Exploiting LoRaWAN for efficient and resilient IoT networks

Gu, Chaojie

2020

Gu, C. (2020). Exploiting LoRaWAN for efficient and resilient IoT networks. Doctoral thesis, Nanyang Technological University, Singapore.

https://hdl.handle.net/10356/143912

https://doi.org/10.32657/10356/143912

NANYANG TECHNOLOGICAL UNIVERSITY
SINGAPORE

EXPLOITING LORAWAN FOR EFFICIENT AND RESILIENT IOT NETWORKS

GU CHAOJIE

School of Computer Science and Engineering

2020

EXPLOITING LORAWAN FOR EFFICIENT AND RESILIENT IOT NETWORKS

GU CHAOJIE

School of Computer Science and Engineering

A thesis submitted to the Nanyang Technological University
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

2020

Statement of Originality

I hereby certify that the work embodied in this thesis is the result of original research, is free of plagiarised materials, and has not been submitted for a higher degree to any other University or Institution.

September 30, 2020 Chaojie Gu.

. . . .

Date GU Chaojie

Supervisor Declaration Statement

I have reviewed the content and presentation style of this thesis and declare it is free of plagiarism and of sufficient grammatical clarity to be examined. To the best of my knowledge, the research and writing are those of the candidate except as acknowledged in the Author Attribution Statement. I confirm that the investigations were conducted in accord with the ethics policies and integrity standards of Nanyang Technological University and that the research data are presented honestly and without prejudice.

	TAN. Rui
September 30, 2020	
....	....
Date	Asst. Prof. Rui Tan

Authorship Attribution Statement

This thesis contains material from one paper published in a peer-reviewed journal and three papers accepted at workshop/conferences in which I am listed as an author.

Chapter 2 contains material from a workshop paper:

- Chaojie Gu, Linshan Jiang, and Rui Tan. “LoRa-Based Localization: Opportunities and Challenges”, in The 1st workshop on Low Power Wide Area Networks for Internet of Things (LPNET) with EWSN 2019, February 25, 2019, Beijing, China. DOI: 10.5555/3324320.3324420.

The contributions of the co-authors are as follows:

- GU Chaojie co-designed the approach, collected and analyzed the data, and prepared the paper drafts.
- JIANG Linshan co-designed the approach and provided theoretic analysis.
- TAN Rui provided the project direction, co-designed the approach, and revised the paper drafts.

Chapter 3 contains material from a conference paper and its extended journal version:

- Chaojie Gu, Rui Tan, Xin Lou, and Dusit Niyato. One-Hop Out-of-Band Control Planes for Low-Power Multi-Hop Wireless Networks. In Proceedings of the 37th Annual IEEE International Conference on Computer Communications (INFOCOM'18), April 15 - 19, 2018, Honolulu, HI. DOI: 10.1109/INFOCOM.2018.8486301.

- Chaojie Gu, Rui Tan, and Xin Lou. One-Hop Out-of-Band Control Planes for Multi-Hop Wireless Sensor Networks. ACM Transactions on Sensor Networks, 2019. Article 40. DOI: 10.1145/3342100

The contributions of the co-authors are as follows:

- GU Chaojie co-designed the approach, implemented and evaluated the approach, analyzed the data, and prepared the paper drafts.
- TAN Rui provided the project direction, co-designed the approach, and revised the paper drafts.
- LOU Xin participated in discussions and helped with the literature review.

- NIYATO Dusit revised the paper drafts.

Chapter 4 contains material from a conference paper:

- Chaojie Gu, Linshan Jiang, Rui Tan, Mo Li, and Jun Huang. Attack-Aware Data Timestamping in Low-Power Synchronization-Free LoRaWAN. In Proceedings of the 40th IEEE International Conference on Distributed Computing Systems (ICDCS'20), Nov-Dec, 2020, Singapore.

The contributions of the co-authors are as follows:

- GU Chaojie co-designed the approach, implemented and evaluated the approach, conducted the experiments, analyzed the data, and prepared the paper drafts.
- JIANG Linshan provided theoretic analysis on the jamming model, and conducted the experiments.
- TAN Rui provided the project direction, co-designed the approach, and revised the paper drafts.
- LI Mo participated in discussions, contributed to the experiments, and revised the paper drafts.
- HUANG Jun participated in discussions and revised the paper drafts.

September 30, 2020 Chaojie Gu.

. . . .

Date GU Chaojie

Acknowledgements

First and foremost I would like to thank Prof. Rui Tan. As my advisor, he gives me great support when I am pursuing Ph.D. degree. For years, Prof. Rui Tan sets an example of critical thinking and diligent working. He teaches me how to cultivate a positive and optimistic approach to research. It is my great fortune to be his student.

I want to express my special thanks of gratitude to my family for the constant support.

I am greatly thankful to Prof. Mo Li, Prof. Dusit Niyato and Prof. Jun Huang for their contributions to my research. I would like to thank my Thesis Advisory Committee members and Oral Defence Panel members, Prof. Arvind Easwaran, Prof. Lilian Wang, Prof. A S Madhukumar, Prof. Lee Bu Sung, Prof. Yeo Chai Kiat. Much appreciation for their valuable comments and feedback.

I am also sincerely thankful to my friends, Haoqing Zhu, Guoxin Huang, Muhan Liu, Hongyi Luo, Zhaohui Hou, Kaijing Luo, Yusheng Ye, and Xiang Ni, for their friendship and emotional support. I am very thankful to my friend, Shuyang Sun, for providing me valuable suggestions on Ph.D. application and research. I am greatly thankful for my friends and colleagues in Computer Networks and Communications Lab (CNCL) for their support and encouragement. It is my honor to work with them.

To my dear family

Abstract

It is estimated that, by 2025, there will be more than 21 billion Internet of Things (IoT) devices deployed in various domains. These massive IoT devices will be interconnected by numerous IoT networks with the Internet as the backbone. The IoT networks will be primarily wireless, ranging from cellular networks, Wi-Fi infrastructures, low-power multi-hop wireless networks (e.g., Zigbee and Bluetooth personal area networks), and the recently emerging low-power wide-area networks. The greatly increased pervasive connectivity owing to the deployment of these IoT networks will foster the next-generation Internet-based innovations. This thesis focuses on exploiting LoRaWAN, a representative low-power wide-area networking technology, to build efficient and resilient low-power wireless IoT networks.

Given the increasingly crowded radio frequency (RF) spectrum, the efficiency of utilizing the finite wireless bandwidth is a primary goal of designing and operating IoT networks. Moreover, the networks' resilience, i.e., their abilities to recover and maintain connectivity and efficiency despite external disturbances such as interference from neighbor RF technologies and even cyber-attacks, is also important to the IoT applications. This thesis aims at studying how the low-power long-range communication capability of LoRaWAN can be exploited to address some of the efficiency and resilience issues in IoT networks. This thesis studies the following two main problems. First, it studies how to use the one-hop LoRaWAN to build out-of-band control planes for the low-power multi-hop wireless networks to improve their efficiency and resilience. Specifically, it exploits the simplicity of LoRaWAN to manage the complexity of multi-hop wireless networks. Second, given LoRaWAN's communication throughput limitation due to the narrow bandwidth and low duty cycle defined in LoRaWAN specification, this thesis studies how to efficiently maintain the common notion of time among all LoRaWAN end devices. In addition, it investigates the potential attacks that aim at disrupting the common notion of time and develops countermeasures for resilience. The details of the two main problems and this thesis' solutions are as follows.

The first part of this thesis addresses the Separation of Control and Data Planes (SCDP) for low-power multi-hop wireless networks using LoRaWAN. SCDP is a desirable paradigm for low-power multi-hop wireless networks requiring high network performance and manageability. Existing SCDP networks generally adopt an *in-band* control plane scheme in that the control-plane messages are delivered by their data-plane networks. The physical coupling of the two planes may lead to undesirable consequences. For example, when a node loses connections with its neighbors, the controller cannot reach the node anymore. Recently, multi-radio platforms (e.g., TI CC1350 and OpenMote B) are increasingly available, which make the physical SCDP possible. To advance the network architecture design, this thesis leverages on the LoRaWAN to form one-hop out-of-band control planes called LoRaCP. Several characteristics of LoRaWAN such as downlink-uplink asymmetry and primitive ALOHA media access control need to be dealt with to achieve high efficiency and good resilience. To address these challenges, a TDMA-based multi-channel transmission control is designed, which features an urgent channel and negative acknowledgment. On a testbed of 16 nodes, LoRaCP is applied to physically separate the control-plane network of the Collection Tree Protocol (CTP) from its Zigbee-based data-plane network. Extensive experiments show that LoRaCP increases CTP's packet delivery ratio from 65% to 80% in the presence of external interference, while consuming little per-node average radio power.

LoRaWAN is promising for collecting low-rate monitoring data from geographically distributed sensors, in which timestamping the sensor data with a common notion of time is a critical system function. The second part of this thesis considers a synchronization-free approach to timestamping LoRaWAN uplink data based on signal arrival time at the gateway, which well matches LoRaWAN's one-hop star topology and releases bandwidth from transmitting timestamps and synchronizing end devices' clocks at all times. However, this thesis shows that this approach is susceptible to a *frame delay attack* consisting of malicious frame collision and delayed replay. In the attack, the attacker records the signal sent by the transmitter and sends a colliding frame to jam the receiver. Then, it replays the recorded signal after an intended delay. Real experiments show that the attack can affect the end devices in large areas up to about $50,000$ square meters. In a broader sense, the attack threatens any system functions requiring timely deliveries of LoRaWAN frames. To address this threat, this thesis proposes a LoRa TimeStamping (LoRaTS) gateway design that integrates a commodity LoRaWAN gateway and a listen-only low-power software-defined radio to track the inherent frequency biases of the end devices. Based on an analytic model of LoRa's chirp spread spectrum modulation, this thesis develops signal processing algorithms to estimate the frequency biases with high accuracy beyond that achieved by LoRa's default demodulation. The accurate frequency bias tracking capability enables the detection of the attack that introduces additional frequency biases. Our approach supports the bandwidth-efficient sync-free time stamping and requires no modifications on the LoRaWAN end devices. Extensive real-world experiments based on a testbed deployed in a university campus show the effectiveness of the proposed approach.

The availability of multiple RF technologies and the mixed use of them in IoT networks create both opportunities and also challenges. The approach designs presented in this thesis demonstrate the exploitation of the new unique features of LoRaWAN in addressing the efficiency and resilience issues of the legacy Zigbee networks and LoRaWAN networks themselves. The author's future work will also follow the same methodology to improve IoT networks.

Contents

List of Figures

List of Tables

List of Acronyms

ADCs	Analog-to-digital Converters
AI	Artificial Intelligence
AIC	Akaike Information Criterion
BLE	Bluetooth Low Energy
CORR	Correlation
COTS	Commercial Off-The-Shelf
CSS	Chirp Spread Spectrum
CTP	Collection Tree Protocol
ETX	Expected Transmission Count
DPR	Data rate to Power consumption Ratio
DSSS	Direct-Sequence Spread Spectrum
ENV	Envelope
FBs	Frequency Biases
FDR	Frame Delivery Ratio
FFT	Fast Fourier Transform
IoT	Internet of Things
LoRaTS	LoRa TimeStamping
LoRaCP	Long-Range Control plane
LPWANs	Low-Power Wide-Area Networks
LTE	Long-Term Evolution
MAC	Media Access Control
MCUs	microcontrollers
MHWNs	Multi-Hop Wireless Networks
MIMO	Multi-Input Multi-Output
MRUP	Multi-Radio Unification Protocol

NAK	Negative Acknowledgment
NTP	Network Time Protocol
PDRs	Packet Delivery Ratios
ppm	parts-per-million
PSDs	Power Spectral Densities
RAT	Radio Timer
RETX	Residual ETX
RF	Radio Frequency
RMSD	Root-Mean-Square Deviation
RPi	Raspberry Pi
RSSI	Received Signal Strength Indicator
RTM	Relative Time Misalignment
SCDP	Separation of Control and Data Planes
SCR	Signal-to-Collision Ratio
SDN	Software-Defined Networking
SDR	Software-Defined Radio
SF	Spreading Factor
SFD	Start Frame Delimiter
SNR	Signal-to-Noise Ratio
SoC	System-on-a-Chip
SSIDs	Service Set Identifiers
TDMA	Time-Division Multiple Access
TCXO	Temperature Compensated Crystal Oscillator
TI	Texas Instruments
TI-RTOS	TI's Real-Time Operating System
WSNs	Wireless Sensor Networks

Chapter 1

Introduction

1.1 Background

Internet of Things (IoT) will be an important global infrastructure. It is estimated that, by 2025, the IoT will have more than 21 billion connected devices [1]. The IoT with pervasive connectivity will support a wide spectrum of applications in various smart systems, such as smart homes, intelligent transportation systems, smart cities, smart agricultures, etc. The IoT infrastructure will also enable the innovations for other unseen applications.

The IoT networks will be primarily wireless, ranging from cellular networks, Wi-Fi infrastructures, low-power multi-hop wireless networks (e.g., Zigbee and Bluetooth personal area networks), and the recently emerging low-power wide-area networks. Wireless IoT networks allow ad-hoc and easy deployment of IoT objects. These wireless technologies vary in communication range, data rate, and power consumption. Depending on the design objectives of the applications and the practical constraints that they face, the IoT system designers can choose one or more wireless technologies to build the IoT networks supporting the many IoT devices that are embedded in various physical environments.

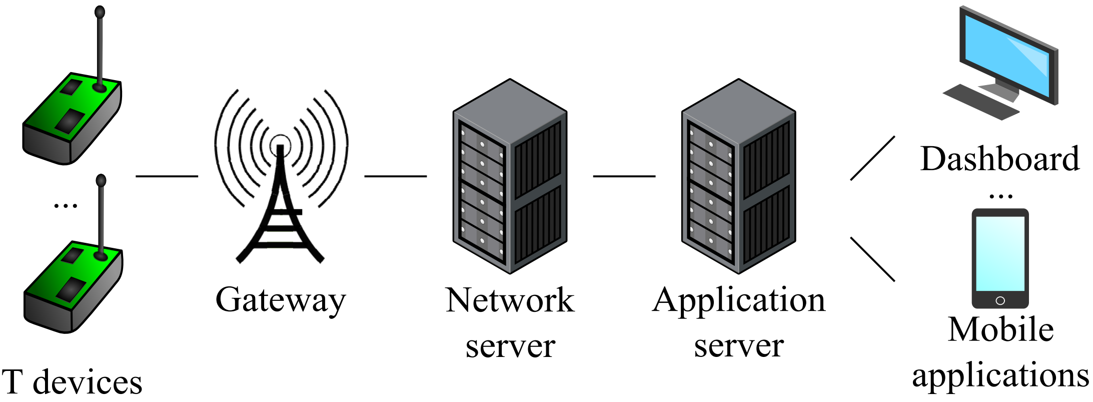

FIGURE 1.1: A typical LoRaWAN network deployment.

1.2 Motivation

Given the increasingly crowded radio frequency (RF) spectrum, the efficiency of utilizing the finite wireless bandwidth is a primary goal of designing and operating IoT networks. IoT networks with better efficiency will allow supporting more IoT devices within the same collision domain. It can reduce the retransmissions so that the lifetime of battery-based IoT devices will be extended. With the efficient IoT networks, artificial intelligence (AI)-empowered IoT (a.k.a., AIoT) can adopt the remote inference scheme [2] to offload more data to the edge/cloud subject to the real-time requirements. Moreover, the networks' resilience, i.e., their abilities to recover and maintain connectivity and efficiency despite external disturbances such as interference from neighbor RF technologies and even cyber-attacks, is also important to the IoT applications. Comparing with wireline networks, wireless networks face much more external disturbances such as wireless signal propagation blockage, interferences from neighboring RF technologies, movement of end devices, etc. Moreover, due to the broadcast nature of wireless communications, wireless IoT networks are susceptible to wireless attacks. In particular, for the wireless networks adopting the multi-hop scheme, loss of critical links may lead to isolated sub-networks and malfunction of the applications running on top of the network. Therefore, improving the network the networks' ability to increase wireless networks survive external interference and even cyber-attacks is a critical research issue.

Low-Power Wide-Area Networks (LPWANs) enable direct wireless interconnections among end devices and gateways in geographic areas of square kilometers. It increases network connectivity as a defining characteristic of the IoT. Among various LPWAN technologies (including NB-IoT and Sigfox), LoRaWAN, which is an open data link layer specification based on the LoRa modulation scheme [3], offers the advantages of using license-free ISM bands, low costs for end devices, and independence from managed cellular infrastructures. LoRaWAN can cover a wide area with a single gateway to form a single-hop topology that simplifies the organization of the IoT network. Figure 1.1 presents a typical LoRaWAN network deployment. A LoRaWAN gateway receives LoRa modulated RF messages from IoT devices and forwards data to the LoRaWAN network server. The network server authenticates the identity of every sensor on the network and the integrity of every message. The application server is responsible for securely handling and interpreting sensor application data. Users can access the applications vis dashboard or mobile applications. This thesis proposes to exploit LoRaWAN to improve the efficiency and resilience of IoT networks. Specifically, it studies two issues as follows.

First, this thesis exploits LoRaWAN's simplicity of one-hop network topology to improve the efficiency and resilience of the low-power multi-hop wireless networks. Centralized control network has been applied in many systems to impose a high network performance. All these centralized network control systems adopt an in-band control scheme, in which the control plane and data plane share the same data communication network. However, this scheme introduces undesirable coupling between the two planes. Loss in critical links will compromise the network performance significantly. The first part of this thesis proposes to apply LoRaWAN to build the out-of-band control plane for low-power multi-hop wireless networks, in which the control plane and data plane are two physically separated communication networks. In this way, the failures in the data-plane network will not affect the control plane, and thus the controller can recover the data-plane network in time. However, the multi-hop data-plane network is more prone to errors due to its fragility, whereas the single-hop control-plane network can deal with errors more easily. Thus, this thesis does not specifically analyze the impact of the control-plane network failures on the system.

Second, this thesis aims at developing an energy-efficient and resilient approach for LoRaWAN to achieve common notion of time in timestamping the data generated by IoT end devices. Maintaining tight clock synchronization of LoRaWAN end devices for data timestamping is costly due to limited bandwidth. The second part of this thesis shows that timestamping based on uplink frame arrival time at the gateway is efficient but can be insecure. The thesis develops and implements an algorithm to track frequency biases to make the system resilient to the delay attack while remains energy-efficient due to the avoidance of costly clock synchronization. This thesis is the first to use frequency bias tracking to achieve a good trade-off between energy efficiency and security.

The following two sections of this chapter introduce the details of the above two problems.

1.3 One-hop Out-of-band Control Planes for Multi-hop Wireless Networks

Many networks will follow the paradigm of multi-hop wireless networks. For instance, wireless meshes are increasingly adopted to interconnect surveillance cameras [4] and vehicles [5]. Wireless sensors have been widely deployed for sensing and control of building environment and energy use. Bluetooth Low Energy (BLE) supports mesh networking [6]. Wireless connectivity is also critical to the vision of Industry 4.0. Utility and manufacturing systems are increasingly adopting wireless metering and monitoring [7].

The main advantage of Multi-Hop Wireless Networks (MHWNs) is that, during the deployment phase, a network can easily scale up to cover a large geographic area. A primary design principle for MHWNs is the use of distributed protocols (e.g., routing [8]), where each node independently performs various networking functions (e.g., data forwarding) based on local information. Thus, the *control plane* (i.e., determination of how to handle packets) and the *data plane* (i.e., carrying out control-plane decisions) of these distributed protocols are jointly implemented at each network node. However, a distributed scheme without the global view often yields suboptimal performance. Moreover, although the distributed scheme may work satisfactorily most of the time thanks to a decade of research, it is often complex and difficult to manage once the network is deployed.

To improve the network performance and manageability, some MHWNs, especially those deployed for mission-critical tasks, have adopted centralized network controls.

For instance, WirelessHART, an MHWN standard that has been adopted in over 53,501 manufacturing systems [7], prescribes centralized routing control based on a global view of the network. It thus better achieves certain performance objectives (e.g., firm/soft real-time packet delivery). Similarly, ISA100.11a, another industry-oriented MHWN standard, also adopts centralized routing control and network management. For the routing in these MHWNs, the control plane is separated from the data plane, in that the routing control is implemented at a centralized node whereas other network nodes follow the routing schedule to forward data packets. However, all these MHWNs adopt *in-band* control planes, i.e., the control-plane messages such as network status reports and routing schedules are delivered by the data-plane networks.

The physical coupling between the control and data planes in the in-band scheme may lead to undesirable consequences. The wireless data-plane network is susceptible to external interference. Deteriorated data-plane links may lead to delayed deliveries or even losses of the control-plane messages, making the network less responsive to data-plane link quality variations. Moreover, when the data plane loses key routing nodes (e.g., due to node hardware/software fault and depletion of battery) or the control plane makes wrong control decisions (e.g., due to design defects or erroneous human operations), the data-plane network may fall apart to disconnected partitions. As a result, restorative network control commands in the control plane may not be able to reach the destination nodes. Recent research has studied protecting the control plane from data-plane faults [9]. However, the solution has limited protection capability against a single link failure only [9].

In light of the in-band scheme's pitfalls, the first part of this thesis studies an *out-of-band* scheme, where the control plane uses a dedicated wireless network different from the data-plane network. The increasingly adopted multihoming and increasingly available dual-band or multi-band radios prepare the IoT hardware platforms for implementing the out-of-band scheme. The prevailing IoT platforms are generally equipped with multiple heterogeneous network interfaces: Raspberry Pi 3 supports Ethernet, Wi-Fi, and BLE; Firestorm [10] supports BLE and Zigbee; Arduino has various add-on boards to support different radios. Some latest IoT platforms are equipped with both short-range and long-range low-power radios. The OpenMote B platform [11] integrates a Texas Instruments (TI) CC2538 Zigbee radio and an Atmel AT86RF215 sub-GHz long-range radio. The LoPy4 platform

[12] offers both a low-power wide-area networking radio (Sigfox or LoRaWAN) and a short-range radio (Zigbee or Wi-Fi). In particular, several recent System-on-a-Chip (SoC) modules integrate both short-range and long-range low-power radios. For example, the TI CC1350 [13] provides both BLE and sub-GHz long-range radio. Note that CC1350 is one of several SoC modules adopted by the SensorTag platform [14]. The TI CC1352R [15] provides both Zigbee and sub-GHz long-radio radios. Given the increasing availability of heterogeneous low-power radios, the first part of this thesis aims to study the network performance and manageability improvements as well as the related overhead by using different radios to form physically separated data- and control-plane networks.

To design the out-of-band control-plane network for MHWNs, Zigbee, Wi-Fi Direct, and the BLE mesh are ill-suited, since otherwise the control-plane network will be yet another multi-hop network that suffers the same manageability and fragility issues as the data-plane network. Cellular networks provide pervasive connectivity. However, the cellular network radios consume excessive power. As measured in [16], the Long-Term Evolution (LTE) user equipment consumes about 2 W for both downlink and uplink transmissions. Instead, we propose to use the emerging LPWAN technologies (e.g., LoRaWAN, Sigfox, Weightless, and other sub-GHz wireless) for the out-of-band control plane. From our measurements in this thesis, two LPWAN technologies have a power consumption of about 0.1 W only. Moreover, owing to the kilometers communication range of LPWAN links, the LPWAN-based control plane can be a one-hop star network, greatly simplifying its deployment and management.

The first part of this thesis investigates the suitability of two LPWAN technologies for designing out-of-band control planes for MHWNs, i.e., LoRaWAN and TI's sub-GHz long-range radio. These two LPWAN technologies use license-free ISM bands and do not rely on managed infrastructures. Extensive measurements are conducted to profile the power consumption, timing performance, and indoor communication performance of the two radios. While both of them can meet the technical requirements for building control planes, the first part of this thesis focuses on using LoRaWAN to prototype the proposed one-hop control plane and gain insights. While the low-power long-range communication capability is the key advantage of LoRaWAN, the following three limiting characteristics of LoRaWAN need to be managed properly. First, a LoRaWAN downlink frame from the controller to a network node must be in response to a precedent uplink frame. Thus, the transmissions of network control commands initiated by the controllers may be postponed to the network node's status reporting. Second, LoRaWAN supports uplink concurrency but no downlink concurrency. This downlink-uplink asymmetry impedes acknowledging each uplink frame, whereas the control plane generally desires reliable message delivery. Third, LoRaWAN adopts the ALOHA Media Access Control (MAC) protocol, which may perform poorly in traffic surges.

To address these issues, first part of this thesis presents the design and implementation of a prototype system called *LoRaCP* (Long-Range Control Plane). Based on our extensive measurements on LoRaWAN's energy and latency profiles, we design *LoRaCP-TxC*, a TDMA-based multi-channel transmission control approach featuring uplink heartbeats, the Negative Acknowledgment (NAK), and an ALOHA-based urgent channel, to manage the transmissions of the control-plane messages. The uplink heartbeats open downlink windows for controller-initiated network commands and maintain network nodes' clock synchronization for TDMA. With NAK, the controller does not need to acknowledge every uplink frame. The urgent channel complements the TDMA channels to mitigate the rigidness of TDMA. On a testbed of 16 nodes, LoRaCP is applied to physically separate the control plane of the Collection Tree Protocol (CTP) [8] from its Zigbee-based data-plane network. The LoRaCP-TxC is implemented in the application layer using the program library of a LoRaWAN platform. Extensive experiments show that LoRaCP increases CTP's packet delivery ratio from 65% to 80% in the presence of external interference, while consuming a per-node average radio power of 0.9 mA only under an operating voltage of 3.3 V, much lower than the active power of many recent sensor platforms' microcontrollers (e.g., 8.6 mA on Firestorm [10]).

1.4 Attack-Aware Synchronization-Free Data Timestamping in LoRaWAN

LoRaWAN is promising for the applications of collecting low-rate monitoring data from geographically distributed sensors, such as utility meters, environment sensors, roadway detectors, industrial measurement devices, etc. All these applications require data timestamping as a basic system service, though they may require different timestamp accuracies. For instance, data center environment condition monitoring generally requires sub-second accuracy for sensor data timestamps to capture the thermodynamics [17]. Sub-second-accurate timestamps for the traffic data generated by roadway detectors can be used to reconstruct real-time traffic maps [18]. In a range of industrial monitoring applications such as oil pipeline monitoring, milliseconds accuracy may be required [19]. In volcano monitoring, the onset times of seismic events detected by geographically distributed sensors require sub-10 milliseconds accuracy to be meaningful to volcanic earthquake hypocenter estimation [20].

There are two basic approaches, namely, *sync-based* and *sync-free*, to data timestamping in Wireless Sensor Networks (WSNs). In the sync-based approach, the sensor nodes keep their clocks synchronized and use the clock value to timestamp the data once generated. Differently, the sync-free approach uses the gateway with wall time to timestamp the data upon the arrival of the corresponding network frame. Based on various existing distributed clock synchronization protocols, multi-hop WSNs mostly adopt the sync-based approach. The sync-free approach is ill-suited for multi-hop WSNs, because the data delivery on each hop may have uncertain delays due to various factors such as channel contention among nodes.

In contrast, LoRaWANs prefer the sync-free approach for uplink data timestamping. Reasons are two-fold. First, different from multi-hop WSNs, LoRaWANs adopt a one-hop gateway-centered star topology that is free of the issue of hop-wise uncertain delays. Specifically, as the radio signal propagation time from an end device to the gateway is generally in microseconds, the LoRaWAN frame arrival time can well represent the time when the frame leaves the end device. As a result, timestamping the uplink data at the gateway can meet the milliseconds or sub-second timestamping accuracy requirements of many applications. Second, if the sync-based approach is adopted otherwise, the task of keeping the end devices' clocks synchronized at all times and the inclusion of timestamps in the LoRaWAN data frames will introduce communication overhead to the narrowband LoRaWANs. Therefore, performance-wise, the sync-free approach well matches LoRaWANs' star topology and addresses its bandwidth scarcity.

However, LoRaWAN's long-range communication capability also renders itself susceptible to wireless attacks that can be launched from remote and hidden sites.

The attacks may affect many end devices in large geographic areas. In particular, the conventional security measures that have been included in the LoRaWAN specifications (e.g., frame confidentiality and integrity) may be inadequate to protect the network from wireless attacks on the physical layer. Therefore, it is of importance to study the potential wireless attacks against the sync-free data timestamping, since incorrect timestamps render sensor data useless and even harmful. For example, when applying LoRa for IoT object localization by triangulation, tiny timestamping error will lead to large localization errors. In the second part of this thesis, we consider a basic threat of *frame delay attack* that directly invalidates the assumption of near-zero signal propagation time. Specifically, by setting up a *collider* device close to the LoRaWAN gateway and an *eavesdropper* device at a remote location, a combination of malicious frame collision and delayed replay may introduce arbitrary delays to the deliveries of uplink frames. Although wireless jamming and replay have been studied extensively, how easily they can be launched in a coordinated manner to introduce frame delay and how much impact (e.g., in terms of affected area) the attack can generate are still open questions in the context of LoRaWANs.

This second part of this thesis answers these questions via real experiments. Our measurements show that LoRa demodulators have lengthy vulnerable time windows, in which the gateway cannot decode either the victim frame or the collision frame, and raises no alerts. Thus, it is easy to launch stealthy attacks by exploiting the vulnerable time windows. In particular, as the attack does not breach the integrity of the frame content and sequence, the attack cannot be solved by cryptographic protection and frame counting. Our experiments in a campus LoRaWAN show that, a fixed setup of a collider and an eavesdropper can subvert the sync-free data timestamping service for end devices in a large geographic area of about $50,000\,\mathrm{m}^2$. In a broader sense, this attack threatens any system functions that require timely deliveries of uplink frames in LoRaWAN. Note that this attack is valid but marginally important in short-range wireless networks (e.g., Zigbee and Wi-Fi) because of the limited area affected by the attack and the difficulty in controlling the attack radios' timing. Differently, it is important to LoRaWANs because it can affect large geographic areas and the timing of the attack radios can be easily controlled due to LoRaWAN's long symbol times.

Therefore, an upgraded sync-free timestamping approach that integrates countermeasures against the attack and meanwhile preserves the bandwidth efficiency is desirable. Moreover, it should only require changes to the gateway. In the second part of this thesis, we aim to develop awareness of the attack by monitoring the end devices' radio Frequency Biases (FBs). Due to the manufacturing imperfections of the radio chips' internal oscillators, each radio chip generally has an FB that is the difference between the frequency of the carrier signal emitted by the chip and the nominal value. A change of FB detected by the gateway suggests the received frame may be a replayed one, since the adversary's replay device superimposes its own FB onto the replayed signal. To access the physical layer, we integrate a low-cost (US$25) Software-Defined Radio (SDR) receiver [21] with a commodity LoRaWAN gateway to form our LoRaTS gateway. We develop time-domain signal processing algorithms for LoRaTS to estimate the FB. Experiments show that (i) with a received Signal-to-Noise Ratio (SNR) of down to $-18\,\mathrm{dB}$, LoRaTS achieves an accuracy of $120\,\mathrm{Hz}$ in estimating FB, which is just 0.14 parts-per-million (ppm) of the channel's central frequency of 869.75 MHz; (ii) the frame replay by an SDR transceiver introduces an additional FB of at least 0.24 ppm. Thus, LoRaTS can track FB to detect the replay step of the frame delay attack. Note that the detection does not require uniqueness or distinctiveness of the FBs across different LoRa transceivers, because it is based on changes of FB.

In summary, LoRaTS supports the bandwidth-efficient sync-free timestamping and requires no modifications on the LoRaWAN end devices. It is a low-cost countermeasure that increases the cost and technical barrier for launching effective frame delay attacks, since the attackers need to eliminate the tiny FBs of their radio apparatuses. LoRaTS strikes a satisfactory trade-off between network efficiency and the security level required by typical LoRaWAN applications.

1.5 Summary of Contributions

Existing IoT networking technologies face two challenges for interacting with massive objects. The first is the efficiency of utilizing the wireless bandwidth. With the development of wireless technology, the frequency spectrum becomes more and more crowded. Moreover, the number of IoT devices is increasing while the bandwidth is limited.

The second is the resilience of the IoT network. The IoT network suffers from external disturbances, such as interference from neighbor RF technologies, propagation blockage, and movement of end devices. Moreover, due to the broadcast nature of wireless communications, wireless IoT networks are susceptible to wireless attacks. Therefore, it is critical to improve the resilience of the IoT network to survive external interference and even cyber-attacks.

To address these two challenges, this thesis makes several contributions summarized as follows:

- It designs an out-of-band control plane for multi-hop wireless networks using LoRaWAN. A set of comparative TinyOS simulations are conducted to study the performance of CTP in an MHWN under the distributed and centralized network controls, as well as in-band and out-of-band centralized network controls. The comparative simulations motivate the design of the one-hop out-of-band control plane. Measurement studies are performed to profile the power consumption, timing accuracy, and indoor communication performance of two LPWAN radios, i.e., LoRaWAN and TI's sub-GHz long-range radio. A LoRaWAN-based out-of-band control plane prototype system called LoRaCP is designed and implemented.

- It implements the frame delay attack against LoRaWANs. The attack will devastate the sync-free data timestamping service and any other system services requiring timely frame delivery. The attack implementation shows the susceptibility of LoRaWAN to wireless attacks. Simulations and experiments show the large sizes of the geographic areas vulnerable to the attack. Based on an analytic model of LoRa's Chirp Spread Spectrum (CSS) modulation and profiling, we show that the bias of the LoRa signal's carrier frequency from the nominal value is an effective radiometric feature. The thesis further designs a time-domain signal processing pipeline to accurately estimate end devices' FBs. The proposed approach can reveal the additional FB introduced by the attack. Thus, the approach can achieve sync-free timestamping with the awareness of frame delay attack.

- This thesis implements the proposed approaches on real devices. For LoRaCP, a testbed consisting of a gateway and 16 end nodes is built. Experiments on the testbed show the network resilience brought by LoRaCP for

CTP. For LoRaTS, this thesis implements the frame delay attack on a campus LoRaWAN infrastructure. Experiments in both indoor and urban environments with different end devices are conducted. Results show that LoRaTS can detect the frame delay attacks that introduce additional FBs. Extensive experiments are also conducted to evaluate the impact introduced by ambient factors. For example, an 87-hour long experiment shows the impact of the ambient temperature on FB and the attack detection performance.

1.6 Thesis Structure

The rest of this thesis is organized as follows. Chapter 2 presents the technical preliminaries of LoRaWAN. Chapter 3 presents LoRaCP, a one-hop out-of-band control plane built using LoRaWAN for multi-hop wireless networks. We first motivate our proposed solution by several comparative simulation studies. Then, we describe the detailed design of our approach. From evaluation results on a real testbed, our approach achieves good performance and energy efficiency. In Chapter 4, by launching frame delay attack on a real LoRaWAN system, we show that the frame delay attack is a real security threat to LoRaWAN. To gain awareness of the frame delay attack, we present LoRaTS, a gateway added with the RTL-SDR to analyze the bias of the uplink carrier frequency to improve the security of gateway-side data timestamping. Chapter 5 concludes this thesis and discusses future work.

Chapter 2

Preliminaries on LoRaWAN

In this chapter, we introduce the preliminaries of LoRaWAN[1]. Section 2.1 presents LPWAN, a new wireless platform. Section 2.2 presents primers of LoRaWAN's physical modulation technique. Section 2.3 introduces LoRaWAN and its characteristics.

2.1 LPWAN

LPWANs are an emerging wireless platform that aims to sustain power-constrained end devices (e.g., those based on batteries or energy harvesting) to operate for years while communicating at low data rates to gateways several kilometers away. LPWAN technologies will largely increase the degree of connectivity of IoT and enable deep penetration of IoT objects into the urban territories. Fig. 2.1 illustrates the comparisons among various wireless technologies in terms of radio power consumption and communication ranges. From the figure, LPWANs (e.g., LoRaWAN [25], Sigfox [26], Weightless-P [27], and NB-IoT [28]) form an important pole in the spectrum of radio power consumption versus communication range.

[1] This chapter is partially published on [22], [23] and [24].

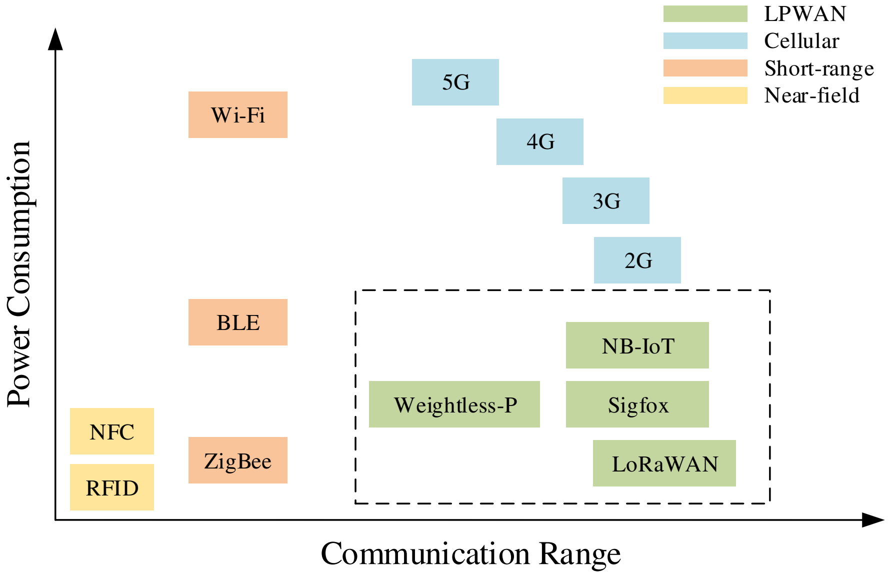

FIGURE 2.1: Power consumption versus communication range for various radios.

Specially, we select LoRaWAN to conduct research because of its use of license-free ISM band, open data link standard, and unmanaged network which means that we do not need to depend on the infrastructure provided by ISP. And LoRaWAN has been widely deployed for many applications as shown in the right figure.

2.2 LoRa Primer

LoRa is a physical layer technique that uses a CSS modulation and operates in sub-GHz ISM bands (e.g., 868 MHz in Europe). In LoRa's CSS, a chirp is a finite-time signal with time-varying instantaneous frequency that swaps the whole bandwidth of the communication channel in a linear manner. Given a certain central frequency, denoted by f_c, an up chirp's instantaneous frequency increases from $f_c - \frac{BW}{2}$ to $f_c + \frac{BW}{2}$, whereas a down chirp's instantaneous frequency decreases from $f_c + \frac{BW}{2}$ to $f_c - \frac{BW}{2}$. The time duration of a chirp is determined by the *spreading factor* and *bandwidth*, which are denoted by SF and BW, respectively. Specifically, the chirp time is given by

$$t = \frac{2^{SF}}{BW}. \tag{2.1}$$

For the EU868 frequency band, there are six spreading factors, ranging from 7 to 12. For example, with $SF = 7$, $BW = 125\,\text{kHz}$, $f_c = 869.75\,\text{MHz}$ and initial phase

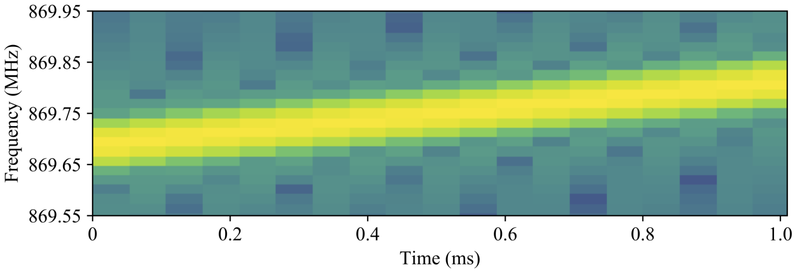

FIGURE 2.2: Spectrogram of an up chirp.

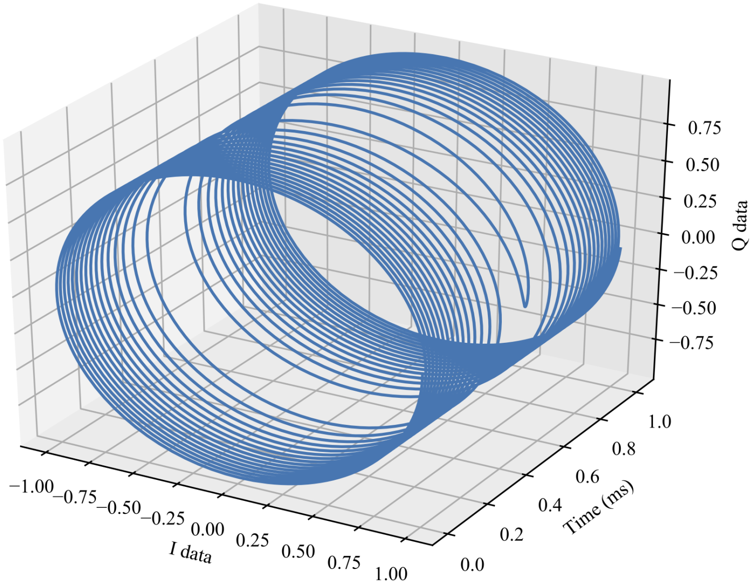

FIGURE 2.3: I and Q data ($\theta = 0$) of an up chirp.

$\theta = 0$, an up chirp's spectrogram is shown in Fig. 2.2. Fig. 2.3 presents the in-phase (I) and quadrature (Q) data of this chirp in the time domain.

2.3 LoRaWAN and Its Characteristics

2.3.1 Introduction of LoRaWAN

LoRaWAN is an open data link layer specification based on LoRa, a proprietary PHY layer technique that uses a Chirp Spread Spectrum modulation and operates in sub-GHz ISM bands (e.g., EU868 MHz and US915 MHz). LoRa admits configuring the ratio between the symbol rate and chip rate by specifying an integer

Spreading Factor (SF) within [6, 12]. Specifically, each symbol is modulated by 2^{SF} chips. A higher SF increases the signal-to-noise ratio and the communication range, but decreases the symbol rate. In this thesis, six SF settings (from SF7 to SF12) are used.[2] The communications using different SFs are orthogonal and thus can be concurrent. LoRa also admits configuring bandwidth and coding rate. These two parameters can affect the communication performance of LoRa. In this thesis, we configure them to be 125 kHz and 4/5, respectively. Our performance profiling can be easily extended to address other settings of these two parameters.

A LoRaWAN network is formed by one or more *gateways* and many *end devices*. The gateway, often Internet-connected, can simultaneously handle the communications with multiple nodes in different channels. LoRaWAN defines three classes (A, B, and C) of end devices. A Class-A device's uplink transmission is followed by two downlink windows (RX1 and RX2). Downlink communications to the node at any other time will have to wait until the next uplink from the node. As Class-A is the most power efficient and supported by any end device, Class-A is chosen to design the out-hop out-of-band control plane.

2.3.2 Characteristics of LoRaWAN

The low-power long-range communication capability is the main advantage of LoRaWAN that makes it promising for control planes of MHWNs. However, we need to keep in mind the following two limiting characteristics of LoRaWAN in the design of the one-hop out-of-band control plane.

- **Downlink-uplink asymmetry:** LoRaWAN is mainly designed and optimized for uplinks from end devices to gateway. For instance, the LoRaWAN concentrator can receive frames from multiple channels simultaneously, whereas it can send a single downlink frame only at a time. Moreover, the Class-A specification requires that any downlink transmission must be unicast, in response to a precedent uplink transmission.

- **Lossy links:** From existing tests [30], with SF12, the frame reception rate is about 80% at a distances of 2.5 km. To build a reliable control-plane network,

[2] According to a LoRa chip's datasheet [29], SF6 is a special setting that is not enabled by default. Thus, we do not use SF6.

the frame losses need to be dealt with properly. Acknowledging each uplink frame is wasteful given the scarce downlink time as discussed earlier.

Moreover, the following two default features of LoRaWAN need to be considered and/or re-engineered in the design of an efficient one-hop out-of-band control plane.

- **ALOHA MAC:** LoRaWAN uses ALOHA that may perform poorly in surges of control plane messages. Moreover, as LoRa does not prescribe carrier sense capability, CSMA is not viable. Time-Division Multiple Access (TDMA) is often adopted for reliability that control planes desire. However, as shown in the first part of this thesis, the implementation of TDMA on LoRaWAN is non-trivial. Moreover, a strict TDMA may result in undesirable delays in transmitting urgent messages. In addition, for the universality of LoRaCP, a transmission control protocol that does not need to modify LoRaWAN's MAC-layer code is desired.
- **Text transmission only:** LoRa admits hexadecimal ASCII string only. To send an integer, it transmits the hexadecimal ASCII code word of each literal character of the integer.

In the design of the one-hop out-of-band control plane (cf. §3.4), the downlink-uplink asymmetry and lossy links will be managed by the NAK mechanism. Moreover, we will design a TDMA-based multi-channel transmission control protocol on top of LoRaWAN's ALOHA MAC. Both features can be implemented in the application layer.

2.3.3 LoRaWAN Primer

LoRa is a physical layer technique that adopts CSS modulation. LoRaWAN is an open data link specification based on LoRa. A LoRaWAN is a star network consisting of a number of *end devices* and a *gateway* that is often connected to the Internet. Gateways are often equipped with GPS receivers for time keeping. The transmission direction from the end device to the gateway is called *uplink* and the opposite is called *downlink*. LoRaWAN defines three classes for end devices, i.e., Class A, B and C. In Class A, each communication session must be initiated by an uplink transmission. There are two subsequent downlink windows. Class A end devices can sleep to save energy when there are no pending data to transmit. Class A adopts the ALOHA media access control protocol. Class B extends Class A with additional scheduled downlink windows. However, such scheduled downlink windows require the end devices to have synchronized clocks, incurring considerable overhead as we will analyze shortly. Class C requires the end devices to listen to the channel all the time. Clearly, Class C is not for low-power end devices. In this thesis, we focus on Class A, because it is supported by all commodity platforms and energy-efficient. To the best of our knowledge, no commodity platforms have out-of-the-box support for Class B that requires clock synchronization.

Chapter 3

One-Hop Out-of-Band Control Planes for Multi-Hop Wireless Networks

This chapter is organized as follows [1]. §3.1 reviews related work. §3.2 presents a number of simulation-based examples to motivate the out-of-band scheme for control planes. §3.3 profiles the performance and overhead of LoRaWAN and TI's sub-GHz long-range radio. Based on the profiling results, §3.4 designs LoRaCP. §3.5 presents the evaluation results of LoRaCP in testbed experiments.

3.1 Related Work

The increasing availability of multihoming and multi-band radios enables researchers to investigate the benefits by leveraging on multiple radios. Existing studies that exploit multiple radios can be broadly divided into two classes of *bandwidth aggregation* and *SCDP*.

Bandwidth aggregation uses multiple network interfaces to transmit/receive data simultaneously to increase throughput. Habak et al. [31] surveyed early bandwidth aggregation literature, and categorized them into solutions at the application, transport, network, and link layers. Early studies include the Multi-Radio

[1] This chapter is partially published on [23] and [24].

Unification Protocol (MRUP) [32] and system designs of multi-radio approaches [33]. The MRUP proposed by Adya et al. [32] runs at the link layer to coordinate the operations of multiple wireless network cards tuned to non-overlapping frequency channels, based on locally available information only. Subsequently, Bahl et al. [33] study various design issues of the multi-radio approaches in the hardware, algorithmic, and protocol aspects. Recent development that is not covered by the survey paper [31] is reviewed as follows. These new studies are divided into two categories. The first category exploits *homogeneous* radios. FatVAP [34] enables a 802.11 wireless card to connect to multiple access points. FastForward [35] uses two 802.15.4 radios operating on different channels, with one receiving and the other forwarding data simultaneously. The second category exploits *heterogeneous* radios. In [36], various trade-offs in designing energy efficient multi-radio platforms are investigated. In MicroCast [37], a group of smartphones cooperate in downloading a video from the cellular and share their downloaded portions through device-to-device links (e.g., Wi-Fi). MultiNets [38] deals with the switching between multiple network interfaces on mobile devices. In [39], Mu et al. optimize the selection of radios and their transmission powers. Recent studies [40, 41] characterize the performance and energy consumption of Multipath TCP through multiple radios of a mobile device. Different from bandwidth aggregation that combines multiple network interfaces in the data plane to increase throughput, SCDP aims to improve network optimality and manageability.

Software-Defined Networking (SDN), with SCDP as its core concept, is a growing momentum in data-intensive networks (e.g., data center and enterprise backbone networks). To avoid the undesirable coupling between the control and data planes, SDN recommends the out-of-band scheme [42]. SCDP can be naturally applied in wireless local area networks and cellular networks, as their topologically centralized access points and base stations can run the control-plane logics for better resource allocation and mobile node handover [43]. However, there is limited research on SCDP in multi-hop wireless networks. An OpenFlow-enabled Wi-Fi mesh was built in [44], where each Wi-Fi card is split into two virtual interfaces with different Service Set Identifiers (SSIDs) and the two planes are two multi-hop networks in their respective SSIDs. In the initial thinking of applying SCDP in wireless sensor networks [45], the design choice of in-band or out-of-band control plane is dubious. To the best of our knowledge, WASP [46] is the only system that implements out-of-band control plane for multi-hop wireless networks. WASP uses Wi-Fi Direct and cellular network of smartphones to form the data and control planes, respectively. Different from WASP, the first part of this thesis focuses on low-power networks with a limited energy budget.

LoRaWAN, an emerging LPWAN technology, has received increasing research in recent years. Existing studies mainly aim to improve the communication performance and battery lifetime of the LoRaWAN end devices. LoRaWAN's communication performance is profiled via field measurements [30, 47, 48]. A recent work [49] provides an extensive and in-depth measurement study in various outdoor environments. The work [50] presents the design of a multi-channel and multi-hop MAC protocol for a LoRa-based wildlife monitoring system. The work [51] presents an approach that can predict LoRaWAN link quality based on multispectral images from remote sensing. The Choir [52] system exploits the diverse frequency biases of the LoRaWAN end devices to decode colliding frames from different end devices. The Charm [53] system exploits coherent combining to decode a frame from the weak signals received by multiple geographically distributed LoRaWAN gateways. In addition to the above studies on improving the scalability and robustness of LoRaWAN networks, several recent studies have proposed various backscatter designs for LoRa [54–57]. Different from the off-the-shelf LoRa nodes that use their own power sources to drive the radio transmissions, a backscatter device generates backscatter signals in response to some activation signal to transmit bits. Thus, backscatter can extend the battery lifetime. Different from these studies, this paper focuses on exploiting LoRa's long-range communication capability to improve the data plane's network performance and resilience against external disturbances.

3.2 Simulation Studies

In this section, comparative simulation studies are conducted to motivate the use of out-of-band centralized network control to improve network performance and resilience. Specifically, §3.2.1 compares the distributed network control scheme and the centralized network control scheme. §3.2.2 compares the in-band centralized network control and the out-of-band centralized network control. All the simulations are conducted in the TinyOS simulator TOSSIM [58].

3.2.1 Distributed versus Centralized Network Control

This section compares through simulations the network performance achieved by the CTP [8] and its centralized variant that is called CTP-SCDP. In §3.5, LoRaCP will be applied to implement CTP-SCDP and evaluated on a testbed. The first part of this thesis uses CTP as the case study network protocol, because CTP has an open implementation and is a standard component of the industry-class TinyOS Production operating system [59]. The results based on CTP will provide insights into the performance improvement by SCDP and showcase the use of LoRaCP to physically separate the control and data planes. The obtained insights are also useful to the SCDP designs of other MHWN protocols.

CTP aims to maintain a minimum-cost routing tree in the presence of dynamic link quality characterized by the Expected Transmission Count (ETX). The cost of a route to the tree root is the sum of the ETXs of the links on the route. A node i estimates the route cost using the Residual ETX (RETX), which is given by $\mathrm{RETX}_i = \mathrm{ETX}_{i,p} + \mathrm{RETX}_p$, where $\mathrm{ETX}_{i,p}$ is the ETX of the link between node i and its parent node p, and RETX_p is node p's RETX. CTP works in a distributed manner, in that each node i selects its parent p from the set of its neighbor nodes $\mathcal{N}$ based on local information only. Specifically, $p = \arg\min_{j \in \mathcal{N}} \mathrm{ETX}_{i,j} + \mathrm{RETX}_j$, where $\mathrm{ETX}_{i,j}$ is estimated based on the transmissions of beacons and data frames; RETX_j is broadcast in node j's beacons.

In CTP, the information about the quality of a link propagates to the whole network during the beaconing process. However, this propagation takes time. Thus, when link quality changes over time, the RETX of any node i cannot capture the latest ETXs of the links on its route to the root. In particular, the closer the links on the route are to the root, node i's knowledge about the links (which is encompassed in RETX_i) is more out-of-date. As a result, CTP may not construct the minimum-cost tree in the presence of time-varying link quality. Differently, in CTP-SCDP, the latest ETXs can be updated to the network controller in time. Specifically, when the change of any ETX of a node exceeds a certain threshold, the node can send the latest ETX to the network controller via the dedicated control-plane network. Upon receiving an updated ETX, the network controller shall recompute the optimal routing and send the changes in routing to concerned nodes via the control-plane network as well.

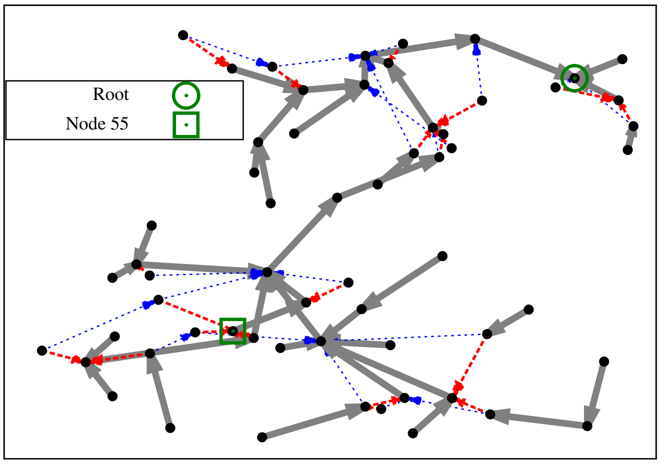

FIGURE 3.1: Node placement and the routing trees constructed by CTP and CTP-SCDP. The solid thick gray links are shared by the CTP and CTP-SCDP trees; the dashed thick red links are on the CTP tree only; the dashed thin blue links are on the CTP-SCDP tree only.

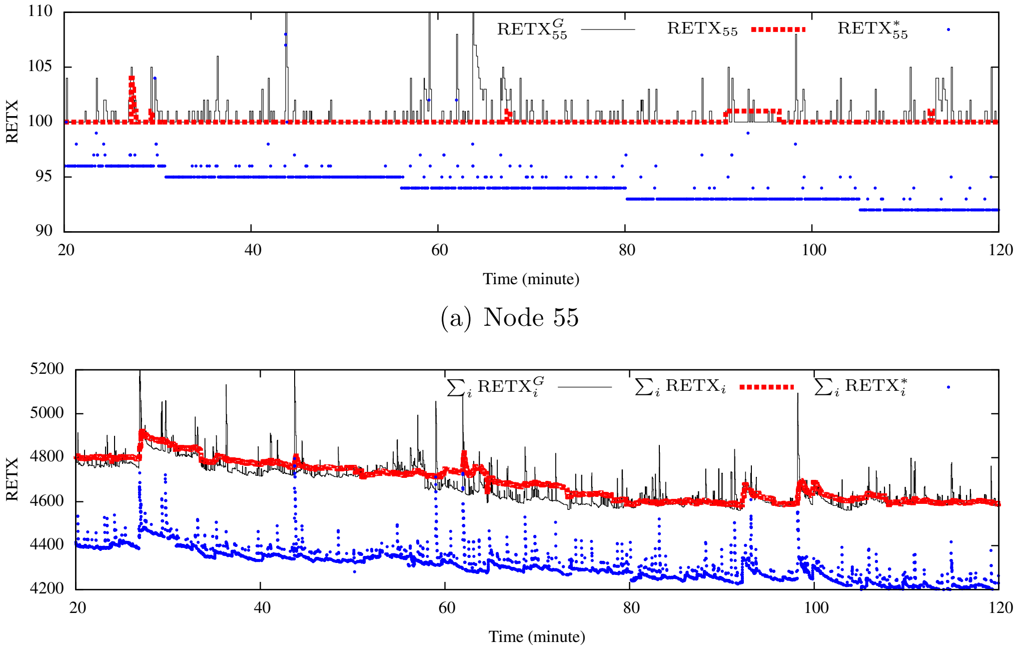

(a) Node 55

(b) All nodes

FIGURE 3.2: The groundtruth RETX and estimated RETX in CTP as well as the optimal RETX if CTP-SCDP is adopted. The results show that CTP cannot build a minimum-cost tree.

Simulations are conducted in TOSSIM to compare CTP and CTP-SCDP. A total of 60 nodes are placed randomly in a 200 m × 200 m region as illustrated in Fig. 3.1. Note that the size of the simulated region is similar to the dimension of a building (190 m long) in which we deploy LoRaWAN and TI's sub-GHz nodes to measure LPWAN's communication performance in §3.3. The density of the simulated nodes is similar to the building environment monitoring applications, e.g., temperature distribution monitoring in data centers and hazard dust/gas concentration monitoring on factory floors. Link gains are generated according to the Euclidean distances between nodes using a tool in TOSSIM. TOSSIM applies a radio propagation model for each node. It can also simulate the RF noises and interference that a node is subjected to. Radios' hardware noise floor is set to be -90 dBm, which is a mild noise level. To simulate CTP-SCDP, a node is added as the network controller, which has sufficiently large link gains with any other nodes, such that the control-plane network is a one-hop star network. The TOSSIM is configured such that the data-plane links do not interfere with the control-plane links. As this one-hop star network is not used to transfer sensor data, this CTP-SCDP system follows the out-of-band control plane scheme. In CTP-SCDP, node i sends the latest $\mathrm{ETX}_{i,j}$ to the network controller. Upon receiving an ETX update, the controller updates a directed graph with the ETXes as the edge costs and recomputes the minimum-cost routing tree using the Dijkstra's algorithm. Then, the controller sends the new parent information to the nodes.

Two sets of simulations are conducted to show the benefits of SCDP. The first set shows the suboptimal performance of CTP. Specifically, CTP and CTP-SCDP run concurrently, but the controller in CTP-SCDP does not send routing control commands to the nodes. Thus, the routing is managed by CTP only. The following evaluation metrics are considered:

1. RETX of node i estimated by CTP (denoted by RETX_i) and the sum of all RETXes (denoted by $\sum_i \mathrm{RETX}_i$);

2. The *ground-truth* RETX of the route determined by CTP for node i (denoted by RETX_i^G), which can be measured as the sum of the latest ETXes of the links on the route obtained by the controller in CTP-SCDP, as well as the sum of all ground-truth RETXes (i.e., $\sum_i \mathrm{RETX}_i^G$);

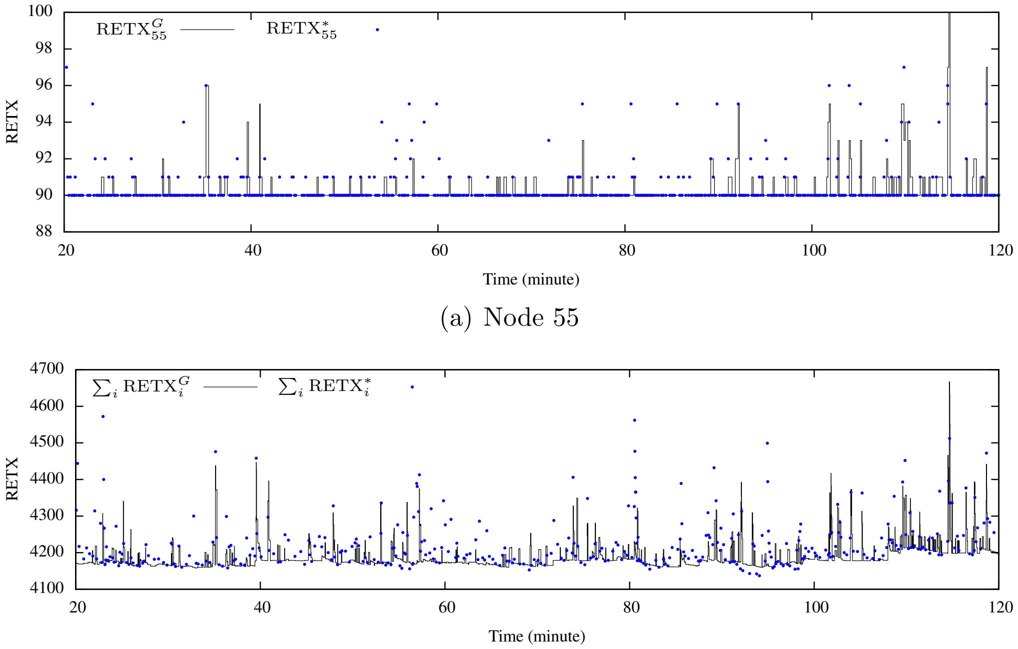

(a) Node 55

(b) All nodes

Figure 3.3: The groundtruth RETX achieved by CTP-SCDP and the true optimal RETX. The results show that CTP-SCDP can build a minimum-cost tree.

3. The minimum RTEX of node i computed by CTP-SCDP (denoted by RETX^*_i) and the sum $\sum_i \mathrm{RETX}^*_i$.

The simulated time duration is two hours, during which each node generates a data packet every eight seconds. Fig. 3.1 shows the routing trees computed by CTP and CTP-SCDP at the end of the simulation. They are different. Fig. 3.2 shows the evaluation metrics for Node 55 and all the nodes over the two hours. It can be seen that, compared with the ground truth (i.e., the solid black curves), CTP's knowledge about the chosen routes (i.e., the dashed red curves) cannot capture many transient changes in the ground truth, because of the information propagation latency in the distributed network control. Compared with the global optimal (i.e., the blue dots), the routes chosen by CTP have higher costs.

In the second set of experiments, only CTP-SCDP runs. Fig. 3.3 shows the results. It can be seen from the figure that the routes chosen by CTP-SCDP generally achieve the minimum costs. The above two sets of simulations show that the centralized network control improves the network performance in dynamic network conditions. Thus, the centralized control enabled by SCDP is desirable for performance-critical networks such as those deployed for industrial applications [7].

3.2.2 In-band versus Out-of-Band Network Control

Our simulations in §3.2.1 demonstrate the underperformance of distributed network control. As discussed in §1.3, a number of mission-critical MHWNs have adopted centralized network control to improve network performance and manageability. They all follow the in-band control plane scheme. However, the physical coupling of the control and data planes generates various challenges. For instance, given the fragile nature of wireless, how to protect the in-band control plane against data-plane link failures is a challenging problem. Recent research has investigated this issue. Nevertheless, existing solutions provide limited protection capability. For instance, the solution proposed in [9], though sophisticated, can handle a single link failure only. The in-band control plane protection under a general setting is still an open issue. In this section, simulations are conducted to compare the CTP systems with in-band and out-of-band centralized network controls, which are referred to as in-band CTP-SCDP and out-of-band CTP-SCDP, respectively. The out-of-band CTP-SCDP in this section is same as the CTP-SCDP in §3.2.1. In the in-band CTP-SCDP, each node will report the link quality information to the network controller using the multi-hop data-plane network. Upon receiving the link quality information from a node, the controller will compute the optimal routing path for this node. Then, the controller disseminates the new routes using targeted messages to the individual nodes. In both the in-band and out-of-band CTP-SCDP, a packet will be dropped after 30 unsuccessful (re-)transmissions to the next hop. In what follows, the in-band and out-of-band schemes are compared in terms of resilience to link noises and node faults, as well as the network convergence speed.

3.2.2.1 Resilience to link noises

The simulation settings in §3.2.1 are adopted in this section. The radios' hardware noise floor is set to be $-90\,\mathrm{dBm}$. The simulated time duration is two hours, during which each node generates a data packet every eight seconds. Fig. 3.4(a) and Fig. 3.4(b) show the evaluation metrics of the in-band control scheme for Node 55 and all the nodes over the two hours. The results show that the in-band control scheme cannot achieve optimal network performance. Compared with the true optimal network performance that is achieved by the out-of-band CTP-SCDP (i.e., the blue dots), the routes chosen by the in-band CTP-SCDP have higher cost. This

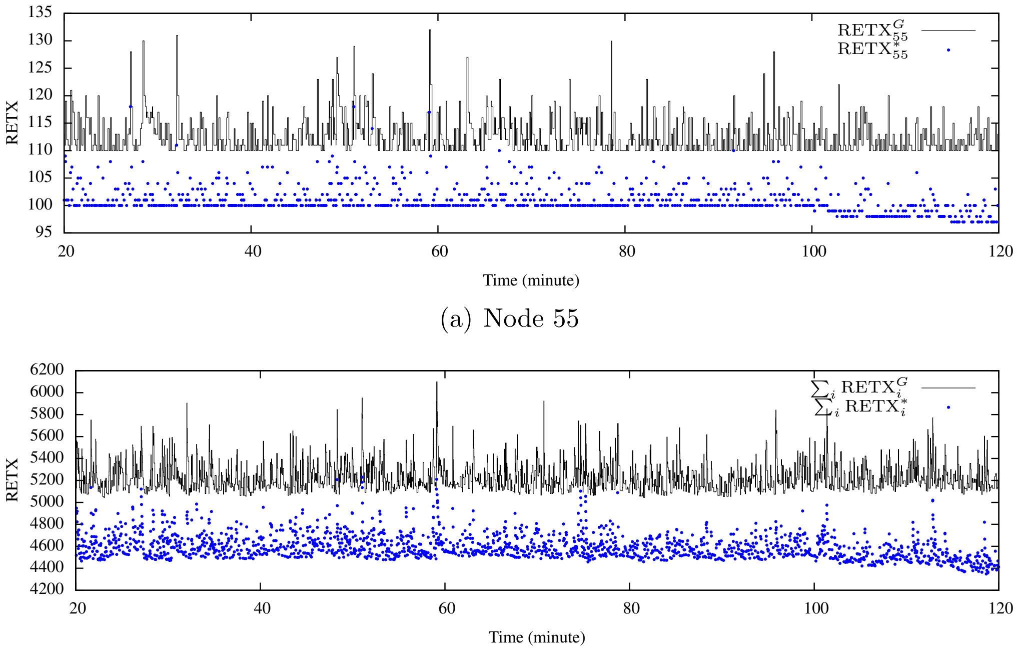

(a) Node 55

(b) All nodes

Figure 3.4: The groundtruth RETX achieved by the in-band CTP-SCDP and the true optimal RETX. The results show that the in-band CTP cannot build a minimum-cost tree.

Table 3.1: The Packet Delivery Ratios (PDRs) achieved by the in-band and out-of-band CTP-SCDP schemes under different settings of radio noise floor.

Noise floor setting (dBm)	-105	-100	-95	-90	-85
PDR of in-band CTP-SCDP	99.97%	99.87%	99.80%	99.80%	8.48%
PDR of out-of-band CTP-SCDP	100%	100%	100%	100%	10.90%

is because, with the in-band CTP-SCDP, the data-plane and control-plane packets delivered by the same data-plane network may collide. The delayed and unsuccessful deliveries of the control-plane packets will lead to performance degradation of the data-plane network. Due to the limited payload size of a TinyOS message, the need of sending new routes to different nodes in separate packets also increases the contention between the data-plane and control-plane networks. Differently, in the out-of-band CTP-SCDP, the data-plane and control-plane packets will not interfere with each other.

The data Packet Delivery Ratios (PDRs) achieved by the in-band and out-of-band CTP-SCDP networks under various settings of the radio noise floor are measured. Table 3.1 shows the results. It can be seen that when the noise floor is from $-105\,\text{dBm}$ to $-90\,\text{dBm}$, both the in-band and out-of-band CTP-SCDP networks

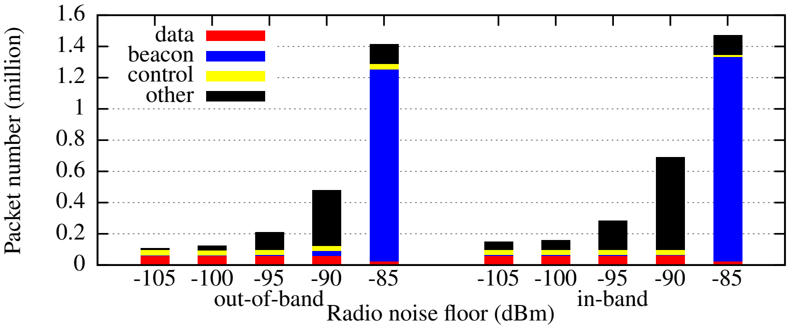

FIGURE 3.5: The breakdown of the total transmitted packets in the out-of-band (left group) and in-band (right group) CTP-SCDP networks under various radio noise floor settings. The bars labeled "data", "beacon", "control" represent the data packets, beacons, and control-plane packets generated by all the nodes; the bar labeled "other" represents the packets (either data or control packets) successfully forwarded and re-transmitted packets by all the nodes.

can deliver (almost) all data packets to the sink. When the noise floor setting is −85 dBm, both networks have similarly low PDRs. This is because the link quality of the data-plane network is too poor to support reliable packet transmissions.

The numbers of packets in the following four categories of the whole network are counted to better understand the differences between the in-band and out-of-band schemes: (1) data packets generated by the nodes, (2) beacons, (3) control packets generated by the nodes, and (4) others including the packets (either data or control packets) successfully forwarded and re-transmitted by the nodes. Fig. 3.5 shows the above counts in the in-band and out-of-band CTP-SCDP networks under various radio noise floor settings. We now summarize the observations from Fig. 3.5 and discuss the reasons of these observations.

- When the noise floor setting increases from −105 dBm to −90 dBm, the increases of re-transmissions are observed. This is due to the deteriorated link quality. Moreover, the in-band CTP-SCDP network always has more re-transmissions than the out-of-band CTP-SCDP network. This is because the data and control planes of the in-band CTP-SCDP network contend for the bandwidth. The delayed or unsuccessful deliveries of the control-plane packets also lead to the sub-optimality of the data collection tree on the data-plane network as observed earlier. The sub-optimality leads to increased total traffic in return, although the in-band CTP-SCDP network maintains nearly 100% PDRs as shown in Table 3.5.

- When the noise floor setting is $-85\,\mathrm{dBm}$, the number of successful re-transmissions is reduced compared with that when the noise floor is $-90\,\mathrm{dBm}$. This is because the networks have excessive dropped packets due to the poor link quality. Moreover, the nodes transmit beacons intensively to track the highly dynamic link quality. Under this noise floor setting, from Table 3.1, both the in-band and out-of-band CTP-SCDP networks have low PDRs. Compared with the results for lower noise floor settings, the data packets both in the out-of-band and in-band CTP-SCDP networks reduce. This is because there are too many data packages to be re-transmitted and only a portion of generated data packets are processed by the end of the simulation.

It can be seen from the above results that, when the network can maintain good PDRs in the presence of link noises, the out-of-band scheme can reduce the cost in forwarding and re-transmitting packets.

3.2.2.2 Network convergence

With a one-hop out-of-band control-plane network, the nodes' local information can be directly transmitted to the controller. In contrast, in the distributed scheme or the in-band SCDP network, the local information propagates to the controller hop by hop to the controller, resulting in longer delays. Therefore, the out-of-band CTP-SCDP can converge to a new routing schedule faster when the network boots or has changes (say due to nodes' movements). A set of experiments are conducted to compare the convergence of the networks under different network control schemes. A metric $D(t_0, t_1)$ is computed to characterize the convergence of the network within a time duration of (t_0, t_1). Specifically, the convergence metric is given by $D(t_0, t_1) = \frac{\sum_{t=t_0}^{t_1} \left(\sum_i \mathrm{RETX}_i^G(t) - \sum_i \mathrm{RETX}_i^*(t)\right)^2}{t_1 - t_0}$, where $\mathrm{RETX}_i^G(t)$ and $\mathrm{RETX}_i^*(t)$ denote the ground-truth RETX of node i at time t and the minimum RTEX of node i at time t, respectively. A smaller $D(t_0, t_1)$ value suggests a better convergence to the minimum-cost tree.

Our CTP-SCDP networks adopt the CTP's adaptive beaconing to keep track of the $\mathrm{ETX}_{i,j}$ (i.e., the ETX between node i and node j). The network convergence highly depends on the frequency of the adaptive beaconing. We now briefly introduce the adaptive beaconing mechanism. More details can be found in [8]. When the

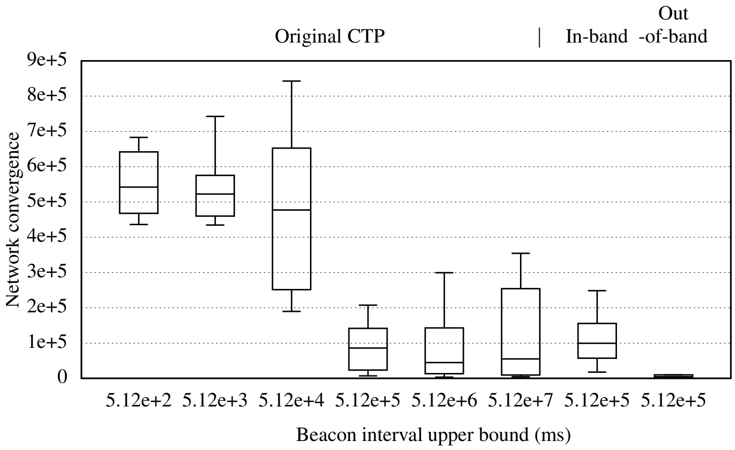

FIGURE 3.6: Measured network convergence metric $D(15\text{th minute}, 120\text{th minute})$ of the original CTP, in-band CTP-SCDP, and out-of-band CTP-SCDP, under various settings of the beaconing interval upper bound.

network is booted or the CTP detects a potential change of ETX based on the (re-)transmissions of the data packets, CTP sends beacons every T_{bl} seconds, where T_{bl} is the lower bound of the beaconing interval. After that, it increases the beaconing interval exponentially up to T_{bu}, where T_{bu} is the upper bound. This adaptive beaconing saves radio energy when the link quality is stable. In CTP, the default settings for T_{bl} and T_{bu} are 128 ms and 5.12×10^5 ms, respectively. The first six error bars in Fig. 3.6 show the measured convergence metric $D(15\,\text{min}, 120\,\text{min})$ of the CTP when the T_{bu} increases from 5.12×10^2 ms to 5.12×10^7 ms and T_{bl} adopts its default setting. Specifically, each error bar shows the distribution of the convergence metric in a total of fifteen 2-hour simulations. It can be seen that the knee point of the average D value is at CTP's default setting for T_{bu} (i.e., 5.12×10^5 ms). With smaller settings for T_{bu}, the CTP network does not converge well (i.e., large D values). This is because, in CTP, both data packets and control packets have the same priority and the highly frequent beacons can result in contentin. With larger settings for T_{bu}, the distributions of the D value are more spread. This is because the CTP network updates the ETXs using less frequent beacons and the network may not react to network condition changes in time. Thus, the default setting of $T_{bu} = 5.12 \times 10^5$ strikes a satisfactory trade-off between the overhead of beaconing and performance in tracking ETXs timely for maintaining low cost.

The last two error bars in Fig. 3.6 show the results for the in-band and out-of-band

CTP-SCDP when the default setting for T_{bu} is applied. It can be seen that the in-band CTP-SCDP network has worse network convergence compared with the CTP adopting the default T_{bu} setting. This is because, compared with the distributed control scheme, the in-band scheme imposes additional overhead of conveying control packets to/from the centralized network controller using the data-plane network. The out-of-band CTP-SCDP achieves the best network convergence.

3.2.2.3 Resilience to node faults

In real deployments, WSN node faults (due to say battery depletion or hardware malfunction) are not uncommon. The distributed network control in general can well handle node faults. For instance, in CTP, the ETXs associated with a faulty node will quickly increase to infinity. As a result, the faulty node will not be selected as the parent node in the route to the sink. Differently, in a pure SCDP network, the decision of switching the parent node shall be made by the centralized network controller. If the SCDP network adopts the in-band scheme, the deliveries of the network status updates and parent node switching decisions may be affected by the node faults. For instance, in the simulations of the in-band CTP-SCDP, we switch off Node 55 shown in Fig. 3.1 in the 80th minute to simulate a node fault event. As a result, the nodes on the sub-tree rooted at Node 55 are disconnected from the network – they cannot send/receive data/control packets to/from the network controller. The network controller can only infer the occurrence of the fault based on the increasing ETXs reported by the neighbors of Node 55 that are still connected with the network controller. To stick to the pure SCDP scheme with in-band control plane, a separate network restoration mechanism will be needed. In particular, the control packets to the disconnected sub-tree cannot be delivered using the old routes. A message flooding may be needed to deliver the control packets. Differently, with an one-hop out-of-band control-plane network, since every node is directly connected with the network controller, new routes can be sent to the affected nodes to isolate the faulty node.

3.2.3 Summary of Simulation Results

Our performance profiling for CTP and its various variants shows that (i) the centralized control can further reduce the cost of the data collection tree compared with the original distributed control scheme (§3.2.1); (ii) the out-of-band SCDP network outperforms the in-band SCDP network in terms of network convergence as well as resilience to link noises and node faults (§3.2.2). These results suggest that the out-of-band SCDP scheme is a promising design for the MHWNs that have certain performance and resilience requirements. The rest of this chapter will study various aspects of the out-of-band SCDP scheme, which include its real implementation using LPWAN technologies, additional energy consumption, and the resulted data-plane network performance improvement under real-world settings.

3.3 LoRaWAN and TI's Sub-GHz Performance Profiling

This section profiles the energy consumption and latency of two LPWAN technologies, i.e., LoRaWAN and TI's sub-GHz, which are important to the design of the one-hop out-of-band control planes.

3.3.1 LoRaWAN Performance Profiling

This section profiles the performance of LoRaWAN using our LoRaCP hardware prototypes. The results are important to the software design of LoRaCP in §3.4.

3.3.1.1 LoRaCP hardware prototypes

Performance profiling is performed based on the following prototype hardware platforms. Each end device integrates a Cooking Hacks LoRaWAN shield and a Raspberry Pi (RPi) 3 Model B single-board computer. The shield has a Microchip RN2483 LoRaWAN chip, an 868 MHz antenna, and interfacing circuits. The shield can be controlled by the RPi using a C++ library from Cooking Hacks. The gateway integrates an RPi and an IMST iC880A LoRaWAN concentrator board [60]. The iC880A board can receive frames over all LoRa channels simultaneously.

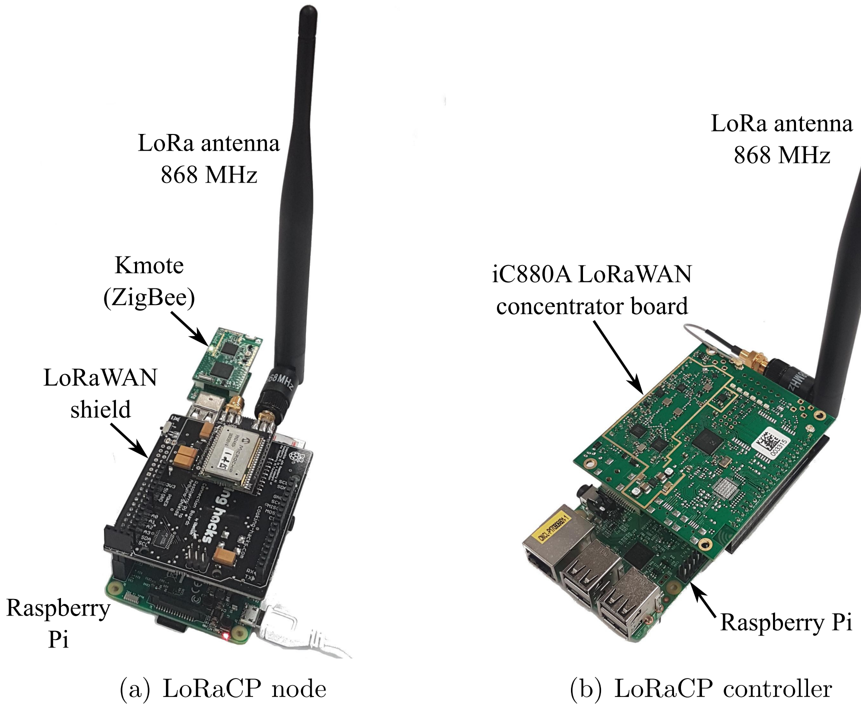

(a) LoRaCP node (b) LoRaCP controller

Figure 3.7: LoRaCP hardware prototypes. (The Raspberry Pi is for fast prototyping only; it will not be needed if LoRa is built into the MHWN platform.)

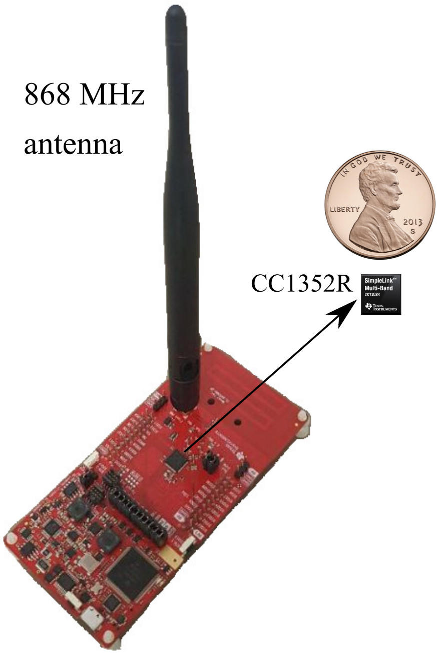

Figure 3.8: CC1352R launchpad used in our measurement study. The launchpad includes various peripherals that will not be needed in real deployments. The CC1352R consists of MCU, Zigbee and sub-GHz radios.

Table 3.2: Current consumption of a LoRaWAN 868 MHz module with bandwidth configured to 125 kHz.

SF	7	8	9	10	11	12
Transmitting current (mA)	39.741	39.861	39.495	39.694	40.441	40.413
Receiving current (mA)	14.196	14.208	14.227	14.240	14.234	14.240

A Zigbee-based Kmote is plugged into a USB port of the RPi of each end device, forming a *LoRaCP node*. The nodes use their Zigbee radios to form the data-plane network. In this thesis, we choose Zigbee because there is a body of MHWN protocols implemented for Zigbee. From now on, the gateway is referred to as *LoRaCP controller*. The controller unnecessarily has a Zigbee radio, since it may not be in the data-plane network. RPi is used to quickly prototype the integration of LoRaWAN and Zigbee. The results of this thesis will suggest that integrating LoRaWAN into the design of MHWN platforms, especially those desiring high network performance and manageability, is valuable. In such designs, the RPi will not be needed. Fig. 3.7 shows our prototypes.

3.3.1.2 Power and energy profiling

From RN2483's datasheet, its current consumption during transmitting, receiving, and sleeping modes with a supply voltage of 3.3 V are 38.9 mA, 14.2 mA, and 0.0016 mA, respectively [61, 62]. We use a Monsoon meter to measure the current supply of the whole LoRaWAN shield after properly jumping the power wires. Monsoon is a power meter that can deliver the real-time power consumption readings to a computer via a USB cable [63]. Table 3.2 shows the measurement results under different SFs. The results are close to RN2483's datasheet, showing that the shield's encapsulating and interfacing circuits consume little power.

A possible concern about LoRaWAN is its low Data rate to Power consumption Ratio (DPR), compared with other low-power radios. For instance, with SF7 in the EU868 MHz band, the DPR is 11 kbps/38.9 mA = 0.28 kbps/mA. In contrast, the DPR for Zigbee is 250 kbps/19.5 mA = 12.82 kbps/mA. However, the severity of this concern should be discriminated regarding the aimed communication range. This is illustrated by an example of moving x bits of data over a distance of L meters by multiple hops. The radio energy used to move the x bits over a hop is $(P_{Tx} + P_{Rx}) \cdot \frac{x}{v}$, where P_{Tx} and P_{Rx} are the transmitting and receiving

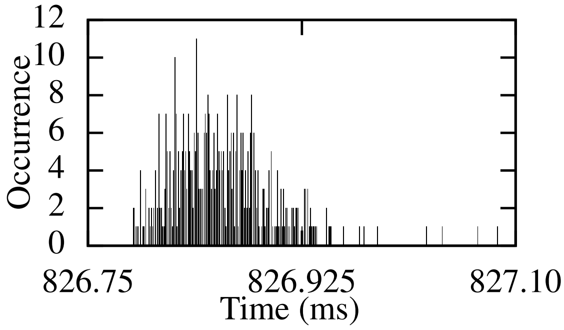

FIGURE 3.9: Radio awaking latency.

powers, respectively; v is the link data rate in bps. Thus, the total energy used by the network's radios to move x bits over L meters is $(P_{Tx} + P_{Rx}) \cdot \frac{x}{v} \cdot \frac{L}{d}$, where d is the typical one-hop transmission range. Considering $L = 1\,\text{km}$, we set $\frac{L}{d}$ to be 1 and 10 for LoRaWAN and Zigbee, respectively. Moreover, we set the data rate v to be 11 kbps and 250 kbps for LoRaWAN and Zigbee, respectively. After applying respective power consumption measurements, LoRaWAN's total radio energy consumption is 2.94 times of Zigbee's. Although the above simplistic energy consumption estimation does not consider other factors like nodes' processor energy consumption and MAC, the result underlines our understanding. While LoRaWAN consumes more energy than Zigbee, it substantially simplifies the control-plane network design due to its one-hop nature. Moreover, the concern of LoRaWAN's higher energy consumption can be mitigated by the fact that the control plane's traffic volume is much lower than the data plane's. For instance, as measured in §3.5, the number of CTP-SCDP's control-plane frames is just about 5% of its data-plane packets. Thus, we believe that, for the control-plane networks, the energy saving by using high-DPR but short-range radios is not worth sacrificing network simplicity.

3.3.1.3 Latency profiling

Under TDMA, the LoRaWAN radio can sleep to save energy while waiting for the next time slot. The time delays in awaking the radio and transmitting a frame are critical to the radio's sleep scheduling and clock synchronization required by TDMA, respectively. The latency in awaking the radio from the RPi using the shield's C++ API is measured. Fig. 3.9 shows the distribution of the awaking

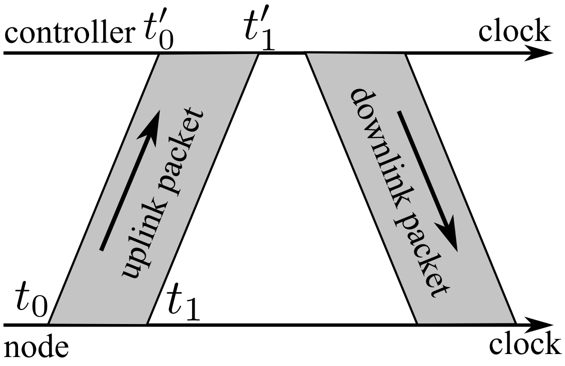

FIGURE 3.10: A communication session.

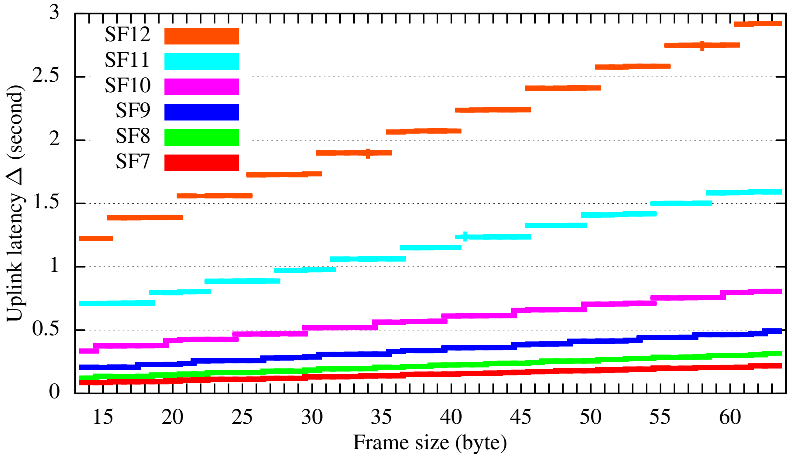

FIGURE 3.11: Uplink latency under different SFs and frame sizes.

latency over 500 tests that give satisfactory statistical significance of the measurement. The mean and standard deviation are 826.9 ms and 0.044 ms, respectively. The small standard deviation suggests that a LoRaCP node can awake the radio punctually for the next TDMA time slot.

Then, the latency in transmitting an uplink frame is measured. Fig. 3.10 illustrates the uplink transmission's timing. The node starts and completes the transmission when its clock values are t_0 and t_1, respectively. The controller starts and completes the reception when its clock values are t'_0 and t'_1, respectively. The t_0, t_1, and t'_1 can be recorded in the LoRaWAN shield's and concentrator's C++ user programs running at their RPis. To measure the uplink latency, the clocks of the node's and controller's RPis are synchronized using the Network Time Protocol (NTP) over an Ethernet that gives sub-ms synchronization accuracy. We define the uplink latency as $\Delta = t'_1 - t_0$.[2] Thus, the latency is determined by the data rate, which further depends on SF, and the frame size. Fig. 3.11 shows the box plots of the measured uplink latency under different SFs and frame sizes. As the latency has little variations under each setting, the boxes and whiskers of the plots are not visible. It can be seen that the latency increases with both frame size and SF, which are consistent with our understanding. Interestingly, for a certain SF, the latency exhibits step changes when the frame size increases. This is because each LoRa frame is a certain number of bits aligned for easy hardware handling. The above measurement results lay a foundation for developing LoRaWAN clock synchronization in §3.4.3.2.

3.3.1.4 Indoor communication profiling

Various existing studies [30, 49] have investigated LoRaWAN's communication performance in outdoor environments. However, many MHWNs are deployed in indoor environments, e.g., a WirelessHART network for a manufacturing system and a building environment monitoring WSN. Thus, we profile the communication performance of LoRaWAN in a large 6-story concrete building. Fig. 3.12 shows the exterior of the building. Along the building's long dimension of about 190 meters, the building has three sections (A, B, C) and two section junctions (J1 and J2). The LoRaCP controller is deployed in Section A1 on the third floor, as illustrated by the circle in Fig. 3.13. Then, a LoRaCP node is carried to different positions inside the building to measure the SNRs. Note that the Microchip RN2483 chip can estimate the SNR of the received signal. In each section, we measure three positions. The numbers in the cells of Fig. 3.13 are the SNR measurements in dB. It can be seen that the SNR measurements are from −19 dB to 6 dB. The SNR in general decreases with the distance between the LoRaCP node and the LoRaCP controller. Considerable SNR reductions are observed in the section junctions J1 and J2. This is because there are complex and steel structures in J1 and J2, such as stair cases, lifts, water facilities, and etc. For RN2483, the minimum SNRs required for reliable demodulation with SF7 to SF12 are −7.5 dB, −10 dB, −12.5 dB, −15 dB, −17.5 dB, and −20 dB, respectively [64]. Thus, in Fig. 3.13, a certain gray color is used to indicate the minimum SF setting that can cover each cell. It can

[2] We do not use t_1 because it contains a non-negligible uncertain delay from the actual completion of the transmission to the LoRaWAN shield's C++ library's periodic pull of the event from the shield's hardware interface.

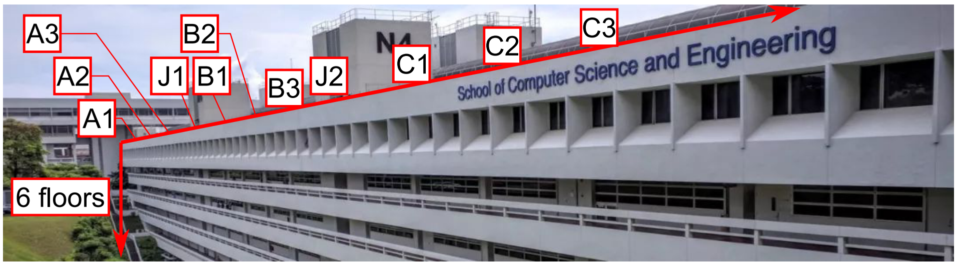

FIGURE 3.12: Exterior of the 6-story building.

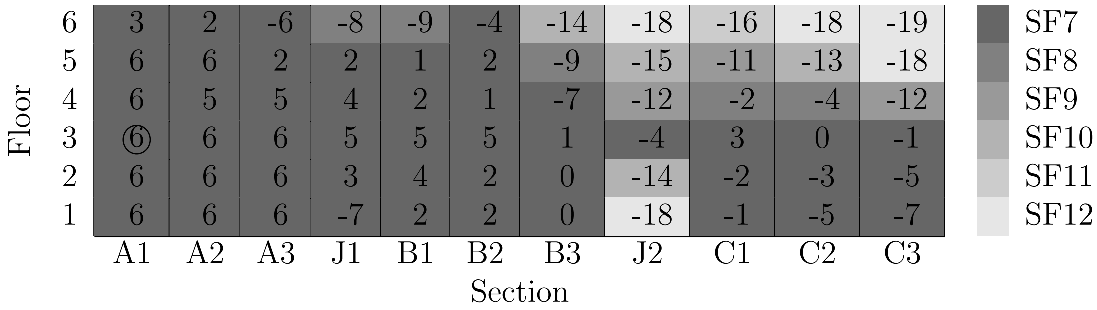

Floor	A1	A2	A3	J1	B1	B2	B3	J2	C1	C2	C3
6	3	2	-6	-8	-9	-4	-14	-18	-16	-18	-19
5	6	6	2	2	1	2	-9	-15	-11	-13	-18
4	6	5	5	4	2	1	-7	-12	-2	-4	-12
3	⑥	6	6	5	5	5	1	-4	3	0	-1
2	6	6	6	3	4	2	0	-14	-2	-3	-5
1	6	6	6	-7	2	2	0	-18	-1	-5	-7

Section

FIGURE 3.13: SNR measurements in a 6-story building and the coverage of LoRaWAN with different SF settings. The numbers in the cells are the SNRs in dB measured by the RN2483 of the LoRaCP node. The LoRaCP controller is located in the cell indicated by a circle. The dimension over the sections is about 190 meters. With SF12, the whole building can be covered.

be seen that most cells can be covered by a one-hop LoRaWAN with SF7. The whole building can be covered by a one-hop LoRaWAN with SF12. The above results show that although the indoor structures (e.g., concrete floors, walls, and steel objects) cause significant signal attenuation, with a large SF setting, a one-hop LoRaWAN network can well cover a building. For a larger building, multiple LoRaWAN gateways can be deployed to make sure every LoRaCP node in the building can communicate with the LoRaCP controller over a single wireless hop only.

3.3.1.5 Discussion

From the measurement results in [49], with line of sight in outdoor environments, LoRaWAN can achieve nearly 100% frame reception ratio when the distances are within 5 km. However, with obstruction, the frame reception will be severely affected. Thus, to apply the proposed one-hop out-of-band control plane design for outdoor MHWNs, extra care is needed to make sure each LoRaCP node can communicate with the LoRaCP controller reliably. Before system deployment, a remote sensing approach proposed in [51] can be used to predict the link quality and plan the LoRaCP node and LoRaWAN gateway positions. After system deployment, more LoRaWAN gateways can be added to fix coverage issues.

3.3.2 Performance Profiling of TI Sub-GHz Radio

In addition to LoRaWAN, this section investigates the feasibility of building one-hop out-of-band control planes using TI's SimpleLink sub-GHz long-range radio.

3.3.2.1 Introduction and characteristics of SimpleLink

SimpleLink [65] is a wireless MCU family made by TI. The SimpleLink covers a spectrum of wireless technologies including Wi-Fi, BLE, Thread, Zigbee, and sub-GHz radios. A SimpleLink platform also integrates an MCU to run user-defined programs. While TI provides a unified proprietary software development kit, the SimpleLink platforms can also run open-source operating systems such as Contiki-NG [66] and RIOT-OS [67]. Among various SimpleLink chips, the CC1350 and CC1352R are best suitable for building the one-hop out-of-band control planes. The CC1350, which has been adopted by the SensorTag platform, consists of a Cortex-M3 MCU, a BLE radio, and a sub-GHz long-range radio. The more recent CC1352R consists of a Cortex-M4F MCU, a sub-GHz long-range radio, and a 2.4GHz radio that supports multiple protocols including Zigbee and BLE. CC1352R has two operating modes for its sub-GHz radio: 2-GFSK mode and SimpleLink long-range mode. The 2-GFSK mode follows the IEEE 802.15.4g standard and adopts a binary Gaussian frequency shift keying modulation that can achieve 50 kbps and 200 kbps data rates with low and high output powers. The SimpleLink long-range mode uses the Direct-Sequence Spread Spectrum (DSSS) modulation that can achieve a maximum speed of 20 ksps (which is similar to LoRa) and a maximum communication range of 20 km. The form factor of the CC1352R is just 7 mm × 7 mm × 0.5 mm. This shows that a miniature platform such as CC1352R may meet the hardware requirements of our proposed one-hop out-of-band control plane design. Since the CC1352R supports Zigbee, it can leverage on the legacy programs of TinyOS and Contiki that are designed for Zigbee, e.g., CTP. Thus, the CC1352R is a promising platform for implementing our proposed one-hop out-of-band control plane design.

In what follows, the energy consumption and latency of CC1352R are profiled using two pre-production launchpads from TI. Fig. 3.8 shows a CC1352R launchpad.

3.3.2.2 Power profiling

From the datasheet, CC1352R's current draw is 24.3 mA at 3.7 V when using the maximum transmitting power (14 dBm). Thus, the transmission power consumption is $24.3 \times 3.7 = 89.91\,\text{mW}$. In the profiling, a voltage of 3.3 V is supplied to the CC1352R launchpad. We use the Monsoon power meter to measure the current draw of the launchpad. Due to the tight integration of the CC1352R chip and the launchpad, we cannot measure the sole power consumption of the CC1352R chip. Monsoon reads 125.8 mA and 161.7 mA when the launchpad is idle and transmitting using the SimpleLink mode, respectively. Thus, our estimate of CC1352R's power consumption in transmission is $(161.7\,\text{mA} - 125.8\,\text{mA}) \times 3.3\,\text{V} = 118.47\,\text{mW}$, which is about 20 mW higher than the datasheet value. A potential reason is that, when the CC1352R chip is in transmission, some supporting circuits on the launchpad are also active and consuming power. Using the same approach, the CC1352R's power consumption in the receiving mode is estimated as $11.1\,\text{mA} \times 3.3\,\text{V} = 36.63\,\text{mW}$. In our future work, the power profiling approaches adopted in [68, 69] can be used to understand CC1352R launchpad's power consumption better.

3.3.2.3 Timing performance profiling

This section profiles the timing performance of CC1352R's sub-GHz radio. Different from the method for measuring the latency of LoRaWAN radio in transmitting a frame, we measure the error in synchronizing the internal clocks of two CC1352R nodes using their sub-GHz radios. Specifically, the two CC1352R nodes are placed close to each other and let them perform a round-trip communication as in NTP. The standard deviation of the one-way transmission delay (i.e., half of the round-trip time) characterizes the error in synchronizing the two nodes. We implement the round-trip timing in TI's Real-Time Operating System (TI-RTOS). In TI-RTOS, the transmission and arrival of the frames are timestamped using Radio Timer (RAT), which is a functional unit on the radio core providing high-resolution timing. From our measurements, the one-way transmission delay has a mean of $532.59\,\mu\text{s}$ and a standard deviation of $2.50\,\mu\text{s}$. Thus, the clock synchronization

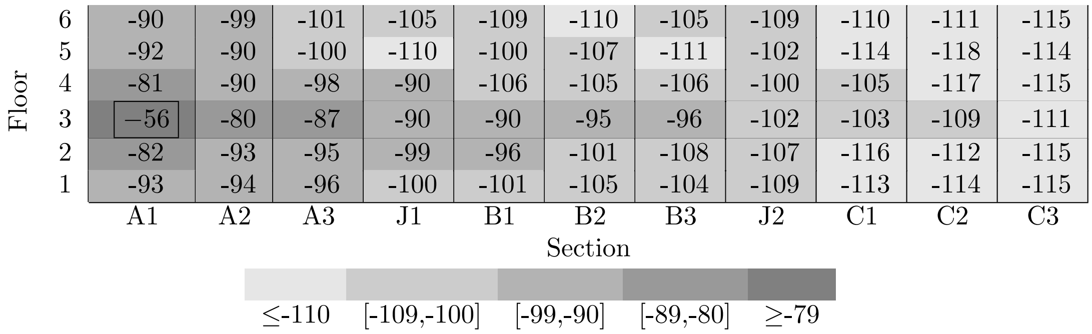

Floor	A1	A2	A3	J1	B1	B2	B3	J2	C1	C2	C3
6	-90	-99	-101	-105	-109	-110	-105	-109	-110	-111	-115
5	-92	-90	-100	-110	-100	-107	-111	-102	-114	-118	-114
4	-81	-90	-98	-90	-106	-105	-106	-100	-105	-117	-115
3	−56	-80	-87	-90	-90	-95	-96	-102	-103	-109	-111
2	-82	-93	-95	-99	-96	-101	-108	-107	-116	-112	-115
1	-93	-94	-96	-100	-101	-105	-104	-109	-113	-114	-115

FIGURE 3.14: RSSI measurements in a 6-story building. The number in a cell is the RSSI in dBm measured by a CC1352R in the cell. Another CC1352R is located in the cell indicated by a box. The dimension over the sections is about 190 meters.

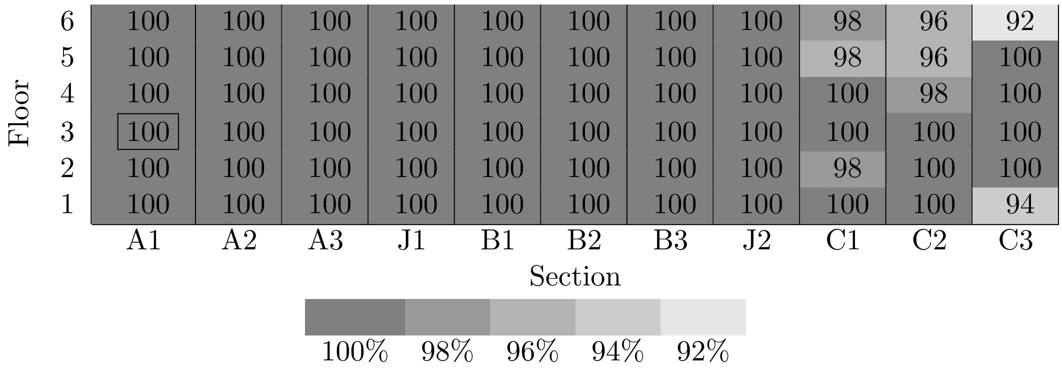

Floor	A1	A2	A3	J1	B1	B2	B3	J2	C1	C2	C3
6	100	100	100	100	100	100	100	100	98	96	92
5	100	100	100	100	100	100	100	100	98	96	100
4	100	100	100	100	100	100	100	100	100	98	100
3	100	100	100	100	100	100	100	100	100	100	100
2	100	100	100	100	100	100	100	100	98	100	100
1	100	100	100	100	100	100	100	100	100	100	94

FIGURE 3.15: FDR measurements in a 6-story building. The number in a cell is the FDR in percentage measured by a CC1352R in the cell. Another CC1352R is located in the cell indicated by a box. The dimension over the sections is about 190 meters.

using CC1352R's sub-GHz radio can achieve microseconds accuracy. Note that as measured in §3.3.1.3, LoRaWAN's frame transmission timing has millisecond-level uncertainty, which results in the milliseconds clock synchronization error as shown in §3.4.3.2. The higher synchronization accuracy achieved by CC1352R is due to its hardware-level timestamping capability provided by RAT. In contrast, the LoRaWAN frame timestamping performed on the RPi is subjected to the overhead of the Raspbian operating system. These results show that CC1352R has good performance in timing the transmissions of frames, which is desirable for implementing the TMDA-based control planes.

3.3.2.4 Indoor communication profiling

To profile the indoor communication performance, a CC1352R node is fixed at a location and another CC1352R node is moved in the 6-story building mentioned in §3.3.1.4. Fig. 3.14 and Fig. 3.15 show the Received Signal Strength Indicator (RSSI) and Frame Delivery Ratio (FDR) measurements when the mobile CC1352R node is in different locations in the building, respectively. In this set of measurements, both nodes adopt a transmission power of 14 dBm, same as the setting used for LoRaWAN. It can be seen from the results that the RSSI attenuates with distance. The CC1352R can achieve satisfactory FDRs throughout the building. These results show that a one-hop SimpleLink network can well cover a building.

3.3.2.5 Discussion

From the above results, the power, latency, and indoor communication performance profiles of CC1352R's SimpleLink long-range radio are similar to LoRaWAN's. Moreover, CC1352R has a built-in 2.4GHz radio that supports Zigbee. Thus, CC1352R is a promising platform for implementing the proposed one-hop out-of-band control plane design. As CC1352R is still in the pre-production phase, the first part of this thesis focuses on using the LoRaCP hardware prototype presented in §3.3.1.1 to implement and evaluate our proposed approach. The implementation of the out-of-band control planes using CC1352R will be similar.

3.4 Design and Implementation of LoRaCP

3.4.1 System Overview and LoRaCP-TxC

The goal of LoRaCP is to use LoRaWAN's uplinks and downlinks to transmit *network reports* from the nodes to the controller and *network commands* from the controller to the nodes, respectively. The network commands have two categories: a *reactive network command* to a node is in response to a precedent network report from the node, whereas an *active network command* is initiated by the controller. All the control-plane transmissions are managed by LoRaCP's transmission control protocol as illustrated in Fig. 3.16, which is called *LoRaCP-TxC*. As discussed

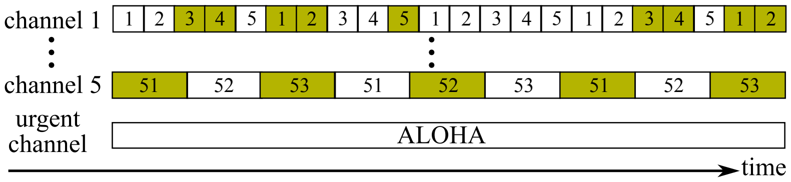

FIGURE 3.16: Illustration of LoRaCP-TxC (shaded blocks mean heartbeat slots).

in §2.3, LoRaWAN has six concurrent uplink channels. Five of them use TDMA, while the remaining one (called *urgent channel*) uses ALOHA to transmit urgent frames. The five concurrent TDMA channels increase the throughput for the network reports. The urgent channel mitigates the rigidness of TDMA and allows the control-plane application developers to deal with urgent situations such as sudden strong interference or even malicious jamming to the data-plane network. As the TDMA channels have different data rates, their time slot lengths can be different to achieve the same maximum frame size. The time slots of a TDMA channel are allocated in a round-robin fashion to the LoRaCP nodes that use the channel. The LoRaCP nodes can be assigned to the TDMA channels to balance their time delays in waiting for the next time slot, while considering the channels' communication ranges and the nodes' distances to the controller.

Certain regions impose duty cycle requirements on the sub-GHz ISM bands used by LoRaWAN. For instance, in Europe, a LoRaWAN end device operating in the EU868 MHz band needs to conform to a duty cycle requirement specified by The European Telecommunications Standards Institute [70]. The duty cycle upper limit can be 0.1%, 1% or 10% depending on the used sub-band. Other regions impose dwell time, i.e., the channel occupancy time. For instance, in North America, the dwell time of the LoRaWAN operating on a 125 kHz bandwidth centered at any frequency must not be longer than 0.4 seconds within any 20-second period [71]. Note that a dwell time requirement can be translated to a duty cycle requirement. For instance, the dwell time requirement mentioned earlier can be translated to a 2% duty cycle requirement. To meet the duty cycle requirement, each node has a minimum waiting time between two consecutive transmissions. Accordingly, the time slot lengths of the channels in LoRaCP-TxC can be designed to meet the duty cycle requirement. Specifically, the time slot length for a channel should satisfy

$$\text{time slot length} \geq \max\left\{\frac{\text{a node's minimum waiting time}}{\text{the number of nodes assigned to the channel}}, 3 \text{ seconds}\right\},$$

where the 3-second time is the minimum time for completing a communication session consisting of an uplink and two optional downlink transmissions. Note that a node will open two continuous 1-second windows waiting for the downlink transmission. The remaining one second in a communication session is for the uplink transmission and data processing. The minimum waiting time can be derived based on the LoRa's on air time calculation method described in [72]. For instance, for the channel using SF7, bandwidth of 125 kHz and coding rate of 4/5, if the payload size is 30 bytes and the duty cycle requirement is 1%, the minimum waiting time is 8.73 seconds. If there are 5 LoRaCP nodes in this channel, the time slot length can be set as $\max\left\{\frac{8.73}{5}, 3\right\} = 3$ seconds.

We now present two features of LoRaCP-TxC that address LoRaWAN's downlink-uplink asymmetry and lossy links.

- **Heartbeat time slots:** When a node has no uplink data to transmit, it can skip its next time slot. However, because any downlink frame must be in response to a precedent uplink frame, LoRaCP-TxC designates periodic heartbeat time slots for each node. For instance, in Fig. 3.16, the shaded blocks represent heartbeat slots. In Channel 1, the heartbeat period is three time slots. A node must transmit an uplink frame in a heartbeat slot. This open a downlink window to maintain the clock synchronization of the node (cf. §3.4.3.2) and send active network commands. The heartbeat period can be set according to the nodes' clock drift rates and the required clock accuracy to avoid TDMA panic. The heartbeats also help the LoRaCP controller be aware of whether a node is still alive.

- **Negative acknowledgment (NAK)**: To deal with frame losses, acknowledging all concurrent uplink transmissions is wasteful because of the downlink-uplink asymmetry. Thus, LoRaCP uses the NAK scheme. In LoRaWAN, the uplink and downlink frames from/to a node have continuously increasing counters, respectively. Thus, both the controller and the nodes can detect if there are lost frames by checking the continuity of the frame counters. If the controller detects lost frames, it sends an NAK using the subsequent downlink transmission to notify the node, which can then use the urgent channel or wait for the next TDMA slot to resend the lost data. The node can also send NAK using the urgent channel or the next TDMA slot to request lost frames. With the NAK scheme, the controller does not need to respond to a node's network report if there are no network commands for the node and no lost frames. This design mitigates the contention for the downlink time.

3.4.2 Software Architectures of LoRaCP Node and Controller

In §3.3.1, we have introduced the hardware prototypes of the LoRaCP node and controller. This section presents their software architectures as illustrated in Fig. 3.17. Note that the software of LoRaCP, including the transmission control protocol LoRaCP-TxC, is implemented in the application layer using the program library of the used LoRaWAN platform. This means that we do not need to modify the LoRaWAN source code. The LoRaWAN's ALOHA MAC is beneath the program library.

3.4.2.1 LoRaCP node

A C++ forwarder program *LoRaCPFwd* runs on the RPi to buffer and forward the data between Kmote and the LoRaWAN shield, while following LoRaCP-TxC. The node parts of the clock synchronization and TDMA are also implemented in *LoRaCPFwd*. The Kmote runs TinyOS. We design a TinyOS module *LoRaCPC* that provides the *AMSend* and *Receive* interfaces to send and receive data to/from the RPi through serial communications. Thus, in our prototype design, the Kmote uses LoRaWAN as a service.

3.4.2.2 LoRaCP controller

The RPi of the controller runs an open-source LoRaWAN server architecture [73] consisting of *packet_forwarder*, *LoRa Gateway Bridge*, *LoRa Server*, and *LoRa App Server*. This architecture, through providing JSON-based APIs to subscribe/send messages from/to the LoRaWAN network, greatly simplifies the design of centralized network control applications. The role of this architecture is similar to that of an SDN controller platform (e.g., OpenDaylight) that facilitates the design of SDN control applications. In the first part of this thesis, the centralized network

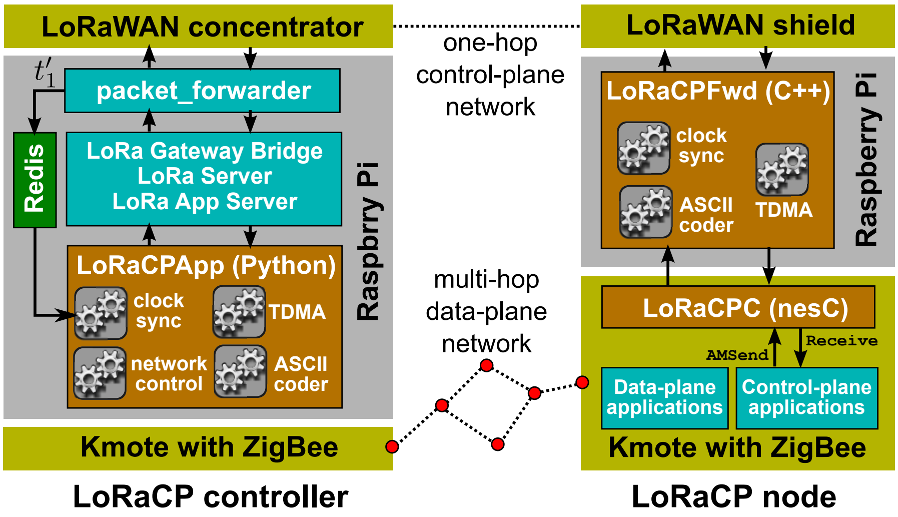

FIGURE 3.17: Software architectures of LoRaCP controller and node. (The illustration includes a Zigbee radio for the controller to be a control-plane sink.)

controls and the controller parts of the clock synchronization and TDMA are implemented in a single Python program called *LoRaCPApp*. Note that the LoRaWAN server architecture [73] supports multiple LoRaWAN gateways. Although the first part of this thesis focuses on a single LoRaCP controller, in the future work, the multi-gateway support can be exploited to develop redundant LoRaCP controllers to improve the system's reliability against a single point of failure.

3.4.3 Implementation of LoRaCP Components

This section provides implementation details of LoRaCP's integer coding scheme, clock synchronization, and TDMA.

3.4.3.1 ASCII coding of integers

In this work, we aim to develop our system using the APIs provided by Cooking Hacks. The LoRaWAN library provided by the manufacture applies a preprocessing algorithm on payload. We design an ASCII coder to bypass this scheme and improve efficiency. ASCII code words are used to directly encode integers for better efficiency. The use of the ASCII code words is because LoRa admits hexadecimal ASCII string only. Since there are 128 ASCII code words, the last one is reserved for separator and the remaining 127 code words to encode an integer.

Specifically, the decimal integer is presented in the base-127 numeral system. For instance, $10000 = 78 \times 127^1 + 94 \times 127^0 = (78, 94)_{127}$. Then, the hexadecimal ASCII string consisting of the digits of the base-127 representation is fed to LoRa. For the above example, the string is "4E5E", much shorter than the "3130303030" in LoRa's default coding. The decoding is simply the reverse process. The above ASCII coder is implemented in *LoRaCPFwd* and *LoRaCPApp* at the LoRaCP node and controller, respectively.

3.4.3.2 Clock synchronization

Clock synchronization is a basis for implementing TDMA. Although there are various existing clock synchronization protocols for MHWNs (e.g., FTSP), if the LoRaCP nodes are synchronized to the controller using the data-plane network, the control plane's TDMA will depend on the data-plane network, incurring the undesirable coupling. Thus, the LoRaCP nodes should be synchronized to the controller using the control-plane network. However, there is still limited research on clock synchronization over LoRaWAN. In our prototype system, the RPi's clock is used as the node's or controller's clock. Although the LoRaWAN devices and the Kmote have their own timers, using the RPi's clock can simplify the evaluation of the accuracy of the LoRaWAN clock synchronization using the RPi's Ethernet interface.

To save the downlink time, LoRaCP does not prescribe dedicated frames for clock synchronization. Instead, LoRaCP piggybacks several bytes to each control-plane frame for clock synchronization. Specifically, each uplink frame is appended with the node's clock value t_0 as illustrated in Fig. 3.10. The controller records its clock value t_1' on completion of the frame reception. The clock offset between the node and the controller, denoted by δ, can be estimated as $\delta = t_1' - (t_0 + \Delta)$, where Δ is the uplink latency presented in Fig. 3.11. Then, the controller piggybacks δ onto the downlink frame as illustrated in Fig. 3.10. Upon receiving δ, the node resets its clock by $t = t + \delta$, where t denotes the node's current clock value. Alternatively, the node's clock advance speed can be calibrated according to δ using a negative feedback loop.

We now discuss several implementation issues of the above clock synchronization approach. First, the LoRaWAN frame header added by the shield has changeable size because the integers in the headers are represented as variable length hexadecimal ASCII strings (cf. §2.3.2). As shown in Fig. 3.11, the uplink latency Δ has a complex relationship with the frame size in different channels. When the LoRaCP controller receives the uplink frame, it checks the actual frame size and the SF used by the node to query the corresponding Δ from the data in Fig. 3.11. Thus, for LoRaWAN clock synchronization, the prior knowledge in Fig. 3.11 is critical. Note that most MHWN clock synchronization approaches are free from this frame size dependence issue because they use dedicated synchronization frames with fixed sizes or the frame size has little impact on transmission latency. Second, we modify *packet_forwarder*, i.e., LoRaWAN concentrators' driver program, to record t'_1, because other components of the LoRaWAN server architecture may suffer software delays. As illustrated in Fig. 3.17, the timestamp t'_1, together with the corresponding source ID and frame ID, are written into a Redis in-memory database and then retrieved by the *LoRaCPApp* to compute δ.

The synchronization accuracy of the above approach is measured using the `ntpdate` tool to check the clock offset between the node and the controller over a local Ethernet network connecting the RPis. The mean absolute synchronization error is 2.9 ms with a standard deviation of 1.7 ms. Given the second-level frame transmission time, such synchronization errors of a few milliseconds are satisfactory.

3.4.3.3 TDMA

The prototype LoRaCP node controls the sleep of the LoRaWAN radio and transmissions of frames based on its RPi's synchronized clock. Specifically, if *LoRaCPFwd* has received a network report from the Kmote, the RPi starts awaking the LoRaWAN radio 850 ms before its next TDMA time slot, transmits the report in the time slot, receives any subsequent network command, re-transmits frames using the urgent channel if an NAK is received. Finally, *LoRaCPFwd* forwards all received network commands to the Kmote. In our current experimental implementation, LoRaWAN channels and time slots are assigned to nodes manually.

3.5 Performance Evaluation

Various testbed experiments are conducted to evaluate our LoRaCP implementation.

3.5.1 Experiment Methodology and Settings

LoRaCP is applied to implement the out-of-band CTP-SCDP presented in §3.2. Specifically, if the Kmote of a LoRaCP node detects a change of ETX with any of its neighbor node, it uses the *LoRaCPC* to send the latest ETX using a network report frame to the LoRaCP controller. Upon receiving the ETX update, the controller's *LoRaCPApp* python program computes the optimal routing and pushes network commands containing new parent node information to the downlink queue of the LoRaWAN server architecture. Upon receiving a network command, a LoRaCP node updates its parent node accordingly. In the data plane, each node generates a data packet every eight seconds.

We conduct experiments on a testbed consisting of a LoRaCP controller and 15 LoRaCP nodes. The LoRa modules use the 868 MHz ISM band. The nodes are placed at the grid points of a lab space. The nodes are evenly divided to use three LoRaWAN channels (SF7, SF8, and SF9). The time slot lengths in these three channels are 3, 4, and 5 seconds, respectively.[3] The controller uses the first downlink window RX1 to transmit network commands. Before the RX1 window, the controller has a *wait time* of one second to compute the network commands, which is generally sufficient. On our 16-node testbed, each LoRaCP has a time slot every 25 seconds or less. For larger networks, to maintain this rotating period for each node, multiple geographically distributed nodes in the same channel can be assigned to use the same time slot, since they unlikely report ETX changes at the same time.

[3] The duty cycles in the three channels are 0.9%, 1.4%, 2.1%. Note that the region in which this thesis's experiments are conducted does not impose duty cycle requirement on the 868 MHz ISM band [74]; it only imposes a transmitting power upper limit. To meet Europe's 1% duty cycle requirement, the approach presented in §3.4.1 can be used to configure the time slot lengths.

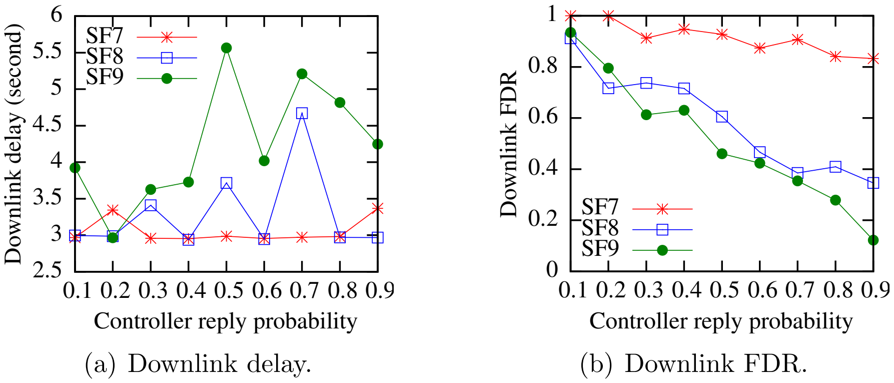

(a) Downlink delay.

(b) Downlink FDR.

Figure 3.18: Control plane communication performance.

3.5.2 Experiment Results

We conduct three sets of experiments: §3.5.2.1 evaluates the control plane communication performance; §3.5.2.2 evaluates the control plane performance of CTP-SCDP; §3.5.2.3 compares CTP and CTP-SCDP.

3.5.2.1 Control plane communication performance

While the concurrent uplink channels increase the throughput for network reports, LoRaWAN's downlink-uplink asymmetry presents a bottleneck for the downlink communications. The downlink performance is evaluated. Specifically, each LoRaCP node transmits a network report every its time slot. Thus, the controller receives frames from the three channels concurrently almost at all the time. It replies to each network report with a certain probability. The frame size of the replies ranges from 29 to 33 bytes. NAK is turned off in these tests.

Fig. 3.18 shows the average control-plane downlink delays and FDRs of different channels versus the probability that the controller replies. The downlink delay is measured as the time duration between i) the controller's *LoRaCPApp* pushes a network command to the LoRaWAN server architecture and ii) the node's *LoRaCPFwd* receives the command. This downlink delay includes the wait time of one second. From Fig. 3.18(a), the average downlink delay does not significantly increase with the controller's reply probability. The average delay ranges from 3 seconds to 5.5 seconds. It increases with the SF, because a larger SF has a lower

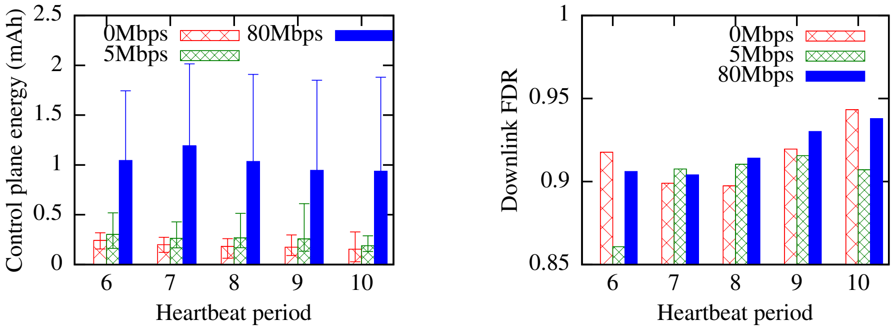

(a) Per-node energy consumption for the control plane in one hour.

(b) Downlink FDR.

FIGURE 3.19: CTP-SCDP control plane performance under Wi-Fi interference against the data-plane network. The error bar represents min and max values.

data rate. Fig. 3.18(b) shows the control-plane downlink FDR versus the controller's reply probability. The FDR decreases with the reply probability. This is because the open-source LoRaWAN server architecture [73] drops frames when it receives excessive frames to be transmitted beyond the downlink throughput. From the results in Fig. 3.18, the downlink bottleneck mainly affects the downlink FDR. Thus, in the remaining experiments, the downlink FDR is used to assess whether the control plane performance is throttled by the downlink-uplink asymmetry.

3.5.2.2 Control plane performance in CTP-SCDP

The performance of CTP-SCDP's control-plane network is evaluated. To create data-plane link quality variations, a laptop placed close to the testbed is used to generate Wi-Fi traffic to interfere with the Zigbee data-plane network. Zigbee radios use Channel 18 and the Wi-Fi AP uses Channel 6, which interfere with each other. On the laptop, `iperf3` is used to generate data traffic at a specified bit rate. This experiment methodology well captures the increasingly crowded 2.4 GHz ISM band used by the Zigbee-/BLE-based data-plane networks. In the presence of the Wi-Fi interference, the CTP-SCDP generates more control-plane messages to report the volatile link ETXes of the data-plane network to the LoRaCP controller.

First, we estimate the energy consumption of each LoRaCP node's LoRaWAN shield by multiplying the transmitting/receiving currents with the measured total times in respective modes. Fig. 3.19(a) shows the error bars of per-node energy consumption by the shield in one hour under different settings of heartbeat period

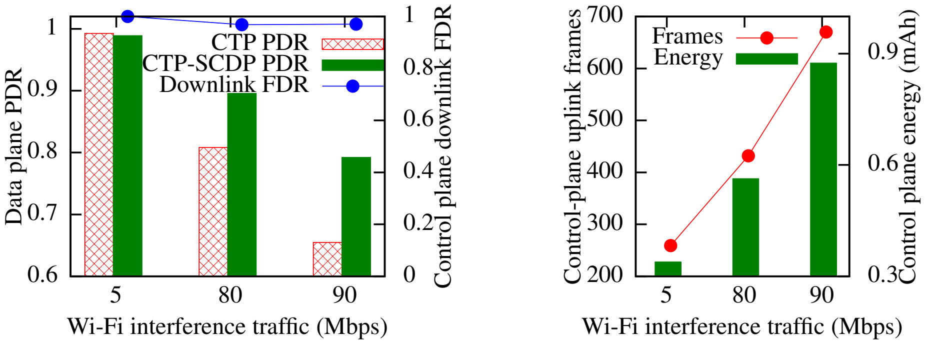

(a) Data-plane PDR and control-plane downlink FDR.

(b) Control-plane uplink frames and per-node energy in one hour.

FIGURE 3.20: Performance comparison between CTP and CTP-SCDP.

and Wi-Fi interference intensity. The control-plane energy consumption increases with the interference intensity due to the increased control-plane messages. When we do not generate Wi-Fi interference, the energy consumption decreases with the heartbeat period. This is because, in the absence of the interference, the link ETXes seldom change and most control-plane messages are the heartbeats. In the presence of interference (i.e., 5 Mbps and 80 Mbps), the energy consumption has no monotonic relationship with the heartbeat period, because the node will utilize the non-heartbeat time slots to report the volatile ETXes. From Fig. 3.19(a), with no and intensive interference (80 Mbps), the per-node power consumption by the control plane averaged over time is about 0.25 mA and 1 mA, respectively, which are comparable to or lower than the power consumption of low-power microcontrollers (MCUs). For instance, the active power of TelosB's MCU is 1.8 mA, whereas the recent Firestorm's MCU consumes 8.6 mA in the common configuration [10].

Second, we measure the average control-plane downlink FDR over all channels. The results are shown in Fig. 3.19(b). Even if the data-plane network experiences intensive interference, the FDR is generally above 90%. Thus, the CTP-SCDP's control plane is still beyond the downlink bottleneck.

3.5.2.3 Comparison between CTP and CTP-SCDP

We load CTP to eight nodes and CTP-SCDP to another eight nodes. CTP and CTP-SCDP run side by side on the testbed, so that they experience almost the same Wi-Fi interference for fair comparison. CTP-SCDP's LoRaCP heartbeat period is 10. Fig. 3.20(a) shows the data plane's PDR, i.e., the ratio of the Zigbee

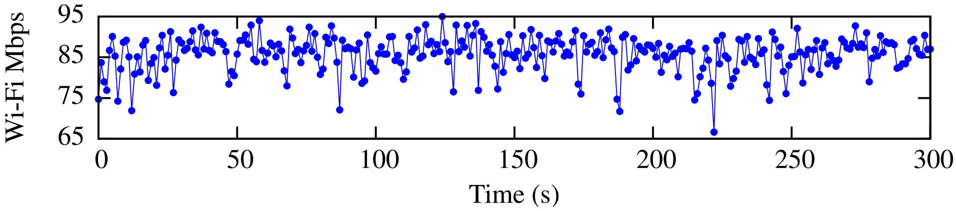

FIGURE 3.21: Wi-Fi data rate fluctuations over time (setpoint: 90 Mbps).

packets received by the data-plane sink over all packets generated by the source nodes. When the Wi-Fi interference intensity is low (5 Mbps), CTP and CTP-SCDP achieve similarly high PDRs. When the interference intensity is 80 Mbps, CTP-SCDP's PDR is 10% higher than CTP's. When the interference intensity is 90 Mbps, CTP's PDR drops to 65%, while CTP-SCDP's is 80%. Note that the actual data rate of the Wi-Fi interference traffic fluctuates over time. Fig. 3.21 shows the actual interference data rate when the setpoint to `iperf3` is 90 Mbps. Moreover, the fluctuation level increases with the setpoint. The data rate deviations are 0.8 Mbps only and up to 20 Mbps for setpoints 5 Mbps and 90 Mbps, respectively. Thus, the control-plane networks experience more dynamic interference with a larger setpoint, resulting in the increasing PDR gain of CTP-SCDP over CTP with the interference intensity setting. This result is consistent with our observation from the simulation study in §3.2.1 that CTP cannot handle dynamic network conditions well.

Fig. 3.20(a) also shows the control-plane downlink FDRs, which are above 97%. This suggests that the control plane is beyond the downlink bottleneck. Fig. 3.20(b) shows the total number of control-plane uplink frames of CTP-SCDP during one hour and the projected per-node energy consumption by the LoRaWAN shield. In the presence of stronger interference, more uplink frames will be transmitted to report the volatile ETXes. With 5 Mbps and 90 Mbps interference, the total numbers of data-plane transmissions (including beacons and forwarded packets) are 5,022 and 10,024, respectively. The corresponding numbers of control-plane uplink frames are just 5.2% and 6.7% of these data-plane transmissions. With strong interference (90 Mbps), the per-node control-plane power consumption averaged over time is less than 0.9 mA, consistent with the results in Fig. 3.19 obtained with 15 nodes.

Chapter 4

Attack-Aware Data Timestamping in Low-Power Synchronization-Free LoRaWAN

This chapter is organized as follows[1]. §4.1 reviews related work; §4.2 describes sync-free data timestamping; §4.3 studies the attack; §4.4 presents LoRaTS and uplink frame arrival time detection approach; §4.5 studies LoRa's FB and uses it to detect attack; §4.6 presents experiment results; §4.7 discusses several issues.

4.1 Related Work

Improving LoRaWAN's communication performance has received increasing research. The Choir system [52] exploits the diverse FBs of the LoRaWAN end devices to disentangle colliding frames from different end devices. Choir uses the dechirping and Fourier transform processing pipeline to analyze FB, which does not provide sufficient resolution for detecting the tiny extra FB introduced by attack (see details in §4.5.1). In the second part of this thesis, based on an analytic model of LoRa's CSS modulation, we develop a new time-domain signal processing algorithm based on a least squares formulation to achieve the required resolution. The Charm system [76] exploits *coherent combining* to decode a frame from the

[1] This chapter is partially published on [75].

weak signals received by multiple geographically distributed LoRaWAN gateways. It allows the LoRaWAN end device to use a lower transmitting power. Several recent studies [55, 77] have devised various backscatter designs for LoRa to reduce the power consumption of end devices. All the studies mentioned above focus on understanding and improving the data communication performance of LoRaWAN [52, 76], or reducing power consumption via backscattering [55, 77]. None of them specifically addresses efficient data timestamping, which is a basic system function of many LoRaWAN-based systems.

LongShoT [78] is an approach to synchronize the LoRaWAN end devices with the gateway. Through low-level offline time profiling for a LoRaWAN radio chip (e.g., to measure the time delays between hardware interrupts and the chip's power consumption rise), LongShoT achieves sub-50 microseconds accuracy, which is echoed by our results on the accuracy of estimating signal arrival time using a different approach. LongShoT is designed for the LoRaWAN systems requiring tight clock synchronization. Differently, we address data timestamping and focus on the less stringent but more commonly seen milliseconds or sub-second accuracy requirements. Our sync-free approach releases the bandwidth from frequent clock synchronization operations.

Security of LoRaWAN is receiving research attention. In [79], Aras et al. discuss several possible attacks against LoRaWAN, including key compromise and jamming. The key compromise requires prior physical attack of memory extraction. In [80], a selective jamming attack against certain receivers and/or certain application frames is studied. Different from the studies [79, 80] that do not consider the stealthiness of jamming, we consider stealthy frame collision. From our results in §4.3.2, the selective jamming in [80] cannot be stealthy because it cannot start jamming until the frame header is decoded and the corruption of payload must lead to integrity check failures. In [81], Robyns et al. apply supervised machine learning for end device classification based on the received LoRa signal. From our measurements, the dissimilarity between the original and the replayed signals is much lower than that among the original signals from different end devices. Thus, the supervised machine learning is not promising for attack detection.

Device identification based on radiometric features has been studied for short-range wireless technologies. A radiometric feature is the difference between the nominal and the measured values of a certain modulation parameter. The work

[82] studied the radiometric features of IEEE 802.11 radios, including symbol-level features regarding signal magnitude and phase, as well as the frame-level feature regarding carrier frequency. In LoRaWAN, the received signal strength is often rather low due to long-distance propagation or barrier penetration. As such, the signal magnitude radiometric feature cannot be used as a radiometric feature. As the phase of LoRa signal is arbitrary, it cannot be employed as a radiometric feature too. In the second part of this thesis, we show that the bias of the LoRa signal's carrier frequency from the nominal value is an effective radiometric feature. This feature can be used to counteract the frame delay attack. Based on LoRa's CSS modulation, we develop a lightweight algorithm that can estimate this feature from the received LoRa signal. It requires a low-cost SDR receiver, unlike the expensive vector signal analyzer [83] used in the work [82].

Our prior work [75] studied the efficiency and security of sync-free data timestamping in LoRaWAN. It presents the initial design of LoRaTS to minitor the bias of the uplink carrier frequency. Based on [75], we make the following three extensions. First, we extend §4.3.2 to present the detailed experiment results on LoRa radio's vulnerable time window to frame collision. Second, in §4.4.3, we model the LoRa signal reception process and investigate the algorithms that can accurately estimate the uplink frame arrival time. Third, in §4.6.3, we present experiment results regarding the impacts of ambient temperature and employment of temperature compensated crystal oscillators for end devices on LoRaTS's attack detection performance.

4.2 Advantages of Sync-Free Timestamping

Data timestamping, i.e., to record the *time of interest* in terms of the wall clock, is a basic system function required by the data collection applications for monitoring. For a sensor measurement, the time of interest is the time instant when the measurement is taken by the end device. Multi-hop WSNs largely adopt the *sync-based* approach. Specifically, the clocks of the WSN nodes are synchronized to the global time using some clock synchronization protocol. Then, each WSN node can timestamp the data using its local clock. WSNs have to adopt this approach due primarily to that the multi-hop data deliveries from the WSN nodes to the gateway in general suffer uncertain delays. Thus, although the clock synchronization introduces additional complexity to the system implementation, it has become a standard component for systems requiring data timestamping. However, the clock synchronization introduces considerable communication overhead to the bandwidth-limited LoRaWANs.

We present an example to illustrate the overhead to maintain sub-10 milliseconds (ms) clock accuracy in LoRaWANs. Typical crystal oscillators in microcontrollers have drift rates of 30 to 50 ppm [84]. Without loss of generality, we adopt 40 ppm for this example. With this drift rate, an end device needs 14 synchronization sessions per hour to maintain sub-10 ms clock accuracy. These 14 sessions represent a significant communication overhead for an end device. For instance, in Europe, a LoRaWAN end device adopting a spreading factor of 12 can only send 24 30-byte frames per hour to conform to the 1% duty cycle requirement [70]. Although the synchronization information may be piggybacked to the data frames, a low-rate monitoring application may have to send the frames more frequently just to keep time. In addition, the data frames need to include data timestamps, each of which needs at least a few bytes. This is also an overhead given the bandwidth scarcity.

To efficiently utilize LoRaWAN's scarce bandwidth and exploit its star topology, the sync-free timestamping approach can be adopted. In this approach, an end device transmits a sensor reading once generated. Upon receiving the frame, the gateway uses the frame arrival time as the data timestamp. The signal propagation time from the end device to the gateway, which is often microseconds, can be ignored for millisecond-accurate timestamping. Compared with the sync-based approach, this sync-free approach avoids the communication overhead caused by the frequent clock synchronization operations and the transmissions of timestamps. Thus, the sync-free approach is simple and provides bandwidth-saving benefit throughout the lifetime of the LoRaWANs.

4.3 Security of Sync-Free Timestamping

The long-range communication capability of LoRaWAN enables the less complex and bandwidth-efficient sync-free timestamping. However, it may also be subject to wireless attacks that can affect large geographic areas. Having understood the

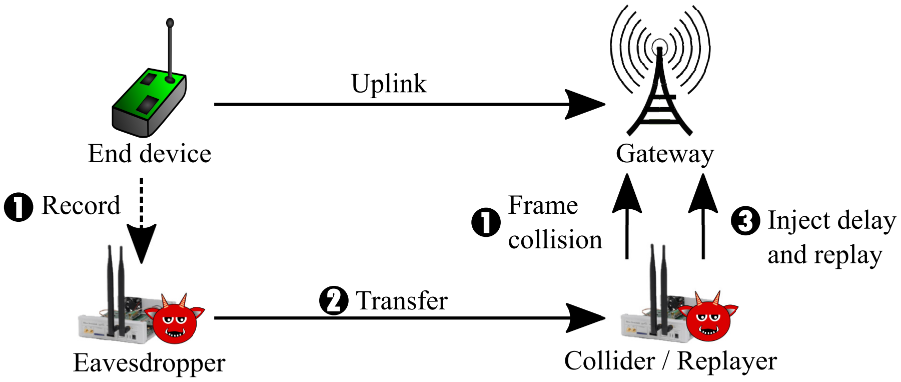

FIGURE 4.1: Steps for implementing frame delay attack.

benefit of sync-free timestamping, we also need to understand its security risk and the related countermeasure for achieving a more comprehensive assessment on the efficiency-security tradeoff. A major and direct threat against the sync-free approach is the *frame delay attack* that manipulates the frame delivery time to invalidate the assumption of near-zero signal propagation delay. We define the attack as follows.

Frame delay attack: The end device and gateway are not corrupted by the adversary. However, the adversary may delay the deliveries of the uplink frames. The malicious delay for any uplink frame is finite. Moreover, the frame cannot be tampered with because of cryptographic protection.

The attack results in wrong timestamps under the sync-free approach. This section studies the attack implementation (§4.3.1), investigates the timing of malicious frame collision (§4.3.2), and studies the size of the vulnerable area in which the end devices are affected by the attack (§4.3.3).

4.3.1 Attack Implementation

4.3.1.1 Implementation steps

Fig. 4.1 illustrates the attack implementation. The adversary sets up two malicious devices called *eavesdropper* and *collider* that are close to the end device and the gateway, respectively. The attack consists of three steps. ❶ At the beginning, both the eavesdropper and the collider listen to the LoRa communication channel between the end device and the gateway. Once the collider detects an uplink frame transmission, it transmits a collision frame. In §4.3.2, we will investigate experimentally a stealthy collision method such that the victim gateway does not raise any warning message to the application layer. Meanwhile, once the eavesdropper detects an uplink frame transmission, it records the radio waveform of the frame. Note that the collider may choose a proper transmitting power of the collision frame such that the collision can affect the victim gateway, while not corrupting the radio waveform recorded by the eavesdropper. ❷ The eavesdropper sends the recorded radio waveform data to the collider via a separate communication link that provides enough bandwidth. ❸ After a time duration of τ seconds from the onset time of the victim frame transmission, the collider replays the recorded radio waveform. Thus, in the second part of this thesis, the *collider* and the *replayer* refer to the same attack device. The above collision-and-replay process does not need to decipher the payload of the recorded frame; it simply re-transmits the recorded radio waveform. As the gateway cannot receive the original frame and the integrity of the replayed frame is preserved, the gateway accepts the replayed frame even if it checks the frame integrity and frame counter. The attack introduces a delay of τ seconds to the delivery of the frame.

We discuss several issues in the attack implementation. First, using a normal LoRaWAN frame to create malicious collision is more stealthy than brute-force jamming, since it may be difficult to differentiate malicious and normal collisions. Brute-force jamming can be easily detected and located. Second, as the adversary delays the uplink frame, how does the adversary know in time the direction of the current transmission? In LoRaWAN, the uplink preamble uses up chirps, whereas the downlink preamble uses down chirps. Thus, the adversary can quickly detect the direction of the current transmission within a chirp time. From our results in §4.3.2, a time duration of one chirp for sensing the direction of the transmission does not impede the timeliness of the collision attack. Third, to increase the stealthiness of the replay attack, the replayer can well control the transmitting power of the replay such that only the victim gateway can receive the replayed frame. Fourth, the attack does not require clock synchronization between the eavesdropper and the collider.

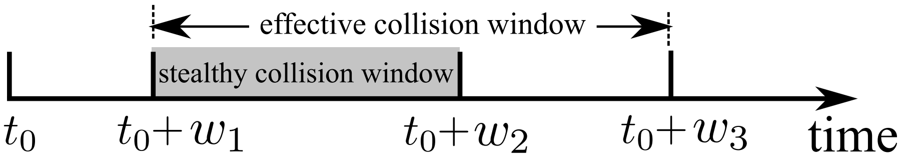

FIGURE 4.2: Collision attack time window.

4.3.1.2 Discussion on a simple attack detector

A simple attack detection approach is to perform round-trip timing and then compare the measured round-trip time with a threshold. However, this approach has the following three shortcomings. First, it needs a downlink transmission for each uplink transmission, which doubles the communication overhead. LoRaWAN is mainly designed and optimized for uplinks. For instance, a LoRaWAN gateway can receive frames from multiple end devices simultaneously using different spreading factors, whereas it can send a single downlink frame only at a time. This is because Class A specification requires that any downlink transmission must be unicast, in response to a precedent uplink transmission. Thus, the round-trip timing approach matches poorly with the uplink-downlink asymmetry characteristic of LoRaWAN. Second, with this simple attack detection approach, it is the end device detecting the attack after receiving the downlink acknowledgement. The end device needs to inform the gateway using another uplink frame that is also subject to malicious collision. Third, as the attacks are rare (but critical) events, continually using downlink acknowledgements to preclude the threat is a low cost-effective solution. In summary, this simple round-trip timing countermeasure is inefficient and error-prone.

4.3.2 Timing of Malicious Frame Collision

In this section, we study the timing of effective malicious frame collision. When investigating the geographic area affected by the attack, the ratio between the powers of the victim signal and the collision signal also needs to be considered. §4.3.3 will jointly consider the collision timing and the signal power ratio. We set up two SX1276-based LoRa nodes as the transmitter and the receiver, which are separated by about 5 m. We use a third LoRa node as the collider against the receiver. The distance between the collider and the receiver is about 1 m. Although the quantified results obtained based on SX1276 are chip specific, the qualitative results (i.e., the trend) are consistent with the general understanding on wireless demodulation. Thus, the qualitative results provide general insights and implications. The gateway-class iC880A LoRaWAN concentrator and an open-source LoRa demodulator that we use in §4.3.3 also exhibit similar trend. In practice, the adversary may conduct experiments similar to those presented below to obtain the required attack timing once they know the model of the victim LoRa chip.

From our experiments, there are three critical time windows (denoted by w_1, w_2, and w_3) after the onset time of the victim transmission (denoted by t_0). These time windows are illustrated in Fig. 4.2. If the onset time of the collision frame is in $[t_0, t_0 + w_1]$, the receiver most likely receives the collision frame only; if it is in $[t_0 + w_1, t_0 + w_2]$, the receiver receives neither frame and raises no alerts; if it is in $[t_0 + w_2, t_0 + w_3]$, the receiver reports "bad frame" and yields no frame content; if it is after $t_0 + w_3$, the receiver can receive both frames sequentially. Therefore, the time window $[t_0 + w_1, t_0 + w_2]$ is called *stealthy collision window* and the $[t_0 + w_1, t_0 + w_3]$ is called *effective collision window*. Note that we view the "bad frame" situation as effective attack, because the receiver cannot differentiate malicious and normal collisions based on the warning message. Moreover, if the preamble is corrupted, there is no way for the receiver to detect the collision.

Our experiments measure w_1, w_2, and w_3 under a wide range of settings including spreading factor and the payload size. Table 4.1 summarizes the results. From the results for w_1, the collision should start after the 5th chirp of the victim frame transmission. Explanation is as follows. (Note that as the demodulation mechanism of used SX1276 is proprietary and not publicly available, our explanations in this section are based on general understanding on wireless demodulation.) First, the receiver has not locked the victim frame's preamble until the 6th chirp. If the collision starts before the 5th chirp of the victim frame, the receiver will re-lock the collision frame's preamble with higher signal strength, resulting in reception of the collision frame. Second, the receiver locks the victim frame's preamble from the 6th chirp and simply drops any received radio data without reporting any error if any of the last three chirps (i.e., the 6th, 7th, and 8th chirps) of the preamble and/or the frame header are corrupted. For the latter case of frame header corruption, the radio chip cannot determine whether itself is the intended recipient and hence

Table 4.1: Collision time windows for SX1276.

Spreading factor S	Chirp time	Preamble time	Payload (byte)	w_1	w_2	w_3
7	1.024	8.2	10	5	28	141
			20	5	38	156
			30	6	41	165
			40	6	54	178
7	1.024	8.2	30	6	41	165
8	2.048	16.4		10	82	208
9	4.096	32.8		22	156	274

* Unit for chirp time, preamble time, w_1, w_2, w_3 is millisecond.

drops the received data. Thus, the collision should start after the 5th chirp of the victim frame.

We can also see that w_2 increases exponentially with the spreading factor. This is because: i) the total time for transmitting the preamble and frame header increases exponentially with the spreading factor; ii) corruption of the payload after the frame header leads to integrity check error and the "bad frame" message. The w_3 is roughly the time for transmitting the victim frame. Thus, if the collision onset time is after $t_0 + w_3$, both the victim and collision frames can be received.

The above experiments show that, there is a time window of more than 20 ms for the collision to corrupt the preamble partially and the frame header such that the victim simply drops the received data and raises no alerts. Collision starting in this window is stealthy. There is also an effective attack window of more than 100 ms. It is not difficult to satisfy such timing requirements using commodity radio devices.

4.3.3 Size of Vulnerable Area

In this section, through simulations and extensive experiments in a campus, we show that by setting up a collider and an eavesdropper at fixed locations, the frame delay attack can affect many end devices in a geographic area. The simulations based on realistic measurements with an open-source LoRa demodulator and a path loss model [51] provide insights into understanding the vulnerable area. The experiments in the campus further capture other affecting factors such as terrain

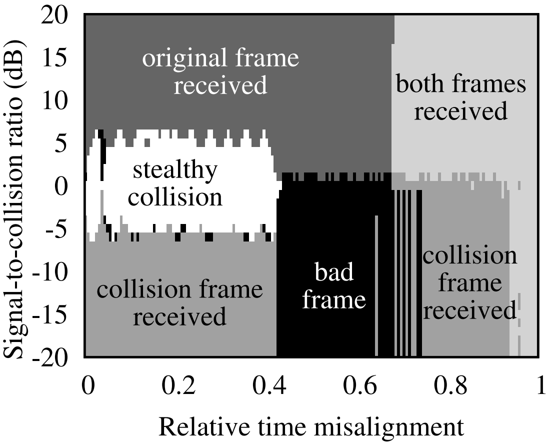

FIGURE 4.3: Result of `gr-lora`'s demodulation under collision with different signal-to-collision ratios and relative time misalignments.

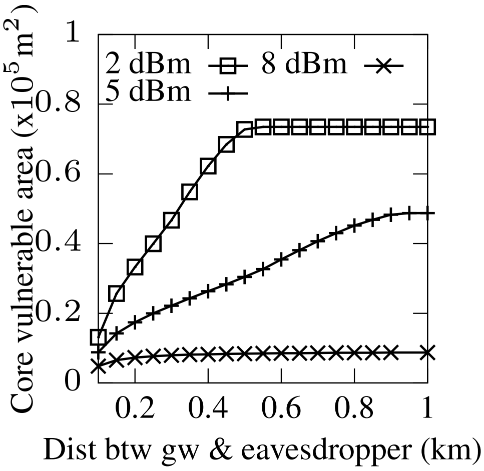

FIGURE 4.4: Core vulnerable area vs. distance between gateway and eavesdropper under various collision powers.

and signal blockage from buildings. In this section, the *core vulnerable area* refers to the geographic area in which the end devices are subject to stealthy collision and successful eavesdropping; the *vulnerable area* additionally includes the area in which the end devices are subject to the collision causing "bad frame" reports and successful eavesdropping.

4.3.3.1 Simulations

To study the vulnerable area, we need to consider the signal path loss and the ratio between the powers of the victim signal and the collision signal at the receiver. We call this ratio *Signal-to-Collision Ratio* (SCR). To characterize attack timing, we define *Relative Time Misalignment* (RTM) as $\frac{\text{collision time lag}}{\text{frame time}}$, where the collision time lag is the time lag of the collision onset from the victim signal onset. In our simulation, the victim and collision frames have identical length but different payload contents. We generate the I and Q waveforms of these two frames using LoRa signal model. We superimpose the two frames' signals to simulate collision. Moreover, we scale the amplitudes of the two signals and time-misalign them to create certain SCR and RTM. The sum signal is processed using an open-source LoRa demodulator `gr-lora` [85]. Fig. 4.3 shows the demodulation results under various SCR and RTM settings. We can see that if RTM is less than 0.4 and SCR at the gateway is within $[-6\,\text{dB}, 6\,\text{dB}]$, the collision is stealthy. The eavesdropped frame can be demodulated if SCR at the eavesdropper is greater than 6 dB.

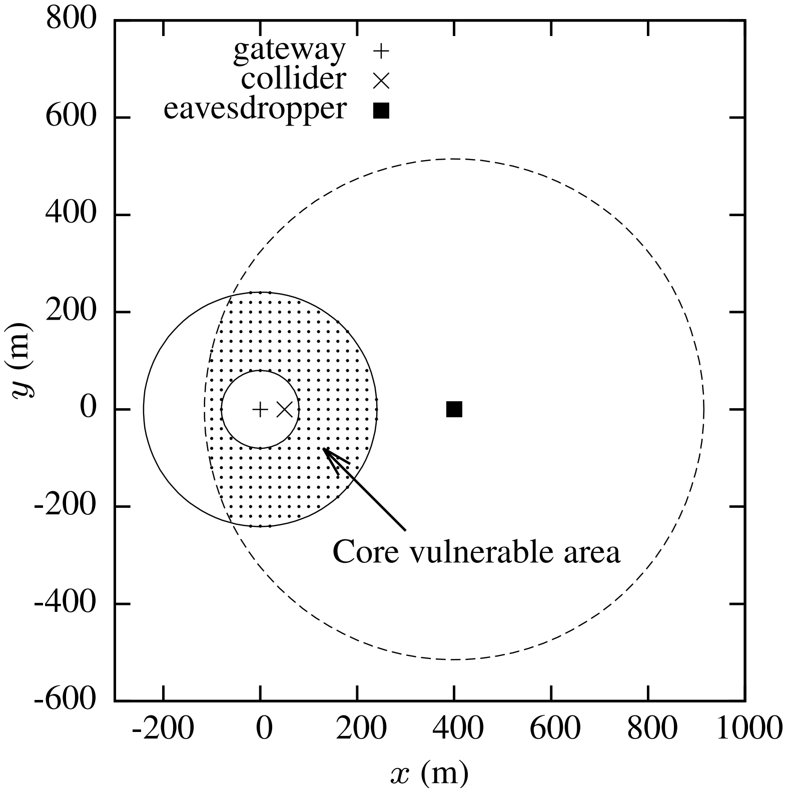

FIGURE 4.5: The core vulnerable area (i.e., the shaded area) with gateway at $(0, 0)$, collider at $(50, 0)$, and eavesdropper at $(400, 0)$. Collider's and victim end device's transmitting powers are 2 dBm and 14 dBm, respectively. End devices in the ring centered at $(0, 0)$ are subject to stealthy collision; end devices in the dashed circle are subject to successful eavesdropping.

We adopt a LoRa signal path loss model for urban areas proposed in [51] based on real measurements. Specifically, the path loss L in dB is given by

$$\begin{aligned}L &= 69.55 + 26.16\log f - 13.82\log h_b - (1.1\log f - 0.7)h_m \\ &\quad + (1.56\log f - 0.8) + (44.9 - 6.55\log h_b)\log d,\end{aligned}$$

where the base of the logarithm is 10, f is LoRa signal's central frequency in MHz, h_m and h_b are the heights of the transmitter and receiver in meters, and d is the distance in kilometers between the transmitter and the receiver. The frame delay attack is successful if the attacker can control RTM below 0.4 and satisfy the following two conditions:

$$-6\,\mathrm{dB} \le P_v - L_{v,g} - (P_c - L_{c,g}) \le 6\,\mathrm{dB}, \tag{4.1}$$

$$6\,\mathrm{dB} \le P_v - L_{v,e} - (P_c - L_{c,e}), \tag{4.2}$$

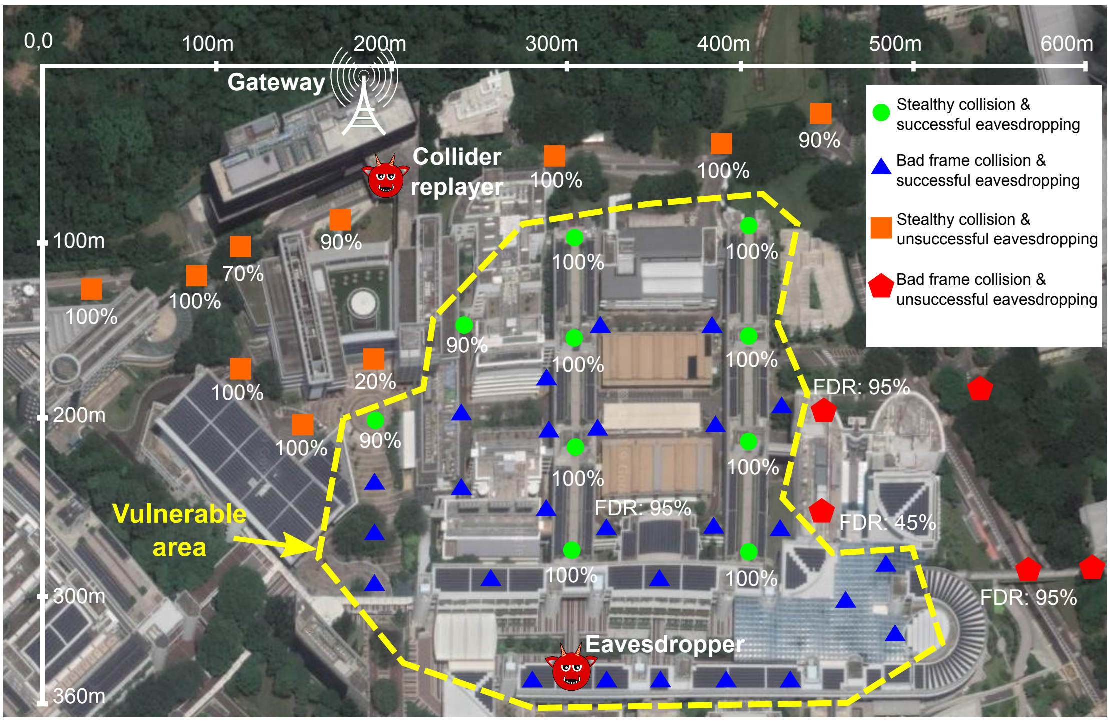

Figure 4.6: Vulnerable area of a campus LoRaWAN. A gateway and a USRP-based eavesdropper are deployed on the rooftops of two buildings. A collider is deployed on an overhead bridge. We carry an end device to each of the marked locations and conduct an attack experiment. The four point shapes represent four types of attack outcomes. (Satellite image credit: Google Map)

where the subscripts v, g, c, and e respectively denote the victim end device, the gateway, the collider, and the eavesdropper; P_x denotes the transmitting power of device x; $L_{x,y}$ denotes the path loss from device x to y. Eq. (4.1) is the condition for stealthy collision; Eq. (4.2) is the condition for successful eavesdropping. The SCR thresholds of 6 dB and -6 dB in Eqs. (4.1) and (4.2) are from Fig. 4.3. Note that our modeling of successful eavesdropping in Eq. (4.2) only considers the case that the signal from the collider at the eavesdropper has a power much higher than the noise floor, so that we can ignore the impact of noise on the eavesdropping.

Fig. 4.5 shows an example of the areas defined by Eqs. (4.1) and (4.2). The collider's and end device's transmitting powers are 2 dBm and 14 dBm. The gateway's altitude is 25 m; the collider, eavesdropper, and end devices have an identical altitude of 0 m. As shown in Fig. 4.5, the ring centered at the gateway is defined by Eq. (4.1); the disk area in the dashed circle is defined by Eq. (4.2). Thus, the overlap between the ring and the disk is the core vulnerable area, which is $62,246\,\text{m}^2$. Then, we vary the distance between the gateway and the eavesdropper (denoted by d_{ge}) and the P_c setting. Fig. 4.4 shows the resulting core vulnerable area. We can see that the core vulnerable area in general increases with d_{ge} and becomes flat after d_{ge} exceeds a certain value. Moreover, among the three P_c settings (i.e., 2, 5, and 8 dBm), $P_c = 2\,\text{dBm}$ gives larger core vulnerable areas. Reason of the above two observations is that the eavesdropper can achieve a larger eavesdropping area due to the weaker collision signal received by the eavesdropper. The core vulnerable area saturates because the eavesdropping area in the dashed circle illustrated in Fig. 4.5 covers the entire ring area when d_{ge} exceeds a certain value. Note that when d_{ge} is very large, the noise power dominates and the core vulnerable area shrinks to zero.

The above simulation results suggest that the location of the gateway is the key information that the adversary needs to obtain. Based on that, the adversary can plan the placement of the collider and eavesdropper to affect a large geographic area. For the LoRaWANs adopting multiple gateways, the adversary can place a collider close to each of the gateways. In practice, the locations of the gateways can be obtained by the adversary in various ways (e.g., social engineering) and should not be relied on for the security of the system.

4.3.3.2 Experiments in a campus LoRaWAN

We conduct a set of experiments in an existing campus LoRaWAN to investigate the vulnerable area in real environments. Note that the LoRaWAN consists of three gateways that can cover the whole campus. Our experiments only involve one of the three gateways, which covers the area shown in Fig. 4.6 that has a number of multistory buildings. The gateway, which consists of an iC880a LoRaWAN concentrator board, a Raspberry Pi, and a high-gain antenna, is located on the rooftop of a building. Both the collider and the eavesdropper consist of a laptop computer and a USRP N210 each. The collider is placed on an overhead bridge attached to the gateway's building. The horizontal distance between the gateway and the collider is about 50 m. The eavesdropper is placed on the rooftop of another building that is about 320 m from the gateway's building. We carry an SX1276-based LoRaWAN end device to each of the locations marked in Fig. 4.6, measure the FDR, and perform an attack experiment. The measured FDRs at all the visited locations are 100%, except the four locations labeled with non-100% FDRs. Thus, the gateway can cover the accessible area shown in Fig. 4.6.

In each attack experiment, the end device's and the collider's transmitting powers are 14 dBm and 8 dBm, respectively. All malicious collisions are effective. The outcomes can be classified into four categories, which are the combinations of the collision results (stealthy collision or "bad frame") and eavesdropping results (successful or unsuccessful). In Fig. 4.6, we use four point shapes to represent the four attack outcomes. The percentage below a location is the ratio of stealthy collisions. We can see that, at most locations close to the gateway and collider, the malicious collisions are stealthy. At the locations in the bottom most part of Fig. 4.6, the collisions cause gateway's bad frame reports. There is a transit region in the middle of Fig. 4.6, in which the collision outcomes are mixed. Note that the visited locations shown in Fig. 4.6 are on the rooftops, in semi-outdoor corridors, or in indoor environments. The indoor/outdoor condition may affect the collision outcome type. At the locations in the area enclosed by the dashed polygon, the gateway can decode the frame that is recorded by the eavesdropper and then replayed by the collider, suggesting that the eavesdropping is successful. Thus, this area is the vulnerable area caused by the attack setup, which is about $50,000\,\mathrm{m}^2$.

Note that the demodulation mechanism of the iC880a concentrator is proprietary and can be different from the open-source LoRa demodulator we used in §4.3.3.1. The actual signal propagation behaviors in the campus LoRaWAN can be much more complex than the model used in §4.3.3. However, the simulation result (Fig. 4.5) and real experiment result (Fig. 4.6) show similar patterns, i.e., the eavesdropping area is around the eavesdropper and the core vulnerable area is a belt region between the gateway and the eavesdropper. Thus, our modeling and simulations in §4.3.3 provide useful understanding on the LoRaWAN vulnerability.

4.4 LoRaTS Gateway

As shown in §4.3, a fixed setup of a collider and an eavesdropper can subvert the sync-free timestamping for many end devices in a large geographic area. This section presents the LoRaTS gateway that supports the bandwidth-efficient sync-free timestamping as an advantage throughout the network lifetime and develops awareness of the frame delay attack. Thus, it strikes a good trade off between efficiency and security. This section presents the LoRaTS gateway. We describes its hardware and software in §4.4.1 and §4.4.2, respectively. §4.4.3 models the

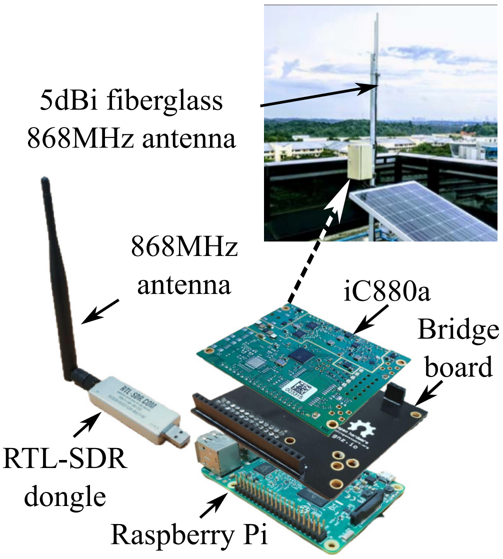

FIGURE 4.7: LoRaTS hardware prototype consisting of Raspberry Pi, iC880a concentrator, bridge board, RTL-SDR USB dongle.

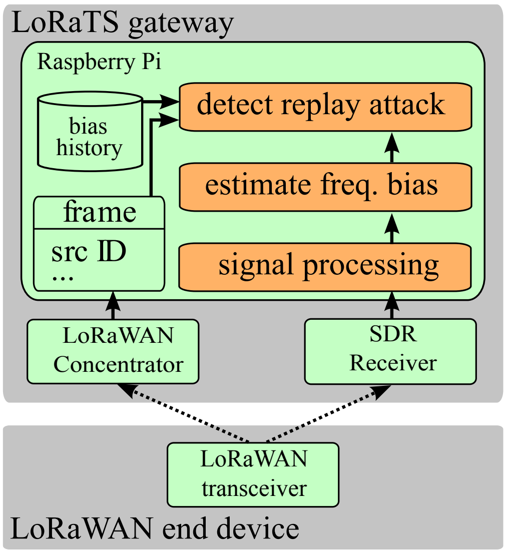

FIGURE 4.8: LoRaTS software. Bottom part is end device; upper part is gateway; solid arrows are local data flows; dashed arrows are transmissions.

reception of LoRa signal using an SDR receiver and presents the algorithm to timestamp uplink frame arrival.

4.4.1 LoRaTS Gateway Hardware

To detect the attack, we integrate an SDR receiver with a LoRaWAN gateway to monitor the physical layer. Various cheap (US$25 only [86]) and low-power SDR receivers are available now. In the second part of this thesis, we use RTL-SDR USB dongles based on the RTL2832U chipset [21], which were originally designed to be DVB-T TV tuners. The RTL-SDR supports continuous tuning in the range of $[24, 1766]$ MHz, which covers the LoRaWAN bands (i.e., 430, 433, 868, 915 MHz). It can operate at 2.4 Msps reliably for extended time periods. Thus, the sampling resolution is $1/2.4\,\text{Msps} = 0.42\,\mu\text{s}$. Our research is conducted based on a LoRaTS hardware prototype that integrates a Raspberry Pi, an iC880a LoRaWAN concentrator, and an RTL-SDR USB dongle. Fig. 4.7 shows the prototype. An 868 MHz antenna is used with the RTL-SDR to improve signal reception.

The SDR receiver is used to capture the radio signal over a time duration of the first two preamble chirps of an uplink frame. The first sampled chirp is used to extract an accurate timestamp (cf. §4.4.3), whereas the second sampled chirp is used to extract the FB of the transmitter (cf. §4.5). The accurate timestamp is a prerequisite of the FB estimation. As only two chirps' radio waveform is analyzed, the Raspberry Pi suffices for performing the computation. Instead of using RTL-SDR, a full-fledged SDR transceiver (e.g., USRP) can be used to design a customized gateway with physical layer access. However, this design loses the factory-optimized hardware-speed LoRa demodulation built in the iC880a concentrator. Moreover, full-fledged SDR transceivers are often 10x more expensive than LoRaTS. The low-cost, low-power, listen-only RTL-SDR suffices for developing the attack detector.

4.4.2 LoRaTS Gateway Software

The upper part of Fig. 4.8 illustrates the software architecture of LoRaTS to detect the attack. It is based on the results in the subsequent sections of this chapter. The uplink transmission from the end device is captured by both the gateway's LoRaWAN concentrator and the SDR receiver. The LoRaWAN concentrator demodulates the received radio signal and passes the frame content to the Raspberry Pi. Signal processing algorithms are applied on the LoRa signal after down-conversion by the SDR receiver to determine precisely the arrival time of the uplink frame, estimate the transmitter's FB, and detect whether the current frame is a replayed one. The replay detection is by checking whether the estimated FB is consistent with the historical FBs associated with the transmitter ID contained in the current frame. Thus, the gateway is aware of the attack and can take necessary actions. Note that LoRaTS uses the SDR receiver to obtain FBs, rather than to decode the frame.

4.4.3 Signal Modeling and Uplink Frame Arrival Timestamping

In this section, we present the modeling of the LoRa signal reception and our approach of detecting the onset time of the first preamble chirp. They form a basis

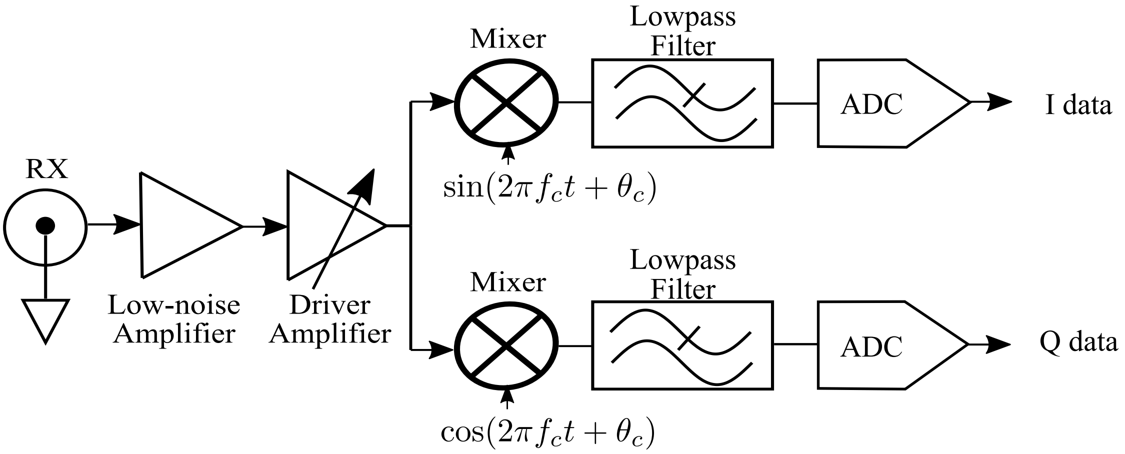

FIGURE 4.9: Analog signal processing in SDR receiver.

for developing the FB estimation algorithms in §4.5.

4.4.3.1 Derivation of I and Q components of LoRa signal

Fig. 4.9 illustrates the essential analog signal processing steps of most SDR receivers to yield the in-phase (I) and quadrature (Q) components of the received radio signal. The SDR receiver generates two unit-amplitude orthogonal carriers $\sin(2\pi f_c t + \theta_{\mathrm{Rx}})$ and $\cos(2\pi f_c t + \theta_{\mathrm{Rx}})$, where f_c is a specified frequency and θ_{Rx} is the phase of the two self-generated carriers. The f_c can be set to be the central frequency of the used LoRa channel. The I and Q components, denoted by $s_I(t)$ and $s_Q(t)$, are

$$\begin{aligned} s_I(t) =& s(t) \cdot \sin(2\pi f_c t + \theta_{\mathrm{Rx}}) \\ =& \frac{A(t)}{2} \left(\cos \left(2\pi \int_0^t f(x) \mathrm{d}x - 2\pi f_c t + \theta_{\mathrm{Tx}} - \theta_{\mathrm{Rx}} \right) \right. \qquad (4.3) \\ & \left. - \cos \left(2\pi \int_0^t f(x) \mathrm{d}x + 2\pi f_c t + \theta_{\mathrm{Tx}} + \theta_{\mathrm{Rx}} \right) \right), \qquad (4.4) \end{aligned}$$

$$\begin{aligned} s_Q(t) =& s(t) \cdot \cos\left(2\pi f_c t + \theta_{\mathrm{Rx}}\right) \\ =& \frac{A(t)}{2} \left(\sin \left(2\pi \int_0^t f(x) \mathrm{d}x - 2\pi f_c t + \theta_{\mathrm{Tx}} - \theta_{\mathrm{Rx}} \right) \right. \qquad (4.5) \\ & \left. + \sin \left(2\pi \int_0^t f(x) \mathrm{d}x + 2\pi f_c t + \theta_{\mathrm{Tx}} + \theta_{\mathrm{Rx}} \right) \right), \qquad (4.6) \end{aligned}$$

The high-frequency components in Eqs. (4.4) and (4.6) are removed by the low-pass filters of the SDR receiver. Thus, the I and Q components after the filtering, denoted by $I(t)$ and $Q(t)$, are given by Eqs. (4.3) and (4.5). They can be rewritten as

$$
\begin{aligned}
&I(t) = \frac{A(t)}{2}\cos\Theta(t), \quad Q(t) = \frac{A(t)}{2}\sin\Theta(t), \\
&\Theta(t) = 2\pi\int_0^t f(x)\mathrm{d}x - 2\pi f_c t + \theta, \quad \theta = \theta_{\mathrm{Tx}} - \theta_{\mathrm{Rx}}.
\end{aligned}
$$

The continuous-time $I(t)$ and $Q(t)$ are then sampled by the analog-to-digital converters (ADCs) to yield the I and Q data. For simplicity of exposition, the analysis in this chapter is performed in the continuous-time domain.

4.4.3.2 CSS reception using SDR receiver

A chirp is a finite-time signal with time-varying frequency that sweeps the channel's bandwidth. Specifically, it can be expressed as $s(t) = A(t)\sin\left(2\pi\int_0^t f(x)\mathrm{d}x + \theta_{\mathrm{Tx}}\right)$, where $A(t)$ and $f(t)$ denote the instantaneous amplitude and frequency of the chirp at the time instant t, $\theta_{\mathrm{Tx}} \in [0, 2\pi)$ is the transmitter's phase that is usually unknown.

A LoRaWAN uplink preamble consists of eight up chirps by default [3]. For a preamble chirp, $f(t) = \frac{W^2}{2^S}\cdot t - \frac{W}{2} + f_c$ for $t \in \left[0, \frac{2^S}{W}\right]$, where W is the channel bandwidth, $S \in \{6, 7, \dots, 12\}$ is the *spreading factor*, and $\frac{2^S}{W}$ is the *chirp time*. The $f(t)$ increases linearly from $(f_c - W/2)$ Hz to $(f_c + W/2)$ Hz over a chirp time. The angle of the preamble chirp can be derived as $\Theta(t) = \frac{\pi W^2}{2^S}t^2 - \pi W t + \theta_{\mathrm{Tx}} - \theta_{\mathrm{Rx}}$. In this chapter, we use a channel with $f_c = 869.75\,\mathrm{MHz}$ and $W = 125\,\mathrm{kHz}$. Fig. 4.10 shows the I data and the spectrogram of an ideal preamble chirp. The parameters for generating Fig. 4.10 are $A(t) = 2$, $\theta_{\mathrm{Tx}} - \theta_{\mathrm{Rx}} = 0$, and $S = 7$. Thus, the chirp time $\frac{2^S}{W}$ is $1.024\,\mathrm{ms}$. To generate the spectrogram, we apply the short-time Fast Fourier Transform (FFT) with 2^S-point Kaiser window and 16-point overlap between two neighbor windows. Thus, the spectrogram consists of 20 Power Spectral Densities (PSDs) over the chirp time of $1.024\,\mathrm{ms}$.

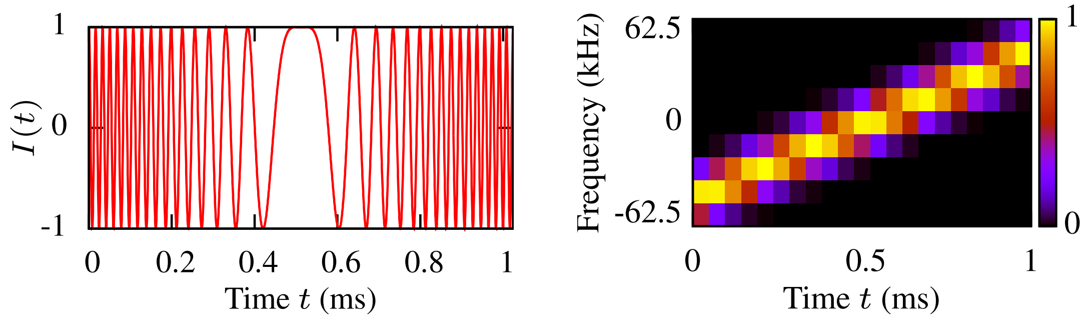

(a) I data

(b) Spectrogram

FIGURE 4.10: I data (θ_{Tx}=θ_{Rx}) and spectrogram of a preamble chirp.

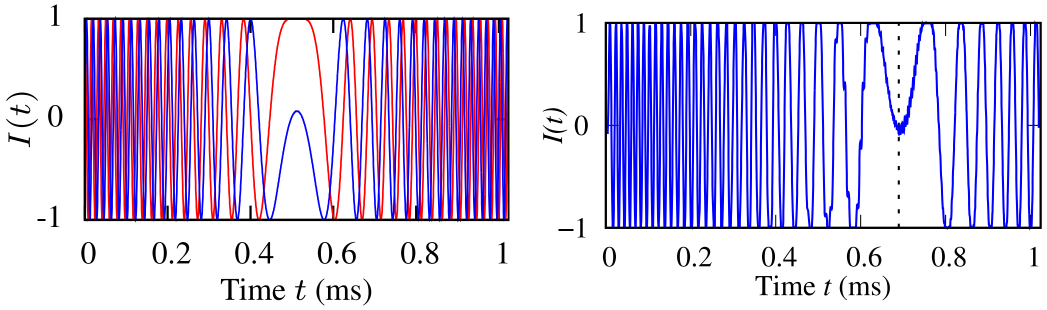

(a) I data with different phases.

(b) Actual I data of a preamble chirp with frequency bias.

FIGURE 4.11: LoRa's I signal waveform is affected by the initial phase and frequency bias (FB).

4.4.3.3 Preamble onset time detection

Detecting the onset time of the preamble is non-trivial. In this section, we discuss the matched filter approach and its inefficacy. Then, we present three other candidate methods.

Matched filter is a widely adopted symbol detection technique. As a coherent detection technique, the matched filter requires that the receiver is phase-locked to the transmitter (i.e., $\theta_{\text{Rx}} = \theta_{\text{Tx}}$) to achieve the best symbol detection accuracy. However, as LoRa adopts time-varying frequency, it is difficult for the SDR receiver to estimate the transmitter's phase θ_{Tx}. As a result, the phase difference $\theta_{\text{Tx}} - \theta_{\text{Rx}}$, which is a critical factor affecting the shape of $I(t)$ and $Q(t)$, will be random. Fig. 4.11(a) shows the ideal $I(t)$ traces of the preamble chirp when the phase difference is 0 and π, respectively. The waveform shapes are different. Thus, we cannot define a template shape required by the matched filter.

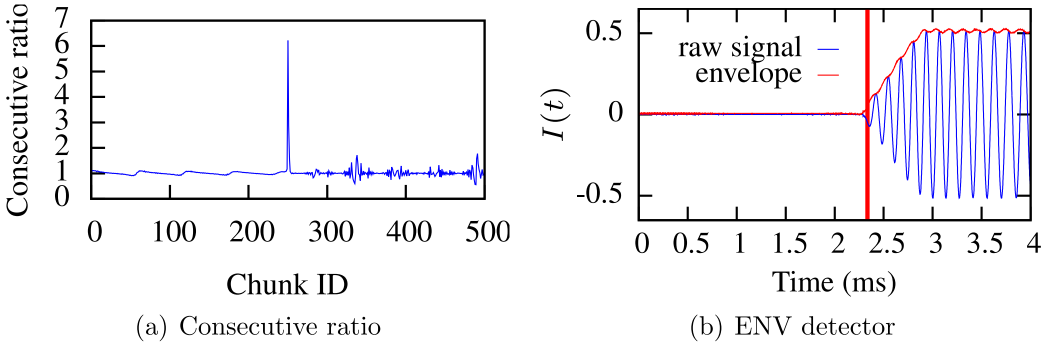

(a) Consecutive ratio

(b) ENV detector

FIGURE 4.12: The intermediate and final results of envelope detector.

For LoRaTS, we consider three parameter-less detectors:

Envelope (ENV) detector: First, we apply the Hilbert transform to extract the amplitude envelope of the I or Q signal. Fig. 4.12(b) shows the extracted amplitude envelope for I data. We adopt the *folding* technique [87, 88] to detect the signal onset time from the amplitude envelope. Specifically, we evenly divide the envelope into chunks of equal length (e.g., 200 samples). Then, we calculate the sum of the absolute values of the amplitudes of all samples in each chunk, which is referred to as *trunk sum*. Lastly, we compute the ratio between the trunk sums of any two consecutive trucks to generate a ratio sequence. As shown in Fig. 4.12(a), the ratio sequence has a single peak. The detector yields the peak's time instant as the preamble onset time. The red vertical line in Fig. 4.12(b) indicates the detected onset time.

Correlation (CORR) detector: The Start Frame Delimiter (SFD) in a LoRa frame consists of two and a quarter down chirps. SFD is used by LoRa receiver for synchronization, because the junction of between the up chirp before SFD and the first down chirp of SFD presents a salient hill peak as shown in the upper part of Fig. 4.15. We can compute the correlation between the spectrograms of the received LoRa signal and a locally generated hill peak template. The maximum of the correlation trace gives the time instant of the hill peak, which can be used to infer the onset time of the LoRa frame. The bottom part of Fig. 4.15 shows the normalized correlation coefficient trace and the detected hill peak time represented by the red vertical line.

AIC detector: The autoregressive Akaike Information Criterion (AIC) algorithm [89] was originally developed to estimate the arrival time of seismic waves with

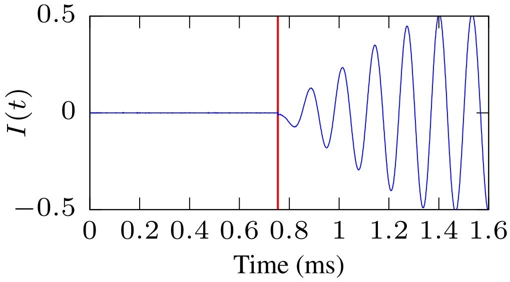

FIGURE 4.13: AIC detector

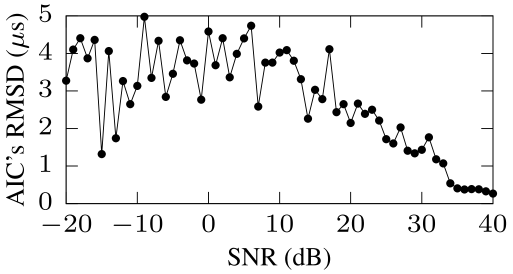

FIGURE 4.14: AIC's RMSD vs. SNR.

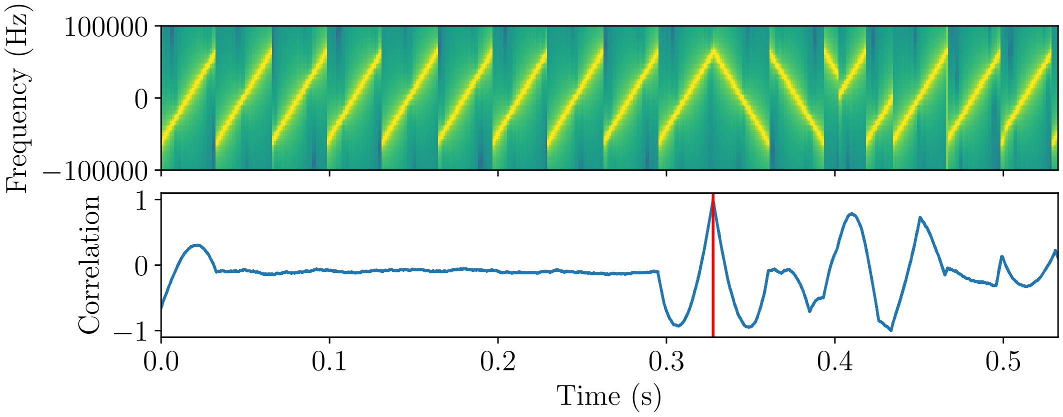

FIGURE 4.15: Correlation detector. The upper part shows the spectrogram of a LoRa frame. The bottom part shows the normalized correlation coefficient between the spectrogram and a locally generated hill peak pattern.

an accuracy of a single sampling point. As the I and Q signals are similar to the seismic waves [20], the AIC is a promising solution for our problem. It works as follows. For each point of the signal as an onset time candidate, two autoregressive models are constructed for the signal segments before and after the onset time candidate. The candidate that gives the largest dissimilarity between the two autoregressive models is yielded as the final result. From Fig. 4.13, AIC can detect the onset time from the signal with a smooth start. From the results in [89], AIC's detection results have a bias of 4 samples. With a sampling rate of 2.4 MHz, the bias is $\mathbb{E}[\epsilon] = \frac{4}{2.4\,\text{Msps}} = 1.67\,\mu\text{s}$ only, where ϵ represents onset time detection error.

4.4.3.4 Evaluation

We conduct experiments to evaluate the performance of the three detectors presented above. As AIC is nearly unbiased [89], we primarily assess the Root-Mean-Square Deviation (RMSD) which characterizes the consistency of the detection results. Due to the difficulty in obtaining the ground truth of the preamble arrival time, we indirectly estimate the RMSD as follows. We place two LoRaTS nodes A and B close to each other such that the signal propagation time is near-zero. Node A initiates a round-trip communication; each of them detects the onset times for both its transmitted and received signals, generating total four onset times. Note that a LoRaTS node's SDR receiver can also capture the signal transmitted by the node's LoRa radio and detect the onset time. Denote by Δ the measured round-trip time based on the detected onset times; denote by ϵ_X^{Tx} and ϵ_X^{Rx} the node X's unknown onset time detection errors for its transmitted and received signals. For the two close nodes A and B, $\Delta = \epsilon_A^{Rx} - \epsilon_A^{Tx} + \epsilon_B^{Rx} - \epsilon_B^{Tx}$ and $\text{RMSD}(\epsilon) = \frac{1}{2}\text{RMSD}(\Delta)$ if the errors are independent and identically distributed. Measurements show that RMSD(ϵ) is $1.21\,\mu$s, $0.64\,\mu$s, and $0.33\,\mu$s for ENV, CORR, and AIC, respectively. Thus, AIC achieves more consistent detection results. Then, we evaluate the impact of random noises on AIC's RMSD(ϵ). We artificially add zero-mean Gaussian noises to the collected high-SNR I and Q traces. Then, we apply AIC on the noise-added traces to detect the preamble onset time. Fig. 4.14 shows the results. Note that the SNR range in Fig. 4.14 can cover realistic SNRs, e.g., 13 dB to -1 dB in a multistory building (cf. §4.6). From Fig. 4.14, the AIC's RMSD(ϵ) is less than $5\,\mu$s when the SNR is down to -20 dB. Thus, AIC achieves robust onset time detection in the presence of strong noises. The rest of this chapter uses AIC.

4.5 Frame Delay Attack Detection

Internal oscillators for generating carriers generally have FBs due to manufacturing imperfection. This section develops algorithms for estimating LoRa transmitters' FBs based on LoRa's CSS modulation and use them to detect the frame delay attack. Note that the existing FB estimation algorithms developed for other radios cannot be ported to LoRa due to different modulation schemes. For instance, the FB estimation for OFDM [90] is apparently not applicable for LoRa CSS. As discussed later, LoRa demodulation's built-in FB estimation technique does not provide sufficient resolution. Thus, highly accurate FB estimation for LoRa CSS is a non-trivial problem.

4.5.1 FB Estimation

This section describes algorithms for estimating the transmitter's FB based on an up chirp in the preamble. First, we analyze the impact of the transmitter's and SDR receiver's FBs (denoted by δ_{Tx} and δ_{Rx}) on the I and Q traces. The up chirp's instantaneous frequency accounting for δ_{Tx} is $f(t) = \frac{W^2}{2^S} \cdot t - \frac{W}{2} + f_c + \delta_{\mathrm{Tx}}$, $t \in \left[0, \frac{2^S}{W}\right]$. The two local unit-amplitude orthogonal carriers generated by the SDR receiver are $\sin(2\pi(f_c + \delta_{\mathrm{Rx}})t + \theta_{\mathrm{Rx}})$ and $\cos(2\pi(f_c + \delta_{\mathrm{Rx}})t + \theta_{\mathrm{Rx}})$. After mixing and low-pass filtering, the I and Q components of the received up chirp can be derived as $I(t) = \frac{A(t)}{2} \cos \Theta(t)$ and $Q(t) = \frac{A(t)}{2} \sin \Theta(t)$, where the angle $\Theta(t)$ is given by

$$\Theta(t) = \frac{\pi W^2}{2^S} t^2 - \pi W t + 2\pi\delta t + \theta_{\mathrm{Tx}} - \theta_{\mathrm{Rx}}, \quad \delta = \delta_{\mathrm{Tx}} - \delta_{\mathrm{Rx}}. \tag{4.7}$$

When $\delta = 0$, the axis of symmetry of $I(t)$ is located at the midpoint of the preamble chirp time. As shown in Fig. 4.11(b), a negative δ causes a right shift of the axis of the symmetry in the time domain, whereas a positive δ causes a left shift.

For a certain SDR receiver, the FB estimation problem is to estimate δ from the captured I and Q traces. We do not need to estimate δ_{Tx}, because for a certain SDR receiver with a nearly fixed δ_{Rx}, a change in δ indicates a change in δ_{Tx} and a replay attack. In fact, FB estimation is a prerequisite of LoRa demodulation. Now, we discuss the incompetence of the LoRa demodulators' built-in FB estimation technique for attack detection. LoRa's CSS scheme evenly divides the whole channel bandwidth of W Hz into 2^S bins, where S is the spreading factor. The starting frequency of a bin corresponds to a symbol state. Since the preamble chirp linearly sweeps the channel bandwidth, its starting frequency can be viewed as the FB. LoRa demodulation firstly applies *dechirping* and then FFT to identify the preamble's and any data chirp's starting frequency bin indexes. The difference between the two indexes is the symbol state. As FFT achieves a resolution of $\frac{1}{x}$ Hz using x seconds of data, the Fourier transform of a chirp with length of $\frac{2^S}{W}$ seconds has a frequency resolution of $\frac{W}{2^S}$ Hz. This is also the resolution of the built-in FB estimation. Thus, for low spreading factor settings, the resolution may be poor. For instance, when $S = 7$ and $W = 125$ kHz, the resolution is 976.56 Hz. As we will show in §4.5.2, this near-1 kHz resolution is insufficient to detect attacks that introduce sub-1 kHz FBs. The colliding frame disentanglement approach Choir [52] also uses the dechirping-FFT pipeline to analyze FB. Thus, it is subject to

Algorithm 1 Timestamping and FB estimation

```
Given: $I$ and $Q$ traces from the SDR receiver, $k = 0$
apply AR-AIC to detect the onset time of the first chirp
$I[n]$ and $Q[n]$ ($n \in [1, N]$) are the $I$ and $Q$ segments from the detected onset
    time with a length of $2^S/W$ seconds
for $i = 1$ to $N$ do
    $t[n] = \frac{n}{N} \cdot \frac{2^S}{W}$, $\Theta_0[n] = \text{atan2}(Q[n], I[n])$
    if $n > 1$ and $\Theta_0[n-1] < -3$ and $\Theta_0[n] > 3$ then
        $k = k - 1$
    else if $n > 1$ and $\Theta_0[n-1] > 3$ and $\Theta_0[n] < -3$ then
        $k = k + 1$
    end if
    $\Theta[n] = \Theta_0[n] + 2k\pi$, $y[n] = \Theta[n] - \frac{\pi W^2}{2^S}(t[n])^2 + \pi W t[n]$
end for
apply linear regression to the data points $\{t[n], y[n] | n \in [1, N]\}$, yield the slope
    divided by $2\pi$ as the FB
```

the insufficient resolution. To achieve higher resolutions, this section presents two time-domain approaches designed based on Eq. (4.7). We also conduct extensive evaluations to compare their performance.

4.5.1.1 Linear regression approach

Eq. (4.7) can be rewritten as $\Theta(t) - \frac{\pi W^2}{2^S}t^2 + \pi W t = 2\pi\delta t + \theta$, which is a linear function of t with $2\pi\delta$ as the slope. Thus, the slope can be estimated by linear regression based on the data pairs $(t, \Theta(t) - \frac{\pi W^2}{2^S}t^2 + \pi W t)$, where $t \in \left[0, \frac{2^S}{W}\right]$, $\Theta(t) = \text{atan2}(Q(t), I(t)) + 2k\pi$, and $k \in \mathbb{Z}$ rectifies the multi-valued inverse tangent function $\text{atan2}(\,\cdot\,,\,\cdot\,) \in (-\pi, \pi)$ to an unlimited value domain. The rectification is as follows. The k is initialized to be 0 when $t = 0$. As t increases, if $\text{atan2}(Q(t), I(t))$ jumps from $-\pi$ to π, k decreases by one; if $\text{atan2}(Q(t), I(t))$ jumps from π to $-\pi$, k increases by one. Note that the traces $I(t)$ and $Q(t)$ where $t \in \left[0, \frac{2^S}{W}\right]$ are the segments of the captured I and Q signals starting from the preamble onset time detected by the AIC algorithm and lasting for a chirp time duration of $\frac{2^S}{W}$ seconds. Algorithm 1 shows the pseudocode of the FB estimation.

Fig. 4.16 shows the intermediate results of the FB extraction. Fig. 4.16(a) shows real $I(t)$ and $Q(t)$ traces of an up chirp emitted by an SX1276-based LoRa transmitter and captured by the SDR receiver. Fig. 4.16(b) shows the $\text{atan2}(Q(t), I(t))$. Fig. 4.16(c) shows the $\Theta(t)$ obtained by rectifying the result in Fig. 4.16(b) with

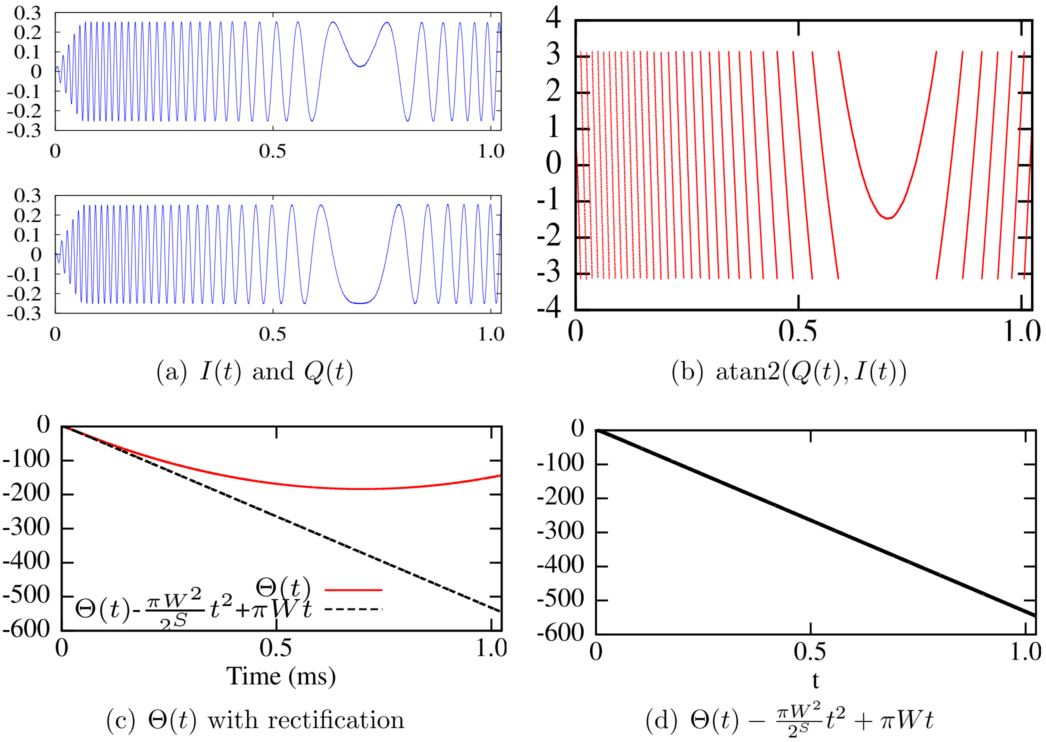

(a) $I(t)$ and $Q(t)$

(b) $\text{atan2}(Q(t), I(t))$

(c) $\Theta(t)$ with rectification

(d) $\Theta(t) - \frac{\pi W^2}{2^S}t^2 + \pi W t$

FIGURE 4.16: Estimating FB from real $I(t)$ and $Q(t)$ traces. All x-axes are time t in ms from the chirp onset time.

$2k\pi$. Fig. 4.16(d) shows $\Theta(t) - \frac{\pi W^2}{2^S}t^2 + \pi W t$. We can see that it is indeed a linear function of time t, which conforms to our analysis. By applying linear regression to the result in Fig. 4.16(d), the FB δ (i.e., the slope of the fitted line divided by 2π) is estimated as -22.8 kHz. Note that the $I(t)$ and $Q(t)$ are the I and Q data traces captured by the SDR receiver for a complete preamble chirp. The preamble onset time detected by AIC is used to segment the I and Q traces to chirps. From our measurements, the first preamble chirp, in general, has an increasing amplitude $A(t)$ after the onset time (as shown in Fig. 4.13), which generates a negative impact on the linear regression accuracy. As the second preamble chirp has more stable $A(t)$, we use the second chirp for the linear regression. Fig. 4.16(c) shows the $\Theta(t)$ computed from real I and Q traces of the second chirp of a preamble emitted by an SX1276-based end device and captured by LoRaTS's SDR receiver. It also shows $\Theta(t) - \frac{\pi W^2}{2^S}t^2 + \pi W t$, which is indeed a linear function of time. As the linear regression approach has a closed-form formula to compute δ, it has a complexity of $\mathcal{O}(1)$.

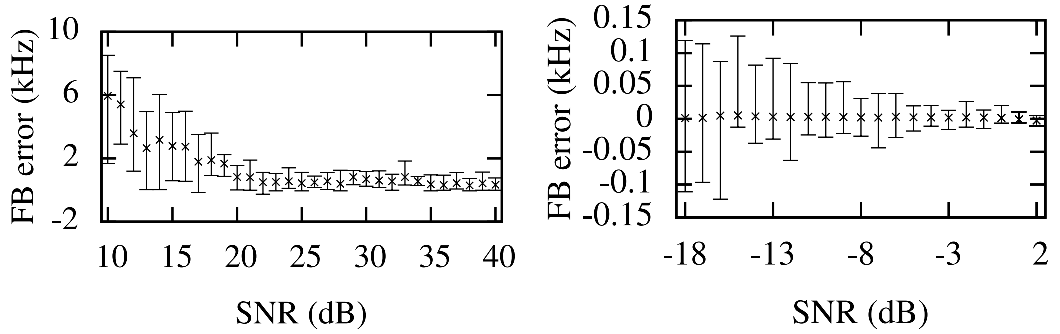

(a) Linear regression

(b) Least squares

FIGURE 4.17: FB estimation errors vs. SNR.

4.5.1.2 Least squares approach

The LoRa signals can be very weak after long-distance propagation or barrier penetration. The LoRa's demodulation is designed to address low SNRs. For SX1276, the minimum SNRs required for reliable demodulation with spreading factors of 7 to 12 are $-7.5\,\mathrm{dB}$ to $-20\,\mathrm{dB}$ [91]. We aim at extracting FB at such low SNRs. We solve a least squares problem: $\arg\min_{\theta_{\mathrm{Tx}}-\theta_{\mathrm{Rx}}\in[0,2\pi),\delta} \sum_{t\in[0,2^S/W]} \left(Q(t) - A\sin\Theta(t)\right)^2 + \left(I(t) - A\cos\Theta(t)\right)^2$, where $Q(t)$ and $I(t)$ are the received Q and I traces; $\Theta(t)$ is given by Eq. (4.7); $A\sin\Theta(t)$ and $A\cos\Theta(t)$ are the noiseless Q and I templates. The above formulation requires that the Q and I templates have an identical and constant amplitude A. As the second preamble chirp can meet this requirement, we use it for FB estimation. The A can be estimated as the square root of the difference between the average powers of the LoRa signal and the pure noise. Let $Q(t) = A\sin\Theta(t) + Z_Q(t)$ and $I(t) = A\cos\Theta(t) + Z_I(t)$, where $Z_Q(t)$ and $Z_I(t)$ are zero-mean random noises in the Q and I traces, respectively. Thus, the average power of the LoRa signal can be derived as $\mathbb{E}\left[Q^2(t) + I^2(t)\right] = A^2 + \mathbb{E}\left[Z_Q^2(t) + Z_I^2(t)\right]$, where $\mathbb{E}\left[Z_Q^2(t) + Z_I^2(t)\right]$ is the average noise power that can be measured when there is no LoRa signal. We use a `scipy` implementation of the differential evolution algorithm [92] to solve the least squares problem. Raspberry Pi uses about 0.7 seconds to solve it.

4.5.1.3 Performance comparison

We compare the FB estimation accuracy of the linear regression approach and the least squares approach. Fig. 4.17 shows the results. For each SNR setting, 20 LoRa I and Q traces with random FBs are generated using the signal model in Eq. (4.7). We also generate 20 noises traces; the magnitude of the noise is controlled to achieve the specified SNR. In Fig. 4.17, each error bar showing the 20%- and 80%-percentiles is from the 20 FB estimation results performed on the sum signals of the generated ideal LoRa signals and noise. From Fig. 4.17(a), the linear regression approach can achieve low FB estimation errors when the SNR is very high (e.g., 40 dB). However, it performs poorly for low SNRs. This is caused by the susceptibility of the inverse tangent rectification to noises. Specifically, as the inverse tangent rectification is based on a heuristic to detect atan2's sudden transitions between $-\pi$ and π, large noises leads to false positive detection of the transitions. Differently, the least squares approach maintains the FB estimation error within 120 Hz (i.e., 0.14 ppm), when the SNR is down to −18 dB. Thus, the rest of this chapter adopts the noise-resilient least squares approach, though it is more compute-intensive.

4.5.1.4 FB measurements for 16 end devices

We use an RTL-SDR to estimate the FBs of 16 SX1276-based end devices. In each test for an end device, the distance between the end device and the RTL-SDR is about 5 m. The error bars labeled "original" in Fig. 4.18 show the results. We can see that the FBs for a certain node are stable and the nodes generally have different FBs. The absolute FBs are from 17 kHz to 25 kHz, which are about 20 ppm to 29 ppm of the nominal central frequency of 869.75 MHz. Some nodes have similar FBs, e.g., Node 3, 8, and 14. Note that the detection of the replay attack is based on the fact that the replayed transmission has a different FB. In other words, the attack detection does not require distinct FBs among different end devices. From Fig. 4.18, we also observe that all nodes have negative FB measurements, which means that $\delta_{\mathrm{Tx}} < \delta_{\mathrm{Rx}}$, where δ_{Tx} and δ_{Rx} are the unknown FBs of the end device and the RTL-SDR. Note that as the RTL-SDR is a low-cost device, it may have a large FB causing the negative relative FB measurements.

4.5.2 Replay Attack Detection

The replayer also has an FB. The error bars labeled "replayed" in Fig. 4.18 show the FBs estimated from the LoRa signals received by the LoRaTS's SDR receiver when

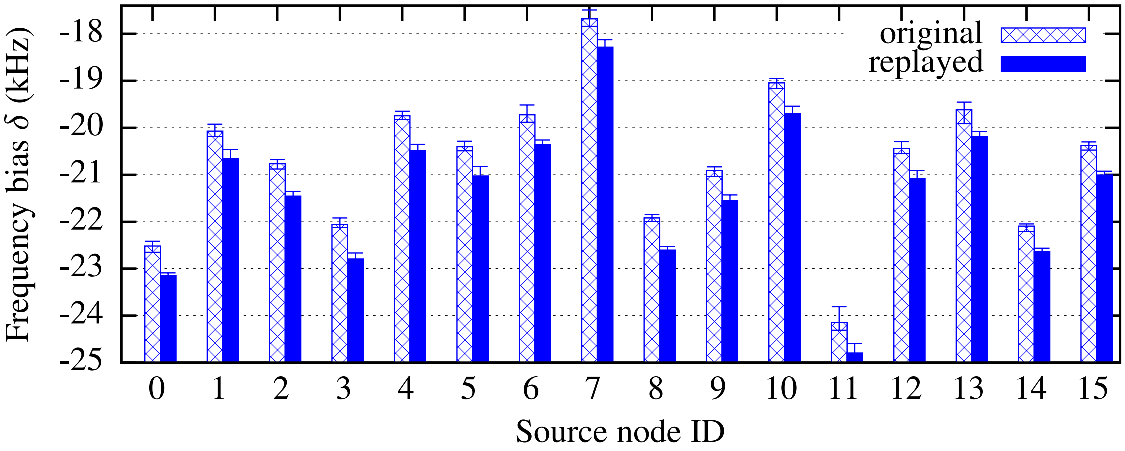

FIGURE 4.18: FBs estimated from the original LoRa signals from 16 end nodes and those replayed by a USRP-based replayer. The error bar shows mean, minimum, and maximum of FBs in 20 frame transmissions.

a USRP replays the radio waveform captured by itself in the experiments presented in §4.5.1. Compared with the results labeled "original", the FBs of the replayed transmissions are consistently lower. This is because the USRP has a negative FB. The average additional FBs introduced by the replayer range from −543 to −743 Hz, i.e., 0.62 to 0.85 ppm of the channel's central frequency. Thus, with the FB estimation accuracy of 0.14 ppm achieved under low SNRs (cf. §4.5.1.2), the additional FBs caused by the replay attack can be detected.

Based on the above observation, we describe an approach to detect the delayed replay. LoRaTS maintains a database of the FBs of the nodes with which it communicates. This database can be built offline or at run time using its SDR receiver in the absence of attacks. To address the end devices' time-varying radio frequency skews due to run-time conditions like temperature, LoRaTS can continuously update the database entries based on the FBs estimated from recent frames. To decide whether the current received frame is a replayed frame, the LoRaTS gateway checks whether the FB of the current received frame is within the acceptable FB range of the end device based on the database. This detection approach is applied after the LoRaTS gateway decodes the frame to obtain the end device ID. The FB estimated from a frame detected as a replayed one should not be used to update the database.

This detection mechanism forms a first line of defense against the frame delay attacks that introduce extra FBs. It gives awareness of the attack that is based on the logistics of collision and record-and-replay. With knowledge of our detector,

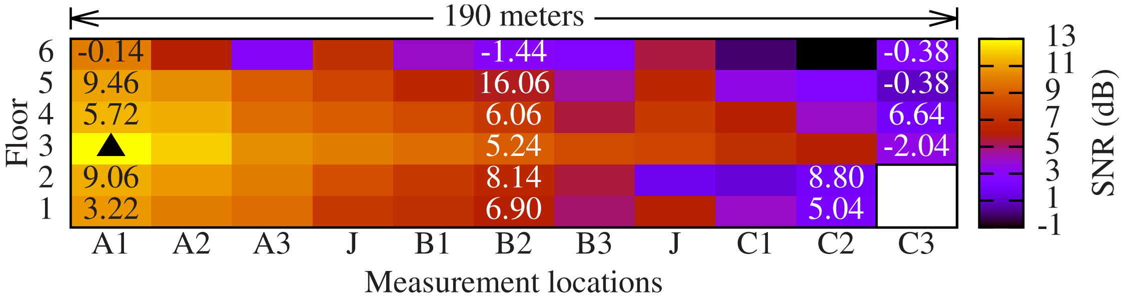

FIGURE 4.19: SNR survey in a building (lateral view) with 3 sections (A, B, C) and 2 junctions (J). The triangle represents the fixed node. The number in a cell is round-trip time measurement in μs excluding propagation time but containing onset time detection error when mobile node is in the cell.

the attackers may invest more resources and efforts to hide their radiometrics. §4.7 will discuss potential approaches to eliminate the extra FBs. While this attack-defense chase is interesting, in this chapter, we focus on showing the vulnerability of sync-free timestamping and propose the FB-based attack detector that forces the attackers to hide their radiometrics with increased cost and technical barriers.

4.6 Experiments

4.6.1 Experiments in a Multistory Building

LoRaWAN can be used for indoor applications, such as utility metering. We conduct a set of experiments to investigate the feasibility of attack and effectiveness of our attack detector in a concrete building with six floors. The building has three sections and two section junctions along its long dimension of 190 meters. Fig. 4.19 illustrates a lateral view of the building. First, we survey the SNR inside the building to understand the signal attenuation. We deploy a fixed LoRaWAN transmitter in Section A on the 3rd floor. Then, we carry an SDR receiver to different positions inside the building to measure the SNR. At each position, we first profile the noise power and then measure the total power when the fixed node transmits. In each section, we measure three positions. The heat map in Fig. 4.19 shows the SNR measurements. We can see that the SNR decays with the distance between the two nodes. The SNRs are from -1 dB to 13 dB. Then, we conduct the following experiments. By default, we set $S = 12$.

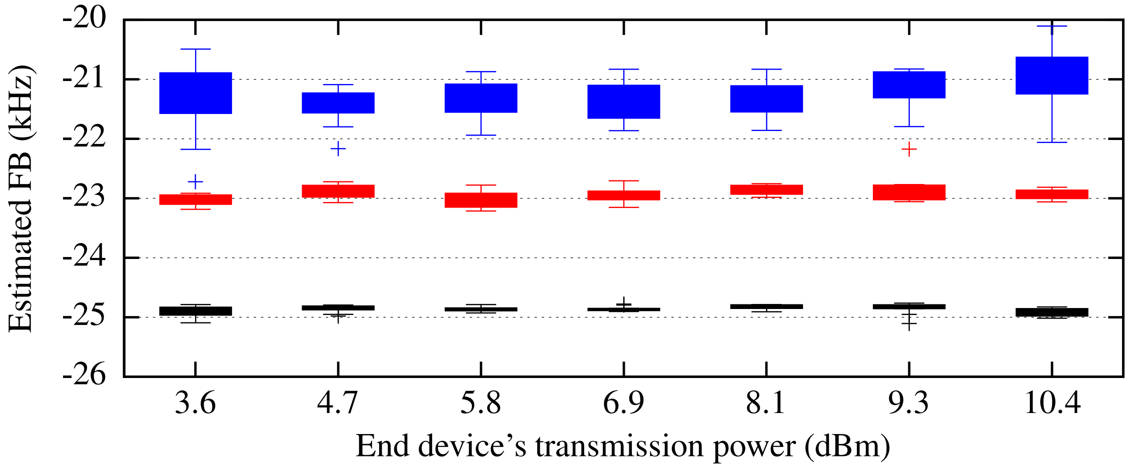

FIGURE 4.20: Estimated FB vs. transmitting power of the end device. Each box plot shows min, max, 25% and 75% percentiles. (1) Bottom row: end device to eavesdropper; (2) Mid row: end device to LoRaTS gateway; (3) Top row: replayer to LoRaTS.

Attack experiments: We deploy an iC880a-based gateway and an SX1276-based end device in Section A1 of the 3rd floor and Section C3 of the 6th floor, respectively. The LoRa signals are significantly attenuated after passing through multiple building floors. If the end device adopts a spreading factor of 7, it cannot communicate with the gateway. A minimum spreading factor of 8 is needed for communications. We deploy two USRP N210 stations as the eavesdropper and the collider, next to the end device and the gateway, respectively. We set the transmitting power of the end device and the collider to be 14 dBm. The malicious collision is stealthy to the gateway; the eavesdropping is successful. Thus, the frame delay attack can be launched in this building.

Onset time detection: We replace the iC880a gateway with our LoRaTS gateway and move the end device in the building. The number shown in a cell of Fig. 4.19 is the measured round-trip time Δ in μs excluding the propagation delay when the end device is at the corresponding location. Note that the propagation delay is calculated based on the estimated Euclidean distance between the gateway and the end device. As Δ contains onset time detection error, it may become negative. From the analysis in §4.4.3.4, the AIC's average RMSD(ϵ) in this building is 2.4 μs only. This result is consistent with that in Fig. 4.14.

Impact of transmitting power on FB estimation: Fig. 4.20 shows the estimated FBs versus the end device's transmitting power under different settings. The bottom row of black box plots are the FBs estimated by the eavesdropper when

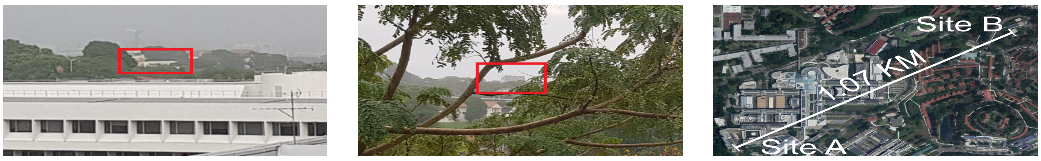

(a) View parklot from rooftop (b) View rooftop from parklot (c) The deployment sites

FIGURE 4.21: Pictures taken at the two sites (Site A: rooftop, Site B: parklot).

the end device transmits the uplink frame with different transmitting powers. The middle row of red box plots are the FBs estimated by the LoRaTS gateway in the absence of the frame collision and replay attacks. Thus, the FBs estimated by the eavesdropper and the LoRaTS gateway are different. This is because that as analyzed in §4.5.1, the estimated FB δ contains the transmitter's and receiver's FBs δ_{Tx} and δ_{Rx}. Note that the eavesdropper and the LoRaTS gateway in general have different FBs. From Fig. 4.20, the end device's transmitting power has little impact on the FB estimation.

Additional FB introduced by replayer: In Fig. 4.20, the top row of blue box plots are the FBs estimated by the LoRaTS when the replayer replays the radio waveform recorded by the eavesdropper. When the end device adopts a higher transmitting power, the replayed signal also has higher power. By comparing the middle and the top rows, we can see that the replay attack introduces an additional FB of about 2 kHz, which is 2.3 ppm of the LoRa channel's central frequency. Therefore, the FB monitoring can easily detect the replay attack. Compared with the results in Fig. 4.18 showing additional FBs of 0.62 to 0.85 ppm, the FBs in this set of experiments are higher. This is because that here we use two different USRPs as the eavesdropper and replayer; their FBs are superimposed.

4.6.2 Outdoor Experiments with Longer Distance

In this set of experiments, we deploy SX1276-based end devices in an outdoor parking lot. We replace the iC880a-based gateway shown in Fig. 4.6 with a LoRaTS gateway. The distance between the end device and the LoRaTS gateway is about 1.07 km. Fig. 4.21 shows the pictures taken at the two sites. The circled construct in a figure is the building where the other site is located in. The collider shown in Fig. 4.6 is also used in this set of experiments. The eavesdropper is deployed at a location about 200 m from the end device. When the transmitting powers

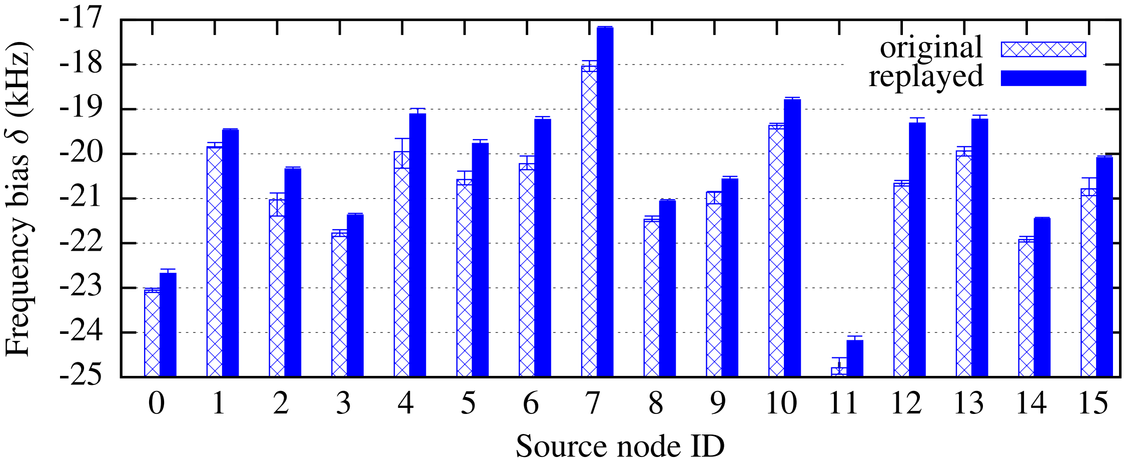

FIGURE 4.22: LoRaTS's FB estimates when the distances between the gateway and the end devices are about 1.07 km.

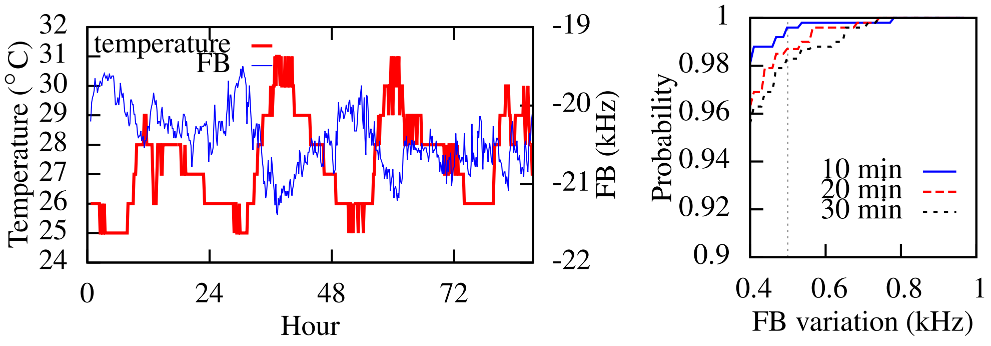

(a) FB of a device & temperature. (b) CDF of FB variation.

FIGURE 4.23: Temporal stability of SX1276's FB over 87 hours.

of the end device and the collider are 14 dBm and 8 dBm, respectively, we can successfully launch the frame delay attack. We also use the round-trip timing approach discussed in §4.4.3.4 to evaluate AIC's performance. The measurements show that AIC's RMSD(ϵ) is 1.29 μs only. This result is better than that obtained in the multistory building because the LoRa signal suffers significant attenuation in the indoor environment. Then, we investigate the additional FBs introduced by the replay attack. Fig. 4.22 shows LoRaTS's FB estimates for the frames transmitted by 16 end devices and the corresponding replays. The extra FBs introduced by the attack is up to 1.76 ppm. Thus, the attack can be detected.

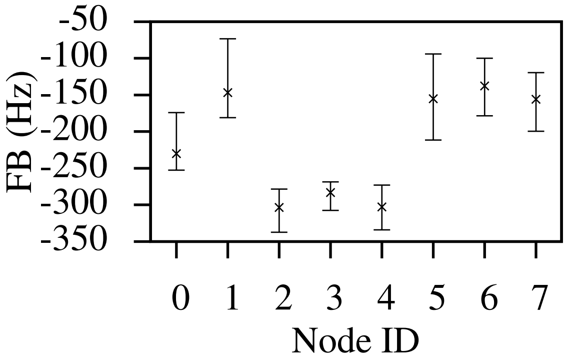

FIGURE 4.24: The FBs of eight SX1262 LoRa chips with TCXOs.

4.6.3 Temporal Stability of FB

FB can be affected by ambient condition such as temperature. We continuously track the FB of an SX1276-based end device for 87 hours to study its temporal stability. We place the end device with a temperature sensor in a semi-outdoor corridor with time-varying temperature. The end device transmits 10 frames every 10 minutes to the LoRaTS gateway as shown in Fig. 4.6, resulting 1,440 frames per day. Fig. 4.23(a) shows the end device's temperature and FB traces. The Pearson correlation between FB and temperature is -0.78. Moreover, the FB has transient variations that can be caused by interference from other communication systems operating in neighbor frequency bands. As LoRaTS detects the attack based on the changes of FB, such transient variations may cause false alarms. Fig. 4.23(b) shows the CDFs of the maximum FB variation if the end device transmits a frame every 10, 20, and 30 minutes. If the attack detection threshold is 500 Hz based on our previous measurements of the additional FB introduced by the attack, from the CDFs, the false alarm rate (i.e., the probability that the FB variation exceeds 500 Hz) is about 0.4%, 1.3%, and 1.7% for the three frame interval settings. The SX1276 used in this thesis does not have Temperature Compensated Crystal Oscillator (TCXO). For LoRa radios with TCXO, the false alarm rate can be further reduced. To verify this, we deploy eight SX1262-based end devices. SX1262 is the next-generation LoRa chip equipped with TCXO [93]. In Fig. 4.24, each error bar shows the 10%- and 90%-percentiles of 150 FB estimation results. We can see that the TCXO can significantly shrink the fluctuation. Specifically, from our measurements, the FB variations are no greater than 250 Hz. In contrast, from Fig. 4.23(a), without TCXO, the FB fluctuation range is up to 1 kHz. Thus, with a detection threshold of 500 Hz, the false alarm rate of our approach for a system based on SX1262 will be near-zero.

4.7 Discussions

Zero-FB attack: To bypass the proposed attack detector, the adversary needs to precisely calibrate its eavesdropper and replayer to have FBs lower than the resolution of our FB estimation algorithm. Such calibration requires a highly accurate (e.g., ppb level) frequency source operating at the channel frequency, which is non-trivial. The GPSDO module of USRP provides a GPS-locked reference clock of 10 MHz with 0.025 ppm accuracy [94]. While the non-integer scaling from 10 MHz to channel frequency may be subject to biases, the additional cost of two GPSDO modules (about US$1,800) is non-trivial for the eavesdropper and replayer to tune frequency accurately. There is also a possibility that the replayer's FB happens to cancel the eavesdropper's FB, rendering the superimposed FB zero. However, relying on such a random incident is an inefficient strategy for the attacker. Overall, the proposed low-cost (US$25 for RTL-SDR) attack detector significantly increases the cost and technical barrier of attack.

Timestamp recovery: Recovering timestamp under attack is challenging and needs further study. A recent concurrent LoRa demodulator [95] may not work for this purpose because it requires time-misalignment between two concurrent frames. The attacker can reduce the time-misalignment.

Chapter 5

Conclusion and Future Work

This thesis exploits LoRaWAN to improve the efficiency and resilience of IoT networks. The first part of this thesis studied how to use LPWAN radios to form one-hop out-of-band control planes for multi-hop wireless networks through extensive simulational and hardware profiling studies, system design, and testbed evaluation. Through an example application of CTP, the simulational study showed the advantages of the one-hop out-of-band control plane design over the distributed network control scheme and the in-band control plane design. The hardware profiling study based on two LPWAN technologies, i.e., LoRaWAN and CC1352R's sub-GHz long-range radio, showed their characteristics that should be considered in the design of the out-of-band control planes. The designed LoRaWAN-based control plane, LoRaCP, is applied to physically separate the control plane of CTP from its Zigbee-based data-plane network. Experiments showed that LoRaCP increases CTP's packet delivery ratio from 65% to 80% in the presence of external interference, while consuming a per-node average radio current of 0.9 mA only with an operating voltage of 3.3 V. This current consumption is comparable to or lower than those of low-power MCUs found on sensor network platforms such as TelosB and Firestorm.

The second part of this thesis showed that sync-free data timestamping, though bandwidth-efficient, is susceptible to the frame delay attack that can be implemented by a combination of frame collision and delayed replay. Experiments show that the attack can affect many end devices in a large geographic area. As the attack does not need to break the cryptographic protection of the frame, existing security measures prescribed by LoRaWAN cannot counteract this threat. To gain attack awareness, this thesis designs a gateway called LoRaTS that integrates a low-power SDR receiver with a commodity LoRaWAN gateway. This thesis models the reception of LoRa signal using an SDR receiver and investigates the algorithms that can accurately estimate the uplink frame arrival time. It develops efficient time-domain signal processing algorithms to estimate the FBs of the end devices. The least squares FB estimation algorithm achieves high resolution and can uncover the additional FBs introduced by the attack. In summary, LoRaTS achieves efficient sync-free data timestamping with awareness of frame delay attack.

In the future, it is interesting to fully implement one-hop out-of-band control planes based on TI's sub-GHz radios. As the Contiki-NG operating system has better support of new hardware, the design of a Contiki-NG code library that provides well-defined interfaces to facilitate the implementation of one-hop out-of-band control planes will be of meaningful contribution. Moreover, how to let LoRa signal survive frame delay attack more than being aware of the attack only is a nontrivial issue. The design of a full-fledged secure system for LoRa will be of valuable contribution.

LoRaWAN has received increasing interest in both academia and industry. Its modulation approach facilitates its robustness to noise and long-range communication capability. Meanwhile, as a trade-off, LoRaWAN has a limited transmission speed. Moreover, governments in different countries and regions set restrictions on duty cycles due to its use of ISM band. As a result, it is desirable to design a scheme that enables multiple LoRaWAN nodes to transmit at the same time using the same spreading factor. While existing studies on LoRaWAN has focused on network connectivity and performance, accurate positioning of LoRa end devices is still largely an open issue. Recently, LoRaWAN is available on 2.4 GHz frequency band, comparing with other 2.4 GHz communication radios (e.g., Wi-Fi, BLE, ZigBee), LoRaWAN outperforms them regarding communication range, power consumption, and deployment cost. However, it is not easy to apply LoRaWAN for accurate positioning due to its narrow-band nature. And LoRaWAN requires gateways to support device positioning. Two possible research directions are using Multi-Input Multi-Output (MIMO) to improve positioning accuracy and using Commercial Off-The-Shelf (COTS) Wi-Fi routers to improve positioning universality. Beyond the technological development, how to incentivize individuals, organizations to deploy their LPWAN systems, and form a global IoT infrastructure is also important. The economical models to drive the global IoT infrastructure are still open questions. Helium [96] pioneers this space by applying cryptocurrency to LPWAN-based IoT networks. Research can be done in this space such as secure sensing data sharing, secure federated learning, and sensing data pricing.

As IoT is a data generation infrastructure, applying machine learning in IoT is desirable. Comparing with the edge or cloud, IoT end devices do not have enough computation power to handle sophisticated patterns. To solve this problem, one approach is offloading data to the backend to execute the deep neural networks for inference. Another approach is making the use of the coalition of the IoT objects to achieve better intelligence (e.g., via federated learning). Both the geographic separation of data sources and the computation power, and the distribution of data among geographically distributed IoT objects need the support of the communication networks. Thus, connected AI is important in IoT due to the above observations. Research can be done to understand the mutual impact of the communication networks and the connected AI.

List of Author's Publications

Conference papers

- **Chaojie Gu**, Rui Tan, Xin Lou, Dusit Niyato, "One-Hop Out-of-Band Control Planes for Low-Power Multi-Hop Wireless Networks", in *Proceedings of the 37th Annual IEEE International Conference on Computer Communications (INFOCOM 2018)*, April 15-19, 2018, Honolulu, HI.
- **Chaojie Gu**, Linshan Jiang, Rui Tan, Mo Li, Jun Huang. Attack-Aware Data Timestamping in Low-Power Synchronization-Free LoRaWAN. In *Proceedings of the 40th IEEE International Conference on Distributed Computing Systems (ICDCS 2020)*, 2020, Singapore.
- Amalinda Gamage, Jansen C. Liando, **Chaojie Gu**, Rui Tan, Mo Li.LMAC: Efficient Carrier-Sense Multiple Access for LoRa. In *The 26th Annual International Conference on Mobile Computing and Networking (MobiCom 2020)*, September 21-25, 2020, London, United Kingdom.

Journal articles

- Qun Song, **Chaojie Gu**, Rui Tan. "Deep Room Recognition Using Inaudible Echos", In *Proceedings of the 2018 ACM International Joint Conference and 2018 International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers (UbiComp 2018)*, October 8-12, 2018, Singapore.
- **Chaojie Gu**, Rui Tan, Xin Lou, "One-Hop Out-of-Band Control Planes for Multi-Hop Wireless Sensor Networks", *ACM Transactions on Sensor Networks*, 2019. Article 40.

- Dixing Xu, Mengyao Zheng, Linshan Jiang, **Chaojie Gu**, Rui Tan, Peng Cheng, "Lightweight and Unobtrusive Data Obfuscation at IoT Edge for Remote Inference", *IEEE Internet of Things Journal*, Special Issue on Artificial Intelligence Powered Edge Computing for Internet of Things.

Workshop papers

- **Chaojie Gu**, Linshan Jiang, Jun Huang, "LoRa-Based Localization: Opportunities and Challenges", in *The 1st workshop on Low Power Wide Area Networks for Internet of Things (LPNET)* with EWSN 2019, February 25, 2019, Beijing, China.
- Mengyao Zheng, Dixing Xu, Linshan Jiang, **Chaojie Gu**, Rui Tan, Peng Cheng. "Challenges of Privacy-Preserving Machine Learning in IoT", *The 1st International Workshop on Challenges in Artificial Intelligence and Machine Learning for Internet of Things (AIChallengeIoT)* with SenSys 2019, November 10, 2019, New York, NY, USA.

Others

- **Chaojie Gu**, Rui Tan, Jun Huang, "Poster Abstract: SoftLora-A LoRa-Based Platform for Accurate and Secure Timing", in *Proceedings of the 18th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN 2019)*, April 15-18, 2019, Montreal, QC, Canada.

Bibliography

[1] The future of iot: 10 predictions about the internet of things, 2020. URL `https://us.norton.com/internetsecurity-iot-5-predictions-for-the-future-of-iot.html`. 1

[2] Dixing Xu, Mengyao Zheng, Linshan Jiang, Chaojie Gu, Rui Tan, and Peng Cheng. Lightweight and unobtrusive data obfuscation at iot edge for remote inference. *IEEE Internet of Things Journal*, 2020. 2

[3] N. Sornin, M. Luis, T. Eirich, T. Kramp, and O. Hersent. Lorawan^{TM}specification (v1.0.2), 2016. 3, 72

[4] Bill Zalud. How mesh networks form the backbone of smart cities, Dec 2013. URL `https://www.securitymagazine.com/articles/84986-how-mesh-networks-form-the-backbone-of-smart-cities`. 4

[5] Veniam. Creating the world's largest network of connected vehicles for smart cities, Apr 2016. URL `https://www.worldwifiday.com/wp-content/uploads/2016/05/3.-PortoCaseStudy_Letter_2016-04-15.pdf`. 4

[6] Inc. Bluetooth SIG. Bluetooth LE: mesh, 2019. URL `https://www.bluetooth.com/bluetooth-technology/topology-options`. 4

[7] Emerson Process Management. Industrial wireless technology, 2020. URL `https://www.emerson.com/en-us/expertise/automation/industrial-internet-things/pervasive-sensing-solutions/wireless-technology`. 4, 5, 25

[8] Omprakash Gnawali, Rodrigo Fonseca, Kyle Jamieson, David Moss, and Philip Levis. Collection tree protocol. In *The 7th ACM conference on embedded networked sensor systems (SenSys)*, pages 1–14. ACM, 2009. 4, 7, 22, 29

[9] Huawei Huang, Song Guo, Weifa Liang, Keqiu Li, Baoliu Ye, and Weihua Zhuang. Near-optimal routing protection for in-band software-defined heterogeneous networks. *IEEE Journal on Selected Areas in Communications (J-SAC)*, 34(11):2918–2934, 2016. 5, 26

[10] Michael P Andersen, Gabe Fierro, and David E Culler. System design for a synergistic, low power mote/ble embedded platform. In *The 15th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 1–12. IEEE, 2016. 5, 7, 52

[11] S.L. Openmote Technologies. Openmote, 2018. URL http://www.openmote.com/. 5

[12] Pycom. Lopy4, 2019. URL https://pycom.io/product/lopy4/. 6

[13] Texas Instruments Incorporated. Ti cc1350, 2018. URL http://www.ti.com/product/CC1350. 6

[14] Texas Instruments Incorporated. Ti sensortag, 2018. URL http://www.ti.com/ww/en/wireless_connectivity/sensortag/. 6

[15] Texas Instruments Incorporated. Cc1352r, 2018. URL http://www.ti.com/product/CC1352R. 6

[16] Anders R Jensen, Mads Lauridsen, Preben Mogensen, Troels B Sørensen, and Per Jensen. Lte ue power consumption model: For system level energy and performance optimization. In *2012 IEEE Vehicular Technology Conference (VTC Fall)*, pages 1–5. IEEE, 2012. 6

[17] Jinzhu Chen, Rui Tan, Yu Wang, Guoliang Xing, Xiaorui Wang, Xiaodong Wang, Bill Punch, and Dirk Colbry. A high-fidelity temperature distribution forecasting system for data centers. In *IEEE RTSS*, 2012. 8

[18] Seri Oh, Stephen G Ritchie, and Cheol Oh. Real-time traffic measurement from single loop inductive signatures. *Transportation Research Record*, 1804 (1):98–106, 2002. 8

[19] Fog computing for industrial automation, 2019. https://www.controleng.com/articles/fog-computing-for-industrial-automation/. 8

[20] Guojin Liu, Rui Tan, Ruogu Zhou, Guoliang Xing, Wen-Zhan Song, and Jonathan M Lees. Volcanic earthquake timing using wireless sensor networks. In *IPSN*, 2013. 8, 75

[21] RTL-SDR, 2019. https://www.rtl-sdr.com/. 10, 69

[22] Chaojie Gu, Linshan Jiang, and Rui Tan. Lora-based localization: Opportunities and challenges. In *Proceedings of the 2019 International Conference on Embedded Wireless Systems and Networks*, pages 413–418, 2019. 13

[23] Chaojie Gu, Rui Tan, Xin Lou, and Dusit Niyato. One-hop out-of-band control planes for low-power multi-hop wireless networks. In *INFOCOM*, 2018. 13, 19

[24] Chaojie Gu, Rui Tan, and Xin Lou. One-hop out-of-band control planes for multi-hop wireless sensor networks. *ACM Transactions on Sensor Networks (TOSN)*, 15(4):1–29, 2019. 13, 19

[25] LoRa Alliance. Lora alliance, 2019. URL https://www.lora-alliance.org/. 13

[26] Sigfox. Sigfox, 2019. URL https://www.sigfox.com/. 13

[27] Weightless SIG. Weightless, 2015. URL `http://www.weightless.org/`. 13

[28] Narrowband internet of things (nb-iot), 2020. `https://www.gsma.com/iot/narrow-band-internet-of-things-nb-iot/`. 13

[29] Semtech. Sx1276/77/78/79 – 137mhz to 1020mhz low power long range transceiver, 2019. URL `https://www.semtech.com/uploads/documents/DS_SX1276-7-8-9_W_APP_V6.pdf`. 16

[30] Paul J Marcelis, Vijay S Rao, and R Venkatesha Prasad. Dare: Data recovery through application layer coding for lorawan. In *IEEE/ACM Second International Conference on Internet-of-Things Design and Implementation (IoTDI)*, pages 97–108. IEEE, 2017. 16, 21, 37

[31] Karim Habak, Khaled A Harras, and Moustafa Youssef. Bandwidth aggregation techniques in heterogeneous multi-homed devices: A survey. *Computer Networks*, 92:168–188, 2015. 19, 20

[32] Atul Adya, Paramvir Bahl, Jitendra Padhye, Alec Wolman, and Lidong Zhou. A multi-radio unification protocol for ieee 802.11 wireless networks. In *The 1st International Conference on Broadband Networks (BroadNets)*, pages 344–354. IEEE, 2004. 20

[33] Paramvir Bahl, Atul Adya, Jitendra Padhye, and Alec Walman. Reconsidering wireless systems with multiple radios. *ACM SIGCOMM Computer Communication Review*, 34(5):39–46, 2004. 20

[34] Srikanth Kandula, Kate Ching-Ju Lin, Tural Badirkhanli, and Dina Katabi. Fatvap: Aggregating ap backhaul capacity to maximize throughput. In *The 5th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, volume 8, pages 89–104, 2008. 20

[35] Gholamhossein Ekbatanifard, Philipp Sommer, Branislav Kusy, Venkat Iyer, and Koen Langendoen. Fastforward: High-throughput dual-radio streaming. In *IEEE 10th International Conference on Mobile Ad-Hoc and Sensor Systems (MASS)*, pages 209–213. IEEE, 2013. 20

[36] Dimitrios Lymberopoulos, Nissanka B Priyantha, Michel Goraczko, and Feng Zhao. Towards energy efficient design of multi-radio platforms for wireless sensor networks. In *The 7th international conference on Information processing in sensor networks (IPSN)*, pages 257–268. IEEE, 2008. 20

[37] Lorenzo Keller, Anh Le, Blerim Cici, Hulya Seferoglu, Christina Fragouli, and Athina Markopoulou. Microcast: Cooperative video streaming on smartphones. In *The 10th international conference on Mobile systems, applications, and services (MobiSys)*, pages 57–70. ACM, 2012. 20

[38] Shahriar Nirjon, Angela Nicoara, Cheng-Hsin Hsu, Jatinder Pal Singh, and John A Stankovic. MultiNets: A system for real-time switching between multiple network interfaces on mobile devices. *ACM Transactions on Embedded Computing Systems (TECS)*, 13(4s):121, 2014. 20

[39] Di Mu, Yunpeng Ge, Mo Sha, Steve Paul, Niranjan Ravichandra, and Souma Chowdhury. Adaptive radio and transmission power selection for internet of things. In *IEEE/ACM 25th International Symposium on Quality of Service (IWQoS)*, pages 1–10. IEEE, 2017. 20

[40] Yeon-sup Lim, Yung-Chih Chen, Erich M Nahum, Don Towsley, Richard J Gibbens, and Emmanuel Cecchet. Design, implementation, and evaluation of energy-aware multi-path tcp. In *The 11th ACM Conference on Emerging Networking Experiments and Technologies (CoNEXT)*, page 30. ACM, 2015. 20

[41] Ashkan Nikravesh, Yihua Guo, Feng Qian, Z Morley Mao, and Subhabrata Sen. An in-depth understanding of multipath tcp on mobile devices: measurement and system design. In *The 22nd Annual International Conference on Mobile Computing and Networking (MobiCom)*, pages 189–201. ACM, 2016. 20

[42] Bruno Astuto A Nunes, Marc Mendonca, Xuan-Nam Nguyen, Katia Obraczka, and Thierry Turletti. A survey of software-defined networking: Past, present, and future of programmable networks. *IEEE Communications Surveys & Tutorials*, 16(3):1617–1634, 2014. 20

[43] Nachikethas A Jagadeesan and Bhaskar Krishnamachari. Software-defined networking paradigms in wireless networks: a survey. *ACM Computing Surveys (CSUR)*, 47(2):27, 2015. 20

[44] Peter Dely, Andreas Kassler, and Nico Bayer. OpenFlow for wireless mesh networks. In *The 20th International Conference on Computer Communications and Networks (ICCCN)*, pages 1–6. IEEE, 2011. 20

[45] Tie Luo, Hwee-Pink Tan, and Tony QS Quek. Sensor openflow: Enabling software-defined wireless sensor networks. *IEEE Communications Letters*, 16 (11):1896–1899, 2012. 20

[46] Murad Kaplan, Chenyu Zheng, Matthew Monaco, Eric Keller, and Douglas Sicker. WASP: a software-defined communication layer for hybrid wireless networks. In *ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS)*, pages 5–15. IEEE, 2014. 20

[47] T. Petri, M. Goessens, L. Nuaymi, L. Toutain, and A. Pelov. Measurements, performance and analysis of lora fabian, a real-world implementation of lpwan. In *IEEE 27th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*, pages 1–7, Sept 2016. 21

[48] Viktor Toldov, J.P. Meijers, Roman Igual-Perez, Riaan Wolhuter, Nathalie Mitton, and Laurent Clavier. Performance evaluation of lora radio solution for prednet wildlife animal tracking project. In *1st Annual LPWAN Conference*, 2016. 21

[49] 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 (TOSN)*, 15(2):1–35, 2019. 21, 37, 38

[50] Viktor Toldov, Laurent Clavier, and Nathalie Mitton. Multi-channel distributed mac protocol for wsn-based wildlife monitoring. In *2018 14th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*, pages 1–8. IEEE, 2018. 21

[51] 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 *The 18th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, 2019. 21, 39, 63, 65

[52] Rashad Eletreby, Diana Zhang, Swarun Kumar, and Osman Yağan. Empowering low-power wide area networks in urban settings. In *The Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*, pages 309–321. ACM, 2017. 21, 55, 56, 77

[53] 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 *The 17th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 60–71. IEEE, 2018. 21

[54] Mehrdad Hessar, Ali Najafi, and Shyamnath Gollakota. Netscatter: Enabling large-scale backscatter networks. In *The 16th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. USENIX, 2019. 21

[55] Yao Peng, Longfei Shangguan, Yue Hu, Yujie Qian, Xianshang Lin, Xiaojiang Chen, Dingyi Fang, and Kyle Jamieson. Plora: a passive long-range data network from ambient lora transmissions. In *The 2018 Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*, pages 147–160. ACM, 2018. 56

[56] Vamsi Talla, Mehrdad Hessar, Bryce Kellogg, Ali Najafi, Joshua R Smith, and Shyamnath Gollakota. Lora backscatter: Enabling the vision of ubiquitous connectivity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 1(3):105, 2017.

[57] Ambuj Varshney, Oliver Harms, Carlos Pérez-Penichet, Christian Rohner, Frederik Hermans, and Thiemo Voigt. Lorea: A backscatter architecture that achieves a long communication range. In *The 15th ACM Conference on Embedded Network Sensor Systems (SenSys)*, page 18. ACM, 2017. 21

[58] Philip Levis, Nelson Lee, Matt Welsh, and David Culler. TOSSIM: Accurate and scalable simulation of entire tinyos applications. In *The 1st international*

conference on Embedded networked sensor systems (SenSys), pages 126–137. ACM, 2003. 21

[59] Tinyos production, 2018. URL https://github.com/tinyprod. 22

[60] iC880A - LoRaWAN concentrator 868mhz, 2019. URL https://wireless-solutions.de/products/radiomodules/ic880a.html. 32

[61] Microchip. RN2483: Low-Power Long Range LoRa Technology Transceiver Module, 2019. URL http://ww1.microchip.com/downloads/en/devicedoc/50002346c.pdf. 34

[62] Microchip. Rn2483, 2019. URL https://www.microchip.com/wwwproducts/en/RN2483. 34

[63] Inc. Monsoon Solutions. Monsoon power monitor", 2019. URL http://msoon.github.io/powermonitor/PowerTool/doc/LVPM%20Manual.pdf. 34

[64] Modtronix Engineering. Wireless SX1276 LoRa module, 2018. URL http://modtronix.com/inair9b.html. 37

[65] Texas Instruments Incorporated. Simplelink, 2019. URL http://www.ti.com/wireless-connectivity/simplelink-solutions/overview/overview.html. 39

[66] Contiki. Contiki-ng, 2019. URL http://www.contiki-ng.org/. 39

[67] RIOT. Riot: The friendly operating system for the internet of things, 2013. URL https://riot-os.org/. 39

[68] Pietro Boccadoro, Michele Barile, Giuseppe Piro, and Luigi Alfredo Grieco. Energy consumption analysis of tsch-enabled platforms for the industrial-iot. In *IEEE 2nd International Forum on Research and Technologies for Society and Industry Leveraging a better tomorrow (RTSI)*, pages 1–5. IEEE, 2016. 40

[69] Pietro Boccadoro, Giuseppe Piro, Domenico Striccoli, and Luigi Alfredo Grieco. Experimental comparison of industrial internet of things protocol stacks in time slotted channel hopping scenarios. In *IEEE International Conference on Communications (ICC)*, pages 1–6. IEEE, 2018. 40

[70] The European Telecommunications Standards Institute. Technical characteristics for low power wide area networks chirp spread spectrum (lpwan-css) operating in the uhf spectrum below 1 ghz, 2019. URL https://www.etsi.org/docdeliver/etsi_tr/103500_103599/103526/01.01.01_60/tr_103526v010101p.docx. 43, 58

[71] Federal Communications Commission. Fcc regulations for ism band devices: 902 - 928 mhz, 2019. URL https://www.semtech.com/uploads/documents/fcc_part15_regulations_semtech.pdf. 43

[72] Semtech. Sx1272/3/6/7/8: Lora modem designer's guide, 2019. URL `https://www.rs-online.com/designspark/rel-assets/ds-assets/uploads/knowledge-items/application-notes-for-the-internet-of-things/LoRa%20Design%20Guide.pdf`. 44

[73] Orne Brocaar. LoRa server system architecture, 2019. URL `https://www.loraserver.io/overview/architecture/`. 45, 46, 51

[74] Info-communications Media Development Authority of Singapore. Spectrum management handbook (issue 1 rev 2.9 - july 2017, 2017. URL `https://www.imda.gov.sg/-/media/imda/files/regulation-licensing-and-consultations/frameworks-and-policies/spectrum-management-and-coordination/spectrummgmthb.pdf?la=en`. 49

[75] Chaojie Gu, Linshan Jiang, Rui Tan, Mo Li, and Jun Huang. Attack-aware data timestamping in low-power synchronization-free lorawan. In *IEEE ICDCS*, 2020. 55, 57

[76] A. Dongare, R. Narayanan, A. Gadre, A. Luong, A. Balanuta, S. Kumar, B. Iannucci, and A. Rowe. Charm: Exploiting geographical diversity through coherent combing in low-power wide-area networks. In *IPSN*, 2018. 55, 56

[77] Mehrdad Hessar, Ali Najafi, and Shyamnath Gollakota. NetScatter: Enabling large-scale backscatter networks. In *NSDI*, 2019. 56

[78] Ceferino Gabriel Ramirez, Anton Sergeyev, Assya Dyussenova, and Bob Iannucci. Longshot: long-range synchronization of time. In *IPSN*, 2019. 56

[79] E. Aras, G. S. Ramachandran, P. Lawrence, and D. Hughes. Exploring the security vulnerabilities of lora. In *CYBCONF*, pages 1–6, June 2017. 56

[80] Emekcan Aras, Nicolas Small, Gowri Sankar Ramachandran, Stephane Delbruel, Wouter Joosen, and Danny Hughes. Selective jamming of lorawan using commodity hardware. In *MobiQuitous*, 2017. 56

[81] Pieter Robyns, Eduard Marin, Wim Lamotte, Peter Quax, Dave Singelée, and Bart Preneel. Physical-layer fingerprinting of lora devices using supervised and zero-shot learning. In *WiSec*, 2017. 56

[82] Vladimir Brik, Suman Banerjee, Marco Gruteser, and Sangho Oh. Wireless device identification with radiometric signatures. In *MobiCom*, 2008. 57

[83] Vector signal analyzer, 2019. `https://www.keysight.com/en/pc-2416877/vector-signal-analyzer?cc=US&lc=eng`. 57

[84] T. Hao, R. Zhou, G. Xing, and M. Mutka. WizSync: Exploiting wi-fi infrastructure for clock synchronization in wireless sensor networks. In *RTSS*, 2011. 58

[85] gr-lora, 2019. https://github.com/rpp0/gr-lora. 64

[86] Amazon. Rtl-sdr blog r820t2 rtl2832u 1ppm tcxo sma software defined radio (dongle only), 2020. https://www.amazon.com/dp/B0129EBDS2. 69

[87] Ruogu Zhou, Yongping Xiong, Guoliang Xing, Limin Sun, and Jian Ma. Zifi: wireless lan discovery via zigbee interference signatures. In *MobiCom*, pages 49–60, 2010. 74

[88] RVE Lovelace, JM Sutton, and EE Salpeter. Digital search methods for pulsars. *Nature*, 222(5190):231–233, 1969. 74

[89] Reinoud Sleeman and Torild Van Eck. Robust automatic p-phase picking: an on-line implementation in the analysis of broadband seismogram recordings. *Physics of the earth and planetary interiors*, 113(1-4):265–275, 1999. 74, 75

[90] Yingwei Yao and Georgios B Giannakis. Blind carrier frequency offset estimation in siso, mimo, and multiuser ofdm systems. *IEEE Transactions on Communications*, 53(1):173–183, 2005. 76

[91] Semtech Ltd. Sx1276/77/78 - 137 mhz to 1020 mhz low power long range transceiver, 2018. http://modtronix.com/prod/components/wireless/sx1276.pdf. 80

[92] Rainer Storn and Kenneth Price. Differential evolution–a simple and efficient heuristic for global optimization over continuous spaces. *Journal of global optimization*, 11(4):341–359, 1997. 80

[93] Wireless SX1262 LoRa module, 2020. URL http://www.ebyte.com/en/product-view-news.aspx?id=437. 87

[94] Gpsdo, 2019. https://www.ettus.com/wp-content/uploads/2019/01/gpsdo-kit_4.pdf. 88

[95] Xianjin Xia, Yuanqing Zheng, and Tao Gu. Ftrack: Parallel decoding for lora transmissions. In *SenSys*, 2019. 88

[96] Helium, 2020. URL https://www.helium.com/. 91