Citywide LoRa Network Deployment and Operation: Measurements, Analysis, and Implications

Shuai Tong[1], Jiliang Wang[1]✉, Jing Yang[1], Yunhao Liu[1], Jun Zhang[2]
[1]Tsinghua University [2]ENNEW Technology Co., Ltd.
tl19@mails.tsinghua.edu.cn, jiliangwang@tsinghua.edu.cn, jing-yan18@mails.tsinghua.edu.cn
yunhao@tsinghua.edu.cn, junzhang@enn.com

ABSTRACT

LoRa, as a representative Low-Power Wide-Area Network (LPWAN) technology, holds tremendous potential for various city and industrial applications. However, as there are few real large-scale deployments, it is unclear whether and how well LoRa can eventually meet its prospects. In this paper, we demystify the real performance of LoRa by deploying and measuring a citywide LoRa network, named *CityWAN*, which consists of 100 gateways and 19,821 LoRa end nodes, covering an area of 130 km^2 for 12 applications. Our measurement focuses on the following perspectives: *(i)* Performance of applications running on the citywide LoRa network; *(ii)* Infrastructure efficiency and deployment optimization; *(iii)* Physical layer signal features and link performance; *(iv)* Energy profiling and cost estimation for LoRa applications. The results reveal that LoRa performance in urban settings is bottlenecked by the prevalent blind spots, and there is a gap between the gateway efficiency and network coverage for the infrastructure deployment. Besides, we find that LoRa links at the physical layer are susceptible to environmental variations, and LoRa and other LPWANs show diverse costs for different scenarios. Our measurement provides insights for large-scale LoRa network deployment and also for future academic research to fully unleash the potential of LoRa.

CCS CONCEPTS

• **Networks** → **Network measurement**.

KEYWORDS

LoRa, Network measurement, Internet of Things

ACM Reference Format:
Shuai Tong, Jiliang Wang, Yunhao Liu, Jun Zhang. 2023. Citywide LoRa Network Deployment and Operation: Measurements, Analysis, and Implications. In *The 21st ACM Conference on Embedded Networked Sensor Systems (SenSys '23), November 12–17, 2023, Istanbul, Turkiye.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3625687.3625796

1 INTRODUCTION

Nowadays, Low-Power Wide-Area Networks (LPWANs) are garnering increasing attention owing to their long communication range and low power consumption, which is suitable for connecting large-scale Internet of Things (IoT). As a representative LPWAN technique, LoRa envisions an anywhere, anything connectivity paradigm for providing low data rate, long-range, and low power connections to various IoT applications. The application market of LoRa has made significant progress in recent years. It is reported that the number of LoRa-connected devices has exceeded 225 million by 2021, and the LoRa device market is expected to reach $11.6 billion by 2028 across a broad range of industries [1, 2].

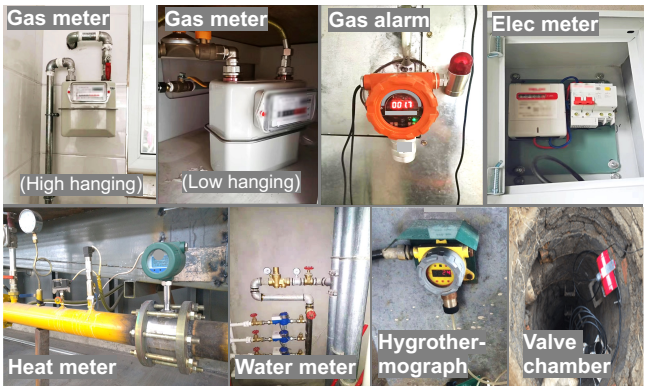

Figure 1: The deployment environments of LoRa based municipal devices in the *CityWAN*.

Despite the huge popularity, one should be cautious that the LoRa technology development is still in a very preliminary stage. The functionality and performance of LoRa in long-term and large-scale systems still need to be measured and validated. Previous literature has presented that LoRa at the current stage has limitations in terms of coverage [3–5], capacity [6–9], and link reliability [10–12]. Therefore, one natural question comes – "At this stage, how far away is LoRa from its prospects and what does it take to reach its vision of connecting massive devices and applications? "

Status Quo and their Limitations. Many empirical studies have been devoted to LoRa to understand its performance and limitations. Liando et al. [13] build a campus-scale LoRa network consisting of 50 end nodes and 4 gateways to measure the energy consumption and network capacity with various LoRaWAN parameters. Ren et al. [14] measure the PHY-layer links and LoRa communication ranges in both LoS (Line-of-Sight) and non-LoS scenarios, using six mobile LoRa end nodes and two gateways. Other LoRa measurements, such as [15–18] and references therein, evaluate the link capacities, channel attenuations, and energy profiles of end nodes with various LoRa deployments.

Fundamental limitations. All previous measurements are performed on small-scale LoRa testbeds with only several gateways

SenSys '23, November 12–17, 2023, Istanbul, Turkiye

ACM ISBN 979-8-4007-0414-7/23/11.
https://doi.org/10.1145/3625687.3625796

Table 1: Number of deployed devices in various smart city applications.

Device Name	Number	Device Name	Number
Gas meter	18765	**Water meter**	591
Heat meter	245	**Gas stove**	17
Gas alarm	143	**Regulator valve**	26
Electricity meter	19	**Hygrothermograph**	2
Intelligent module	8	**Pressure sensor**	3
SOS Emergency button	1	**Gate magnetism**	1

and tens of end nodes. They mostly simulate large-scale LoRa connections by having end nodes transmit packets at high duty-cycles. The scalability of LoRa networks when connecting massive real IoT devices has not been verified. Also, testbeds in previous measurements are not designed for real-world applications. Their deployment environments and data transmission modes are very different from the real-world applications. The functionality and performance of LoRa in long-term and large-scale real applications still need to be measured and validated.

Overview of our Measurement. In this paper, we report a long-term and full-aspect measurement on the *CityWAN* large-scale LoRa network. Our measurement differs from all previous works in the scale of the evaluated network and whether the connected nodes are used in real applications. We build the network to connect citywide municipal devices, which consists of 100 gateways and 19,821 LoRa nodes, covering an area of 130 km^2 and serving 12 kinds of smart city applications as shown in Table 1. The deployment environments of LoRa end devices are shown in Figure 1. We focus on the deployment, performance, and operation challenges of the large-scale network for real-world smart city applications. We collect more than 389M transmission records with fine-grained cross-layer network information from a five-month measurement. We evaluate the performance of the LoRa network from the application layer to the physical layer, revealing the performance gap between the LoRa coverage and gateway efficiency. We profile the energy consumption of LoRa transmitters and compare LoRa with other LPWAN technologies in terms of the end node lifetime and system costs under different scenarios and user densities. The purpose of this work is to provide guidance for large-scale LoRa deployment and show insights for future academic research on LoRa network optimization.

Measurement perspectives. Our measurement aims to demystify LoRa performance from four major perspectives:

(i) Application performance (Sec. 3). LoRa enables IoT applications with long-range and low-power communications. However, the long-term performance of applications on large-scale LoRa networks still needs to be measured and validated. We evaluate the performance of LoRa applications from high-level data transports down to the physical layer link behaviors. We also summarize network characteristics and analyze the bottleneck that restricts LoRa application performance.

(ii) Infrastructure deployment (Sec. 4). Gateways in LoRa are customized by users, which is considered one of the advantages of LoRa. The installation and maintenance of gateways are costly and affect the coverage of LoRa networks. Thus, it is of great significance to optimize the gateway deployment. In this paper, we investigate the impact of gateway deployment on network performance by measuring gateway efficiency for the citywide LoRa network. Based on the measurement, we reveal a gap between the gateway efficiency and network coverage, which inspires us to optimize the future deployment of LoRa networks.

(iii) Link characteristics (Sec. 5). LoRa claims to provide long-range communication links, i.e., five kilometers in urban areas and 15 kilometers or more in rural areas, by its physical layer chirp spreading spectrum (CSS) mechanism. However, nodes in real-world applications are deployed in diverse environments, such as underground or surrounded by metal materials. The practical performance of LoRa links faces many attrition factors, e.g., signal interference, rich multi-path, and link attenuation of different shields. To understand how these factors manifest LoRa links in practice, we perform link measurements over different transmission times, channels, gateway redundancies, and surrounding environments. We analyze the root cause of link fluctuations and present innovations for improving the reliability of LoRa links.

(iv) Energy Consumption and System Costs (Sec. 6). Energy consumption and costs are first-order concerns of IoT devices as they are battery-powered and should be low-cost for large-scale deployment. We measure the power profile of LoRa devices and build an economic model for the deployment of LoRa networks, considering the whole lifecycle of IoT applications. We compare the system costs of LoRa, Sigfox, and NB-IoT in urban and rural scenarios. Our measurement identifies the key costs of LPWAN deployments and answers the advantages and disadvantages of each technology in different application scenarios.

Summary of insights. Our measurement leads to several major insights, which we summarize as follows:

(i) Our measurement reveals that packet loss and coverage blind spots are much more prevalent for real urban applications than in testbed measurements. Blind spots exist even when gateways are densely deployed in the network region. Many blind spots happen with node deployments in indoor scenarios, especially when the nodes are deployed on the bottom floors or surrounded by shields. We also find that the blind spots in urban applications are spatially scattered, making it expensive and inefficient to improve the node coverage by adding more gateways.

(ii) As for the infrastructure deployment, we find that the importance of different gateways in the network is very different, where 40 critical gateways can cover 95.3% of the end nodes, but connecting the rest 4.7% of nodes needs 60 more gateways. This indicates a gap between the gateway efficiency and network coverage, where the full coverage demands a large number of low-utilization gateways to connect sparse users at edge areas or blind spots.

(iii) We find that LoRa links vary dramatically across transmission times, channel frequencies, and surrounding environments. We first identify that the time-related link fluctuation is due to the spectrum overlapping of Digital Terrestrial Multimedia Broadcast (DTMB) signals and LoRa bands. Then, for links of different channels, we find that transmission SNRs show short-term aggregation and long-term fluctuation in each channel due to the rich multi-path from urban reflections. Multi-path signals over different frequencies have different wavelengths, leading to constructively superposition

at some channels and destructively cancellation at others. Finally, we verify that the installation environments of end nodes have a significant impact on the link performance, as nodes in real-world applications may be deployed in hard-to-reach locations.

(iv) Our energy consumption measurement shows that LoRa has the best power efficiency among different LPWAN technologies for low data rate IoT applications. This is mainly due to the low protocol overhead of LoRa transmissions. We further estimate the system costs of large-scale LPWANs and find that LoRa is cost-effective for urban low-density scenarios, but NB-IoT is cost-effective for rural and high-density scenarios. This inspires us to apply mixed technologies to build IoT systems with the best cost efficiency.

Contributions. To our knowledge, this work represents the first measurement study for the deployment and operation of a citywide LoRa network for real-world applications. Our main contributions can be summarized as follows: *(i)* Deploying and operating a large-scale LoRa network for connecting municipal devices in real-world applications. *(ii)* Quantitative characterization of LoRa transmission and coverage profile in urban settings, providing guidance for large-scale network deployment. *(iii)* Locating the bottleneck of LoRa performance by fine-grained connectivity measurements. *(iv)* Identifying the gap between gateway efficiency and network coverage to inspire network deployment optimization. *(v)* Profiling LoRa links under various transmission times, channel frequencies, and environments. *(vi)* A detailed measurement of the LoRa energy profile and cost comparison to other LPWAN technologies. *(vii)* Releasing the collected dataset and analysis scripts as open source at https://github.com/iot-book/CityWAN.

2 STUDY METHODOLOGY

Network Deployment. We build the *CityWAN* LoRa network to connect IoT devices in an urban area of 130 km^2. The network consists of 100 gateways and 19,821 end nodes, where gateways are with SX1301 [19] and end nodes use SX1278 [20]. We deploy the gateways on the roofs of tall buildings with the minimum surrounding obstacles. We also set two USRP gateways which can collect physical layer samples for fine-grained LoRa link measurement.

We use the network to connect twelve smart city applications. The number of devices for each application is shown in Table 1, where gas meters occupy the main body of the network, accounting for 94.7% of end nodes. In uplink transmissions, end nodes send data to gateways, and gateways pass the data to the network server (NS) via backbone network connections. The NS then forwards the data to the application server (AS) as necessary. In turn, the NS sends downlink messages, either for network management or on behalf of the AS, through gateways to end nodes. Figure 1 shows the deployment environments of various municipal devices. Most nodes are deployed inside buildings, and some are embedded in cabinets or underground. For example, gas meters and water meters are mostly low-hanging in cabinets close to the ground. This brings significant attenuation and rich multi-path to wireless signals. To accommodate the harsh environment, we adopt an over-provisioning gateway deployment to ensure coverage, where the average distance between two neighboring gateways is 0.72 km, and the minimum gateway distance is 0.2 km.

At the protocol layer, we apply the LoRaWAN for MAC and upper network stacks. We configure end nodes in Class A, where all transmissions are node-initiated and a downlink transmission happens only following an uplink transmission. Each successful uplink transmission is followed by an ACK generated by the network server and forwarded by a selected gateway. Most nodes in our network transmit regularly for application data collection. By default, end nodes report their data readings with four uplink packets every day. Each data packet is modulated in SF10 with a Code Rate (CR) of 4/8. The inter-packet interval for two adjacent packets in the same report is five seconds, and each packet has 64-byte payloads. For each uplink transmission, end nodes randomly select a target carrier from eight predefined LoRa channels with center frequencies ranging from 475.1 MHz to 476.5 MHz.

Data Collection. We need to collect fine-grained network data for in-depth LoRa performance measurement. However, data collection in the citywide network is very challenging. The first challenge is the inability to get physical access to end nodes. As shown in Figure 1, end nodes in our network are mostly deployed in difficult-to-access locations. Getting physical access to an end node can take a deployment manager days or weeks to schedule a time when the user is available. Therefore, we have to recover the end node information, such as uplink attempts and packet loss, only based on the records at the gateway side. We resolve this challenge by taking functionalities of the LoRaWAN protocol. LoRaWAN uses the "*NbTrans*" parameter to control retransmissions of node uplinks. When "*NbTrans*" is set to 1, there is no physical layer retransmission. Then, at the gateway, we extract the frame counters (i.e., *fCnt*) of uplink packets. The frame counters for each end node should be a monotonically increasing sequence whose value indicates the number of uplink attempts, and numbers skipped in the sequence correspond to lost packets of that device. We process the potential packet loss at the application layer by designing a transmission acknowledgment mechanism at the NS.

Second, for end-to-end link estimation, we need to estimate the power of LoRa signals and interference, respectively. LoRa gateways provide SNR and RSSI estimations of each uplink packet. However, as LoRa supports high-sensitive reception for signals that are upto 20 dB under the noise floor, the estimated RSSI cannot reflect the actual signal power of LoRa packets. To settle this, we use the metric of Expected Signal Power (ESP) for revealing the actual signal strength even with extremely low SNRs. We borrow from the deduction in [21] and estimate the ESP from RSSI and SNR of received LoRa packets based on the independence between signals and interference/noise:

$$ESP_{[dBm]} = RSSI_{[dBm]} + SNR_{[dB]} - 10 \cdot \log_{10}\left(1 + 10^{0.1 SNR_{[dB]}}\right)$$

We use the ESP estimation for evaluating the LoRa coverage as well as the end-to-end link performance. We collect other network information such as transmission times, channel frequencies, and device IDs for all records at the gateway side. We also collect the geographic locations and installation environments for all gateways and some end nodes for analyzing the environmental impact on network performance.

General Network Statistics. Figure 2(a) shows the number of end nodes in three main applications over time. The node number increases consistently during the whole measurement period (around

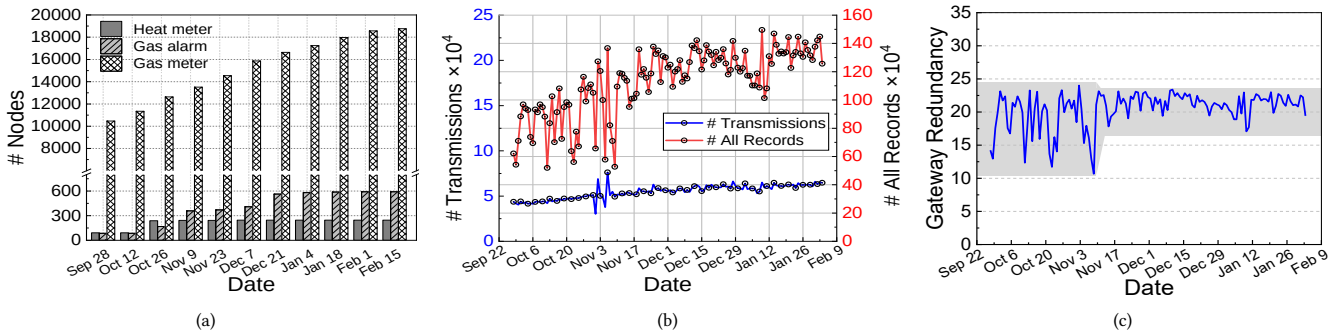

Figure 2: General statistics of *CityWAN*: (a) Number of devices for various applications over the measurement period; (b) Number of transmissions at end nodes and number of receiving records at gateways; (c) Gateway redundancy estimated by ratios of gateway-side records to node-side transmissions.

35% growth), as the deployment of municipal facilities is an ongoing activity. We collect a total of 389,899,615 transmission records from 19,810 user devices and 100 gateways over the five-month network measurement. Figure 2(b) shows the number of uplink transmissions at the end node side and the number of receiving records at the gateway side. There are fluctuations in the number of gateway records from Sep 25 to Nov 7, which is mainly due to the unexpected record loss at the NS database by a configuration fault. We diagnosed and fixed it on Nov 7, and after that, the number of gateway records increases much more steadily with the increase of network devices. We calculate the ratio between the numbers of gateway receptions and node transmissions for evaluating the receiver redundancy of uplink transmissions as shown in Figure 2(c), where each transmission is received by 16 to 24 gateways on average. This indicates that uplink transmissions in our network have high receiver redundancies, where each packet is received by multiple gateways due to the high density of gateway deployment in the urban area.

3 APPLICATION PERFORMANCE

This section characterizes the performance of applications running on the large-scale *CityWAN* LoRa network. We intend to illustrate how well large-scale LoRa can support different IoT applications and to provide realistic experiences for application users. Our observations range from the high-level data transmission performance down to behaviors at the physical layer. From those observations, we summarize the network characteristics and explore reasons that restrict the application performance in large-scale LoRa networks.

3.1 Data Transmission Performance

We measure the data transmission performance to show the capability of LoRa for data collection. We first evaluate the packet loss of each end node based on the sequence numbers of their uplink packets. Then, we look down to the physical layer to evaluate the SNR performance of each transmission. The following presents our observations:

(1) Poorly connected nodes exist even with dense gateway deployments. Figure 3 shows the CDF of packet loss rate for three different applications. We can clearly see the difference in the distribution of packet loss rates between nodes in different applications. Nodes of gas meters have the best data transmission performance of all applications, where the packet loss rate for 94.4% of nodes is lower than 5%, and 99.1% of nodes have a rate lower than 10%. The median packet loss rates for heat meters and gas alarms are 8.2% and 7.9%, respectively. However, there still exist end nodes whose data cannot be reliably collected, where 197 gas meters, 10 heat meters, and 18 gas alarms have packet loss rates over 10%. Such a situation is not usual especially considering the dense deployment of gateways, where an uplink transmission can be received by 20 redundant receivers on average. We make a deep analysis and reveal the root cause for transmission failures of these poorly connected nodes in Sec. 3.2.

(2) There is a long tail in the SNR distribution of LoRa transmissions. Figure 4 shows the SNR distribution of transmissions in three different applications. For each transmission, we exclude redundant records at multiple gateways and only use the SNR records that are finally forwarded to the application server. Due to the dense deployment of gateways, most transmissions have high SNRs, e.g., 92.3% of gas meter transmissions, 88.5% of gas alarm transmissions, and 58.8% of heat meter transmissions have the SNRs in range of 0 to 15 dB. However, many transmissions are occurring far away from the central part of the distribution, showing a long tail in the low-SNR region between −20 and 0 dB. There are two reasons for those low-SNR transmissions. The first is the burst interference in the unlicensed LoRa band where the superposition of interference can lead to decreasing SNRs. The second is weak links of poorly connected nodes. We analyze the roots for these low-SNR transmissions in Sec. 3.2.

In addition, we observe that transmissions in the gas metering application have better reliability than other applications, considering both the packet loss and transmission SNRs. This is because the gas metering nodes are mostly installed in urban housing estates which have dense gateways nearby. Besides, we prioritize the connectivity to gas meters by adding gateways to regions with dense node installations, as gas metering is an essential municipal service. Such results indicate that an intensive gateway deployment can certainly improve the application reliability but is difficult to achieve complete coverage for all nodes. We map the SNR distributions to packet delivery rates (PDR) to show the application layer functionality as in Figure 5. We estimate the PDR at different SNRs

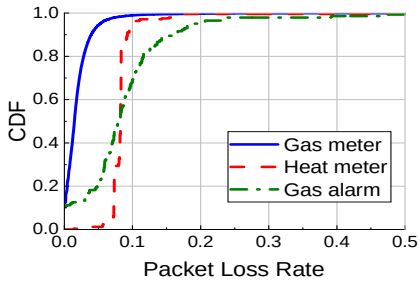

Figure 3: CDF of packet loss rates for different applications.

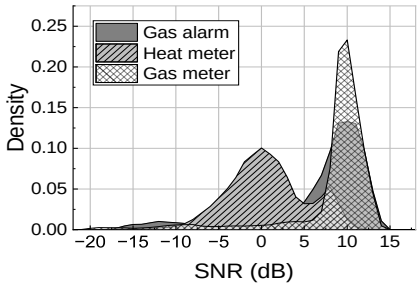

Figure 4: SNR distribution of transmissions in different applications.

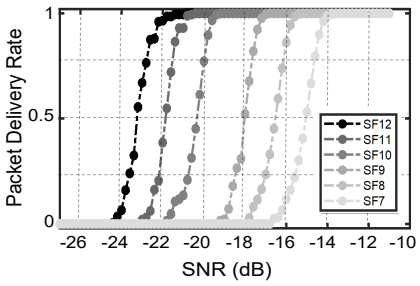

Figure 5: Packet delivery rate of standard LoRa for various SFs and SNRs.

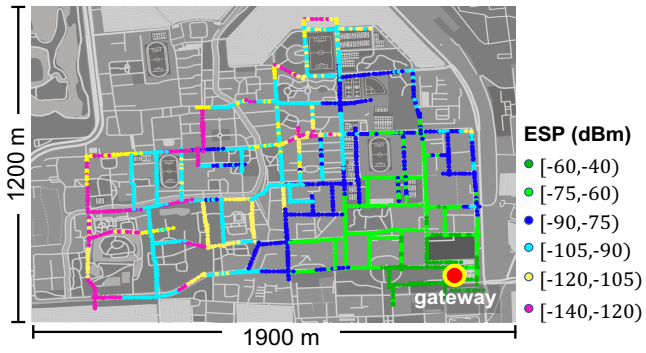

Figure 6: Fine-grained coverage measurement for the ESP map in a typical urban environment.

by connecting an SX1278 LoRa transmitter to an SX1301 gateway with an adjustable attenuator. As the SNR decreases, the PDR first remains constant and then undergoes a quick drop when the SNR falls below the demodulation boundaries. This indicates that end nodes with SNRs below the demodulation boundaries will be disconnected from the network, severely affecting the application performance.

3.2 Locating the Performance Bottleneck

This section studies the fundamental factor that impacts the data transmission performance and locates where the packet loss anomaly takes place. To understand the spatial characteristic of LoRa transmissions, we perform a fine-grained measurement by focusing on a typical urban area. Figure 6 shows the bird view of the measurement region. We do not rely on the data of real applications because their spatial density and transmission frequency are too low to satisfy the fine-grained measurement. Instead, we use a GPS-equipped LoRa node to traverse the measurement region, and thus estimate transmission links at different locations. We use a USRP gateway deployed on the roof of an office building for collecting uplink transmissions. We configure the LoRa node to periodically report its locations to the gateway by sending LoRa packets. The hardware and radio configurations are the same as devices in the gas metering application. Then, at the gateway, we obtain node locations as well as the corresponding ESP values based on the uplink transmission records.

Figure 6 shows all the ESP samples over the whole experiment region. Transmissions around the gateway have strong signals, where the minimum ESP is higher than −100 dBm for measurements in the one-kilometer range. As the distance extends, the measured ESP decreases due to signal attenuation and shadowing. However, we observe that the link quality does not vary monotonically as the distance changes. Some points close to the gateway have worse links compared to positions that are further away. We name such positions of poor network connections as *blind spots*. To show this more clearly, we measure links from farther distances. The results are shown in Figure 7. Generally, the ESP decreases as the link distance extends. Yet a long-distance link is not necessarily worse than a near-region one. *It proves that the gateway coverage is non-isotropic and blind spots exist even for links near the gateway.* The presence of *blind spots* explains why some nodes have high packet loss and some transmissions are with low SNRs even with a high gateway density. A similar problem is observed in cellular and satellite networks, which are known as *urban canyons* or *street canyons* [22, 23]. The shadowing and multi-path effects of tall buildings beside urban streets lead to poor signal reception. However, we observe that blind spots in LoRa follow a more scattered distribution in spatial. As shown in Figure 6, there are multiple isolated regions of blind spots within one kilometer of the gateway, each with a diameter of only tens or even a few meters. This is because LoRa signals use low-frequency carriers and are more susceptible to diffraction and interference during signal propagation.

3.3 Lessons and Implications

(i) Blind spots is non-negligible in LoRa networks, i.e., 7.1% of our evaluated locations have PDR lower than 50%. Although previous literature has shown that the coverage radius of LoRa in the urban environment can reach several kilometers [13], we find that the blind spots are prevalent even when end nodes are close to the gateway. The problem is much more pronounced for indoor applications where nodes are installed with a large number of signal reflectors and shields in the surrounding area. This makes LoRa challenging for applications that require high transmission reliability and complete coverage.

(ii) Blind spots in the urban scene are scattered. Covering nodes in these spots by replenishing gateways will be both expensive and inefficient. Blind spots in urban environments cover small areas but are scattered everywhere. Deploying gateways at each blind spot

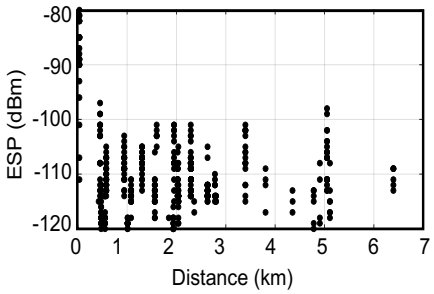

Figure 7: LoRa estimated signal power over different distances: ESP fluctuates with increasing of the link distance.

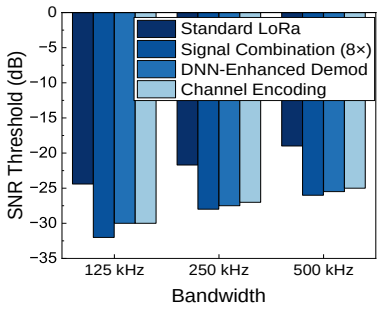

Figure 8: Weak signal demodulation: SNR thresholds for different demodulation mechanisms.

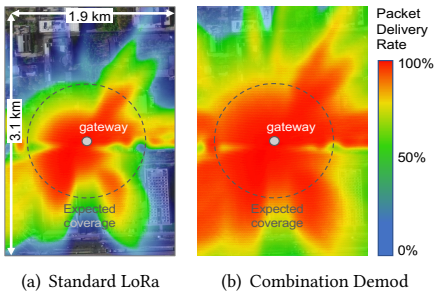

Figure 9: Coverage areas of standard LoRa and the signal combination based demodulation mechanism.

costs a very high capital expenditure and increases the difficulty of network maintenance. The spatial discreteness of blind spots is mainly due to the rich reflection of LoRa signals in urban environments, where the link multi-path of two adjacent locations varies significantly. As a result, LoRa networks with different layouts will all be affected by the blind spots, no matter whether end nodes are uniformly distributed or densely gathered in center areas, as the blind spots are only determined by surrounding physical environments. This problem should be carefully resolved before applying LoRa for reliable network coverage in urban scenarios.

(iii) Weak signal demodulation can help connect nodes in blind spots. A native idea to combat transmission failures in blind spots is to enhance the signal over weak links. We can coherently combine signals of multiple repeated LoRa chirps to improve the signal strength. To achieve this, we program the LoRa node to generate LoRa packets with each symbol repeated multiple times. Then, we modify the gateway to demodulate weak signals by concentrating the energy of multiple repeated symbols. Recent literature validates the effectiveness of signal combinations for boosting weak signal power. For example, Choir [3] launches simultaneous transmissions of multiple co-located transmitters for signal enhancement over the propagation channel. XCopy [24] boosts signal strength by generating continuous retransmissions of the same content and adds up their signals at the gateway. We validate the performance of the coherent signal combination with commodity LoRa nodes and USRP gateways. We estimate the SNR threshold for successful demodulation of signal combination and compare it with the standard LoRa, DNN-enhanced LoRa demodulator [12], and active channel interfering based encoding [11]. The DNN demodulator [12] exploits the feature abstraction ability of deep learning for weak signal demodulation, which we implement on dedicated gateways with strong computability. The channel encoding [11] approach transmits data bits on low-SNR links by selectively interfering with other LoRa transmissions. The resulting SNR thresholds of different bandwidths are shown in Figure 9. All approaches improve the weak signal demodulation ability compared with the standard LoRa. The signal combination approach achieves 7 ∼ 9 dB SNR gains which can be further improved by gathering more repeated symbols. We also evaluate the transmission range and coverage performance of the combination mechanism using physical layer signal samples collected in Sec. 3.2. Figure 9(a) and (b) shows the packet delivery rate of standard LoRa and signal combination demodulation in the same area. The results show that we can improve the LoRa network coverage by extending the signal transmission range and connecting blind spots with the weak signal demodulation mechanism.

4 INFRASTRUCTURE DEPLOYMENT

This section investigates the impact of gateway deployment on network performance. We aim to provide experience for network operators who deploy and maintain network infrastructures. We measure the efficiency of gateways in our citywide network and reveal the performance gap between gateway efficiency and network coverage.

4.1 Gateway Efficiency

We first answer the question – *"How many gateways are used by each end node in the current network deployment?"* Different from cellular technologies such as 4G/5G and NB-IoT, end nodes in LoRa are not associated with a specific gateway. A LoRa uplink transmission will be received by multiple neighboring gateways. Then, each gateway decodes and reports the message to the network server (NS), where the NS chooses the best-SNR message to forward to the application server (AS) and discards all other duplicate ones. In this measurement, we account for three numbers of gateways for each node: *(1) Actually Used Gateways:* the number of gateways whose message has been chosen for forwarding to the AS; *(2) Average Reached Gateways:* the average number of gateways that a node can reach per transmission; *(3) Maximum Reached Gateways:* the maximum number of gateways that a node can reach in all of its transmissions. Figure 10 shows the CDF of the three gateway numbers for each end node. The median numbers of the three kinds of gateways are 7, 18, and 35, respectively. This indicates that for nodes in the current network, the number of reachable gateways is far more than the actual need for uplink transmissions. Many gateways are involved in redundant reporting as most of their messages will be discarded at the NS.

Our second question is – *"Do all gateways have equal importance in the network?"* For answering this, we classify gateways into two categories (replaceable or irreplaceable) depending on whether all traffic on a gateway can be replaced by other gateways. We perform the classification under different application demands, i.e., the *maximum transmission intervals*. Nodes in our network report uplink data every day. At the application layer, however, it is acceptable

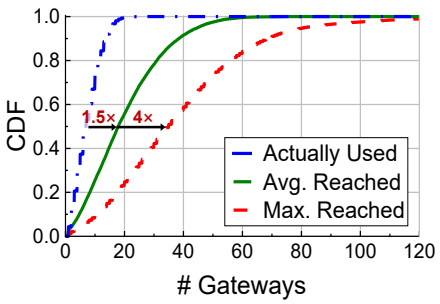

Figure 10: Numbers of gateways that end nodes can connect.

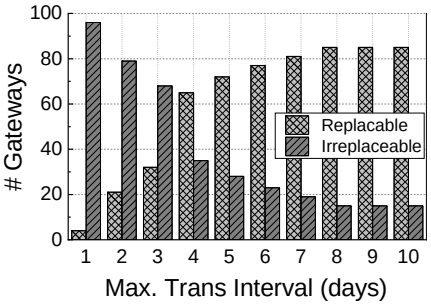

Figure 11: Gateway classification for different interval demands.

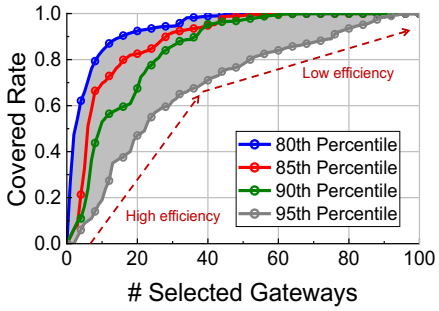

Figure 12: Node covered rate over different number of applied gateways.

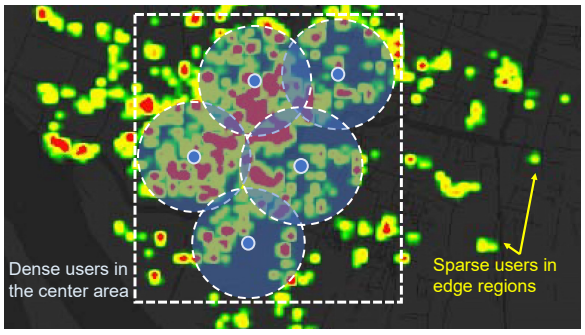

Figure 13: User distribution map of the urban scenario.

for a node to have packet loss, as long as the packet is received before a predefined maximum transmission interval. The shorter the interval, the more reliable the link is required by the application. Figure 11 shows the distribution of the two categories of gateways under different maximum transmission intervals. We have two observations from the measurement: *(1) Under a reasonable maximum transmission interval, most gateways in the network are replaceable by others.* For example, when relaxing the maximum interval to one week, 81 out of 100 gateways are replaceable, i.e., a breakdown on any of the 81 gateways does not affect the running of the network. This is due to the high redundancy in the gateway deployment, where coverage areas of different gateways overlap significantly. *(2) Some gateways play an irreplaceable important role in the network.* No matter how the maximum interval changes, traffics on some gateways can never be replaced by any other gateways, e.g., 15 gateways always being irreplaceable even if the maximum interval reaches 10 days. This shows that the importance of different gateways in the network is very different. A few critical gateways can affect links to a mass number of end nodes.

4.2 Efficiency and Coverage Gap

Now, we have shown that the importance of gateways in the network is inequable, and traffics on some gateways can be carried by other gateways. Based on this observation, we ask the question – *"What is the minimum number of gateways needed to cover all end nodes in a network?"* We model the gateway selection as a classical set cover problem (SCP). Given a set of elements $\{1, 2, ..., n\}$ (called the universe whose items correspond to end nodes) and a collection $\mathbb{S}$ of m sets whose union equals the universe, the SCP is to identify the smallest sub-collection of $\mathbb{S}$ whose union equals the universe. The decision version of an SCP is NP-complete. Therefore, we use a greedy algorithm to find an approximation of the minimum gateway set in polynomial time. We choose gateways according to one rule: at each stage, choose the gateway set that contains the largest number of uncovered end nodes. We consider a node is covered if its PDR is higher than a predefined threshold (i.e., 80%, 85%, 90%, and 95% in our measurements). Figure 12 shows the node covered rate over different numbers of selected gateways with four PDR thresholds. As the number of selected gateways increases, the number of covered nodes grows rapidly at first and then slows down. We have two findings: *(1) A small number of gateways can cover most nodes in the network.* For example, when setting the PDR threshold to 90%, 40 gateways can cover 95.3% end nodes. This is due to the fact that most end nodes in urban applications are densely deployed. Therefore, a few carefully selected gateways can cover a large number of dense users. *(2) As the number of gateways increases in the network, the coverage gain brought by the new gateway decreases.* For example, in the 90% PDR threshold case, when the number of gateways increases to more than 40, each newly added gateway can cover only 15 extra nodes on average. This is mainly due to the uneven spatial distribution of users. Figure 13 shows the user distribution in a typical urban region, where most users are gathered in the center area and some sparse users are located at the edges. Dense users in the center area can be covered by a small number of carefully selected gateways. In contrast, users in edges and blind spots need more gateways as they are sparsely distributed and far apart from each other. Covering those sparse disconnected users requires a lot of extra gateways, which sacrifice the utilization efficiency of gateways. This measurement shows a gap between the gateway efficiency and network coverages, where a complete coverage expects a large number of low-utilization gateways for connecting sparse end nodes in edges or blind spots.

4.3 Lessons and Implications

(i) The efficiency and coverage gap leads to unbalanced gateway utilization when connecting end nodes. Load balancing among different gateways should be considered for the infrastructure deployment. Our measurement shows that unbalanced gateway utilization is common in large-scale networks when the operator takes the coverage as the first consideration and applies the over-provisioning

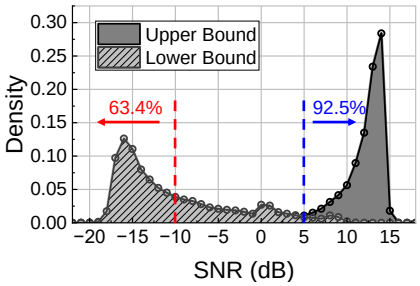

Figure 14: Upper and lower bounds of SNRs for each link.

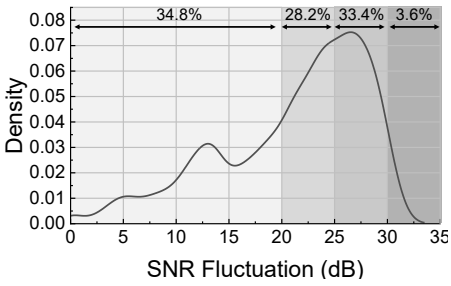

Figure 15: Distribution of SNR fluctuations in all links.

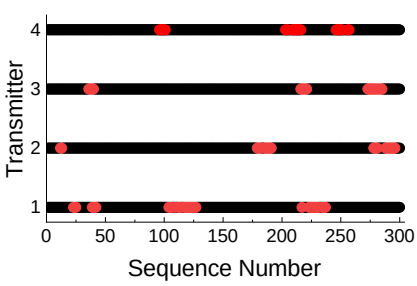

Figure 16: Bursty loss pattern of LoRa transmissions.

infrastructure deployment. Some gateways in the network play an irreplaceable important role with high utilization, while some gateways are less used. The unbalanced gateway loads affect the reliability of the network, where failures or interference on several pivotal gateways will lead the whole network to get into trouble. The deployment of LoRa gateways is user customized. Therefore, we should carefully design the deployment parameters, such as site locations and antenna directions, to balance the gateway overloads and thus improve the network reliability.

(ii) Low-cost connectivity to sparse users is the key to bridging the efficiency and coverage gap. The measurement shows that the over-provisioning gateway deployment is inefficient due to the efficiency and coverage gap. To reduce the infrastructure costs as well as promising coverage, we can leverage peer nodes with reliable connections to ferry data for weak links, and thus connect edge users and blind points at low costs. By applying this idea, we only need a small number of gateways at a few carefully selected locations (such as in areas with dense users), and use the connected end nodes to forward transmissions for uncovered edge users. Such forwarding nodes can also be specially deployed by network operators (e.g., as micro gateways), which cover much smaller areas than normal gateways but are low-cost. Previous studies have validated implementing the packet forwarder on SX1278 LoRa nodes while keeping compatibility with the LoRaWAN protocol [25]. Measurements in [26] show that the range of device-device LoRa communication in urban scenarios can reach 500 m, which will be sufficient for connecting sparse edge users and blind spots. More technique challenges, such as low-power neighbor discovery and collision avoidance, will be our future work.

5 LINK PERFORMANCE

In this section, we estimate link performance between LoRa nodes and gateways. We first present the overall picture of LoRa links in terms of packet loss and physical layer SNRs. Then, we zoom in on the link variation over transmission time, channels, and environmental impacts.

5.1 Overall Picture

Based on our long-term link measurements in the citywide LoRa network, we have the following observations:

(1) Most links fluctuate significantly during the measurement period. To reveal the reliability of link connection in LoRa, we estimate the upper and lower boundaries of transmission SNRs in each link, in Figure 14. There are 92.5% of links with a maximum SNR over 5 dB, showing superb connectivity. But there are also 63.4% of links that have experienced very low SNRs below −10 dB. The distribution of SNR fluctuations for all links is shown in Figure 15. SNR fluctuations of 65.2% of links are higher than 20 dB, showing drastic link variation during the measurement. Meanwhile, some links show much more stable connection performance, where 21.4% of links have SNR fluctuations below 10 dB. For the rest of this section, we explore reasons for such fluctuations by estimating link variations over different times, channels, and environments.

(2) Packet loss in LoRa exhibits a clear bursty pattern. We then explore the packet loss characteristics for LoRa links based on sequence numbers of reception records. We identify packet loss by detecting missing elements in sequence numbers. Figure 16 shows the sequence numbers from four typical gas meter nodes. We observe a clear bursty pattern in packet loss of LoRa links. This is mainly caused by the burst interference from coexisting protocols on ISM bands, which inspires us to carefully design the LoRa retransmission mechanism to avoid time-persistent interference. As LoRa works in the ISM band, there exist rich interference and packet corruptions over the open channel. A possible solution to this problem is to apply carrier sensing before LoRa transmissions [26, 27].

5.2 Link Variation over Time

To verify link performance over different transmission times, we group packet records according to the hours they are transmitted. We extract the SNR and ESP records from each group and estimate their distributions. Figure 17 shows the SNR distribution of transmissions over 24 hours. We find that transmissions from 0 to 6 AM have better SNRs, which are 5dB higher than the average SNR of the rest 18 hours. To show the time-related link variation more clearly, we perform fine-grained measurements on a LoRa link with continuous transmissions. We configure a gas meter to continuously transmit uplink packets every 140 seconds, and plot its SNR traces in Figure 18. There are significant SNR fluctuations over time, where transmissions between 0 and 6 AM have better SNRs. This observation makes a lot of sense, especially considering that we need to decide the appropriate transmission time for end nodes in smart city applications. To dig out the cause of link fluctuations, we perform long-term channel tracking by analyzing possible interference sources. As a result, we find that the spectrum of the Digital Terrestrial Multimedia Broadcast (DTMB) in our city is partially overlapped with the LoRa band. Therefore, it

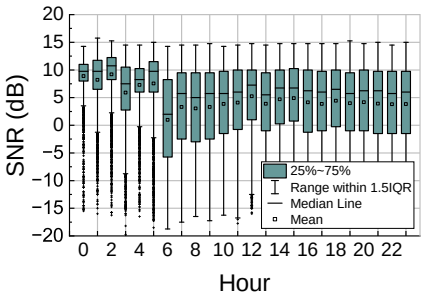

Figure 17: SNR distribution for transmissions at different times.

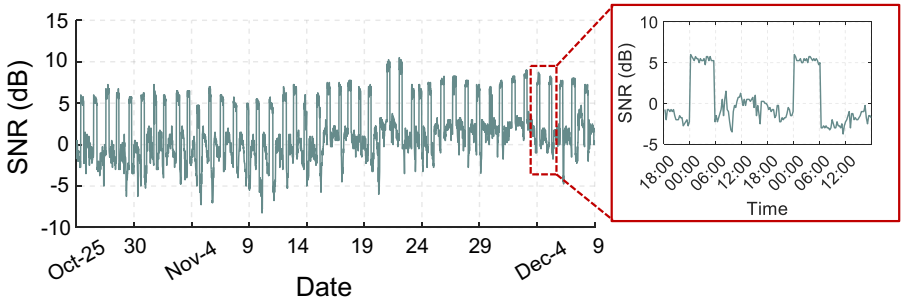

Figure 18: Fine-grained SNR measurements on a typical LoRa link, showing link fluctuations with different transmission times of a day.

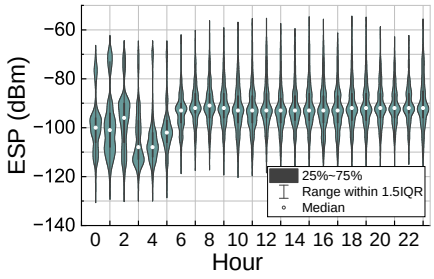

Figure 19: ESP distribution at different times of the day, which increases after 6 AM due to interference.

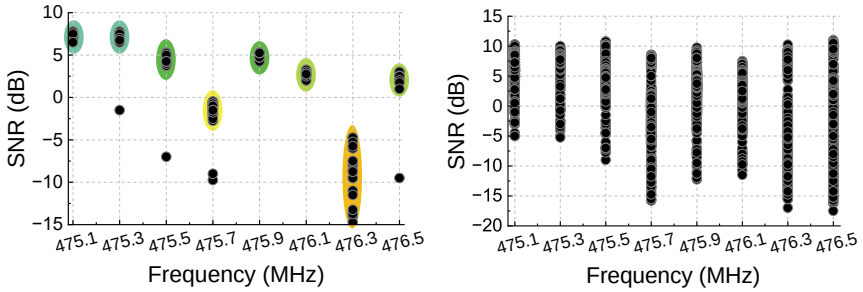

Figure 20: Link SNRs over channels show the trend of (a) aggregation for the short-term measurement in a day and (b) fluctuation for the long-term measurement in a month.

introduces significant interference and leads to 5 dB SNR loss for LoRa transmissions over its working time (i.e., between 6 and 24). Figure 19 shows the ESP distribution for different times of the day, ESPs during the day are significantly higher than ESPs at night due to the DTMB interference.

5.3 Link Variation over Channels

Uplink transmissions in our network randomly take place at eight channels, ranging from 475.1 MHz to 476.5 MHz. Therefore, we verify the impact of channels by measuring transmissions over different frequencies. We exclude hardware impacts by using the fixed end node and gateway and avoid DTMB interference by only collecting transmissions between 0 and 6 AM. Figure 20 shows the SNR distribution of transmissions over different channels. *We observe that the measured SNRs are aggregated in the short term and fluctuating in the long term for links at different channels.* Specifically, in the short-term measurement with records collected in a day, we observe that SNRs on each channel are mostly aggregated, but have significant diversity across different channels. This is mainly caused by the multi-path from the rich urban reflections. Multi-path signals over different channels have different wavelengths, leading to constructive superposition at some channels and destructive cancellation at others. On the contrary, for the long-term measurement with records over a month, we observe significant SNR fluctuations in all eight channels. Such fluctuations are mainly due to the change of multipath from surrounding environments.

5.4 Impact of Installation Environment

We now evaluate the impact of installation environments on link performance. We collect information of both the physical environment and the radio environment around end nodes and gateways. The following are our observations:

(1) Installation height of both end nodes and gateways has an impact on the link performance, where nodes near the top floors have much better connections. Figure 21 presents the number of connected gateways for nodes on different floors. Nodes deployed above 20 floors can connect to 10 to 30 gateways on average. This is because links to nodes of high floors travel through fewer occlusions, such as the concrete and ground, and thus suffer less propagation attenuation. We also observe nodes in lower floors (i.e., under 20) all behave similarly well, each with connections to 5 to 10 gateways. This violets our conventional wisdom for radio system deployment recommends that the higher the node deployment, the better the communication performance. The reason for this situation is that buildings in the measurement region are mostly over 20 stories, leading to similar occlusion conditions for nodes on lower floors. As an exception, nodes on the first floor perform poorly as they suffer more interference from the ground reflection. We correspond the number of reached gateways to the packet loss rate in Figure 22, which indicates that top-floor nodes that can reach more gateways usually have lower link packet loss rates.

(2) Interference on ISM bands presents frequency and spatial aggregation. Figure 23 shows measurements of radio environments over 100

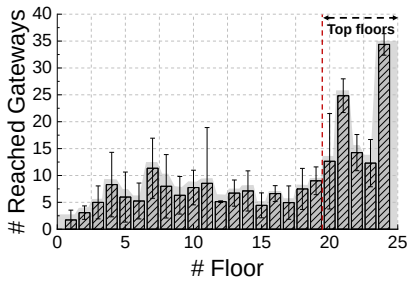

Figure 21: Number of connected gateways at different floors.

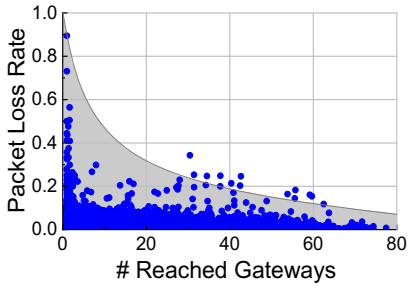

Figure 22: Packet loss rates versus number of reached gateways.

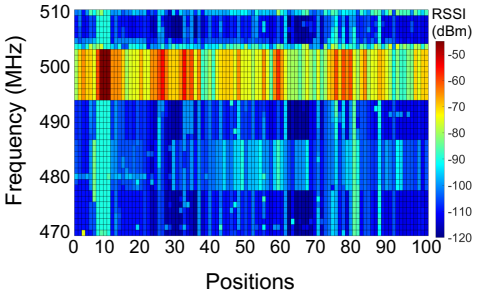

Figure 23: Spectrum measurement over different frequencies and positions.

locations in the network region. We perform the measurement by suspending LoRa transmissions and using a spectrograph for scanning the spectrum. The results show strong interference between 494 and 503 MHz, which is mainly due to the DTMB signal as we discussed in Sec. 5.2. Besides, we observe that some locations suffer persistent interference at all frequencies. This is mainly due to the physical proximity to high-power radio sources whose energy can leak out over a wide spectrum.

5.5 Lessons and Implications

(i) LoRa links vary dramatically across different channels in multipath-rich urban environments. Optimal channel selection is a crucial issue that should be studied for LoRa transmission optimization. Communication over the optimal radio channel can effectively improve the link reliability, supporting higher data rate and battery life of end nodes. Early research studies optimal channel prediction based on temporal and spatial correlation with long-term statistics of channel occupancy and signal power [28–31]. Recent works disentangle physical links with multiple receivers and ascertain the optimal channel accordingly [10]. All existing works make trade-offs in time, costs, and prediction accuracy. How to achieve real-time optimal channel prediction is still an open problem.

(ii) Interference on ISM bands significantly affects LoRa transmissions. Spectrum sensing mechanisms should be studied to avoid potential interference in LoRa. Cognitive radio and spectrum sensing are primarily aimed at identifying vacant frequency bands to minimize interference. However, existing mechanisms are all energy intensive and are unsuitable for LoRa. Future work should focus on low-power spectrum sensing to improve LoRa performance in interfering bands.

6 ENERGY AND SYSTEM EXPENDITURE

6.1 Energy Consumption

Energy consumption of end nodes is a first-order concern of smart city applications, where many devices are powered by batteries. We measure the power profile of end nodes in our network under various data transmission demands and compare LoRa with several other LPWAN technologies, i.e., Sigfox, licensed NB-IoT, and LTE-M, in terms of the battery lifetime, to provide a sense of the power tradeoff between these technologies. We measure the current profile of a LoRa node and plot the currents for all radio access phases in

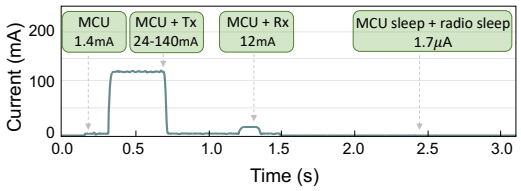

Figure 24: Current consumption profile for a complete LoRa transmission. The device is powered at 3.6 V.

Figure 24. The node stays in sleep mode when it is not transmitting or receiving. The radio transmission consumes the highest amount of energy by a large margin. Thus, any optimization of battery life must focus on reducing the energy of transmissions. In particular, saving a small amount of packet energy can translate to tens or hundreds of days of prolonged node life.

Based on the current measurement, we calculate the energy consumption of a node by summing up the energy of all radio access phases. Then, we estimate the battery lifetime of the node by assuming a typical battery capacity of 5000 mAh. Figure 25 shows the estimated battery lifetime of a LoRa node under various data rates and uplink transmission demands. The lifetime decreases as the node transmits more uplink packets per day or with lower data rates (i.e., higher SF and smaller bandwidth). LoRa decides transmission data rates depending on the coverage level of the node, where poorly connected nodes tend to use higher SF and smaller bandwidth, suffering significantly shorter lifetimes.

We then compare the lifetime of LoRa with other LPWAN technologies. We calculate the energy consumption of three other technologies based on the estimation in the prior literature [16]. Table 2 shows the estimated battery life of four technologies with transmission rates from 90 bytes every day to 256 bytes every hour. We estimate the lifetime of NB-IoT and LoRa under two different data rates and label the corresponding link budget of each set in the table. *We observe that LoRa has the best power efficiency for all applications.* Sigfox has a very low data rate, and thus it must fragment payloads to multiple packets, resulting in low energy efficiency. Cellular-based NB-IoT and LTE-M protocols are also power-intensive due to their complicated physical layers and access control mechanisms. LoRa has the best energy efficiency due to the low protocol overhead of LoRa transmissions. *However, the advantage of LoRa decreases in high data rate applications.* When

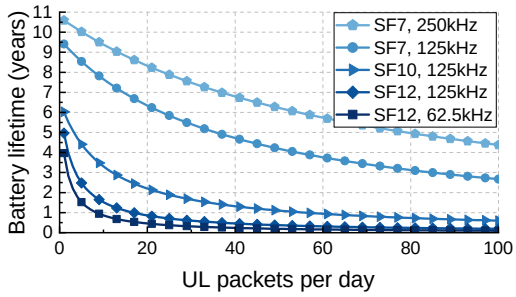

Figure 25: Node battery lifetime with different data rates and uplink transmission demands.

Table 2: Lifetime estimation for different network technologies across various application demands.

Network Technology	End Node Lifetime (years)			
	Heat Meter 90 Bytes Per Day	Gas Meter 256 Bytes Per Day	Hygrometer 64 Bytes Per Hour	Regulator 256 Bytes Per Hour
LoRa (143 dB)	9.5	6.7	3.8	1.6
NB-IoT (144 dB)	5.3	3.5	1.4	1.1
Sigfox (155 dB)	3.4	2.3	0.9	0.6
LoRa (157 dB)	4.5	2.9	1.7	0.8
NB-IoT (164 dB)	2.5	1.7	0.7	0.6
LTE-M (164 dB)	2.2	1.4	0.5	0.5

the application demands bulk data transmissions, LoRa takes a long transmission time due to the limited data rate it can support.

6.2 System Costs

This section identifies the key costs of the LPWAN deployment and answers the question: "What are the advantages and disadvantages of each technology to build a network in different scenarios?" We first investigate the system cost considering the whole lifecycle of IoT applications. As presented in Table 3, we divide the system costs into three categories, i.e., deployment expenditure, operation expenditure, and maintenance expenditure, based on the running phases of IoT applications. We take the cost values of NB-IoT and Sigfox from the previous literature [32] and calculate the battery replacement costs based on the lifetime estimation in Sec. 6.1. We consider an urban area of 300 km^2 and a rural area of 10,000 km^2. For the urban area, we consider high-density deployments with 200 users per km^2, and low density with 5 users per km^2. While for the rural area deployments, we consider 10 users per km^2 as high density and 0.1 users per km^2 as low density. For each scenario, we estimate the number of required sites (i.e., gateways or base stations) based on the coverage and network capacity of previous estimations [33].

The system costs of different scenarios are presented in Figure 26, where we have two observations: *(1) LoRa is cost-effective for the urban scenario and low user density scenes.* LoRa has a low gateway cost, so it has an advantage when applied in urban scenarios where many gateways are needed for combating high signal attenuation.

Table 3: System Expenditures.

	Cost	LoRa	NB-IoT	Sigfox
Deployment Expenditure	User Equipment ($)	4-6	6-12	4
	Site build (K$)	2.1	21	10.5
	Site lease (K$/year)	0.4-1.1	3.7-8.4	0.9-1.1
Operation Expenditure	Spectrum (K$/kHz/site)	0	0.001	0
	Electricity (K$/year)	0.1	1	1
Maintenance Expenditure	Battery replacement ($/node/year)	0.8	1.4	1.1
	Facilities maintenance (relative to deployment expenditure/year)	20%	10%	15%

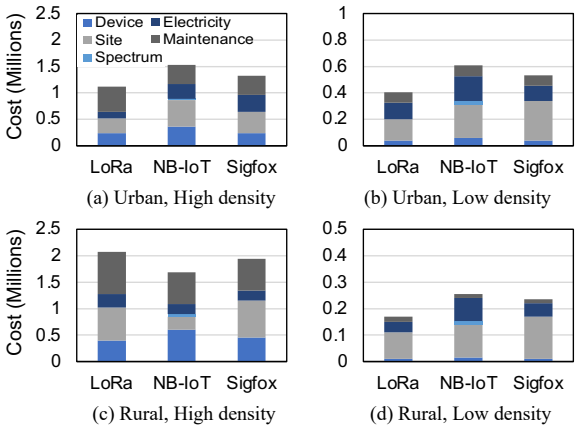

Figure 26: Cost evaluation for three LPWAN technologies under different scenarios and user densities.

At the same time, the low gateway cost also makes LoRa suitable for low device density applications, where the network cost per user device becomes relatively low. *(2) NB-IoT shows premium coverage and capacity performance by using cellular base stations that are widely spread over urban and rural areas.* Cellular-based NB-IoT has the longest coverage range and the highest capacity per site. Therefore, it is suitable to provide connections to dense user devices in open large areas, e.g., rural scenarios. These observations inspire us to choose suitable network technologies to optimize the system performance as well as minimize system costs.

6.3 Lessons and Implications

(i) LoRa data rate plays a critical role in the energy consumption of end nodes, where nodes with frequent high-SF transmissions suffer significantly short lifetimes. Thus, link optimization and adaptive data rate should be studied to prolong the node lifetime. The data rate setting in the current LoRaWAN is threshold based and fails to achieve optimal energy efficiency. Recent research proposes dynamic data rate control based on link decomposition and SNR prediction [10, 34]. However, they require either long-term link estimation or intensive hardware such as multiple gateways. Lightweight and real-time LoRa data rate adaption mechanisms should be studied for optimizing the energy efficiency of end nodes.

(ii) Mixed technologies should be considered to optimize the deployment cost of smart city applications. Our measurements show

that various technologies have advantages and disadvantages in different application scenarios. This inspires us to use mixed technologies for lowering the deployment cost. For example, in an urban scenario, we can use LoRa gateways for connecting dense users in the center area, and use NB-IoT for connecting sparse users in edges or blind spots.

7 RELATED WORK

LoRa measurements. Substantial empirical studies have been devoted to LoRa to understand its performance and limitations. We categorize representative topics as follows: *(i) Link Performance.* Extensive LoRa measurement works, such as [13, 15, 35] and references therein, study the characteristics of LoRa links and verify their performance in terms of communication ranges, capacities, attenuation models, and anti-interference capabilities. These studies focus on a single transmission link between a LoRa node and the gateway. Thus, only a few end nodes and one or two gateways are needed for the evaluation. *(ii) Network Operation.* Other works focus on the functionality of the whole LoRa network with different environments [15, 17, 35], underlying parameters [10, 13], network densities [34, 36], and MAC configurations [14, 16]. These measurements are based on dedicated LoRa deployments with a few gateways and end nodes. Some emulate large-scale LoRa by theoretical models or simulations. The performance of large-scale LoRa networks in real applications is still unclear. *(iii) Energy Consumption.* Many works characterize the energy profiles of LoRa nodes by evaluating the battery life [13, 37–39]. These works focus more on the impact of LoRa modulation settings, such as SF and bandwidth, and lack comparison with other technologies in different scenarios. Other works compare the energy management of LoRa with cellular techniques for physical layer modulation mechanisms, MAC protocols, and state machines [39, 40]. *(iv) Application Measurements.* Many works explore using LoRa to build IoT applications for data collection and instruction delivery. Measurements and evaluations are performed on these LoRa-based IoT applications, including smart agriculture [41], industrial automation [42], supply chain management [43], etc. Recent study evaluates LoRa performance for different application scenarios, which compares wild forest sensing and urban truck tracking applications with both static and mobile end nodes [44]. [32] compares LoRa infrastructure with other wide-area wireless technologies for various application scenarios. The system closest to ours in scale is Helium [45, 46], which presents the decentralized LoRa network architecture with crowdsourced and incentive-guided infrastructure deployment. Helium consists of hundreds of thousands of geographically dispersed end nodes and gateways (named Helium Hotspots) deployed ad-hoc by users all over the world. Studies on Helium [45] focus on hotspot implementation, user incentive mechanisms, and application layer traffic statistics. Few link and network layer characteristics of Helium networks are observed as the nodes and gateways are crowdsourced and the information of underlying layers is inaccessible by network operators.

LoRa optimization. Many recent research efforts propose optimization mechanisms for LoRa networks, including ways to mitigate collisions [6–9, 47–49], improve parameter allocations [10, 34, 50, 51], build up enhanced coding mechanisms [11, 12, 52, 53], and design a variety of channel sensing and estimation solutions [27, 54–57]. Extensive works focus on mitigating LoRa collisions and disentangling interfering transmissions via LoRa physical features such as hardware imperfections [3], temporal signal characteristics [6, 8, 58, 59], energy distributions [47, 60], and frequency domain features [7, 9]. Other works improve LoRa transmissions by selecting optimum parameters (e.g., radio powers [34], spreading factors [50], bandwidth [51], and channel frequencies [10]) or adapting specially designed coding mechanisms, such as distributed and scattered CSS [52, 61]. Finally, many works focus on carrier sensing and link optimization and propose MAC protocols like CSMA for LoRa optimization [26, 27, 54].

Our study differs from the previous works in the following two aspects. First, previous field studies and measurements were conducted with small-scale and dedicated LoRa deployments of only a few end nodes and gateways. We measure the LoRa performance with a citywide large-scale network and uncover its performance challenges in real-world applications. Second, although previous works also study the link and coverage performance of LoRa, their measurements just provide a rough understanding of signal propagation. In contrast, we not only perform a detailed study to understand the LoRa characteristics at different layers but also study the facts that bottleneck LoRa performance. Our experience in deploying citywide networks provides realistic lessons for both application users and network operators in using and deploying LoRa networks.

8 CONCLUSION

This paper presents a full-fledged study on the deployment and operation of a citywide LoRa network, named *CityWAN*. The network is featured by its large scale and real-world usage, which consists of 19,810 end devices and 100 gateways, serving 12 smart city applications. We perform long-term and in-depth measurements in network performance spanning multiple aspects, including application performance, infrastructure deployment, link characteristics, and energy overheads. Based on the fine-grained network information, we reveal some critical issues that hamper the performance of LoRa networks. Our measurements show the LoRa performance in urban settings is bottlenecked by the prevalent blind spots. As for the gateway deployment, there is an efficiency gap in the gateway utilization and network coverage. We also find that LoRa physical links are susceptible to radio and physical environmental variations, and LoRa and other LPWANs show diverse costs and energy efficiencies for different scenarios. Our measurement provides feasible guidance to large-scale LoRa network deployment and also points to directions for academic research to unleash the LoRa potential for future IoT applications.

ACKNOWLEDGEMENTS

We sincerely thank the anonymous reviewers and our shepherd for their valuable feedback. This work is in part supported by National Key R&D Program of China 2022YFC3801300, National Natural Science Foundation of China (U22A2031, No. 61932013, 62172250). Jiliang Wang is the corresponding author.

REFERENCES

[1] Brand Essence Research. Lora and lorawan devices market size, share & trends. 2022.

[2] LoRa Alliance. Lorawan deployments achieve market leadership. 2022.

[3] Rashad Eletreby, Diana Zhang, Swarun Kumar, and Osman Yağan. Empowering low-power wide area networks in urban settings. In *Proceedings of ACM SIGCOMM*, Los Angeles, CA, USA, August 21-25, 2017.

[4] Adwait Dongare, Revathy Narayanan, Akshay Gadre, Anh Luong, Artur Balanuta, Swarun Kumar, Bob Iannucci, and Anthony Rowe. Charm: exploiting geographical diversity through coherent combining in low-power wide-area networks. In *Processings of ACM/IEEE IPSN*, Porto, Portugal, April 11-13, 2018.

[5] Artur Balanuta, Nuno Pereira, Swarun Kumar, and Anthony Rowe. A cloud-optimized link layer for low-power wide-area networks. In *Proceedings of ACM MobiSys*, Toronto, Canada, June 16-19, 2020.

[6] Muhammad Osama Shahid, Millan Philipose, Krishna Chintalapudi, Suman Banerjee, and Bhuvana Krishnaswamy. Concurrent interference cancellation: Decoding multi-packet collisions in lora. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, pages 503–515, 2021.

[7] Shuai Tong, Jiliang Wang, and Yunhao Liu. Combating packet collisions using non-stationary signal scaling in lpwans. In *Proceedings of ACM MobiSys*, Toronto, Canada, June 16-19, 2020.

[8] Xia Xianjin, Zheng Yuanqing, and Gu Tao. Ftrack: Parallel decoding for lora transmissions. In *Proceedings of ACM SenSys*, New York, NY, USA, November 10-13, 2019.

[9] Shuai Tong, Zhenqiang Xu, and Jiliang Wang. Colora: Enabling multi-packet reception in lora. In *Proceedings of IEEE INFOCOM*, Online, July 6-9, 2020.

[10] Akshay Gadre, Revathy Narayanan, Anh Luong, Anthony Rowe, Bob Iannucci, and Swarun Kumar. Frequency configuration for low-power wide-area networks in a heartbeat. In *Proceedings of USENIX NSDI*, Online, Februray 25-27, 2020.

[11] Shuai Tong, Zilin Shen, Yunhao Liu, and Jiliang Wang. Combating link dynamics for reliable lora connection in urban settings. In *Proceedings of ACM MobiCom*, New York, NY, USA, 2021.

[12] Chenning Li, Hanqing Guo, Shuai Tong, Xiao Zeng, Zhichao Cao, Mi Zhang, Qiben Yan, Li Xiao, Jiliang Wang, and Yunhao Liu. Nelora: Towards ultra-low snr lora communication with neural-enhanced demodulation. In *Proceedings of ACM SenSys*, New York, NY, USA, Nov 10-14, 2021.

[13] Jansen C Liando, Amalinda Gamage, Agustinus W Tengourtius, and Mo Li. Known and unknown facts of lora: Experiences from a large-scale measurement study. *ACM Transactions on Sensor Networks*, 15(2):1–35, February 2019.

[14] Yidong Ren, Li Liu, Chenning Li, Zhichao Cao, and Shigang Chen. Is lorawan really wide? fine-grained lora link-level measurement in an urban environment. In *2022 IEEE 30th International Conference on Network Protocols (ICNP)*, pages 1–12. IEEE, 2022.

[15] Weitao Xu, Jun Young Kim, Walter Huang, Salil S Kanhere, Sanjay K Jha, and Wen Hu. Measurement, characterization, and modeling of lora technology in multifloor buildings. *IEEE Internet of Things Journal*, 7(1):298–310, 2019.

[16] Ghena Branden, Adkins Joshua, Shangguan Longfei, Jamieson Kyle, Levis Phil, and Dutta Prabal. Challenge: Unlicensed lpwans are not yet the path to ubiquitous connectivity. In *Proceedings of ACM Mobicom*, Los Cabos, Mexico, October 21-25, 2019.

[17] Michael Rademacher, Hendrik Linka, Thorsten Horstmann, and Martin Henze. Path loss in urban lora networks: A large-scale measurement study. In *2021 IEEE 94th Vehicular Technology Conference (VTC2021-Fall)*, pages 1–6. IEEE, 2021.

[18] Zhenqiang Xu, Shuai Tong, Pengjin Xie, and Jiliang Wang. From demodulation to decoding: Toward complete lora phy understanding and implementation. *ACM Transactions on Sensor Networks*, 18(4):1–27, 2023.

[19] Semtch. Sx1301 datasheet. *Available: https://www.semtech.com/*.

[20] Semtch. Sx1278/77/78/79 datasheet. *Available: https://www.semtech.com/*.

[21] Ahmed Abdelghany, Bernard Uguen, Christophe Moy, and Dominique Lemur. On superior reliability of effective signal power versus rssi in lorawan. In *2021 28th International Conference on Telecommunications (ICT)*, pages 1–5. IEEE, 2021.

[22] Zhidan Liu, Jiancong Liu, Xiaowen Xu, and Kaishun Wu. Deepgps: Deep learning enhanced gps positioning in urban canyons. *IEEE Transactions on Mobile Computing*, 2022.

[23] P. Misra and P. Enge. *Global Positioning System: Signals, Measurements, and Performance*. Ganga-Jamuna Press, 2011.

[24] Xia Xianjin, Chen Qianwu, Hou Ningning, Zheng Yuanqing, and Li Mo. Xcopy: Boosting weak links for reliable lora communication. In *Proceedings of the ACM MobiCom*, 2023.

[25] Jing Yang, Zhenqiang Xu, and Jiliang Wang. Ferrylink: Combating link degradation for practical lpwan deployments. In *Proceedings of IEEE ICPADS*, 2021.

[26] Raghav Subbaraman, Yeswanth Guntupalli, Shruti Jain, Rohit Kumar, Krishna Chintalapudi, and Dinesh Bharadia. Bsma: scalable lora networks using full duplex gateways. In *Proceedings of the 28th Annual International Conference on Mobile Computing And Networking*, pages 676–689, 2022.

[27] Amalinda Gamage, Jansen Christian Liando, Chaojie Gu, Rui Tan, and Mo Li. Lmac: Efficient carrier-sense multiple access for lora. In *Proceedings of ACM MobiCom*, Online, September 21-25, 2020.

[28] A Canavitsas, LAR Silva Mello, and M Grivet. White space prediction technique for cognitive radio applications. In *2013 SBMO/IEEE MTT-S International Microwave & Optoelectronics Conference (IMOC)*, pages 1–5. IEEE, 2013.

[29] Hao Chen, Lingjia Liu, Thomas Novlan, John D Matyjas, Boon Loong Ng, and Jianzhong Zhang. Spatial spectrum sensing-based device-to-device cellular networks. *IEEE Transactions on Wireless Communications*, 15(11):7299–7313, 2016.

[30] Wenchi Cheng, Xi Zhang, and Hailin Zhang. Full-duplex spectrum-sensing and mac-protocol for multichannel nontime-slotted cognitive radio networks. *IEEE Journal on Selected Areas in Communications*, 33(5):820–831, 2014.

[31] Junil Choi, David J Love, and Patrick Bidigare. Downlink training techniques for fdd massive mimo systems: Open-loop and closed-loop training with memory. *IEEE Journal of Selected Topics in Signal Processing*, 8(5):802–814, 2014.

[32] Mohammad Istiak Hossain and I. Markendahl. Jan. Comparison of lpwan technologies: Cost structure and scalability. *Wireless Personal Communications*, 121(1):887–903, 2021.

[33] Mads Lauridsen, Istvan Z. Kovacs, Preben Mogensen, Mads Sorensen, and Steffen Holst. Coverage and capacity analysis of lte-m and nb-iot in a rural area. In *Proceedings of IEEE VTC-Fall*, 2016.

[34] Yinghui Li, Jing Yang, and Jiliang Wang Wang. Dylora: Towards energy efficient dynamic lora transmission control. In *Proceedings of IEEE INFOCOM*, Online, July 6-9, 2020.

[35] Juha Petajajarvi, Konstantin Mikhaylov, Antti Roivainen, Tuomo Hanninen, and Marko Pettissalo. On the coverage of lpwans: range evaluation and channel attenuation model for lora technology. In *2015 14th international conference on its telecommunications (itst)*, pages 55–59. IEEE, 2015.

[36] Ferran Adelantado, Xavier Vilajosana, Pere Tuset-Peiro, Borja Martinez, Joan Melia-Segui, and Thomas Watteyne. Understanding the limits of lorawan. *IEEE Communications magazine*, 55(9):34–40, 2017.

[37] Sezana Fahmida, Venkata P Modekurthy, Mahbubur Rahman, Abusayeed Saifullah, and Marco Brocanelli. Long-lived lora: Prolonging the lifetime of a lora network. In *2020 IEEE 28th International Conference on Network Protocols (ICNP)*, pages 1–12. IEEE, 2020.

[38] Taoufik Bouguera, Jean-François Diouris, Jean-Jacques Chaillout, Randa Jaouadi, and Guillaume Andrieux. Energy consumption model for sensor nodes based on lora and lorawan. *Sensors*, 18(7):2104, 2018.

[39] Ritesh Kumar Singh, Priyesh Pappinisseri Puluckul, Rafael Berkvens, and Maarten Weyn. Energy consumption analysis of lpwan technologies and lifetime estimation for iot application. *Sensors*, 20(17):4794, 2020.

[40] 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, 2019.

[41] Ayesha Siddique, Bhakti Prabhu, Aishwarya Chaskar, and Rasika Pathak. A review on intelligent agriculture service platform with lora based wireless sensor network. *Life*, 100:7000, 2019.

[42] Mattia Rizzi, Paolo Ferrari, Alessandra Flammini, Emiliano Sisinni, and Mikael Gidlund. Using lora for industrial wireless networks. In *2017 IEEE 13th international workshop on factory communication systems (WFCS)*, pages 1–4. IEEE, 2017.

[43] Aristeidis Karras, Christos Karras, Georgios Drakopoulos, Dimitrios Tsolis, Phivos Mylonas, and Spyros Sioutas. Saf: a peer to peer iot lora system for smart supply chain in agriculture. In *Artificial Intelligence Applications and Innovations: 18th IFIP WG 12.5 International Conference, AIAI 2022*, pages 41–50. Springer, 2022.

[44] Ana Elisa Ferreira, Fernando M Ortiz, Luis Henrique MK Costa, Brandon Foubert, Ibrahim Amadou, and Nathalie Mitton. A study of the lora signal propagation in forest, urban, and suburban environments. *Annals of Telecommunications*, 75:333–351, 2020.

[45] 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.

[46] Helium. Helium engineering blog. app version 3.2.0. In *https://engineering.helium.com/ 2021/05/18/app-version-320.html, Accessed May*, 2021.

[47] Zhenqiang Xu, Pengjin Xie, and Jiliang Wang. Pyramid: Real-time lora collision decoding with peak tracking. In *Proceedings of IEEE INFOCOM*, Online, May 10-13, 2021.

[48] Zhenqiang Xu, Shuai Tong, Pengjin Xie, and Jiliang Wang. Fliplora: Resolving collisions with up-down quasi-orthogonality. In *Proceedings of IEEE SECON*, Online, June 22-25, 2020.

[49] Qian Chen and Jiliang Wang. Aligntrack: Push the limit of lora collision decoding. In *2021 IEEE 29th International Conference on Network Protocols (ICNP)*, pages 1–11. IEEE, 2021.

[50] Weifeng Gao, Zhiwei Zhao, and Geyong Min. Adaplora: Resource adaptation for maximizing network lifetime in lora networks. In *Proceedings of IEEE ICNP*, Online, October 13-16, 2020.

[51] Liu Li, Yao Yuguang, Cao Zhichao, and Zhang Mi. Deeplora: Learning accurate path loss model for long distance links in lpwan. In *Proceedings of IEEE INFOCOM*,

Online, May 10-13, 2021.

[52] Mehrdad Hessar, Ali Najafi, and Shyamnath Gollakota. Netscatter: Enabling large-scale backscatter networks. In *Proceedings of USENIX NSDI*, Boston, MA, USA, February 26-28, 2019.

[53] Shuai Tong, Yangliang He, Yunhao Liu, and Jiliang Wang. De-spreading over the air: long-range ctc for diverse receivers with lora. In *Proceedings of the 28th Annual International Conference on Mobile Computing And Networking*, pages 42–54, 2022.

[54] Justin Chan, Anran Wang, Arvind Krishnamurthy, and Shyamnath Gollakota. Deepsense: Enabling carrier sense in low-power wide area networks using deep learning. In *ArXiv*, 2019.

[55] Xiong Wang, Linghe Kong, Zucheng Wu, Long Cheng, Chenren Xu, and Guihai Chen. Slora: towards secure lora communications with fine-grained physical layer features. In *Proceedings of ACM SenSys*, Online, November 16-19, 2020.

[56] Silvia Demetri, Marco Zúñiga, Gian Pietro Picco, Fernando Kuipers, Lorenzo Bruzzone, and Thomas Telkamp. Automated estimation of link quality for lora: A remote sensing approach. In *Proceedings of IEEE IPSN*, Montreal, Canada, April 16-18, 2019.

[57] Yuxiang Lin, Wei Dong, Yi Gao, and Tao Gu. Sateloc: A virtual fingerprinting approach to outdoor lora localization using satellite images. In *Proceedings of ACM/IEEE IPSN*, Online, April 21-24, 2020.

[58] Xiong Wang, Linghe Kong, Liang He, and Guihai Chen. mlora: A multi-packet reception protocol for lora communications. In *Proceedings of IEEE ICNP*, Chicago, Illinois, USA, October 7-10, 2019.

[59] Zhe Wang, Linghe Kong, Kangjie Xu, Liang He, Kaishun Wu, and Guihai Chen. Online concurrent transmissions at lora gateway. In *Proceedings of IEEE INFOCOM*, Online, July 6-9, 2020.

[60] Bin Hu, Zhimeng Yin, Shuai Wang, Zhuqing Xu, and Tian He. Sclora: Leveraging multi-dimensionality in decoding collided lora transmissions. In *Proceedings of IEEE ICNP*, Online, October 13-16, 2020.

[61] Xie Pengjin, Li Yinghui, Xu Zhenqiang, Chen Qian, Liu Yunhao, and Wang Jiliang. Push the limit of lpwans with concurrent transmissions. In *In Proceedings of IEEE INFOCOM*, 2023.