Low-Power Wide-Area Networks 101

SCT
SIGNAL CRAFT TECHNOLOGIES

Introduction

Low-power wide-area networks (LPWAN) will play a crucial role in the impending advance of the Internet of Things (IoT). These wireless networks will become the primary backbone of IoT and a cost-effective way to connect a huge number of remote devices spread over vast areas, all while using power so efficiently that batteries in remote devices will be capable of lasting years—even decades.

LPWANs' data rate, payload size, and latency have been intentionally restricted in favor of long operating range and ultra-low power consumption, both of which far exceed what can be achieved with more traditional wireless technologies.

This is possible thanks to radically new physical layer designs, aimed at achieving very high receiver sensitivity. For example, while the nominal sensitivity of ZigBee and Bluetooth receivers is about -125 dBm and -90 dBm, respectively, some LPWAN receivers achieve sensitivities as low as -142dBm.

In this paper, we will discuss some of the key characteristics that should be considered to help select the appropriate LPWAN technology for a given IoT application.

Key System Attributes

To fully appreciate the benefits LPWAN can offer, it's important to gain an understanding of what attributes set LPWANs apart from other wireless technologies. Each LPWAN technology employs a variety of techniques to achieve the distinguishing features of any LPWAN: long operating range, very low power, massive unit capacity, advanced security, and low cost of operation. Some of the more sophisticated LPWAN technologies use a multitude of techniques to achieve these goals, while the simpler LPWAN technologies may only apply a handful.

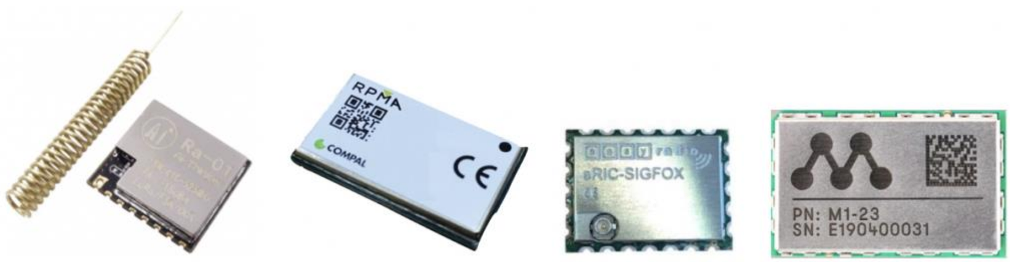

Examples of Off-The-Shelf LPWAN Radio Modules

LPWAN technologies are quickly gaining popularity because of their ability to provide very long coverage ranges using relatively low powered transmitters.

Long Operating Range

LPWAN technologies are quickly gaining popularity because of their ability to provide very long coverage ranges using relatively low powered transmitters. Real-world range performance in the order of 10 to 15 kilometers in rural areas and two to five kilometers in urban areas are absolutely achievable with battery operated nodes transmitting at 100mW or less. LPWAN system architectures use a multitude of techniques to achieve such impressive operating ranges.

Broadcast Frequency

The majority of LPWAN solutions have chosen to operate at frequency bands below 1 GHz. The Friis equation tells us that signals at a lower frequency will propagate further than at a higher frequency. It is also known that a lower frequency will penetrate more readily into structures and experience less multipath fading.

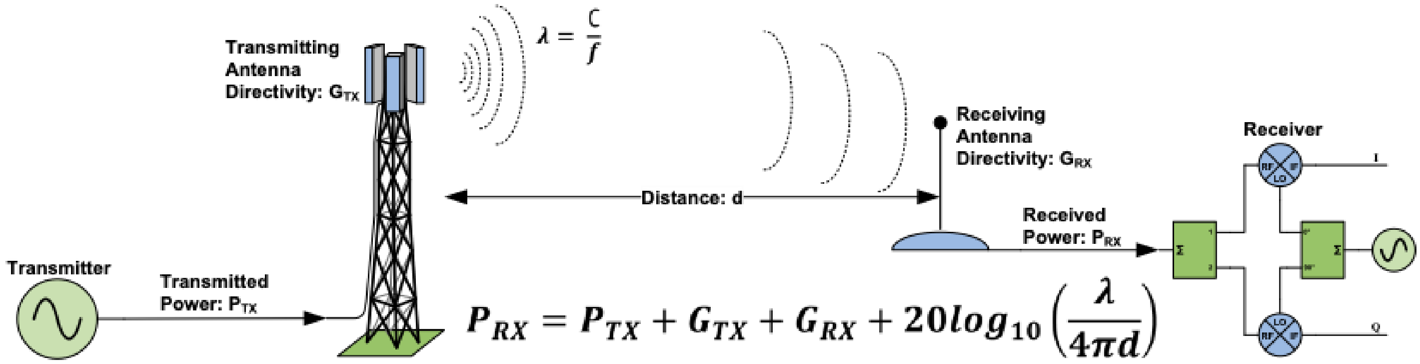

Friis Path Loss Equation

Low Data Rates

A receiver's ability to reliably decode messages is related to the ratio of signal power to noise power at the demodulator input and in turn, received noise power is proportional to the bandwidth of the channel that the message is being passed through. Narrowing a channel reduces the amount of noise entering the demodulator and correspondingly less signal power is required to correctly decode a message. Being able to decode weaker signals means that a receiver can be placed at greater distances from the transmitter.

The trade-off of using narrow channels is that they have very limited capacity for information. Therefore, most LPWAN technologies transmit messages at extremely low data rates relative to more traditional wireless technologies. Ultra-Narrowband (UNB) and Spread Spectrum are the LPWAN industry's most popular means of transmitting low data rate messages.

Ultra-Narrowband

Some LPWAN technologies have adopted the strategy of using an Ultra-Narrow Band (UNB) channel to send low rate data. Messages are transmitted at rates as slow as 100 bits per second (bps), and such low transmission rates allow for channel bandwidths as narrow as 100 Hz. Restricting the channel to only 100 Hz drops the thermal noise floor at -154 dBm. LPWAN technologies that exploit UNB architectures have published receiver sensitivity as low as -142 dBm. Taking advantage of UNB can enable a simple and inexpensive transceiver design. SigFox, for example, is a proponent of the UNB strategy.

The magic of processing gain allows a spread signal below the noise floor to be decoded.

Spread Spectrum

The second technique commonly used by LPWAN technologies is called Spread Spectrum. With Spread Spectrum, the narrowband energy from low rate data is deliberately spread over a much wider frequency band than would otherwise be required. The ratio between the spread bandwidth and the original signal bandwidth is called the "processing gain."

When the wideband signals are de-spread at the receiver, it is converted back to the original narrow bandwidth of the low-rate message. The magic of processing gain allows a spread signal below the noise floor to be decoded.

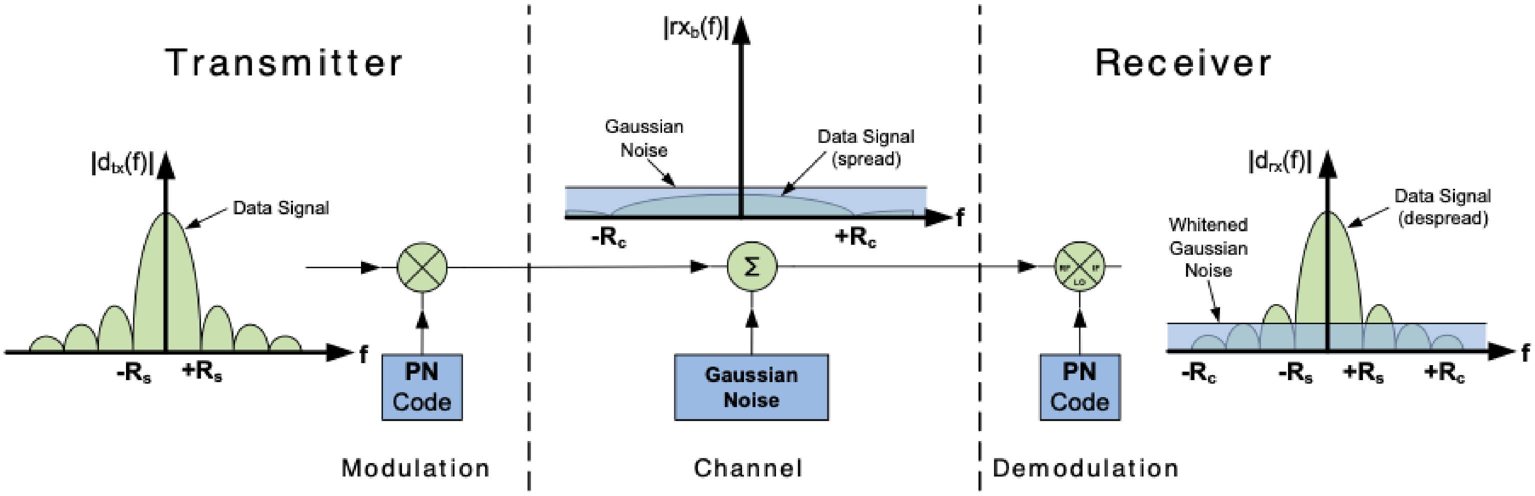

Impact of Gaussian Noise in a Spread Spectrum Transmission

(Source: J. Meel, "Spread Spectrum Introduction," Spread Spectrum Scene)

The process of spreading and de-spreading a message does not inherently improve sensitivity, but a transmission that is spread will be more resilient to interference and multi-path fading than one that is not.

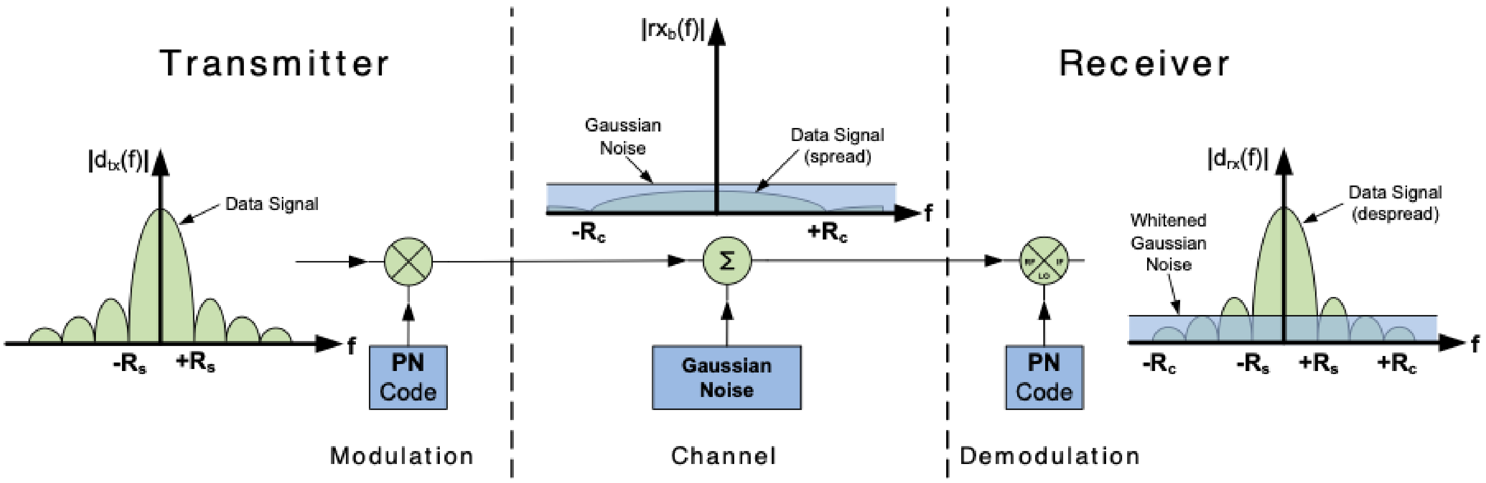

Mitigation of Interference in a Spread Spectrum Transmission

(Source: J. Meel, “Spread Spectrum Introduction,” Spread Spectrum Scene)

Besides the improved resilience of the transmission, the other reason for choosing to spread a transmission is that in certain jurisdictions, spread spectrum signals are allowed to transmit at much greater power levels than non-spread transmissions, leading to increased operating ranges.

On its own, spreading a narrowband signal over a wide bandwidth would be an inefficient use of spectrum, but a unique and orthogonal spreading sequence is applied to each signal, which allows signals that are superimposed on the same channel to be independently decoded. Ultimately, this results in a higher overall network capacity.

> By selecting an LPWAN technology that supports operation in a less restrictive operating band, network operators can leverage increased transmit power, if needed, to greatly increase coverage areas.

While spread spectrum techniques will enhance many aspects of the signal transmission between Access Point (AP) and node, it comes at some expense in terms of complexity and power consumption. Direct-Sequence Spread Spectrum (DSSS) and Chirp-Spread Spectrum (CSS) are two popular techniques of signal spreading used by various LPWAN technologies. LoRaWAN uses CSS and Ingenu RPMA uses DSSS.

Link Budget

A radio’s link budget is essentially the difference between maximum available transmit power at one end of the link and the minimum receive power needed to decode a message at the other end of the link. The greater the link budget, the longer a particular LPWAN’s operating range will be.

Certain jurisdictions and operating bands have quite restrictive limits on allowable transmit power levels. By selecting an LPWAN technology that supports operation in a less restrictive operating band, network operators can leverage increased transmit power, if needed, to greatly increase coverage areas. Be aware that this strategy drains batteries faster than others, but only for those devices operating at the furthest reaches of the coverage area.

Some published link budgets vary from as little as 149 dB (SigFox, USA) to 172 dB (Ingenu RPMA, USA). To illustrate the importance of link budget, the 23 dB difference between these competing technologies would roughly translate into a difference in operating range of about 0.9 km (coverage area=2.5 km^2) for the SigFox budget and an operating range of 7.6 km (coverage area 181 km^2) for the RPMA budget. These values are purely estimates but have been calculated based upon propagation models that are typical of real-world performance.

Interference

A reduction of interference sources inherently leads to higher communication reliability at a longer operating range. LPWAN technologies that have chosen to operate in sub 1GHz bands avoid exposure to the interference of the heavily used 2.4 GHz band in which Wi-Fi, Bluetooth, and numerous other types of consumer products operate.

Forward Error Correction

Forward error correction (FEC) is a message coding technique that is used to detect and correct errors occurring within the data of a received message. This capability allows received messages to be correctly decoded at lower signal power levels than would otherwise be possible without FEC, which translates to an increase operating range and/or saves precious bandwidth by reducing the need for message retransmissions.

FEC requires additional coding bits, which will increase the overall message size slightly, any LPWAN technology using FEC coding would be carefully architected to strike a desirable balance of performance.

Diversity

Many LPWAN technologies use diversity techniques to mitigate channel impairments such as multi-path fading and interference, improving signal-to-noise and signal-to-interference ratios. In urban environments, multi-path and interference issues can be especially pronounced, so applying diversity techniques will lead to increase coverage range.

Many LPWAN technologies use diversity techniques to mitigate channel impairments such as multi-path fading and interference, improving signal-to-noise and signal-to-interference ratios.

As an example, SigFox transmits a message three times, each at different frequencies and times. This means transmission is diversified in two ways, frequency and time, so on average, the system will benefit from improved range performance in certain operating environments. The down-side is that this methodology will inherently reduce the data capacity of the network by a factor of three.

Roaming

The majority of IoT applications monitor 'things' having limited or no movement within a coverage region. However, there are some applications—such as transportation—where the 'thing' (i.e. node) may be nomadic, moving from region to region, or possibly throughout the world.

Many of the proprietary IoT technologies have limited support for roaming nodes. There is a plethora of private and public network operators in existence, each with limited geographical coverage, but the vast majority do not have roaming agreements in place. Currently, LPWAN technologies based on cellular or satellite networks tend to be better suited to applications requiring the capability of global roaming.

Satellite vs. Terrestrial

The lion's share of IoT revenue will be derived from regions of the world where there is relatively high population density, therefore it makes economic sense to deploy conventional terrestrial-based LPWAN solutions in those cases. That said, there is still a significant number of opportunities in more remote locations that would benefit from the efficiencies that IoT connectivity can bring.

Although most of the current terrestrial-based LPWAN technologies offer amazingly large coverage areas, it is still far from being economically feasible to use terrestrial LPWAN to provide coverage throughout the entire earth. Fortunately, numerous companies are introducing satellite-based LPWAN technologies that will provide truly global coverage. Astrocast, EchoStar, Eutelsat, Hiber, Inmarsat, Iridium, Globalstar, Kepler, Kineis, Lacuna, Myriota, Orbcomm, and Swarm are a few of the many companies that plan to offer satellite-based IoT services. Most of these companies are building up their global coverage networks using low-cost, non-geostationary, Low Earth Orbit (LEO) cube satellites.

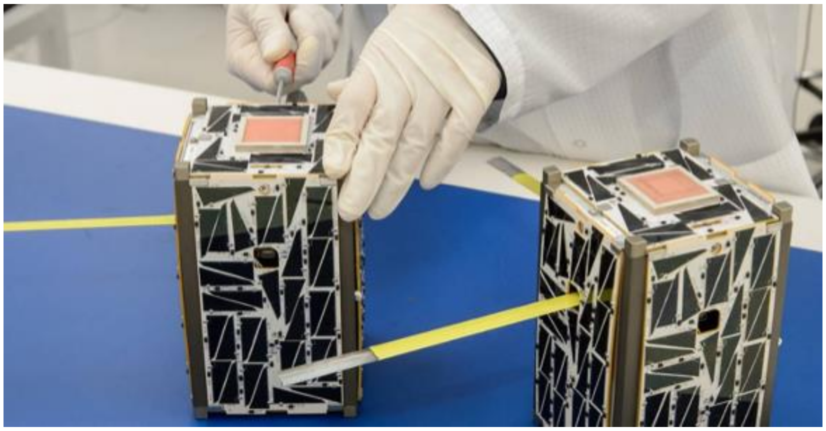

Low-Cost Cube Satellites (Source: https://www.nasa.gov/image-feature/nodes-integration)

Satellite LPWAN services will provide global IoT coverage and stand to revolutionize applications such as environmental monitoring, agriculture, public infrastructure management, shipping/logistics, energy, and any other type of application involving wide-area remote sensing.

These Satellite LPWAN services will provide global IoT coverage and stand to revolutionize applications such as environmental monitoring, agriculture, public infrastructure management, shipping/logistics, energy, and any other type of application involving wide-area remote sensing.

Energy Efficiency

All LPWAN technologies have been optimized to achieve previously unheard of levels of energy efficiency. While some LPWAN technologies may prove to be more energy efficient than others, generally speaking, LPWAN devices will outlast a conventional wireless device by one or two orders of magnitude for given battery capacity. To achieve such impressive battery lifecycles, LPWAN technologies have employed a variety of power-saving techniques. We discuss a few of the more popular methods in this section.

Low-Hardware Duty Cycles

All LPWAN technologies leverage hardware duty-cycling to achieve the exceptionally long battery life that the market demands. Hardware duty cycling is a strategy of disabling power-hungry circuits whenever possible to conserve the battery’s energy. The goal is to always operate in the lowest possible energy consumption state required for the tasks at hand.

Node circuitry can be modularized based on function, with each module having independent on/off power control. The circuit modules are typically divided into functions such as receiver, transmitter, sensors, memory, clocks, etc. Functions such as receiving and transmitting are typically heavy energy consumers, and as a result, LPWAN protocols have been tailored to support extremely low duty cycles of transceiver utilization.

Scheduling

Being able to pre-arrange transactions between a node and AP is a crucial technique LPWANs can use to save power. Knowing when to expect transmissions from the AP and conversely, when the AP is expecting a transmission from the node saves significant energy. The node can enter into an extremely low power state and wait for the next scheduled transaction event.

To minimize node energy consumption, an LPWAN protocol must be optimized to keep overhead to the absolute minimum needed to maintain network stability while still accommodating application data requirements.

Lightweight Medium Access Controls

The more overhead associated with a protocol, the more time a node will spend in a higher power operating state. To minimize node energy consumption, an LPWAN protocol must be optimized to keep overhead to the absolute minimum needed to maintain network stability while still accommodating application data requirements.

Many LPWAN medium access controls (MACs) forgo tight timing synchronization between APs and nodes to minimize protocol overhead and its associated power consumption. Doing so allows the nodes to also make use of low-cost oscillators. However, time drift leads to scheduling uncertainty; therefore, nodes must prematurely wake into their high-power states to ensure that a scheduled event is not missed. This results in more energy being consumed than would be the case if tight timing was maintained.

Network Topographies

Mesh Topography

A mesh-network topology makes it possible to geographically extend the reach of short-range wireless technologies by repeating transmissions through intermediary-nodes. However, as the network expands, node messages must pass through an ever-increasing number of intermediary-nodes to reach a distant AP. The burden of servicing pass-through traffic will deplete a node's battery life from years to months, or worse. The control traffic necessary to service a multi-hop networking structure can become quite considerable and will impair the network's capacity for payload data. Multi-hop routing yields longer communication delays, and these latencies can prove to be quite imbalanced and unpredictable among the devices.

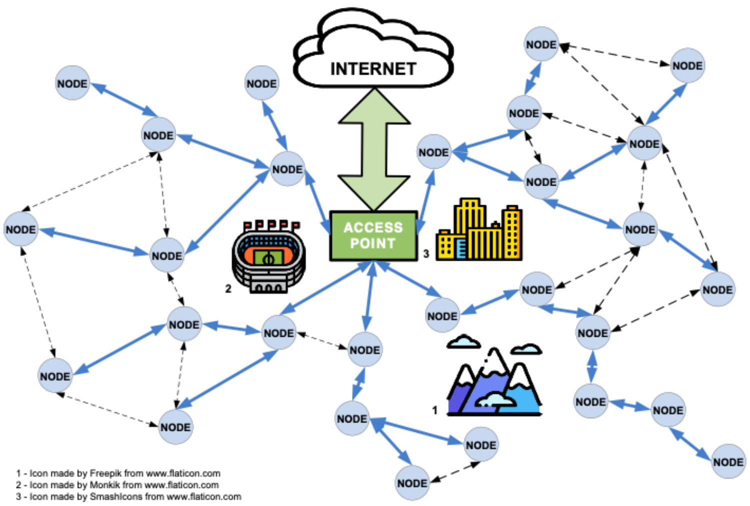

Illustration of a Mesh Wireless Network Topology

Star Topography

Star networks provide a single, always-on AP for direct communications to all nodes within the AP's coverage area. Star networks are by far the most common topology type used by proprietary LPWAN technologies. LPWANs have been architected to send relatively small amounts of data over very long distances and therefore are well suited to this simple network structure.

In this topology, nodes need not waste precious energy relaying random data traffic to and from neighboring nodes. Each node can preserve its limited energy resources strictly for its own data transactions.

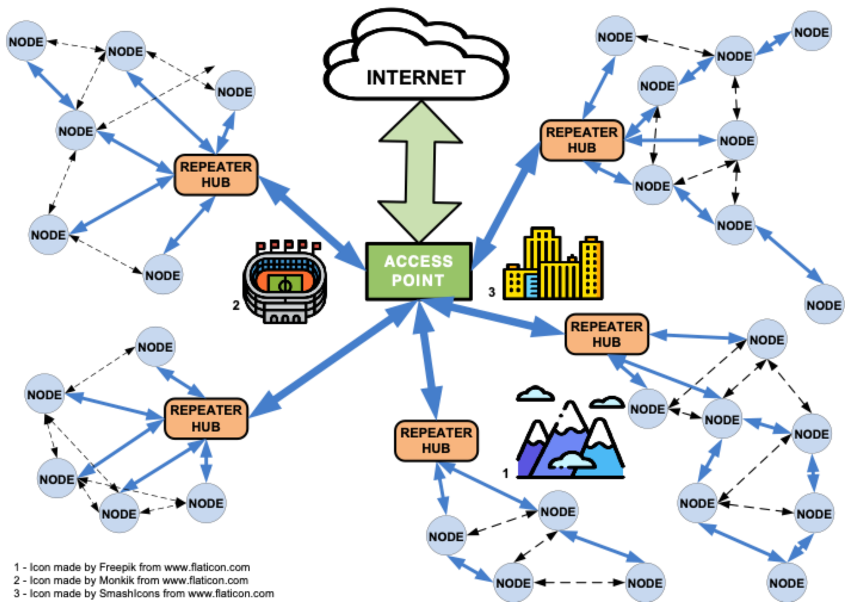

Star Wireless Network Topology

Adaptable Operating Modes

Some LPWAN technologies have the capability to change the node operating characteristics in order to place the node into the most power-efficient operating state required to ensure good quality uplink/downlink communications. Parameters that can be adjusted are things such as transmit power, spreading factors, bit rates, and modulation types.

To take advantage of dynamically adaptable operating parameters, the LPWAN technology must support bi-directional communication. SigFox and Weightless-N are both essentially unidirectional technologies, so they are not able to take advantage of these types of energy optimization techniques.

Capacity and Scalability

Densification

As IoT applications continue to grow exponentially, the need to support massive numbers of nodes located within geographically small areas will become a critical concern for many LPWAN networks. Some LPWAN technologies take a very simplistic approach to addressing high-density applications. If network capacity gets overwhelmed, the network operator has no choice other than installing additional APs to break the geographical region down into smaller coverage areas, each with a more manageable number of nodes.

However, installing more infrastructure equipment raises capital and operating expenses for the network operator. To remain competitive, network operators need to minimize the Capex and Opex expenses per node. More sophisticated LPWAN technologies have put a great deal of thought into being able to support massive numbers of nodes with a minimum amount of infrastructure equipment. In the long run, these types of LPWAN technologies could end up being the most cost-effective choice for urban and suburban coverage areas.

The more sophisticated LPWAN technologies can vary parameters such as transmit power, spreading factor, operating channel, modulation type, or bit rate. Leveraging these techniques allows networks to pack more nodes on to a network while maintaining consistent quality of service. How much a given LPWAN technology can manipulate these system parameters could be severely limited by restrictions of its operating band or the asymmetry of the link itself. In these cases, the LPWAN technology has no choice but to stick to very simple measures, or no measures at all, to improve bandwidth efficiency.

Regional Restrictions

Nearly all LPWAN technologies have solutions that will work globally, however, those technologies designed to operate in the sub-GHz bands tend to have a complicated mix of regional variants. Some global regions are quite restrictive in terms of the available frequency spectrum, channel utilization and transmit power. These jurisdictional limitations can significantly impair the capability and capacity of a network, so these restrictions should be well understood before settling on a given LPWAN type.

Operation in the 2.4 GHz ISM band or one of the 5G cellular bands have the least restrictions placed upon them, so operation in these operational bands can be fully optimized for capacity.

From the perspective of jurisdictional restrictions, operation in the 2.4 GHz ISM band or one of the 5G cellular bands have the least restrictions placed upon them, so operation in these operational bands can be fully optimized for capacity.

Channel Access Strategies

Carrier Sense Multiple Access

To avoid transmission collisions, channel access schemes such as Carrier Sense Multiple Access (CSMA) are popular in other wireless technologies such as WLANs. However, these schemes would require more sophisticated protocols to avoid breaking down in densely populated LPWAN networks.

CSMA protocols allow nodes to transmit data on an as-needed basis, but they must listen to the channel for the vacancy from other transmissions prior to attempting their own transmission. If the channel appears to be clear, the node will transmit its message to the AP. The AP confirms the reception of every node's message with an acknowledgment message (ACK). The node uses this ACK to gauge whether or not its message was successfully received. If a corresponding ACK does not arrive, the node will wait a random period of time before initiating a message retry process. (i.e. listen, transmit, wait for ACK)

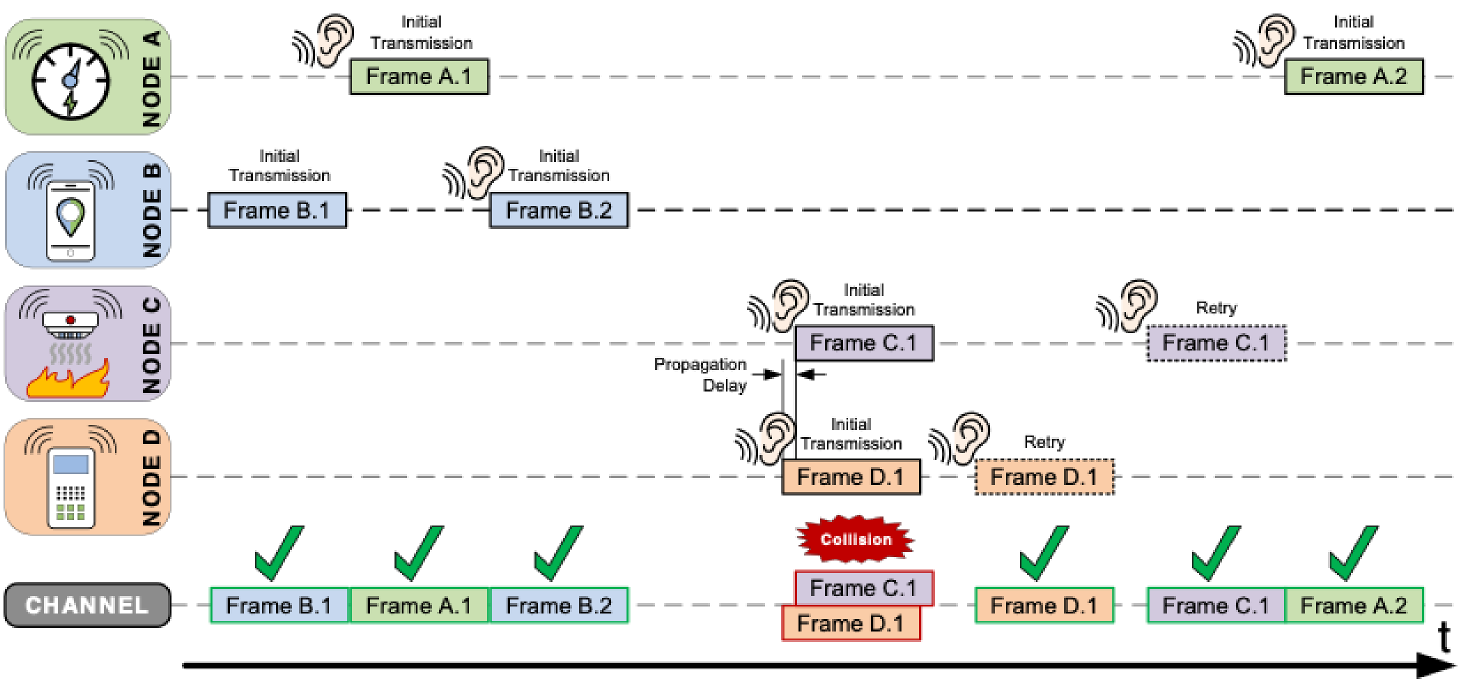

An Illustration of Carrier Sense Multiple Access Protocol

ALOHA

Several LPWAN systems such as LoRaWAN resort to basic random-access schemes such as ALOHA. In an ALOHA system, a node will transmit data as needed and wait for an acknowledgment (ACK) from the AP that its message has been received. If no ACK arrives, the node will wait a random amount of time before resending the data transmission.

ALOHA is a very simple scheme, and thus keeps the node design inexpensive, but can lead to inefficient use of bandwidth resources when a network gets busy. Additionally, noticeably degraded node battery life will result as nodes are forced to remain in high power operating states for longer periods of time as data traffic becomes congested on a heavily loaded network.

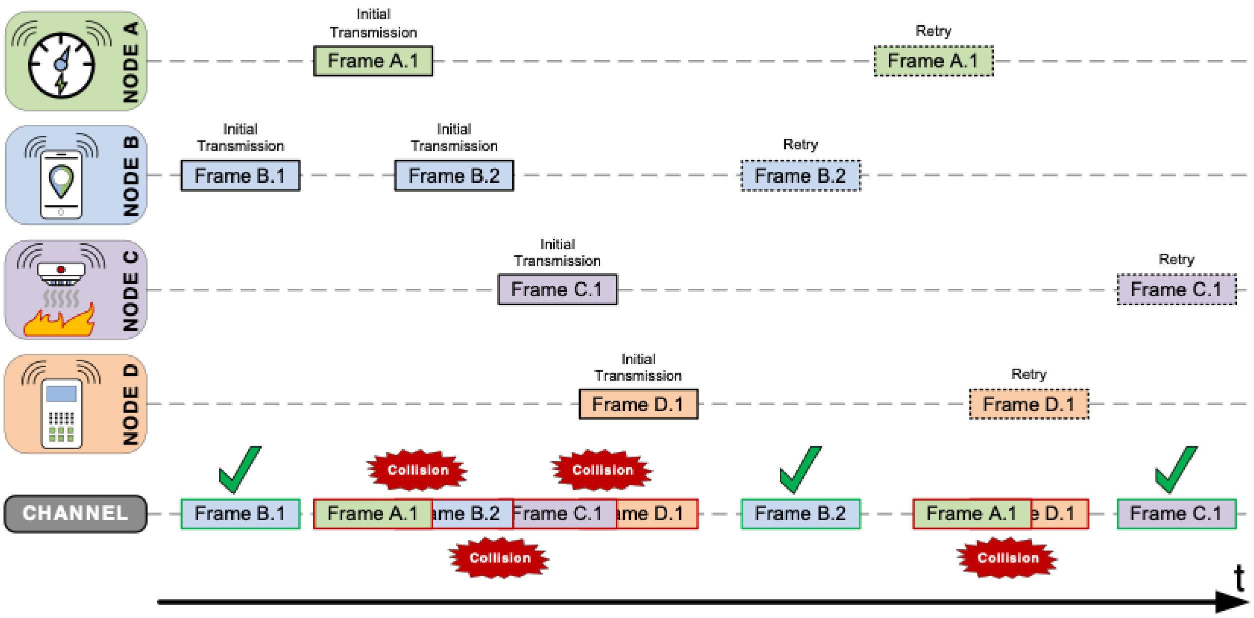

An Illustration of ALOHA Protocol

An ALOHA scheme is only feasible if the network supports bi-directional communications. SigFox and Weightless-N both lack sufficient downlink traffic capabilities to support message ACKs from the AP. SigFox mitigates uplink message uncertainty by repeating each message on three random frequencies and times, regardless of whether it is required or not. The intent being that through redundancy a message will hopefully arrive successfully at the AP. This approach clearly has a negative impact on precious bandwidth resources and node battery life, while still not guaranteeing a successful exchange of data.

Time Division Multiple Access

Some LPWAN technologies such as NB-IoT and Ingenu RPMA invested the implementation of Time Division Multiple Access (TDMA) protocols that allocate bandwidth resources more efficiently in a controlled time slot manner. These more sophisticated protocols require more processing resources to manage the network protocol but are better suited to supporting massive numbers of nodes within a network. In densely populated networks, the additional processing load is generally offset by the reduced time that the node spends in higher power states, thus maintaining critical energy efficiency.

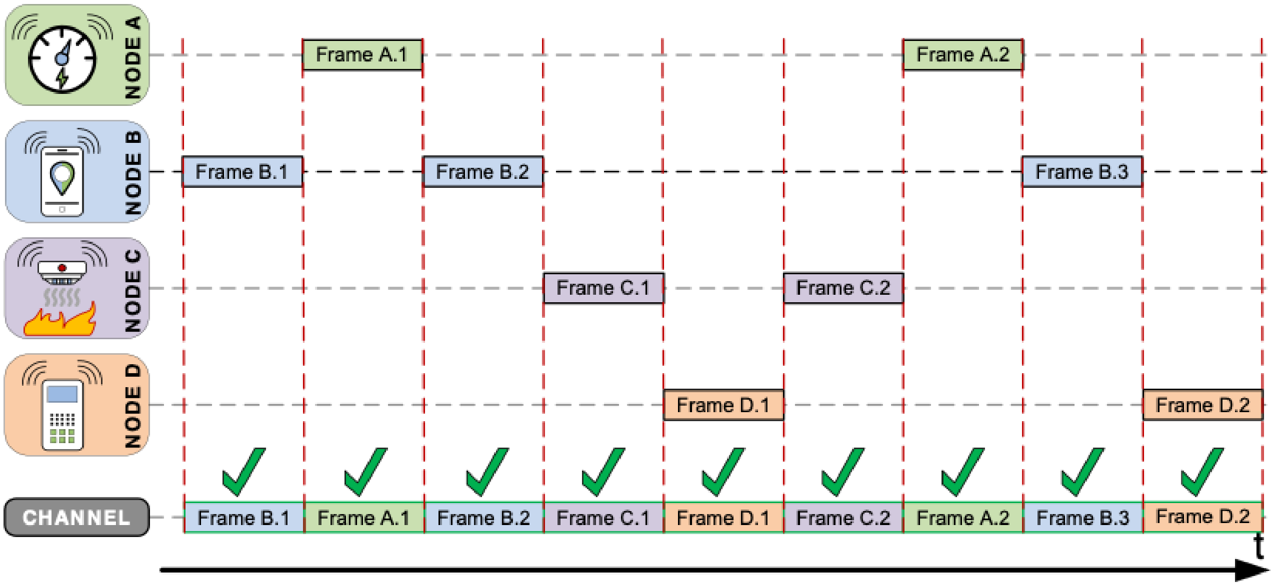

An Illustration of TDMA Protocol

Total Cost of Ownership

Economies of Scale

Certain LPWAN technologies are currently more widely adopted than others and therefore would stand to benefit from economies of scale. As production volumes increase, history shows that this leads to a corresponding drop in the price. Large scale market acceptance of a particular technology also attracts competitors, which leads to a number of benefits for the consumer, such as improved quality, more features, and lower prices.

LPWAN industry is still in its infancy, so first-to-market technologies may be the current volume leaders, but as new technologies emerge, they could prove to be disruptive and end up displacing the first-to-market technologies in terms of mass acceptance.

That said, the LPWAN industry is still in its infancy, so first-to-market technologies may be the current volume leaders, but as new technologies emerge, they could prove to be disruptive and end up displacing the first-to-market technologies in terms of mass acceptance.

Node Hardware Simplicity

As a general rule, simplicity equates to the lowest cost, so the most simplistic LPWAN technologies are likely to have the lowest cost application solutions. Any IoT applications should consider the most minimalistic LPWAN technology that suitably addresses all of the application requirements.

Minimal Infrastructure

Some LPWAN technologies will be inherently better suited to addressing certain applications than others. The best technology solution to support massive numbers of nodes with minimal infrastructure might be totally different than the best technology for a low-density application. A key takeaway would be to select a technology that will best address the application's coverage and capacity needs with a minimum of infrastructure.

License-Free Operation

Not all LPWAN technologies have chosen to operate in license-free bands. For instance, LTE-M, NB-IoT, and others leverage the existing cellular infrastructure and therefore will operate in a licensed band. License ownership for these bands has been limited to a few very large telecom companies with the financial resources needed to afford the extremely high license costs. The high costs of these licenses must be recouped from the customers using these networks and ultimately usage fees will be higher than if there was no network license involved. This high cost of entry reduces competition, which is generally not a good outcome for consumers of the services.

Ease of Deployment, Operation, and Maintenance

An important consideration when selecting any LPWAN technology is the cost associated with deploying, operating and maintaining the network infrastructure. A sophisticated, but more expense LPWAN technologies may be capable of covering larger geographic regions and supporting greater numbers of nodes than an inexpensive, but simplistic competing technology. Ultimately, it could end up being more cost-effective to use a more sophisticated LPWAN technology for larger network applications due to the reduced cost of deploying and maintaining the network infrastructure.

Security

Data Encryption

Most LPWAN technologies implement some form of symmetric key cryptography. A unique and secret key is shared by the AP and node and is used to encrypt data before transmission and decrypt data upon reception.

Network Join Authentication

LPWAN technologies can also employ mutual authentication techniques between a node and an AP as part of the network join procedure. These measures ensure that only genuine and authorized devices will be connected to genuine and authentic networks.

Spread Spectrum

Those LPWAN technologies that employ spread spectrum techniques make the signal much more difficult to detect and eavesdrop upon. It also makes the signal more resilient to interference and denial of service attacks by means of frequency jamming.

Over-The-Air Firmware Updates

An essential security feature of any LPWAN technology should be the ability to upgrade the firmware (FW) of nodes over-the-air (OTA). Some LPWAN technologies are unable to implement an OTA FW update and this would leave that network vulnerable to future security risks or bugs. If a threat to network security was discovered, and an OTA update facility was unavailable, network operators may have no other choice than to recall all nodes from the field for a firmware update. This process could prove to be very time consuming and costly.

SigFox and Weightless-N both fall into the category LPWAN technologies that are not capable of doing OTA FW updates.

How SignalCraft Can Help

It's easy to see why LPWANs are a cost-effective way to provide remote monitoring and control of massive numbers of devices, often spread across huge areas. Having the ability to manage devices remotely will lower operating costs and raise efficiencies for those businesses that choose to take advantage of the IoT.

It can be a quite daunting task to wade through the numerous competing LPWAN technologies and select a suitable solution for your application. SignalCraft has years of wireless product design experience and specific application experience with various LPWAN technologies. We help fast-track our customers through the process of selecting an appropriate LPWAN technology and integration of that technology into their IoT application.

SignalCraft offers the engineering services needed to realize fully customized Industrial IoT (IIoT) designs, and we also have an off-the-shelf line of IIoT products capable of providing remote monitoring and control functionality for some of the most popular industrial interfaces. We developed this platform of products to be agnostic to any specific LPWAN technology so they can be readily adapted to whatever technology is appropriate. Product SKUs are available for both typical and hazardous operating environments.

Visit our Website to learn more about our IoT product family or contact us to discuss a custom

About SignalCraft Technologies

We build brilliantly designed, high frequency digital and RF products, 100% in-house from the ground up to your specs and schedule. From leading global test brands to industrial communications startups, technical leaders trust SignalCraft as their wireless product development partner.

signalcraft.com