Policy implications of scaling low-power wide-area networks over white spaces

Devin Gund*

May 2017

Abstract

Low-power wide-area networks (LPWAN) have the potential to revolutionize the Internet of Things by providing a reliable and low-cost solution for communication between embedded devices. However, the lack of currently available wireless spectrum increases interference and limits the capacity of these networks to scale. To address this limitation, I propose the integration of unused white space spectrum to boost the bandwidth and capabilities of LPWAN. Additionally, I discuss the policy implications and technical challenges of scaling LPWAN across the United States, designing mechanisms for bottom-up coordination of contributors.

1 Introduction

With the growth of the Internet of Things (IoT), connecting devices to collect and exchange data is a priority. Applications including infrastructure monitoring, precision agriculture, smart embedded systems, and emergency communications would benefit greatly from network advancements that connect potentially thousands of devices to source and consolidate information. Significant challenges remain with implementing wide-area wireless networks for low-power communication that can scale with large numbers of devices. Current mainstream wireless standards, such as Bluetooth, 802.11 (e.g. Wi-Fi), and 802.15.4 (e.g. Zigbee), generally cover less than 100 meter ranges and are not power efficient [1, 9]. To increase coverage, these technologies often utilize mesh networks of multiple transmitters and receivers, at the expense of increased cost and energy usage [2].

Recent radio advancements have driven the field of low-power wide-area networks (LPWAN), which trades off data throughput for significant increases in range. LPWAN architectures, such as OpenChirp, utilize narrow radio bands that are reserved for unlicensed communication [3]. However, larger areas of unused spectrum, known as white spaces, could potentially be used to increase bandwidth and capacity of LPWANs.

In this paper, I propose the integration of white spaces into OpenChirp and other LPWAN technologies. Building on this contribution, I discuss regulatory recommendations and policy implications for scaling and coordinating LPWAN across the United States.

*Devin Gund is an undergraduate in Electrical & Computer Engineering and International Relations & Politics at Carnegie Mellon University. Website at dgund.com.

2 Low-power wide-area networks

LPWAN has been developed as a solution for connecting devices that need to send small amounts of data over a long distance, while using very little power. Transmissions use sub-GHz spectrum, primarily the 433 MHz and 915 MHz industrial, scientific and medical (ISM) radio bands that are reserved for unlicensed use [15]. Utilization of these frequencies, rather than the 2.4 GHz ISM band used by Wi-fi and Bluetooth, leads to ranges of up to 10 kilometers and battery life between 5 and 10 years for LPWAN nodes [11]. The trade-off is that the radios operate at a lower data rate, although this is not an issue for small sensors and other devices with low data requirements [11].

2.1 LoRa

LoRa comprises the physical and data link layers used for long-range communication. LoRa radios transmit data using a form of Chirp Spread Spectrum (CSS) modulation, trading data rate for channel sensitivity within a fixed bandwidth [13]. The design implements a variable data rate, generally between 0.3 kbps and 50 kbps, which allows the system to optimize for long-range and low-power [3, 18].

2.2 LoRaWAN

LoRaWAN is an open network protocol that manages communication between gateways and end-node devices, representing the network and transport layers. LoRaWAN is responsible for pairing devices to gateways, assigning encryption keys, and managing the channel, data rate, and power for all devices. Devices in the network are asynchronous and transmit when they have data available to send [11, 13]. In general, data transmitted by a node device is received by multiple gateways, which forward the data packets to a centralized network server. This server filters duplicate packets, performs security checks, and manages network attributes. The packets may then be forwarded to relevant application servers [13].

LoRaWAN utilizes three different device classes to balance the need for energy efficiency and downlink communication latency. Class A devices are the most energy efficient, allowing for bi-directional communication through two short downlink receive windows after an uplink transmission. Class B devices open additional receive windows at scheduled times, determined by a time-synchronized beacon from the gateway. Class C devices have receive windows open almost continuously, except when the device is transmitting [13, 12].

2.3 OpenChirp

OpenChirp is a prototype end-to-end LPWAN architecture developed by researchers at Carnegie Mellon University's WiSE Lab. OpenChirp builds on LoRaWAN by adding an application interface, user management framework, services for data serialization in transmission, and improved data storage [3]. The primary objective is to simplify the user experience to add and operate new devices in the network, as well as improving performance for communities that share bandwidth and locality. The group has also developed LoRaBug, an extensible and open-source end-node running custom-built firmware to optimize memory efficiency and energy consumption. [3].

3 White spaces

The transition from analog to digital television broadcasts freed up a substantial amount of spectrum in the VHF/UHF band (i.e. 54 MHz to 698 MHz) [19]. These previously-allocated but unused frequencies are known as white spaces. In a 5-0 ruling on November 4, 2008, the Federal Communications Commission (FCC) adopted rules to allow unlicensed radio transmitters to operate in white spaces [7, 6, 5]. Similar regulations are being adopted in other countries including Canada and the United Kingdom [16].

These white spaces exist in lower frequencies than the ISM bands used for Bluetooth and Wi-Fi, resulting in significantly better propogation characteristics over long distances. The lower frequencies also allow for transmission with very little signal decay through obstacles. Importantly, white spaces are less crowded and offer more spectrum, reducing the impact of interference [16]. Altogether, white spaces exhibit extremely high potential to augment the capacity of LPWAN networks.

4 Regulations

In order to properly ascertain the policy implications of scaling LPWAN over white spaces, it is imperative to review the regulations governing devices and transmissions. In the United States, the FCC is responsible for governing the use of low-power, non-licensed transmitters [14]. This overview of regulations is not exhaustive and should be read in conjunction with Title 47, Part 15 of the Code of Federal Regulations for more information.

4.1 LPWAN regulations

Subpart C of Part 15 contains the FCC regulations for transmitters operating in the ISM bands. Transmitting devices may not cause harmful interference, and any device must accept interference received, including interference that may cause undesired operation [14, 17].

Part 15.247 defines transmission compliance through frequency hopping schemes or digital modulation schemes. In frequency hopping schemes, which are employed by LoRaWAN, the transmitter and receiver hop in synchronization between several channels in a pseudo-random fashion [13]. If the 20 dB bandwidth is less than 250 kHz, the system must use at least 50 channels, with an average channel dwell time not exceeding 400 mS within a 20 second period. If the 20 dB bandwidth is between 250 kHz and 500 kHZ, then the system must use at least 25 channels, with the average channel dwell time not exceeding 400 mS within a 10 second period. The maximum output power is 24 dBm (0.25 W) for systems with 25 to 50 channels and 30 dBm (1 W) for systems with more than 50 channels. The power in any 100 kHz bandwidth outside of the current channel must be at least 20 dB below that of the highest level within the channel [14, 17].

Part 15.249 defines limits on radiated emissions. The maximum field strength is 50 mV/m, or 0.5 mV/m for harmonic components, measured at 3 meters from the transmitting device. Radiated emissions other than harmonics must be attenuated by at least 50 dB below the level of the fundamental or to general limits [14, 17].

LoRaWAN complies with these regulations by operating sixty-four 125 kHz uplink channels from 902.3 MHz to 914.9 MHz in 200 kHz increments. There are also eight 500 kHz uplink channels in 1.6 MHz increments from 903 MHz to 914.9 MHz. The eight 500 kHz downlink channels from 923.3 MHz to 927.5 MHz. Thee maximum output power is 30 dBm, but the specification states that 20 dBm is sufficient for most devices [14, 13, 12].

4.2 White space regulations

Subpart H of Part 15 contains the FCC regulations for transmitters operating in white spaces. To determine available white spaces at a location, a device must either sense the medium before transmitting or check a cloud-hosted geo-location database, either periodically or upon moving more than 100 meters [14, 16].

Part 15.709 defines the maximum output power for white space devices. The maximum EIRP per 6 MHz for fixed devices is 36 dBm (4 W), or 40 dBm (10 W) in less congested areas, in white space TV bands, the 600 MHz service band, and the 608 MHz to 614 MHz band. For fixed devices, antenna height cannot exceed 30 meters above the ground [14].

5 Proposal: LPWAN over white spaces

Noting the considerable potential of utilizing white space spectrum to improve improve network bandwidth and capacity, I propose an extension of the LoRaWAN protocol to provide LPWAN over white spaces. In this section, I delineate the functionality of the proposed system, using the OpenChirp architecture as a reference.

5.1 Determining open spectrum

In order to minimize interference and comply with FCC regulations, transmitting devices must utilize only the white space channels available at their locations. In keeping with OpenChirp's philosophy of end-node simplicity I propose a system whereby the LoRaWAN network server manages checking the white space databases and communicating available channels to client devices. This would require the network server to know the location of all base stations and end-nodes. This not difficult, assuming the devices are fixed and the location acquired through GPS or documented during setup. The limitations of LoRaWAN prevent precise localization of end-nodes at this time [3, 11].

This system eliminates the costly need for each end-node to check a database and the potential complications of handling checks on multiple gateways. This would also reduce the requirement for listen-before-talk MAC protocols, although such functionality would still be convenient for further interference avoidance [11]. OpenChirp's current gateway design also includes a GPS radio for localization and a software-defined radio, which could be utilized for white space spectrum sensing [3].

5.2 Communicating open spectrum

If the determination of available white space spectrum is performed by the network server or gateways, this information must be transmitted to the end-nodes in order to synchronize communication frequencies. Initially, communication should occur through the default ISM channels specified by LoRaWAN, as this spectrum is guaranteed to be available for use. If there is no open white space spectrum, the ISM band may be used permanently.

If there is available white space, the gateway must notify the end-nodes to transition to that spectrum. For Class A devices, this notification must occur during one of the two downlink receive windows following an uplink transmission. Both transmissions would occur using the same channels from previous transmissions (i.e. the ISM band if this is the initial communication). After receiving the notification, the end-node would switch to the communication channel frequencies specified by the gateway, similar to the existing MAC commands for modifying channels (i.e. NewChannelReq) [12]. For Class B and C devices, this downlink notification does not need to follow an uplink transmission.

5.3 Utilizing open spectrum

Once available white space has been confirmed and communicated to end-nodes, the final step is to utilize the spectrum. Ideally the spectrum should be divided into enough channels to support LoRaWAN frequency-hopping, which would then behave in an equivalent fashion to the current system, only over different frequencies. If the spectrum cannot support enough channels, the LoRaWAN Hybrid digital modulation mode could be utilized [13].

6 Policy implications

Following my proposal for LPWAN over white spaces, it is essential to examine the policy recommendations and implications for scaling and coordinating networks.

6.1 White space policy

The FCC has benefited many industries across the United States in the first significant increase in unlicensed spectrum below 5 GHz since the 1980s [6]. As usage increases and demand for bandwidth grows, the FCC should utilize new technologies to improve spectrum efficiency of white spaces. Improved methods of calculating predicted radio interference should be employed to increase the amount of white space spectrum available while maintaining the current level of protection for licensed providers [20]. Additionally, enhanced use of detailed topographic data can also expand available white space spectrum [21].

6.2 Security

LoRaWAN transmissions are secured through end-to-end AES encryption. LoRaWAN also uses separate cryptographic keys for application and network sessions in order to diminish the fallout from an exposed key [13]. However, as LPWAN consists of many devices spread over a large area, it is difficult to secure the physical devices from all attack vectors. End-nodes are designed to be low-cost and low-power, which precludes them from containing complex security systems. OpenChirp partially resolves this with secure configuration of end-nodes over a Bluetooth connection [3]. Nonetheless, an infallible security protocol for LPWAN remains an open problem. Incorporating a secure element on-device for tamper resistant key storage would improve security. As perfect physical security cannot be achieved, end-nodes will likely remain the primary vulnerability in LPWAN. Thus, the main cybersecurity focus should be mitigating problems from compromised devices, through solutions such as continuous monitoring and remote key updates.

6.3 Scaling the network

With the ultimate goal of nationwide LPWAN use, it is important to explore the proper motivations and regulations to scale such networks across the United States. OpenChirp's target is to support an open community-driven network, with gateways and end-devices independently manageable by their owners [3]. This emerging network category is known as inverse infrastructures, emphasizing self-organization, user-driven development, and decentralized control [4]. As the network scales, community standards must be instituted for guidance. In this bottom-up infrastructure, successful mechanisms to induce coordination should should entice and unify individuals rather than impose authority. Areas in need of improvement should be listed in a centralized location to task new contributors, in a similar fashion to open-source software development. Markers of experience or expertise may also be established to elevate individuals who have invested more time or effort into the project.

LPWAN communities face unique scaling challenges in that new contribution provides negative utility after a certain point. While it is initially crucial to encourage the operation of new gateways, a high density of gateways and devices in a local area increases interference and hurts network capacity. Thus, it is extremely important to design policies and standards that incentivize contribution but disincentivize overcontribution and its negative effects.

6.3.1 Incentivizing contribution

A principal concern of self-organizing infrastructures is how to incentivize contribution and altruistic action in the interest of the entire organization. These mechanisms take the form of benefits earned for favorable behavior or costs incurred for unsatisfactory behavior.

Peer-to-peer protocols, such as BitTorrent, use both types of incentives in order to entice users to upload data to share with their peers, known as seeding, instead of solely downloading. The primary mechanism is the choking algorithm, which gives transmission preference to users who upload data at high rates. This algorithm incentivizes sharing, as peers that do not contribute are not able to attain high download rates. Free-riders who never upload are therefore penalized [10]. Additionally, reward structures, such as micro-transactions or point-based schemes, have been shown to successfully incentivize sharing [8].

FON, a wireless internet access initiative, encourages users to share their personal access points through membership benefits. Normal users pay a small fee for internet access. Network members share their access with all users and in return are allowed to use all access points at no cost. Additionally, members can collect half of the access fee that users pay [4].

An LPWAN implementation could utilize similar mechanisms to incentivize contribution to the network, specifically the operation of new gateways. Use of the network could be gated by micro-transactions or a point-based scheme, deterring free-loaders by requiring a cost to use the system. Users who operate gateways, thereby increasing network capacity and range, would be given free access and the ability to generate revenue from peer use. Importantly, the cost of a gateway should be minimized in order to encourage participation.

6.3.2 Disincentivizing interference

While LPWAN growth is dependent on incentives to operate gateways at new locations, crowded use of gateways and devices in a limited region can cause interference and must be deterred. Continuing the previous example, the community project could update a map of areas offering LPWAN access and other locations that require gateways. In areas with sufficient coverage, installing a new gateway would not grant a user with any free access or profits, effectively removing their incentive to do so. However, an individual could still invest in the operation of a gateway in an area of need, granting them with free use of the network everywhere. Together, this deters the redundant operation of gateways in the same location, but still encourages investment to further scale the network. Additionally, if it becomes necessary to limit the number of transmitting devices in a local area, demand-based pricing for network use could be implemented. This would discourage overuse of the network in any one area, spreading out load and preserving high capacity.

7 Conclusion

In this paper, I presented a proposal for the integration of white space spectrum allocation into OpenChirp and other LPWAN technologies. I also discussed the policy implications and technical challenges of scaling and coordinating LPWAN across the United States. LPWAN has the potential to reshape the entire structure of telecommunications by providing a low-cost and low-power method of long-range communication. The use of white space spectrum will greatly increase the bandwidth and capabilities of LPWAN. Future work should focus on technical challenges of implementation, specifically the feasibility of covering a variety of ISM and white space frequencies with an embedded radio.

8 Acknowledgements

This work was supported by the Carnegie Mellon University WiSE Lab and the Department of Electrical & Computer Engineering. I would like to thank Professor Anthony Rowe for providing research mentorship and access to OpenChirp materials.

References

[1] N. Baker. "ZigBee and Bluetooth: strengths and weaknesses for industrial applications". In: *Computing and Control Engineering* 16 (2 2005), 20–25(5). ISSN: 0956-3385. URL: http://digital-library.theiet.org/content/journals/10.1049/cce_20050204.

[2] Nilanjan Banerjee et al. "Relays, Base Stations, and Meshes: Enhancing Mobile Networks with Infrastructure". In: *Proceedings of the 14th ACM International Conference on Mobile Computing and Networking*. MobiCom '08. San Francisco, California, USA: ACM, 2008, pp. 81–91. ISBN: 978-1-60558-096-8. DOI: 10.1145/1409944.1409955. URL: http://doi.acm.org/10.1145/1409944.1409955.

[3] Adwait Dongare et al. "OpenChirp: A Low-Power Wide-Area Networking Architecture". In: (2017).

[4] Tineke M Egyedi, Jos LM Vrancken, and Jolien Ubacht. "Inverse infrastructures: Coordination in self-organizing systems". In: *Standardization and Innovation in Information Technology, 2007. SIIT 2007. 5th International Conference on*. IEEE. 2007, pp. 23–36.

[5] Federal Communications Commission. *Report and Order*. 2015. URL: https://apps.fcc.gov/edocs_public/attachmatch/FCC-15-99A1.pdf.

[6] Federal Communications Commission. *Second Memorandum Opinion and Order*. 2010. URL: https://apps.fcc.gov/edocs_public/attachmatch/FCC-10-174A1.pdf.

[7] Federal Communications Commission. *Second Report and Order and Memorandum Opinion and Order*. 2008. URL: https://apps.fcc.gov/edocs_public/attachmatch/FCC-08-260A1.pdf.

[8] Philippe Golle et al. "Incentives for sharing in peer-to-peer networks". In: *Electronic Commerce*. Springer, 2001, pp. 75–87.

[9] J. S. Lee, Y. W. Su, and C. C. Shen. "A Comparative Study of Wireless Protocols: Bluetooth, UWB, ZigBee, and Wi-Fi". In: *IECON 2007 - 33rd Annual Conference of the IEEE Industrial Electronics Society*. 2007, pp. 46–51. DOI: 10.1109/IECON.2007.4460126.

[10] Arnaud Legout et al. "Clustering and sharing incentives in bittorrent systems". In: *ACM SIGMETRICS Performance Evaluation Review*. Vol. 35. 1. ACM. 2007, pp. 301–312.

[11] Link Labs. *A Comprehensive Look at Low Power, Wide Area Networks*. 2016. URL: https://www.link-labs.com/lpwan.

[12] LoRa Alliance. *LoRaWAN Specification*. 2016. URL: https://www.lora-alliance.org/For-Developers/LoRaWANDevelopers.

[13] LoRa Alliance. *LoRaWAN: What is it?* 2015. URL: https://www.lora-alliance.org/portals/0/documents/whitepapers/LoRaWAN101.pdf.

[14] *Radio Frequency Devices, 47 C.F.R. § 15*. 2017. URL: http://www.ecfr.gov/cgi-bin/text-idx?node=pt47.1.15.

[15] Usman Raza, Parag Kulkarni, and Mahesh Sooriyabandara. "Low power wide area networks: An overview". In: *IEEE Communications Surveys & Tutorials* (2017). DOI: 10.1109/COMST.2017.2652320.

[16] Abusayeed Saifullah et al. "SNOW: Sensor Network over White Spaces". In: *Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems CD-ROM*. SenSys '16. Stanford, USA: ACM, 2016, pp. 272–285. ISBN: 978-1-4503-4263-6. DOI: 10.1145/2994551.2994552.

[17] Semtech. *FCC Regulations for ISM Band Devices: 902 - 928 MHz*. 2006. URL: https://www.semtech.com/images/promo/FCC_Part15_regulations_Semtech.pdf.

[18] Semtech. *SX1276/77/78/79 - 137 MHz to 1020 MHz Low Power Long Range Transceiver*. 2016. URL: http://www.semtech.com/images/datasheet/sx1276.pdf.

[19] Stephen J. Shellhammer, Ahmed K. Sadek, and Wenyi Zhang. "Technical challenges for cognitive radio in the TV white space spectrum". In: *2009 Information Theory and Applications Workshop*. 2009, pp. 323–333. DOI: 10.1109/ITA.2009.5044964.

[20] Gabriel P. Villardi et al. "Multilevel Protection to Broadcaster Contour and Its Impact on TV White Space Availability". In: *IEEE Transactions on Vehicular Technology* 66.2 (2017), pp. 1393–1407. ISSN: 0018-9545. DOI: 10.1109/TVT.2016.2566675.

[21] Gabriel P. Villardi et al. "Primary Contour Prediction Based on Detailed Topographic Data and Its Impact on TV White Space Availability". In: *IEEE Transactions on Antennas and Propagation* 64.8 (2016), pp. 3619–3631. ISSN: 0018-926X. DOI: 10.1109/TAP.2016.2580164.