A COMPREHENSIVE LOOK AT

Low Power, Wide Area Networks

For 'Internet of Things' Engineers and Decision Makers

Low power wide area networks (LPWAN) are not a new phenomenon

HOWEVER, THEY ARE BECOMING MORE POPULAR DUE TO THE GROWTH OF THE INTERNET OF THINGS (IoT).

LPWAN is often used when other wireless networks aren't a good fit—Bluetooth and BLE (and, to a lesser extent, WiFi and ZigBee) are often not suited for long-range performance, and cellular M2M networks are costly, consume a lot of power, and are expensive as far as hardware and services are concerned.

LPWAN technology is perfectly suited for connecting devices that need to send small amounts of data over a long range, while maintaining long battery life. Some IoT applications only need to transmit tiny amounts of information—a parking garage sensor, for example, which only transmits when a spot is open or when it is taken. The low power consumption of such a device allows that task to be carried out with minimal cost and battery draw.

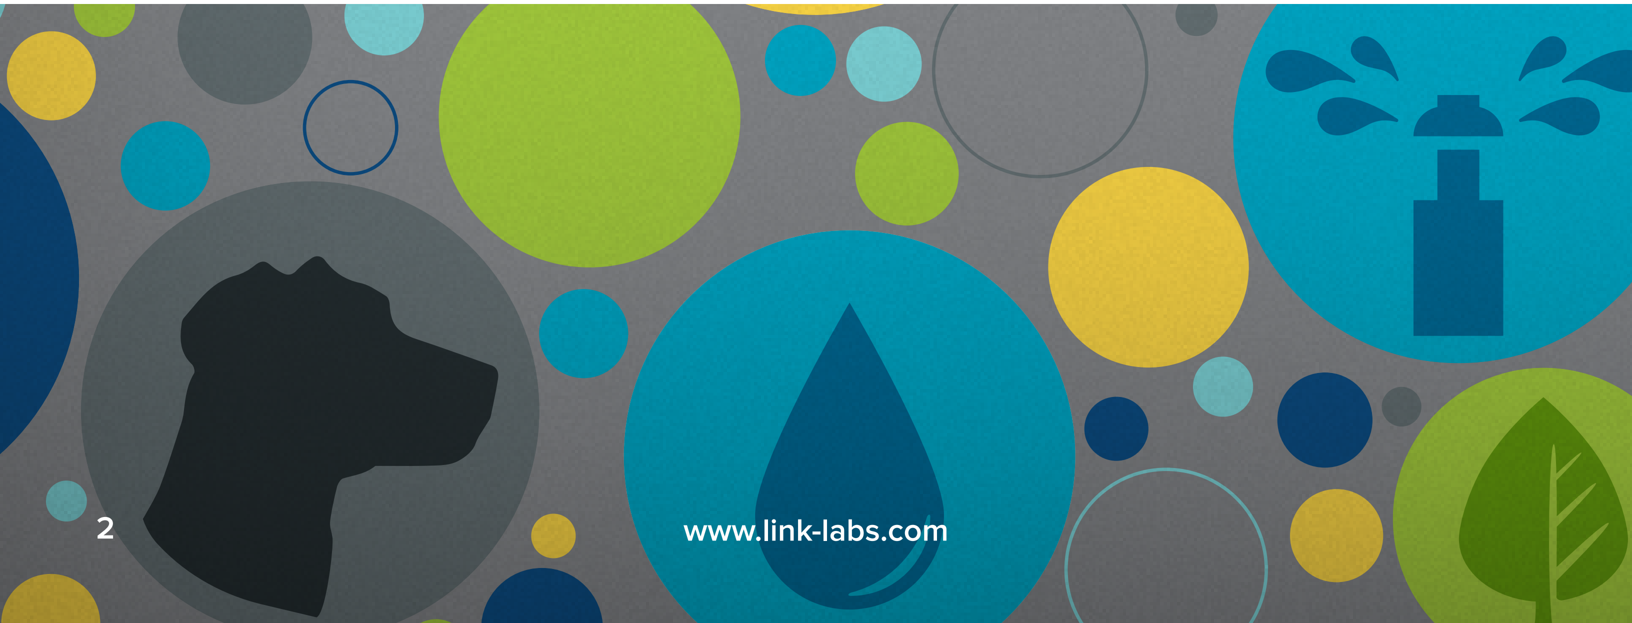

LPWAN FEATURES

LONG RANGE
The end-nodes can be up to 10 kilometers from the gateway, depending on the technology deployed.

LOW DATA RATE
Less than 5,000 bits per second. Often only 20-256 bytes per message are sent several times a day.

LOW POWER CONSUMPTION
This makes very long battery life, often between five and 10 years, possible.

WITHIN THIS DOCUMENT, WE'LL BE LOOKING AT:

- Use cases where LPWAN technologies are best suited.
- Nine fundamental LPWAN concepts.
- Five main LPWAN technologies.

There are two main areas where LPWAN technologies are best suited

1 FIXED, MEDIUM- TO HIGH-DENSITY CONNECTIONS

In cities or buildings, LPWAN technologies are a great alternative to cellular M2M connections. Some examples include smart lighting controllers, distribution automation (smart grid), and campus- or city-focused GPS asset tracking.

2 LONG LIFE, BATTERY-POWERED APPLICATIONS

When a longer range is needed than legacy technologies can provide, LPWAN can be a good fit. Examples include wide-area water metering, gas detectors, smart agriculture, and battery-powered door locks and access control points.

THE SWEET SPOT FOR LPWAN

Different wireless technologies address application-specific needs with changes in modulation and frequency schemes. Long-range applications with low bandwidth requirements that are typical for IoT applications are not supported well by these existing technologies.

SOURCE: PETER R. EGLI, 2015, http://www.slideshare.net/PeterREgli/lpwan

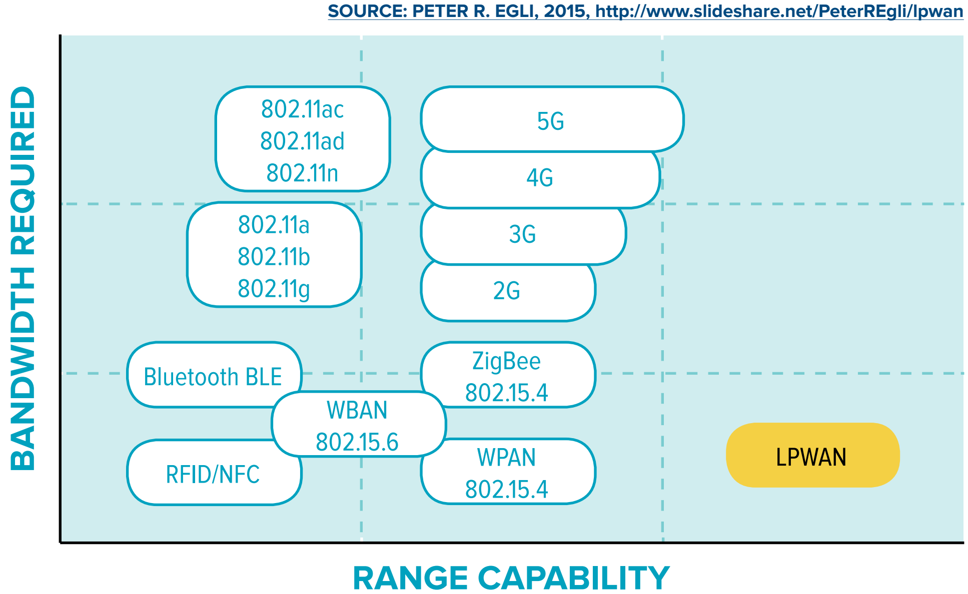

Network Configuration

MESH TOPOLOGY

Many Link Labs LPWAN customers have previously tried to solve their wireless connection **problems with mesh topology networks** like ZigBee. They struggled with mesh network solutions because the link budgets for these connections are very limited due to high data rates and low receiver sensitivities. Some ZigBee connections have trouble sending data more than 20-30 meters away because the power coming from the transmitter is lost too quickly. Additionally, others have been surprised by the amount of mesh infrastructure required to actually build a reliable network.

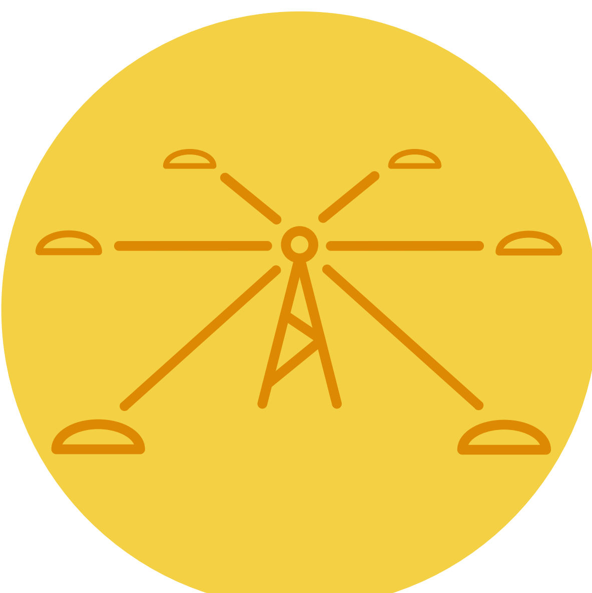

STAR TOPOLOGY

Instead of a mesh topology network, most **LPWAN technologies use a star topology network.** Similar to WiFi, the endpoints of star networks are connected directly to the access point. Link Labs can use a repeater to easily fill in gaps in coverage, which, for most applications, is a good middle ground in terms of latency, reliability, and coverage.

Fundamental LPWAN Concepts

Range Vs. Data Rate

TO ACHIEVE LONG RANGE IN WIRELESS COMMUNICATIONS, YOU NEED A LARGE LINK BUDGET.

In other words, when you transmit a signal, it needs enough energy to be detected when it's received. Because a certain amount of power is lost along the way as it propagates through space and materials in between, there is a baseline amount needed to transmit the signal properly.

LPWAN technologies generally operate with about 140-160 decibels (dB) of total path, which can add up to many miles of range in the right circumstances. This is primarily achieved by high receiver sensitivities. Receiver sensitivities of more than -130 dBm are common in LPWAN technologies, compared with the -90 to -110 dBm seen in many traditional wireless technologies. Technologies with -130 dBm can detect signals 10,000 times weaker than technologies with -90 dBm, so you can see how this is important for LPWAN.

The slower the modulation rate, the higher the receiver sensitivity can be. This comes down to the Shannon-Hartley theorem, or Information Theory, which states that the energy per symbol or energy per bit is the main lever to change the possibility of a message being heard. By slowing the modulation rate by half, you are putting twice as much energy into each symbol; thus, you are increasing the link budget, or receiver sensitivity, by double (3 dB).

Sigfox[1] is an example of how modulation rate and range are connected. Sigfox transmits data using a standard radio transmission method called binary phase-shift keying (BPSK). Its modulation rate—300 bps—is extremely slow in a modern sense. But due to this slow modulation rate, it's able to get great range with fewer base stations. In the U.S., Sigfox modulates at a higher rate, because otherwise it would not be able to meet the FCC Part 15 requirement that the maximum time a transmission can be on the air is 0.4 seconds.

Symphony Link™[2] from Link Labs uses a continuously variable data rate, which adjusts the modulation rate in accordance with the channel fade estimate from the end node's perspective. In other words, we put enough energy into each signal to make the link.

Processing Gain

THE TECHNICAL DEFINITION FOR PROCESSING GAIN IS THE RATIO OF THE RADIO FREQUENCY (RF) BANDWIDTH TO THE UNSPREAD BANDWIDTH, USUALLY EXPRESSED IN DECIBELS.

Here's a simple way to think of it: imagine you're sitting in front of a TV screen, and all you see is static. That static can be thought of as a visual representation of noise. Now let's assume you are able to press pause on your TV remote, freeze the static, put a transparency to your TV screen, and color in all the black pixels until you had an exact replica of the static at that moment. If you

1. http://www.sigfox.com
2. http://www.link-labs.com/symphony-module/

then decided to label this transparency as "Static X," you could press play again and, with the transparency in hand, watch the static until you saw a frame that looked similar to your drawing. Once this happened, you could say that someone had transmitted Static X.

When applied in more realistic terms, this processing gain illustration shows that when a signal is mixed across the RF spectrum, it is only detectable when you have processed all of the noise and are looking at it with a filter. Negative dB signal-to-noise ratio (SNR) means that the signal is below the noise floor: **it can't be seen with a simple receiver unless you are looking for it.** This, in a nutshell, is processing gain.

As another example, Sigfox technology is a BPSK narrowband signal with very narrow channel sizes. Weak signals are more easily detected in a narrow channel, since the noise floor is effectively lower than it is for wider band signals like LoRa. This is because noise is spread throughout the spectrum. If your receiver bandwidth is smaller, then the noise level is smaller too. However, traditional frequency-shift keying (FSK) signals—which transmit information through the frequency changes of a carrier wave—have no "processing" or "coding" gain. This means they must have a positive signal-to-noise ratio of around 10 dB to detect the signal.

When coding is used, a signal can be detected up to around -20 dB SNR. For code-division multiple access (CDMA) signals like Ingenu's or chirp spread spectrum (CSS) modulations like LoRa, the effects of a higher receiver noise floor are mitigated by processing gain. For the most part, coded signals are better than narrowband signals in terms of minimum detectable energy, but there are some drawbacks associated with them, which we will discuss below.

Noise Vs. Bandwidth

AS REFERENCED DURING OUR DISCUSSION ON PROCESSING GAIN, THE NOISE FLOOR OF A RECEIVER IS SET BY TWO THINGS: THE BANDWIDTH AND THE NOISE.

Think of it this way: if you look through a pinhole, you see less light than if you look through a paper towel roll. That same logic can be applied to radios.

A narrowband channel of 100 Hz has a thermal noise floor of about -154 dBm, which means that if you require a 10 dB SNR, your ideal receiver (and theoretical maximum sensitivity) would get -144 dBm (unless you use coding). If you use coding in that channel, then you can get below the noise floor to the same sensitivity—but coding a signal requires bandwidth to spread the energy across.

Symphony Link uses a 125 kHz channel size, which has a thermal noise power of -124 dBm. Since we can achieve up to 20 dB of coding gain, our theoretical maximum sensitivity is also -144 dBm.

Interference

BASED ON THE COMPARISON BETWEEN NOISE AND BANDWIDTH, YOU UNDERSTAND THAT THEORETICAL PERFORMANCE BETWEEN A NARROWBAND CHANNEL AND A CODED CHANNEL IS THE SAME.

But many people in the LPWAN space disagree on which technology is better when it comes to noise. (For example, this article demonstrates Texas Instruments' opinion[3] on long-range RF communication.)

3. http://www.ti.com/lit/wp/swry006/swry006.pdf

Narrowband noise can be thought of as actual narrow signals sitting above the noise floor. If you're a narrowband system (like Sigfox) looking at a tiny 100 Hz channel, the channel next door can be loud, and it still won't affect you. If it's an interferer in the same channel, however, you're going to get clobbered.

For **narrowband noise**, which much of the 900 MHz ISM band interference is, if a narrowband signal gets "clobbered" (i.e., a signal lands right in the channel), it has very poor blocking performance. However, narrowband systems can have more than 50 dB adjacent channel rejection.

Wide-band noise is like widespread noise that effectively raises the noise floor. Narrowband interference is less of a problem for a wide-band, coded system, since it just adds to the overall noise in the band.

Wondering which system is better? Frankly, we don't know, because it depends on the specifics of the environment. This is widely discussed within the LPWAN space, but we won't make any speculations within this paper.

Licensed Vs. Unlicensed

MOST CURRENT LPWAN TECHNOLOGIES USE AN UNLICENSED BAND. SIGFOX AND LINK LABS BOTH USE THE 900 MHZ ISM BAND IN THE U.S. AND THE 868 MHZ BAND IN EUROPE. INGENU USES THE 2.4 GHZ BAND.

All the technologies mentioned above work just as well in licensed bands. In fact, they work better because there is less interference from other users. So what's the issue? When using licensed bands, you have to re-tool the MAC scheme to deal with different channel size, spacing, etc. For instance, if you go to a licensed spectrum, you'd probably have less than 1 MHz of spectrum, whereas in unlicensed bands, you could get 26 MHz.

Both coded and narrowband signals work well in licensed spectrums, but using them becomes an issue of spectral efficiency and capacity more than anything else. (We'll discuss this in depth in the section on orthogonality). Essentially, the challenge is to pack as much data flow into the band as possible. The FCC Part 15 and ETSI rules go out the window as well, because as a license holder, you have much more freedom to use your spectrum to your advantage.

However, the issue of licensed vs. unlicensed spectrums may not be an issue much longer. In August 2015, the GSMA (Groupe Speciale Mobile Association)—a group made up of mobile operators—announced that it plans on standardizing LPWAN technology on a licensed spectrum by early 2016. This push has received endorsements and backing from companies like "AT&T, Bell Canada, China Mobile, China Telecom, China Unicom, Deutsche Telekom, Etisalat, KDDI, NTT DOCOMO, Ooredoo, Orange, Singtel, Telecom Italia, Telefonica, Telenor, Telstra and Vodafone," according to an article put out by Telecom TV[4]. We're very interested to see if this breaking LPWAN news moves forward and on GSMA's proposed timeline.

Sub-GHz Spectrum Availability Worldwide

THIS IS PROBABLY THE BIGGEST DRAWBACK FOR LPWAN TECHNOLOGIES IN THE 900/868 MHZ BAND.

Every country has different rules about using the sub-GHz spectrum. There are generally two

4. http://www.telecomtv.com/articles/iot/mobile-operators-look-to-take-charge-of-iot-friendly-lpwan-development-12760/

camps: those that follow Europe (868 MHz), and those that follow the U.S. (915 MHz). The 915 MHz band is available only in about a third of the world, and some countries don't have any bands available. In fact, many countries have added special caveats that make standardization nearly impossible. Until this issue is resolved, there is no globally available band for LPWAN technologies like there is at the 2.4 GHz level (for Bluetooth and WiFi).

At Link Labs, we've attempted to solve this problem by allowing our radios to scan for an access point in several bands, which tells the endpoint how to behave and where to transmit. This is part of the reason why widespread standardization in sub-GHz will face big problems. (We suggest reading this guide on worldwide sub-GHz bands[5] for more information.)

Localization Capabilities

MEASURING THE LOCATION OF AN RF SIGNAL IS DONE BY ESSENTIALLY CONVERTING THE TIME OF ARRIVAL INTO A DIRECT PATH LENGTH. SO TO MEASURE TIME, YOU ABSOLUTELY MUST BE ABLE TO DETECT A DIRECT PATH.

There are two things that go into the location (or really the time of arrival) of an RF signal: enough **power** to detect the direct path and enough **bandwidth** to resolve the multipath reflections from the direct path.

Imagine you're in the living room, and someone is in the bedroom with a strobe light on. You can't see this person, but you can see the strobe lights, because the light is bouncing around and refracting off the walls. In a discussion on network localization capabilities, that type of light transmission would be considered a multipath (or non-direct) path. Most LPWAN technologies are *not* received on the direct path; they are received on the multipath channels. This is a good thing for data reception, since weak signals that have bounced frequently are still received, but it also means that LPWAN technologies are not ideal for localization. Unfortunately, no amount of averaging can change this—it's simply the laws of physics. Averaging only helps if something is moving (in space and frequency), and in most LPWAN systems, neither of these are happening.

On the other hand, if you're using radio waves, you *must* be able to detect the direct path. If you're unable to do so, you'll be potentially creating a huge area of uncertainty. If you used the wrong measurements in your calculations, for example, you could end up with a kilometer-wide uncertainty circle.

Signal bandwidth is required because the ability to determine the difference in path length between two signals (say, the direct path, and a multipath reflection) is a function of signal bandwidth. A narrowband signal (100 Hz) could never be used for accurate time-based measurement, and even a 125 kHz LoRa signal only has a multi-path resolution ability of about 1 km. That means if there are any reflected paths with a length of less than one kilometer different from the direct path, the measurement will not be accurate.

Because of these limitations, we encourage those who need localization capabilities to look into GPS, WiFi, or proximity-based RFID.

5. https://www.scribd.com/doc/273973068/UHF-Regulations-Sub-GHz-ISM

Orthogonality

WHEN TWO LINES ARE *ORTHOGONAL*, IT MEANS THEY ARE BOTH RIGHT ANGLES. THE RF WORLD "HIJACKED" THIS TERM, IF YOU WILL, TO MEAN TWO SIMULTANEOUS SIGNALS THAT ARE BOTH DETECTABLE.

So ***orthogonality* is detecting multiple data streams in the same channel and at the same time.** This is a feature of a coded channel, and it offers a solution for getting back good spectral efficiency for wider band systems. Because coded signals are spread across a larger swath of spectrum, those signals take up more frequency real estate. Narrowband signals, however, can pack quite a bit of traffic into that same bandwidth. If there are multiple coded streams simultaneously on the air, you buy back some (though usually not all) of the spectral efficiency you give up with coding. FSK systems cannot detect more than one signal at a time, and if two signals use the same channel at the same time, only the stronger signal will be decoded, if at all.

Importance of MAC Protocols

Ultimately, much of the value of a LPWAN technology is not the underlying RF characteristics, assuming that the "link is closed." The ability to create a network, control it, and offer bi-directional data flow is what matters most for end users. The limitations or features of one MAC implementation compared with another are very important to understand.

HERE ARE SOME LPWAN FEATURES THAT LINK LABS' SYMPHONY LINK MAC LAYER PROVIDES:

- Uses repeaters to fill in coverage
- Real Time Adaptive Data Rate
- Open Standard
- International Roaming Support (Multi-Band)
- 100% Acknowledged Messages
- Over-the-air Firmware Upgrades
- Multicast Message Groups
- Flexible Downlink Capability
- Scalable Capacity
- Low Downlink Latency
- Uplink Power Control
- Real Time Quality of Service
- Handover
- Interference Avoidance
- Supports Internet Disconnected Operations
- Supports high jitter (SATCOM) connections
- Supports 1W Uplink Transmissions under FCC
- Uplink-Downlink Collisions Prevented
- MAC Layer Packetization and Retry
- Fixed MTU Size

Top LPWAN Platforms & Technologies

BELOW ARE DESCRIPTIONS OF THE SIX KEY PLAYERS IN LPWAN ALONG WITH A LIST OF THEIR TECHNOLOGIES' ADVANTAGES AND DISADVANTAGES.

Link Labs

Link Labs builds hardware that supports the LoRa Alliance standard, but also developed a proprietary system to provide more advanced functionality. Symphony Link is the Link Labs[6] LPWAN solution. It uses the LoRa PHY (Physical Layer), but not the standard LoRaWAN MAC architecture. It has a range that is 100 times that of WiFi and, like most of the other listed competitors, is far more cost-effective than cellular networks. The Symphony Link gateway is an eight-channel sub-GHz base station that is ideal for industrial or municipal monitoring applications, like building management[7] or smart indoor and outdoor lighting[8]. The Link Labs transceiver modules allow developers to bring long-range communications to their devices. These devices operate in the 915 MHz ISM band, and they are ETSI certified for use in the 868 MHz band.

POSITIVES

- **Many miles of range** are possible due to high sensitivity (up to -137 dBm).
- **Flexible technology** is capable of licensed or unlicensed deployment from 150 MHz to 1 GHz.
- **Most sophisticated MAC functionality.**

CONSIDERATIONS

- **Requires LoRa chipsets and Symphony Link™-specific software.**

6. http://www.link-labs.com
7. http://www.link-labs.com/iot-solutions/enterprise/
8. http://www.link-labs.com/smart-lighting/

Nwave

Nwave[8] runs off an ultra narrowband (UNB) radio, which operates in sub-1 GHz ISM bands. Like Link Labs, it operates a star topology, allowing for direct base station communication. Nwave has what it brands as "advanced demodulation techniques," which are meant to allow its network to coexist with other radio technologies without additional noise.

POSITIVES
- **Similar to Sigfox, with a better MAC-layer implementation.**

CONSIDERATIONS
- **Less is known about this technology.**

Ingenu (formerly On-Ramp)

Ingenu[9] wireless technology is called Random Phase Multiple Access (RPMA). It differentiates itself from the market with this flexible network system. Like Nwave, it has good MAC-layer implementation. Ingenu believes RPMA is the protocol by which a device network standard should be built and, as such, is a founding member of the IEEE 802.15.4k task group, which is dedicated to low-energy infrastructure monitoring.

POSITIVES
- **Very good technology stack.**
- **Very high capacity.**
- **Good commercial traction.**

CONSIDERATIONS
- **It uses 2.4 GHz,** a band with a lot more interference.
- **Propagation loss is significantly more** at its higher frequency.
- **Its underlying technology takes much more processing power** (and thus more actual power), so it's not as well suited for battery-powered applications.

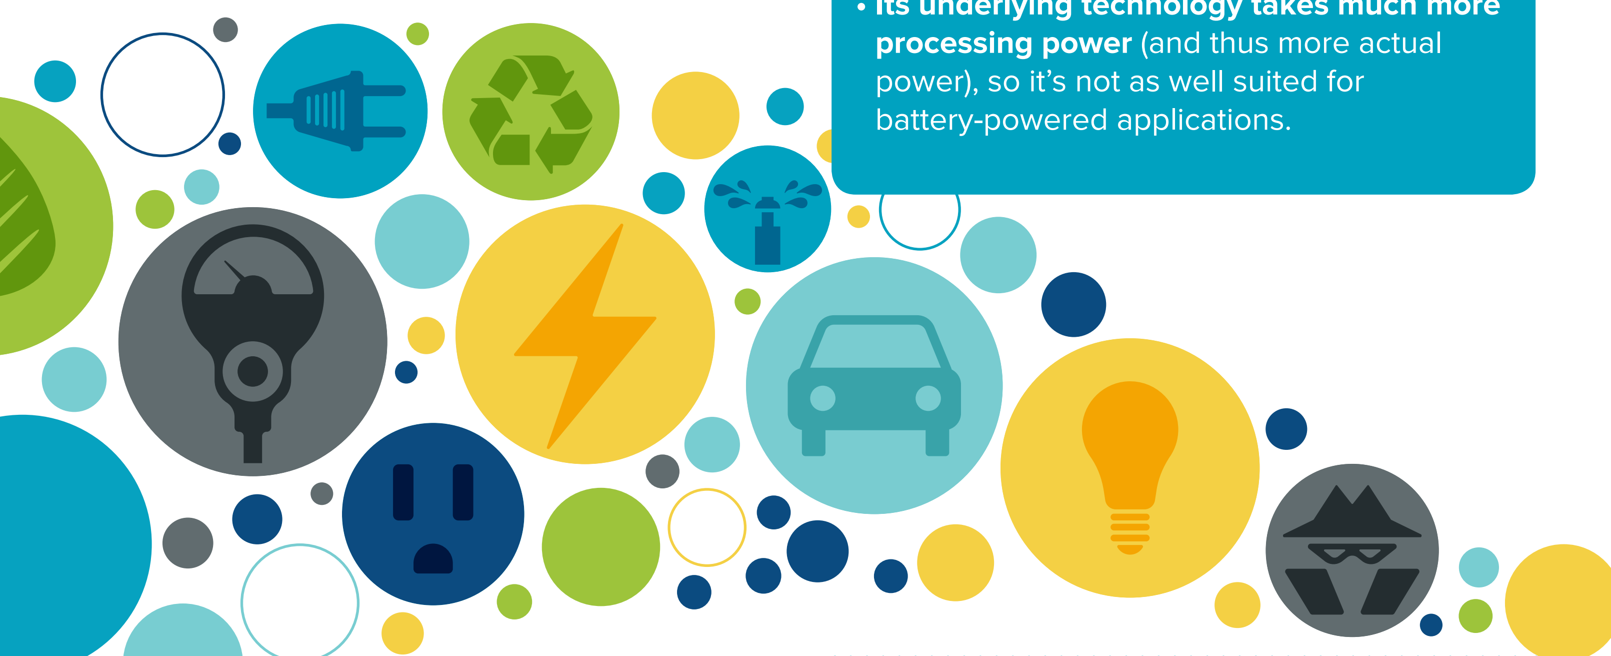

8. http://www.nwave.io
9. http://www.ingenu.com

Sigfox

Sigfox[10] sets up antennas on towers (like a cell phone company) and receives data transmissions from devices like parking sensors or water meters. These transmissions occur in the 868 or 915 MHz bands, as we discussed earlier.

Sigfox's wireless systems send very small amounts of data (12 bytes) very slowly (300 baud) using BPSK. The long-range capabilities of Sigfox are accomplished as a result of very long and very slow messages. As mentioned above, Information Theory states that the slower you transmit a message, the easier it is to "hear."

This technology is a good fit for any application that needs to send small, infrequent bursts of data. Things like basic alarm systems, location monitoring, and simple metering are all examples of one-way systems that might make sense for this network. In these networks, the signal is typically sent a few times to "ensure" the message goes through. While this works, there are some limitations, such as shorter battery life for battery-powered applications. Sigfox announced an upgrade to its original system in 2015 that allows for guaranteed message acknowledgment for up to four messages a day.

10. http://www.sigfox.com

POSITIVES

- **Sigfox has gained a lot of traction** in the LPWAN space, and it's deploying in a lot of areas.
- **It has a good ecosystem of radio vendors,** like Texas Instruments, Silicon Labs, Axom, and others that support its technology.

CONSIDERATIONS

- **Currently, Sigfox offers limited traffic profiles, which cause download traffic to be constrained.** This limits the end user to 15 bytes of traffic at a time with about 10 messages and only four acknowledgements a day. So, it's aimed at very simple devices, which would be insufficient for some companies.
- **The company has faced challenges in moving its technology into the U.S. market.** Under FCC Part 15, the maximum time a transmission can be on the air is 0.4 seconds. Since Sigfox transmissions last three seconds or so, this has required a new architecture, and it is likely the reasonit has been slower to deploy in the U.S. than promised. The frequency band in the U.S. is also subject to much higher levels of interference than the band used in Europe.

Weightless

Weightless[11] is an open standard. It believes a global standard can be achieved by allowing for open software innovation. Like the other technologies here, Weightless protocols operate in sub-1 GHz unlicensed spectrum.

POSITIVES

- **Its three open standards provide the end user with more choices.** Weightless-N offers a simple one-way directional standard with a very long (10-year) battery life and a low overall cost. Weightless offers two-way communication, but it has a shorter battery life and higher network cost. Weightless-W is the most extensive option, and it runs off of unused TV spectrum, but has some drawbacks (see consideration at right).
- **Weightless has an open ecosystem,** meaning there's more open software and vendors available. It runs the Weightless Special Interests Group (SIG), a nonprofit organization formed to develop its open standards, and test upcoming technologies.

CONSIDERATIONS

- Its most extensive open standard, Weightless-W, **has a shorter battery life** (three-five years) **and a higher cost** for both the terminal and the network.
- Like Nwave, less is known about this technology.

11. http://www.weightless.org/

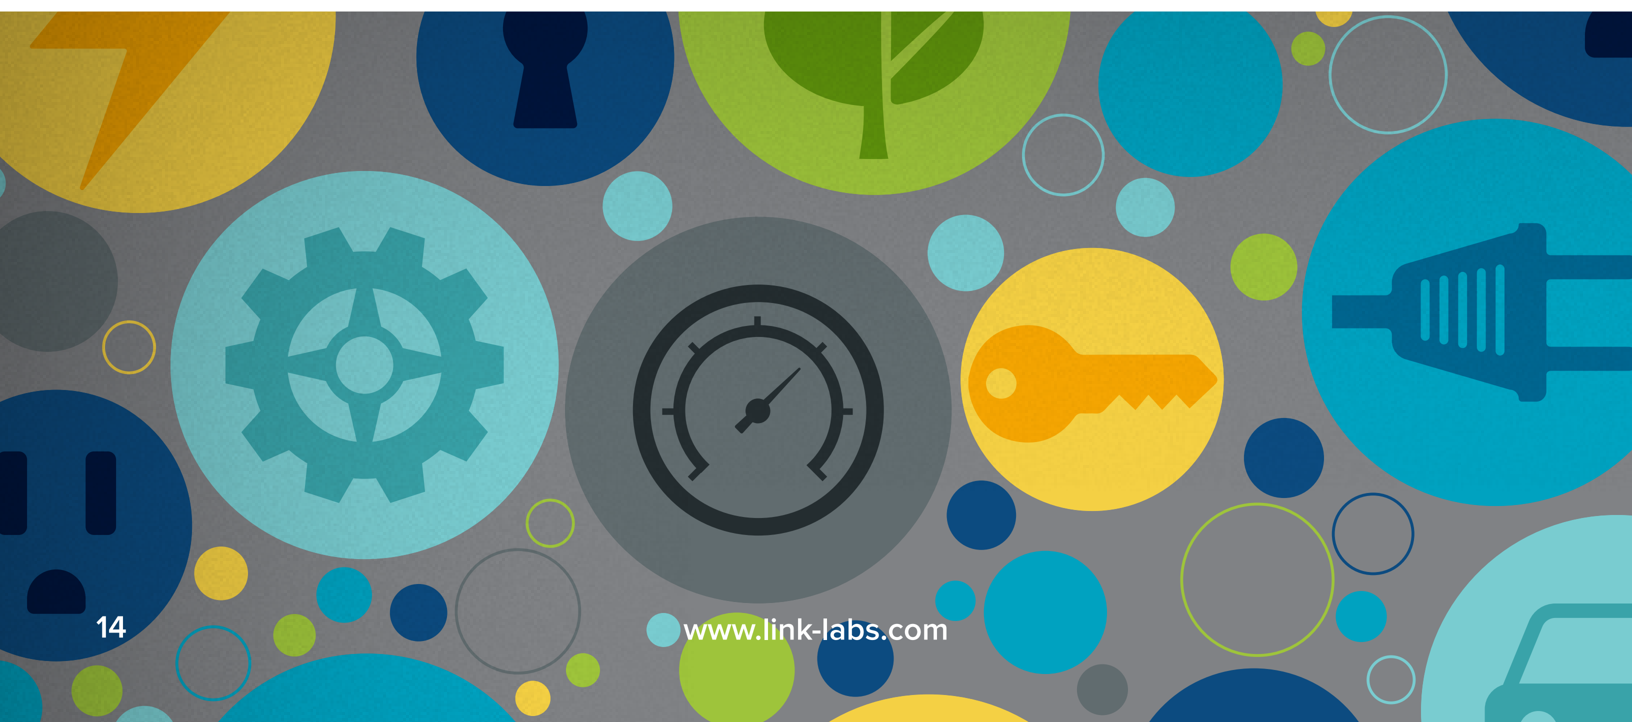

LoRa Alliance

Like Weightless, the LoRa Alliance[12] promotes an open standard for LoRa-based networks called LoRaWAN. This standard was developed by Semtech, the owner of the underlying chip technology, IBM Research, and Actility. Note: Link Labs is a founding member of the LoRa Alliance, and we use LoRaWAN technology for customers when it is a good fit.

POSITIVES

- **Its three open standards provide the end user with more choices.** LoRaWAN offers functionality that is very similar to Sigfox, making it ideal for sensor devices.
- **LoRaWAN has an open ecosystem,** meaning there's more open software and vendors available.

CONSIDERATIONS

- **The LoRaWAN standard lacks features that are important for some customers**, including roaming, packetization and retry, disconnected operations, quality of service, firmware upgrades over the air, and repeaters.
- In order to use LoRaWAN, **the network server software must be run "in the cloud,"** which requires a subscription from a network server vendor.
- **Semtech is currently the only vendor of chips,** though it has announced an agreement with STMicroelectronics to manufacture chipsets.

12. https://www.lora-alliance.org/

In Conclusion

Low power wide area networks will continue to revolutionize the way companies do business by allowing them to collect data and control devices in ways that were economically impossible before. As the technology companies described in this report begin to help companies solve their customers' problems, we're certain more resources will be invested in the space, which will lead to further advancements in LPWAN technology and applications.

Want to learn more about Symphony Link?

Try out the range calculator here.

www.link-labs.com/symphony-link/

LinkLabs

For additional questions:
info.link-labs.com/contact
+1 (202) 524-1390

Learn about our other products:
link-labs.com/

