

Wireless connectivity for the Internet of Things: One size does not fit all



Nick Lethaby
IoT Ecosystem Manager
Texas Instruments

In the rapidly growing Internet of Things (IoT), applications from personal electronics to industrial machines and sensors connect wirelessly to the internet. Covering a wide variety of use cases in various environments and serving diverse requirements, no single wireless standard can adequately prevail.

With numerous standards deployed in the market, spread over multiple frequency bands and using different communication protocols, choosing the right wireless connectivity technology for an IoT application can be quite challenging. In this paper, we will review the predominant wireless connectivity technologies, discuss their key technical concepts and engineering trade-offs, and provide guidelines for selecting the right wireless technology for different applications. We will focus specifically on wireless technologies that operate in the industrial, scientific and medical (ISM) band where spectrum use is free, rather than technologies like cellular where the purchase of licensed spectrum drives up cost.

Frequency bands and worldwide regulations

Agencies such as the Federal Communications Commission (FCC) in the U.S. and the Conference of Postal and Telecommunications Administrations (CEPT) in Europe regulate radio transmissions worldwide. These organizations allocate frequency bands for specific use and drive radio transmitter standards and certification plans. The usable spectrum in most regions is designated as “licensed,” which means that users need to buy a license from the local regulator to operate a radio transmitter in a designated frequency channel. A familiar example for licensed frequency-band use is cellular communication. Government auctions worldwide sell spectrum bands to mobile operators to regulate commercial frequency allocation.

The International Telecommunication Union’s Radio communication (ITU-R) Sector, which coordinates the shared global use of the radio spectrum, has reserved several frequency bands for industrial, scientific and medical (ISM) applications. ISM bands are unlicensed, and vary slightly from country to country. Used by wireless communication systems such as remote controls, cordless phones and Wi-Fi®, popular ISM bands in recent years have included 433MHz, 868MHz, 915MHz, 2.4GHz and 5GHz. **Figure 1** shows a map of popular ISM bands allocated worldwide. The 2.4GHz band became very popular because it is allowed for unlicensed use in all regions. This band’s ubiquity makes development and distribution of 2.4GHz-based products across nations easier.




Figure 1. Worldwide unlicensed frequency bands.

As a general rule, higher-frequency bands offer more channels and more bandwidth, and thus can serve larger networks and drive more data throughput. Lower-frequency radio waves propagate better than higher-frequency waves and can thus achieve better range, especially inside buildings.

Communication protocols

Communication systems use a set of rules and standards to format data and control data exchange. The most common model in data communication systems is the Open Systems Interconnection (OSI) model, which breaks communication into functional layers and enables easier implementation of scalable and interoperable networks. The OSI model has seven layers; layers two through five of this model are shown in Figure 2 along with how these map to the different protocols of a wireless the TCP/IP stack. Wireless connectivity protocols are generally implemented at layer 2, with layer 1 representing the physical layer such as Ethernet or, in the case of wireless connectivity, the modulation implementation.




Figure 2. Simplified OSI model (left) and an example of a TCP/IP protocol stack (right).

The link layer converts bits to radio signals (and vice versa), takes care of data framing for reliable wireless communication, and manages access to the radio channel. In the TCP/IP example in **Figure 2**, Wi-Fi is the link layer protocol.

The link layer is often broken down into two sublayers for wireless protocols: the logical link control (LLC) layer and the media access control (MAC) layer. Some wireless protocols, such as 802.15.4, only standardize the MAC layer.

The network layer addresses and routes data through the network. IP is the internet's network layer protocol, providing an IP address to devices

and carrying IP packets from one device to another. The transport layer generates communication sessions between applications running on two ends of a network. It allows multiple applications to run on one device, each using its own communication channel. TCP is the internet's predominant transport protocol.

The application layer is responsible for data formatting and governs the data flow in an optimal scheme for specific applications. A popular application layer protocol in the TCP/IP stack is Hypertext Transfer Protocol (HTTP), which was created to transfer web content over the internet.

A layered network implementation introduces complexity and requires more code and memory. It also introduces data overhead because every layer requires additional framing and control messages. On the other hand, layered networks enable more flexibility and scale – so much so that most wireless networks today use a scheme similar to the one shown in **Figure 2**. One example of a simple network design, with no or little layering, is a proprietary application protocol running directly over the physical layer provided by a simple radio transceiver. A proprietary implementation like this can be very efficient, but it is mainly used in simple, single-function networks.

To IP or not to IP

The Internet of Things (IoT) is all about connecting things to the internet. Devices that (directly) connect to the internet must use the IP suite in order to exchange data with devices and servers. Nevertheless, devices in a local network can use non-IP protocols to communicate within the local network.

Non-IP devices connect to the internet through an internet gateway. The gateway communicates with local devices using a non-IP method on one

side, and with other devices on the internet using IP on the other side. The gateway is an application layer gateway, because it needs to strip down the data coming in from the local network and restructure it with a TCP/IP stack in order to enable communication with an internet service.

The benefit of using the IP in IoT node devices is that it allows the use of a network (IP)-layer gateway. A common example of an IP-layer gateway would be a Wi-Fi access point or router that connects Wi-Fi devices to the internet. An IP-layer gateway is more flexible than an application layer gateway because you can modify or add applications to devices without changing anything on the gateway.

For example, the reason you can install a new internet application on your laptop without changing anything in your home network is because your laptop runs a TCP/IP stack, and home routers (using wired Ethernet or Wi-Fi) are only manipulating data at the link layer.

In many cases, a local network is for one specific application. For example, alarm systems with wireless sensors form a local network with one function – sending sensor activation messages to the alarm controller. If the alarm controller is relaying messages to your phone over the internet, then it uses TCP/IP. But the communication between the sensors and the alarm controller does not have to be based on TCP/IP. In fact, in most cases today, it is not.

Other often-overlooked advantages of the TCP/IP stack come from the unprecedented widespread use of internet applications. The number of application protocol implementations created for TCP/IP over the last 20 years is overwhelming. You can use many protocols in a local network, even if internet connectivity is not required. Reusing existing, proven protocol implementations can dramatically shorten development time. Diagnostics,

management and commissioning tools for IP networks already exist and can help shorten development cycles as well.

One of the disadvantages of the TCP/IP stack is that it is fairly complex and large, and therefore requires more processing power and memory, implying more expensive devices. This complexity also results in larger data packets that consume more power during send and receive. For these reasons, many simple networks (such as the wireless sensors in an alarm system) choose to implement simpler and usually proprietary protocols.

As silicon technology advances, processing power and memory become more available and affordable, making TCP/IP communication more attractive even for small and simple networks. With the integration of a full TCP/IP stack into products like the Texas Instruments (TI) SimpleLink™ family, expect more and more applications to move from proprietary protocols to IP-based protocols, enabling flexible internet connectivity and faster development cycles.

Network range

As **Figure 3** shows, a network's range is typically categorized into four classes: personal area networks (PANs), local area networks (LANs), neighborhood area networks (NANs) and wide area networks (WANs).




Figure 3. Different ranges and applications for PANs, LANs, NANs and WANs.

PANs are usually wireless and cover a range of about 10m. A common wireless PAN is a smartphone connected over Bluetooth® to a handful of accessories such as a wireless headset, watch or fitness device. Wireless PAN devices usually have low radio transmission power and run on small batteries.

LANs are either wired or wireless (or a combination of the two). Wireless LANs (WLANs) usually cover a range up to 100m. A predominant example is a home Wi-Fi network providing internet access to personal computers, smartphones, TVs, and IoT devices like thermostats and home appliances.

NANs are usually wireless and can reach more than 25km. They transmit high power levels, but relay relatively low data traffic. An example NAN is a smart grid network that transmits electric meter readings from homes to utility companies using a proprietary protocol over a 900MHz radio.

Finally, WANs are spread across a very large area – as big as the entire globe. The internet is considered a WAN and comprises a complex mix of wired and wireless connections.

Network topology and size

Wireless networks can also be categorized by their topology – the way nodes in a network are arranged and how they connect to each other. The two fundamental network topologies are star and mesh, as depicted in **Figure 4**. In a star topology, all of the nodes connect to one central node, which is typically also used as the gateway to the internet.

A popular example of a star topology is a Wi-Fi network, where the center node is called an access point (AP) and the other nodes are called stations.



Figure 4. Mesh topology (left) and star topology (right).

In a mesh network, every node can connect to multiple other nodes. One or more nodes in the network serve as an internet gateway. In Figure 4, every node in the network connects to every other node.

In real life, a mesh topology is simpler. A popular example of a mesh network is a Zigbee® Light Link network, where multiple lights form a mesh network to extend the reach in large buildings. One of the Zigbee nodes is called a coordinator, and it usually also serves as an internet gateway.

The benefit of a mesh topology is that it can extend a network range through multiple hops while maintaining low radio transmission power. A mesh topology can also achieve better reliability by enabling more than one path to relay a message through the network. However, routing nodes in a mesh network often requires additional processing power and memory to support the routing function, making mesh networks more complex and costly to design. Plus, mesh networks can exhibit longer delays when routing a message from a remote node through the mesh compared to star networks.

Network size, or the maximum number of simultaneously connected devices, is also an important consideration in system design. Some technologies, like Bluetooth, support up to 20 connections; other technologies, like Zigbee, can support thousands of connections.

Standards and interoperability

One of the biggest challenges in communication systems is interoperability – meaning the ability of devices from different vendors to exchange data. Addressing this challenge is the main goal of the many standards organizations that define specifications and testing procedures. Recalling the OSI network model in Figure 2, some standards define one or several network layers, while others define entire end-to-end network specifications.

The Institute of Electrical and Electronics Engineers (IEEE) is a nonprofit organization formed in 1963. It has a considerable focus on communication and radio engineering, and one of its well-known contributions to networking technology is the IEEE 802.x family of standards. For example, 802.3 defines the Ethernet specification, which governs most wired computer networks today. 802.11 defines the WLAN specification, which is the baseline of the Wi-Fi standard, and 802.15.4 defines the wireless PAN standard used in Zigbee, 6LoWPAN and Wireless Highway Addressable Transducer Protocol (HART). 802.x standards define only the link layer of the network.

The Internet Engineering Task Force (IETF) is an open standards organization formed in 1986 that is responsible for developing internet standards, particularly the TCP/IP suite. IETF specifications are established through the publication of draft specifications under the title “request for comments” (RFCs). After an RFC, IETF members conduct multiple reviews and make edits until the specification is approved to “best current practice” status. RFCs define thousands of internet standards. A few examples are RFC 791, which describes the IPv4 protocol; RFC 793, which describes the TCP protocol; and RFC 2616, which defines the HTTP/1.1 protocol.

The IETF, like the IEEE, does not run certification programs. In other words, vendors cannot get recognition from either the IETF or IEEE that their products comply with any of their standards. The IEEE and IETF standards only set good practices, so other organizations adopt IETF and IEEE standards and use them to create certification programs. These organizations often adopt just a portion of an IEEE or an IETF standard for various reasons.

The Wi-Fi Alliance, Bluetooth Special Interest Group (SIG), Thread Group and Zigbee Alliance manage certification programs to assure interoperability between wirelessly connected devices. All of these organizations give member companies the option to take products through an interoperability test plan, which grants rights to use the Wi-Fi, Bluetooth, Thread or Zigbee logo.

So far, we have reviewed some of the key concepts in wireless connectivity and discussed engineering trade-offs in wireless connectivity system design. Now, let's cover the predominant wireless connectivity technologies in the industry and their applications in more detail.



Wi-Fi technology, based on the IEEE 802.11 standard, was developed as a wireless replacement for the popular wired IEEE 802.3 Ethernet standard. As such, it was created from day one for internet connectivity.

Although Wi-Fi technology primarily defines the link layer of a local network, it is so natively integrated with the TCP/IP stack that when people say they are using Wi-Fi, they implicitly mean that they are also using a TCP/IP for internet connectivity.

Riding on the huge success of smartphones and tablets, Wi-Fi has become so ubiquitous that people

often refer to it as just "wireless." Wi-Fi APs are deployed today in most homes, as well as in almost all offices, schools, airports, coffee shops and retail stores. The huge success of Wi-Fi is largely due to the remarkable interoperability programs run by the Wi-Fi Alliance and to the increasing demand for easy and cost-effective internet access.

Wi-Fi is already integrated into all new laptops, tablets, smartphones and TVs. Taking advantage of its existing ubiquity, Wi-Fi is also widely used in IoT applications that can leverage installed Wi-Fi infrastructures without custom gateways.

Most Wi-Fi networks operate in the 2.4GHz ISM band. Wi-Fi can also operate in the 5GHz band, where more channels exist and higher data rates are available. Wi-Fi networks have a star topology, with the AP as the internet gateway. The output power of Wi-Fi is high enough to enable full in-home coverage in most cases. In large buildings, more than one AP is often deployed in different locations to increase network coverage, since multipath conditions may cause dead spots. To overcome dead signal reception spots, some Wi-Fi products include two antennas for diversity.

Wi-Fi and TCP/IP software are fairly large and complex. For laptops and smartphones with powerful microprocessors (MPUs) and large amounts of memory, this imposes no issue. Until recently, adding Wi-Fi connectivity to devices with little processing power such as thermostats and home appliances was not possible nor cost-effective. Today, silicon devices and modules embed the Wi-Fi software and TCP/IP software inside the device. These new devices eliminate most of the overhead from the MPU and enable wireless internet connectivity with the smallest microcontroller (MCU). The increasing level of integration in these Wi-Fi devices also eliminates the need to have radio design experience

and reduces the barriers of Wi-Fi integration.

The latest Wi-Fi silicon devices apply advanced sleep protocols and fast on/off times to reduce the average power consumption dramatically. Since most IoT products do not need the maximum data rates that Wi-Fi offers, clever power-management design can efficiently draw bursts of current from the battery for very short intervals and keep products connected to the internet for over a year using two AA alkaline batteries. Today, you can buy a Wi-Fi-based sports watch that uploads workout data to the internet.

Most Wi-Fi APs claim support for as many as 250 simultaneously connected devices. Enterprise-grade APs can support even more connections, while other popular consumer APs handle no more than 50.

Wi-Fi is the most ubiquitous wireless Internet connectivity technology today. Its high power and complexity has been a major barrier for IoT developers, but new silicon devices and modules reduce many of these barriers and enable Wi-Fi integration into emerging IoT applications and battery-operated devices.



Ericsson invented Bluetooth technology in 1994 as a standard for wireless communication between phones and computers. The Bluetooth link layer, operating in the 2.4GHz ISM band, was previously standardized as IEEE 802.15.1, but that IEEE standard is no longer maintained. The Bluetooth SIG controls the Bluetooth standard.

The original Bluetooth is today commonly referred to as Classic Bluetooth, to distinguish it from Bluetooth low energy. Classic Bluetooth became very successful in mobile phones – so much so that all mobile phones today, even entry-level phones, have Bluetooth connectivity. The main use case that made Classic Bluetooth initially popular was hands-free phone calls with headsets and car kits. As mobile

phones became more full-featured, more use cases like high-fidelity music streaming and health and fitness accessories evolved.

Classic Bluetooth is a PAN technology primarily used today as a cable replacement for short-range communication and is primarily a point-to-point or star network topology. It uses the 2.4GHz ISM band supports data throughput up to 2Mbps, with up to eight connected devices. If a Wi-Fi connection is available, as it typically the case in a mobile device, Classic Bluetooth can operate over Wi-Fi, known as Bluetooth High Speed, and provide transfer rates of up to 25Mbps. The technology is fairly low power; devices typically use small rechargeable batteries or two alkaline batteries.

Bluetooth low energy, previously known as Bluetooth Smart, is a more recent addition to the Bluetooth specification. Nokia originally created Bluetooth low energy, and the Bluetooth SIG included it in the Bluetooth 4.0 standard in 2010 to enter the low-power IoT space.

Bluetooth low energy also uses the 2.4GHz ISM band, but is not compatible with Classic Bluetooth. Bluetooth low energy uses 40 2MHz-wide channels, whereas Classic Bluetooth uses 79 1MHz-wide channels. Compared to Classic Bluetooth, Bluetooth low energy greatly reduces the power consumption of Bluetooth devices (at the cost of lower data throughput), and enables lengthy operation using coin-cell batteries. Bluetooth low energy -also offers a beaconing capability and location-based services. Enjoying the same near universal mobile device support as Classic Bluetooth, Bluetooth low energy has proved very popular, triggering an explosion of new applications in spaces as diverse as fitness, toys and automotive. It is now the main driving force behind new Bluetooth standards.

One of the strengths of the Classic Bluetooth standard is that it includes application profiles,

such as for audio/video remote control and heart-rate monitors. These profiles define in detail how applications exchange information to achieve specific tasks and include comprehensive Bluetooth SIG-defined certification programs to ensure excellent interoperability in the market. Bluetooth low energy also offers generic attribute (GATT) profiles. GATT provides a structured list that defines the services, characteristics and attributes of a given node. A key benefit of GATT is that it provides a standard mechanism that can implement new profiles, enabling Bluetooth low energy to quickly adapt to new applications. Classic Bluetooth profiles are now advertised as “adopted in the GATT” and are available to Bluetooth low energy.

Bluetooth 4.1, introduced in 2013, further improved power consumption by allowing devices to specify custom reconnection timeout intervals rather than a fixed timeout period, enabling devices to spend much longer in ultra-low-power states. Devices could act as both hub and end point simultaneously so that peripherals could communicate without a host. A phone-specific enhancement of Bluetooth 4.1 improved coordination of Bluetooth and 4G radios so that both could simultaneously operate at maximum performance.

The Bluetooth 4.2 specification, introduced in 2014, included higher performance; greater range, privacy and security enhancements; and simpler internet connectivity. By increasing the packet capacity from 27 bytes to 251 bytes, data throughput reached up to 800Kbps. Bluetooth low energy 4.2 added industrial-strength security with elliptic curve cryptography (ECC)-based key management and Advanced Encryption Standard (AES)-counter with cipher block chaining message authentication code (CCM) cryptography for message encryption. To facilitate connecting Bluetooth low energy devices to the internet, Bluetooth low energy 4.2 added IP Support Profile and Bluetooth low

energy internet gateways via GATT. Smart internet gateways simplified the development of gateways to IPv4 networks, which constitute the majority of today’s internet. The IP Support Profile added IPv6 connectivity to Bluetooth low energy devices.

Compared to Bluetooth 4.2, Bluetooth 5, the current standard introduced in 2016, quadruples the range, doubles the speed (up to 2Mbps) and provides an eightfold increase in data broadcasting capacity by increasing the advertising data length. Bluetooth 5 offers a choice of data rates – 2Mbps, 1Mbps, 500Kbps, and 125Kbps – with the lower data rates offering longer ranges. The increases in range and data rates make Bluetooth low energy increasingly attractive in nontraditional segments such as industrial data loggers or smart energy meters, given Bluetooth low energy’s inherent advantage of built-in compatibility with mobile devices. Its mobile device connectivity offers an excellent avenue for data display and retrieval, internet connectivity, and initial provisioning and configuration.

In 2017, the SIG released the mesh profile and mesh model specifications. The mesh specification enables the use of Bluetooth low energy for many-to-many device communications in home automation, sensor networks and other applications, while enabling longer-range communication using intermediary nodes. These new mesh standards are compatible with both the Bluetooth 5 and Bluetooth 4.x standards.



ZigBee®

Zigbee is an open standard designed, promoted and maintained by the Zigbee Alliance, a group of more than 400 companies from different industries. Zigbee is designed to provide a whole connectivity solution for device interoperability and cloud connectivity.

Zigbee is based on the IEEE 802.15.4 link layer and operates in the 2.4GHz ISM band. Its networking layer has been designed with mesh-topology operations in mind from the ground up; providing the ability to scale the network geographically through multihop operations, as well as fault tolerance and increased reliability as backup paths are created through the mesh.

One of Zigbee's major benefits is that it provides a complete solution that enables true device interoperability between different manufacturers. In fact, the Zigbee protocol suite incorporates the Zigbee cluster library: a standard library of device types, data models and behaviors built by original equipment manufacturers (OEMs) operating in different domains and proven in actual deployments for many years.

The Zigbee protocol suite also includes standard commissioning, security, network and device-management procedures. Various device types can join and be authenticated in the network, and be factory-reset or decommissioned in an interoperable way, guaranteeing end-to-end device interoperability from the start of device operation and seamlessly integrating with data collectors or hubs. The broad availability of Zigbee solutions is a further advantage for OEMs that own entire ecosystems and want to sell services to their customers that leverage existing products and solutions.

Zigbee-based applications mostly target the smart home and smart building domains, with focused segments in lighting and home control and security. Service providers like Comcast, Deutsche Telekom and Alibaba endorsed Zigbee as the protocol of choice when introducing their home automation services to consumers, and lighting manufacturers like Philips and Osram have a whole line of ZigBee-based wireless (LED) products.

Zigbee effectively uses the 250kbps bandwidth to convey both application data to operate devices

(such as a switch toggling a light) and network management procedures like mesh and routing management with a very small energy footprint.

Devices using Zigbee are typically low data rate emitting sensors and actuators and – with the low-power performance of the underlying 802.15.4 link layer radio and the efficiency of a whole protocol suite designed to transport application and management data in short frames – vastly minimizing air time and therefore power consumption. Zigbee also supports the Green Power feature set, an extension to the Zigbee specification for even more energy-constrained devices. Thanks to its Green Power addition, battery-less devices scavenging the environment for power can operate in a ZigBee network and interoperate with other devices.

A rigorous certification program managed by the Zigbee Alliance guarantees interoperability between Zigbee devices, verifying device-type behavior and functionality from an end-product perspective and ensuring that products from different manufacturers can operate together.



6LoWPAN is an acronym for IPv6 over Low Power Wireless Personal Area Networks. 6LoWPAN enables very small, low-power devices to directly communicate with any other IP-based (Ethernet or Wi-Fi) server or device on the internet. The standard was formalized under RFC 6282 in September 2011. As indicated by the RFC title, the 6LoWPAN standard only defines an efficient adaptation layer between the 802.15.4 link layer and a TCP/IP stack. Since the 802.15.4 link layer has multiple optional modes, various 6LoWPAN solutions from different vendors are not interoperable at the local network level. These interoperability issues have slowed 6LoWPAN's adoption and were a key driver behind the Thread standardization effort (see the next section).

Like Wi-Fi, 6LoWPAN networks require an IP-layer gateway to access the internet. Since 6LoWPAN only supports IPv6 and most of the deployed internet today is still using IPv4, a 6LoWPAN gateway typically includes an IPv6-to-IPv4 conversion protocol that gives 6LoWPAN nodes and applications direct access to the internet.

6LoWPAN implementations exist for both the 2.4GHz and 868MHz/915MHz ISM bands. 6LoWPAN offers the advantages of 802.15.4 – a mesh network topology, large network size, reliable communication and low power consumption – combined with the benefits of IP communication. A popular, proven, open-source 6LoWPAN implementation is the Contiki open-source operating system (OS), originally created by Thingsquare.



Google, Samsung and a number of silicon vendors who wanted to directly connect low-power IoT nodes to IP networks formed the Thread Group to define a standard that ensures device interoperability in the smart home space. Most recently, Thread announced the extension of its standard to commercial building applications. Thread appears to have established itself as the future solution for IP over 802.15.4 compared to Zigbee IP, which has attracted very limited interest.

Thread is based on the IEEE 802.15.4 2006 specification and operates in the 2.4GHz ISM band, similar to Wi-Fi, Bluetooth low energy and Zigbee. It provides data rates of up to 250kbps and can support as many as 250 devices in one local network mesh. Thread inherits the advantages of 802.15.4 and 6LoWPAN such as low-power operation, IPv6 addressability and mesh networking. To maintain low power and cost advantages, Thread specifies lower-overhead higher-level protocols such as User Datagram Protocol (UDP) and Datagram Transport

Layer Security (DTLS). Thread builds on this base feature set to add superior security, reliability, scalability and ease of commissioning.

Thread provides a secure commissioning process through easy-to-use mobile applications. Each Thread device has a unique quick response (QR) code that ensures that only authorized devices can join the network. When a customer adds a new device to their network, they scan the QR code and then enter the Thread network passcode in the mobile application. The commissioning node then establishes a secure DTLS connection to send an AES-128 key to the device. This key enables the device to join the network and is also used for subsequent communication.

Thread achieves high reliability through a self-healing mesh implementation that avoids a single point of failure. Thread defines several device types: border routers (gateways), routers, router-eligible end devices (REEDs) and sleep end devices. Border routers are Thread devices in the network with one additional IP networking interface besides 802.15.4 (such as Wi-Fi, Ethernet or cellular) to provide a path from the Thread network to the internet. Thread supports multiple border routers in a Thread network to increase reliability. REEDs are essentially back-up routers that a Thread network can transparently upgrade to routers if a router fails. Similarly, the role of a leader in a Thread network can be reassigned to another router if the original leader becomes unavailable. This self-healing mesh capability is based on Mesh Link Establishment (MLE), a protocol designed as part of the Thread implementation. MLE messages distribute or exchange the addresses of neighboring devices, routing information, updates to network data, security material and device information such as sleep cycles, enabling a network to dynamically reconfigure itself in response to any failures.

Unlike traditional 6LoWPAN solutions that typically use higher-level protocols like the IPv6 Routing Protocol for Low-Power and Lossy Networks (RPL) for routing, Thread routing is optimized for device-to-device communication. This results in lower latency and therefore enables a network to support more devices.

A Thread Group-managed certification program enables interoperability between Thread devices and guarantees that end products coming from different OEMs can interoperate together. Many product manufacturers use Google's OpenThread, an open-source implementation of the Thread networking protocol. OpenThread supports several silicon platforms, including wireless MCU devices from TI.

WirelessHART

Wireless HART and Wireless IO Link both operate in the 2.4GHz ISM band and are designed for industrial applications. In particular, they are designed to operate in environments with high levels of electromagnetic interference. Both protocols require gateways to connect to IP networks.

IO-Link (International Electrotechnical Commission [IEC] 61131-9) provides factory automation systems with a standard wired mechanism for real-time communication to sensors and actuators. An IO-Link sensor provides significantly more information, configurability and control than a traditional analog or digital sensor.

Wireless IO Link is a wireless communication option for IO-Link sensors. The communications range is short (20m maximum), with an emphasis on very low latency, since sensors and actuators require fast responses. Wireless IO Link implements radio-channel blacklisting to ensure that it can offer reliable real-time communication while coexisting with already installed wireless systems such as Wi-Fi networks.

The HART communication protocol is well established in the process automation space, where it uses wired analog two-way communication between a host application and smart field instruments, providing access to diagnostics, configuration and process data. Wireless HART allows full wireless connectivity to HART-enabled devices, enabling real-time communication of device diagnostics and process data to asset-management systems. Wireless HART is based on the 802.15.4 standard and offers both star and mesh topologies. It is deterministic with a range of several tens of meters and is intended for low-frequency information updates.

Sub-1GHz

While the 2.4GHz band is the most well-known ISM band, almost all countries make lower ISM frequency bands available, such as at 433MHz, 868MHz and 915MHz. These bands are commonly known as Sub-1 GHz. Many industrial applications today use proprietary protocols running over radio transceivers operating in the Sub-1 GHz band. The radio transceiver provides the link layer of the network (or often just the physical layer [PHY]). The OEM implements the rest of the network protocol.

Many utility companies have created proprietary NANs to relay meter readings to a neighborhood collection point. Other popular applications for Sub-1 GHz radios are security systems and industrial control and monitoring. To connect to the IoT, Sub-1 GHz systems need an application-layer internet gateway.

Sub-1 GHz solutions can offer very low power and long ranges with a simple star topology. Note that there is a direct correlation between range, data rates and power consumption. The longer the transmission range, the lower the data transmission rate must be.

Low data transmission rates in turn lead to increased power consumption, as the device must be active longer to transmit data.

Despite demonstrated transmission ranges of up to 100km, a maximum range of around 25km is more typical in real-world deployments, even with very low data rates. Particular implementations requiring higher data rates or lower power may offer a range of only a few kilometers.

Although the long range and low-power attributes of the Sub-1 GHz band are very attractive for IoT applications, widespread adoption of Sub-1 GHz in the IoT space has been hindered by the need for OEMs to have radio frequency (RF) protocol expertise to deploy an application. Although factors such as an extensive range of PHYs (rather than a single one, as in Wi-Fi and Bluetooth low energy) preclude the standardization of protocol stacks, a number of off-the-shelf solutions do eliminate the need for in-house RF protocol experience. We'll discuss commercial network solutions such as Sigfox and 802.15.4 stacks from Sub-1 GHz silicon providers below. In addition, the Contiki OS solution mentioned in the 6LoWPAN section supports Sub-1 GHz.

802.15.4g-based protocols

Although the 802.15.4 standard was originally based on the 2.4GHz band, in 2011, the g extension expanded its operation to the Sub-1 GHz spectrum. This enabled Sub-1 GHz radio silicon vendors to provide 802.15.4-based protocol stacks that enable customers with no RF protocol design experience to connect Sub-1 GHz nodes to the internet. Similar to 6LoWPAN, these stacks are not interoperable, as 802.15.4 lacks standard link-layer implementation.

The specifics of each 802.15.4g solution will vary between silicon providers, an overview of TI's implementation summarizes the key points. The TI-15.4 stack operates as star network topology and

provides a line-of-sight range (depending on data rate) of up to several kilometers between the node and the gateway. Supported data rates are 5kbps, 50kbps and 500kbps, with exceptionally low-power performance allowing several years of operation on a coin-cell battery.

The TI 15.4-Stack includes a proprietary link layer to perform duties such as network management that are necessary but not defined in the 802.15.4 standard. A standard 802.15.4 AES-CCM provides security, enhanced by message integrity check and passphrase. Wireless Smart Ubiquitous Network (Wi-SUN)-based frequency hopping provides more robust transmission.

Since a 15.4 network requires a gateway to connect to the internet, a reference implementation of a gateway must take the 15.4 packets (typically using a virtual serial interface for the radio) and make them available through TCP/IP-based network applications, from which the data is easily sent to the cloud. It must also provide some network formation and management functions that interact with the link-layer protocol on the nodes.



Sigfox is a network provider, conceptually similar to a mobile operator such as T-Mobile or Verizon. The Sigfox network is aimed at IoT applications such as remote sensing that transmit small amounts of data relatively infrequently. Compared to cellular, Sigfox offers significant cost and battery-life advantages for such applications.

Sigfox's network technology uses ultra-narrowband techniques in the Sub-1 GHz spectrum (specifically the 915MHz ISM band in the U.S. and the 868MHz band in Europe) to provide long-range connectivity. The Sigfox network provides very long-range

connectivity (up to 30km in rural areas) between the base station and the node using a star network topology. To achieve such a long range, Sigfox must transmit at very low data rates (up to 100bps using 12 byte packets, and no more than 140 messages per node per day).

Although the Sigfox network is bidirectional, the payload for transmissions from the base station to the nodes is very low, which limits the range of usable systems. However, this approach is sufficient for a number of IoT applications, such as a connected water meter or smoke detector.

Sigfox was originally launched in Europe and has been slower to deploy in the U.S. market. FCC regulations limit the maximum time a transmission can be on the air to 0.4s. Since Sigfox data rates required a transmission time of around 3s, meeting the FCC regulation required some re-architecture.

Conclusion

There are many wireless technologies in the world. Each one has benefits, and no one is perfect.

The question that you need to answer is, “Which technology, or combination of technologies, is the best one for my application?” We hope that this paper has helped you better understand the popular wireless technologies for IoT, and their strengths and weaknesses.

Table 1 maps connectivity technologies to range, throughput, power consumption, network topology and desktop/mobile device compatibility. Use these data points as only one part of your decision process, since there are many overlapping areas.

Additional considerations that go beyond the scope of this paper include cost, ease of integration and security. We see great improvement on total solution cost and ease of integration in many new products that use wireless connectivity, but you should consider cost and integration efforts further in the context of specific applications.

Security aspects of IoT applications include the supported capabilities of each of the protocols, as well as additional hardware and software considerations, also beyond the scope of this paper.

	Wi-Fi*	BLE/ Bluetooth 5	Thread	Sub-1 GHz: TI 15.4	Sub-1GHz: Sigfox	Zigbee
Data throughput	Up to 72Mbps	Up to 2Mbps	Up to 250kbps	Up to 200kbps	100bps	Up to 250kbps
Range**	100m	Up to 750m	100m via mesh	4km	25km	130m LOS
Power consumption	Up to 1 year on AA batteries	Up to years on a coin-cell battery	Up to years on a coin-cell battery	Up to years on a coin-cell battery for 1km range	Up to years on a coin-cell battery for limited range	Years on a coin-cell battery
Topology	Star	Point-to-point/Mesh	Mesh & Star	Star	Star	Mesh & Star
IP at the device node	Yes	No	Yes	No	No	No
PC, mobile OS support	Yes	Yes	No	No	No	No
Infrastructure widely deployed	Yes, Access Points	Yes, smart phones	No	No	No	No

*Single stream 802.11n Wi-Fi MCUs may support lower throughput than peak physical capacity of the network.

**LOS = Line Of Sight. For range, note that maximum data rates are often not available at the longest range.

Table 1. Some of the key considerations that will influence the choice of wireless protocols for a specific application, such as data rate, range and power.

TI's SimpleLink wireless connectivity portfolio

TI's SimpleLink MCU platform of low-power wireless and wired MCUs for the broad embedded market – makes it easier to develop and connect anything to the IoT. Spanning more than 10 standards and technologies including Wi-Fi, Bluetooth low energy, Zigbee, Sub-1 GHz, 6LoWPAN, Thread, Ethernet and more, SimpleLink products help manufacturers add wireless connectivity to any design, for anyone.

Texas Instruments and the IoT

With the industry's broadest IoT-ready portfolio of wired and wireless connectivity technologies, MCUs, processors, sensors and analog signal-chain and power solutions, TI offers cloud-ready system solutions designed for IoT accessibility. From high-performance home, industrial and automotive applications to battery-powered wearable and portable electronics or energy-harvested wireless sensor nodes, TI makes developing applications easier with hardware, software, tools and support to get anything connected within the IoT.

Important Notice: The products and services of Texas Instruments Incorporated and its subsidiaries described herein are sold subject to TI's standard terms and conditions of sale. Customers are advised to obtain the most current and complete information about TI products and services before placing orders. TI assumes no liability for applications assistance, customer's applications or product designs, software performance, or infringement of patents. The publication of information regarding any other company's products or services does not constitute TI's approval, warranty or endorsement thereof.

A011617

The platform bar is a trademark of Texas Instruments.
All other trademarks are the property of their respective owners.

IMPORTANT NOTICE FOR TI DESIGN INFORMATION AND RESOURCES

Texas Instruments Incorporated ("TI") technical, application or other design advice, services or information, including, but not limited to, reference designs and materials relating to evaluation modules, (collectively, "TI Resources") are intended to assist designers who are developing applications that incorporate TI products; by downloading, accessing or using any particular TI Resource in any way, you (individually or, if you are acting on behalf of a company, your company) agree to use it solely for this purpose and subject to the terms of this Notice.

TI's provision of TI Resources does not expand or otherwise alter TI's applicable published warranties or warranty disclaimers for TI products, and no additional obligations or liabilities arise from TI providing such TI Resources. TI reserves the right to make corrections, enhancements, improvements and other changes to its TI Resources.

You understand and agree that you remain responsible for using your independent analysis, evaluation and judgment in designing your applications and that you have full and exclusive responsibility to assure the safety of your applications and compliance of your applications (and of all TI products used in or for your applications) with all applicable regulations, laws and other applicable requirements. You represent that, with respect to your applications, you have all the necessary expertise to create and implement safeguards that (1) anticipate dangerous consequences of failures, (2) monitor failures and their consequences, and (3) lessen the likelihood of failures that might cause harm and take appropriate actions. You agree that prior to using or distributing any applications that include TI products, you will thoroughly test such applications and the functionality of such TI products as used in such applications. TI has not conducted any testing other than that specifically described in the published documentation for a particular TI Resource.

You are authorized to use, copy and modify any individual TI Resource only in connection with the development of applications that include the TI product(s) identified in such TI Resource. NO OTHER LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE TO ANY OTHER TI INTELLECTUAL PROPERTY RIGHT, AND NO LICENSE TO ANY TECHNOLOGY OR INTELLECTUAL PROPERTY RIGHT OF TI OR ANY THIRD PARTY IS GRANTED HEREIN, including but not limited to any patent right, copyright, mask work right, or other intellectual property right relating to any combination, machine, or process in which TI products or services are used. Information regarding or referencing third-party products or services does not constitute a license to use such products or services, or a warranty or endorsement thereof. Use of TI Resources may require a license from a third party under the patents or other intellectual property of the third party, or a license from TI under the patents or other intellectual property of TI.

TI RESOURCES ARE PROVIDED "AS IS" AND WITH ALL FAULTS. TI DISCLAIMS ALL OTHER WARRANTIES OR REPRESENTATIONS, EXPRESS OR IMPLIED, REGARDING TI RESOURCES OR USE THEREOF, INCLUDING BUT NOT LIMITED TO ACCURACY OR COMPLETENESS, TITLE, ANY EPIDEMIC FAILURE WARRANTY AND ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT OF ANY THIRD PARTY INTELLECTUAL PROPERTY RIGHTS.

TI SHALL NOT BE LIABLE FOR AND SHALL NOT DEFEND OR INDEMNIFY YOU AGAINST ANY CLAIM, INCLUDING BUT NOT LIMITED TO ANY INFRINGEMENT CLAIM THAT RELATES TO OR IS BASED ON ANY COMBINATION OF PRODUCTS EVEN IF DESCRIBED IN TI RESOURCES OR OTHERWISE. IN NO EVENT SHALL TI BE LIABLE FOR ANY ACTUAL, DIRECT, SPECIAL, COLLATERAL, INDIRECT, PUNITIVE, INCIDENTAL, CONSEQUENTIAL OR EXEMPLARY DAMAGES IN CONNECTION WITH OR ARISING OUT OF TI RESOURCES OR USE THEREOF, AND REGARDLESS OF WHETHER TI HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

You agree to fully indemnify TI and its representatives against any damages, costs, losses, and/or liabilities arising out of your non-compliance with the terms and provisions of this Notice.

This Notice applies to TI Resources. Additional terms apply to the use and purchase of certain types of materials, TI products and services. These include; without limitation, TI's standard terms for semiconductor products (<http://www.ti.com/sc/docs/stdterms.htm>), [evaluation modules](#), and [samples](http://www.ti.com/sc/docs/sampterm.htm) (<http://www.ti.com/sc/docs/sampterm.htm>).

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265
Copyright © 2017, Texas Instruments Incorporated