Throughput, Coverage and Scalability of LoRa LPWAN for Internet of Things

Asif M. Yousuf, Edward M. Rochester, Behnam Ousat and Majid Ghaderi
Department of Computer Science
University of Calgary

Abstract—**LoRa is a leading Low-Power Wide-Area Network (LPWAN) technology for Internet of Things (IoT). While LoRa networks are rapidly being deployed around the world, it is important to understand the capabilities and limitations of this technology in terms of its throughput, coverage and scalability. Using a combination of real-world measurements and high fidelity simulations, this paper aims at characterizing the performance of LoRa. Specifically, we present and analyze measurement data collected from a city-wide LoRa deployment in order to characterize the throughput and coverage of LoRa. Moreover, using a custom-built simulator tuned based on our measurement data, we present extensive simulation results in order to characterize the scalability of LoRa under a variety of traffic and network settings. Our measurement results show that as few as three gateways are sufficient to cover a dense urban area within an approximately 15 Km radius. Also, a single gateway can support as many as 10^5 end devices, each sending 50 bytes of data every hour with negligible packet drops. On the negative side, while a throughput of up to 5.5 Kbps can be achieved over a single 125 KHz channel at the physical layer, the throughput achieved at the application layer is substantially lower, less than 1 Kbps, due to the network protocols overhead.**

I. Introduction

A. Background and Motivation

The Internet of Things (IoT) is an emerging paradigm in which everyday objects are equipped with Internet connectivity, enabling them to collect and exchange information. With the increasing popularity of IoT devices, it is estimated that by 2025, around 30 billion IoT devices will be deployed around the world, a quarter of which will be connected to the Internet using Low-Power Wide-Area Network (LPWAN) technologies [1]. LPWANs are emerging wireless technologies that complement traditional cellular and short range wireless technologies to address diverse requirements of IoT applications. LPWAN technologies offer long-range connectivity for low power and low rate devices, not provided by legacy technologies. Specifically, LPWAN technologies are considered for those applications that are delay tolerant, do not need high data rates, and typically require low power consumption [2].

Currently, there are several competing LPWAN technologies on the market, such as LoRa [3], Sigfox [4], RPMA [5], Telensa [6], and Weightless [7], each employing a different technique to achieve long-range low-power operation. These technologies are required to provide connectivity for a massive number of heterogeneous IoT devices scattered over a wide geographic area, where devices may communicate over distances exceeding 10 Km [8]. Such a requirement is defined by major applications foreseen for LPWAN, among which are the automotive and intelligent transportation systems (incident report and alerts, fleet management, *etc.*), metering applications (*e.g.*, electrical, water and gas consumption monitoring, medical metering and alerts) and smart homes (*e.g.*, thermostat control and security systems) [8], [9].

As such, one of the main technical challenges with LPWAN technologies is *scalability* as these technologies are required to provide connectivity for a massive number of IoT devices. To keep the complexity of constructing and maintaining the network low, LPWAN technologies often rely on a star topology in which end devices directly communicate with a few so-called *gateways* in a single-hop manner. However, given the fact that LPWANs cover a large geographic area, a large number of end devices have to share the wireless medium. This has profound consequences for the scalability of such networks and naturally raises the question of how many devices can be supported in the same area without dissatisfying application quality of service (QoS) requirements.

While there is extensive push toward development and standardization of LPWAN technologies, there are only a few sporadic works analyzing the performance and scalability of LPWANs for IoT applications (*e.g.*, [2], [10], [11]). In this paper, we focus on LoRa[1], one of the leading LPWAN technologies, and conduct a comprehensive analysis of its key performance metrics including throughput, coverage and scalability. Our objective is to provide concrete data points regarding the performance of LoRa in order to help assess its suitability for IoT applications. The main focus of the paper is on experimental evaluation of the throughput and coverage of LoRa.

B. Related Work

There are few works in the literature on the performance of LPWAN technologies, specifically LoRa, as they are

---

[1]While LoRa refers to the physical layer technology used in LoRa networks a.k.a LoRaWANs, we use the term LoRa throughout the paper to refer to LoRa LPWANs.

relatively new and still under active development and standardization. In the following, we will briefly review a few recent works that are more relevant to our work.

Overviews. An overview of various LPWAN technologies, including LoRa and SigFox, is provided in [12]. The authors qualitatively (*i.e.*, based on technology specifications) compare various LPWAN technologies using metrics such as network topology, hardware cost, and theoretical throughput. The main conclusion is that there is no *one size fits all* solution with each of the technologies having their pros and cons. The work in [13] compares LoRa and RF Mesh technologies in the context of smart grid applications.

Measurements. Petajajarvi et al. [10] focused on the Doppler effect on the performance of LoRa in different mobility scenarios. They showed that mobility of the receiver or transmitter could significantly degrade the communication quality, and hence the coverage of the network. In another work Petajajarvi et al. [14] focused on the range evaluation and channel attenuation model for LoRa technology. The measurement study was conducted in Oulu, Finland over the EU ISM 868 MHz band. The work in [11] provides an overview of various LPWAN technologies and presents the results of a measurement study consisting of a single-cell LoRa deployment in Padova, Italy over the EU ISM band. The measurement experiments show a coverage range of about 2 Km in an urban environment.

DIY Experiments. One of the major obstacles in deploying IoT applications is the cost of building and operating the required communication infrastructure. With the rise of Do-It-Yourself (DIY) electronics (*e.g.*, Arduino and Raspberri PI) and open software projects (*e.g.*, Linux and LMiC), there is a growing list of DIY LoRa network deployments. For instance, Pham [15] presented a low-cost LoRa network for small to medium size IoT deployments. However, no measurements or performance data was reported. Another DIY deployment is reported in [2], where the authors built a LoRa network using off-the-shelf components to build LoRa end devices and gateways. The gateways, however, were deployed indoor. Nevertheless, their measurements indicate that LoRa is able to achieve great indoor coverage even in a harsh propagation environment consisting of a high-rise concrete and steel building.

Simulations. A simulation model for assessing the scalability of a single LoRa gateway is presented in [16]. With two physical end nodes, the authors determine the intra-technology interference behavior and later use this information in their simulator. Using simulations, it is shown that when the number of devices increases to 1000 per gateway, the packet loss rate increases to over 30% [16]. The simulations, however, consider the EU specifications for LoRa, which impose restrictions on the number of channels (*i.e.*, only 8 channels) and radio duty cycling (at most 10%) compared to the North American specifications (please see Section II for more details). In another work,

Martin *et al.* [17] developed a simulator to study the LoRa communication behavior and scalability. However, similar to [16], this work also considers EU specifications at 868 MHz ISM band where 7 spreading factors are available (as opposed to only 4 in North American specifications).

C. Our Work

In this work, we present a comprehensive analysis of LoRa's key performance metrics including throughput, coverage and scalability using real world measurements and detailed simulations. Our work is different from the works described above as we consider North American LoRa specifications and use a city-wide commercial LoRa deployment to conduct our measurements. Almost all of the above works considered European deployments and specifications, where the radio frequency regulations are different from those in North America. For example, while in Europe, only 8 channels at the 868 MHz ISM band are available to LoRa devices, in North America, LoRa Alliance specifies 72 dedicated uplink channels at the 915 MHz ISM band [3]. Thus, it is unclear if the conclusions made in the above works regarding the performance of LoRa remain valid in North America. It is this gap that this work is trying to fill.

Our contributions can be summarized as follows:

- We characterize theoretical and real-world LoRa throughput, while considering different physical layer configurations. We characterize throughput achieved by an end device both at the radio and application level, where all networking and processing overhead is taken into consideration.
- We characterize real-world LoRa coverage by collecting an extensive set of measurements from a city-scale commercial deployment with multiple state-of-the-art gateways. We characterize coverage in three scenarios, namely indoor coverage, outdoor coverage and mobility coverage.
- We characterize LoRa scalability using simulations. We develop a detail LoRa simulator considering North American specifications. The simulator implements many aspects of a LoRa network including channel access, wireless propagation, packet collision, and capture effect. Our simulator can easily simulate networks consisting of hundreds of thousands of devices and multiple gateways.

D. Paper Organization

An overview of LoRa technology is presented in Section II. Our LoRa testbed is described in Section III. LoRa throughput is analyzed in Section IV. Characterization of LoRa coverage is presented in Section V. LoRa scalability analysis is presented in Section VI. Finally, Section VII, concludes the paper.

II. OVERVIEW OF LoRa TECHNOLOGY

LoRa (Long Range) is an LPWAN technology developed by Semtech Corporation [18]. To keep the complexity of

the network low, LoRa relies on a star topology in which end devices directly communicate with a few gateways in a single-hop manner. Gateways in turn forward data received from end devices to a central network server (see Fig. 1). Gateways and end-devices communicate with each other using different frequency channels and data rates, where the selection of a particular data rate provides a trade-off between communication range and message duration.
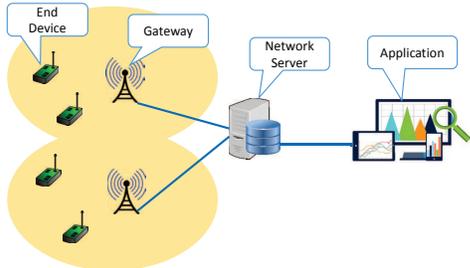
Fig. 1: Typical LoRa network architecture.

In recent years, LoRa has attracted a significant amount of attention due to its inherent ability to efficiently trade communication range for high data-rates, which in return enables it as a compelling communication technology for IoT applications at an urban scale. Semtech specifications define three major components of LoRa networks, namely the physical (PHY) layer, link layer, and the network architecture [19].

A. PHY Layer

LoRa implements Chirp Spread Spectrum (CSS) with integrated Forward Error Correction (FEC) [3]. Due to this design, end devices using different data rates do not interfere with each other. It also operates over multiple channels which increases the capacity of the network. LoRa networks operate in unlicensed ISM frequency band, which for North America is the frequency band $902 - 928$ MHz with center frequency of 915 MHz. For this band, the LoRa specifications define 64 channels of 125 KHz bandwidth from 902.3 to 914.9 MHz in 200 KHz increments. There are an additional eight 500 KHz uplink channels in 1.6 MHz increments from 903 MHz to 914.9 MHz. This brings the total number of uplink channels to 72 channels, although the eight 500 KHz channels are overlapping with the remaining 64 channels. There are eight downlink channels, each 500 KHz wide starting from 923.3 MHz to 927.5 MHz.

Compared to the European regulations, the Federal Communications Commission (FCC) allows a higher peak power of 1 Watt (30 dBm) if the bandwidth of the channel is at least 500 KHz. For lower bandwidths, the LoRa device has to implement Frequency Hopping (FH) with a maximum dwell time of 400 msec per channel. This makes the lowest LoRa date rates not usable, as transmitting the packet preamble alone takes more than 400 msec.

In addition to the above, one must decide on the spreading factor (SF) and coding rate (CR) used by end devices. Such variables are consequential for robustness to interference and time on air of the transmissions. LoRa uses orthogonal spreading factors, which enable multiple packets with different SF's to be transmitted over the same channel concurrently, in return improving network efficiency and throughput. For European deployments, SF is between 7 and 12, while North America specifications define SF between 7 and 10, affecting the time it takes to transmit a packet.

LoRa also implements a form of FEC, which permits the recovery of information in case of transmission errors. Applying FEC requires additional coding data to be included in each transmitted packet, where the amount of coding data is determined by the coding rate. Depending on which CR is selected, one may attain an additional robustness in the presence of interference, with the available options being $\{4/5, 4/6, 4/7, 4/8\}$.

LoRa packet structure at the physical layer includes a preamble, an optional header and the data payload. The preamble is used to synchronize the receiver with transmitter. Optional header contains payload length in bytes, Forward Error Correction (FEC) code rate of the payload and header CRC. The optional header is always protected with the FEC of the lowest (*i.e.*, most robust) coding rate of 4/8.

B. Link Layer

The link layer of LoRa LPWAN networks is referred to as LoRaWAN. The MAC layer that operates on top of the LoRa PHY layer is defined in LoRaWAN specifications. It distinguishes between three end-device classes, namely class A, B, and C, where B and C class devices are required to be compatible with class A devices. Class A devices are optimized for power consumption, where a device receives downlink messages only immediately after an uplink transmission, by opening two short receive windows. In addition to the two receive windows defined for class A devices, class B devices open extra downlink receive windows at scheduled times, where time is synchronized with beacons transmitted by the gateway. Class C devices, on the other hand, continuously keep the receive window open, only closing the window when transmitting.

The channel access mechanism in LoRaWAN is pure ALOHA, in which an end device accesses the channel without sensing the channel for ongoing communications. This is to further prolong device battery life by avoiding spending energy for listening to the communication channel as done, for example, in CSMA-based WiFi networks.

C. Network Architecture

LoRaWAN networks are organized in a star-topology with each gateway directly receiving messages from multiple end-devices. Gateways are connected to a network server and use TCP/IP protocols to communicate with the network server. Each end-device may adjust its data rate manually or using adaptive data rate (ADR) [19]. The network server implements ADR and determines the

optimal data rate to be used by each end device. Since end devices broadcast their messages, the same message may be received by multiple gateways who will forward the message to the network server, where the redundant messages are filtered. Within this network architecture, the network server is also responsible for security, diagnostics and, if so desired, acknowledgments [19].

III. MEASUREMENT SETUP

In this section, we describe the LoRa network where our measurements are conducted. We also briefly describe the equipment used in the measurements.

A. Network Location

The network is deployed in Calgary, Alberta, Canada. Specifically, three commercial-grade gateways (denoted by G1, G2 and G3) are mounted on radio masts located in the close vicinity of the downtown area of the city as shown in Fig. 2. On the figure, the measurement points are marked with I1-I5 (for indoor measurements) and O1-O6 (for outdoor measurements). Gateway locations allow to
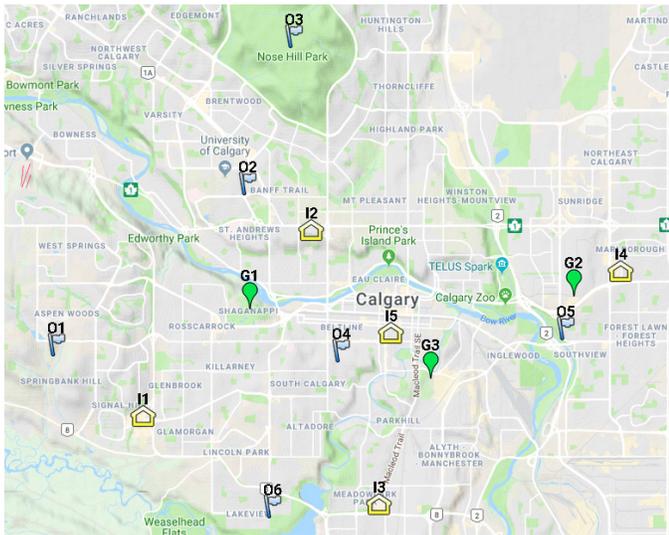
Fig. 2: Gateway locations and measurement points.

have a minimally obstructed line of sight (LOS) to at least one of the gateways from most parts in the city, except for the dense urban area of the downtown, where skyscrapers create large obstacles even at distances of several hundred meters.

B. Gateways

The gateways operate at 915 MHz ISM band. Each gateway supports 64 uplink channels (125 KHz each) and is equipped with two omni-directional antennas with 8 dBi gain. Gateways allow for full-duplex operation on each of the two antennas. Each gateway is mounted on a radio mast that provides an extra height gain of $\approx$ 150 meter, allowing for a wider coverage range. Gateways support the newest LoRaWaN specification (*i.e.*, ver. 1.1 [19]) and run proprietary software for diagnostics and recovery.

C. End Devices

LoRa end devices used in our measurements are depicted in Fig. 3. These devices were built in our lab and manually programmed to operate over different channels.

(a) Node 1
Arduino UNO/RFM95W.

(b) Node 2
Arduino M0/RFM95W.

(c) Node 3
Mbed FRMD/SX1276MB1LAS.

(d) Node 4
Arduino M0 Pro/RFM95W.
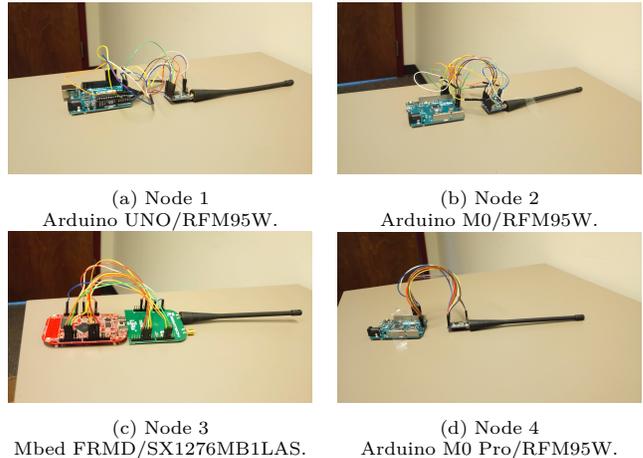
Fig. 3: LoRa end devices used in measurements.

Each end devices is equipped with an omni-directional antenna with 3 dBi gain. The output power for the transmitter is set to 23 dBm. Since our main goal is to measure the coverage of the network, the spreading factor (SF) is set to the highest, *i.e.*, SF 10, with 125 KHz channel bandwidth. During measurements, each end device broadcasts its messages as soon as possible following pure ALOHA. Each message consists of a LoRaWAN pre-defined header and hard-coded payload of size 1 byte. Arduino-based devices (Nodes 1, 2 and 4 in Fig. 3) run the LMiC library [20], which is a slightly modified version of the LoRaWAN implementation by IBM. The Mbed-based device (Nodes 3 in Fig. 3) runs the reference implementation of LoRaWAN provided by Semtech.

IV. THROUGHPUT CHARACTERIZATION

The throughput of a LoRa end device depends on its transmission mode, where each mode is specified by a combination of spreading factors, coding rates and channel bandwidth. To characterize the throughput of LoRa devices, two approaches, namely theoretical and experimental, were followed. First, using the Semtech's published specs for LoRa [19], we compute the maximum throughput achievable for each mode of operation, *i.e.*, combination of channel bandwidth, coding rate and spreading factor. This is the theoretical throughput and provides an upper bound on the maximum transmission rate that can be achieved by a device under ideal channel and interference conditions. Following the theoretical analysis, we then perform real-world measurements to characterize device throughput.

A. Theoretical Throughput

Using Semtech's published specifications for LoRa modulation [19], the relationship between the achieved data

(a) Theoretical transmission rate. (b) Measured transmission rate at PHY. (c) Measured transmission rate at APP.
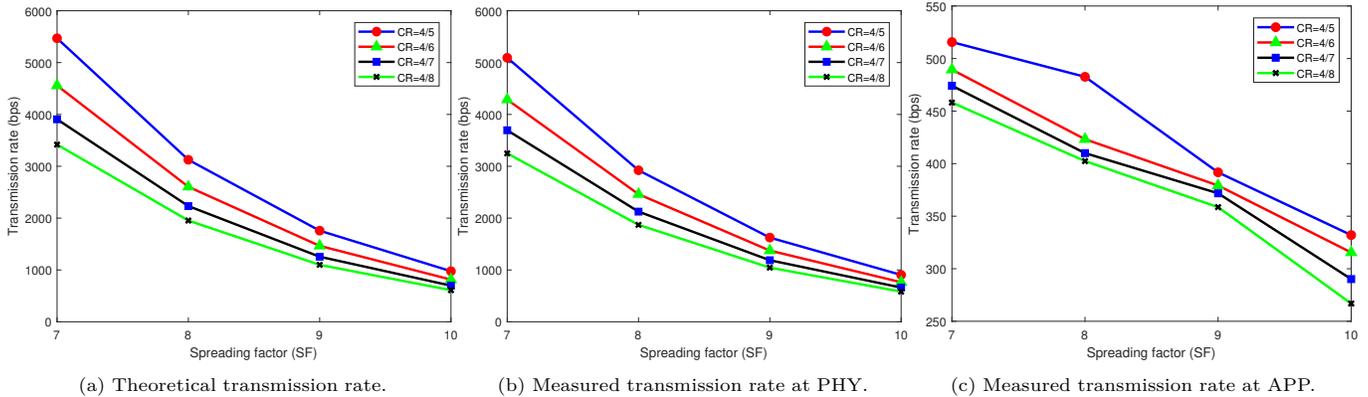
Fig. 4: Throughput characterization. PHY and APP refer to measurements at radio and application level, respectively. APP throughput is an order of magnitude smaller than PHY throughput due to protocol and application overheads that significantly contribute to increased packet transmission times at the application level.

rate, channel bandwidth, spreading factor and coding rate is expressed as,

$$\text{Transmission Rate} = SF \times \frac{BW}{2^{SF}} \times CR, \qquad (1)$$

where, SF is the spreading factor, BW is the channel bandwidth and CR denotes the coding rate.

Fig. 4(a) shows the theoretical throughput for channel bandwidth 125 KHz, as this channel bandwidth is commonly used in LoRa deployments. As expected, the combinations (SF=7, CR=4/5) and (SF=10, CR=4/8) result in the highest and lowest transmission rates, respectively. Specifically, the highest possible rate is 5468 bps, while the lowest rate is 610 bps. Note that these transmission rates are the raw data rate at the physical layer. The application layer throughput is lower due to LoRaWAN protocol overhead. In particular, each LoRaWAN packet has a minimum header length of 13 bytes.

B. Measured Throughput

For the measurements, as recommended by Semtech, we transmit maximum sized packets (255 bytes) and measure how much time it takes for each packet transmission to finish. Each point on the plot is based on the average of 100 packet transmissions for each SF and CR combination.

The results are depicted in Fig. 4(b) and Fig. 4(c), where PHY and APP transmission rates are defined as follows:

- *PHY*: To compute the PHY throughput, the transmission time of the packet is measured at the radio transmitter. In this case, the transmission time does not include the time spent on packet creation, encryption and any other application specific operation.
- *APP*: To compute the APP throughput, the transmission time of the packet is measured at the application level and includes the time from when the packet creation starts until the transmission is completed.

From Fig. 4(b), we observe that the highest and lowest PHY transmission rates are given by 5092 bps and 581 bps, respectively. These numbers are remarkably close

to the numbers presented in Fig. 4(a). The application layer transmission rate, on the other hand, is significantly lower than the theoretical numbers. Moreover, it can be seen that there is not much difference between the APP transmission rate across different combinations of CR and SF. Specifically, we see a 10x difference between the lowest and highest PHY rates. However, when looking at APP rates, the highest and lowest rates are 515 bps and 266 bps, respectively, which represent a 2x difference only. This means that while the radio behaves differently with respect to CR and SF, the high level application and protocol operations affect the achievable throughput significantly.

V. Coverage Characterization

Our objective in this section is to characterize the coverage of LoRa in an urban environment (*i.e.*, the City of Calgary) under three different scenarios, namely *Indoor*, *Outdoor* (where end devices are stationary) and *Mobile* (where end devices are placed in a moving car).

A. Coverage Criteria

To characterize coverage at each location, we calculate the *packet delivery ratio* (PDR) at that location. PDR is the ratio of the number of packets successfully received at the network server over the total number of packets transmitted by an end device. To calculate the packet delivery ratio at each location, four end devices operating over different channels are used to transmit packets simultaneously. Each device transmits 200 packets back-to-back. We then compute the average PDR at each location using the calculated packet delivery ratio of all devices.

B. Outdoor Measurements

Outdoor stationary measurements were conducted at a number of locations across the city. For the ease of exposition, we only present the measurement results for six locations identified on the map as O1 to O6 (see Fig. 5). The results are summarized in Table I.

To visualize the network coverage, we imposed the PDR

results for the gateways on Google maps to create a heatmap of the aggregate network coverage across all gateways, as shown in Fig. 5.
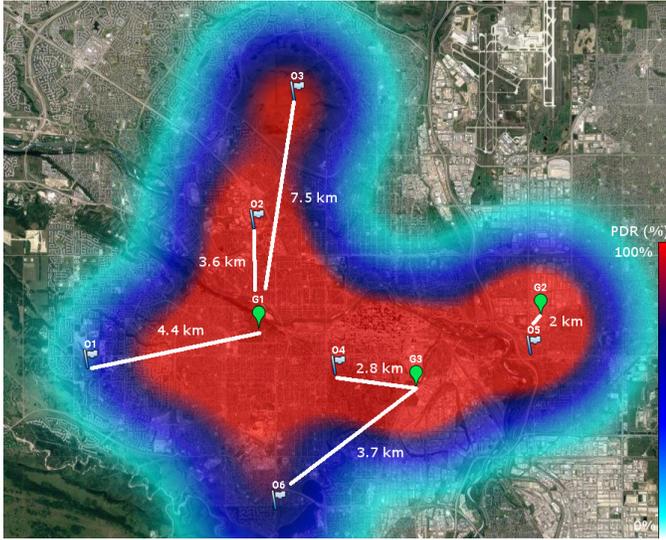
Fig. 5: Network coverage heatmap.

From the Table I, we can see that even at a distance of 7.5 Km, 95.3 % PDR is achieved between O3 location and gateway G1. At the same time, we see that even though the gateways deployed in this network are placed close to the downtown area, the coverage of the network extends far beyond to the edges of the city. We speculate that with more optimal gateway placement, even 3 gateways would be enough to cover the whole city. These results confirm that effective long-range communication (≥ 10 Km) can be achieved using LoRaWAN even in urban environments.

TABLE I: Outdoor measurement results.

	Average PDR (%)		
Locations	G1	G2	G3
O1	2.6	1.9	1.4
O2	97.1	0.6	18.4
O3	95.3	3.1	0
O4	97.6	75.8	0.1
O5	77	98.1	99.9
O6	25.9	25.1	48.9

C. Indoor Measurements

In indoor measurements, the RF signals had to propagate through different materials to reach gateways. For locations I1-I4, signals had to propagate through concrete walls and metal piping within the buildings. The location I5 allowed for access to the window, thus signals had to propagate only through glass and concrete walls on the other side of the building.

Table II shows the PDR values for each of the locations with the last column summarizing average PDR over all gateways. Notice how drastically PDR changes by moving end devices indoor. When devices were located away from the windows and closer to the middle of the building,

surrounded by concrete walls and other obstacles, even from I1, which is located close to G1 gateway, PDR drops to as low as 23%.

Another major factor is the altitude at which the devices are located, as can be seen for I5. While from the bottom floor no transmissions were received, as soon as devices were placed on the top floor of the building, the PDR went as high as 49%. Another observation is the difference between I2 and I4 PDRs. Although, I4 and I2 are relatively close to G2 and G1, respectively, their PDRs are significantly different (PDR at I2 is almost 0). The reason is the altitude difference between the locations, where there is an ≈ 170 meter difference in altitude between I2 and I1.

Based on our results, although it is feasible to use LoRaWAN in indoor scenarios, more gateways are required to provide effective communication as many factors such as the building construction materials and line-of-sight obstacles significantly affect signals, sometimes even resulting in 0% PDR.

TABLE II: Indoor measurement results.

	Average PDR (%)			
Locations	G1	G2	G3	All
I1	23.0	0.0	0.0	23.0
I2	0.0	0.0	0.0	0.0
I3	0.0	6.0	53.0	57.6
I4	0.0	96.6	0.0	96.6
I5 (Top Floor)	26.6	1.3	49.5	53.4
I5 (Bottom Floor)	0.0	0.0	0.0	0.0

D. Mobility Measurements

In this scenario, the end devices were put on the front seat of a car, which was driven on the routes depicted in Fig. 6. The routes were driven in both directions. Collected measurement results were then grouped into *low-speed* (average movement speed of 50 Km/h) and *high-speed* (average movement speed of 80 Km/h) route results.
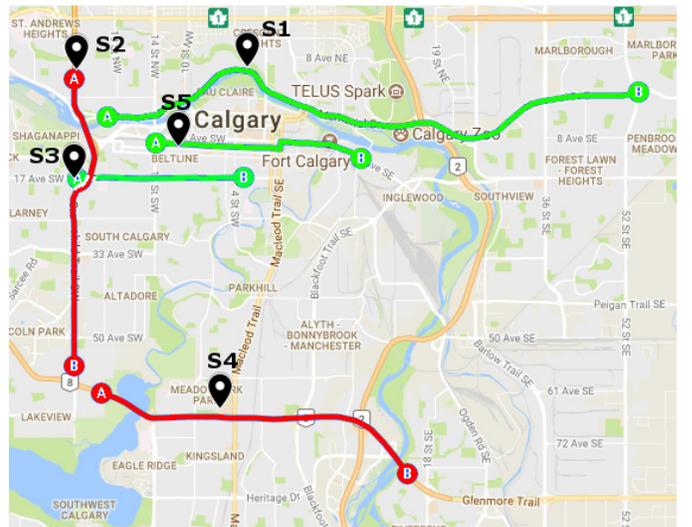
Fig. 6: Routes for mobility measurements. Low-speed and high-speed routes are colored as green and red on the map.

Additionally, Table III summarizes the results for stationary reference point measurements when devices transmit from within the car, while the car is stopped. These results allow us to see how transmitting from inside the car with no mobility affects PDR. We can clearly see that transmitting from inside the car lowers PDR. Even for S2 location that is only a couple kilometers away from the gateway with clear LOS, PDR is only about 86.6%.

TABLE III: Reference point measurements.

Location	Average PDR (%)		
	G1	G2	G3
S1	70.3	27.4	14.2
S2	86.6	4.7	-
S3	71.9	7.9	29.2
S4	17.8	19.1	45.5
S5	89.1	-	13.3

Table IV summarizes the average PDR for high-speed and low-speed routes. Notice that, although high-speed and low-speed routes are located in close proximity of G1, low-speed routes have higher PDR for G2. This can be explained by the fact that most of the low-speed routes are located in urban, downtown area and so, even though G1 is closer, it is obstructed by buildings throughout the measurement process, making it easier for transmissions to reach G2.

TABLE IV: Mobility measurement results.

Routes	Average PDR (%)			
	G1	G2	G3	All
Low-Speed (50 km/h)	34.5	62.4	22.9	71.4
High-Speed (80 km/h)	64.4	2.9	2.4	66.8

The last column of Table IV summarizes average PDR across all gateways for high-speed and low-speed routes. We observe that although low-speed routes have higher PDR on average, the difference in our measurements was not significant at only about 5%. It can be concluded that other factors such as the topology of the city and LOS obstructions as well as the actual location of devices (*i.e.*, inside the car vs. outside the car) play a more significant role in determining PDR than the speed of movement.

VI. Scalability Analysis

In this section, we analyze the scalability of LoRa using our custom-built LoRaWAN simulator. While live network measurements provide useful insight about the throughput and coverage of LoRa, it is very difficult to change the parameters of the production network. Moreover, it is very expensive (and perhaps infeasible) to characterize the network performance *at scale*, when a large number of end devices are deployed in the network. Therefore, to further study the performance of LoRa and assess its scalability, we have developed a simulator that is capable of simulating a LoRa network consisting of multiple gateways and a large number of end devices. The simulator parameters are tuned based on the actual measurements conducted in the network.

A. LoRaWAN Simulator

In this subsection, we describe the design and specifications of the LoRaWAN simulator.

1) Simulator Design: The simulator is designed as a discrete-event simulator and implemented using Java programming language[2]. The simulator considers all the current features of LoRaWAN specification [19]. The simulator is configured to simulate LoRaWAN operation based on North American specifications. It allows full configuration of gateways, end devices and network parameters, as described below:

- **Gateways:** The simulator accepts an input configuration file for gateways that describes the coordination of gateway locations, and their PHY parameters such as the trasmit power and number of channels.
- **End Devices:** The simulator accepts an input configuration file to specify the location and PHY as well as application-level parameters of each end device. For PHY parameters, one can specify the transmit power, spreading factor, coding rate, preamble and payload size, and operating channel. For application parameters, one can specify, for each device, the type of traffic (*e.g.*, deterministic or stochastic) and the parameters of the specified traffic model such as the inter-arrival time of packets.
- **Network Parameters:** The simulator allows specifying a variety of network parameters such as the parameters of the propagation channel model including path-loss exponent and shadowing parameters.

2) Packet Reception Model: To determine if a packet is correctly received at a gateway (*i.e.*, the packet can be successfully decoded), the gateway calculates the received signal strength indicator (RSSI) associated with the packet and compares it with the sensitivity threshold of the LoRa radio receiver used at the gateway. The received RSSI is calculated using the following relation,

$$RSSI = PX + GL + PL(d), \qquad (2)$$

where PX is the transmit power of the end device in dB, GL combines all gains and losses in the tramsmit/receive path in dB, and $PL(d)$ represents the path loss in dB assuming that the distance between the end device and gateway is d meters. The calculated RSSI is then compared with the sensitivity thresholds reported in Table V. The values in this table are extracted from the specifications of the popular LoRa transceiver RFM95/96/97/98(W) [21]. These values can be easily changed for other transceivers.

3) Wireless Propagation Model: To calculate the path loss, we implemented the log-normal shadowing model, where the parameters of the model are estimated from our

---

(a) Capacity with SF7.   (b) Capacity with SF10.
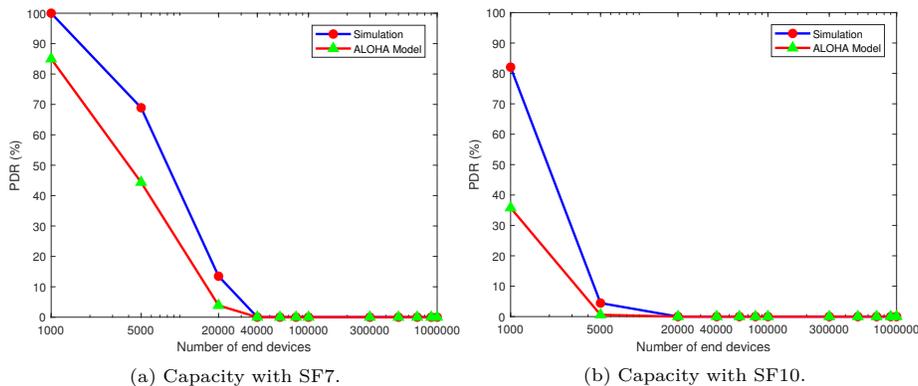
Fig. 7: Effect of packet capture on single-channel gateway capacity. Simulation results consider capture effect to decode more packets resulting in higher capacity compared to the ALOHA model.

TABLE V: RFM95/96/97/98(W) LoRa receiver sensitivity at 125 KHz bandwidth.

Spreading Factor (SF)	RSSI (dBm)
7	-123
8	-126
9	-129
10	-132

measurement data. The path loss can be calculated from the following formula:

$$\text{PL}(d) = \text{PL}(d_0) + 10\alpha \log_{10}(d/d_0) + X, \qquad (3)$$

where $\text{PL}(d_0)$ (in dB) is the reference path loss value at distance d_0 meters from the gateway, α is the path loss exponent, d is the distance between the end device and gateway in meter, and X is the random shadowing modeled as a zero mean log-normal variable with standard deviation σ_x dB. Using our measurement data, we considered the reference distance d_0 as 1000 meter. With this information, in the urban environment, we calculated $\text{PL}(d_0)$=130.12 dB, $\alpha = 2.1$, and $\sigma_x = 7.79$ dB.

4) Packet Collision Model: When multiple LoRa transmissions arrive at the gateway at the same time on the same channel with the same bandwidth (BW) and spreading factor (SF), there are several conditions which determine whether the gateway can decode one or multiple signals or nothing at all. The LoRa packet structure consists of a preamble, an optional header and data payload. The preamble is used to synchronize the receiver with the transmitter, the header contains the payload length in bytes, FEC code rate of the payload and header CRC. Following the analysis presented in [16], in our simulator, we determine the collision behavior and capture effect using the following rules:

1) For more than one concurrent receptions at gateway, if the interfered transmission has non overlapping preamble and header reception time, and the interferer RSSI is less than or equal to the interfered RSSI, interfered packet will be received successfully.

2) If the interferer RSSI is greater than 6 dB, interfered packet will be lost even if the interfered packet has non-overlapping preamble and header reception.

3) Both interferer and interfered packet will be lost if there is no non-overlapping preamble and header reception.

B. Scalability Analysis

For scalability, we focus on the *capacity* of a single gateway, where the capacity of a gateway is defined as the number of end devices that can be supported by a gateway at a pre-specified PDR. As such, the capacity depends not only on network parameters, *e.g.*, number of channels, but also on the traffic load generated by end devices.

For the ease of exposition, the simulated network consists of one 64-channel gateway. End devices are distributed uniformly randomly around the gateway in such a way that there is no packet loss due to propagation distance as our focus is on the gateway capacity under ideal conditions. The transmit power of each device is set to 23 dBm. For each experiment, we consider a one day simulation run time, where each data point is obtained as the average of 25 simulation runs. Table VI summarizes the default parameters that are used in simulations unless otherwise specified.

TABLE VI: Default end device configuration.

Parameter	Value
Tx Power	23 dBm
CR	4/5
BW	125 KHz
Payload Size	50 Bytes
Packet Inter-Arrival Time	20 Minutes

1) Effect of Collisions and Capture: To study the effect of the collision model implemented in the simulator on the gateway capacity, we have compared the simulation results with theoretical results derived from the analysis of pure ALOHA. We have developed a model to calculate gateway capacity by extending the standard analysis of ALOHA [22] to consider LoRa network specifications. The
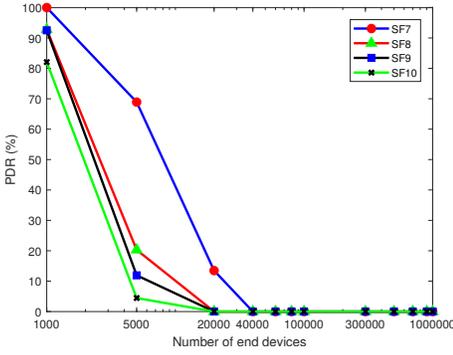
Fig. 8: Effect of spreading factor on single-channel gateway capacity. SF7 supports significantly more devices compared to the other SFs.
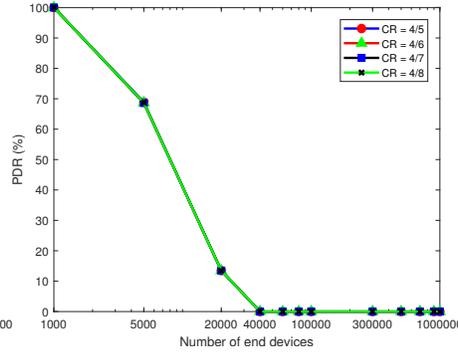
Fig. 9: Effect of coding rate on single-channel gateway capacity. All coding rates achieve almost the same performance.
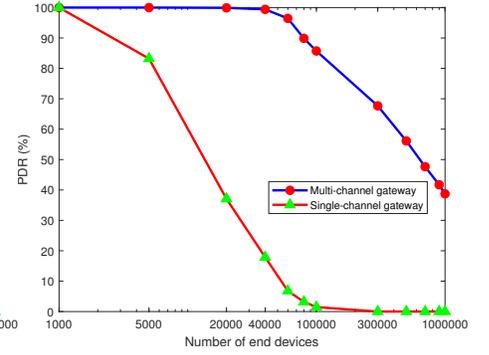
Fig. 10: Effect of number of channels on gateway capacity. The capacity increases almost linearly as the number of channels increases.

ALOHA model considers all concurrent transmissions as *collisions*, while the simulator tries to decode them, as described earlier.

In this experiment, we consider a single-channel gateway. We show the results for SF7 and SF10, in Figs. 7(a) and 7(b), as these two spreading factors have the highest and lowest data rates, respectively. As expected, as the number of devices increases, the packet delivery ratio (PDR) decreases in both cases. We observe that slightly better PDR is achieved for LoRa simulator compared to the ALOHA model, as the simulator considers the capture effect. Notice that for the highest data rate (*i.e.*, SF7), almost 69% PDR is achieved for 5000 end devices, while for the same number of end devices, only 5% PDR is achieved for the lowest data rate (*i.e.*, SF10). The reason is that it takes much longer to transmit packets at lower rates, resulting in more collisions.

2) Effect of Spreading Factor: Capacity of a single channel gateway is evaluated in this experiment. The results are presented in Fig. 8. As expected, increasing the spreading factor (*i.e.*, decreasing transmission rate) results in lower PDR. This is because increasing SF results in longer transmission times, which cause more collisions and keep channels busy for longer period of times.

An unexpected observation is the difference between the effect of spreading factors. Specifically, SF7 (*i.e.*, the fastest SF) supports significantly more devices compared to other SFs. This can be justified with respect to the relation between the transmission rate and spreading factor. As expressed by (1), the transmission rate is proportional to the product of SF and $1/2^{SF}$, which is dominated by the term $1/2^{SF}$ for higher SFs. In other words, for higher SFs, the transmission rate drops exponentially fast, hence the significant differences between the performance of SF7 and other SFs.

3) Effect of Coding Rate: Capacity of a single channel gateway is evaluated in this experiment. The results are presented in Fig. 9. We observe that changing the coding rate does not have a noticeable effect on PDR. This can

be explained by looking at the effect of coding rate on throughput in Fig. 4(c). While decreasing the coding rate results in more robust transmissions (*i.e.*, lower decoding errors), this is compensated for by the slight decrease in transmission rate (*i.e.*, higher air time).

4) Effect of Number of Channels: To increase gateway capacity, most LoRa gateways support multiple channels. LoRa specs for North America allow 64 non-overlapping uplink channels at 125 KHz each. In this experiment, we evaluate the capacity of a multi-channel gateway with 64 channels and compare it with that of a single-channel gateway. Each end device chooses its spreading factor and coding rate randomly from the set4 spreading factors and 4 coding rates. For the multi-channel gateway, each device also randomly chooses one of the 64 channels. The results are depicted in Fig. 10. As expected, the gateway capacity increases linearly proportional to the number of channels.

5) Effect of Payload Size: In this experiment, we consider a multi-channel gateway with 64 channels. We change the size of the payload for end devices and measure the corresponding PDRs. The results are presented in Fig. 11. From the figure, it can be seen that for smaller number of end devices, the payload size does not affect the gateway capacity as the gateway is not saturated. As the number of devices increases and the gateway becomes saturated, however, increasing the payload size decreases the capacity, as expected. Specifically, for 10^6 devices, the PDRs achieved for each payload size are separated by about 10%, which is quite significant.

6) Effect of Message Inter-Arrival Time: In this experiment, we consider a multi-channel gateway with 64 channels. The results are presented in Fig. 12. As expected, sending more packets (*i.e.*, shorter inter-arrival time) results in lower gateway capacity.

7) Effect of QoS Requirements: In this experiment, we directly compute the capacity of the multi-channel gateway for different required PDRs. We change the traffic load and compute the gateway capacity by gradually increasing the number of end devices until the target PDR cannot be
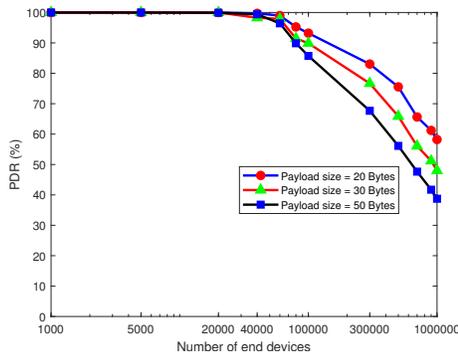
Fig. 11: Effect of payload size on gateway capacity. As the number of end devices increases, the effect of payload size becomes more prominent.
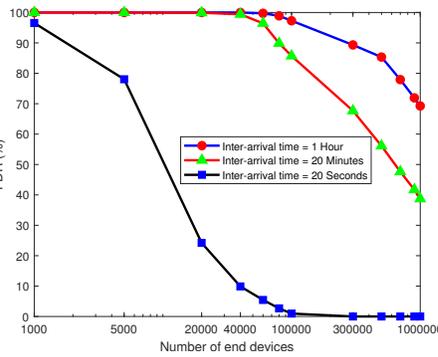
Fig. 12: Effect of message inter-arrival time on gateway capacity. As with payload size, the effect is more pronounced with high number of devices.
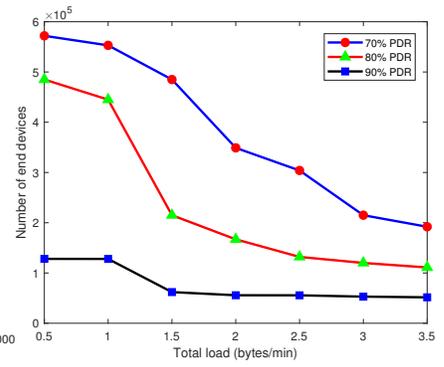
Fig. 13: Effect of QoS requirements on gateway capacity. Very high PDR requirements lead to a drastic reduction in gateway capacity.

satisfied. The number of devices at that point is referred to as the capacity of the gateway. The results are presented in Fig. 13. It can be seen that for high PDR requirements, the effect of traffic load on gateway capacity is marginal. In contrast, for lower PDR requirements, the effect of traffic load on the capacity is significant. Specifically, at traffic load of 1 byte/min, while 5.7×10^5 devices can be supported at 70% PDR, only 1.1×10^5 devices can be supported at 90% PDR, which is a significant reduction in capacity.

VII. Conclusion

In this paper, we analyzed key performance metrics of LoRa including throughput, coverage and scalability using live measurements and simulations. The measurements were conducted using a city-wide LoRa deployment. Our measurements indicated that LoRa end devices can achieve throughputs that are suitable for low-rate IoT applications, while enjoying the long range offered by LoRa technology. We observed varying coverage quality outdoors and indoors with the main impediment being the high density building obstacles in the urban setting. Our results alos indicated that LoRa networks are highly scalable. In fact, a single gateway can support hundreds of thousands end devices assuming that the traffic load of each device is low.

Acknowledgment

References

[1] Nokia, "LTE evolution for IoT connectivity," Accessed: August, 2017. [Online]. Available: https://tools.ext.nokia.com/asset/200178

[2] A. M. Yousuf, E. Rochester, and M. Ghaderi, "A low-cost lorawan testbed for iot: Implementation and measurements," in *Proc. IEEE 4th World Forum on Internet of Things*, 2018.

[3] LoRa Alliance. [Online]. Available: https://www.lora-alliance.org

[4] Sigfox. [Online]. Available: https://www.sigfox.com

[5] Ingenue RPMA. [Online]. Available: https://www.ingenu.com/technology/rpma

[6] Telensa. [Online]. Available: http://www.telensa.com

[7] Weightless. [Online]. Available: http://www.weightless.org

[8] Nokia, "LTE M2M - optimizing LTE for the Internet of Things," Accessed: August, 2017. [Online]. Available: https://novotech.com/docs/default-source/default-document-library/lte-m-optimizing-lte-for-the-internet-of-things.pdf

[9] S. Andreev *et al.*, "Understanding the IoT connectivity landscape: A contemporary M2M radio technology roadmap," *IEEE Com. Mag.*, vol. 53, no. 9, 2015.

[10] J. Petäjäjärvi *et al.*, "Performance of a low-power wide area network based on LoRa technology: Doppler robustness, scalability, and coverage," *Journal of Distributed Sensor Networks*, vol. 13, no. 3, 2017.

[11] M. Centenaro, L. Vangelista, A. Zanella, and M. Zorzi, "Long-range communications in unlicensed bands: The rising stars in the IoT and smart city scenarios," *IEEE Wireless Communications*, vol. 23, no. 5, 2016.

[12] U. Raza, P. Kulkarni, and M. Sooriyabandara, "Low power wide area networks: An overview," in *IEEE Communications Surveys & Tutorials*, 2017.

[13] H. Filho, J. Filho, and V. Moreli, "The adequacy of LoRaWAN on Smart Grids: A comparison with rf mesh technology," in *Proc. IEEE Smart Cities*, 2016.

[14] J. Petäjäjärvi, K. Mikhaylov, A. Roivainen, and T. Hänninen, "On the coverage of LPWANs: Range evaluation and channel attenuation model for LoRa technology," in *Proc.14th International Conference on ITS Telecommunications*, 2015.

[15] C. Pham, "Building low-cost gateways and devices for open LoRa IoT test-beds," *Journal of Testbeds and Research Infrastructures for the Development of Networks and Communities*, 2017.

[16] J. Haxhibeqiri *et al.*, "Lora scalability: A simulation model based on interference measurements," *Sensors*, vol. 17, no. 6, 2017.

[17] M. Bor, U. Roedig, T. Voigt, and J. M. Alonso, "Do lora low-power wide-area networks scale?" in *Proc. 19th INternational Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, 2016.

[18] Semtech Corporation. [Online]. Available: http://www.semtech.com/

[19] Semtech, "LoRaWAN specification v1.1." [Online]. Available: https://www.lora-alliance.org/technology

[20] LMiC-Arduino Library, Accessed: August, 2017. [Online]. Available: https://github.com/matthijskooijman/arduino-lmic

[21] H. Electronic, Accessed: February, 2018. [Online]. Available: http://www.hoperf.com/upload/rf/RFM95_96_97_98W.pdf/

[22] J. Kurose and K. Ross, *Computer Networking: A Top-Down Approach*. Pearson, 2017.