fig. 2b**). CSS has been proven over decades of use in military and space communication, where its resistance to interference and high coding gain has provided reliable ultra-low communications over long distances.

### The Transmission Benefits of LoRa

To further enhance signal integrity, LoRa uses a wide transmission band (typically 125 kHz), which makes it extremely resistant to channel noise, long-term relative frequency, Doppler effects, and fading.

LoRa also adds a unique capability called adaptive data rate (ADR) to increase transmission distance. Geoff Mulligan, creator of the LoRa protocol, explained, "ADR

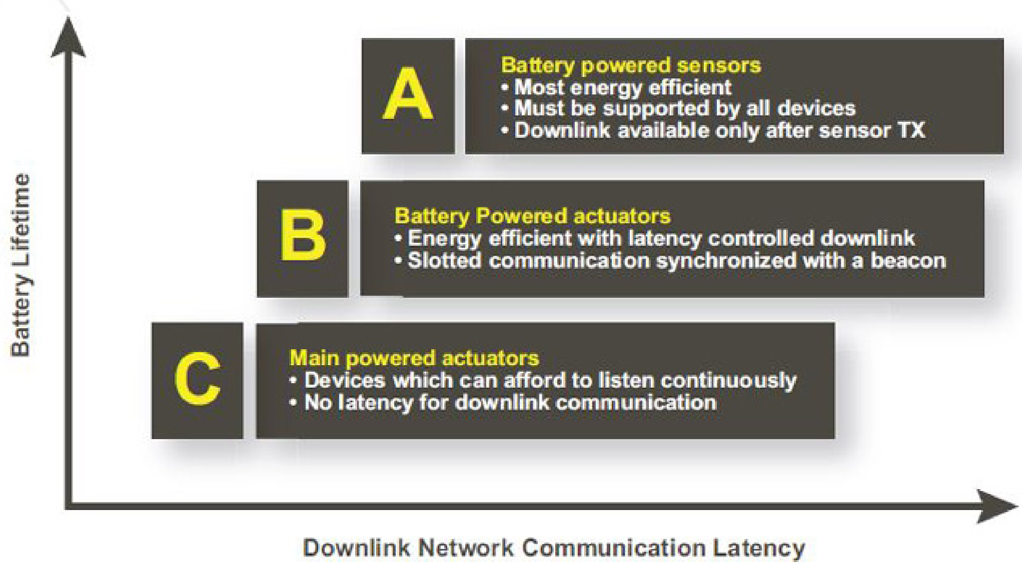


**3. The LoRaWAN technology stack supports three distinct classes of endpoint devices.**
LoRa Alliance

allows the LoRaWAN network to automatically and independently choose the most appropriate output power and data rate for optimal and efficient communications by each device, while controlling the SNR (signal-to-noise ratio) and link budget. Devices that are closer to the base station can transmit at higher data rates, while devices further away transmit at lower rates."

Mulligan added, "This transmission model makes it possible to match the power and speed requirements of a particular application, rather than using a 'one-size-fits-all' approach. By using different ADRs (SF1 through SF12), devices can transmit up to 50 kb/s or over distances up to 20 km with link budgets up to −148 dB. LoRa's robust link-management scheme enables it to support lower data rates across even larger distances."

Because different applications have different communication requirements, the LoRaWAN specification supports three different "classes" of devices (**Fig. 3**):

- Class A *(Asynchronous)* devices are generally battery-operated sensors that communicate on an as-needed basis. Examples include temperature and motion sensors that send data when triggered by some external anomaly, such as a high or low temperature or sensed motion. In addition, these devices use the "Queen's Protocol," meaning they can only be spoken to after they have first spoken (i.e., wait for a response after transmitting). Class A devices offer the most power-efficient mode, enabling field deployments of sensor devices that operate for many years on a single battery.
- *Class B* devices are also generally battery-operated, but with the added capability of enabling the application to send data to the sensor on a predefined time interval (the beacon period). This still provides very-low-power operations, while allowing

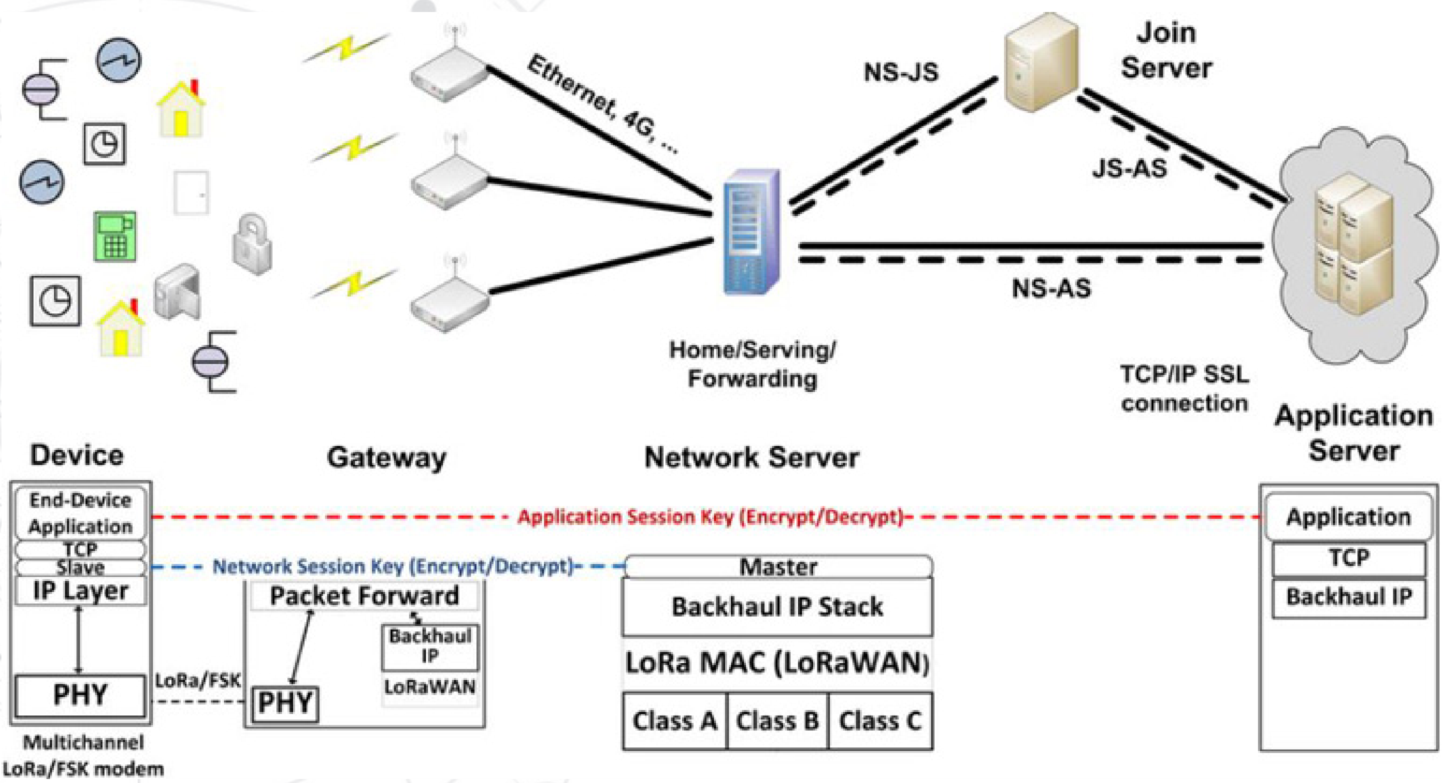


**4. The LoRaWAN standard uses a star topology and AES security.** ResearchGate for efficient two-way communications that don't require the application to wait for communication from the sensor or end device. Using a Class B device can reduce the battery lifetime, but at the same time, it significantly reduces latency and provides deterministic communication.

- *Class* C devices are "mains-powered" in most cases, or have some sort of external power source (e.g., photovoltaic). These end points are "always on" and listening for messages from the application. A typical use for a Class C device is to provide communications for an actuator or controller. Here, the application can send commands or data to the end point whenever necessary, thereby nearly eliminating the latencies of other device classes, but with higher power consumption.

To make the most effective use of its wireless protocol, the LoRa Alliance chose a star-type network topology. It's similar to the topology used by Wi-Fi networks (with endpoints and access points), with one significant difference: All access points receive messages transmitted by end points (**Fig. 4**).

That very important difference gives LoRaWAN some unique capabilities. First and foremost, the LoRaWAN protocol makes use of the fact that radio transmissions are "broadcast," meaning many devices can and will receive the transmitted message.

Unlike the unicast messaging protocol used by ZigBee, Bluetooth, and Wi-Fi, the LoRaWAN protocol provides a more robust communication path and eliminates the single points of failure that are common in these other systems. And, by designing in the concept of multiple independent receivers, the LoRaWAN protocol can provide intrinsic support for endpoint geolocation without resorting to power- hungry and expensive GPS radios. Instead, LoRaWAN nodes can apply a technique called time difference of arrival (TDOA) to triangulate the precise location of any endpoint.

## LoRaWAN Built with Bottom-Up Security

Security was a primary consideration during the development of the LoRaWAN protocol. As a result, security and intrusion resistance is baked into its DNA.

The protocol's architecture utilizes two separate AES keys (Advanced Encryption Standard, developed by NIST and adopted by the U.S. government). The first key (the network session key, as seen in Figure 4) is used to protect and secure data during transmission across the LoRaWAN network. The second key (the application session key) is used to protect and secure the data end-to-end from the endpoint to the sensor application.

The purpose for using two distinct keys is to permit devices to roam from one network to another. The network session key is known only to the endpoints and the local network, while the application session key is only known to the endpoint and end-sensor application.

In this way, a device can move from one network to another (roam), and the network can ensure that packets not be intercepted or inserted into the local network (that the network is secure). At the same time, the end device and application can be assured of the data's confidentiality and integrity, no matter which network is used to transfer the data.

*Part 2 of this series looks at how LoRaWAN can be used to bring the intelligence of the IoT to agriculture, commerce, and a wide range of other public infrastructure applications.*

**Read more articles in the TechXchange: Wireless IoT Technologies.**

**References**

1. “LoRa and LoRaWAN: A Technical Overview,” Semtech, 2019.
2. Umber Noreen, Ahc`ene Bounceur, Laurent Clavier, “A Study of LoRa Low Power and Wide Area Network Technology,” ATSIP’2017.
3. “What Is LoRa?,” LoRa Alliance.
4. LoRaWAN Technical Specifications, LoRa Alliance.

Additional white papers on wireless security, LPWA technology, applications, and other topics may be found on the LoRa Alliance website.

*to view this article online,* ☞ *click here*

☞ **BACK TO TABLE OF CONTENTS**


CHAPTER 6:

# The Benefits of Wirelessly Monitoring Tank Levels

STEPHEN J. MRAZ, former Editor, *Machine Design*

**Wireless sensors and data networks can give managers accurate, up-to-date information to make better planning and operational decisions.**

Monitoring and managing the contents inside the hundreds of thousands of storage tanks, totes and containers located around the globe in a wide variety of locations and environments can help owners and asset managers improve efficiencies and increase productivity and profitability. This is true for containers installed indoors or outdoors, above ground or below it, and regardless of what the contents are.

A relatively new and effective way to carry out that task is to use wireless connectivity and radar and ultrasonic sensors. They can overcome common challenges related to container location, contents and quantity to give users timely access to accurate level information for better resource allocation and lower costs.

## Traditional Methods

Several traditional and low-tech methods and tools such as visual inspections and measuring sticks have long been used to measure levels inside containers. They are easy to implement and require little or no investment in equipment. But they are also reactive, imprecise, prone to errors and labor intensive. And those simple methods can get more costly and time consuming depending on how many tanks must be inspected, how far away the tanks are and the difficulty—even risks—in accessing the tanks. And if technicians must travel to get to the tanks, there may be wear-and-tear on vehicles, travel expenses, and lost productivity of staff.

Traditional level measurements also rely on manual reporting. Data from each tank is collected and recorded by staff at the site. This data is typically logged onto a spreadsheet or entered into a management system when the staff member returns to the office or can access a network connected PC or another device. In many cases, tank data is out-of-date by the time it has been logged.

### More Accurate Measurements

Using magnetostrictive floats is a common measurement option that does not depend on human labor. However, these devices must contact the container's contents to determine their levels. This limits the mechanical life of the device. This also may not be acceptable in applications where contact with measurement tools could contaminate the contents. Additionally, prolonged contact may degrade the performance and accuracy of the float.

**READ MORE: Switch Tips: Ultrasonic level switches**

Radar and ultrasonic sensors, on the other hand, use electromagnetic or sound waves to detect objects and tank levels, so they do not need to contact the materials they are measuring. They do not rely on light for detection, so they are largely unaffected by conditions that typically challenge photoelectric sensors such as target color, reflectivity, transparency and wet environments. And if the ultrasonic or radar sensors generate proportional analog outputs, they can deliver more accurate measurements which can easily be converted into liters, gallons, feet or meters; this is a preferred method for monitoring tank and container levels.

### Collecting Data

Connecting measurement tools to wireless networks is typically less costly than doing so on wired networks. Eliminating the need for wires lets companies cut the costs of installation, alterations and expanding the network. These costs include hardware and labor costs as well as time.

In a basic setup, a wireless node connects to a measuring device, such as an ultrasonic sensor installed on the container. Each node is bound to a wireless controller via an integrated gateway.

Next, a site survey verifies the connection between deployed nodes and the gateway. The system is then configured and inspection parameters and alarm thresholds set. Lastly, an Ethernet or cellular connection is established, giving remote access to the data collected from the tanks to authorized users. A system like this can monitor many containers and be set up with minimal time.

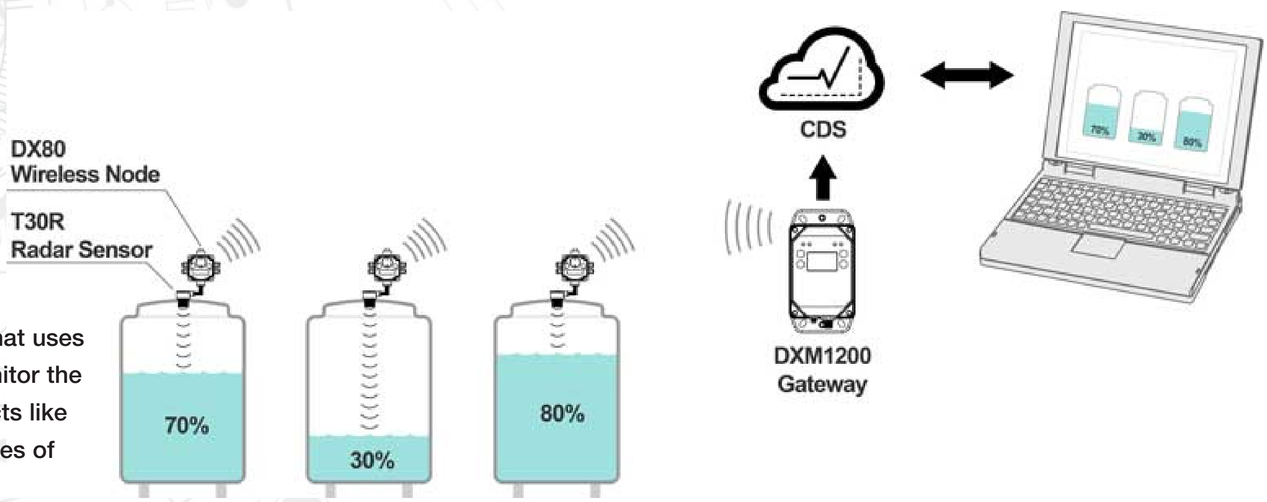


This is a typical setup that uses radar to detect and monitor the level of fluids (or products like grain and sand) in a series of tanks. Banner Engineering

## Wireless Benefits

Wireless tank monitors give managers access to timely, accurate information via remote access to the network, providing significant gains in efficiency and productivity. So, rather than sending staff to a site based on forecasted expectations of container levels or by following an arbitrary schedule, managers can strategically plan visits, as well as routes and resources, based on actual needs. This makes it easy to keep tanks filled and processes running without overstocking a site.

In collection or filling applications, this information can prevent overfills and ensure staff and equipment do not need to be sent to a site to empty, move, or remove under-filled containers. For companies providing either of these services, this can mean added capacity from staff and equipment to better serve customers and attract new ones. Additionally, accurate data collected over time can lead to more accurate forecasts that inform business and support decisions.

Remote monitoring can reduce the chances of problems getting out of control. Sudden changes in use, equipment failures or other unexpected events can dramatically change content levels. An empty tank can affect equipment that depends on its contents, and processes may come to an unexpected halt. Overfills can be wasteful, can damage equipment and may have environmental consequences. A wireless system can be easily set to alert key personnel via text or email if tank levels are outside established parameters. This lets workers respond to a problem before it has time to become critical.

Businesses change over time and companies may need to add or remove containers

**This Banner ultrasonic sensor measures tank levels and transmits the data wirelessly.**
Banner Engineering

to keep up with new demands. Wireless networks are versatile, highly scalable and can be easily adapted to meet these changes. Containers can be moved as often as needed without altering the system or disconnecting the node and measuring device. Nodes and sensors measuring levels in new containers can be added to a network, or a new network can be deployed to accommodate a new group of tanks, all with relative ease and without incurring substantial costs or large commitments of time and labor.

**READ MORE: Switch Tips: Float-level switches**

Although wireless monitoring provides numerous benefits, there are also many reasons facilities may choose to keep or to implement wired sensors. Wired and wireless monitors can also co-exist. The right solution depends on the use and customer needs.

If the liquid level in a facility's tanks constantly and rapidly changes, it may need a faster level sampling rate than wireless monitoring can provide. Or maybe wireless technology won't work well due to obstructions or distance between assets. With wired hardware, there's never a worry about signal strength or sample rates.

For applications that rely on a wired approach, an overlay network, such as Banner Engineering's Snap Signal, can bring tank-level data onto an industrial network or the cloud. An overlay network connects sensors of all signal types, discrete or analog, letting facility mangers focus on existing tank-level sensors without disrupting the legacy controls. If plant managers need to measure level, temperature and humidity, they can simply add the corresponding sensor. Overlay components convert each of these signals to a smart serial protocol so they can all communicate on a common network. A gateway can be added to gather information in one place and send it wherever it is needed.

**READ MORE: Measuring Liquid Levels with LVDTs**

Tank monitors that use radar or ultrasonic sensors in a wireless network can be set up quickly, cost-effectively, and without large commitments of human labor or infrastructure changes. They can monitor many containers and be adapted to meet changing needs and scaled to accommodate new containers. Automated alerts, for example, can be set up to notify staff to take action in case of potential problems, minimizing the chance of emergencies developing.

Data generated by these monitors is more accurate and can be accessed from anywhere and at any time. This lets managers make more informed decisions about assets, how they are used, when they should be serviced, and what staff and resources are needed to service them. Finally, if a wired solution works better than wireless for the facility, an overlay network such as Banner's Snap Signal offers a way to connect sensors of any signal type and convert them to a common communications protocol on the facility's network or in the cloud.

For more information about wireless sensors for remotely monitoring tank levels, click here.

*to view this article online,* ☞ *click here*

☞ **BACK TO TABLE OF CONTENTS**