UNIVERSITY OF CALIFORNIA SAN DIEGO

# Non-Traditional Infrastructure for Wide-Area Wireless Connectivity

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Computer Science (Computer Engineering)

by

Alexander Yen

Committee in charge:

Professor Patrick Pannuto, Chair
Professor Aaron Schulman
Professor Geoffrey Voelker
Professor Xinyu Zhang

2026

Copyright

Alexander Yen, 2026

All rights reserved.

The Dissertation of Alexander Yen is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2026

# DEDICATION

To anyone that has been a part of my life – I am who I am because of you in some way.

# EPIGRAPH

"I don't know how much value I have in this universe, but I do know that I've made a few people happier than they would have been without me, and as long as I know that, I'm as rich as I ever need to be."

Robin Williams

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

***To my family:***

Thank you for supporting and loving me in every way possible, even if I don't outwardly express how much I love you all. It is truly a privilege to be around a family that is constantly growing, changing, understanding, and accepting. Words cannot describe how fortunate I am to be in an environment where I have the unconditional love from each and every one of you, and I am so grateful to have gone on this life journey with you all.

To Auntie Alice and Uncle Gordon, who have been like parents during my graduate studies in California. To Christina and Phil: my closest cousins that are always there for me no matter whatever happens in life. To Steven and Mike: the silliest, troll, but genuine brothers who always have my back. To my Mom, whom I aspire to be in terms of a caring person and a partner. You're my role model and the person I look up to the most even if I don't outwardly say it, and you constantly teach me how to be a better person to family. To my Dad, whom I'm so proud of and is constantly changing and self-reflecting to treat us with even more love and care. You're the one that inspires me to never stop changing and growing in a positive way.

Finally, to NaiNai, the most admirable, caring, and lovable woman in this entire world.

Not a day goes by that I don't miss you.

***To my collaborators and mentors:***

To balex, Dhananajay Jagtap, Anthony Quiroga, Sylvia Imanirakiza, Zeal Shah, and Jason Zhang – thank you for being amazing collaborators. From walking around Hilcrest with a sketchy shoebox contraption to shipping pelican cases all around the country, spending late nights setting up ACam/stereo camera equipment to collect data, and pushing out papers together, I am so fortunate to have worked with you throughout my graduate career. Thank you for letting me grow alongside you.

To Subin Yoon, Ethan Morris, and Joe Polastre, who made my final PhD project dream come true; launching weather balloons and flying in a SR22T are experiences that not many

graduate students get to have!! I am truly so thankful to have worked with you during my graduate studies, and you have all turned my mentally burdening, demanding, and tough fifth year into something exceptionally memorable.

Finally, to my mentors Professor Goeckel, Jay Taneja, Aaron Schulman, Geoff Voelker, and Pat Pannuto.

- To Professor Goeckel, who helped me cope with early imposter syndrome symptoms and was always willing to give me pep talks.
- To Jay, who gave me a chance and lit my career in academia.
- To Aaron, who was like my second advisor and helped fuel my crazy ideas.
- To Geoff, my unpaid graduate school therapist and the undisputed GOAT in CSE.
- (To Pat, well, you'll have your spiel later.)

You all have indisputably had the most significant influence and impact on my academic journey, and you are the people I aspire to be as a mentor towards others.

***To all of my friends:***

Where do I even begin? My graduate studies have priceless memories solely because of each and everyone of you. It has been an unforgettable experience and one that is a defining chapter of my life. There are too many people to thank and so many things to talk about, but what I write about below is only the tip of the iceberg in terms of what I value between myself and each and every one of you.

A core memory of my first moments at UCSD were meeting Alex Liu and Alex Snoeren for Syslunch, socially distanced in a room in 2020, all masked up. From that first interaction, I made my first set of friends UCSD playing tennis with (Alex) Liu, George Sullivan, and Anil Yelam. It was one of the only hobbies we could do at that time to hang out and have some social interaction, and they were a defining part of my first year during Covid. After being vaccinated, hangouts with friends became exponentially more frequent, and my friendships at UCSD started to feel a lot more like home.

Jaxon Wang and Tatsuki Koga were amongst the first people that I slowly got to know over the years, initially as random roommates and now as friends. Dhananjay Jagtap, Liv Weng, Alisha Ukani, and Zac Blanco were a few of the closest friends I'd made during spring-summer of 2021, and those friendships still last strong till today. Audrey Ariniello-Randall and Zac Blanco are the fam that got me into biking, and while dying going up Mount Soledad is not necessarily a fun experience, it's one that lasts throughout time. Chris Ye, Chester Holtz, Sam Chen, and Captain Anthony Ostuni are my swimming crew mates that I've reluctantly drowned alongside with. Keegan Ryan and Stew Grant are a few of the climbing buddies were great spend time with at the gym who also never judged me for my beginner level climbing skills. Amanda Tomlinson and Stew are homies that were almost always down for a chill night of pool on a Friday night. Zhenghua Ma was my first Chinese teacher and sparked my initial confidence in speaking Mandarin more openly with other friends and people. Lydia Li and Mike Bula are like family that would always be there for me if I needed it. Katherine Izhikevich has brought a lot of fun hangouts and times in recent years. Rosie Meng is a nice person to be around who is genuinely caring towards everyone around her. Seoyoung Kweon and Ben Du are folks that brought about a lot of crazy shenanigans in my final time here at UCSD. Alex Liu, Chris Priebe, Nikolai Vogler, Daniel Spokoyny, and Shuheng Li are tennis buddies that have made my love for the sport stronger and stronger. Yibo Guo is someone who was always down to do an overnight trip to Anza Borrego for stargazing. Sumanth Rao is one of my closest friends after convincing him to continue his PhD at UCSD. Ben is also my fellow, adopted brother who is always up for any social activity; it was truly a loss when you left SD and I didn't have anyone to accompany me to Yun Tea late at night to do more work.

Lastly, I'm grateful for all of my high school friendships that continue strong till this day. Sid Limaye, Jeffrey Shao, Nonu Bajaj, Mark Guan, and Willy VVu are friendships that have lasted over a decade already and continue to stay strong against the test of time.

Each and every relationship is special to me in some form or way. I often think about you all from time to time even if we are separated by physical distance or do not talk often. Thank

you for all being there for me, and I value each and every one of you.

***To my advisor, Pat:***

Just as how you mentioned that Prabal took a chance on you, I feel the same way that you took a chance on me. I historically have suffered from issues with self-confidence; I always never felt good enough about myself no matter how hard I tried. But you're the one the gave me this chance and this opportunity to grow, and I took that opportunity. I took that opportunity to grow and challenge myself in ways I never thought would ever happen to me; I experienced and achieved goals that I thought were always impossible and unreachable; and I started to believe in myself and believe that I'm capable of defining my own path; that I'm capable of pursuing whatever direction I want to head in. *What a wild ride this has been!* I've had so many different and unique experiences, and this is all in part due to you and this chance you've given me. Things were definitely not all sunshine and rainbows, but is any graduate degree ever that way? Especially since the last 1-2 years of my PhD have been the most mentally toughest parts of my grad school career, it was definitely hard to remember all those priceless times when the last moments felts so tough and challenging. But after reflecting back on my time here at UCSD, the journey, experience, and memories I've had (and made) is irreplaceable.

Thank you.

***Paper Acknowledgements:***

Chapter 3, in full, is a reprint of the material as it appears in Dhananjay Jagtap*, Alex Yen*, Huanlei Wu, Aaron Schulman, and Pat Pannuto. Infrastructure: Usage, Patterns, and Insights from "The People's Network". In Measurement Conference (IMC '21), November 2021. The dissertation author was the primary investigator and author of this paper.

## VITA

2020 Bachelor of Science, University of Massachusetts Amherst

2023 Master of Science, University of California San Diego

2026 Doctor of Philosophy, University of California San Diego

## PUBLICATIONS

Dhananjay Jagtap*, Alex Yen*, Huanlei Wu, Aaron Schulman, and Pat Pannuto. 2021. Federated Infrastructure: Usage, Patterns, and Insights from "The People's Network". In ACM Internet Measurement Conference (IMC '21), November 2021.

## FIELDS OF STUDY

Major Field: Computer Science (Computer Engineering)

Studies in Computer Science and Engineering
Professors Patrick Pannuto, Aaron Schulman, and Geffrey Voelker

Studies in Eletrical and Computer Engineering
Professor Xinyu Zhang

ABSTRACT OF THE DISSERTATION

Non-Traditional Infrastructure for Wide-Area Wireless Connectivity

by

Alexander Yen

Doctor of Philosophy in Computer Science (Computer Engineering)

University of California San Diego, 2026

Professor Patrick Pannuto, Chair

LoRa is a new wireless technology that enables applications requiring long range, low power, low data-rate connectivity. Although LoRa has gained widespread academic interest, LoRa struggles to gain widespread, societal attention and adoption a decade after its inception. Arguably, the problem with the wireless technology's adoption is that it is not clearly differentiated from existing wireless technologies. That is, alternative wireless technologies today (e.g. Wi-Fi, cellular, Bluetooth Low Energy) already—or will—achieve what LoRa currently has to offer.

This dissertation claims that LoRa requires effective infrastructure deployment to realize the potential of LoRa technology and to deliver new applications. Specifically, deploying LoRa

transceivers on non-terrestrial infrastructure (e.g. airplanes, unmanned balloons) can enable a new realm of wide-area, low power applications that can span entire metropolitan areas at fractions of the cost of traditional infrastructure. This dissertation posits that the ability of LoRa to achieve (very) long ranges of communication (i.e. $>$100km) is an undervalued aspect of LoRa that can be exploited via non-terrestrial infrastructure.

# Chapter 1

# Introduction

LoRa is a new wireless technology that is meant to serve low-power applications while enabling **Lo**ng **Ra**nge communication. The tradeoff with these two benefits regards a reduction in data bandwidth – the LoRa technology adopts the Chirp Spread Spectrum (CSS) modulation scheme, essentially trading off redundancy (i.e. less throughput/goodput) for range. This technology was specified to cater towards new applications that do not require high data throughput but instead value low power connectivity. One major class of applications involves the Internet of Things (IoT)—networked devices, sensors, appliances, or *Things* in general that can be controlled via wireless communication—in which applications might range from smart homes (e.g., smart locks, cameras, bulbs, thermostats) to smart cities (e.g. smart trash cans, meters) and even smart monitoring (e.g. wildfire detection, earthquake detection). IoT serves as an opportunity to extend and incorporate the fluid usage of networked devices within our everyday lives, customizing and augmenting the ability by which we can interact with our environments. LoRa is envisioned as one solution to low power, wide-area connectivity that can enable the deployment of more IoT devices within our world.

However, in the time since LoRa's inception and standardization into the LoRa Wide Area Network (LoRaWAN) protocol in 2015, the expected explosion of new IoT applications and devices is orders of magnitudes less than expected. A critical problem that prevents this explosion regards the "gateway problem" – IoT devices require convenient, low power access to

a gateway (i.e. a receiver node that is connected to the Internet). But unfortunately, "The Internet of Things Has a Gateway Problem" [90], "Still Has a Gateway Problem" [88], and continues to have a gateway problem to this day. While sources estimate(d) one trillion IoT devices by 202X or 203X [17, 71], not even approximately twenty billion IoT devices exist globally today [59]. A decade later, LoRa has not appeared to have significantly solved the gateway problem.

A fundamental problem is that LoRa does not—and will not—replace current-day wireless technologies that already support a wide range of IoT applications. In the case of a smart bulb, Wi-Fi is suitable to control these home-based devices that do not have low-power requirements (i.e. the bulbs are directly plugged into sockets). On the other hand, smart meters—which often collects and transmits utility data (i.e. electricity, water) in real-time—have low-latency and reliability requirements, and cellular is low power (enough) to be suitable for this application and coverage already exists today [74]. While LoRa could could also be used for smart bulbs, smart meters, and similar types of applications, existing applications already rely on existing wireless technologies to function. Furthermore, the onset of newer cellular protocols (e.g. New Radio RedCap [73]) that are specifically cantered towards cheap, low power, reliable, low latent, and pervasive connectivity is likely to outcompete LoRa in a race to low power, wide-area connectivity. It is difficult to say whether LoRa is an overwhelming contender to solving the gateway problem and low power, long range connectivity. The question then becomes, where should we leverage the LoRa technology and what are the new set of applications uniquely enabled by LoRa that extend into and beyond IoT?

A core insight with respect to the usefulness of a wireless technology regards the necessary, supporting infrastructure. That is, if the cellular technology were not built in a methodical way with billions of capital spent on towers and radio equipment, cellular would not be the reliable and usable technology we know today. While a country-wide LoRaWAN can be achieved at similar cost, there is no clear application, use case, or consumer for such a scaled network that would amortize the upfront costs to build such a network. Additionally, the intended application space (i.e. IoT) is meant to be relatively cheap, and an expensive and thoroughly

built LoRaWAN—that is dedicated to serve cheap applications—is counterintuitive. Hence, the supporting infrastructure for LoRa needs to be (1) cheap and (2) must serve connectivity that differs from existing wireless technologies (e.g. WiFi, cellular, Bluetooth Low Energy).

## 1.1 Insights

Two undervalued, unexploited characteristics about LoRa are (1) the ability to provide ad-hoc, cheap coverage and (2) the ability to provide exceptionally wide-area coverage (e.g. $>$10,000km$^2$ or 4,000mi$^2$). I discuss how the combination of these two characteristics opens up new opportunities and applications for the LoRa technology.

### 1.1.1 Ad-Hoc Infrastructure

Ad-hoc LoRaWAN infrastructure and coverage is achievable today as shown by The Things Network (TTN) [80] and The Helium Network [29], two crowdsourced LoRaWANs that provide general-purpose, global, LoRa-based coverage. In either network, arbitrary participants (i.e. people) can purchase transceiver equipment for just a few hundred dollars and set up their own device to participate within TTN or the Helium Network. In essence, these networks capitalize on the fact that any person can purchase a LoRaWAN kit [5, 2, 6, 8] and establish their own, personalized LoRaWAN. Additionally, not only is LoRaWAN equipment affordable to purchase and relatively easy to set up, a LoRaWAN's core network is significantly simpler to build than that of cellular's core network. The simplicity of a LoRaWAN compared to a cellular network (i.e. Radio Access Network) is an intentional design choice and a clear advantage over the cellular technology when it comes to realizing real-world, wide-area coverage without a large, central authority. A real-world application that relies on this ad-hoc feature of LoRa pertains to smart agricultural monitoring, in which devices (e.g. soil moisture sensors, smart irrigation) are integrated into crops to better automate and effectively promote crop health and yield. For areas that lack (reliable) cellular coverage, ad-hoc LoRa infrastructure can be a potential solution accessible and practical for individual operators. This capability to enable ad-hoc infrastructure

is a key, under-discussed characteristic that differentiates LoRa from infrastructures of other wide-area wireless technologies.

### 1.1.2 Long(er) Range

Another under-appreciated value is how LoRa can transmit over not just long-distances (a well-discussed, defining feature of LoRa) but over *very* long distances. Traditionally, Semtech and many other researchers [49, 56, 60] estimated a theoretical range of roughly 5km in urban settings and 15km in rural settings. This range is the same order of magnitude as cellular, and thus it appears that LoRa does not provide a significant range advantage over cellular. However, particular cases have shown that LoRa's transmission ranges can reach up to 832km [75] and 1336km [76] [1]. These transmission distances were achieved with infrastructure decisions that maximize uninterrupted line-of-sight (LOS) communication, specifically via unmanned balloon and over seawater. While the achievable ranges of approximately 1000 kilometers likely pushes the boundaries of LoRa's normal communication range, the (nearly) 100x improvement in communication range questions whether we are deploying LoRaWAN infrastructure *inefficiently*.

While LOS communication is a well understood and recognized desire for wireless communication—i.e., less obstruction implies less signal degradation—this universal understanding is not currently adapted to be deployed with LoRa. As a result, one of the major observations and hypothesis that this dissertation targets is that a primary, under-acknowledged strength of LoRa lies in the fact that in comparison to other wireless technologies, LoRa achieves a disproportionately significant gain if communication range in LOS versus non-LOS (NLOS) conditions.

### 1.1.3 Non-Terrestrial Infrastructure for Long Range Communication

Because of the lower resource (e.g. energy and compute consumption), requirements for base stations, LoRa is more easily and readily deployable on a variety of non-terrestrial

[1] Unfortunately, these cases did not discuss whether modifications to the protocol were made (e.g. transmit power, coding gain)

infrastructure today. These readily available infrastructures include unmanned balloons (e.g. weather balloons) as well as airplanes (e.g. commercial airplanes), and they operate at altitudes with LOS access to a broad area of land, serving as potential carriers for ad-hoc infrastructure that can enable long(er) range communication. Furthermore, these infrastructures are accessible. They are either cheap to deploy (i.e. an unmanned balloon costs less than $100 to launch) or already exist (i.e. commercial airplanes already fly millions of people a day). There are certainly challenges to deploying on either vehicle—unmanned balloons require scheduled Notice to Air Missions (NOTAMs) to the Federal Aviation Administration (FAA), and logistically deploying equipment and verifying lack of radio frequency (RF) interference is non-trivial. However, a step forward into this space requires us to empirically measure how well the LoRa technology would work on these non-terrestrial vehicles.

## 1.2 Thesis

I defend the following thesis: *By deploying LoRa on non-terrestrial infrastructure, it is possible to (1) economically and rapidly enable new low power, pervasive applications that (2) span entire metropolitan-wide coverage from a single transmitter with out-of-the-box configurations.*

I first discuss our general observations that support the case for LoRa's ad-hoc deployment. In this work, I measure the performance and efficacy of LoRa in NLOS and LOS scenarios with respect to the Helium Network. To do so, I deploy transmitters on ground vehicles (e.g. car, truck) and aerial vehicles (e.g. balloon, airplane) to measure the efficacy of NLOS and LOS conditions, respectively. We purchase a commercial hardware kit that is equipped with LoRa-compatible communication capabilities and deploy many nodes into the real-world for wide-area measurement purposes.

I then empirically demonstrate that we can enable LoRa coverage to entire, metropolitan areas by deploying on non-terrestrial infrastructure (e.g. unmanned balloons, airplanes) that

allows LoRa to communicate in LOS conditions over many hundreds of kilometers of distance. By deploying on these non-terrestrial infrastructures, we can enable longer transmission distances than we currently expect out of the technology. I.e., current transmission ranges via ground-to-ground communication lie between 5-15 kilometers (approximately 3-9 miles), versus air-to-ground is observed to reach at least 1,000 km [76]. By increasing the range of communication, we can thus achieve greater scale and broadness of coverage.

Lastly, I explore how deploying transmitters on these non-terrestrial infrastructure (e.g. balloons, airplanes) allows us to achieve broad coverage at orders-of-magnitude less cost than existing infrastructure (i.e. terrestrial towers or satellites). The primary observation is that cheap-to-deploy-on non-terrestrial infrastructures exist today. Unmanned balloons are already pervasively used today and are cheap, disposable infrastructure – the National Weather Service (NWS) deploys weather balloons in various parts of the United States twice a day to collet atmospheric data for our weather forecasts. So long as unmanned balloons follow Federal Aviation Administration (FAA) regulations with proper Notice to Air Missions (NOTAMs), they can be deployed in an ad-hoc fashion for less than $100 per deployment setup. Similarly, commercial airplanes fly at all times of the day (although the density or quantity of airplanes in the sky typically exhibits a diurnal effect). Commercial airplanes are another piece of underutilized—relatively cheap—infrastructure that can enable LoRa to transmit in LOS conditions. As such, non-terrestrial infrastructure exists today that can cheaply enable LoRa at long ranges of communication.

## 1.3 Challenges

In order to empirically measure the performance of LoRa on non-terrestrial infrastructure, we must either build a metropolitan-wide LoRaWAN or have access to one that is available. The upfront costs to build a wide-area LoRaWAN are non-trivial, particularly if the primary use case is just for empirical measurement. This is where a crowdsourced network (e.g. The Helium

Network, The Things Network) can be an opportunistic solution towards empirically measuring a distributed LoRaWAN without demanding the upfront costs to build such a measurement testbed. In essence, these distributed networks allow arbitrary people to purchase equipment and deploy the network themselves, typically with receiver equipment residing within a person's home. By allowing others to purchase and build the network, the network entities (that maintain and control the network core), can leave the focus and burden of deploying base stations to random users. Within the United States, we measure using the Helium Network for the work in this dissertation because of its large availability within the country, providing a potential solution towards evaluating LoRa at scale.

However, reliably estimatin ground truth is one challenge and limitation of this work. The Helium Network is composed of various participating Helium network operators, and the veracity of the network coverage depends on the behavior of these actors. The behavior of actors can range from (1) malicious and deceiving (i.e. those that are not truthfully providing LoRa coverage to the network) to (2) honest and well-behaved participants (i.e. those that participate ethically well), to (3) well-intentioned but ineffective (e.g poor placement of radio frequency hardware), to (4) forgetful and inactive. While a crowdsourced LoRaWAN enables a wide-scale measurement opportunity on the LoRa technology, we cannot control the behaviors of actors and the validity of the extant infrastructure. The Helium Network is aware of this problem and deploys mechanisms that minimize malicious—or falsely but innocuously reporting) users. As a result, the Helium Network already has policies that attempts to verify the existences and participation from various users. However, this also implies that our results will never reflect *exactly* the ground truth of what we measure. Chapter 5 discusses how we extrapolate a measurement instrument from this uncontrolled, physical testbed, as well as how we account for its limitations in our results and analysis.

## 1.4 Contributions

There are three major contributions in this work, based largely on the measurement of LoRa in different infrastructure contexts. These three contributions pertain to a deeper understanding in the (1) working of crowdsourced LoRaWAN infrastructure (i.e. The Helium Network), (2) an evaluation of LoRa and the Helium Network with respect to other wireless technologies and their corresponding infrastructures, and (3) measurement of LoRa (via Helium) on non-terrestrial infrastructure.

To preface the contributions, I develop a simple, measurement framework that can measure LoRa in a variety of communication environments; this framework is built on top of the Helium Network, which is a crowdsourced LoRaWAN that is widely deployed in the United States by arbitrary users. In our framework, we deploy a transmitting device that will send LoRa packets to the receiving Helium network. After the network receives our packets, we can then retrieve the received packets for analysis. The Helium network enables a large scale measurement opportunity, in which we can broadcast packets to all nearby Helium base stations and deduce insights from the retrieved data.

**In Chapter 2, I discuss background knowledge on the LoRa technology as well as related work.** Background knowledge includes details on the modulation scheme, Chirp Spread Spectrum, which elucidates understanding on why LoRa is able to transmit over very long links. The related work discusses previous analysis of crowdsourced LoRaWANs and various efforts to measure LoRa in different environmental contexts.

**In Chapter 3, I discuss initial findings and insights from a crowdsourced LoRaWAN, the Helium Network.** I look into how well Helium's crowdsourced infrastructure fares as an alternative to expensive, planned infrastructure (e.g. cellular). In this chapter, we dive into the network topology, functionality, usage, and efficacy of this distributed LoRaWAN. By using their own cryptocurrency—the Helium Network Token (HNT)—to incentivize user participation to provide network coverage, the Helium Network adopts and interesting infrastructure model

to provide wide-area LoRa coverage while also providing unique data on the underlying performance of the network. The major points of discussion regard (1) whether the infrastructure model works and can function as a viable LoRaWAN, (2) insights towards network growth and usage, and (3) a small scale measurement study on the performance and reliability of the network.

**In Chapter 4, we extend beyond our initial measurements of the Helium Network to a country-scale measurement comparison between LoRa, cellular, and BLE wireless technologies as well as (more importantly) how well and full the supporting infrastructure works cross country.** Specifically, the corresponding, supporting infrastructures to cellular, BLE, and LoRa wireless technologies that we measure are Verizon, Apple Find My, and Helium. We physically pack a smartphone, an ESP32 BLE device, and a Heltec LoRa transmitter that send out cellular, BLE, and LoRa packets at roughly identical timing intervals (i.e. 10 minutes) from a single case. This case is shipped to various points across the country via truck (i.e. FedEx), which allows us to evaluate the performance and reliability of each wireless technology (and corresponding infrastructure) in a variety of different built environments. We measure the packet reception rate (PRR) of each wireless technology for a particular shipment (i.e. San Diego to New York) and use that as the main metric to compare between the performance and reliability of each wireless technology and their supporting infrastructure.

**Lastly, Chapter 5 distills the performance of LoRa on non-terrestrial infrastructure.** We use similar measurement frameworks as done so in Chapter 3 and Chapter 4 and deploy LoRa transmitters on balloons and an airplane. This allows us to test our hypothesis of LoRa's performance in LOS conditions, which enables us to understand the performance of LoRa in these environments. This leaves us with answers regarding the potential usefulness (or lack of usefulness) for LoRa in these scenarios. Evaluating LoRa's performance within the context of non-terrestrial infrastructure allows us to answer this dissertation's thesis and whether LoRa on non-terrestrial infrastructure enables future opportunities for not only LoRa's pervasiveness to grow but also newer applications that cannot exist with existing wireless technologies and their

corresponding infrastructures.

# Chapter 2

# Background and Related Work

This chapter first discusses background related to the LoRa technology, specifically the modulation scheme, Chirp Spread Spectrum. Then, it dives into the related work with regards to (1) the Helium Network and measuring a crowdsourced LoRaWAN, (2) cross-country measurements and comparisons between LoRa and the Helium Network versus other wireless technologies and corresponding infrastructures, and (3) LoRa deployment on non-terrestrial infrastructure.

## 2.1 Background on LoRa Modulation: Chirp Spread Spectrum

Chirp Spread Spectrum (CSS) is a modulation technique that encodes data into symbols (or chirps) by which physical bits (or chips) are sent in an increasing or decreasing sweeping frequency within a given bandwidth; this modulation technique is adopted by LoRa. To explain modulation scheme in more detail, this section will first run through a detailed example. An upchirp (with no frequency offset) is a symbol that is increasing between $f-(bw/2)$ to $f+(bw/2)$, where $f$ represents the center frequency, and $bw$ represent the bandwidth. If the center frequency is 903.9MHz (i.e. $f = 903.9MHz$) and the bandwidth is 125kHz (i.e. $bw = 125kHz$), then $f-(bw/2) = 903.775MHz$ to $f+(bw/2) = 903.025MHz$. Then, the frequencies between $903.775-903.025MHz$ will be discretized into steps, which can be chosen

between $2^7$ steps to $2^{12}$ steps; for this example, $2^7$ steps will be chosen. To encode data with each upchirp, an offset frequency is applied to the center frequency. This offset frequency is what encodes the actual data for a symbol. In this case of $2^7$ steps, the offset frequency is calculated by $[Symbol\_Value] * bw/128$, which is the bandwidth divided by the total number of steps. The actual symbol value is encoded with a value that ranges between 0 and $2^7$. A frequency offset for a symbol value of 0 would thus be calculated as $0 * 125kHz/128 = 0$, thus meaning that a symbol value of 0 would correspond to 0kHz. The resulting upchirp would sweep from $f - (bw/2) = 903.775MHz$ to $f + (bw/2) = 904.025MHz$, again discretized into $2^7$ (i.e. 128) distinct frequency steps. For a symbol value of 100, the resulting offset frequency would be calculated as $100 * 125kHz/128 = 97.65625kHz$. The upchirp sweeping would thus start at $903.775 + 0.09765625 = 903.87265625MHz$ and sweep to the upper bound of $f + (bw/2) = 904.025MHz$. Then, the frequency sweep would wrap back around from $f - (bw/2) = 903.775MHz$ and sweep up to the frequency offset of $903.87265625MHz$. (Without going into the details for receiver demodulation, a receiver would be able to demodulate and decode this upchirp to retrieve the symbol value of 100.) This example illustrates a real-world example of how CSS modulation works and how data is encoded into each symbol.

Using the example above, this paragraph then discusses the relevant terminology for understanding how CSS enables long range communication. Spreading Factor (SF) is what dictates the frequency step size. A SF value of 7 would correspond to $2^7$ steps, and a SF value of 12 would correspond to $2^{12}$ steps; the official name for a step is a "chip." As such, each SF value encodes the SF value's number of bits into a symbol; a symbol encoded via SF7 will encode 7 bits of information into a symbol, whereas a symbol encoded via SF12 will encode 12 bits of information into a symbol. However, the real purpose of SF allows for signal robusticity – each step (or chip) is allocated the same amount of time, so an upchirp sent with $2^7$ steps will take less time to send than an upchirp with $2^{12}$ steps. However, because there are more physical chips sent with SF12, there is introduced redundancy into our signal and are reducing the goodput of our symbol. That is, the ratio of encoded bits per symbol is $7/128 \approx 0.0547$ for SF7 and

$12/4096 \approx 0.00293$ for SF12.

The CSS modulation technique portrays a tradeoff: a lower SF (e.g. SF7) uses less physical bits to send a symbol but enables higher goodput and faster datarate, whereas a higher SF (e.g. SF12) enables better symbol robusticity by introducing more physical chips per each symbol. Hence, what this tradeoff really underscores is redundancy for reliability versus higher throughput. That is, more redundancy will enable more robustness, whereas less redundancy can allow the signal to achieve higher throughput.

The bulk of this dissertation is more focused on enabling LoRa at longer ranges. That is, this dissertation focuses on how to provide wireless coverage to more physical area hence focuses on measuring LoRa via the highest SF value. In the United States, the highest SF value is SF10; SF11 and SF12 are illegal due to Federal Communication Commission (FCC) regulations on maximum dwell time [57].

## 2.2 Related Work

There is limited related work on (1) measuring and analyzing crowdsourced LoRaWANs, (2) a comparison of wireless technologies and infrastructures for wide-area wireless connectivity, and (3) deployment and measurement of LoRa on non-terrestrial infrastructure.

### 2.2.1 Analyzing Crowdsourced LoRaWANs

For related work on analyzing crowdsourced LoRaWANs, the only measurement study prior to this dissertation's measurement study on the Helium Network is a study on The Things Network (TTN) in 2017. Blenn et al. collect data from TTN gateways but focus their analyses on LoRa-PHY specific properties [15]. As such, this dissertation presents a first measurement study on a crowdsourced LoRaWAN, the Helium Network.

### 2.2.2 Comparison Study of Different Wireless Technologies and Infrastructures

Elkhodr et al. discuss the tradeoffs of low-power, wireless technologies for IoT applications with regards to application requirements and how various technologies might fit these constraints [24]. They compare a wide range of factors (e.g. power, range), however these comparisons are largely theoretical and qualitative. For instance, LoRa's power is regarded as "very low," whereas BLE's power is higher than LoRa's at "low." Montori et al. similarly compare a wide range of wireless technologies for IoT applications, although the analysis is again more similar to a survey [53]. There is a major focus in the explanations of each technology but little or no direct, measured comparison. Again, this survey is primarily based on theoretical assumptions.

### 2.2.3 LoRa on Non-Terrestrial Infra

#### Opportunistic Wireless Comms from Aircraft and Balloons

Ahmad et al. were among the first to propose aircraft providing wireless networking coverage with Wi-Fly [10]. This work bolsters the potential of this dissertation's work, as it shows airplanes can be a viable platform for low-power wireless communication, and their ADS-B data show that aircraft can provide opportunistic WiFi coverage spanning entire continents. This work shows that flying platforms are particularly well-suited for providing timely emergency broadcasts, and the LoRa protocol provides much larger coverage area of $\sim$30 km compared to the $\sim$2 km they encountered with WiFi. Another project that inspired this work is Iizuka et al. that propose using balloons for airdropping sensors to measure various atmospheric data (e.g. eye of typhoon) as an alternative to expensive, high altitude platform stations (HAPS) [40]. In their study, they launch a single balloon that flies over the ocean 100 km away, transmitting and receiving data via LoRa to a device on the ground, using the ground-to-balloon link as a way to transfer atmospheric data as well as a signal to airdrop the sensor. Similarly, Mirza et al. use LoRa to track a balloon's location while collecting environmental data [51]. Both of these

prior works provide motivation for exploring LoRa's capability as a broadcast service because of the long-range links from aircraft they could achieve. However, both only study point-to-point links, rather than evaluating the wide-area broadcast coverage which is the focus of this work. Project Loon demonstrated high-altitude balloons could provide wide-area cellular coverage [83]. However, their wireless links were established with complex cellular base stations, whereas this dissertation shows that a low-cost LoRa transmitter can provide broadcast coverage to at least 1,200 $km^2$. Also Loon's larger payload required much larger specialized balloons, while this work focuses on using low-cost weather balloons for rapidly deployable wide-area broadcasting.

**Measuring Terrestrial LPWAN Coverage**

The Things Network (TTN), a crowdsourced LoRa IoT deployment, was first measured in 2017, in which Blenn et al. analyze the uses of transmitting nodes and receiving gateways participating in TTN [16]. They look at various LoRa factors such as Spreading Factor usage, payload sizes, Signal-to-Noise Ratio (SNR) distribution, and more. Jagtap et al. conducted the first measurement study on the Helium Network in 2021 [41], elucidating the incipient workings of the network as well as performing a small scale measurement study in urban and sub-urban environments, measuring coverage availability and reliability. Following their measurement study in 2023, Tong et al. deploy a city-wide scale study of LoRa—*CityWAN*—by deploying 100 receiving gateways and nearly 20,000 transmitting nodes [79]. They find that coverage reliability can still be inconsistent despite expected coverage. More importantly, they note that while forty critical gateways can cover 95.3% of end nodes, 4.7% of nodes—roughly 900 nodes—need sixty more gateways. However, they mention that these 900 nodes are in blind spots with nodes deployed in indoor scenarios or are wirelessly shielded from signals. These works demonstrate all of the aspects of LoRa that can affect coverage and reliability, and they all show that LoRa has lower than expected coverage area when used only for terrestrial links. The line-of-sight coverage provided by an aircraft overcomes many of the problems that were experienced in these prior studies, however even broadasking from the sky still does not provide complete coverage

and reception rate, and it is likely due to the same factors investigated in this prior work.

**Measuring Satellite LoRa Coverage**

There is also a wide variety of work studying LoRa communication between ground stations and satellites. However, this line of work is differs with ours such that (1) satellites in the LEO orbit must address the Doppler affect from high-speed movement or (2) satellites in the MEO and GEO orbit must transmit at high power (i.e. watts) to reach terrestrial, ground stations. TinyGS [9] is an example application of this work, today, which uses LoRa-based nodes as ground stations to satellite telemetry messages. However, LoRa is not transmitting from the ground back to these satellites, and satellites usually need to transmit on the order of watts for the nodes on the ground to receive data. A major difference between this work and existing LoRa-satellite work is that this paper measures how well LoRa works straight out of the box with no modified configurations.

# Chapter 3

# Helium: Federated Infrastructure

## 3.1 Background

This chapter discusses the operations, functionality, and behavior of the Helium Network. More specifically, this includes information on the network topology and functioning components, network decentralization, coverage volatility and dependencies, and coverage simulations and empirical measurements. These initial insights provide context on the validity and viability of the Helium Network for further measurements in later chapters.

## 3.2 Helium – Quick Overview

Helium is a new wireless provider. The initial focus of Helium is on building crowd-sourced hotspot infrastructure to provide wide-area LoRa coverage. LoRa is one of several new radio technologies competing to provide coverage for low-power, low-bandwidth edge devices (IoT-class devices). In contrast to traditional wireless infrastructure, Helium does not own the deployed LoRa base stations. Rather, individuals purchase, deploy, and maintain Helium hotspots in exchange for payment for coverage and for the data that their hotspot ferries. The Helium company built and sold the initial hotspots, maintains firmware for new third-party, mass-produced hotspots, develops supporting cloud infrastructure, and develops the blockchain that underpins the network.

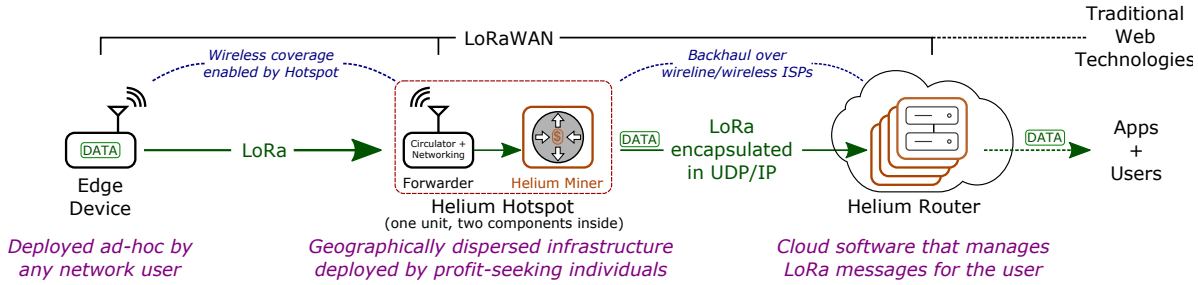


**Figure 3.1.** Helium Overview. Edge devices broadcast packets over the LoRa PHY. One or more nearby infrastructure nodes—hotspots—recover this packet. Helium Routers pay hotspots to release packets from edge devices they wish to recover data from. Routers must be reliably available as the LoRa protocol requires that any downlink responses from the router arrive at the edge device within one second.

### 3.2.1 Helium from a Basic User's Perspective

This subsection explains how Helium works for end users. To deploy a LoRa device that uses the Helium network, such as a sensor, a user first registers a new application with the Helium Console (a Helium-provided cloud service that acts as a LoRaWAN router plus Helium wallet). Users could deploy their own Console equivalent, but as shown later in this chapter's analysis, the vast majority do not and simply rely on the Helium Console. Next, users deposit money in their Console account to pay for future data. Then, they must register a new device with the Console, which gives configuration parameters for the device's networking stack (blindly copied `#defines` prepended to a Helium library). After this, the device is ready to deploy: it can now send and receive packets from the Helium network. In the field, packets sent by devices are received by hotspots, who forward their data to the Console, which pays hotspots and provides packet payloads to application users via HTTP (or numerous other means).

### 3.2.2 Components of the Helium Network

Summarized in Figure 3.1, the Helium network is built on top of the the LoRaWAN network architecture. LoRaWAN is a cloud-based protocol that routes LoRa packets between wireless devices and cloud services. The primary difference from LoRaWAN is that Helium makes it possible to have crowdsourced hotspots to bridge cloud services with the LoRa wireless

network; LoRaWAN requires application owners to deploy and operate their own hotspots. Helium adds routing by overloading identifiers in LoRaMAC that normally identify the device and its owning application.

The following describes each component of the Helium network:

**Routers** are cloud servers which manage the LoRaWAN protocol across devices and applications. They are responsible for authenticating devices and receiving messages from devices. Hotspots find Helium-compliant routers by looking up device owners using packet metadata and a filter list in the Helium blockchain (in contrast to standard LoRaWAN, where gateways have one, statically configured router). Helium routers must also negotiate with hotspots to pay for data. This is done via the `state_channel` mechanism described later. Anyone can host their own Helium compatible LoRaWAN router by purchasing device and application identifiers via a transaction on the blockchain.

**Edge Devices** are the LoRa-enabled end wireless devices. They can be embedded in products such as sensors, collars, or tags and can be used for a variety of applications including environmental, pet, or asset monitoring. Devices are pre-provisioned with a Device End User Identifier (EUI), an Application EUI, and an App key. These are used during Over The Air Activation (OTAA), part of LoRaWAN, to authenticate to a LoRaWAN Router (i.e., a cloud service). Helium differs from LoRaWAN in that EUIs are first used to look up Helium Organizationally Unique Identifiers (OUIs). These determine which router to authenticate to.

**Packet Forwarders** are LoRa wireless modules in hotspots. These use the industry-standard Semtech packet forwarder [67] to relay edge device radio packets to and from the co-resident miner in the hotspot (in regular LoRaWAN, these are forwarded directly to the owning router; the miner is new from Helium). They are generally made up of a LoRa concentrator, which is a radio transceiver capable of operating on multiple parallel subbands and a supporting processor which runs the packet forwarder.

**Miners** transmit encapsulated LoRa packet payloads to and from Helium routers and

maintain the Helium blockchain. A miner is usually a small embedded Linux device.[1] Miners receive Helium tokens (HNT) as rewards for device data transit, network coverage validation, and blockchain consensus activities. They send traffic to LoRaWAN routers on the internet at large using any UDP/IP capable backhaul-like wireline (e.g., Cable and DSL) or wireless (e.g., LTE) Internet access networks.

**Hotspots** are physical boxes with a packet forwarder and a miner. Hotspots make up the majority of deployed Helium infrastructure to date.[2] In principle, forwarders and miners could be separated, as the connection is simply an IP link. In practice, they are co-located to overcome a short-term engineering limitation:

> "The protocol between the gateway and the server is purposefully very basic and for demonstration purpose only, or for use on private and reliable networks. There is no authentication of the gateway or the server, and the acknowledges are only used for network quality assessment, not to correct UDP datagrams losses (no retries)."[3]

**Helium Console** – The Console is a cloud service provided by Helium which acts as a Helium router. Instead of buying an OUI and running their own router, users deploying end devices can opt to use the Console. At this time, the Console charges users wholesale for data communicated with sensors and has no overhead cost.

**Validators** are a new class of Helium node, currently in testing, that will take over blockchain maintenance events from miners. The proposal (HIP25) to introduce validators was ratified in January 2021. Validators solve well-known scalability and robustness challenges for blockchains and use well-known staking techniques. As validators were not yet live during the period under study (through May 2021), they are not considered extensively. We note validators here as they appear as special-case miners on the blockchain and will influence some of the

[1] e.g. in the Rak Wireless v1.5 hotspot, the miner is an off-the-shelf Raspberry Pi 4 Model B – see https://www.youtube.com/watch?v=BUrAb2GoHhc for more detail.

[2] Very early, Helium permitted a small batch of "DIY Hardware." New hotspots must be from certified vendors to earn rewards on the network.

[3] From https://github.com/Lora-net/packet_forwarder/blob/master/PROTOCOL.TXT

analysis of deployed hotspots later. Our focus in this study is primarily Helium as a wireless network, as opposed to the maintenance and operation of the Helium blockchain.

**Handlers** is a new term that is define to describe individuals who own and maintain Helium hotspots. In contrast to users, who send and receive data using the network, handlers and their hotspots do not generate any data; they simply forward it and receive payment. It is possible for individual entities to act as both users of the network and handlers who help deploy the network (i.e., we can own two hotspots and around a dozen devices).

### 3.2.3 Proof of Coverage

To be useful as a wireless carrier, Helium needs to provide geographically wide-area coverage. As a decentralized wireless network not owned and guaranteed by any one entity, there needs to be a mechanism for the network to prove that a hotspot is providing coverage. Helium uses an algorithm called Proof of Coverage (PoC) to prove the location of a hotspot. At random, one hotspot will act as a *challenger* and will randomly select another hotspot, referred to as either the *challengee* or *transmitter*, to send a wireless packet with an encrypted secret. Any hotspot that is in wireless range can *witness* this packet by reporting its contents to the challenger. Hotspot challenges are not geographically coordinated and can be acted on any other hotspot in the world. They do not target and prove any specific region has coverage, rather they stochastically validate every node in the network's coverage over time.

Each of the challenger, challengee, and witnesses receive HNT rewards after the PoC event. Challenger rewards are fixed (the reward is in creating and administering the challenge), but challengees' and witnesses' rewards scale with the quality of coverage (in simple terms here, more witnesses are better). As this chapter will show, there have been and still are numerous attempts to cheat this metric, and it continues to evolve.

When deployments are sparse, i.e., if a hotspot cannot "see" any other hotspots, then hotspots can only earn PoC rewards for challenge construction. This aims to incentivize hotspot density such that there are not holes or gaps in coverage; however, there also is decaying of

rewards if hotspots are too dense. Qualitatively, this seems to work, as improving coverage rewards is often discussed on Helium forums and message boards. Incentives and coverage quantitatively are discussed further in Sections 3.7 and 3.8.

### 3.2.4 HNT and DC

For general purpose users of the network, Helium aims to provide a stable pricing model. Miners are rewarded with newly minted Helium tokens (HNT). HNT is a traditional cryptocurrency, whose value has ranged from $8.32–19.70 USD in the month of May, 2021. Users planning deployments require more stable pricing—transit costs cannot unexpectedly double overnight. For this, Helium introduces Data Credits (DC), whose value is fixed at $0.00001 USD per 1 DC. All payments are done using fixed-cost DC.

Together, DC and HNT create a "burn-and-mint equilibrium." In the long run, the intent is to stabilize the price of HNT and tether it to network usage. In the short term, DC provide a stable, deterministic payment model for network users, while network operators receive incentives in the more speculative form of HNT. Notably, this model enables the Helium network to "pay" handlers without incurring traditional capital outlay. The details of this new crypto-economic model are only provided for background, analysis is beyond the scope of this chapter;this workonly study the network infrastructure-related aspects of Helium.

## 3.3 Data Sources & Methodology

The Helium network is a dynamically growing network composed of handlers that deploy hotspots, (ideally) many users that deploy edge devices, and a modest number of advanced users that deploy their own routers (cloud endpoints). While the distributed nature makes direct inspection of network activity challenging, the pay-per-access design means that most transactions are ultimately recorded in a publicly accessible ledger: the Helium blockchain.

Most of this work's analysis stems from an examination of the history of all transactions on the blockchain. While anyone can download and parse the blockchain—easiest done by

running the miner Docker container locally—most of the chapter's analysis takes advantage of a replica of the blockchain continuously extracted, transformed, and loaded into a database by the Decentralized Wireless Alliance (DeWi). This database also monitors the Helium p2p network, which is also used for this work's analyses. Details on the DeWi database are available on their Discord.[4] This work spot-verified several of our own transactions and our hotspots' p2p records to ensure they appear correctly in this database. We also perform some controlled measurements in Section 3.8 to measure the actual performance of the network.

The Helium blockchain is a fast-moving target. New blocks are minted every 60 s, and as of this writing, roughly 1,000 new hotspots are being added daily. Unless otherwise noted, measurements in this chapter reflect the state of the network as of late May, 2021.

There are 20 native Helium blockchain transactions [32]. The transactions most relevant to this chapter's analysis are as follows:

**add_gateway** adds a new hotspot to the network. It includes the hotspot ID, owner ID, location, and time when it was added.

**assert_location** allows an established hotspot to change its location. To discourage frequent moving of hotspots, this transaction carries a 1,000,000 DC fee ($10 USD).

**PoC_request/PoC_receipt** is created every PoC challenge and indicates the work done to validate network coverage.

**state_channel_open** creates a sidechain[5] for a router to use to receive packets. This stakes DC in advance, which allows packet delivery without waiting for individual payment transactions.

**state_channel_close** settles payments. Spent DC are burned, miners that ferried data receive HNT proportional to total network traffic, and unspent DC are returned to the router.

This blockchain data as a whole for initial insights and investigatory guidance. The first real entry to the blockchain was recorded on July 29, 2019. The vast majority of transactions recorded

[4] discord.com/channels/404106811252408320/769659586205712424/786375046930104341

[5] See Singh et al. for background and detailed discussion of sidechains [70]

on the blockchain consist of PoC requests and receipts. Out of 59,092,640 total transactions, 58,619,153 are carried out only to provide proof for the network accuracy and validity. Since the inception of the Helium network, approximately 99.2% of all blockchain transactions are PoC transactions.

## 3.4 What does "decentralized wireless" look like in practice?

The analysis begins by looking at hotspots, which make up the majority of the deployed infrastructure of Helium. The objectives are to understand where hotspots are deployed, who is deploying them, and what happens to deployed hotspots over time.

### 3.4.1 Where are Hotspots?

Whenever a hotspot's location is first published to the Helium network or changed, Helium records a transaction that contains the hotspot location encoded via the H3 geospatial indexing system [82]. The H3 encoding system encodes locations to hexagonal areas on a map; if a hotspot lies within borders of the hexagon, the hotspot's latitudinal and longitudinal coordinates are mapped to that specific H3 identity. H3 hexes allow for different "resolutions," which encode location with varying precision. Hotspot locations are stored at resolution 12, which are hexagons with an average edge length of 9.4 m and average area of 3.1 $m^2$; for location analysis, all hotspots are assumed to be located at the center of their hex.[6] The H3 location is decoded to retrieve hotspot latitude and longitude, which is used in all subsequent analyses.

Observing how hotspot locations change gives insight into whether the policy and rewards set out by the Helium network improve real-world coverage. Figure 3.2 looks at the absolute

[6]One may note from these averages that H3 hexes are not regular hexagons. This is a result of H3's use of gnomonic projection. For resolution 12 hexes, the min and max area are 1.9 $m^2$ and 3.7 $m^2$ respectively. In practice, hotspot locations are set using manually entered GPS fixes from a "set pin" interface on a companion phone app. Resolution 12 hexes are likely more precise than these fixes for the majority of use cases. As this chapter's analyses are primarily concerned with distances of several hundred meters or more, a few meters imprecision in hotspot location is not significant.

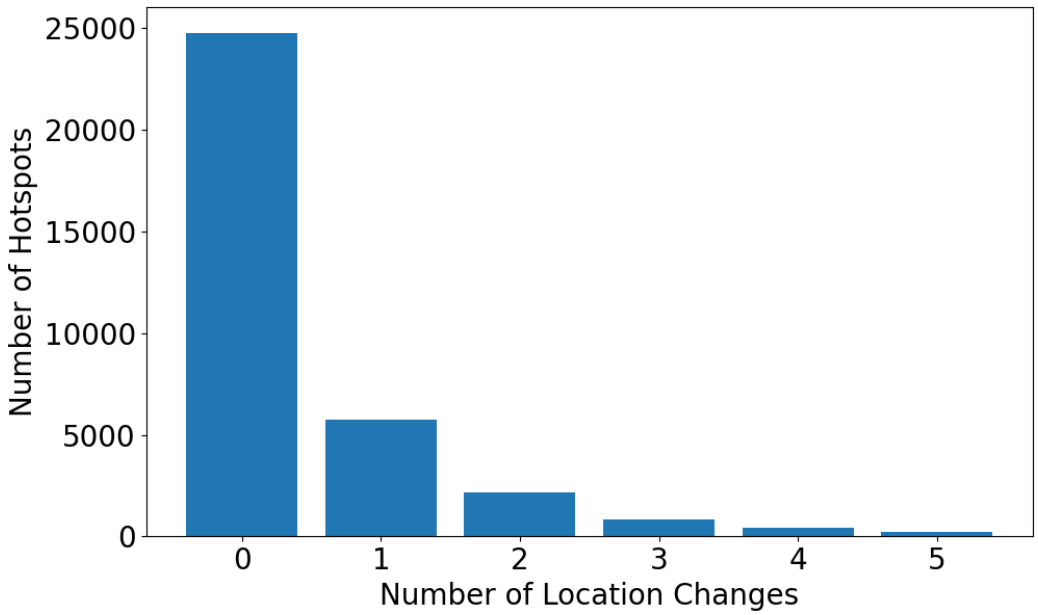


**Figure 3.2.** Location changes per hotspot. 71.9% of hotspots never move once deployed. 94.8% of hotspots do not move more than two times, and only 1% of hotspots move more than five times.

number of moves per hotspot and find that the vast majority of hotspots either do not move at all or move no more than two times. The Helium network permits hotspots to move up to two times for "free" (the Helium company pays the `assert_location` fee). One explanation for the large number of "first-move" events may be new hotspot owners first testing the hotspot in an easily accessible location and then moving the hotspot to a permanent deployment location.

Studying all hotspot moves allows to test this hypothesis. In Figure 3.3, there are two broad categories of moves. The first are a large number of short-distance moves, likely indicative of a local test-then-deploy scenario. The second are a large number of long-distance moves. Longer moves are visualized in Figure 3.3c.

There are two classes of long-distance movement. The first are hotspots that assert their location with a default (0,0) latitude and longitude—the large cluster in the ocean just below West Africa—and then later move to their true location. Hotspots occasionally move to (0,0) but do not stay there; currently, aside from hotspots who initialized at (0,0), there are no online hotspots that have moved to and remain at (0,0). 331 (89%) of (0,0) location assertions were first

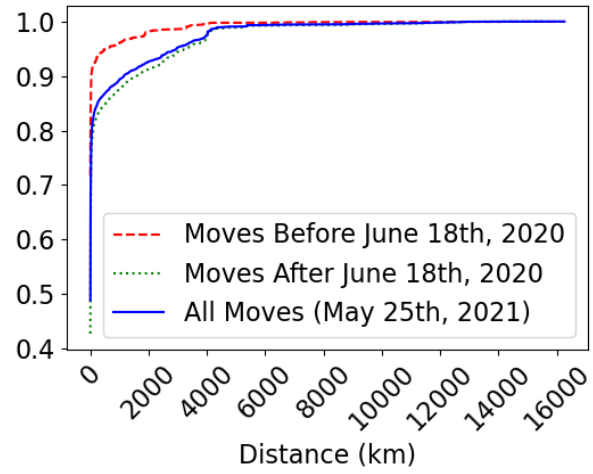


**(a)** CDF of moves (all).

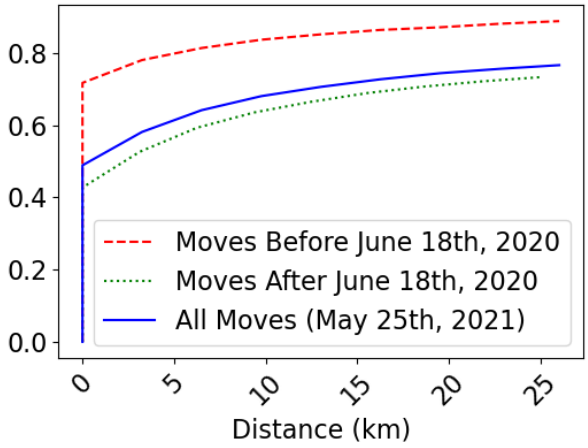


**(b)** CDF of moves (zoomed).

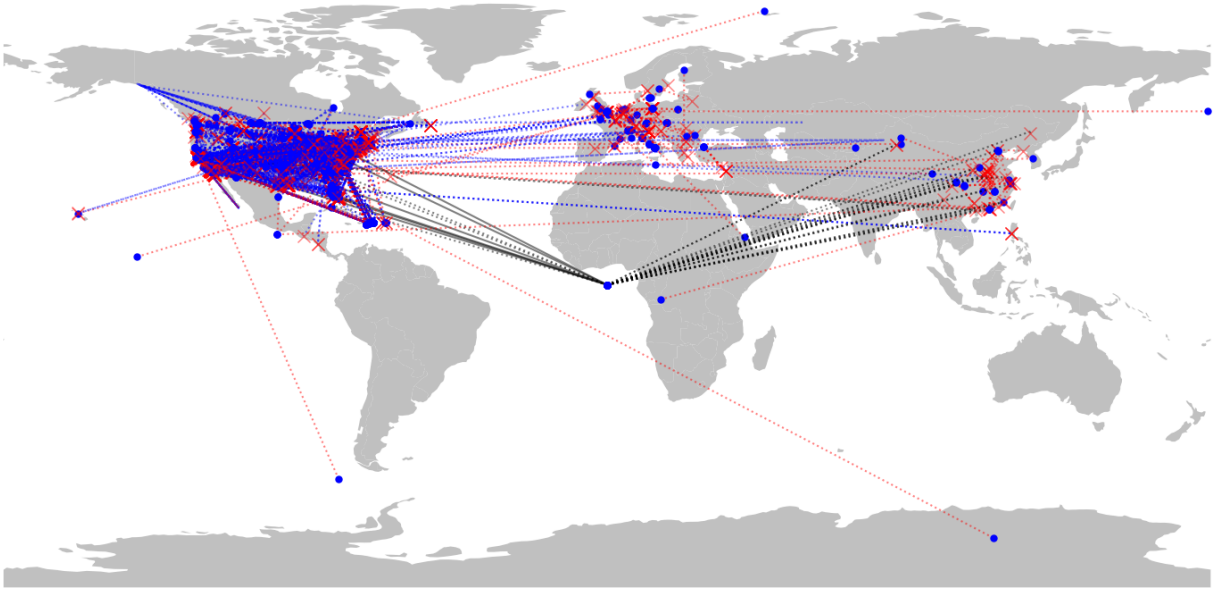

**(c)** Hotspots move from blue circles to red "X" marks.

**Figure 3.3.** CDF of all moves and map of hotspot location changes greater than 500 km. Figure 3.3c shows every location change from initial to final destination distances that are greater than 500 km away. Blue dotted lines indicate moves that departed the US; red dotted lines indicate location initialization outside of the US. Black dotted lines indicate moves from the initial location at (0,0) to anywhere else in the world, and black solid lines indicate a hotspot whose location changed to (0,0).

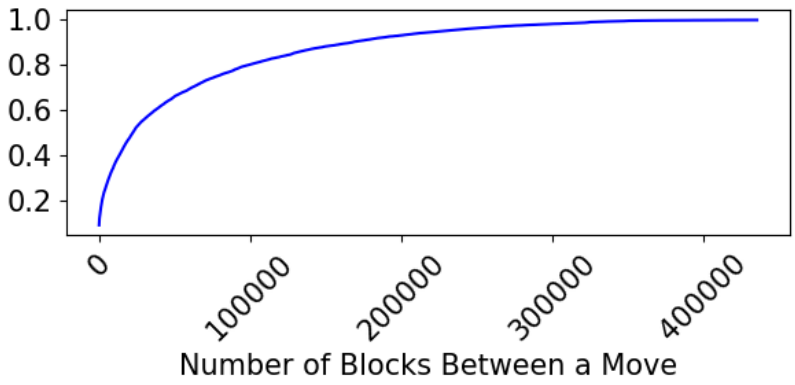


**Figure 3.4.** CDF of block intervals between hotspot relocations. Blocks are minted roughly once a minute.

time assertions, which suggests that most of the initial assertions at (0,0) were accidental (e.g. no GPS fix); the total number of assertions at (0,0) is 372. The remaining 41 assertions at (0,0) were attempted relocations. This could have also been accidental, a test out of curiosity, people attempting to game the system by clustering hotspots at a fake location, or possibly Helium developers testing validator nodes (which appear in the blockchain as hotspots that never transmit packets).

The second trend is a non-trivial "flow" of hotspots from the US to international destinations, particularly Europe. These moves are likely linked to resale of hotspots, and the heavy US-export to Helium's initial sales restriction to only the US market.

This work also studied the typical timing interval between a hotspot's relocation. Figure 3.4 shows a CDF regarding the number of blocks between a hotspot's relocation in to observe the time interval between location assertions; one block approximately corresponds to one minute. 17.9% of hotspot relocations occur within a day, 35.8% of relocations occur within a week, and 63.2% of relocations occur within a month. The remaining 36.8% of hotspot relocations have a timing interval of more than a month.

Finally, in the analysis of hotspot moves, there are a small number of outliers. For example, one hotspot moved twenty times. It is possible this is an individual with multiple residences bringing their hotspot along with them, perhaps in service of providing local coverage for their

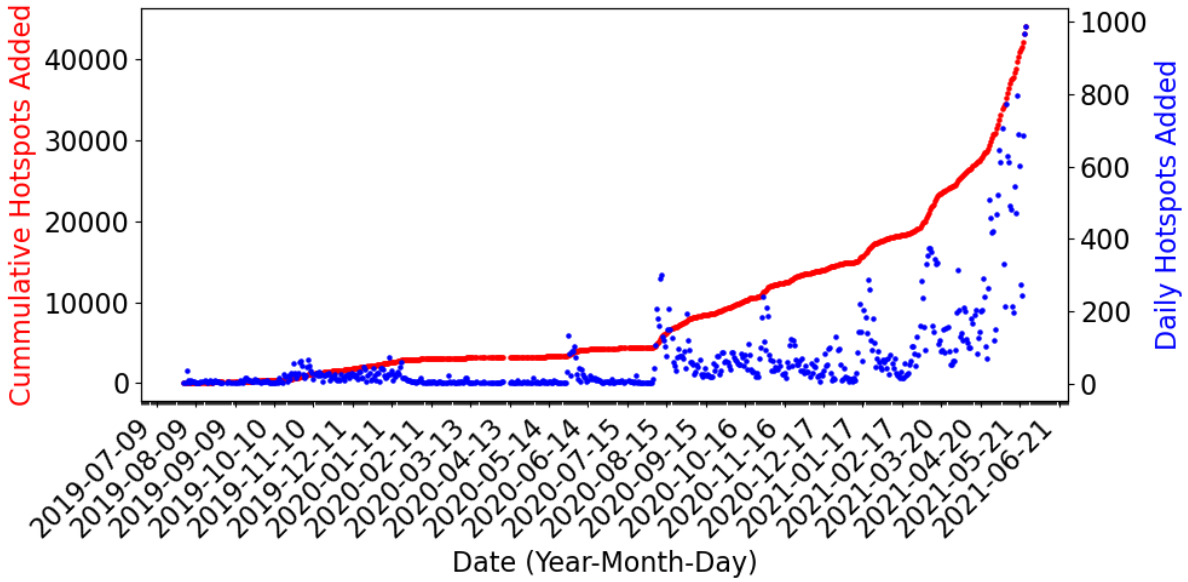


**Figure 3.5.** Helium network cumulative and daily growth. Qualitatively, growth seems mostly limited by hotspot availability. New production runs ('batches') are quickly placed into service.

own edge devices. However, note that the owner of this hotspot has well over 800,000 HNT, possibly posing as a business entity or developer account.

### 3.4.2 How Fast is Helium Growing?

Figure 3.5 shows the cumulative and daily number of hotspots that were added to the Helium network. For coverage discussions, this count is refined and draws a distinction between "connected" and "online" hotspots. A connected hotspot is one that has *ever* connected to the Helium network (this is the number displayed on various Helium status pages), while an online hotspot is one that is still connected and active (defined as fully synced and participating in PoC challenges). On March 7th, 2021, the total number of connected hotspots was about 20,000, however, only approximately 16,000 hotspots were online. On May 26th, 2021, the total number of connected hotspots was about 44,000, with 34,000 hotspots online.

This chapter also aims to understand where new hotspots are deployed. On March 7th, 2021, approximately 15,000 online hotspots located within the United States and 1,000 online hotspots located outside of the United States. On May 26th, 2021, approximately 20,000 online

hotspots in the US and about 14,000 outside of the US. While international coverage is growing rapidly, the network launched in the US in summer 2019 and was not available outside the US until summer 2020.

### 3.4.3 Who is Deploying Hotspots?

Every hotspot has a designated owner, or more precisely, a wallet that receives the rewards earned by the hotspot. There is an exponential decay relationship between the number of hotspots a single owner owns and the number of owners; approximately 5,700 owners (62.1%) own only one hotspot, about 1,300 owners (14.6%) own two hotspots, about 600 owners (7%) own three hotspots. There are about 9,000 unique owners total, of which 83.7% own 3 or fewer hotspots and 10.3% own 5 or more. As of May, 2021, the maximum number of hotspots owned by a single owner is 1,903. Contrast this with March, 2021, when the maximum was 160. The majority of hotspots are located in the US, and the second most common region is western Europe. While there are a modest number of large-scale owners, it is fair to claim that ownership (by unique wallets at least) of the Helium hotspot infrastructure today is decentralized.

This work then investigates several of the larger hotspot pools. A common inference from HNT balances over time is that owners which are using Helium in service of a real-world, end application engage in a large number of data transactions and have thousands to tens of thousands of HNT in their account. In contrast, owners which do not take part in data transactions generally have low HNT in their accounts as they frequently encash their HNT. These owners appear to be using the network as means for gaining profit rather than supporting their own edge devices. Shown next is a few examples of each of these classes of owner.

**Bulk Owners – Commercial**

This subsection begins from large owners on the Helium blockchain and works backwards to identify the owning company. This procedure may be used to identify entities using Helium infrastructure solely with information available on the public blockchain or other easily-

accessible public sources.[7] A similar approach can identify unannounced users of the network. It is emphasized that Helium does not explicitly provide receiver or transmitter anonymity. However, the design of the network also does not explicitly identify the users that are performing transactions. It is easy to identify publicly announced application traffic in the blockchain. This may be problematic for applications intended to be private.

**Careband** is a small startup in Chicago which specializes in developing wander-management wearables to detect patient movement, especially for those who have dementia. Their main office is located in 222 West Merchandise Mart Plaza. By looking at hotspots within that area, there is one owner ID which owns 25 hotspots mostly in and around in Chicago city and some individual hotspots in rest of the United States. It is assumed these hotspots are owned by Careband or provide coverage to Careband customers.

**nowi** is another startup that uses the Helium network to support water monitoring systems for multi-family property owners. They have one testimonial from Edworks LLC Property management. They are registered in Stonington, Connecticut. There are 19 owners that collectively own 61 hotspots in Stonington as of September 17, 2021. Out of them, 9 owners in Stonington own multiple hotspots and regularly sends at least tens of data packets every couple of of hours. This indicates that there are devices communicating their data over the Helium network.

**Bulk Owners – Mining Pools**

Owners which own multiple hotspots and carry out frequent data transactions are potentially service providers who manage these hotspots to provide coverage to their products. Other owners, such as the one in Figure 3.6, own multiple hotspots which were geographically distributed but do not engage in data transactions. These owners might be running hotspots to earn coverage rewards.

[7]The examples presented here have publicly announced their use of Helium via the Helium blog [33, 34], which is believed to mitigate any potential harm from using them.

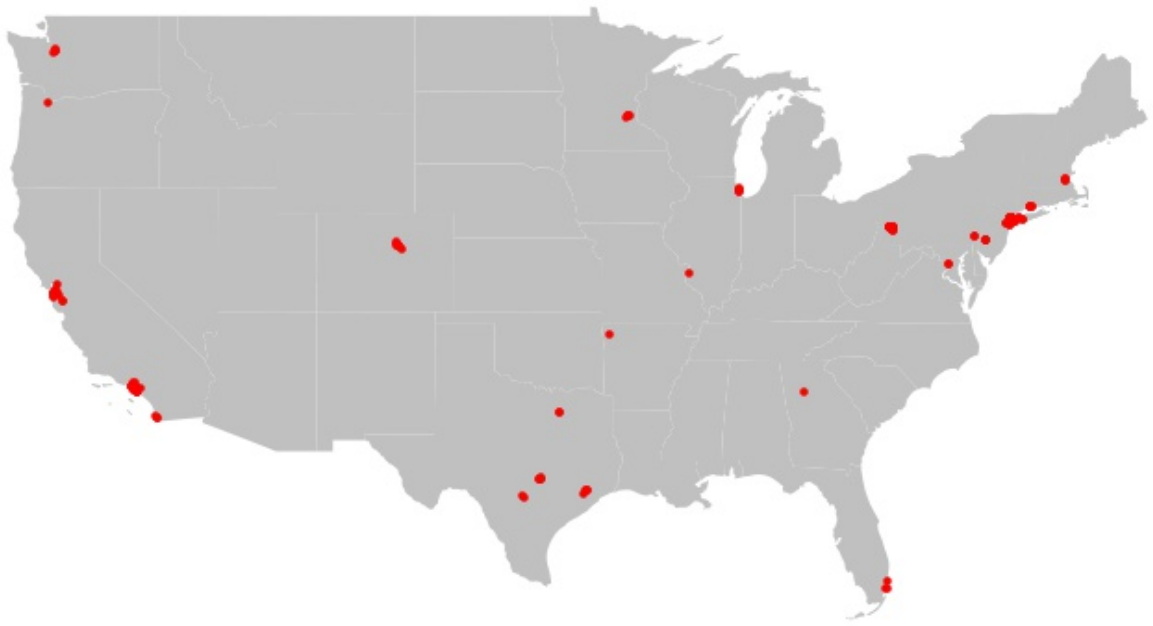

**Figure 3.6.** Hotspot distribution of one larger owner.

Ownership clusters in cities are popular. In one case, two owners own 144 hotspots and 136 hotspots in the Denver, Colorado area. The hotspots appear to be distributed evenly around the city. Deploying hotspots too geospatially close to each other reduces reward benefits; hotspots yield better mining efficiency when placed reasonably far away from each other. This incentive of improved rewards from improved network coverage appears to be important to these types of owners.

### Resale Market

Enthusiasm for Helium, coupled with the global electronics shortage [84] as well as the general manufacturing latency of new-to-market products, has resulted in a shortage of hotspots. The original Helium hotspot sold for $500 USD and newer, Helium-sanctioned third-party hotspots cost around $300-400 USD. An informal survey of the resale market finds a median price of $989 USD for a Helium hotspot among the top twenty eBay listings (min: $405, max: $6,500).

Some of the resale market is new, unopened hotspots. The sales of those hotspots cannot

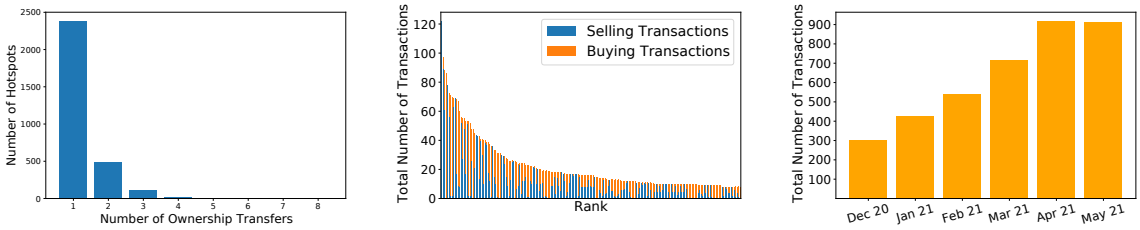


**(a)** Ownership transfers per hotspot.

**(b)** Number of hotspots transferred per owner.

**(c)** Hotspot transfer transactions over time.

**Figure 3.7.** Resale market analysis. About 95.4% of the total hotspots transferred do not change owners more than 2 times as seen in a. This shows that once hotspots are transferred, they do not have a tendency to change ownership. In b, the 200 owners—which have participated in the most hotspot transfers—either purchase or sell hotspots. These owners account for about 10% of the total owners who have participated in carrying out these transactions. Over time there is growth in the number of owners participating in these hotspot transfer transactions as seen in c. There were a total of 3,819 such transactions over a span of 6 months.

be tracked. However, Helium also supports a `transfer_hotspot` transaction, which allows one user to sell an established hotspot to another. Figure 3.7 analyzes hotspots and owners involved in resale transactions. About 8.6% of the total hotspots deployed are transferred to another owner. Over 95.8% of hotspot transfer transactions transfer 0 DC between buyer and seller, which suggests that the majority of resale payments take place using an off-chain marketplace, such as eBay; although it is not clear what ensures sellers actually relinquish ownership of the hotspot to buyers in such markets.

## 3.5 How Much is Built on Helium Today?

Next, this work aims to understand what types of users and applications are running on the Helium network today. This work's analysis reveals that currently Helium remains highly speculative, with more handlers deploying hotspots than users using Helium to ferry data. This is perhaps not surprising. The design of Helium's coverage-based reward model is to break into the chicken-and-egg problem of 'no users to pay for infrastructure' and 'no infrastructure to support possible users.' Still, there is a small but steady growth in what appears to be real-world application traffic, that suggests that 'if [Helium] can build it, they will come.'

### 3.5.1 How does payment-for-data actually work?

Data transfer transactions are not recorded immediately and directly to the main Helium blockchain. Instead, short-lived "state channels" aggregate batches of packets. The purpose of state channels is to permit fast and scalable payment for individual packet transfers, a critical facet of the Helium microtransaction model.

To receive data, routers must first open a state channel. A state channel open transaction stakes DC to pay for packets that may eventually be transferred and sets a deadline some number of blocks[8] after which the channel will be closed. Hotspots that receive a wireless packet from an edge device use the metadata in the packet to look up the owning router and send an offer to the router to buy the packet; the offer includes packet metadata but not yet payload contents. A router purchases a packet by sending back a signed offer to buy, at which point the hotspot releases the packet.

Routers are responsible for closing state channels after they expire. The state channel close transaction (should) include every offer to buy a packet made by the router. If a router signed an offer but never received the packet, it omits that offer from the close transaction. When a hotspot that did send data is left out of a close transaction, there is a 10-block grace period for the hotspot to a submit signed demand that amends the closing. If a hotspot lies about sending data, routers have no recourse but to add the hotspot to a blocklist and not make future offers to purchase its packets. As the value of individual data transactions is quite low and the duration of state channels is short, this discovery of malicious hotspots does not pose significant economic risk to routers.

For the analysis of data transfer behavior, then, there are limitations by the resolution of state channel transactions. State channel duration is decided by each router independently. As explained in the next section, however, nearly all traffic to date is sent on state channels with

[8]While not stated in any documentation, the blockchain implementation (github.com/helium/blockchain-core/blob/9011de7537ecfd737074b85b7b16e7d8e1ceef00/src/transactions/v1/blockchain_txn_state_channel_open_v1.erl#L208) limits this to a minimum of 10 blocks (~10 min) and a maximum of one week (exact block count is derived from current block time).

a 240-block (roughly 2 hour) duration, which is the granularity with which subsequent traffic analyses are able to operate.

One additional thing to note here is that it is possible for multiple hotspots to receive the same packet. While there is sufficient information in the packet metadata that a router can identify duplicate receptions, it can still choose to buy as many copies of a packet as it wishes. Observing payment flow then is a measure of data transferred between hotspots and routers, which may overestimate the actual flow of data from edge devices, depending on how often routers choose to purchase duplicate copies.

### 3.5.2 Who is running Helium routers?

Setting up a router and payment processor requires non-trivial technical expertise. Routers must be continuously online and responsive. The LoRaMAC between edge device and gateway has two acknowledgment windows, at precisely 1 s and 2 s after a packet transmission. The LoRaWAN protocol dictates that routers are responsible for sending acknowledgments if requested (as it is the router which much choose which gateway should send the acknowledgment packet if multiple gateways hear the original uplink packet from the device). Thus the cloud service must (1) learn of a proffered packet, (2) return a signed commitment to pay, (3) receive payload data, (4) generate an acknowledgment, and (5) send a signed commitment to pay for acknowledgment to a hotspot in under 1 s (or, with less reliability 2 s) for each data packet.

As of May 2021, there are only ten OUIs registered. OUI 1 and OUI 2 are registered to the Helium company. Of all state channel open/close transactions, 81.18% belong to OUI 1 and OUI 2.

As a (currently) free service, the Helium company provides the Helium Console, which is both a Helium router as well as an interface for provisioning and managing devices. The Console includes numerous integrations that allow data collected from sensors to flow to other services such as cloud database providers and mapping systems. Console users are required to buy DC for their devices, but this is purchased and used at-cost.

From an analysis perspective, this monopolistic router limits direct insight into application users, as all data transfer payments are from the Helium Console OUI, rather than from individual users or applications. To fund user accounts with DC on the Console, users can either burn their own HNT with the Console wallet as the destination—a transaction which is visible per-user in this analysis—or they can make a credit card purchase, in which case the Console will acquire and burn HNT using its own account. In practice, DC are so cheap that funding events are rare. Indeed, long before beginning this study, our research group made a one-off $10 USD purchase of DC (which is the minimum purchase amount permitted by the Console) in November 2020 to support another experiment. As of this writing, we have used less than 15% of this purchase, despite regular use of the network by multiple research projects and extensive use during the experiments for this work.

### 3.5.3 How much actual data is sent over Helium?

With the caveats on resolution outlined, namely state channel batching, opaque duplicate purchases, and a centralized OUI operator, the next direction is to analyze the behavior of data on the Helium network.

#### Device Data Traffic

This work then observes the trends of the device data traffic since 2019 to identify how much of the network is used for data transfers. Figure 3.8 gives a macro view of data activity. Most of data transfers in the earlier blocks were carried out by the Helium Router. However, the data transfers carried out by third party routers have recently started to increase, evidence of increased usage of the Helium network by end applications.

#### Fake Data & Arbitrage

One exceptional situation occurred on August 12, 2020, which is the date that DC payments first went live on the network. There is a sharp rise in the data transfer between August 12, 2020 and Sep 6, 2020 and a sudden drop after that. Prior to this date, data transfer was free,

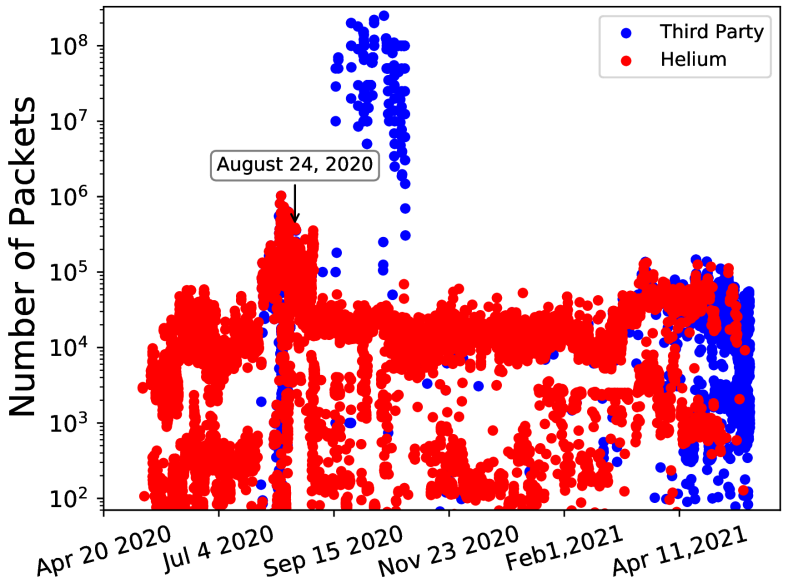


**Figure 3.8.** Packets transfer analysis. This shows the number of packets paid for with each state channel closing transaction, sorted by block. The primary trend is data paid for by OUI 1 and OUI 2, the Helium Console, which closes a state channel roughly every 120 blocks. At around 1 $^{\text{min}}/_{\text{block}}$, aggregate user traffic is approaching 14 $^{\text{packets}}/_{\text{second}}$ across the whole network.

and mining rewards that would have been allocated to data transfer were instead allocated to PoC activity. When DC payments and miner data transfer rewards first went live, there was no cap on the reward one received for network data transfers. Every epoch, 32.5% of newly minted HNT was divided among hotspots that ferried data, in proportion to the amount of data they carried – essentially, more data transfers would fetch you more rewards. Recall, however, that the cost of data is fixed, creating an arbitrage opportunity among DC/USD and HNT. Users were thus gaming the network by spamming packets to devices they owned to increase their shared of mined HNT. The arbitrage was stopped on August 24, 2020 with the implementation of HIP 10 [44], though it took slightly longer for the spam packets to fall off the network. This event remains the largest sustained volume of data traffic carried by the Helium network to date.

## 3.6 Meta-Infrastructure

This section takes a more holistic view of the Helium network and networking at large. The goal is to understand what infrastructure the Helium infrastructure relies on, and whether

**Table 3.1.** Top 15 ISPs used for hotspot backhaul.

| | ISP | Number of Hotspots | | | |
|---|---|---|---|---|---|
| **1** | Spectrum | 2497 | **9** | Sky UK | 199 |
| **2** | Comcast | 1922 | **10** | Telefonica | 199 |
| **3** | Verizon | 1590 | **11** | CenturyLink | 188 |
| **4** | Cablevision | 450 | **12** | TELUS | 185 |
| **5** | AT&T | 338 | **13** | RCN | 154 |
| **6** | Virgin Media | 333 | **14** | Frontier | 146 |
| **7** | Cox | 314 | **15** | Google Fiber | 142 |
| **8** | Level 3 | 202 | | | |

there are hidden points of centralization in this otherwise decentralized network. Despite the wide array of individuals deploying hotspots, the Helium network has potential choke points.

### 3.6.1 What ISPs do Hotspots Rely On?

One such example of a choke point is a region's reliance on ISPs. As all Helium hotspots are currently miners, they all participate in one, large p2p network.[9] The zannotate utility [22] is used together with Route Views data to identify the ASN of all non-relayed hotspots (i.e. hotspots with public IP addresses) connected to the p2p network.

In total, Helium hotspots are deployed in 454 ASNs. Figure 3.9 shows the complete distribution of these ASNs and Table 3.1 shows the top 15 ISPs derived from the ASNs using CAIDA's as2org dataset [20]. The most widely-used ISP is Spectrum with 2,497 hotspots, and the second most used ISP is Comcast, but it hosts significantly fewer hotspots. Verizon comes in a close third. Most hotspots are on Verizon's wireline network, but surprisingly 30 of the 1,590 hotspots are backhauled through Verizon wireless. There are also hotspots that use a cloud provider, such as Digital Ocean (72 hotspots) and Amazon (44 hotspots) rather than a last-mile provider. Potentially, these hotspots are validators, as they appear as hotspots on the blockchain.

Another statistic looked into was the percentage of unique ASNs within a city. The number of locally unique ASNs is important as relying on one ISP could cause a regional outage

[9]Mote that there is a limited window of time for this analysis. With the impending launch of validator nodes, hotspots will have the option to convert to so-called "light" nodes. Only the validators will maintain a fully connected p2p graph, and thus only they will have access to the network information of some hotspots in the future.

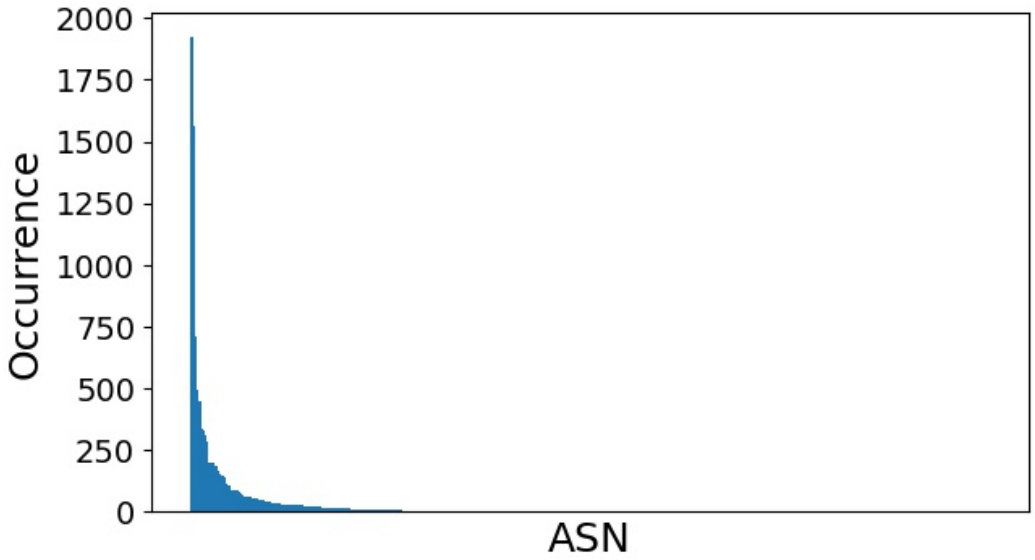


**Figure 3.9.** Distribution of ASNs for hotspots (with public IPs). Sorted by the number of hotspots per ASN, the overwhelming majority of hotspots hang off of just a few networks, although there is a very long tail of ASNs with just one or two hotspots.

if that ISP goes out. In total, there are 3,958 cities with at least one hotspot. Out of those cities, there is a total of 1,588 cities that relied on only one ASN, with 414 of those cities having at least 2 hotspots. These included cities such as Palma, Spain (with a total of 76 hotspots), Mesa, Arizona (13 hotspots), and Rome, Italy (12 hotspots). An example of an outage that may have had a large impact on Helium was the 2020 Spectrum outage in Los Angeles [19]. For a few hours, Spectrum customers across the city lost Internet access. This could have taken down 291 out of the 333 hotspots (87%) in Los Angeles.

### 3.6.2 Relay Analysis

One side-effect of adoption by individuals and smaller operators is that many hotspots are on network connections, such as residential home networks, that do not provide public IPs to all devices. Depending on NAT (or firewall) configuration, hotspots may not be able to accept inbound connections. `libp2p`, which Helium uses to form its network, addresses this with "Circuit Relays" [47].

When a hotspot cannot directly communicate, it opens a persistent connection with

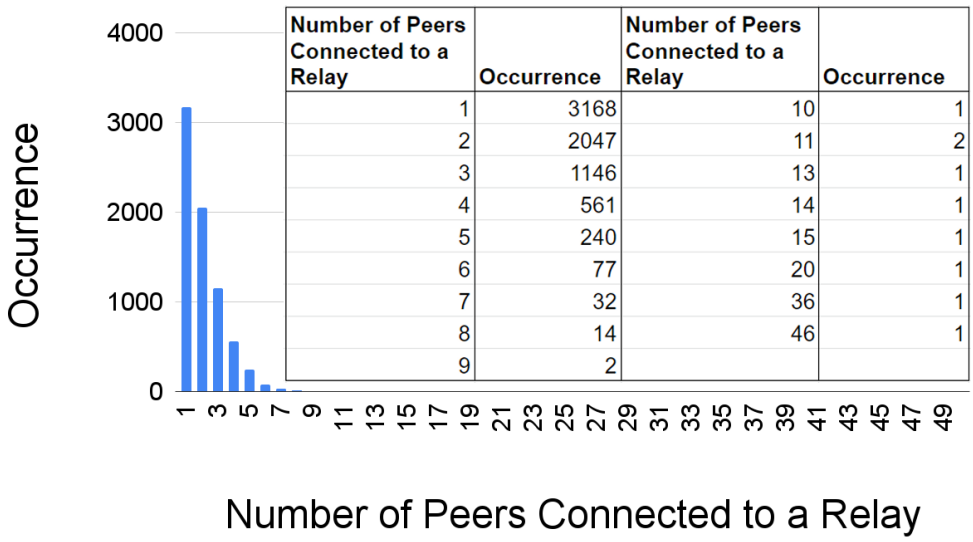


| Number of Peers Connected to a Relay | Occurrence | Number of Peers Connected to a Relay | Occurrence |
|---|---|---|---|
| 1 | 3168 | 10 | 1 |
| 2 | 2047 | 11 | 2 |
| 3 | 1146 | 13 | 1 |
| 4 | 561 | 14 | 1 |
| 5 | 240 | 15 | 1 |
| 6 | 77 | 20 | 1 |
| 7 | 32 | 36 | 1 |
| 8 | 14 | 46 | 1 |
| 9 | 2 | | |

**Figure 3.10.** Relay nodes with *n* peer nodes. Most hotspots relay only a few nodes. The cause of high-relay nodes is unknown.

another hotspot on a less restrictive network to relay messages and data. Peerbook entries are formatted in two ways: `/p2p/relay_node_hash/p2p-circuit/p2p/peer_node_hash` for hotspots who rely on a relay node and `/ip4/ipv4_address/tcp/port` for hotspots that have public IPs and accessible ports. Using this information allows to study relay prevalence and behavior.

First, the prevalence of relays is surprising. Of the 27,281 hotpots with non-empty listening addresses, 55.48%—more than half the network!—are relayed. This heavy reliance on relay nodes increases the meta-infrastructure risks identified in the previous section, as relayed nodes are beholden to their relaying device.

Next, Figure 3.10 looks at how relays are distributed among hotspots. While the majority of relaying nodes support just one or two peer nodes, there are a few who relay as many as 46 other nodes. It cannot be determined why these nodes relay such a large number of peers. One hypothesis investigated was whether these are the hard-coded seed peers[10] that ship with the hotspot firmware image, but the high count relay nodes do not match these IPs.

The next hypothesis is that hotspots bias towards a geographically nearby peer. Such a

[10] `/ip4/35.166.211.46/tcp/2154,/ip4/44.236.95.167/tcp/2154` from https://github.com/helium/router/blob/master/.env-template#L2

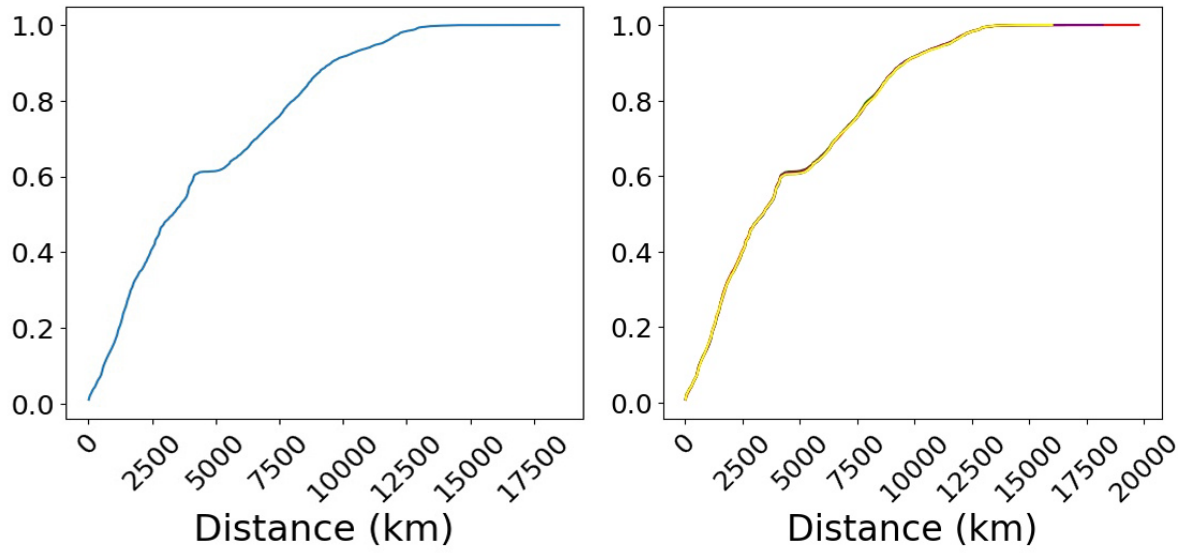


**(a)** Actual.

**(b)** 5 random assignments.

**Figure 3.11.** Relay to peer node distance, actual & simulation. Peers choose relays randomly, without geospatial consideration.

design could be problematic for local robustness: if many geospatially clustered nodes rely on the same relay peer, then coverage reliability for that whole area would fall to the reliability of the relaying node. At the same time, ignoring location also can create problems for a globally distributed network, particularly one which requires multiple round trip communications in under 1 s to support LoRaMAC acknowledgments.

The asserted location data of each hotspot is used to compute the distance between each peer and its relay node and graph the distribution as a CDF, shown in Figure 3.11a. The majority of distances are below 5,000 km. The minimum distance is 0.46 km and the maximum is 18,491.10 km. While this suggests peer selection is random, given the non-uniform geospatial distribution of hotspots, distance alone is insufficient. Next, the list of all relays and relayed nodes are run through multiple trials which randomize the assignment of peers to relays, as shown in Figure 3.11b. With this analysis reinforces confidence that the Helium network does in fact assign peers randomly to relay nodes.

## 3.7 Governance By Incentive

Because the Helium network is decentralized, it cannot directly affect change on the deployed infrastructure. Instead, the network uses economic incentives to motivate changes in user behavior. While these are relatively stable, "Helium Improvement Proposals," or HIPs, can change the rules of the Helium blockchain.[11] In principle, HIPs create economic incentives for hotspot owners to change their behavior. This section looks at examples of the efficacy and inefficacy of government by incentive on the Helium network.

### 3.7.1 Case Study 1: Silent Movers

As mentioned in Section 3.2.3, to test location, any hotspot can send a challenge every 480 blocks to any other hotspot to request that the "challengee" hotspot prove its location. This case aimed to identify any potential challengees with supposed witnesses that are physically impossible. To do so, hotspots' asserted location and the location of where they witnessed for another hotspot are matched. While there were not many unique offenders, there was one common offender whose asserted location was across the country from its witness location. This hotspot is referred to as Joyful Pink Skunk.[12]

Joyful Pink Skunk's last `assert_location` transaction was on April 11, 2021, when it reasserted its location from the state of Florida to the state of Pennsylvania. This transition appears to be honest as the next time the network selected Joyful Pink Skunk as a challengee, which was witnessed by hotspots located in Pennsylvania. Starting May 2, 2021, hotspots in New York, NY and Brooklyn, NY became valid witnesses for the Joyful Pink Skunk. At the time of finishing this work, it still has not reasserted its location, and it witnesses hotspots in the state of New York.

Joyful Pink Skunk never reasserted its location when it moved to New York. Normally, hotspots should be incentivized to update their location to earn PoC rewards. Yet, from the

[11] For details on the HIP process, see: https://github.com/helium/HIP.

[12] Name anonymized to protect hotspot identity.

challenge receipts, Joyful Pink Skunk is receiving HNT regardless of whether its current position matches its last asserted location.

Moreover, hotspots do not have to provide an accurate location at all. This is evident through the hotspot Striped Yellow Bird[12] whose only `assert_location` puts it in Spokane, Washington, but all of its challenge receipts place it in San Francisco, California. Nevertheless, it is still rewarded HNT for providing coverage in an area that is about 1,150 km away from its purported location.

*Takeaway:* If location is not properly considered in the rewarding process, hotspot owners have little to no incentive to keep their location accurate. Shown in Section 3.8, inaccurate locations impede coverage modeling. The $40 USD cost to re-assert location is designed to promote stable spatio-temporal coverage by deterring hotspot moves, but it does little if owners can skip reporting moves.

### 3.7.2 Case Study 2: Lying Witnesses

FCC regulations limit transmitters to +36 dBm EIRP. Yet some witnesses claim an RSSI as high as 1,041,313,293 dBm (presumably either from a buggy radio driver or a misguided attempt to earn more rewards for witnessing "well"). While this value is easily dismissed, it exemplifies that the current PoC model relies on witnesses reporting their RSSI truthfully, while RSSI is easily forged. Colluding, modestly geospatially clustered nodes could easily gossip challengee secrets to increase the number of challenges (plausibly!) "witnessed," and in turn the gossip clique all earn more rewards.

The blockchain implementation has checks that attempt to use RSSI to establish whether a witness is "valid" (and should thus receive PoC reward payment). Real-world RSSI can exhibit enormous variation [23, 81], however, which inevitably makes such heuristics brittle. Ultimately, there are misaligned incentives here. The network wants witnesses to honestly report RSSI to better estimate coverage, while witnesses want to report whatever RSSIs maximize their likelihood of being rewarded for witnessing (independent of whether they are an honest recipient

of the challengee's packet).

*Takeaway:* RSSI is an unreliable, imprecise, and unstandardized measure. Tying reward payments to it will only incentivize gaming the metric. Users with uncharacteristic, but honest RSSIs will be frustrated by unfairly lost revenue and expert manipulators (with access to the cheating detection algorithm running on the public blockchain) will always be able to defeat heuristics.

## 3.8 Empirical Testing

This chapter's final measurements and to answer the following question: (how well) does Helium actually work? While devices can be deployed and data can be recovered, there are significant limitations today in the reliability of Helium. One of the largest challenges for persons considering the own deployment is the absence of a meaningful coverage model – will Helium cover *my* system? Extant blockchain incentives are used to derive implicit coverage models, but find these are quite imperfect, which may imply that the models in this chapter are too simplistic, that current incentives do not sufficiently promote meaningful coverage, or some mixture of both.

### 3.8.1 Basic functionality

For a first test, the best-case scenario is considered. We own an (unmodified) original-batch Helium hotspot, which is attached to the campus backhaul network (on a subnet that grants public IPs and access to arbitrary ports). A ST B-L072Z-LRWAN1 LoRaWAN development board is used as it was the first development board explicitly supported by Helium, and it was at the top of the list of platforms on the Helium Quickstart Guide.

The device is loaded with a basic app which sends an incrementing counter. The app is a free-running send, which attempts to send another packet immediately after the prior packet response.[13] This app is run for about 24 hours and see a packet reception ratio of 68.61%. There

[13]LoRaMAC has two acknowledgment windows, one second and two seconds after transmit. If every packet were ACK'd on the first try, this app would send one packet per second; if no packet were ever ACK'd, it would

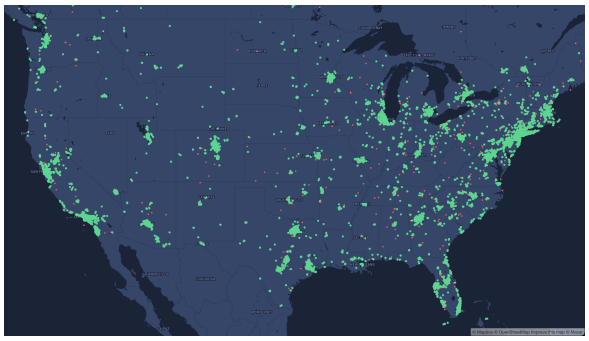

**(a)** Coverage as reported by Helium.

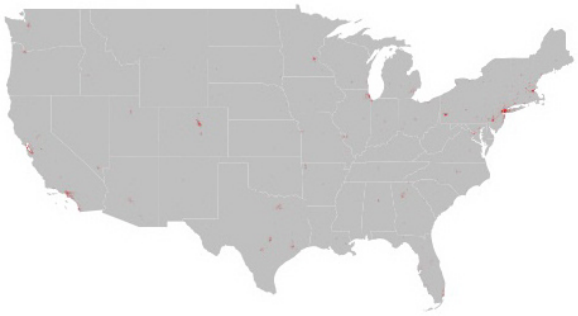

**(b)** Coverage estimate using 300 m cutoff.

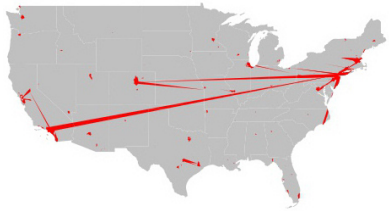

**(c)** Coverage from witness convex hulls.

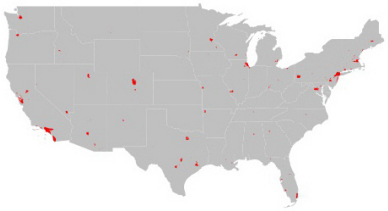

**(d)** Coverage from witness convex hulls (removing witnesses that are more than 25 km away from challengee).

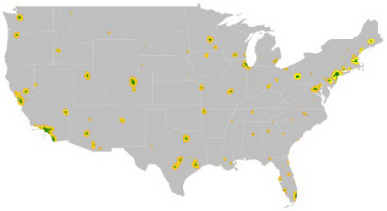

**(e)** Revised convex hull coverage map that factors in the hotspots that make up the vertices of the convex hull and RSSI.

**Figure 3.12.** Estimates of Coverage.

are occasional outages in the network of around 2 hours where no data reaches the cloud. But in between these outages, almost all the packets transmitted are received. This experiment was carried out between 18 May, 2021 and 19 May, 2021. During this same time a new firmware was released [36] which is possibly why there is lower PRR due to the network outages.

To try to remove this firmware confound, this experiment is re-run in September 2021. The sensor is also relocated to a residential neighborhood with a much greater density of hotspots. Despite these changes, unreliable performance is still observed, with an overall PRR of 73.2% across three trials. There are no significant gaps. 83.5% of missed packets are single-misses (i.e. packets before and after were received), 92.2% are single- or double-misses, and the longest sequential run is a single instance of 34 consecutive missed packets.

send one packet every two seconds.

### 3.8.2 Coverage

At the end of the day, the most important question for a wireless infrastructure provider is the quality and availability of service. The Helium network is expanding quickly. As of this writing, Helium is averaging an addition of 1,000 new hotspots per day [52] (a claim that is verified in Figure 3.5). For the studies of coverage, the state of the network is considered as of May 23, 2021.

#### Coverage Models

The blockchain records hotspot locations, but LoRa is intended as a "long-range" wireless technology. This means a model needs to be developed to go from a list of hotspots to expected geospatial coverage. These coverage analyses are focused on the United States, as Helium's initial launch was restricted to US-only territories. While the network is seeing rapid international expansion, the US remains its most established market, which makes it a better representation of current best-case capabilities.

**From `helium.com`:** Helium provides a "Coverage Map" at explorer.helium.com/coverage. Figure 3.12a is a screenshot from that website. Green dots represent online hotspots while red dots represent offline hotspots. While the map is a good representation of hotspot locations, this view of coverage can be misleading as the dots indicating hotspots always render at the same size, and thus individual hotspots appear to cover more and more area as the map is zoomed out and the landmass underneath a hotspot-dot grows.

**Density Incentive:** HIP 15 specifies that hotspots within 300 meters of each other cannot act as a witness for one another. This is to promote wide-area coverage by discouraging hotspots from clustering too closely together. The implication then is that a hotspot should be able to provide coverage to any device within a 300 m radius. Figure 3.12b shows the Helium coverage of the contiguous US using this 300 m radius model. Compared to the Helium coverage map in Figure 3.12a, the area of coverage for the 300 m radius approach is barely perceivable. The total

percentage of contiguous US landmass covered by the 300 meter radius model is 0.09295%.

**Witnesses:** As a reminder, witnesses are hotspots that report challenge packets transmitted by a challengee to confirm the challengee's location. There are two types of witnesses: a valid witness and an invalid witness. A witness is marked valid unless it is deemed invalid by satisfying one of the following criteria:

- is too close to the challengee (<300m)
- has too high of an RSSI (several heuristics)
- has too low of an RSSI (several heuristics)
- is pentagonally distorted (rare artifact of H3 distance)
- claims capture on the wrong channel (impossible)

While the 300 m radius coverage map provides a better understanding of the actual coverage, it to be too conservative. To alleviate this fact, valid witnesses of a challengee are used to infer an empirical measure of coverage. For each challenge, a convex hull is drawn around the challengee and its valid witnesses and assume coverage of the interior of this hull. This coverage model is overlayed on top of a contiguous US map in Figure 3.12c.

One problem that becomes evident from the convex hull model is that some of the "valid" witnesses should in reality not be valid (indeed, debugging this model led to the examples for Section 3.7.1). It is possible to remove questionable witnesses by using a more realistic max distance to craft a more realistic estimate of coverage. Murata, a top LoRa radio vendor, suggests that the realistic range is, "... more than 10 km, between 15 to 20 km" [54].[14]

The distribution of distances for all purportedly valid witnesses are shown in Figure 3.13. For the revised convex hull model, a generous 25 km cutoff is chosen, after which "valid" witnesses are rejected and exclude them from the hulls. Figure 3.12d shows this revised model, with coverage now covering about 0.5723% of the total contiguous US landmass.

**Witness RSSI:** While more generous, the basic convex hull model is still too conservative. In particular, it does not factor in coverage by the hotspots that make up the exterior vertices of

[14]Geography influences this greatly. As example, there appears to be honest hotspots (likely with advanced antennas mounted at high-altitude) in Chicago and western Michigan that witness successfully at ranges of 60-110 km across Lake Michigan.

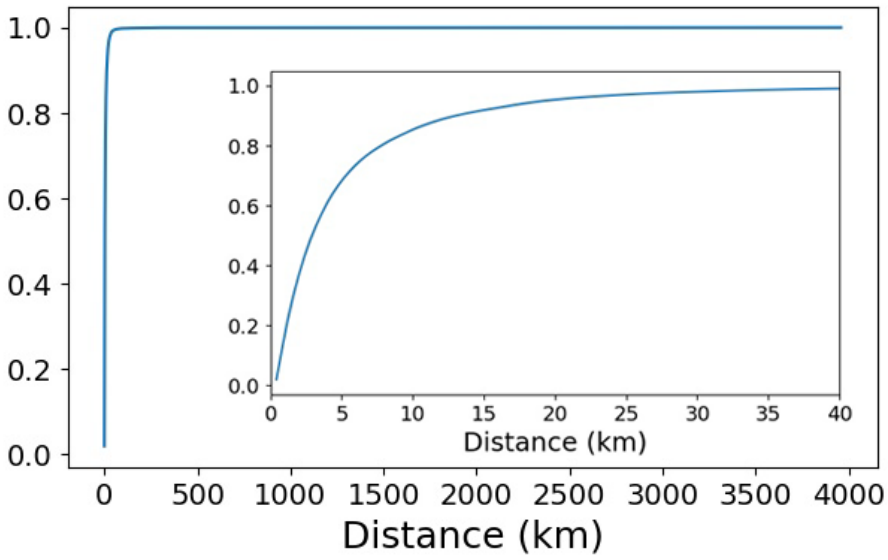


**Figure 3.13.** CDF of all the valid witness' distance with an inset CDF of the distance interval from 0 km to 40 km.

the convex hulls. The model is then revised one final time by including a radial vertex hotspot coverage model and an RSSI coverage model.

The radial coverage by the vertices is simple: find the distance from the vertex witness to the challengee and use the result as the radial coverage for that witness. As for the RSSI model, take the vertex witness RSSI and "grow" the witness–challengee radius using the standard free space path loss model, $d = 10^{\frac{w-s}{20}}$ where $d$ is the distance to lengthen the radius, $w$ is the witness's RSSI value, and $s$ is the sensitivity of the device hoping for coverage. $s$ is set to be a constant -134 dBm as that is the receiver sensitivity of the recommended ST LoRa hardware platform.

Figure 3.12e shows the final result of this model. The green regions come from the revised convex hull model. The yellow areas come from the added radial coverage. The almost-invisible red trim around the yellow regions comes from the RSSI coverage. Zoomed out, it is difficult to see the impact of RSSI coverage. Figure 3.14 shows the distribution of RSSIs reported by witnesses from 2021-05-18 to 2021-05-22. At the median -108 dBm, the RSSI step adds only an additional 20 m of coverage range. With this coverage model, the network achieves 3.3032% coverage of the contiguous US.

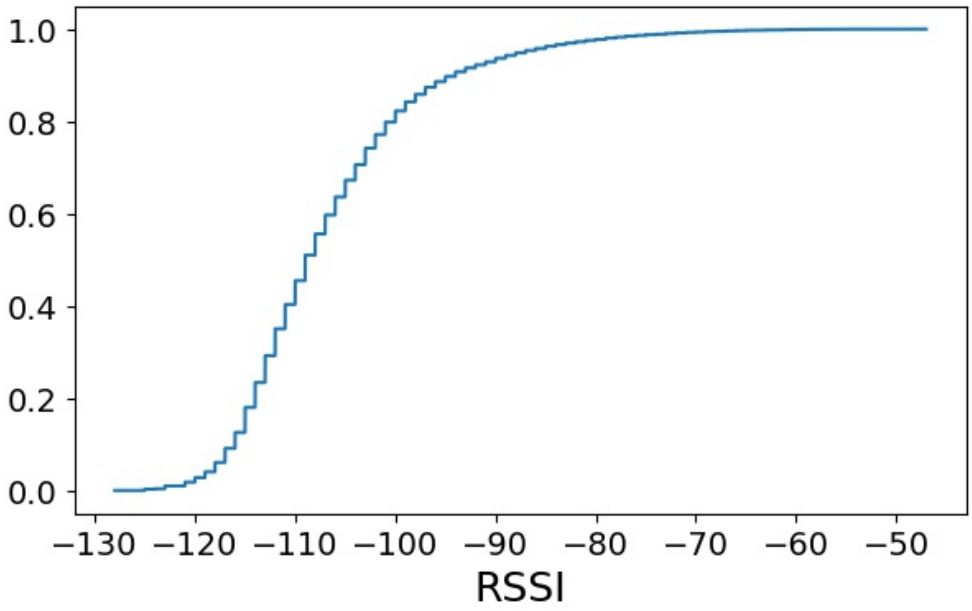


**Figure 3.14.** CDF of RSSI values recorded by witnesses during PoC requests from 2021-05-18 to 2021-05-22.

### Measured Coverage

Two real-world experiments with results are shown in Figure 3.15a and Figure 3.15b to observe the empirical coverage of hotspots. Neighborhood walks are preplanned through areas with varying hotspot density. While walking, an edge device is carried while running the counter app described previously. GPS coordinates and a timestamp are added to the app payload. Packets are also logged to an SD card when sent, to create a record which can be compared to what was received by the cloud. Overall, these walks have a Packet Reception Rate (PRR) of 72.9% and 77.6% respectively.

This experiment investigates whether the HIP 15 promise holds and explains when and where losses occur. Predicting reception when within 300 m of a hotspot is accurate 55.5% of the time, while predicting no reception outside of the radius is accurate for 79.6% of packets.

The LoRa protocol includes an acknowledgment mechanism, where an edge device can request an ACK response from its owning router. At the edge device, failure to receive an ACK is recorded as a NACK. Note that the LoRa PHY is asymmetric; said simply, uplink

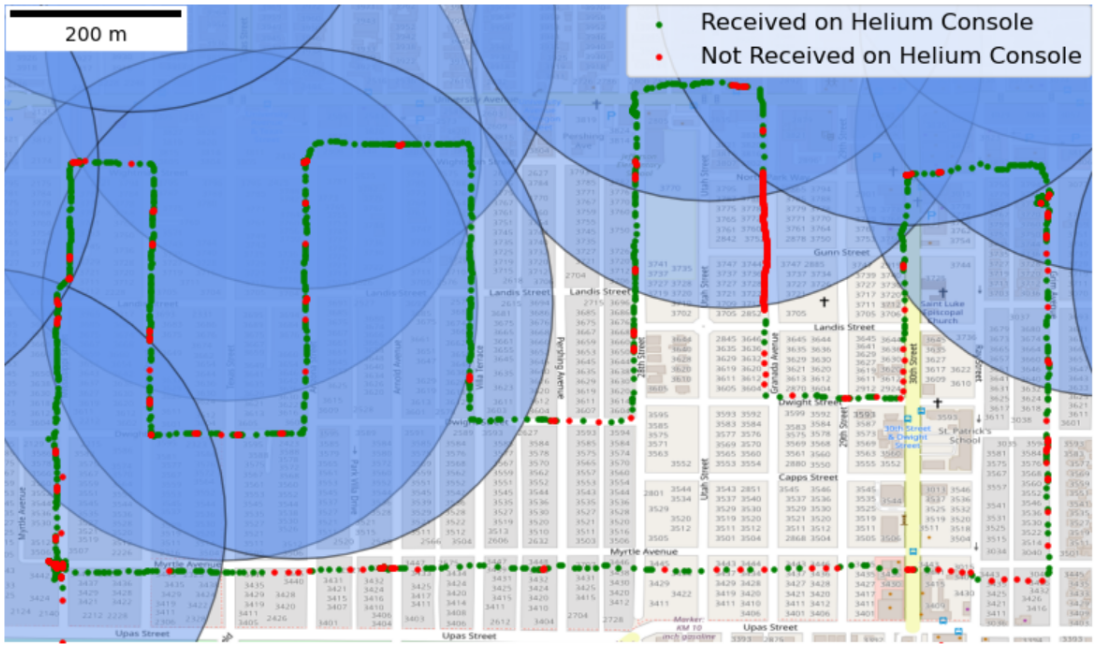


**(a)** Urban.

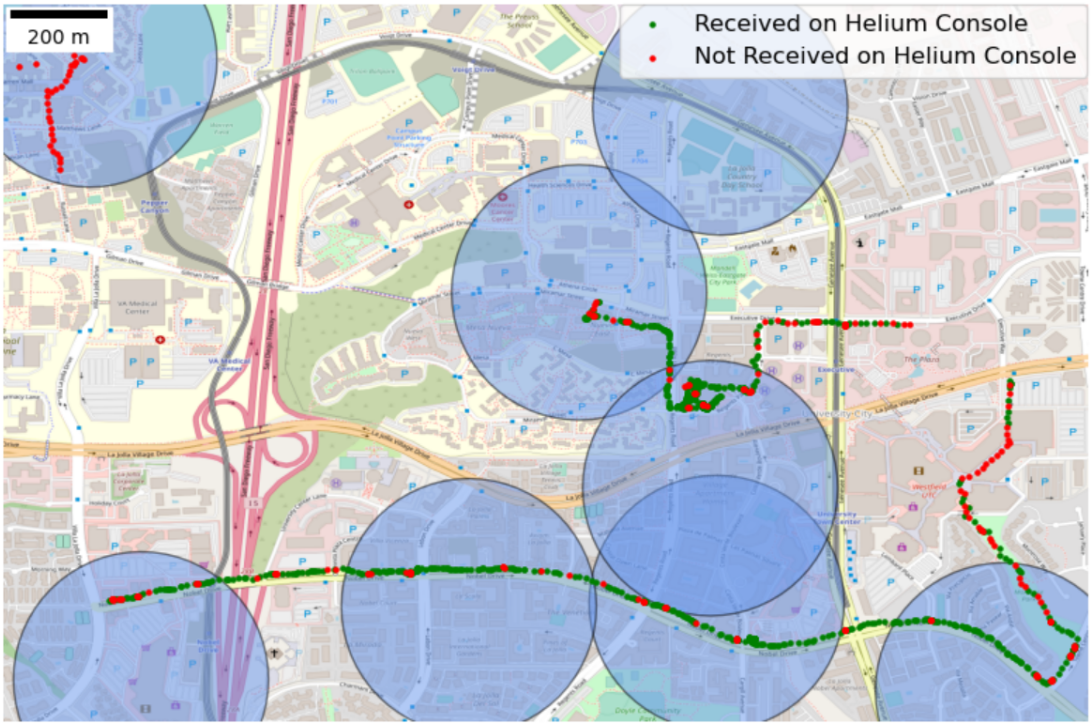


**(b)** Suburban.

**Figure 3.15.** Empirical coverage testing. Green dots represent packets sent by the device and received on the cloud. Red dots are packets that were sent but not received on the cloud. Transparent blue circles show 300 meter hotspot "coverage" radii (hotspots are at the center of the circles). These experiments were carried out by walking outside while attempting to send packets from an edge device to a logging application attached to the Helium Console.

**Table 3.2.** LoRa ACK/NACK Validity from Figure 3.15a

| | Packets Sent | Correct ACK | Correct NACK | Incorrect ACK | Incorrect NACK |
|---|---|---|---|---|---|
| Count | 2393 | 1106 | 986 | 0 | 301 |
| Percent | 100% | 46.2% | 41.2% | 0% | 12.6% |

**Table 3.3.** LoRa ACK/NACK Validity from Figure 3.15b

| | Packets Sent | Correct ACK | Correct NACK | Incorrect ACK | Incorrect NACK |
|---|---|---|---|---|---|
| Count | 1027 | 585 | 237 | 0 | 205 |
| Percent | 100% | 57.0% | 23.1% | 0% | 20.0% |

(edge→gateway) is easier than downlink (gateway→edge) [85]. This means the cloud may record data which the edge device thinks it needs to retransmit. Statistics are run regarding ACK and NACK validity as shown in Table 3.2 and Table 3.3. There were no false ACK messages and many false NACK messages—packets received on the cloud but recorded as NACKs by the edge device.

There are many factors unknown to us regarding the reliability and coverage of the Helium network. This is only a basic experiment to explore network reliability. Future analyses should deploy their own routers, own more gateways to monitor traffic, and extract more diagnostics from the edge device LoRa stack to enable rich root cause analysis.

## 3.9 Summary

This chapter discusses initial insights into the workings of the Helium Network. The network continues to expand rapidly to provide broad connectivity for commodity edge devices, and the growth of the network is in part because Helium is providing well-designed infrastructure and incentives. Federated (i.e. crowdsourced) infrastructure does indeed work, although the uncontrolled and unplanned deployment of hotspots does not necessarily ensure reliable or predictable coverage. This chapter summarizes the viability of the Helium Network as a large scale LoRaWAN. The next chapter aims to evaluate how well the Helium Network—and federated infrastructure—performs in relation to other wireless technologies and their counterpart

infrastructures.

## 3.10 Acknowledgements

Chapter 3, in full, is a reprint of the material as it appears in Dhananjay Jagtap*, Alex Yen*, Huanlei Wu, Aaron Schulman, and Pat Pannuto. Infrastructure: Usage, Patterns, and Insights from ”The People’s Network“. In Measurement Conference (IMC ’21), November 2021. The dissertation author was the primary investigator and author of this paper.

# Chapter 4

# Wide-Area Wireless Measurements

## 4.1 Background

This chapter takes the initial empirical measurements on the Helium Network, extending the small scale measurement study at the end of Chapter 3 to a country-wide study of LoRa and the Helium Network. Specifically, we compare LoRa and the Helium Network against two technologies and infrastructures: (1) cellular via Verizon radio equipment and towers and (2) BLE via Apple devices and the Find My network. The infrastructures via Verizon and Apple represent two extremes, in which Verizon represents a multi-billion dollar, costly infrastructure that is meticulously planned, deployed, and maintained for guaranteed coverage and reliability, and Apple devices—which people have purchased for their personal use—being an opportunistic, unplanned infrastructure that uses these Apple devices to backhaul data as described further in Section 4.2.1. In terms of cost, cellular (via Verizon) represents the costly and expensive infrastructure that guarantees usability, whereas BLE (via Apple devices and Find My) represents "cheap," opportunistic infrastructure that can provide coverage wherever Apple devices exist. As a result, this chapter aims to elucidate how LoRa (via Helium) compares to these two extremes of expensive and "cheap" infrastructures.

**Table 4.1.** A comparison of capabilities of BLE [43], cellular [18], and LoRa [63] as well as the approximate magnitude of their infrastructure.

| | Frequency | Range | Energy per byte | Maximum data rate | Est. num. of base stations worldwide | Medium Access control | Reliability expectation |
|---|---|---|---|---|---|---|---|
| **Cellular**[†] | 700 MHz–2.6 GHz | ∼10–15 km | $10 \sim 10,000\,\mu J$ | 5 Mbps | ∼41.6 million | (C/T)DMA | High |
| **LoRa** | 900 MHz (in N.A.) | ∼5–15 km | $10 \sim 100\,\mu J$ | 21.9 kbps | ∼400,000 (Helium) | ALOHA | Medium |
| **BLE Advertisement** | 2.4 GHz | ∼0.03 km | $1 \sim 10\,\mu J$ | 100 bps | ∼1.8 billion (Find My) | ALOHA | (Variable) |

[†] This study uses traditional 2G-5G cellular. For IoT applications, $^{energy}/_{byte}$ is more often the higher end of this range as 2G-5G networks expect larger data volume to amortize cost. Emerging CAT-M and NB-IoT protocols aim operate at lower power points, but are not available to low-volume users in North America yet.

## 4.2 Background on Wireless Technologies and Infrastructures

We first describe the technical details of each of the technologies and infrastructures we consider (i.e., BLE/Apple Find My, cellular/Verizon and LoRa/Helium) to backhaul mobile data. We set up a notion of expected reliability based on both existing infrastructure scale and underlying technology to predict the relative success of each technology. We then discuss the trade-offs of each technology's reliability with respect to energy consumption. A summary of these two primary comparisons are shown in Table 4.1.

### 4.2.1 Reliability

A wireless technology is only as good as the infrastructure that supports it. If iPhones did not exist, neither would Find My or AirTags. If the scale of cellular infrastructure were only a quarter of what it is today, coverage would deteriorate. In short, the success and reliability of a wireless technology is influenced by the size and scale of the supporting infrastructure, the cost of which is controlled by the design and capability of the wireless technology.

We approximate the size of each technology's supporting infrastructure in Table 4.1. The largest infrastructure are the Apple devices that support the Find My network, at roughly 1.8 billion gateways. This infrastructure is unique in that Apple uses existing products (iPhones, iPads, and MacBooks) to backhaul data, rather than having purpose-made gateway devices. Without deploying any new gateway infrastructure, Apple is able to achieve wide-area coverage using extant devices at nearly zero added monetary cost.

On the other hand, billions of dollars have been invested into the cellular infrastructure over its lifetime. With approximately 40 million base stations around the world, we expect cellular to have high reliability. Although orders of magnitude smaller than Apple's ad-hoc Find My infrastructure, cellular's decades of planned deployment allow for stable, wide-area coverage.

Helium, the largest LoRa deployment to date (which is an order of magnitude smaller than cellular), explores "federated infrastructure" by inviting average citizens to deploy gateways themselves. Helium peaked at nearly 1 million gateways, and today sees around 400,000 active gateways globally. This opt-in infrastructure is most densely deployed in North America, western Europe, and eastern China.[1]

**BLE & Apple Find My**

In recent times, BLE is becoming a common technology for data backhaul. Apple AirTags [12] are one of the most successful examples, with the underlying Find My network made usable for developers through tools like OpenHaystack [31] and TagAlong [14]. Other consumer tracking applications like Tile [78], Life360 [48], and Samsung SmartTags [64] also successfully leverage BLE for backhaul. Similarly, industrial IoT users track cargo, vehicles, and products through Teltonika [77] and Milwaukee [50].

TagAlong [14] showed how Apple's Find My network can be used to leverage the Apple hardware ecosystem as 'data mules' for unbounded, arbitrary data transmission. While BLE has a comparatively limited transmission range, the sheer number of gateways (e.g. iPhones) dwarfs that of cellular and LoRa and can make up for this downside. Coverage exists ephemerally in the presence of Apple device users, and therefore is expected to be strong in urban areas but may suffer in more desolate and unpopulated locations.

[1]The Things Network is beginning to approach a similar scale of deployment. However, it has low adoption rates in North America, where most of the measurements are conducted.

### Cellular & Verizon

In cellular deployments, a centralized entity plans and deploys infrastructure to provide wide-area coverage in a calculated manner. Cellular providers are incentivized to deploy their infrastructure in locations where devices and people exist in order to gain more customers. Recent measurement studies have shown that across different cellular device platforms, as many as 33 failures (on average) can occur on a single device over a period 8 months, with failures lasting up to 3.1 minutes [46]. Xu et al. found that 5G can have spotty performance in outdoor environments, particularly when switching from an outdoor to an indoor environment [86]. Cellular also suffers from the highest device-side energy consumption and operating cost. Despite these struggles, given cellular's historical success over the past few decades, it is expected that cellular has the highest overall reliability between the three technologies.

Lastly, note that the major cellular providers in North America are roughly equivalent. Verizon is an arbitrary choice of convenience.

### LoRa & Helium

LoRa is a relatively new low-power wide-area network (LPWAN) technology that is emerging as a common option for backhauling data in IoT applications. With a transmission range and energy consumption comparable to and lower than cellular, respectively, LoRa is considered to be a more suitable option for IoT edge devices. LoRa's advertised range of communication is what makes this technology a compelling choice to backhaul data. Semtech [69] claims that LoRa coverage can reach up to 5 kilometers in urban environments and 15 kilometers in rural environments; they more recently state that LoRa can achieve a 30 mile range in outdoor environments [68].

While LoRa's range of communication is a highlight of the technology, this all hinges on the reliability of its supporting infrastructure. The largest LoRa gateway deployments are uncontrolled and crowdsourced in a manner similar to BLE. For instance, the Helium network uses cryptocurrency incentives to convince the general public to install gateways that

improve coverage areas and transit data. Given the general notion of LoRa's far range and robustness to interference in the wild, it is expected that LoRa performs equally to cellular's reliability. However, as a relatively new technology, LoRa lacks the proven track record of strong performance that cellular boasts. Musaddiq et al. deployed a measurement study of wetland conservation monitoring, and showed a packet delivery rate of 99.9% with nearby, available infrastructure [55]. It is not necessarily expected that the Helium infrastructure works as well for a mobile deployment.

### 4.2.2 Energy

Energy is often the defining feature shaping the deployment of mobile IoT. Optimizing for maximum battery life influences everything from where a system is deployed, to device form factor, and everything in between. It is therefore a primary consideration when selecting a technology for data backhaul.

#### BLE

BLE is one of the least energy intensive wireless technologies. DeBruin et al. measured that a BLE advertisement message consumes 70 μJ of energy from cold boot to transmit [21]. With such minor energy costs, commodity BLE devices are able to broadcast out BLE advertisement packets without needing to worry too much about taxing battery life.

#### Cellular

In exchange for higher reliability and wider area coverage, sending data via cellular requires more energy. Although more IoT-esque cellular options like LTE-M and NB-IoT attempt to address this issue, cellular is still on the high end of the energy consumption spectrum. Yang et al. measured the power demand of NB-IoT explicitly: at 3.3 V, an uplink transmit packet can use 0.92 J of energy [87]. From Table 4.1, cellular has a higher energy usage per byte by an order of up to 100x magnitude in comparison to LoRa and up to 1000x compared to BLE.

**Table 4.2.** Data collected on local and server side for each technology. Count values, timestamps and locations are collected from all platforms (except for those where it was unavailable).

| | Local | | | Server | | |
|---|---|---|---|---|---|---|
| | **BLE** | **Cellular** | **LoRa** | **BLE** | **Cellular** | **LoRa** |
| **Count** | Count value | Packet number | Count value | Count value | Packet number | Count value |
| **Time** | Time since boot | Unix timestamp | Unix timestamp | Unix timestamp | Unix timestamp | Unix timestamp |
| **Loc.** | - | Cell tower IDs | GPS | GPS (reporter) | GPS (cell tower) | GPS (gateway) |
| **Misc** | - | - | - | Report timestamp | - | SNR |

#### LoRa

LoRa sits between BLE and cellular in terms of energy per byte. At a fraction of the cost of cellular's power usage, LoRa is theoretically able to achieve similar range. While there is an extra layer of consideration for range versus data rate for LoRa (e.g. Spreading Factor, coding rate), LoRa appears to fulfill a nice balance between reasonable energy usage and range [63]. Mathur et al. performed LoRaWAN energy analysis for traffic sensing and found that for a 30 byte packet at SF12, roughly 93mJ of energy is consumed in the most extreme case. [2]

#### Summary

BLE is 1000x less energy than LoRa, 10000x less than cell.

## 4.3 Methodology

The end-to-end measurement system consists of a portable sensor platform shown in Figure 4.2 and a software pipeline described in Figure 4.1. The sensor platform contains three devices measuring the performance of BLE, cellular, and LoRa respectively: an ESP32-WROOM-32D, two different Android smartphones (Samsung A71 and Pixel 6a), and a Heltec CubeCell GPS-6502. Four total sensor platforms each consist of these three devices enclosed in a padded Pelican case to protect the devices from damage and impact. A battery is also enclosed[3] in the Pelican case to power the ESP32-WROOM-32D and CubeCell GPS-6502; the Android smartphones do not require an external power source. Table 4.2 summarizes the data collected

[2]Note that as a North America-based study, packets are limited to SF10 at about 370ms each.

[3]The safety and ethics of shipping batteries is discussed in Section .1.1.

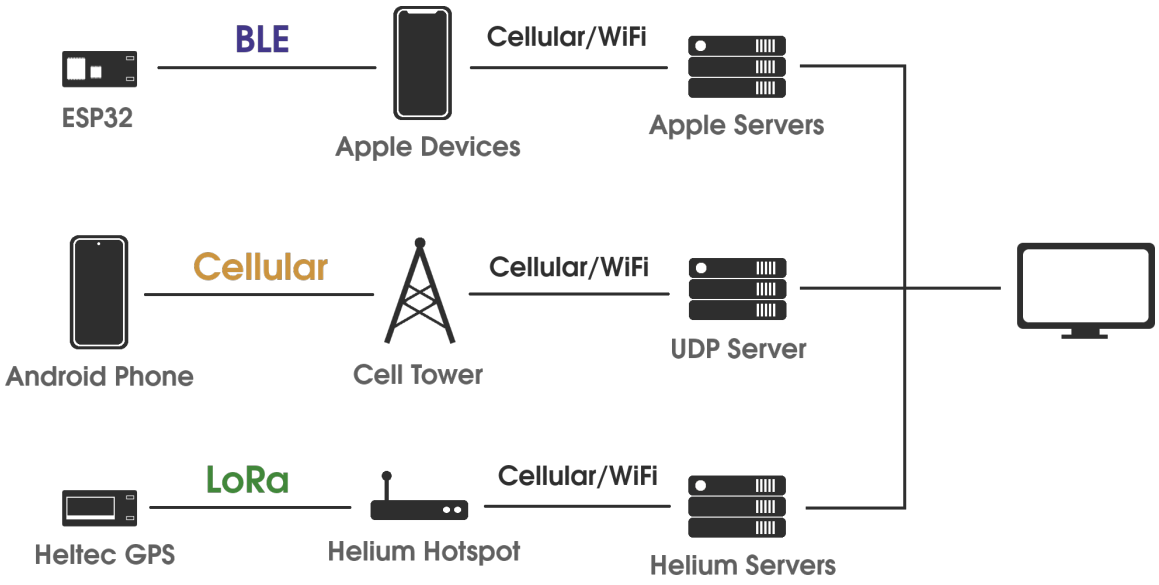


**Figure 4.1.** The data pipeline for sending and receiving BLE, cellular, and LoRa packets.

**Figure 4.2.** The sensor platform containing a BLE-enabled microcontroller, a cellular modem, and a GPS- and LoRa-enabled microcontroller.

from each wireless technology on both the local and server side of the software pipeline.

For country-scale shipments, data is sent every ten minutes to preserve battery life. While the sensor platforms are collecting data for up to a week at a time, it is also desired to get the most fine-grained data possible; ten minutes achieves a balance of device longevity and measurement granularity. For indoor- and transportation-based experiments, a sending rate of two minutes is chosen since battery life is less of a concern. This allows to collect finer-grained data than from the cross-country shipment experiments.

### 4.3.1 BLE Data Collection

Bluetooth data is sent via BLE advertisement messages using Apple's Find My protocol [13]. Leveraging existing Find My infrastructure, AirTags are mimiced with ESP32 devices running the open-source TagAlong firmware [14]. This gives greater visibility than would be afforded by standard AirTags, and all encrypted location reports from nearby Apple devices that "see" the BLE advertisements are collected. These location reports contain the broadcasted public key, the GPS location of the device that published the report, the timestamp of when the BLE advertisement was received, and the timestamp of when the location report was uploaded to Apple's servers. For BLE, (1) the count value (which corresponds to the public key broadcasted from the device) and (2) the time since the device booted is logged locally every time a packet is sent out. The Unix timestamp is logged manually before a deployment to eventually synchronize local time-since-boot with absolute time.

### 4.3.2 Cellular Data Collection

For cellular, Verizon and O2 networks (in the United States and Germany, respectively) are used to send out UDP packets to a desktop machine runing `tcpdump`. The UDP packets contain the MCC, MNC, LAC and cell tower ID of the cell tower the phone is connected to at the time of transmission. This information is also logged locally on the phone. Timestamps are logged when the packet is sent from the phone, and the timestamp when the packet reaches the

UDP server.

### 4.3.3 LoRa Data Collection

The Helium Network is used to measure wide-area LoRa deployment. In the LoRa packets, the following are sent: a counter value, the last available GPS location of the board, and a timestamp of when the packet is sent. This same information is logged locally and stored in flash on the microcontroller. Nearby Helium gateways receive the packet, and forward the data over the Internet to Helium's servers, where can be retrieved through Helium's Console API. All packet are purchased from all gateways that receive the LoRa packet, as opposed to the default of only purchasing from the first gateway to receive the packet. This allows to retrieve information about all gateways in range of the transmitting board. For each gateway, the LoRa payload, the name of the gateway, the timestamp of when it received the LoRa packet, the RSSI and SNR, and the gateway's (self-reported) GPS location are received.

LoRa provides more fine-tuned configuration to developers than BLE and cellular, most importantly with its Spreading Factor (SF) options. Spreading factor is a direct relationship between transmission robustness and payload size and energy; smaller payload sizes can result in higher packet delivery rate. As the focus of this chapter stresses reliability for each technology, the highest spreading factor available in North America—SF10—is chosen for all of the measurement studies.

### 4.3.4 Experimental Design

Eleven experiments are performed overall to measure existing wide-area wireless infrastructure. In five of the experiments, the sensor platform is shipped to different areas around the United States via Fedex; our team only touches the contents of the box at the source and destination. The packages are shipped via ground shipping and shipments last anywhere from two days to a full week. Two of the experiments are stationary "indoor" experiments, in which the sensor platform is left in the basement and third floor within the computer science building

**Table 4.3.** Packet Reception Ratio (PRR) percentage for BLE, cellular, LoRa and GPS for each measurement experiment. Experiments are grouped into "shipping", "indoor", and "transportation" categories in Table 4.3a, Table 4.3b, and Table 4.3c respectively. The backhaul technology (i.e. not GPS) with the best PRR for each experiment is shown in blue.

| | **CAM-LA$_0$** | **CAM-LA$_1$** | **CAM-AUS** | **SEA-NYC** | **NYC-CAM** |
|---|---|---|---|---|---|
| **BLE** | 79.4% | 79.8% | 13.3% | 12.3% | 40.0% |
| **Cellular** | **96.2%** | **98.1%** | **76.3%** | **65.6%** | **80.4%** |
| **LoRa** | 74.6% | 37.9% | 15.2% | 51.1% | 24.2% |
| *GPS* | *69.4%* | *41.9%* | *08.1%* | *12.3%* | *54.6%* |

**(a)** Shipping

| | **Building$_0$** | **Building$_1$** |
|---|---|---|
| **BLE** | **84.7%** | 36.6% |
| **Cellular** | 60.2% | 64.2% |
| **LoRa** | 0.0% | **96.6%** |
| *GPS* | *00.0%* | *00.0%* |

**(b)** Indoor

| | **US Train** | **DE Train** | **US City** | **Hike** |
|---|---|---|---|---|
| **BLE** | **98.8%** | **79.9%** | 77.2% | 24.1% |
| **Cellular** | 83.6% | 24.5% | 94.2% | 90.9% |
| **LoRa** | 64.8% | 23.7% | **98.3%** | **97.0%** |
| *GPS* | *95.9%* | *00.7%* | *49.8%* | *67.6%* |

**(c)** Transportation

for multiple hours. Finally, there are four "transportation" measurements, in which the sensor platform is manually carried along via foot and car (Hike), foot, bus, and light rail (US City), and heavy rail (US Train, DE Train).

Note that in each of the shipment experiments, this research is not in control of gateway deployments. That is, this research does not control whether there is a LoRa gateway, cellular base station, or Apple device nearby the sensor platform, nor do does this research plan paths based on observation of potential infrastructure coverage prior to deployment. This research do not adjust or interfere with any hardware in the measurement platform during experiments; the boxes were kept shut for the duration of each experiment. All of these experiments aim to

provide empirical results to gain more information on how these wireless technologies—and their respective infrastructure—perform in uncontrolled environments.

## 4.4 Cross-Technology Results

Given the measurements were performed in uncontrolled environments, there are myriad variables that affect the results of the measurements. This section first discusses each technology's reliability performance overall. Then it specifically address reliability as it relates to time (both within experiments and between experiments), location, and population.

### 4.4.1 Overall Reliability Comparisons

Aggregate reliability statistics across all of this work's experiments are first shown. Note that these aggregates hide underlying sources of skew, such as packages sitting in shipping depots for hours at a time. These instances are delved later in Section 4.4, but the overall results give an initial sense of the performance of each technology. Eleven measurements performed with the portable measurement platform and are categorized into three categories: "shipping," which captures moving and long distance scenarios, "indoor," which captures stationary and short distance scenarios, and "transportation," which captures moving and short distance scenarios. For the five shipment measurements, BLE has an average PRR of 38.2%, cellular an average of 72.2%, and LoRa an average of 37.4%. The average PRR for the two indoor experiments are 46.0% for BLE, 63.5% for cellular, and 77.7% for LoRa. Finally the average PRRs across the four transportation experiments are 67.3%, 90.0%, and 91.8% for BLE, cellular, and LoRa, respectively. Individual measurement results are shown in Table 4.3.

#### Shipping

Unsurprisingly, it is observed that cellular yields the highest reliability in the shipping experiments, where the path the sensor platform followed was in the control of Fedex, usually being transported along major highways and through cities. This affirms the notion that if an

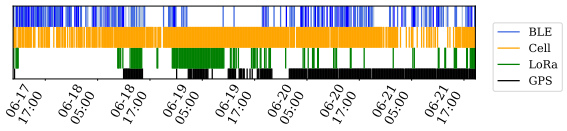


**(a)** NYC-CAM experiment from Table 4.3a.

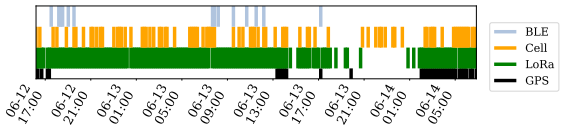


**(b)** A subset of the SEA-NYC experiment from Table 4.3a.

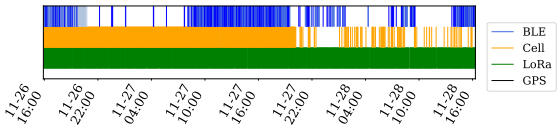


**(c)** Building$_1$ experiment in Table 4.3b. Note, this is an intentional deployment next to the Helium gateway in our computer science building.

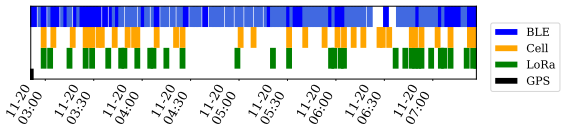


**(d)** DE Train experiment in Table 4.3c.

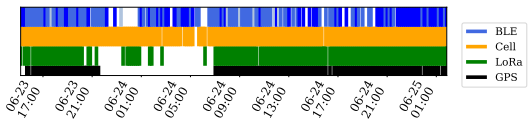


**(e)** CAM-LA$_0$ shipment from Table 4.3a.

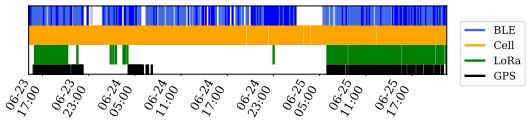


**(f)** CAM-LA$_1$ shipment from Table 4.3a.

**Figure 4.3.** Case studies of packets received over time, by technology. A colored rectangle indicates a packet was received server-side for that technology during that time slice. Since BLE packets do not have to be uploaded to the server immediately, the darker blue BLE rectangles indicate packets were uploaded with lower delay (darkest blue represents a latency of $<$10 minutes).

application requires high data reception reliability and there is limited knowledge or control of deployment locations, the sensible choice for data backhaul is cellular. Cellular is not without fault though, as there are still gaps in coverage. These can be seen in Figure 4.6 and Figure 4.3b, where cellular coverage starts to become spotty in the Pacific Northwest region, and LoRa and BLE are able to get packets through where cellular cannot.

**Indoor**

There are certain scenarios in which LoRa and BLE outperform cellular. In the first stationary experiment, the measurement platform was intentionally placed in a basement with intended poor cellular network connectivity. Here BLE prevailed, as the nearby Apple devices of computer lab visitors were able to receive packets and upload them to the server when they regained cellular or Wi-Fi coverage. The the latency of BLE backhaul is explored in Section 4.5.2.

For the second experiment, the measurement platform was left in an upper floor hallway near offices and meetings rooms. With direct visibility and close proximity to gateways, LoRa was able to provide strong coverage, while BLE struggled likely due to the short times that passing Apple devices stayed within range of the transmitter. Overall, cellular did not perform as strongly for the indoor stationary measurements as it did in shipment measurements, with BLE and LoRa instead taking the lead.

**Transportation**

All three technologies saw higher PRRs overall in the four transportation measurements, with BLE and LoRa performing the best. However, cellular is still fairly reliable, with the exception of the "DE Train" experiment in Europe. BLE achieved 80-100% reliability on the two long-distance train measurements, where the devices were within close range of hundreds of people for hours at a time. For the other two transportation measurements, LoRa was able to beat out cellular to provide the strongest reliability, although only by roughly 4-8% PRR reliability.

**Summary**

Cellular (unsurprisingly) performs the best on average, with BLE and LoRa shining in more niche situations. Understanding where, when, and why performance varies requires introspection into variation over time and geographical space.

### 4.4.2 Temporal Analysis

**Within-Experiment Analysis**

Figure 4.3 shows break out details of each technology's performance across the duration of selected experiments. Packets are binned into ten minute intervals and plot a bar within a time interval if a packet for a certain wireless technology was received in that interval. As a result, the more densely colored the row, the more reliable the technology. Figure 4.3a is representative of the common case, where cellular is the most reliable technology for the majority of the experiment. Figure 4.3b observes part of an experiment in which LoRa performs significantly better than cellular. This was when the package was shipped from one package distribution center in the Pacific Northwest to another. For the majority of the time, the package resided within a shipping depot where there was spotty cellular coverage, but LoRa was still able to get packets through.

This brings up an interesting question of whether a key opportunity for LoRa is for application users to deploy gateways that can compensate for a lack of cellular coverage. In Figure 4.3c, the measurement setup is placed next to a Helium gateway we deployed to provide coverage in a "spotty" cell area, with LoRa's performance obviously skyrocketing as a result.

While BLE provides good coverage during business hours, its diurnal pattern reveals the weakness of relying on human-attached data mules in spaces unlikely to have 24/7 occupancy. The root cause of cellular variability within the "indoor" experiments is unknown, but is qualitatively a consistent phenomenon according to building occupants. LoRa allows us to directly ameliorate this issue without involving a third-party infrastructure provider, simply by deploying

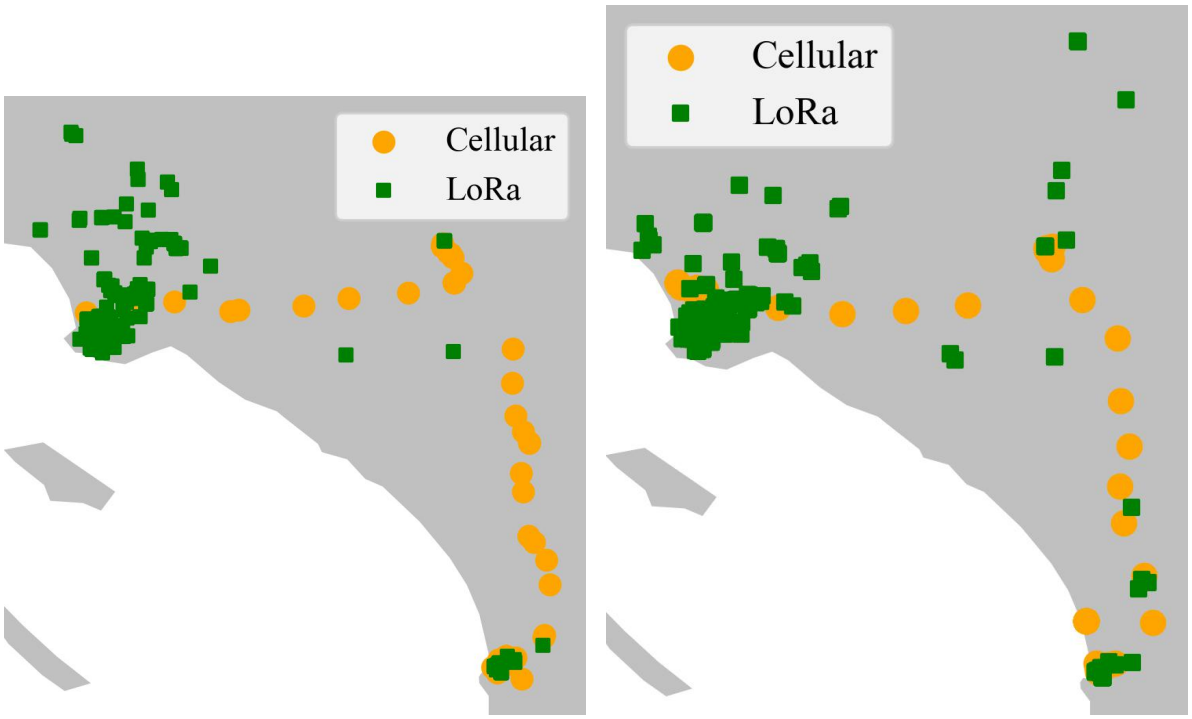


**(a)** CAM-$LA_0$ shipment displayed geographically. **(b)** CAM-$LA_1$ shipment displayed geographically.

**Figure 4.4.** Repeated shipments performed in southern California.

our own gateway.

### Between-Experiment Analysis

Evaluated next is the performance between repeated experiments from the same source to end destinations; while performance can vary across time within one experiment, it can also be drastically different in the same location across different experiments. This is particularly evident in the third shipment in Figure 4.6. This segment shows data collected in June 2023; an identical shipment was also performed in May 2023. In May, very few BLE packets were received in this geographical area, but in June, there was dramatically better coverage for the same region. The likely explanation, is the truck driver or a nearby travelling vehicle had an iPhone that was present throughout the entire shipment.

There are two takeaways from this specific instance. The first takeaway is that BLE can be surprisingly good, but highly unpredictable in terms of its reliability and coverage, even

within the same location. However, the second takeaway is that perhaps this is a merit of BLE and data-muling – it is not tremendously difficult to provide BLE coverage so long as there are people present. This does not work for the remote, desolate areas that are ridden of people, but for applications that are centered around high population density, BLE can be a reliable, low-power alternative to cellular. This second takeaway is shown in both train experiments shown in Figure 4.7b and Figure 4.7a, in which the measurement setup is physically surrounded by many people, some of whom happened to carry Apple devices with them. This "opportunistic" idea—in which hopefully there is a nearby Apple device that can backhaul mobile IoT data—can actually be quite deterministic when the edge device is in a location known to have high foot-traffic.

Another example highlighting the variability of performance in the same location over time is shown in Figure 4.3e and Figure 4.3f. In these experiments, two identical sensor platforms are sent from the same location at the same time, to the same destination. Table 4.3 observes similar reliability results between CAM-$LA_0$ and CAM-$LA_1$ shipments for BLE and cellular. However, there is a striking *40%* reliability difference for LoRa between these shipments. The reasoning for this difference in performance is actually less packets received at the final destination in LA, not via transit; this is affirmed through Figure 4.4, in which both figures show few packets received between the lower right and upper left regions.

### 4.4.3 Geographical Analysis

Figure 4.6 shows the "cross-country" performance of the backhaul technologies. Cellular covers the largest amount of the measurement paths, with BLE and LoRa seeing increased coverage mainly around cities. Cellular also had its weak spots, mainly in between California and the Pacific Northwest, and parts of the Southwest.

While most of the shipment measurements covered mainly flat and lower-elevation environments (i.e. where most highways are), the Hike experiment highlighted how different terrains can affect performance. Shown in detail in Figure 4.5, the Hike experiment involved the sensor

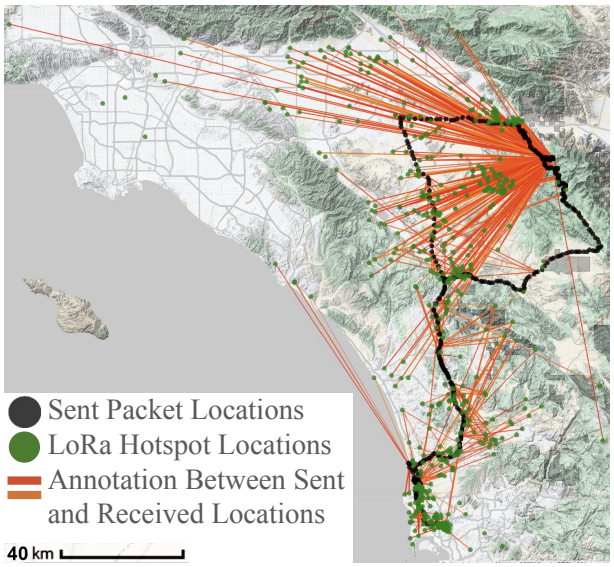


**Figure 4.5.** Map of packets received during a hike up a mountain near CA Metro. As shown in Table 4.3, LoRa achieve a higher PRR rate of 97% over cellular, which achieves 90% PRR. While these differences are not of great magnitude, LoRa happens to work well in clear line-of-sight conditions with limited wireless reception.

platform being taken up a 4,000 ft (1.2 km) mountain. From the top of the mountain, dozens of gateways were able to receive the packets from up to 100 km away. This success exemplifies the types of environments where LoRa shines– high elevation or places with otherwise clear line of sight, such as flat bodies of water or empty deserts.

### 4.4.4 Population Analysis

BLE works well in areas that have a high density of Apple devices that can backhaul data. As such, there is a natural correlation between population density and the likelihood that an Apple device is nearby. There is a correlation between the number of packets received in urban versus rural areas with Helium.

To dive further into the effect that urban versus rural locations have on reliability performance, the performance of the three technologies is compared to how densely locations are populated. The Kontur population dataset is used for 400 meter H3 hexagons from June 2022 [26], and then classify these regions as urban or rural, according to the US Department

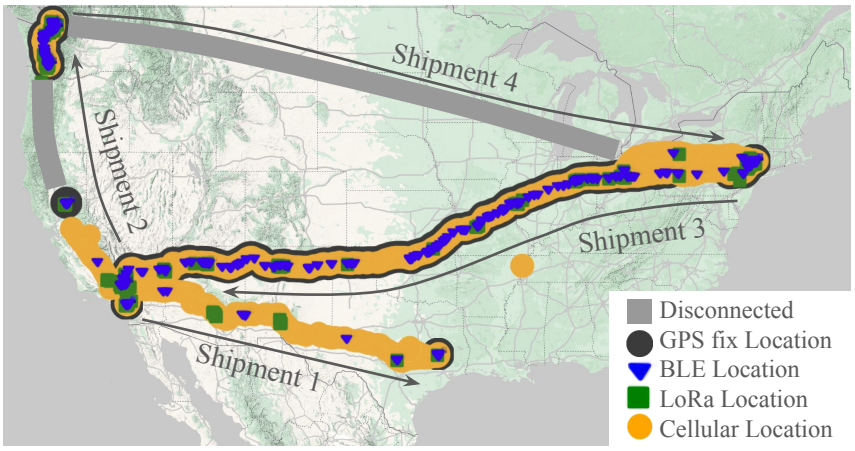


**Figure 4.6.** A map of selected shipments is shown here within the United States for BLE, cellular, and LoRa. The first shipment spans from southern California to eastern Texas (CAM-AUS in Table 4.3a). The second shipment spans from southern California to the Pacific Northwest. The third shipment spans from New York to southern California (NYC-CAM in Table 4.3a). The fourth shipment was mostly disconnected but started from the Pacific Northwest and ended in New York (SEA-NYC in Table 4.3a). The points overlayed on top of each other are the locations of the packets that were received; the different size of the points are only for improved visibility.

of Agriculture's (USDA) definitions. The USDA Economic Research Service specifies that an urban region is classified as either 2,000 housing units or 5,000 residents that live within one square mile of area [65].

As device-side GPS reliability is poor, and the location of the receiving gateway is used to classify measurement locations as urban or rural. Received packet locations are plotted for cellular and LoRa over a population density map in Figure 4.9a and Figure 4.9b for the CAM-$LA_1$ experiment in Table 4.3a; BLE is not mapped as the Apple devices that act as gateways are mobile. While cellular works regardless of environment, LoRa packets sporadically get through with seemingly little relation to whether the measurement location is urban or rural.

This is further investigated, starting with overall statistics of what environments packets are received in for each wireless technology in Table 4.4. BLE and LoRa are both more likely to be received in urban areas over rural areas, while cellular appears to receive packets in both environments relatively equally. Surprisingly, LoRa appears to be more heavily influenced by urban versus rural than BLE.

**Table 4.4.** Number of packets received in urban and rural environments. It would be ideal to analyze where packets were *sent*, but GPS location is not always available. Therefore statistics on received packets are recorded. BLE and LoRa packets are primarily received in urban environments, while cellular receives more packets in rural environments than urban. While there is a one-to-one ratio with packets sent and received for cellular, there is a one-to-many ratio of packets sent and received for BLE and LoRa.

| | Total | Urban (%) | Rural (%) |
|---|---|---|---|
| BLE | 91498 | 60020 (65.6%) | 31478 (34.7%) |
| Cellular | 16661 | 6773 (40.6%) | 9888 (59.3%) |
| LoRa | 34448 | 25251 (73.3%) | 9197 (26.7%) |

## 4.5 BLE (Apple) Versus LoRa (Helium)

While cellular works the best reliability-wise, BLE and LoRa are both lower power options that sacrifice reliability for device longevity. As a result, the tradeoffs between BLE and LoRa are explored more concretely to determine which one is "better". When looked at carefully, the choice to use BLE over LoRa might be more obvious than expected.

### 4.5.1 Reliability and Energy Favors BLE

Revisiting the results in Table 4.3, BLE is just as reliable—or even more reliable—than LoRa in seven out of eleven experiments. Although the difference in reliability between these technologies is not enough to choose one over the other, look back to the contexts in which these technologies are used. When determining which technology to choose for device longevity, BLE has a clear advantage. As originally discussed in Section 4.2.2, one LoRa packet is roughly energy equivalent to over 1,000 BLE packets.

This lower energy cost also comes with a benefit of "cheap" packet retransmission. In Section 5.2, remember that the comparisons between BLE, cellular, and LoRa were all performed on a one-to-one packet transmission ratio. As a result, even if multiple packets were transmitted for BLE for every LoRa packet, BLE's edge as a low-power, wireless technology still holds true in comparison to LoRa. Sending more packets can also influence reliability in the event that the first packet is not received. Zachariah et al. find that in a stationary environment, only

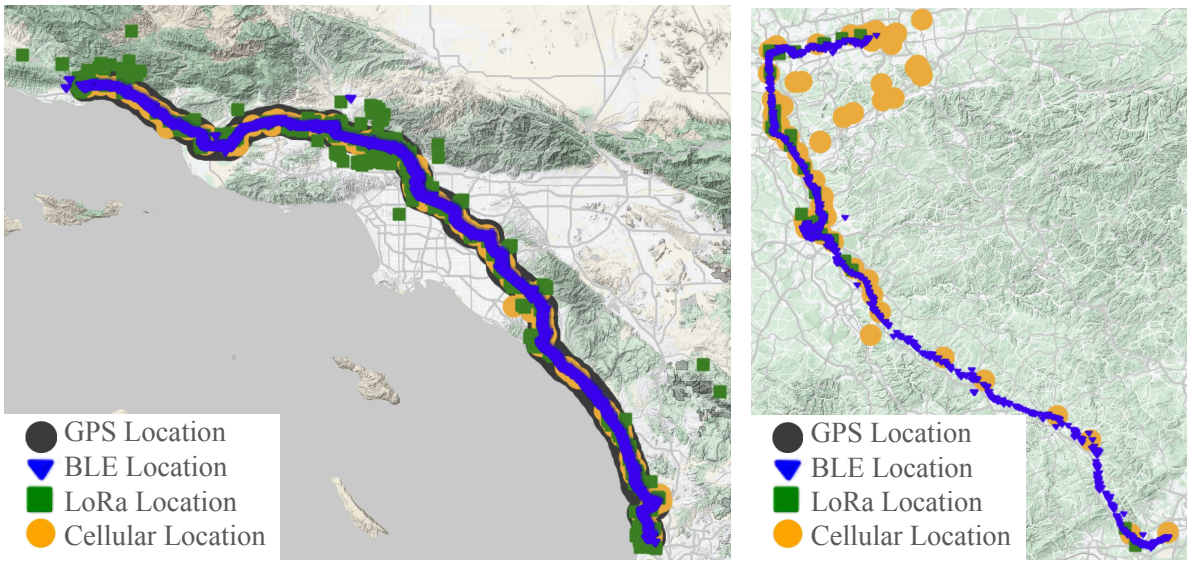


**(a)** Packets received during a train ride in southern California.

**(b)** Packets received during a train ride across North-Rhine Westphalia in Germany.

**Figure 4.7.** Train measurement data collected in southern California and Germany.

three repetitions of the same packet are required to achieve a Data Reception Rate (DRR) of over 99% [89]. Similarly, Tagalong finds in a small-scale study that five repetitions of the same packet achieve a DRR of 95% in a mobile setting (versus a DRR of 65% for a single packet transmission) [14].

Furthermore, BLE already works in areas that LoRa—via Helium—tries to provide coverage in. In Section 4.4.4, it is shown that BLE's and LoRa's reliability is influenced by population density and observe that most packets are received in urban versus rural areas. This extends generally to places where people tend to be present (e.g. trains), as highlighted in Figure 4.7. Shown more concretely in Figure 4.3d and the DE Train PRR results in Table 4.3c, BLE can be a relatively reliable technology given that it can backhaul data in environments more densely populated with people and without traditional network coverage. As such, considering that BLE already works in these areas well and has a clear advantage energy-wise, BLE presents itself as the favored choice to backhaul mobile IoT data over LoRa in the commmon case.

### 4.5.2 Packet Latency

The functionality of Apple AirTags is split into two processes: (1) an Apple device picking up a nearby BLE advertisement and (2) processing and uploading this data to the cloud. There is often a delay between these two steps, as Apple devices can receive BLE advertisements at any time, but may not always have service in order to upload the data. Hence, data backhauling via Apple Find My has a different latency profile than traditional networking. Figure 4.8 shows a CDF of BLE latency for all of the measurement data (capped at twenty-four hours). Roughly 80% of all the data is received within an hour; the median value is roughly fifteen minutes. After that, there is a long tail of decreasingly infrequent, latent packets. The most latent packet in the collected dataset arrives four days after its transmission.

Applications that use Find My and similar BLE systems to backhaul data need to be relatively latency tolerant. Time-sensitive applications (e.g. monitoring dangerous materials for spills) are not suited to this style of backhaul infrastructure, as there is no guarantee that information would arrive promptly. For applications that perform long term data collection where no immediate actions need to be taken, Find My-esque backhaul systems are more appropriate. Also, note that latent packets are more valuable than no packets at all. In Figure 4.3d, most of the received packets for BLE are lighter colored, meaning that packets are delayed when received via Apple's Find My servers. In comparison to cellular via the O2 network, latency can also be a benefit and perk of this methodology to backhaul mobile IoT data.

## 4.6 Is the Problem LoRa or Helium?

LoRa (via Helium) performed worse than anticipated given the volume of supporting infrastructure. Before diving further into discussion of whether the problem lies within the technology's use or the infrastructure's deployment, this section discusses how Helium entices users to participate. Helium's policies incentivize users to provide coverage in areas in which coverage is not redundant. However, a key observation from the analysis is that the highest

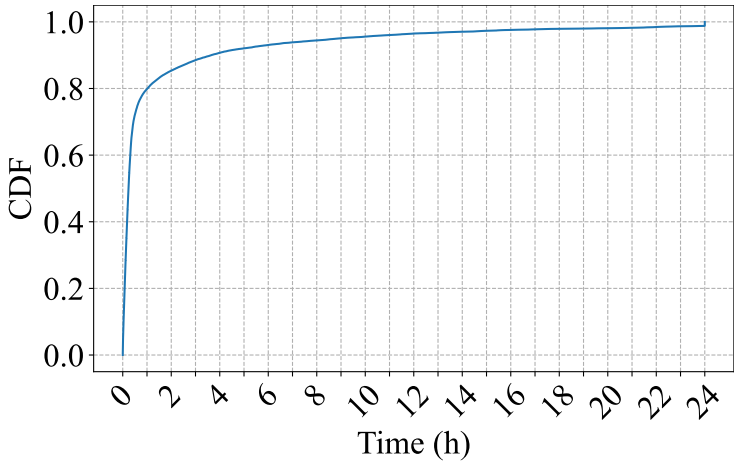


**Figure 4.8.** BLE latency of all the data collected which consists of approximately 140,000 packets. Latencies are capped at twenty-four hours for display purposes.

reliability is shown to be in major cities. As a result, it is important to consider whether reliability is tied to areas with higher population density, and whether Helium is providing the wrong incentives [38, 39] to provide widespread, reliable coverage and infrastructure.

To determine whether the problem lies in the wireless technology or infrastructure more concretely, data from a specific case (i.e. CAM-$LA_1$ from Table 4.3a) is analyzed to better understand the failures in LoRa (Helium) reliability through comparisons to cellular (Verizon) reliability.

### 4.6.1 Infrastructure and Coverage Analysis

In Section 4.4.4, LoRa's (Helium's) data reception is closely tied to urbanized areas, whereas cellular packets are sent and received in areas regardless of population density. To highlight this observation more visually, all received packets are plotted on a physical map with underlying population environments of urban or rural in Figure 4.9a and Figure 4.9b for cellular and LoRa, respectively. Here, more packets are received in rural locations for cellular, whereas more packets were received in urban areas for LoRa. The next question aims to answer whether these areas with human population contain supporting infrastructure for cellular and LoRa technologies as shown in Figure 4.9c and Figure 4.9d, respectively, and find that there is

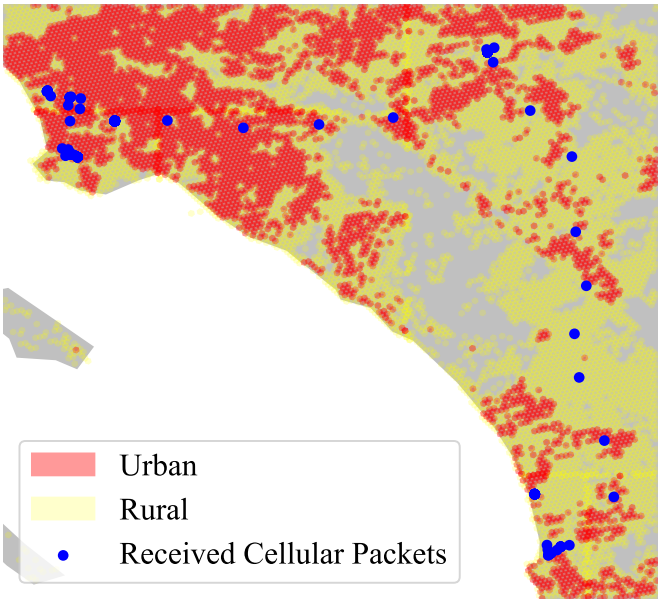


**(a)** Cellular overlayed on a population map. Here, 79 packets were received in urban areas, and 351 packets were received in rural areas.

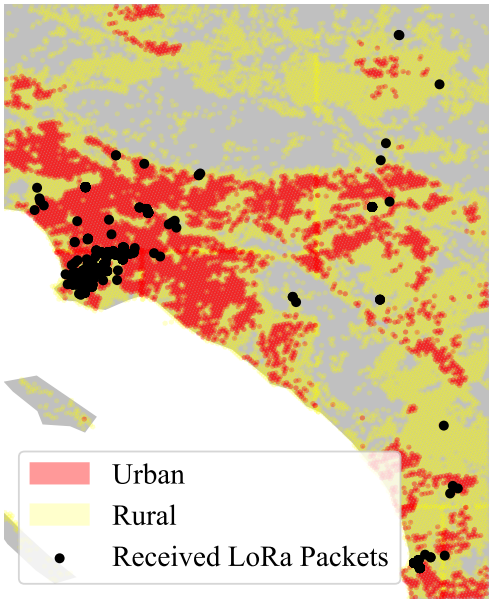


**(b)** LoRa overlayed on a population map. Here, 770 packets were received in urban areas, and 220 packets were received in rural areas.

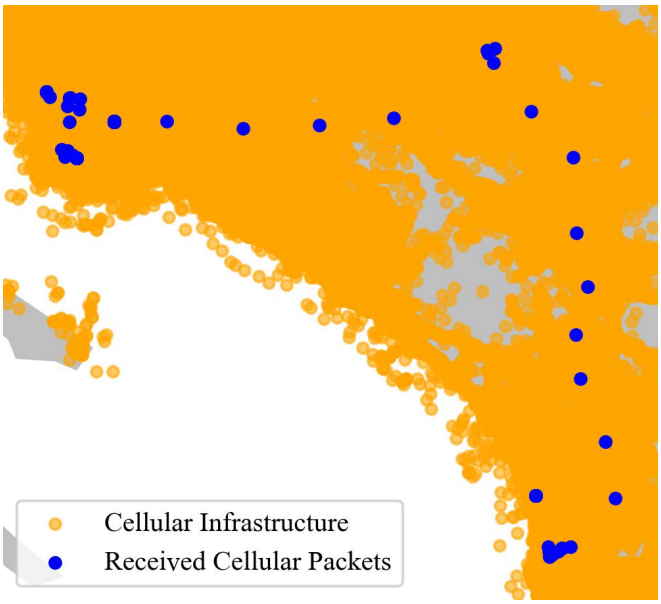


**(c)** Cellular points overlayed on cellular infrastructure.

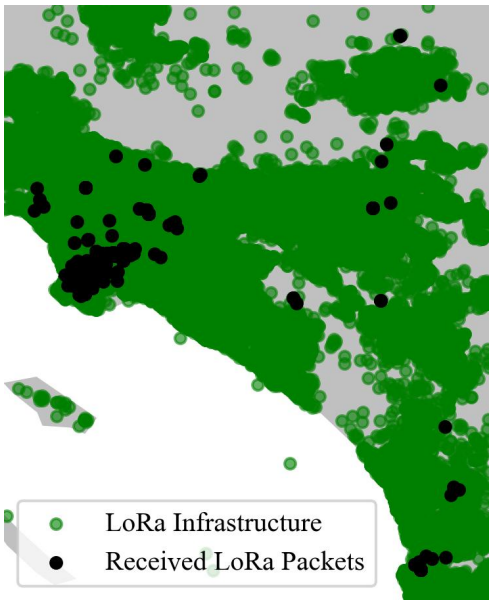


**(d)** LoRa points overlayed on Helium infrastructure.

**Figure 4.9.** Received data for cellular and LoRa in southern California (CAM-LA$_0$ experiment in Table 4.3) overlayed on a population density mapping.

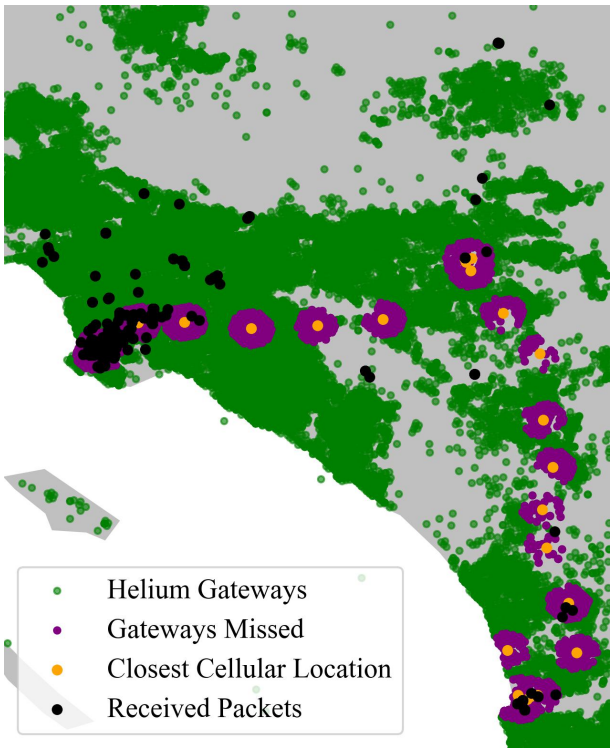


**Figure 4.10.** Missed coverage of LoRa in southern California overlayed on a coverage map. While the data was collected in June 2023 and the underlying coverage map is a snapshot from April 2024, there is an assumption that the coverage is not drastically different between these time frames.

indeed supporting infrastructure for both technologies. However, the results in Table 4.3 observes that while reliability for cellular is at 98.1% PRR, reliability for LoRa is only 37.9% PRR.

The next hypothesis is that available coverage does not necessarily imply reliable coverage. Cellular and LoRa coverage maps were obtained from OpenCelliD [45] and Helium's Hotspot API [35]. All received packets are overlayed on these coverage maps in Figure 4.9c and Figure 4.9d, respectively. While the map shows ample coverage for LoRa, many packets were not received in areas that were theoretically well-covered. To look at the magnitude of this issue, missed coverage is examined in Figure 4.10. Since not all sent packets have GPS available, a sent packet's locations is approximated based on the nearest cell tower. It is assumed that cell towers are deployed closely enough that when a cellular packet is received, the packet's sent location is nearby.

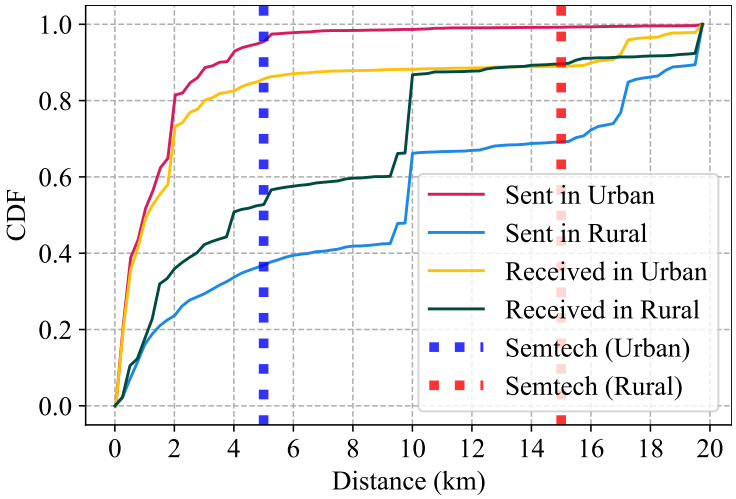


**Figure 4.11.** CDF of packets that are sent and received in urban and rural environments. Semtech specification of urban communication distance is approximately five kilometers, shown as a blue dotted line, and fifteen kilometers for rural areas shown as a red dotted line. In addition, an upper limit of twenty kilometers is induced to better elucidate what might realistically be achieved in uncontrolled environments. Table 4.5 includes all of the data and does not cap at twenty kilometers.

**Table 4.5.** Approximate distance in sent and received urban and rural environments. All measurements are taken from known sent locations with available GPS location and received gateway locations.

| Environment | # of Samples | Median (km) | Mean (km) | STD (km) |
|---|---|---|---|---|
| Sent in Urban | 8762 | 0.733 | 1.377 | 2.848 |
| Sent in Rural | 3138 | 10.125 | 11.376 | 16.766 |
| Received in Urban | 9100 | 0.742 | 2.833 | 6.869 |
| Received in Rural | 2798 | 4.194 | 7.752 | 11.505 |

From this approximate location, highlighted in purple are all gateways that were nominally within range of the sending device but did not receive the LoRa packet. This nominal range is based on Semtech's claim of a transmission range of approximately five kilometers in urban areas [68].

The expectation or hope is that at least one gateway in the purple region is able to receive the LoRa packet, yet there were none. In fact, within each five kilometer radius around the approximated sent packet's location, there are hundreds of available gateways, yet many packets that were sent were never received.

In an attempt to explain this wide chasm between expectation and reality of LoRa performance, this work takes a deeper look into the received data. Historically, Semtech claimed that LoRa is able to have communication distances of up to five kilometers in urban areas and fifteen kilometers in rural areas [69]; more recently, Semtech claims that LoRa should communicate up to almost fifty kilometers in outdoor environments [68]. Through empirical analyses of these claims, the median communication ranges are all below Semtech's original claims and expectations, which is shown in Figure 4.11.

These numbers are more concretely shown in Table 4.5, where median and mean values for communication in urban environments are only about one kilometer, and the median and mean values for communication range in rural environments are approximately ten kilometers. These limited communication ranges partially explain the missed coverage in Figure 4.10; a five kilometer radius is perhaps too tall of an expectation. When the same analysis is run for one kilometer coverage areas, the number of available gateways dwindles down to a few dozen at best.

Overly optimistic range claims do not fully account for the difference in performance, as there is still coverage available where most of these packets are sent. The PHY layer characteristics and then investigated in the received data, plotting the SNR of all received packets in Figure 4.12. When packets are received SNRs are low, with the peak number of SNRs residing between -10 and -15 SNR. This shows that although packets are received, they are just barely making it through. Since Helium infrastructure is already being pushed to its PHY layer limits, the results in Figure 4.10 make more sense.

### 4.6.2 Discussion

#### Who Is At Fault—LoRa or Helium?

LoRa underperformed our expectations. This sub section tries to answer whether this is an artifact of how Helium deploys infrastructure or is something more fundamental to LoRa as a communication technology. As Helium invite arbitrary lay-people to deploy gateways, dubious

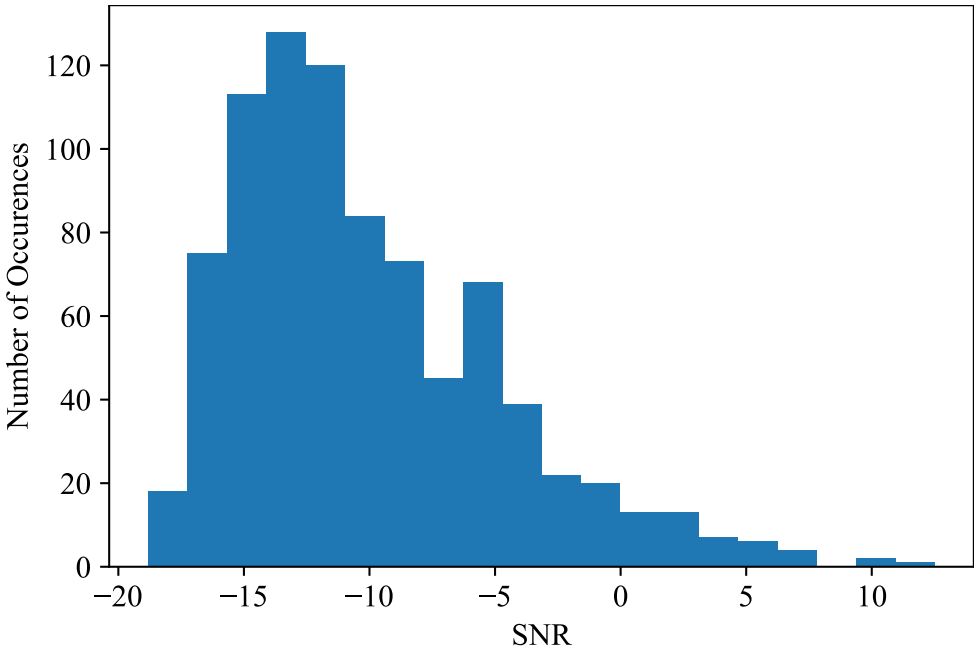


**Figure 4.12.** Histogram of Signal-to-Noise Ratio (SNR) of all LoRa packets received for the CAM-LA$_1$ shipment in Table 4.3 and Figure 4.9.

deployment could be the root cause of poor coverage reliability. However, the odds that not one Helium gateway in these regions was deployed well—amidst *hundreds of active* gateways—that did not receive the packet is an unlikely situation.

Indeed, *CityWAN* is a research project that deploys a more controlled case study of LoRa gateways within a defined area to provide coverage for static, IoT devices [79]. They observe a similar result, where some nodes simply cannot be covered by anything shy of a dedicated, adjacent gateway. Likely, the urban canyon effect is so strong that the coding gains of LoRa over cellular are ultimately moot in real-world urban environments, and thus infrastructure requirements for LoRa is functionally the same as the extant cellular infrastructure. As a result, since the cellular infrastructure already exists and the only benefit that LoRa potentially provides is lower communication power—which emerging cellular standards aim to improve upon—is building new infrastructure to support LoRa coverage really worth it?

**What is the Right Way to Deploy LoRa?**

Given that (1) deployments in dense, populated areas suffer from a variety of factors both through physical and wireless interference and (2) BLE and cellular can already cover these areas relatively well if energy efficiency or reliability is valued, respectively, perhaps the LoRa community should re-imagine the applications in which LoRa is really useful.

One uniquely successful use case for LoRa is agriculture applications [28, 42, 27, 61, 72, 11]. One (or few) gateway(s) can cover multiple kilometers of area for agriculture sensors. Commercial cellular coverage follows customers, and a landowner cannot easily add their own cellular coverage where companies do not see value in building out infrastructure. LoRa equipment is comparatively cheap and easy to build and deploy. Indeed, the $\text{Building}_0$ dataset in Table 4.3 "fills a gap" in local cell coverage with a LoRa gateway we deployed. Perhaps then general purpose backhaul for mobile IoT is not the best use case for LoRa, particularly since the effort to provide last-mile connectivity for every device can be cumbersome and inefficient [79] and ultimately simply duplicates existing cellular infrastructure.

## 4.7 Summary

This chapter find that the current best choices to backhaul mobile IoT data are BLE and cellular for device longevity and reliability, respectively. With respect to reliability, while it is clear that cellular (via Verizon) is the best choice, BLE (via Apple) is surprisingly effective today and is even the preferred choice over LoRa. While there is much enthusiasm for LoRa as an emerging infrastructure for IoT, our results suggest that in many cases, LoRa via the Helium Network is not actually necessary. LoRa does fill a unique niche, however, as an end-user-deployable, wide-area wireless solution for the case where sufficient coverage does not already exist.

The next chapter will discuss how non-terrestrial infrastructure can enable the LoRa technology as well as new applications that can benefit from wide-area coverage from the above.

# Chapter 5

# LoRa on Non-Terrestrial Infrastructure

In this chapter, we discuss the empirical measurements of LoRa on non-terrestrial infrastructure (i.e. balloons, airplanes). Specifically, we discuss within the context of a single application: wide-area broadcast messaging. Our measurements are only unidirectional, i.e., we only transmit from the deployed device on non-terrestrial infrastructure. Hence, we can only discuss our measurements from the context of transmitting (i.e. broadcasting) from the air. However, this chapter aims to showcase the opportunity and advantage of deploying LoRa on non-terrestrial infrastructure.

## 5.1 Background

We provide background information about broadcast systems, what has been done with LoRa technology and broadcsat.

### 5.1.1 LoRa

LoRa is a relatively new LPWAN technology, commonly used to backhaul data in IoT applications. LoRa's appeal lies in its ability to provide a transmission range comparable to cellular technology, but with lower energy consumption [3]. These benefits come at a cost, namely, lower data bandwidth ranging from bits- to kilobits-per-second [66]. The tradeoff to lower bandwidth rewards with long transmission ranges and greater robustness to noise and signal interference.

There are various factors that affect the resilience, range, and reliability of LoRa. Besides standard radio frequency (RF) factors like coding rate, transmit power, and antenna gain, a core feature of the LoRa PHY layer is Spreading Factor (SF) [62], which can be varied to favor either robustness or data throughput. LoRa provides six SF values to select from: SF7 to SF12 in integer increments. Generally, a higher spreading factor (e.g. SF12) results in more redundancy per payload bit, allowing for larger transmission ranges at the cost of lower data throughput. Conversely, a lower spreading factor results in less redundancy per payload, restricting to smaller transmission ranges but enabling higher data throughput.

## 5.2 Methodology

Our goal is to evaluate the feasibility of using LoRa transmitters to rapidly deploy wide-area data broadcasts across entire metro areas. We assume that the operator has pre-defined messages they want to broadcast to an entire metro area (e.g., “Do not drink the water”, or ”Go to this location”). We evaluate the inherent broadcast capability that LoRa’s chirp-spread spectrum robust modulation provides using a COTS LoRa transmitter, with a COTS antenna, and high-redundancy spreading factor configuration to allow for extremly low SNR (-10 dB) reception. We then measure the range and reliability of rapidly deployed LoRa broadcasts by carrying the transmitter on two different airborne vehicles. These broadcast to LoRa receivers on the ground in an existing wide-area deployment of $\sim$7,000 LoRa base stations that report received packets to the cloud.

We now describe this measurement setup, and what conclusions about broadcast capability it will allow us to make.

### 5.2.1 LoRa Broadcast Transmitter

We evaluate the reach and reliability of wide-area LoRa transmissions using a standard, regulation-compliant LoRa radio with its stock omnidirectional 3 dBi stub whip antenna. The goal is to evaluate if the LoRa protocol’s long-range, robust, but low-power physical-layer

protocol is sufficient to provide wide-area coverage. The specific hardware platform we used is the Heltec CubeCell GPS-6502 board [4], which can transmit LoRa packets and simultaneously record GPS location. The CubeCell board runs firmware that sends out a LoRa packet every 6–7 seconds.

Our goal is to assess how the LoRa protocol's unique low-SNR modulation can be used for wide-area broadcast communication. Therefore, we keep its transmission parameters mostly in their default configuration, specifically 125 kHz bandwidth, coding rate of $4/5$, and transmit power of 20 dBm. We also follow the protocol's recommended settings for long-range communication; we transmit packets with spreading factor SF10, which is the highest-available spreading factor for use in North America [58]. To determine which transmitter location and altitude lead to successful reception at each receiver, the firmware encodes the transmitter's real-time GPS location and altitude in each packet it sends.

### 5.2.2 Airborne Broadcasting Platforms

We study two different rapidly deployable airborne vehicles as platforms for LoRa broadcasting: (1) an unmanned balloon flying over San Diego, California, and (2) a private plane flying over Las Vegas, Nevada.

#### Weather Balloons

We evaluated the most inexpensive wide-area high-altitude broadcast platform that was practical by launching the LoRa transmitter on unmanned balloons. We flew three LoRa unmanned balloon deployments over southern California between February and May of 2025; a picture of one of our balloon deployments is shown in Figure 5.1. The first advantage of a weather balloon is that it can be launched at any time, quickly (i.e. filling the balloon with helium takes only a few minutes), and inexpensively (i.e. each launch costs about $200 in total). The second advantage is it can be carried to an extremely high altitude: reaching an altitude of at least 15 km (i.e. lower stratosphere) was common in our experiments. The first challenge


**Figure 5.1.** Weather balloon broadcast platform. We attached a LoRa transmitter to three weather balloons and flew them up to 15 km above San Diego, California. The climb rate was 0.020 meters per second.

with weather balloons is they have a very limited payload capacity; if the payload is too heavy they will not fly at all. Fortunately, we discovered that a LoRa transmitter and its batteries is a typical payload for weather balloons; our LoRa broadcasting setup (e.g. LoRa transmitter, encasement, batteries) is only an 80 g payload, similar to a weather radiosonde which is about 60–80 g. The second challenge is that balloons' climb rate is dictated by the payload capacity and the volume of gas it holds; the longer it takes to climb, the longer it will take for broadcasts to reach the entire metro area it intends to cover. We found that approximately twenty cubic feet (about 585 liters) of helium could lift our 80 g LoRa setup at a rate of 20 centimeters per second. The third challenge with weather balloons we observed is that at above 15 km, a typical LoRa transmitter and its batteries will get too cold to operate reliably. Therefore, the bulk of the balloon's broadcast capability comes from its initial launch period up until it reaches 15 km, at which point the LoRa transmitter often just shuts down until the balloon reaches a high enough altitude ($\sim$20 km in our study) that it pops and falls to the ground. The final challenge with weather balloons is they are uncontrolled – each launch yields a different flight pattern. Therefore balloons may not provide consistent and reliable broadcast coverage. Additionally, the LoRa


**Figure 5.2.** Small aircraft broadcast platform. We carried a LoRa transmitter on a Cirrus SR22 and flew it around Las Vegas, Nevada at altitudes from 1.7 km–3.3 km. The climb rate was 6 meters per second.

antenna is constantly tumbling because it is just tethered by the balloon with a string, which can cause an inconsistent radiation pattern aimed at the ground, resulting in a low Packet Reception Rate (PRR) for broadcast receivers.

**2. Aircraft**

We evaluated LoRa at various altitudes with a controlled flight at altitudes of 1.7 to 4.5 kilometers in Las Vegas using a private plane (Cirrus SR22 Turbo) in April 2025. The primary advantage of aircraft is their predictable flight pattern and altitude; these make it possible for broadcasts to have more consistent coverage and reliability. Indeed, even the LoRa antenna of the transmitter is stable (rather than tumbling like on the balloon string). Unlike balloons, small aircraft are expensive to operate at $624 per hour, they also require an available aircraft and a trained pilot to be available when the broadcast is needed, and this could limit their availability for emergencies.

**Regulations**

We followed all local regulations for all aerial experiments; for instance, prior to balloon deployments, we submitted a Notice to Air Missions (NOTAM) to the US Federal Aviation Administration (FAA). Further discussion of the ethics of transmitting from aerial platforms are

discussed in Section .1.2.

### 5.2.3 Wide-area LoRa Reception Testbed

To measure the coverage and reliability of LoRa broadcasts, we co-opted a large deployment of LoRa receivers that can receive our broadcast packets and collect them centrally. Specifically, we use ∼7,000 Internet-connected LoRa gateways in Southern California and Las Vegas in the Helium Network that were intended to be used as public infrastructure for IoT connectivity [41, 30]. Helium gateways automatically forward the packets they receive back to an Internet server; we obtain all received packets from all gateways using Helium's cloud API. Helium allows us to receive the same packet from multiple gateways simultaneously by enabling the option to purchase packets from all gateways, which we enable over the default of only purchasing from the first gateway to report a received packet. For each gateway, we collect the received LoRa packet's signal quality including RSSI and SNR. We also record each received packet's location context, including the name of the gateway, the timestamp of when it received the LoRa packet, the RSSI and SNR, and the gateway's (self-reported) GPS location. A gateway only reports a captured packet if it was received entirely correctly.

**Baseline**

Helium gateways are crowdsourced so they may not be all operational [41]. Therefore, we need to establish a baseline of what receivers are functional to receive broadcasts during our experiments, especially for evaluating the completeness of the receivers covered by airborne broadcasts. We establish this baseline by wardriving with the same LoRa transmitter used in the aerial setup in San Diego and Las Vegas. Then, we observe what Helium gateways receive these transmissions from the ground, they serve as our baseline set of 596 active receivers in San Diego and ∼528 in Las Vegas.


**(a)** Receiver coverage in San Diego from a single weather balloon deployed in the San Diego area.

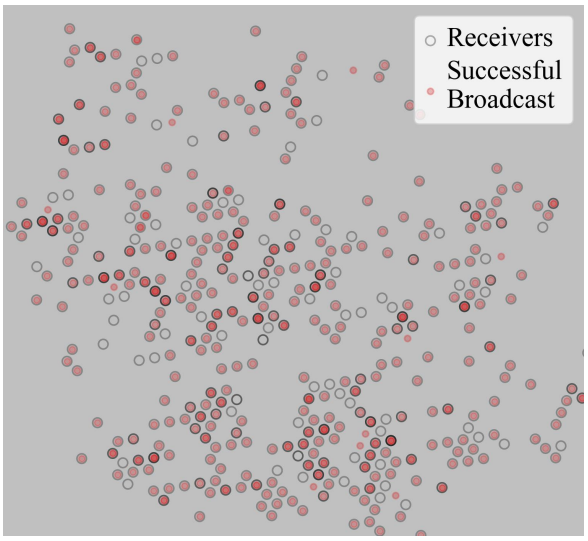


**(b)** Receiver coverage in Las Vegas from an aircraft flying in the area.

**Figure 5.3.** Coverage over San Diego via balloon and coverage over Las Vegas via aircraft. The black, hollow circles represent active receivers that were validated from our driving data collection, and the red dots represent a successful broadcast that was received by our transmitting platform. The alpha values were set to 0.3 each to represent overlapping receivers within the same area.

## 5.3 Results

To evaluate how well LoRa can work for wide-area, emergency broadcast messaging, there are three central questions we need to answer: (1) how full of coverage can we expect to achieve, (2) how timely can we achieve full coverage, and (3) how can we ensure that our broadcasting infrastructure can be reliably received by receivers below? We discuss these results with respect to weather balloon and aircraft broadcasting infrastructures in San Diego and Las Vegas, respectively. In all of our results, we evaluate our data based on the ground truth data we collected via war-driving.

### 5.3.1 San Diego and Las Vegas City Coverage

We first aim to determine how full of coverage we can achieve via a balloon and aircraft. Figure 5.3 shows coverage from both San Diego and Las Vegas, in which we used balloons and an aircraft to measure the percent of coverage from actively participating receivers. For San Diego, we achieve 74.8% of coverage for our first balloon experiment and then 73.6% and 73.5% of coverage for our second two balloon experiments, respectively; the total number of receivers we compare to for this coverage experiment is 596. We find that there is roughly 88-89% overlap in receiver similarity between our first deployment and the second two each. Between the two balloon deployments in May, we see roughly 99-100% overlap in receivers that heard broadcast messages from both balloons; only launched thirty minutes apart, the balloons traveled similar trajectories. The similarity in overlapping coverage between each of the balloon indicates that that similar coverage is attainable despite the balloons traveling in both different and nearly identical paths.

We then observe what percentage of coverage we can obtain via an aircraft to see how well an aircraft can perform at providing broadcast coverage down to receivers below. The primary difference between our aircraft and balloon deployments is that we can test the attainable coverage on a controlled path at different altitudes. Adding altitude as a variable factor allows us to better understand what altitude(s) we should ideally fly at to achieve the best coverage possible. For Las Vegas, we fly at three different altitudes—1.7 km, 3.3 km, and 4.6 km—and achieve 68.4%, 77.8%, and 79.7% of city coverage, respectively; from driving, we validate a total of 528 receivers. While it appears that the best coverage is achieved from the highest flying altitude of 4.6 km, we also observe that the difference in coverage between the aircraft at 3.3 km altitude is not too different from the coverage at 4.6 km. In fact, when we compare the subset of receivers between the two coverages, we find a 97-100% overlap in receivers that heard our broadcasting messages at both altitudes. The largest difference is observed when we compare the coverages attained at 1.7 versus 3.3 km altitude and 1.7 versus 4.6 km, in which we see a

difference of roughly 50-60 of receivers that were missed at 1.7 km altitude but hit at 3.3 and 4.6 km altitudes. As a result, we observe is a more significant coverage gain between an aircraft flying at 1.7 and 3.3 km altitude versus 3.3 and 4.6 km altitude. From this finding, we can see that an altitude between 3.3 and 4.6 km is likely enough to achieve the best achievable coverage, and we validate this observation by confirming that there is roughly a 97-100% overlap of receiver between the achievable coverages at 3.3 km and 4.6 km.

We also notice that there is a relatively random distribution of receivers that never hear our broadcast messages at all. This is shown geographically in each city in Figure 5.3a and Figure 5.3b. In general, there does not appear to be a particular area that suffers from poor reception; all the receivers that never received a single broadcast packet appears to be relatively distributed across the city. However, there is consistently a subset of receivers that never hear our broadcasted messages despite multiple, repeated trials. There are various hypothetical reasons for this occurrence, such that Helium receivers go offline at a rate unknown to us, Helium receiver are not designed to receive LoRa packets from the sky, or that some gateways are inherently in hard-to-reach areas (i.e. the center of the first floor of a ten-story building). Nonetheless, we still observe that a majority of receivers were still able to receive our broadcast messages in each city.

### 5.3.2 Timeliness of Broadcast Coverage

Next, we determine how fast we can achieve best-scenario coverage by broadcasting messages in the sky on balloons and an aircraft. In Section 5.3.1, we determined that balloons and an airplane allowed us to achieve 70-80% city coverage in San Diego and Las Vegas. We now ask, how fast does it take a balloon or aircraft to achieve such coverage?

We first determine how quickly a balloon can achieve coverage in a San Diego. We show the percentage of coverage—with respect to baseline receivers—for balloon and aircraft deployments in Figure 5.4. “B0” represents our first balloon deployment in February 2025, and “B1” and “B2” represent our two balloon deployments in May 2025. For each of our balloon deployments, we find that the majority of coverage is achieved within the first ten to fifteen

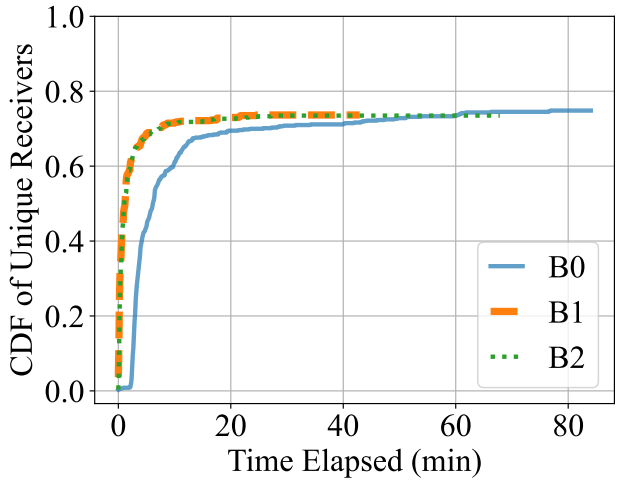


**Figure 5.4.** CDF of percent coverage from balloons with respect to baseline receivers as a function of time.

minutes since launch. However, it takes roughly an hour to reach reach an additional twenty receivers. As a result, we find that balloons can be a relatively timely solution to deploy and broadcast messages to the majority of receivers within a single city.

We then observe how quickly an aircraft can be used to achieve maximum coverage within Las Vegas. These results are shown in Figure 5.5, in which we see how different altitudes can factor towards the timeliness of achieving maximum coverage. We first note that for each altitude, we add a buffer of time that represent that time needed to climb to that altitude from ground level. Starting from 1.7 km of altitude, we notice that although the altitude does not take significantly long to reach, there is a steady increase in coverage with respect to time. This indicates that because the aircraft is not at a high enough altitude, it needs to geographically travel to farther distances in order to attain maximum geographical coverage. In contrast, we find that if that aircraft is able to climb to an altitude of at least 3.3 km, we observe a gain in the percentage of baseline receiver coverage compared to the broadcast coverage from the aircraft flying at an altitude of 1.7 km. As such, there is likely an added benefit of an aircraft transmitting at altitudes higher than 1.7 km to provide fuller broadcast coverage. However, when we compare the broadcast coverage of the aircraft at 3.3 km versus 4.6 km, we interestingly find that there is little to no added benefit to broadcasting at 4.6 km altitude over Las Vegas, specifically. Due to

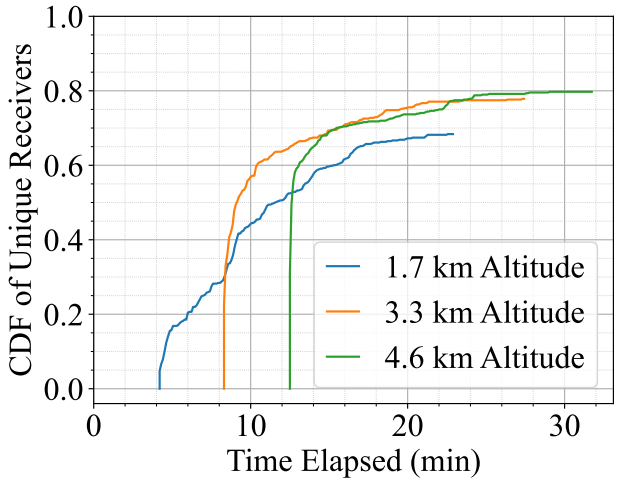


**Figure 5.5.** CDF of percent coverage from an aircraft with respect to baseline receivers as a function of time. Because the aircraft had already reached altitude by the time we start evaluating what percent of receivers are covered, we add a delay to each altitude's starting time based on the time approximate time it would take to reach that altitude.

the approximate time needed to climb from 3.3 km to 4.6 km altitude, we see that the curves eventually align and overlap, meaning that there is minimal gain achieved from transmitting at 4.6 km altitude versus 3.3 km altitude.

### 5.3.3 Reliability

Lastly, we discuss how to achieve better reliability of receiver reception for receivers that successfully heard our broadcast message. In this section, we presented results on the fullness of achievable coverage as well as the time needed to achieve maximum coverage. However, there is variability in the number of broadcasts our receivers receive, which can range from a few broadcasts to multiple hundreds of broadcasts. For receivers that do not consistently and reliably receive our broadcast messages, this would hinder the effectiveness of a real-world, emergency broadcast system. However, there is a relatively simple solution to this problem: increase the number of packet resends to increase the likelihood of successful packet reception. Figure 5.6 shows a CDF of missed packets between sequence numbers. We observe that the majority (i.e. 90%) number of miss packets are between five and fifteen. This indicates that a broadcast message should be sent at least $n$ number of time to achieve a certain probability that a

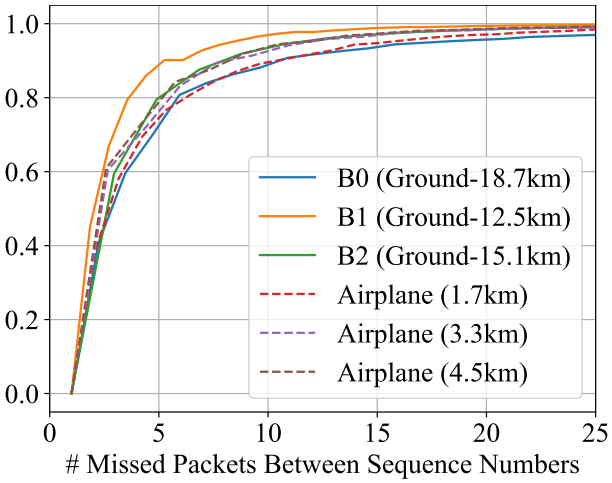


**Figure 5.6.** Number of packet resends to ensure higher reliability. The CDF represents the number of missed packets between sequence numbers for our unmanned-balloon and airplane experiments. For example, if sequence number one and five were received but two through four are missing, then this calculation would result in three missed packets. We see that roughly 80% of all missing gaps in packet reception are roughly six packets. We ignore and cut off the tail end values, which are hotspots that did not have connectivity for the majority of a measurement study.

receiver successfully received our packet (i.e. five to fifteen packet resends should give us around 90% of reliability that any receiver would hear our broadcast message).

Figure 5.6 tells us that that the majority of failed receptions are small, and we can fix this by increasing the number of packet resends to ensure higher reliability of successful broadcast reception.

## 5.4 What's Next?

What do these measurement results mean for the viability of LoRa as a an emergency broadcast technology? We discuss how our findings can inform how the LoRa technology can be used for country-wide, emergency broadcasting by evaluating the potential coverage from commercial airplanes broadcasting messages to the ground below.

We simulate and imagine what LoRa broadcast coverage could look like at scale today using airplanes as the primary provider for broadcast messages from the sky. As a case study

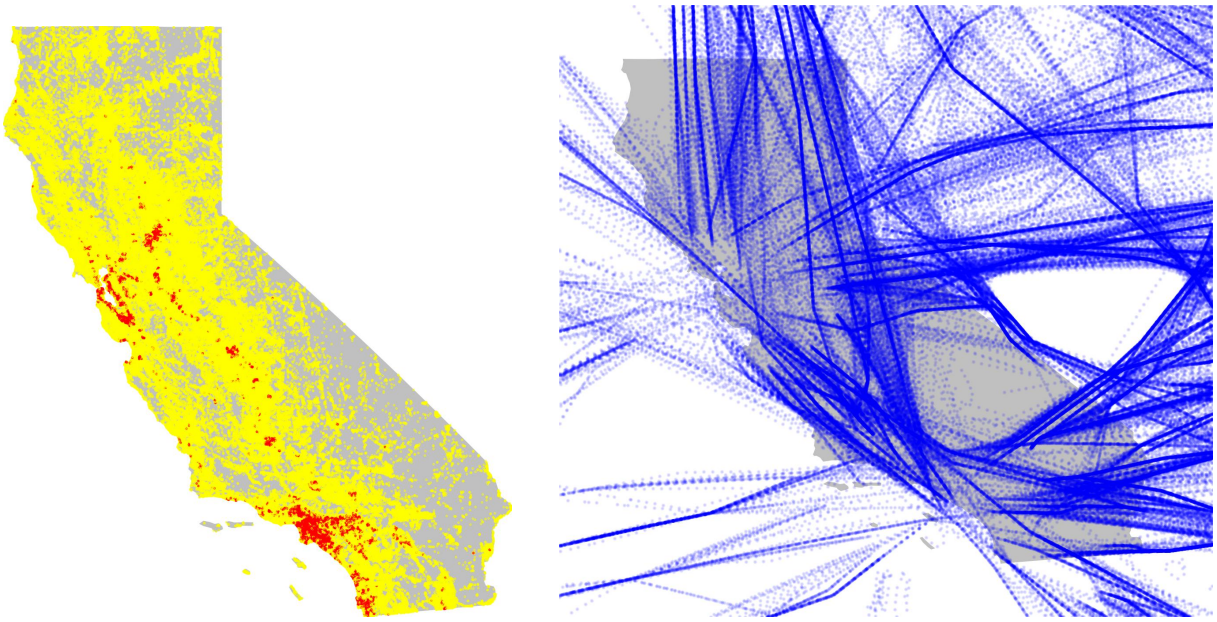

**(a)** Kontur Population. Locations with population overlayed on a map of California from the Kontur population dataset on June 30th, 2022. Red and yellow colors represent urban and rural areas, respectively.

**(b)** ADS-B Data. All locations of airplanes above 30,000 feet (about 9km) according to aircraft's barometric altitude by minute granularity. All points are from a 24-hour time period on October 1st, 2024, all overlayed onto a map of California.

**Figure 5.7.** The Kontur population dataset [26] and ADS-B data [25] overlayed on a map of California. From the Kontur population dataset, the red and yellow colors are divided by greater than and less than 1,250 people within a location, respectively, which denotes urban or rural areas. This is calculated by definition from the USDA, which says that the minimum qualification for an urban area is 5,000 people per square mile of area as of 2020 [65].

(and for simplicity), we observe what state-wide coverage would look like in California, only. To do so, we merge two datasets to analyze how airplane infrastructure (i.e. ADS-B data) translates to coverage on the ground wherever there is population (i.e. Kontur population dataset). Figure 5.7 shows both datasets, with population hexagons (via H3 indexing) overlayed on a map of California and ADS-B datapoints of airplanes overlayed on a map of California as well. As we see in Figure 5.7a, the majority of California is still classified as rural; there is no definition for sub-urban environments to distinguish between rural and sub-urban. However, we see in Figure 5.7b that there lies an opportunity to cover most areas on the ground if we can assume a fixed coverage radius of transmission, assuming that every (commercial) airplane is able to broadcast LoRa data today; only airplanes above 30,000 feet are shown as they likely represent

commercial airplanes, which have fixed and determines schedules.

As a result we simulate coverage of airplanes in the sky over areas with population on the ground in Figure 5.8. For each area on the ground that contains population, we look at the range of all airplane positions above 30,000 feet in altitude within a certain longitudinal and latitudinal range (i.e. 100km, 200km). We then distill two sets of analyses, which are total, temporal coverage of each populated area as well as the maximum latency by which populated areas have no coverage. Figure 5.8a and Figure 5.8b show 24-hour coverage over all of California assuming 100km and 200km coverage radii, respectively. As we can see in the 100km radius coverage model, most portions of California have cover throughout the day except for Death Valley as well as the most north-western part of California. The reason for this can be visually observed in Figure 5.7b, in which there is a gaping "hole" by which airplanes fly around and not through the region. However, as 100km is a relatively conservative estimate of what LoRa in LOS communication can achieve distance-wise, we also show a more optimal and opportunistic coverage model with a 200km radius, in which we then observe that the minimum time of 24-hour coverage changes from 3.8 hours in the 100km radius model to 19.5 hours in the 200km radius model.

We then ask what the worst-case latencies of no-coverage areas in California are. Figure 5.8c and Figure 5.8d show maximum latency models of all populated areas in California with respect to 100km and 200km coverage radii, respectively. For the 100km latency model, we gain see that most maximum latencies occur in Death Valley and most of northern California, with the maximum latency being roughly five hours. However, the maximum latency goes down to around two hours when we extend the coverage radius to 200km, with only northern California having relatively high latencies still. We importantly note that these latencies are actually due to a diurnal effect of airplane flights; intuitively, there are less airplanes flying across the state—and country—at night. Hence, these maximum latencies also occur during the night as well. However, we have shown that there are other options—i.e. unmanned balloons—that can fill in gaps of missing coverage in an ad-hoc fashion as well. We take each of our balloon measurements and

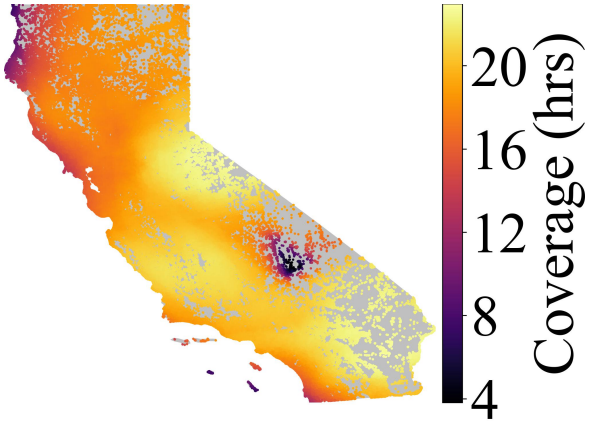


**(a)** 24-hour coverage assuming a 100km broadcast radius.

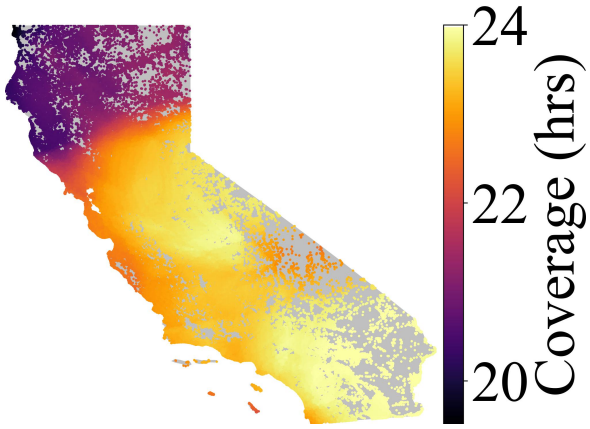


**(b)** 24-hour coverage assuming a 200km broadcast radius.

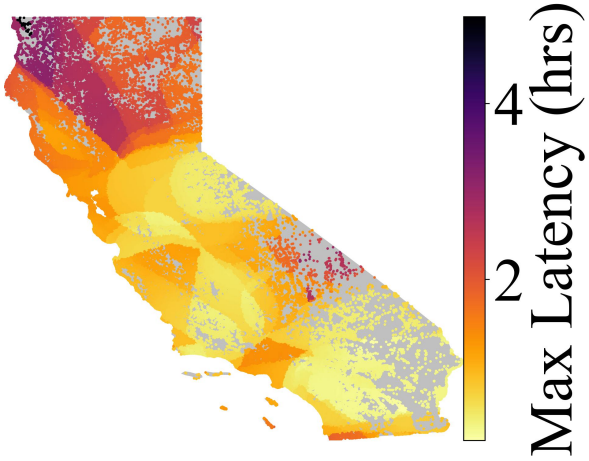


**(c)** Max latency coverage assuming a 100km broadcast radius.

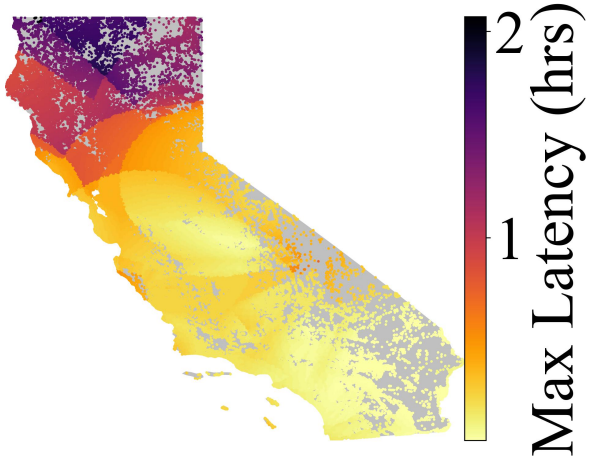


**(d)** Max latency coverage assuming a 200km broadcast radius.

**Figure 5.8.** Analysis of wireless coverage from ADS-B airplane data on October 1st, 2024 covering any location that has human population in California according to the June 30th, 2022 Kontur population dataset. We filter ADS-B data by anything above 30,000 feet—9.144km—of altitude, primarily targeting commercial airplanes flying at cruising altitude. To correlate ADS-B airplane locations with coverage on the ground, we analyze all airplanes within a certain kilometer radius (i.e. 100km, 200km) of each location with population in California. From then, we calculate the amount of time each location on the ground has coverage for, as well as the maximum time latencies by which there are gaps in coverage (up to one-minute of time granularity). Lastly, we ignore terrain in these coverage mappings as our analyses in Section 5.3 do not allow us to predict coverage with different elevations of terrain involved.

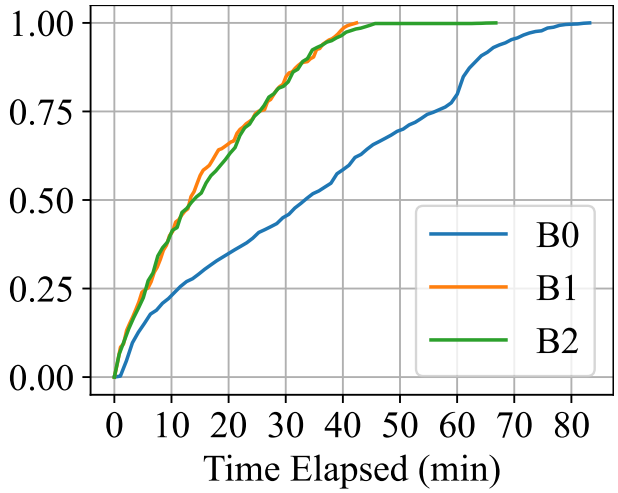


**Figure 5.9.** CDFs of total elapsed time to broadcast to all receivers in our measurements. Each experiment yielded rougly 5,000 uniquely discovered receivers, all covering a range of over one hundred kilometers.

show a CDF of the time it took to broadcast to all receivers in Figure 5.9. As we can observe, B1 and B2 takes around forty minutes to broadcast to all unique receivers on the ground; within the first fifteen minutes, the balloons are already broadcasting to thousands of unique receivers. The weather conditions, weight, amount of helium, and altitude traveled for B0 were different than B1 and B2, which lead it to take roughly forty minutes longer to achieve maximum coverage. Regardless, we show that a weather balloon as an alternative solution is likely able to cover areas that have high latencies of unavailable coverage or lower amounts of daily coverage within 1-1.5 hours of at least a 100km coverage radius.

Lastly, we remind that our results regarding LoRa transmission distances are from configurations straight out of the box – no optimization played a role in towards the range of our measured transmission distances, which underscores the fact that coverage and latencies regarding our 200km radius model is likely achievable today.

## 5.5 Discussion

While our results indicate LoRa can achieve wide-area broadcast coverage, our results also raise several unanswered questions about how well LoRa broadcast would work under

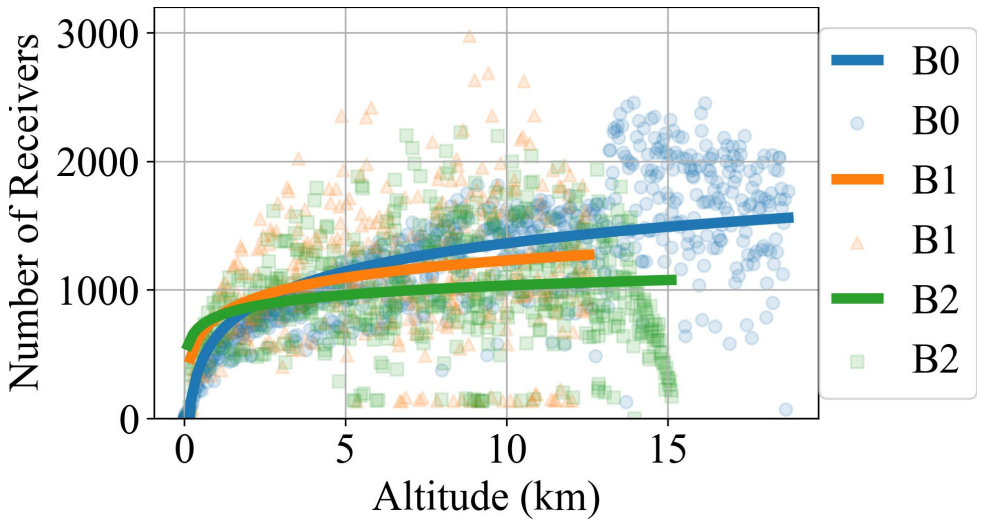


**Figure 5.10.** Number of receivers for every packet transmission with respect to altitude for our three unmanned-balloon measurements. Each dot in the background scatterplot represents a single transmitted packet. We overlay fitted lines over the scatter plot for each balloon experiment that show an exponentially decreasing trend between the relationship of altitude and number of receivers.

different LoRa configurations than we tested, and also how it would perform from different broadcast platforms (e.g., scaling nearby mountains, or driving across the metro area). We also provide a discussion about how LoRa emergency broadcast could handle imperfect broadcast coverage to get closer to nearly 100% coverage. We provide a discussion about each of these issues in turn:

### 5.5.1 Altitude and Broadcast Coverage Area

In Section 5.3.2, we found that the higher altitude a balloon goes, it improves coverage area. In this discussion we investigate the relationship between altitude and coverage area in detail to observe if there is any distinct benefit to the balloon reaching extremely high altitudes (above 10 km). Figure 5.10 illustrates the relationship between altitude of the balloon and how many receivers received packets. We see that the balloon gains coverage of receivers exponentially as the balloon rises to 2 km, but above that there is a much less clear relationship between altitude and coverage. On launch B0 above 12 km, the ballon gained a significant amount of covered

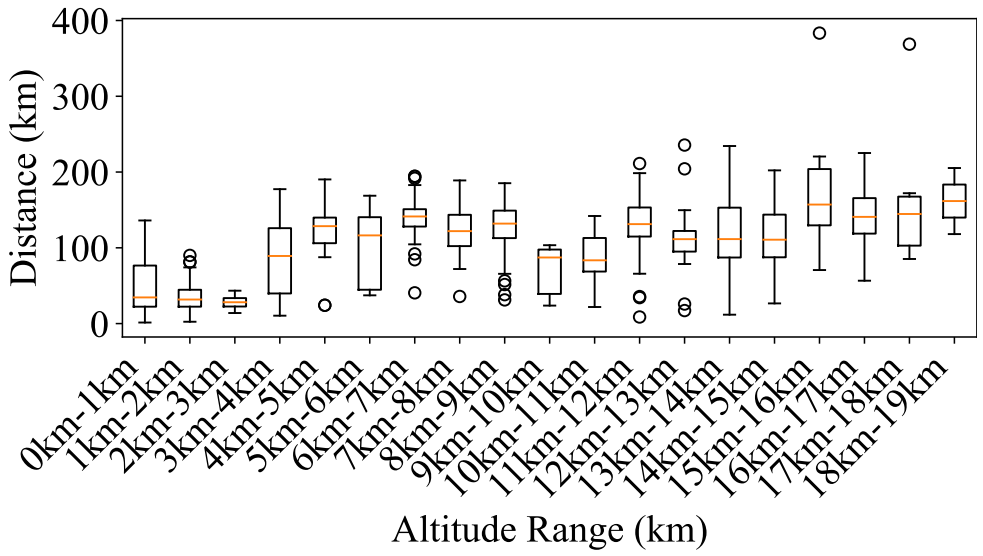


**Figure 5.11.** Boxplot showing for B0 launch, binned by 1 km altitude ranges, the distance the receivers were first contacted at a particular altitude. There are two distinct jumps in how far away receivers are covered at 4–5 km and 11–12 km altitude.

receivers, almost an increase of 100% compared to 2 km. Whereas on launch B1 that significant increase was not seen, and in launch B2 there was a distinct drop off in covered receivers at the highest altitudes (though this is likely due to low temperature and hardware failure). The question remains though, is the coverage area expanding as altitude increases, or are there just more receivers in the same area. Figure 5.11 shows that for the B0 launch there are two distinct altitudes that lead to a large jump in distance of broadcasts making it to far away receivers. At 3–4 km altitude the ballon starts to reach receivers 100 km away, and at 11–12 km, it reaches close to 200 km away. Figure 5.12 shows the geographical coverage of launch B0, and we see that the receivers only first seen above 15 km altitude are predominately in the neighboring Los Angeles metro area $\sim$ 200 km north of San Diego. This indicates that a single balloon may be able to provide redundant coverage of nearby metro areas, but the time to rise to that high altitude, and the uncertain reliability, would still justify sending up a separate balloon in the Los Angeles metro area at the same time.

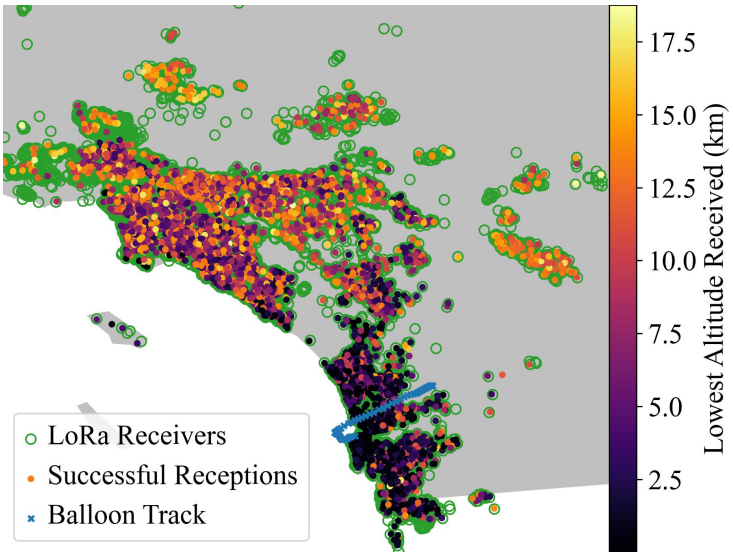


**Figure 5.12.** Minimum alt. a receiver heard Balloon B0's broadcasts. Above 15 k feet the coverage extended well beyond San Diego into the greater Los Angeles metro area.

### 5.5.2 Spreading Factor Affecting Coverage

This only measures the maximum Spreading Factor configuration of LoRa (i.e. SF10) available in the United States, but there are some downsides to using such a high-redundancy transmission rate. The maximum payload size for SF10 is eleven bytes; the maximum payload size for SF7 is 242 bytes. Thus, there is a 22x size difference in the maximum message that can be sent in payload SF7 and SF10. A full length SF7 maximum payload size of 242 bytes will occupy similar airtime to a SF10 message. However, to isolate the effects of spreading factor, we ensure that the payload sizes are identical, even if this results in a difference of airtime (and hence, more or less interference for a packet). In a subset of our balloon experiments we also capture SF7 data in parallel with our SF10 (for B1 and B2 in our May deployments). All configurations (i.e. payload size, transmit power, antenna gain, coding rate, and frequency bandwidth) for the SF7 device are identical to the SF10 device to enable a one-to-one comparison of Spreading Factor only. These results are shown in Figure 5.13. As expected, at SF10 we see that there generally are more receivers of each broadcast packet across all altitudes, this is expected as SF10 has higher redundancy than SF7 and can operate at lower SNR. However, SF10 also has

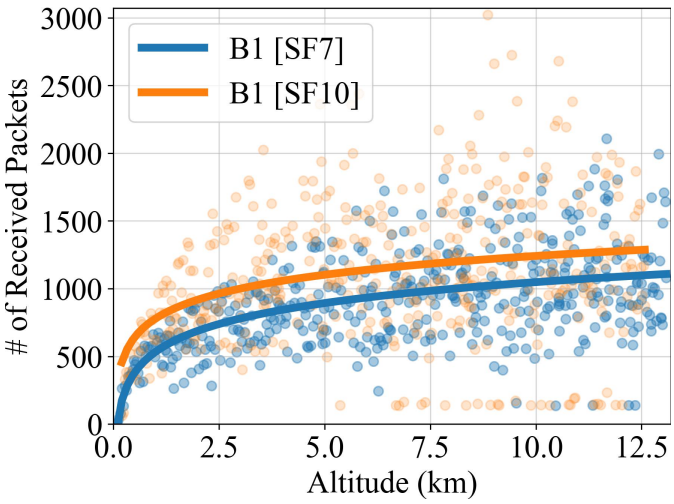


**Figure 5.13.** Number of packet receptions versus altitude for a SF7 and SF10 device.

more broadcasts that completely failed to reach any receiver, particularly at high altitudes, where SF7 broadcasts generally still have more consistent coverage. This could be caused by higher altitudes producing a weaker signal that can not overcome the interference of other interfering transmissions in LoRa's unlicensed 900 MHz band.

Figure 5.14 shows the distribution of distance and SNR on the same plot for received broadcasts at SF7 and SF10 in one ballon launch. We see the expected offset in the SNR distributions, in which SF10 is received often at 9dB lower SNR than SF7. As such, we may expect that SF10's lower SNR reception translates to longer transmission distances of SF10 compared to SF7. However, the distributions of transmission-receive distances in Figure 5.14 are surprisingly similar. As a result, there is data that supports using a lower spreading factor than SF10 (i.e. SF7) for wide-area broadcasts from aircraft, particularly for providing coverage in one metro area. Lower spreading factor is possibly desirable for broadcasts because it significantly increases how many bytes can be sent in one broadcast message. However, as mentioned earlier, a higher-byte payload for SF7 will induce more interference, potentially leading to lower probability and success of reception.

Future work is required to elucidate the tradeoffs between range, payload size, and reliability with respect to Spreading Factor. From our analyses in Figure 5.14, the most no-

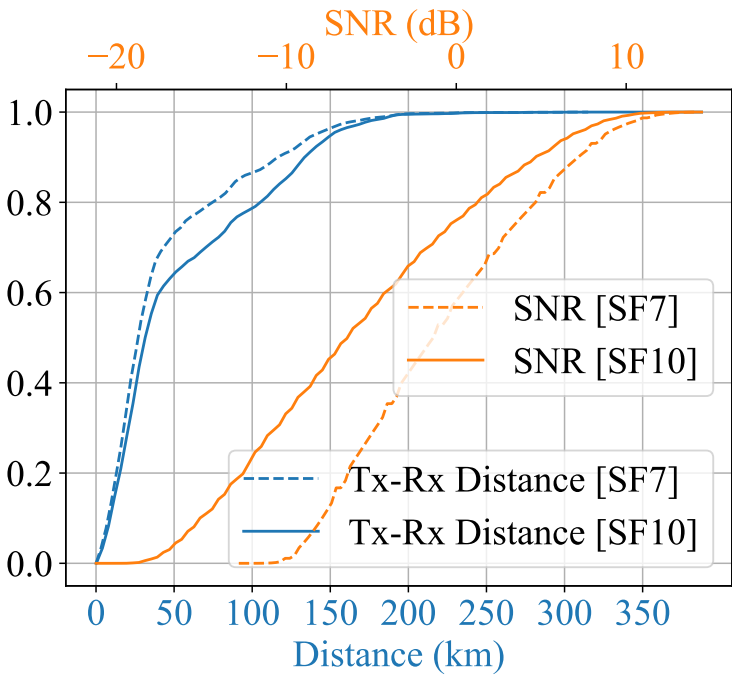


**Figure 5.14.** CDFs of transmission distance and SNR distributions for both SF7 and SF10 for our measurements in Figure 5.13.

ticeable difference between SF7 and SF10 is decibel gain, which directly controls the range of communication and the ability to decode messages as extremely low SNR. Less clear is if there is a correlation of payload sizes between different Spreading Factors (e.g. SF7, SF8, SF9, and SF10) and range and reliablity of communication.

### 5.5.3 Alternative LoRa Broadcast Platforms

In this work we focus on deploying LoRa broadcasting from aircraft; the reason being that they provide clear line-of-sight that provides extremely long communication distances at low transmission powers for LoRa. However, there are other ways of getting line-of-sight, including placing transmitters on a mountain neighboring a metro area, or even just by driving the transmitter across the metro area. We now explore each of these alternative options.

**Broadcasting LoRa from a tower or mountain:**

We tested broadcasting LoRa from a tall 500 m mountain next to San Diego to see what coverage it achieves compared to the balloon launches. For a head-to-head comparison, we

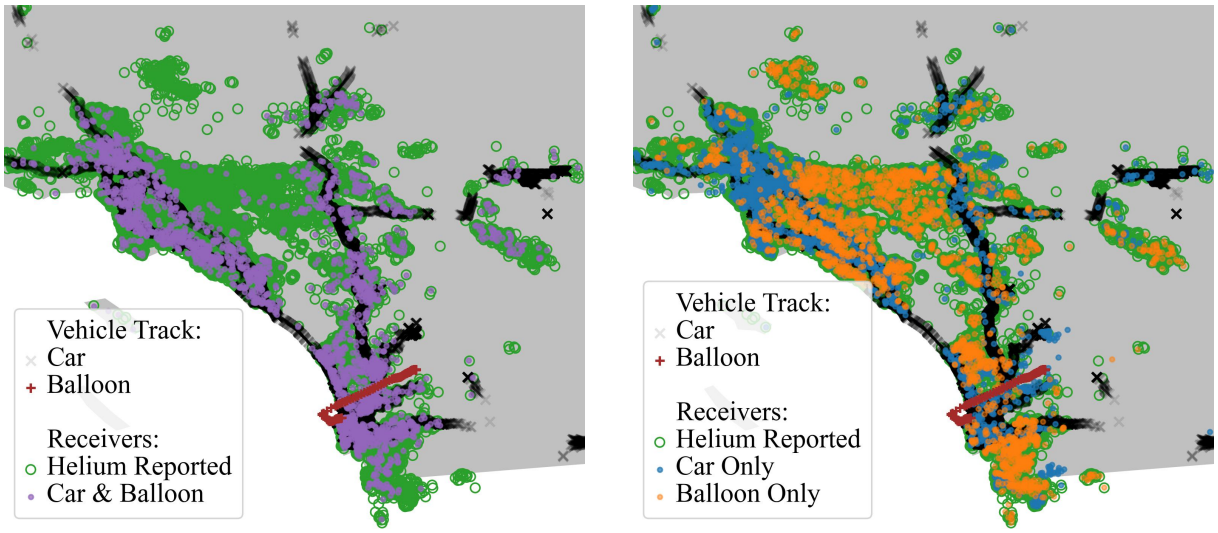


**(a)** Helium receivers heard both by our car- and balloon-transmitting platforms.

**(b)** Helium received by either our car- or balloon-transmitting platforms.

**Figure 5.15.** Wide-area coverage of LoRa transmissions from the ground and the air via car and balloon.

evaluated how many of the baseline (i.e., wardriving) receivers could be covered by broadcasts from the balloon and the transmitter on the mountain. From the mountain, broadcasts could contact 45% of the baseline set of receivers in San Diego, whereas the balloon experiments achieved ∼72% coverage. The balloon not only provides a significant improvement in broadcast coverage, but it also is much more rapidly deployed than hiking or driving a LoRa transmitter to the top of the mountain after all existing broadcast transmitters have been disabled or destroyed. Also, many metro areas do not have nearby tall mountains that can be used as a broadcasting platform.

**Broadcasting LoRa by wardriving:**

For metro areas without a nearby mountain, the only alternative to rapidly deploying broadcasts with aircrafts is to wardrive around the metro area with a car equipped with a LoRa transmitter. This potentially could achieve similar coverage as the aircraft, and without the additional complexity of flight, in exchange for a longer latency for devices receiving the broadcast. Figure 5.15a shows the coverage intersection of both cars and balloons. We see

that this coverage area is mostly around the roads that were driven. However, there are many receivers that exist on the Helium network that were not reached by either the wardriving or balloon transmissions. Figure 5.15b shows the receivers that were covered exclusively by either the car or the balloon. This shows that the ballon has a much wider coverage area than the wardriving car. The car however appears to have better broadcast penetration into the areas near the roads where the balloon has much more spotty coverage. These results indicate that for wide-area coverage balloons are generally going to be preferred compared to cars, but some sort of local transmission, like that from a car, could patch up the holes in coverage where the balloon is unable to penetrate its broadcasts through buildings into the indoor areas where receivers may be located.

### 5.5.4 Can LoRa Broadcast hit 100% coverage?

To achieve 100% coverage with airborne LoRa broadcasting, there needs to be additional relaying of the messages locally between LoRa receivers on the ground. Fortunately, all LoRa devices are both receivers and transmitters, therefore it's possible to build a relay system, such as: when a device receives a broadcast, it then immediately turns around and locally re-broadcasts it to all other nearby receivers. This type of LoRa rebroadcasting system is actually already available today as part of the Meshtastic [37] LoRa mesh messenger system: every text message sent is rebroadcasted by up to $N$ Meshtastic devices so it can achieve wide coverage and reach the intended destination.

## 5.6 Summary

This chapter shows that standard low-power LoRa is capable of inexpensively achieving wide-area broadcasts when LoRa transmitters are mounted to aerial vehicles. As opposed to land-based vehicles and infrastructure, vehicles in the sky (e.g. balloons, airplanes) are underutilized and can provide provide rapid deployment of broadcast infrastructure to reach otherwise disconnected receivers and issue critical commands. The barrier-to-entry to rapidly

deploy a weather balloon is extremely feasible and can be used today, and our small-aircraft deployments show that there is very little added infrastructure that would be needed to deploy LoRa broadcasters on whole fleets of commercial aircraft. There is a significant amount of improvement that can be made to make aerial LoRa deployments have wider broadcast coverage area and higher reliability.

While a limitation of this work regards unidirectional communication, this work shows that LoRa on non-terrestrial infrastructure can be an option to enable the LoRa technology and a new variety of applications that prior technologies and infrastructures could not accomplish before.

# Chapter 6

# Conclusion and Future Work

This dissertation aims to explore different ways by which we can enable the LoRa technology and newer applications that cannot exist with current wireless solutions. Indeed, one major takeaway lies in the observation that LoRa exhibits significant performance benefits (i.e. range of communication) in line of sight communication via non-terrestrial infrastructure versus non-line of sight communication via ground-based communication. This dissertation first shows that while federated infrastructure (i.e., via the Helium Network) is potentially one start to enabling the pervasive usage of LoRa, it lacks pervasive distribution across the wide-area, only servicing reliable coverage in more densely populated areas. The Helium Network—as unplanned, crowdsourced infrastructure—is unlikely able to compete head-to-head against cellular infrastructure. Surprisingly, BLE via Apple devices and the Find My network performs comparatively similar to LoRa via the Helium Network and is a significantly cheaper infrastructure to deploy, leaving the Helium Network in question of its widespread effectiveness as a LoRaWAN today.

However, this dissertation also shows that non-terrestrial infrastructure is one potential avenue that can enable the LoRa technology and newer applications. This work shows that deploying LoRa transmitters on aerial vehicles (i.e. weather balloons, airplanes) enables the LoRa technology to broadcast to entire metropolitan areas, vastly increasing the area of communication via air versus via ground. Non-terrestrial infrastructure enables longer range of communication

in line of sight conditions, hence allowing LoRa to achieve broadness of coverage over large areas. The one caveat to this finding regards how broadness does not imply fullness. That is, the coverage is not thorough, and communicating to hard-to-reach areas is still one challenge and limitation of this work. However, there are known solutions to this problem (i.e. mesh networking) that will accommodate last-mile connectivity from aerial vehicles above.

There are two directions from this dissertation to explore regarding the impact and usefulness of LoRa on non-terrestrial infrastructure. The first direction is to evaluate multi-directional communication in the air (i.e. receiving as well). This work of LoRa on non-terrestrial infrastructure only focuses on transmissions; evaluating non-terrestrial reception would also enable deeper understanding into how LoRa on non-terrestrial infrastructure would enable non-terrestrial networking. That is, not only evaluating data transmission but also data reception allows a fuller picture on how well a non-terrestrial LoRaWAN would function today. Furthermore, because of the relative cheapness of the infrastructure, this future work would better evaluate and understand whether LoRa can have a place amidst other wireless infrastructures today as an economical alternative.

The second direction regards combining LoRa on non-terrestrial infrastructure with mesh networking. Nodes on the ground can act as relays for other nodes, which particularly addresses the communication issue between the air and hard-to-reach areas on the ground. As discussed in Chapter 5, emergency broadcasting functions already perform this way to achieve fuller coverage. As a result, the future direction for this work to enable LoRa on non-terrestrial infrastructure would also be to add mesh networking to the infrastructure and LoRaWAN as well.

# .1 Ethical Considerations

This research does not contain human subjects. The two ethical situations to consider pertains to shipping batteries via commercial shipping and FAA regulations with respect to broadcasting wireless signals in the air.

## .1.1 Measurement via Commercial Shipping

Our measurement approach is inspired by Zhang et. al.'s design from IMC'21 [91]. The primary concern they identified for shipping powered devices is "a dangerous evolution of heat, or hav[ing] the risk of catching fire while in transit."

Our smartphones follow the same design and risk profile as Zhang et. al. Beyond this, we add two microcontroller platforms—an ESP32 and a HelTec CubeCell. These devices operate at a power point $100-1000\times$ lower than the phone and provide no thermal risk. We power these platforms using commercial external battery packs. Heat is proportional to the amount of power drawn off the battery pack, which again is minimal for the microcontrollers. These UL-certified, consumer-grade products also include built-in, automatic thermal shutdown.

## .1.2 FAA and FCC Regulations

With respect to unmanned balloons, we followed the following rules in FAA Part 101 [7], including submitting a Notice to Air Missions (NOTAM) at least six hours prior to launch. A parachute was not needed, and multiple parties told us that the drag induced from the popped balloon would be enough slow down the payload to prevent "hazard to persons or property not associated with the operation" [7]. With respect to airplanes, FAR 91.21 states that "[any] other portable electronic device that the operator of the aircraft has determined will not cause interference with the navigation or communication system of the aircraft on which it is to be used" [1]. Furthermore, an FAA Aviation Safety Inspector wrote via email that a "transmitter such a ham radio operating on or near the 900MHz range would have to be approved by the airline or charter operator before being used on a flight." In our airplane experiment around [City], we checked and ensured that our transmitting, LoRa device did not interfere with our airplane's navigation or communication system prior to departure.

# Bibliography


[1] 91.21 Portable electronic devices. https://www.ecfr.gov/current/title-14/chapter-I/subchapter-F/part-91/subpart-A/section-91.21. Date Accessed: May, 2025.

[2] Blues starter kit for lorawan. https://shop.blues.com/products/blues-starter-kit-lorawan. Date Accessed: October 2025.

[3] Cellular vs. LoRaWAN for IoT deployments: Why the business case matters more than the tech. https://www.particle.io/iot-guides-and-resources/particle-cellular-iot-vs-lorawan/. Date Accessed: May, 2025.

[4] CubeCell GPS-6502 (HTCC-AB02S). https://heltec.org/project/htcc-ab02s/. Date Accessed: April, 2025.

[5] Easily-assembled iot kits and bundles now in stock from rakwireless. https://store.rakwireless.com/collections/kits-bundles?srsltid=AfmBOoom1se8GdV1f6eHjVaOVrM-ZZONFuftKL4QSS-eSHpDqsk61flC. Date Accessed: October 2025.

[6] Lora/lorawan iot kit v3. https://www.dragino.com/products/lora/item/237-lora-lorawan-iot-kit-v3.html. Date Accessed: October 2025.

[7] PART 101—MOORED BALLOONS, KITES, AMATEUR ROCKETS, AND UNMANNED FREE BALLOONS. https://www.ecfr.gov/current/title-14/chapter-I/subchapter-F/part-101. Date Accessed: May, 2025.

[8] Sparkfun digi x-on kit for lorawan® (north america). https://www.sparkfun.com/sparkfun-iot-node-lorawan-gateway-kit-north-america.html. Date Accessed: October 2025.

[9] TinyGS. https://tinygs.com/. Date Accessed: May, 2025.

[10] Talal Ahmad, Ranveer Chandra, Ashish Kapoor, Michael Daum, and Eric Horvitz. Wi-fly: Widespread opportunistic connectivity via commercial air transport. In *Proceedings of the 16th ACM Workshop on Hot Topics in Networks*, HotNets '17, page 43–49, New York, NY, USA, 2017. Association for Computing Machinery.

[11] Mohamed A. Ahmed, Jose Luis Gallardo, Marcos D. Zuniga, Manuel A. Pedraza, Gonzalo Carvajal, Nicolás Jara, and Rodrigo Carvajal. Lora based iot platform for remote monitoring of large-scale agriculture farms in chile. *Sensors*, 22(8):2824, April 2022.

[12] Apple. https://www.apple.com/airtag/. Online; accessed 8-May-2024.

[13] Apple. Find my - apple, 2023.

[14] Alex Bellon, Alex Yen, and Pat Pannuto. TagAlong: Free, wide-area data-muling and services. In *Proceedings of the 24th International Workshop on Mobile Computing Systems and Applications*, HotMobile '23, page 103–109, New York, NY, USA, 2023. Association for Computing Machinery.

[15] Norbert Blenn and Fernando Kuipers. Lorawan in the wild: Measurements from the things network, 2017.

[16] Norbert Blenn and Fernando Kuipers. Lorawan in the wild: Measurements from the things network, 2017.

[17] Robert Bogue. Towards the trillion sensors market. *Sensor Review*, 34(2):137–142, March 2014.

[18] Cable.co.uk. Worldwide mobile data pricing 2022.

[19] CBSLA Staff. Internet outage hits thousands of southern california spectrum customers. https://losangeles.cbslocal.com/2020/07/30/internet-outage-southern-california-spectrum-customers/, 7 2020. Accessed Sep 2021.

[20] Center for Applied Internet Data Analysis. Mapping autonomous systems to organizations: CAIDA's inference methodology. https://www.caida.org/archive/as2org/, 8 2020. Accessed Sep 2021.

[21] Samuel DeBruin, Branden Ghena, Ye-Sheng Kuo, and Prabal Dutta. Powerblade: A low-profile, true-power, plug-through energy meter. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*, SenSys '15, page 17–29, New York, NY, USA, 2015. Association for Computing Machinery.

[22] Zakir Durumeriç. ZAnnotate. https://github.com/zmap/zannotate/, 2017. Accessed: May 2021.

[23] Olakunle Elijah, Sharul Kamal Abdul Rahim, Vitawat Sittakul, Ahmed M. Al-Samman, Michael Cheffena, Jafri Bin Din, and Abdul Rahman Tharek. Effect of weather condition on lora iot communication technology in a tropical region: Malaysia. *IEEE Access*, 9:72835–72843, 2021.

[24] Mahmoud Elkhodr, Seyed Shahrestani, and Hon Cheung. Emerging wireless technologies in the internet of things: a comparative study. 2016.

[25] ADSB Exchange. https://samples.adsbexchange.com/readsb-hist/.

[26] The Humanitarian Data Exchange. Kontur population: Global population density for 400m h3 hexagons. https://data.humdata.org/dataset/kontur-population-dataset, 2022.

[27] Jonathan Gresl, Scott Fazackerley, and Ramon Lawrence. Practical precision agriculture with lora based wireless sensor networks. In *SENSORNETS*, pages 131–140, 2021.

[28] Sebastián Gutiérrez, Israel Martínez, Jorge Varona, M. Cardona, and Ricardo Espinosa. Smart mobile lora agriculture system based on internet of things. In *2019 IEEE 39th Central America and Panama Convention (CONCAPAN XXXIX)*, pages 1–6, 2019.

[29] Aḿeer Haleem et al. Helium a decentralized wireless network. Technical report, Helium Systems Inc., November 2018.

[30] Amir Haleem, Andrew Allen, Andrew Thompson, Marc Nijdam, and Rahul Garg. Helium: A Decentralized Wireless Communication Network. *RS Open Journal on Innovative Communication Technologies*, 4(13), dec 9 2024. https://rs-ojict.pubpub.org/pub/4i0gbpd8.

[31] Alexander Heinrich, Milan Stute, and Matthias Hollick. OpenHaystack: A Framework for Tracking Personal Bluetooth Devices via Apple's Massive Find My Network. In *Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, WiSec '21, page 374–376, New York, NY, USA, 2021. Association for Computing Machinery.

[32] Helium. Helium blockchain primitives. https://docs.helium.com/blockchain/blockchain-primitives. Accessed: May 2021.

[33] Helium. New helium partner – careband. https://blog.helium.com/helium-careband-make-covid-19-contact-tracing-affordable-4dcc90eccc4, 2020. Accessed May, 2021.

[34] Helium. New helium partner – nowi. https://blog.helium.com/nowi-is-using-the-peoples-network-to-stop-leaks-and-water-waste-a00bb6aa280f, 2020. Accessed May, 2021.

[35] Helium. Hotspot api. https://docs.helium.com/architecture/hotspot-api/, 2024.

[36] Helium Engineering Blog. App version 3.2.0. https://engineering.helium.com/2021/05/18/app-version-320.html. Accessed May, 2021.

[37] Kevin Hester and the Meshtastic Community. Meshtastic: Off-grid decentralized lora mesh networking. https://meshtastic.org/, 2025. Open-source project providing long-range, low-power communication via LoRa mesh networks.

[38] https://github.com/Carniverous19. Hip17: Hex density based transmit reward scaling. https://github.com/helium/HIP/blob/main/0017-hex-density-based-transmit-reward-scaling.md, 2020.

[39] https://github.com/mawdegroot and https://github.com/riobah. Hip 104: Fine-tune hip-17 parameters to tackle density. https://github.com/helium/HIP/blob/main/0104-finetune-hip17-parameters-to-tackle-density.md, 2023.

[40] Tatsuya Iizuka, Naoto Endo, Hiroshi Matsubara, and Masaki Hisada. Low-cost airdropping sensor system based on light balloons and lora communication. In *2024 IEEE 100th Vehicular Technology Conference (VTC2024-Fall)*, pages 1–6, 2024.

[41] Dhananjay Jagtap, Alex Yen, Huanlei Wu, Aaron Schulman, and Pat Pannuto. Federated Infrastructure: Usage, Patterns, and Insights From "The People's Network". In *Proceedings of the 21st ACM Internet Measurement Conference*, pages 22–36, 2021.

[42] Mookeun Ji, Juyeon Yoon, Jeongwoo Choo, Minki Jang, and Anthony Smith. Lora-based visual monitoring scheme for agriculture iot. In *2019 IEEE Sensors Applications Symposium (SAS)*, pages 1–6, 2019.

[43] Philipp H. Kindt, Daniel Yunge, Robert Diemer, and Samarjit Chakraborty. Energy modeling for the bluetooth low energy protocol. *ACM Trans. Embed. Comput. Syst.*, 19(2), mar 2020.

[44] Abhay Kumar. Hip10: Usage-based data transfer rewards. https://github.com/helium/HIP/blob/master/0010-usage-based-data-transfer-rewards.md, 2020. Accessed: May 2021.

[45] Unwired Lab. Opencellid cell tower database, 2023.

[46] Yang Li, Hao Lin, Zhenhua Li, Yunhao Liu, Feng Qian, Liangyi Gong, Xianlong Xin, and Tianyin Xu. A nationwide study on cellular reliability: measurement, analysis, and enhancements. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, SIGCOMM '21, page 597–609, New York, NY, USA, 2021. Association for Computing Machinery.

[47] libp2p. Circuit relay. https://docs.libp2p.io/concepts/circuit-relay/. Accessed: May 2021.

[48] Life360. Life360: The #1 family locator app & safety membership. https://www.life360.com/. Online; accessed 8-May-2024.

[49] Kais Mekki, Eddy Bajic, Frederic Chaxel, and Fernand Meyer. A comparative study of LPWAN technologies for large-scale IoT deployment. *ICT Express*, 5(1):1–7, March 2019.

[50] Milwaukee. https://www.milwaukeetool.com/Products/48-21-2301. Online; accessed 8-May-2024.

[51] Mohamed Mirza, Yasser Asrul Ahmad, and Teddy Surya Gunawan. Low altitude balloon iot tracker. In *2017 IEEE 4th International Conference on Smart Instrumentation, Measurement and Application (ICSIMA)*, pages 1–6, 2017.

[52] Frank Mong. Adding over 1k [helium] hotspots per day. https://twitter.com/fmong/status/1395478353629499393. Frank Mong is the COO of Helium.

[53] Federico Montori, Luca Bedogni, Marco Di Felice, and Luciano Bononi. Machine-to-machine wireless communication technologies for the internet of things: Taxonomy, comparison and open issues. *Pervasive and Mobile Computing*, 50:56–81, October 2018.

[54] Murata. LoRa FAQ. https://www.murata.com/support/faqs/lpwa/lora/hardware/0008. Accessed: May 2021.

[55] Arslan Musaddiq, Neda Maleki, Francis Palma, David Mozart, Tobias Olsson, Mustafa Omareen, and Fredrik Ahlgren. Internet of things for wetland conservation using helium network: Experience and analysis. In *Proceedings of the 12th International Conference on the Internet of Things*, IoT '22, page 143–146, New York, NY, USA, 2023. Association for Computing Machinery.

[56] Nitin Naik. LPWAN technologies for IoT systems: Choice between ultra narrow band and spread spectrum. In *2018 IEEE International Systems Engineering Symposium (ISSE)*. IEEE, October 2018.

[57] The Things Network. Us902-928 mhz band. https://www.thethingsnetwork.org/docs/lorawan/regional-parameters/us915/. Date Accessed: November, 2025.

[58] The Things Network. US902-928 MHz Band. https://www.thethingsnetwork.org/docs/lorawan/regional-parameters/us915/.

[59] IoT Business News. State of IoT 2024: Number of connected IoT devices growing 13% to 18.8 billion globally, September 2024.

[60] Olatinwo, Abu-Mahfouz, and Hancke. A survey on LPWAN technologies in WBAN for remote health-care monitoring. *Sensors*, 19(23):5268, November 2019.

[61] Antonino Pagano, Daniele Croce, Ilenia Tinnirello, and Gianpaolo Vitale. A survey on lora for smart agriculture: Current trends and future perspectives. *IEEE Internet of Things Journal*, 10(4):3664–3679, 2023.

[62] Vit Prajzler. What is LoRa: The fundamentals. https://medium.com/@prajzler/what-is-lora-the-fundamentals-79a5bb3e6dec, October 2023. Date Accessed: April, 2025.

[63] Peter Ruckebusch, Spilios Giannoulis, Ingrid Moerman, Jeroen Hoebeke, and Eli De Poorter. Modelling the energy consumption for over-the-air software updates in lpwan networks: Sigfox, lora and ieee 802.15.4g. *Internet of Things*, 3-4:104–119, 2018.

[64] Samsung. https://www.samsung.com/us/mobile/mobile-accessories/phones/galaxy-smarttag2-black-ei-t5600bbegus/. Online; accessed 8-May-2024.

[65] Austin Sanders and John Cromartie. What is rural? https://www.ers.usda.gov/topics/rural-economy-population/rural-classifications/what-is-rural/, 2024.

[66] Semtech. Frequently Asked Questions. https://www.semtech.com/design-support/faq/faq-lorawan/P20.

[67] Semtech. LoRa network packet forwarder project. https://github.com/Lora-net/packet_forwarder, 2017.

[68] Semtech. https://www.semtech.com/products/wireless-rf, 2024.

[69] Semtech. Lora and lorawan: A technical overview. https://lora-developers.semtech.com/uploads/documents/files/LoRa_and_LoRaWAN-A_Tech_Overview-Downloadable.pdf, 2024.

[70] Amritraj Singh, Kelly Click, Reza M. Parizi, Qi Zhang, Ali Dehghantanha, and Kim-Kwang Raymond Choo. Sidechain technologies in blockchain networks: An examination and state-of-the-art review. *Journal of Network and Computer Applications*, 149:102471, 2020.

[71] Philip Sparks. The route to a trillion devices. https://community.arm.com/cfs-file/__key/telligent-evolution-components-attachments/01-1996-00-00-00-01-30-09/ARM-_2D00_-The-route-to-a-trillion-devices-_2D00_-June-2017.pdf, 2017.

[72] S. J. Suji Prasad, M. Thangatamilan, M. Suresh, Hitesh Panchal, Christober Asir Rajan, C. Sagana, B. Gunapriya, Aditi Sharma, Tusharkumar Panchal, and Kishor Kumar Sadasivuni. An efficient lora-based smart agriculture management and monitoring system using wireless sensor networks. *International Journal of Ambient Energy*, 43(1):5447–5450, July 2021.

[73] Carolyn Taylor. A glimpse into redcap nr devices. https://www.3gpp.org/technologies/nr-redcap-glimpse, July 2023.

[74] The Telit Cinterion Team. Cellular iot for smart meters: The future of utilities. https://www.telit.com/blog/cellular-iot-smart-meters/, April 2024. Accessed: October 2025.

[75] The Things Network Global Team. Lora world record broken: 832km/517mi using 25mw. https://www.thethingsnetwork.org/article/lorawan-world-record-broken-twice-in-single-experiment-1, April 2020. Date Accessed: May 2025.

[76] The Things Network Global Team. New lora world record: 1336 km / 830 mi. https://www.thethingsnetwork.org/article/new-lora-world-record-1336-km-830-mi, September 2023. Date Accessed: May 2025.

[77] Teltonika. https://teltonika-gps.com/. Online; accessed 8-May-2024.

[78] Tile. https://www.tile.com/en-us. Online; accessed 8-May-2024.

[79] Shuai Tong, Jiliang Wang, Jing Yang, Yunhao Liu, and Jun Zhang. Citywide lora network deployment and operation: Measurements, analysis, and implications. In *Proceedings of the 21st ACM Conference on Embedded Networked Sensor Systems*, SenSys '23, page 362–375, New York, NY, USA, 2024. Association for Computing Machinery.

[80] TTN. The things network.

[81] Jamie S.C. Turner, M.F. Ramli, L.M. Kamarudin, A. Zakaria, A.Y.M. Shakaff, D.L. Ndzi, C.M. Nor, N. Hassan, and S.M Mamduh. The study of human movement effect on signal strength for indoor wsn deployment. In *2013 IEEE Conference on Wireless Sensor (ICWISE)*, pages 30–35, 2013.

[82] Uber. H3.

[83] Frank Uyeda, Marc Alvidrez, Erik Kline, Bryce Petrini, Brian Barritt, David Mandle, and Aswin Chandy Alexander. Sdn in the stratosphere: loon's aerospace mesh network. In *Proceedings of the ACM SIGCOMM 2022 Conference*, SIGCOMM '22, page 264–280, New York, NY, USA, 2022. Association for Computing Machinery.

[84] Bindiya Vakil and Tom Linton. Why we're in the midst of a global semiconductor shortage. https://hbr.org/2021/02/why-were-in-the-midst-of-a-global-semiconductor-shortage. Accessed: May 2021.

[85] Benny Vejlgaard, Mads Lauridsen, Huan Nguyen, Istvan Z. Kovacs, Preben Mogensen, and Mads Sorensen. Coverage and capacity analysis of sigfox, lora, gprs, and nb-iot. In *2017 IEEE 85th Vehicular Technology Conference (VTC Spring)*, pages 1–5, 2017.

[86] Dongzhu Xu, Anfu Zhou, Xinyu Zhang, Guixian Wang, Xi Liu, Congkai An, Yiming Shi, Liang Liu, and Huadong Ma. Understanding operational 5g: A first measurement study on its coverage, performance and energy consumption. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication*, SIGCOMM '20, page 479–494, New York, NY, USA, 2020. Association for Computing Machinery.

[87] Deliang Yang, Xianghui Zhang, Xuan Huang, Liqian Shen, Jun Huang, Xiangmao Chang, and Guoliang Xing. Understanding power consumption of nb-iot in the wild: tool and large-scale measurement. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, MobiCom '20, New York, NY, USA, 2020. Association for Computing Machinery.

[88] Thomas Zachariah, Neal Jackson, and Prabal Dutta. The internet of things still has a gateway problem. In *Proceedings of the 23rd Annual International Workshop on Mobile Computing Systems and Applications*, HotMobile '22, page 109–115, New York, NY, USA, 2022. Association for Computing Machinery.

[89] Thomas Zachariah, Neal Jackson, Branden Ghena, and Prabal Dutta. Reliable: Towards reliable communication via bluetooth low energy advertisement networks. In *Proceedings of the 2022 International Conference on Embedded Wireless Systems and Networks*, EWSN '22, page 96–107, New York, NY, USA, 2023. Association for Computing Machinery.

[90] Thomas Zachariah, Noah Klugman, Bradford Campbell, Joshua Adkins, Neal Jackson, and Prabal Dutta. The internet of things has a gateway problem. In *Proceedings of the 16th*

*International Workshop on Mobile Computing Systems and Applications*, HotMobile '15, page 27–32, New York, NY, USA, 2015. Association for Computing Machinery.

[91] Zesen Zhang, Alexander Marder, Ricky Mok, Bradley Huffaker, Matthew Luckie, K C Claffy, and Aaron Schulman. Inferring regional access network topologies: Methods and applications. In *Proceedings of the 21st ACM Internet Measurement Conference*, IMC '21, page 720–738, New York, NY, USA, 2021. Association for Computing Machinery.