Combating Link Dynamics for Reliable LoRa Connection in Urban Settings

Shuai Tong, Zilin Shen, Yunhao Liu, Jiliang Wang
School of Software and BNRist, Tsinghua University, Beijing, P.R. China
{tl19, szl17}@mails.tsinghua.edu.cn, {yunhao, jiliangwang}@tsinghua.edu.cn

ABSTRACT

LoRa, as a representative Low-Power Wide-Area Network (LPWAN) technology, can provide long-range communication for battery-powered IoT devices with a 10-year lifetime. LoRa links in practice, however, experience high dynamics in various environments. When the SNR falls below the threshold (e.g., in the building), a LoRa device disconnects from the network. We propose Falcon, which addresses the link dynamics by enabling data transmission for very low SNR or even disconnected LoRa links. At the heart of Falcon, we reveal that low SNR LoRa links that cannot deliver packets can still introduce interference to other LoRa transmissions. Therefore, Falcon transmits data bits on the low SNR link by selectively interfering with other LoRa transmissions. We address practical challenges in Falcon design. We propose a low-power channel activity detection method to detect other LoRa transmissions for selective interference. To interfere with the so-called interference-resilient LoRa, we accurately estimate the time and frequency offsets on LoRa packets and propose an adaptive frequency adjusting strategy to maximize the interference. We implement Falcon, all using commercial off-the-shelf LoRa devices, and extensively evaluate its performance. The results show that Falcon can provide reliable communication links for disconnected LoRa devices and achieves the SNR boundary upto 7.5 dB lower than that of standard LoRa.

CCS CONCEPTS

• **Networks** → **Network protocol design**.

KEYWORDS

Internet of Things, LPWAN, LoRa, selective interfering

ACM Reference Format:
Shuai Tong, Zilin Shen, Yunhao Liu, Jiliang Wang. 2022. Combating Link Dynamics for Reliable LoRa Connection in Urban Settings. In *The 27th Annual International Conference on Mobile Computing and Networking (ACM MobiCom '21), January 31-February 4, 2022, New Orleans, LA, USA.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3447993.3483250

1 INTRODUCTION

Recent years have witnessed the widespread deployment of Low-Power Wide-Area Networks (LPWANs) for connecting large-scale devices in the Internet of Things (IoT). As a representative LPWAN technique, LoRa is adopted by a wide spectrum of low data rate, large scale, and long lifetime IoT applications, including intelligent agriculture [1], animal tracking [2], disaster rescue [3], warehouse management [4], etc. One of the key features of LoRa is the long-range communication capability of up to tens of kilometers [5–7]. However, LoRa links in practice suffer dramatically from dynamic environments. When the SNR falls below the threshold for the lowest data rate of LoRa (i.e., SF12 and CR=4/8), the LoRa link has no data rate, and the device disconnects from the network. In this paper, we ask –"Can we provide reliable communication for commodity LoRa clients even if they are not reachable over the standard LoRa protocol?"

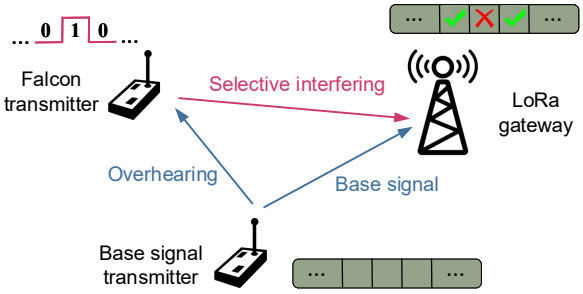

Figure 1: An illustration of Falcon: A Falcon transmitter delivers data bits to a LoRa gateway by selectively interfering the base signal transmissions.

Existing Approaches. Reliable connections are very important before applying LoRa as the main technique for connecting massive IoT devices. Although there is a large collection of literature for improving the reliability of LoRa connections, they have limitations for requiring multi-receiver collaboration or cloud assistance. For example, Charm [8] and Nephalai [9] adopt the cloud radio access network (C-RAN) for exploiting the spatial diversity gain on LoRa coverage enhancement. They require 2-8 collaborative gateways for collecting and offloading PHY samples to cloud servers. OPR [10] gathers MAC symbols from 2-6 gateways to correct corrupt bits on the cloud. Chime [11] estimates wireless channels on 4-6 gateways to identify the optimal operating frequency. Choir [12] collaborates with tens of LoRa clients to transmit coherent data to reach a remote LoRa gateway.

Fundamental limitations: Existing weak signal decoding approaches [8–12] aim to recover original coding information from the received weak LoRa signals. Those approaches mainly focus on enhancing and decoding the low-SNR signal. They have strong requirements on hardwares such as multi-gateway collaborations and cloud assistance, which incurs noneligible deployment costs. Moreover, deploying multiple gateways to maximize the coverage

ACM MobiCom '21, January 31-February 4, 2022, New Orleans, LA, USA

ACM ISBN 978-1-4503-8342-4/22/01…$15.00
https://doi.org/10.1145/3447993.3483250

area, known as a k-coverage problem, is not trivial. This work aims to provide reliable communication links for disconnected LoRa clients with as little cost as possible from the interference perspective.

Our Approach. This paper proposes Falcon, a physical layer mechanism to enable communication links for low-SNR LoRa clients that cannot connect to gateways using the standard LoRa protocol. At the heart of Falcon, we reveal that LoRa signals that are too weak to be decoded can cause interference to other LoRa transmissions. Based on this observation, Falcon provides emergency links for unreachable LoRa clients by making them selectively interfere with other LoRa transmissions on the same channel. Fig. 1 illustrates the principle of Falcon. To start up the transmission, Falcon leverages a base signal transmitter to generate reference packets. This transmitter can be either a LoRa client or a gateway as long as it can emit base signal packets in the form of chirps. Then, when a low-SNR LoRa client (namely Falcon transmitter) has to communicate to the gateway with Falcon, it overhears the base signal over the channel and synchronizes its own time and frequency with the base signal. For delivering data bits, the Falcon transmitter either does nothing, representing bit '0', or it corrupts a base signal packet, representing bit '1'. LoRa gateways extract data bits from Falcon transmitters by determining the SNR of the interfered signal or whether a base signal packet can be successfully decoded or not. Therefore, Falcon can enable standard-compliant communication for low-SNR LoRa clients using COTS LoRa devices without requiring any receiver collaboration or hardware modification.

Design Challenges. We incorporate several novel designs for Falcon to address practical challenges. First, we have to enable power-efficient base signal detection for Falcon transmitters powered by batteries. The main challenge is the limited power budget of Falcon transmitters and the low-SNR property of base signals. We design an effective method for based signal detection using the LoRa Channel Activity Detection (CAD). Second, we need to ensure the interfering signals from Falcon transmitters are synchronized with base signals both in time and frequency. The interfering performance is inherently prone to synchronization errors. We propose a reception-based approach that achieves time and frequency synchronization with only a single reference reception. Third, we need to interfere with the so-called interference resilient LoRa signal. We have to provide effective packet corruption for Falcon transmitters to enable selective packet interfering. We model the superposition of base signals and Falcon interferences and then present how to maximize the base signal deformation by adopting an adaptive frequency adjusting strategy for Falcon transmitters.

We implement Falcon using commodity off-the-shelf (COTS) LoRa devices and conduct extensive experiments to evaluate its performance. Our results show that Falcon significantly pushes the coverage of LoRa under all evaluating situations. In summary, our contributions are as follows:

- To the best of our knowledge, Falcon is the first end-to-end protocol that runs on top of COTS LoRa devices, including both the commercial LoRa clients and gateways, for providing reliable communication links to connect unreachable LoRa clients.
- To address practical challenges, we propose a CAD-based channel detection approach to detect the existence of base signals with low power consumption and high accuracy. We accurately measure and remove the impact of time and frequency offsets on packets corruptions by a reception-based synchronization algorithm and an adaptive frequency adjusting strategy.
- We implemented Falcon using COTS LoRa devices and evaluated its performance extensively. The evaluation results show that Falcon can receive from low-SNR LoRa clients with signals below the standard LoRa SNR boundary upto 7.5 dB.

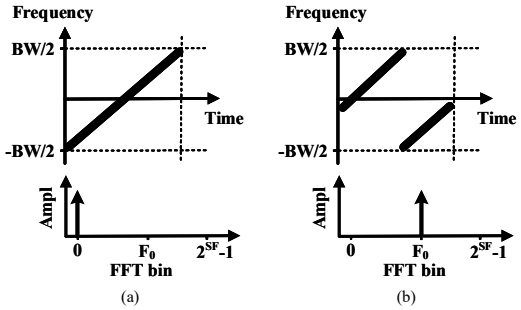

Figure 2: LoRa physical layer modulation mechanism with chirp spread spectrum.

2 BACKGROUND

PHY of LoRa. LoRa employs the Chirp Spreading Spectrum (CSS) modulation at the physical layer. It modulates the signal into up-chirps/down-chirps of linear increasing/decreasing frequencies, occupying the entire channel band. As shown in Fig 2(a), given the channel bandwidth BW, the frequency of the base up-chirp $C(t)$ increases linearly from $-\frac{BW}{2}$ to $\frac{BW}{2}$. LoRa encodes data bits into symbols by shifting initial frequencies of base up-chirps. For encoding SF bits with a chirp, LoRa defines 2^{SF} different shifted initial frequencies. After frequency shifting, all frequencies higher than $\frac{BW}{2}$ will align down to $-\frac{BW}{2}$, as shown in Fig 2(b). The length of an encoded chirp is determined by the SF and BW, which is $T_{chirp} = 2^{SF}/BW$. As the SF increases or the BW decreases, the length of each chirp extends, leading to lower data rates. The LoRa receiver demodulates encoded up-chirps in a two-step process. First, it dechirps symbols by multiplying them with a standard down-chirp. The multiplied down-chirp is the conjugate of the base up-chirp; therefore, the multiplication results in a single tone with frequency equaling to the initial frequency of the encoded up-chirp. Then, the receiver performs the FFT on the dechirped signal, leading to an energy peak in the associated FFT bin. The encoded data can be recovered by determining which FFT bin the energy peak lies in.

MAC of LoRa. LoRaWAN is a widely adopted MAC for LoRa networks. It offers three types of connections, i.e., Class A, B, C, to satisfy different IoT applications. The difference between those three connection types is how and when end devices receive downlink messages. Class A devices receive downlink messages following each uplink transmission. Class B devices receive after uplink transmissions and on each time-synchronized beacon. Class C devices open receiving windows continuously.

3 LORA CORRUPTION BASICS

SNR Boundary for LoRa Demodulation. This section explores the theoretical SNR boundary for the standard LoRa demodulation.

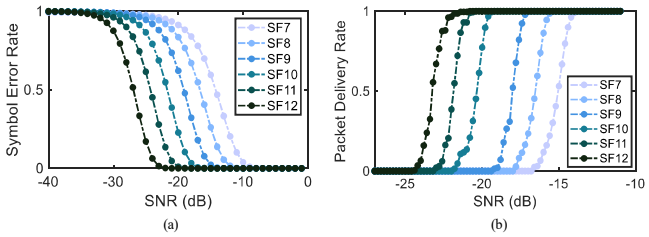

Figure 3: The SER performance for standard LoRa demodulation across SNRs and SFs: (a) SERs from theoretical model. (b) PDR estimated by commercial LoRa gateways.

SF	SNR Boundary	SF	SNR Boundary
7	-11.0dB	**10**	-19.0dB
8	-13.6dB	**11**	-21.7dB
9	-16.3dB	**12**	-24.4dB

Table 1: SNR boundaries for the standard LoRa across SFs.

We consider a LoRa chirp $C(t)$ transmitted over an AWGN channel. The received samples can be expressed as

$$r[n] = \alpha_s C[n] + w[n], \qquad 0 \le n < N$$

where α_s is the amplitude of the received signal, $w[n]$ depicts the zero-mean complex Gaussian noise with a variance of $2\sigma^2$, and N is the number of samples of a LoRa chirp. With perfect synchronizing, the dechirping and FFT of the received signal lead to an energy concentration in a single FFT bin, and all the remaining $N-1$ bins have zero energy of $C(t)$. After FFT, the noise is transformed to the frequency domain and still follows the Gaussian distribution but with a variance of $2\sigma^2 N$ according to Parseval's theorem [13]. Therefore, the value of each FFT bin is distributed as

$$Y_k \sim \begin{cases} \mathcal{CN}(0, 2\sigma^2 N), & k \in \mathcal{S}\backslash s, \\ \mathcal{CN}(\alpha_s N, 2\sigma^2 N), & k = s. \end{cases} \tag{1}$$

where $\mathcal{S} = \{0, 1, \cdots, 2^{SF} - 1\}$ is a collection of all FFT bins each corresponding to a potential encoded symbol, and s is the transmitted symbol. A symbol error occurs if any of the $|Y_k|$ for $k \in \mathcal{S}\backslash s$ exceeds the value of $|Y_s|$, i.e., $\max_{k\in\mathcal{S}\backslash s} |Y_k| > |Y_s|$. According to Eq. 1, Y_s is normally distributed, and thus its absolute value follows the Rice distribution, i.e., $|Y_s| \sim \mathcal{R}(\alpha_s N, \sigma\sqrt{N})$. Therefore, the probability density function (PDF) for the energy in bin s is

$$f_{|Y_s|}(x) = \frac{x}{\sigma^2 N} e^{\frac{-(x^2+(\alpha_s N)^2)}{2\sigma^2 N}} I_0\left(\frac{x\alpha_s}{\sigma^2}\right) \tag{2}$$

where $I_0(\cdot)$ is the modified Bessel function of the first kind with order zero, and the energy of the remaining FFT bins follows the Rayleigh distribution. The cumulative distribution function (CDF) of the maximum energy of the remaining bins, i.e., the CDF of $|Y_{max}| = \max_{k\in\mathcal{S}/s} |Y_k|$, is

$$F_{|Y_{max}|}(x) = \left(1 - e^{-x^2/(2\sigma^2 N)}\right)^{N-1} \tag{3}$$

Therefore, given a specific encoded chirp s, the probability of symbol error is

$$P(\hat{s} \neq s|s) = \int_{x=0}^{+\infty} \left(\left(1 - F_{|Y_{max}|}(x)\right) \times f_{|Y_s|}(x) \right) dx \tag{4}$$

Given any encoded symbol s, the probability of symbol error is identical. Thus, Eq. 4 calculates the expected symbol error rate (SER) for

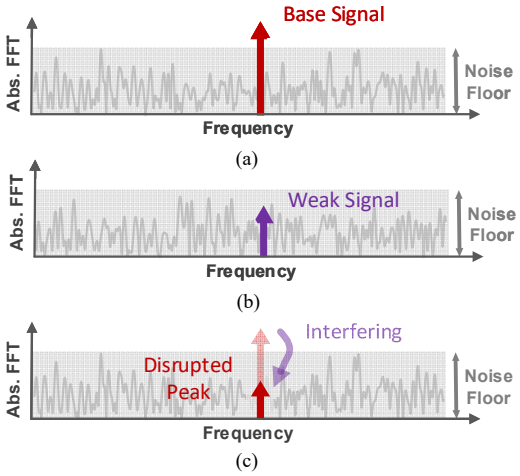

Figure 4: An illustration of signal interfering: FFT peaks of weak LoRa signals can cause interference on base signals.

LoRa demodulation given the signal energy and noise levels. The theoretical SER curves across different received SNRs and chirp SFs are plotted in Fig 3(a). Symbol errors for all SFs undergo a stepped rise as the SNR decreases. The SNR boundaries for all SF configurations when the SER curve exceeds 10% are shown in Table. 1. We verify the SNR boundaries with commercial LoRa devices. We use an SX1278 LoRa client as the transmitter and estimate packet delivery rates (PDR) on an SX1301 gateway. The two devices are wired connected with an adjustable attenuator. Thus, we can evaluate the PDR under different SNR conditions by varying the link attenuation. The evaluation results are plotted in Fig 3(b). As the SNR decreases, the PDR first remains constant and then undergoes a quick drop when SNR falls below the demodulation boundaries. This paper strives to make LoRa signals with SNR below the demodulation boundaries capable of delivering data bits to the gateway.

Interfering LoRa Transmissions. This section describes how LoRa signals below the SNR boundaries can interfere with other LoRa transmissions on the same channel. LoRa receivers demodulate received chirps by identifying the highest energy peak in the frequency domain. A successful transmission requires the energy peak of received chirps to be higher than any other noise peak, as shown in Fig 4(a). Otherwise, the receiver incorrectly picks up the highest noise peak as the decoding result, and this leads to symbol errors, as in Fig 4(b). However, signals drowned in the noise still have an impact on the channel. During the dechirping and FFT, the energy of each weak LoRa chirp is efficiently concentrated, although this energy peak cannot be detected in conventional LoRa demodulation. When a weak LoRa chirp and a base chirp transmit over the same channel simultaneously, the energy peak of the weak chirp can introduce significant cancellation on the peak of the base chirp, making the latter unresolvable, as shown in Fig. 4(c).

We numerically analyze how weak LoRa signals cause interference on other LoRa transmissions, namely base signals, for the rest of this section. Suppose the weak signal is received with an SNR of α dB, below the SNR boundary for standard LoRa demodulation. The base signal is received at β dB, above the boundary and can be demodulated. The power of these two signals can be expressed as

$$P_\alpha = P_w 10^{\frac{\alpha}{10}}, \quad P_\beta = P_w 10^{\frac{\beta}{10}}$$

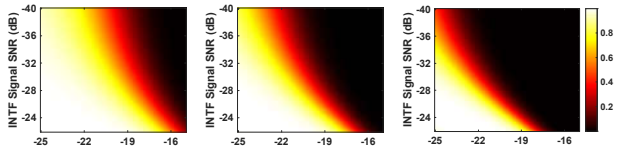

(a) SF10 (b) SF11 (c) SF12

Figure 5: Demonstration of Falcon's effectiveness on signal interfering: symbol error rates (SERs) versus base signal SNRs and interference SNRs across SFs.

where P_w is the power density of the channel noise. Therefore, when the energy peaks of these two signals superposed destructively, the SNR of the resulted signal can be written as

$$SNR_{sup} = 10\lg\left(\frac{\left(\sqrt{P_\alpha}-\sqrt{P_\beta}\right)^2}{P_w}\right) = 10\lg\left(10^{\frac{\alpha}{10}}+10^{\frac{\beta}{10}}-2\times 10^{\frac{\alpha+\beta}{20}}\right) \quad (5)$$

We translate the SNRs of the superposed signal to the expected SERs based on the analytical model in Eq. 4 and plot the SER results under different SNR conditions in Fig 5. We can see that (1) weak signals significantly increase the SERs of base signals by the destructive superposition, and (2) base signals with lower SNRs are more vulnerable to weak signal interference as their peaks are close to the noise output. As the SNR of the base signal approaches the lowest boundary for normal LoRa demodulation, the SNR range for effective interference expands, i.e., Falcon signals with low SNRs can cause more base signal errors. We also compare the Falcon's interfering capabilities for different LoRa configurations (i.e., SFs) in Fig 5. We define the effective interfering area in Fig 5 as the region where base signals are above the normal demodulation boundary and Falcon signals can cause interference. There exist effective interfering areas for all three figures, showing interference is effective under all SF configurations. As the SF increases, the effective interfering area also expands.

Summary. (1) A weak signal, which cannot be decoded due to low SNR, can create non-negligible interference to normal LoRa transmissions. (2) This inspires us to decode a weak LoRa packet by using it to interfere with other LoRa transmissions actively.

4 SYSTEM OVERVIEW

Falcon is a software solution that aims to provide reliable links for low-SNR LoRa devices that cannot reach gateways with the standard LoRa protocol. Falcon achieves this by exploiting the interfering capability of weak LoRa signals to other LoRa transmissions at the same channel. We observe that the SNR requirement for packet interfering is much lower than that of the standard LoRa demodulation. Thus, Falcon enables low-SNR LoRa devices to convey data bits to LoRa gateways via selectively interfering with LoRa packets from base signal transmissions.

At the high level, the overall system design of Falcon is as follows (see Fig. 6): A base signal transmitter, either a LoRa client or a gateway, delivers LoRa packets to a remote LoRa gateway with adaptive transmitting powers meeting the lowest signal demodulation requirement. The base signal transmitter can be specified at the system initialization or selected during system operating. When

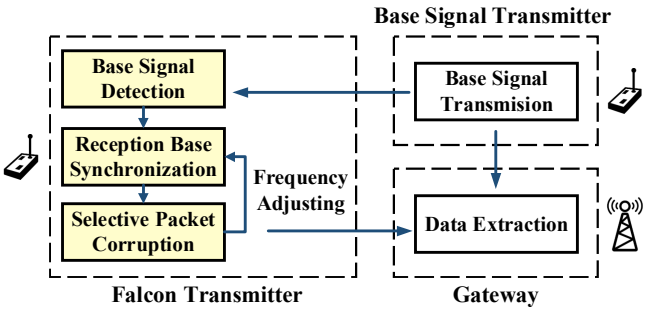

Figure 6: Architecture of Falcon.

initially specified, the base signal transmitter is manually registered to the LoRa server and broadcasts to all clients. Otherwise, the LoRa server selects the optimum base signal transmitter from LoRa clients or gateways that are idle and have sufficient power. A Falcon transmitter, which cannot reach the gateway, detects the base signal and switches to the active interfering mode. The Falcon transmitter first synchronizes its time and frequency with the base signal. Then it either does nothing or corrupts the base signal. If the Falcon transmitter does nothing, the packets of the base signal are successfully decoded at the gateway. Otherwise, the base signal should have a low SNR and is even undecodable. In this way, the Falcon transmitter can successfully convey its data to remote LoRa gateways by selectively corrupting the base signal. The receiver can decode the conveyed data by interpreting the interfering results.

Falcon aims to provide emergency links to LoRa clients to combat the link dynamics. When a client has data to transmit, it will first try the regular LoRa transmission. Only when the regular LoRa fails, it switches to the Falcon mode and waits for the base signal transmitter to boot up the Falcon transmissions. The gateway recognizes a LoRa client as a Falcon transmitter when the client is disconnected from the gateway for more than a time threshold. When multiple Falcon transmitters try to use this system, their transmissions are scheduled by the gateway to avoid collisions. Before each round of base signal transmission, the gateway specifies a Falcon transmitter to transmit at the following slot in a broadcast message. Due to the asymmetry link between clients and gateways, downlink messages from the gateway can be delivered with a higher transmit power. Thus, the control message from the gateway can be overheard by Falcon transmitters, avoiding multiple transmitters interfering with the same base signal simultaneously.

The rest of the paper addresses the key challenges in designing the three main aspects of the Falcon's architecture:

(1) Power-efficient Base Signal Detection: Falcon takes the existing LoRa transmissions as the base signal for conveying data bits. Thus, it has to detect the base signal transmissions accurately. This incurs two challenges in practice. First, battery-powered LoRa clients have limited power budgets. They cannot afford continuous packet detection. Second, base signals of LoRa can be under the noise floor of the receiver. Therefore, conventional RSS-based channel detection cannot be used. Falcon exploits the LoRa Channel Activity Detection (CAD) for detecting base signal transmissions with low power consumption. Sec. 5 describes our approach.

(2) Time and Frequency Synchronization: Falcon should efficiently synchronize the time and frequency between the Falcon

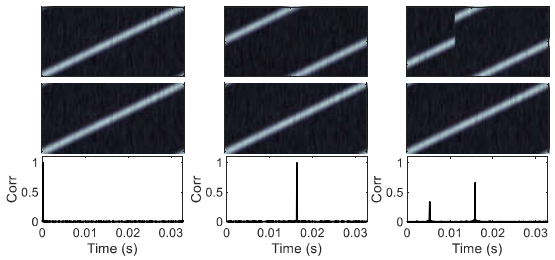

Figure 7: CAD cross-correlation on various LoRa chirps: (a) preamble chirps aligned with base chirp, (b) preamble chirps with time offset, (c) payload chirps with time offset.

transmitter and the base signal transmitter. For effective LoRa interfering, each pair of Falcon interfering chirp and base chirp should be demodulated into the same FFT bin, making the two peaks interfere with each other. This requires a high synchronizations precision. For example, for LoRa chirps of SF12 and BW = 125kHz, the synchronization precision should be at least 8μs in time or 30 Hz in frequency. The key challenge comes from the channel asymmetry, where signals from Falcon transmitters cannot reach gateways or base signal transmitters. Further, LoRa transmissions, unlike Wi-Fi or cellular, are narrowband, making the synchronization even more challenging. Sec. 6 addresses these challenges.

(3) Interfere Interference-Resilient LoRa Packet: LoRa packets are designed to be interference resilient. Thus, we need to design the interfering method carefully. We need to resolve the time/frequency drifts of low-cost LoRa devices and uncontrollable oscillator phase shifts. We adopt an adaptive frequency adjusting approach to address these challenges. Sec. 7 describes our solution.

5 EFFICIENT PACKET DETECTION

This section describes how to detect LoRa packets for Falcon transmitters, which allows them to switch from low-power sleeping mode to active interfering mode when there is a base signal packet.

5.1 Base Signal Detection

Falcon explores using LoRa CAD for detecting base LoRa signals. LoRa CAD is a recently released mechanism for detecting LoRa chirps. It searches the LoRa channels following a two-step process: (1) First, the receiver sets the RX radio to the desired configuration, including frequency, BW, SF, and captures I/Q samples from the channel for a period of time. (2) Then it switches to the processing stage and searches for a strong cross-correlation between the received samples and a locally generated base up-chirp. An incoming LoRa packet is detected if the cross-correlation exceeds a predetermined threshold. LoRa CAD concentrates the energy of received LoRa signals by performing correlation, and thus it detects LoRa chirps even under extremely low SNRs.

LoRa CAD is claimed to be able to detect chirps from both the preambles and payloads [14]. In practice, however, we observe a significant difference when using LoRa CAD to search for preamble chirps and payload chirps. We conduct experiments on the cross-correlation to better understand the behavior of CAD. Specifically, we compute the cross-correlation between a generated base up-chirp and three different LoRa chirps recorded by the USRP. When

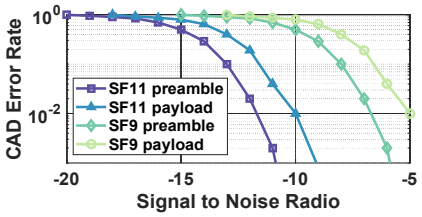

Figure 8: Channel Activity Detection (CAD) Error Rates.

the incoming signal is a preamble chirp strictly aligned with the generated chirp, as shown in Fig. 7(a), the cross-correlation results in a single peak locating at zero which signifies a detection. Then, if the incoming signal is a preamble with a time offset to the generated chirp, as shown in Fig. 7(b), the cross-correlation produces a peak at different time instances but of the same significance. Finally, if the incoming signal is a payload, as shown in Fig. 7(c), the generated chirp overlaps with two consecutive payload chirps. This leads to multiple peaks in the cross-correlation, each with lower significance compared to peaks of preamble chirps. The experiment results indicate that the cross-correlation for payload chirps has a worse detection ability compared to the preamble chirps. The reason is that the consecutive payload chirps are usually with different initial frequencies, and thus their energy cannot be concentrated during the cross-correlation. We also evaluate the effectiveness of LoRa CAD for detecting practical LoRa transmissions under different levels of SNRs and SFs. We count the number of successful detections to calculate the CAD error rates and show the results in Fig. 8. The SNR requirement for detecting payload chirps is obviously higher than that of detecting preambles, which coincides with the behavior of our cross-correlation analysis.

To enable Falcon transmitter to detect the presence of both LoRa preambles and payloads, we design a specialized packet structure whose payloads consist of the continuous base up chirps. By thoroughly exploring the LoRa encoding procedure, we find LoRa payload chirps, except for the header and CRC bits, can be set to any value by carefully designing the data contents. This finding inspires us to design packets with continuous base chirps as payloads; therefore, they can be robustly detected by CAD even during the payload transmission. Based on this, Falcon senses the base signal from LoRa channel by periodically turning on the LoRa CAD, detects base chirps from the LoRa preambles and payloads, and switches to active interfering mode whenever the CAD cross-correlation exceeds a predetermined threshold. The CAD-based carrier sensing avoids continuous RX radio opening and detects packets without full-fledged demodulation. It achieves high detection accuracy with little energy consumption.

5.2 Cooperate with LoRa MAC

Falcon carrier sensing can cooperate with the existing LoRa MAC protocol, i.e., LoRaWAN, for improving the efficiency of base signal detection. Devices in LoRaWAN Class C have continuous receiving windows, and thus they can perform Falcon carrier sensing by receiving downlink base signals at any time. Devices in LoRaWAN Class B open periodical receiving windows based on time-synchronized beacons from gateways. In Class B, the base signal transmitters send base packets tightly following the time-synchronization beacons. Therefore, Falcon transmitters can detect

the channel occupancy periodically through a time scheduling scheme while keeping in low-power sleeping mode between every two time-synchronized beacons. We emphasize that Class B can best cooperate with power-limited Falcon clients to optimize the energy efficiency for carrier sensing. It also ensures that Falcon transmitters can be reliably waked up even when base signals are received with low SNRs. Devices in LoRaWAN Class A open receiving windows only after each uplink transmission. Thus, they cannot detect channel activities in real time. For a node in Class A, it can first transits to Class B or C for efficient carrier sensing.

6 SYNCHRONIZATION

For selectively signal interfering, Falcon transmitters have to be synchronized with the base signal transmitter both in time and frequency. We present a reception-based approach on commodity devices to synchronize both the time and frequency of LoRa transmissions. At high level, our synchronization algorithm is simple. To initialize the synchronization, the base signal transmitter first acquires the medium and transmits base signals. Upon detecting the reference packets, the Falcon transmitter switches to active mode and adjusts its local time and frequency. Finally, the Falcon transmitter turns on the TX radio and joins the base signal transmission with a high synchronization accuracy.

Several challenges should be carefully resolved for synchronization in Falcon. First, there are hardware processing delays for Falcon clients to synchronize and replay on the arrival of base signals. This processing delay, taking tens of milliseconds, can lead to time offsets of several chirps. Second, signals from the Falcon transmitter and base signal transmitter arrive at the receiver through different paths with different propagation delays. For typical LoRa links with communication ranges of $0.1 \sim 10km$, the propagation delay is $0.3 \sim 30\mu s$. Finally, the local oscillators of the Falcon transmitter and base signal transmitter are not synchronized and introduce different carrier frequency offsets. According to the SX1278 datasheet [15], the maximum tolerated frequency offset between transmitter and receiver can be ±25% of the bandwidth. In the following, we focus on measuring and compensating for these delays and offsets.

(a) Hardware Processing Delay. The hardware processing delay is the time required for a LoRa device to decode the received packet, arouse a hardware interrupt, and switch from the radio RX to TX. Our measurements show that the time for packet decoding is decided by the coding rate and data length of the received packet. The interrupt delay and radio switching time are nearly constant for a specific hardware platform. Therefore, we compensate the hardware processing delays for each received packet based on the pre-estimations. Any remaining time delay can be treated as a special frequency offset as $C(t-\tau) = e^{-j2\pi k\tau}C(t)$. We will illustrate how to measure and remove it in the CFO processing part.

(b) Propagation Delay. The propagation delay is the time of flight for the signal from the transmitter to the receiver. Falcon leverages the location information of clients and gateways for estimating the propagation delay. For devices with fixed locations, the coordinates of the devices can be obtained during the initialization, which is usually required in real systems. For mobile devices, we propose to estimate the geographic coordinate of each Falcon transmitter by hearing from multiple distributed LoRa gateways using a TDoA based algorithm. The TDoA based localization approach can estimate the locations of the nodes very roughly (with a positioning accuracy of $600m$ corresponding to a delay estimation error within $2\mu s$) by receiving beacons from distributed gateways in very low computation complexity. Existing systems [16, 17] show that LoRa TDoA algorithms can achieve a positioning accuracy of $300 \sim 600$ meters in open areas, which translates to $1 \sim 2\mu s$ error in propagation delays satisfying our accuracy requirements.

(c) Carrier Frequency Offset. The carrier frequency offset (CFO) is induced by the frequency mismatch between the TX and RX oscillators. It introduces frequency biases on the received LoRa chirps and impacts the demodulation. Fortunately, commodity LoRa chips, such as SX1278, can measure such frequency biases. The key innovation is to leverage the special structure of LoRa packets, i.e., base up chirps in preambles and base down chirps in SFDs. The estimated frequency bias, LoRa Frequency Error Indicator (FEI), is stored as a 20-bit value accessible from registers 0x28 to 0x2A for SX1278. The FEI can be accurately measured as long as the frequency bias is in the region of ±25% of the bandwidth. Note that the hardware processing delays also introduce frequency offsets to the incoming LoRa chirps. The FEI value is a combination of the CFO and the remaining time delays. Base on this, Falcon transmitters cancel the frequency biases by extracting FEIs from local registers and adjusting their local frequencies accordingly.

Summary. Falcon estimates all sorts of time delays and frequency offsets on commodity devices. These delays and offsets are compensated by adjusting the transmitting time and radio frequency of Falcon transmitters. Our experiment results in Sec. 9.1 show that the proposed approach achieves time and frequency synchronization with the accuracy of $\pm 10\mu s$ and $\pm 25Hz$, respectively.

7 INTERFERING THE BASE LORA SIGNAL

Falcon transmitters deliver data bits to remote LoRa gateways by selectively corrupting the base signal packets. Ideal packet corruption expects the FFT peaks of base signals to be canceled by the interfering signal. This requires the interfering peaks to be strictly aligned in frequency and have opposite phases. Two facts, however, challenge the packet corruption: (1) There are random time and frequency drifts between base signals and interfering signals, which is $\pm 10\mu s$ in time and $\pm 25Hz$ in frequency. These drifts come from the hardware fluctuations and environment variations, which cannot be thoroughly eliminated by the reception-based synchronization. (2) Low cost LoRa radio chips do not support delicate phase control. In this section, we first model the superposition of the base signal and interfering chirps, showing how time/frequency drifts and unpredictable phases impact the packet corruption performance. Then, we propose an adaptive frequency adjusting approach that can maximize the base signal deformation and thus enable effective packet corruption.

7.1 Signal Superposition Model

We first model the signal superposition considering the unpredictable synchronization drifts and phase differences. During the Falcon transmission, the base signal transmitter and the Falcon node send packets with the same contents. Those two transmitters have a small time/frequency drift and a random phase difference. Both time and frequency drifts lead to initial frequency shifts of

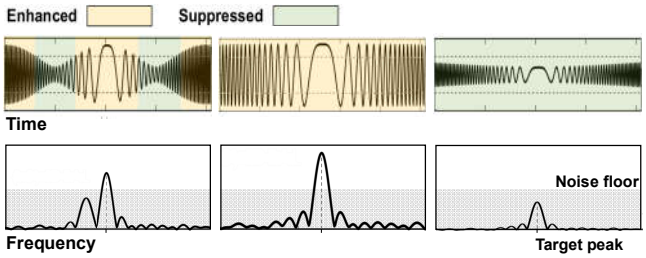

Figure 9: Effects of chirp superposition: (a) Ineffective interference with a significant frequency offset, (b) Ineffective interference when the two chirps are in-phase, (c) Effective interference when the two chirps have opposite-phase.

received LoRa chirps. Assume two LoRa chirps of different initial frequency shifts, i.e., f_1 and f_2, are superposed with a random phase difference, the superposition can be written as

$$\begin{aligned} y(t) &= A_1 \sin\left(2\pi(f_1 + F_t)t + \varphi_1\right) + A_2 \sin\left(2\pi(f_2 + F_t)t + \varphi_2\right) \\ &= A \sin\left(2\pi(\frac{f_1 + f_2}{2} + F_t)t + \frac{\varphi_1 + \varphi_2}{2} + \phi\right) \end{aligned} \tag{6}$$

where

$$A = \sqrt{A_1^2 + A_2^2 + 2A_1A_2\cos[2\pi(f_1 - f_2)t + (\varphi_1 - \varphi_2)]}$$

$$\phi = \arctan\left(\frac{(A_1 - A_2)}{(A_1 + A_2)} \tan\left(\pi(f_1 - f_2)t + \frac{\varphi_1 - \varphi_2}{2}\right)\right)$$

A_i and φ_i are the amplitude and the initial phase of the superposed chirps. F_t is the instantaneous frequency of encoded chirps, which increases linearly in time. Eq. 6 shows that the superposition for two chirps of different frequency shifts results in a chirp with a periodic varying amplitude. The varying trend is determined by both of the frequency difference (i.e., $f_1 - f_2$) and the phase difference (i.e., $\varphi_1 - \varphi_2$). Given the length of a chirp (i.e., $T = 2^{SF}/BW$), when the frequency difference is larger than the reciprocal of the chirp length, i.e., $f_1 - f_2 > 1/T$, the superposition will cause periodical change (enhancement and suppression) for the base signal, as shown in Fig. 9(a). In such a case, two superposed chirps are transformed to peaks at different FFT bins, which can be identified and disentangled in the frequency domain, leading to packet corruption failures. Even when the frequency difference is within a very small range, i.e., peaks of both base signals and interfering signals falling in the same FFT bin, an improper initial phase of the interfering signal can also lead to ineffective interference, as shown in Fig. 9(b). Only when the two chirps are well aligned and have opposite phases, the base chirp can be effectively suppressed, as shown in Fig. 9(c).

To summarize, for effective packet interfering, the interfering signal should (1) have a frequency difference no larger than $1/T$ with base chirps, and (2) have an opposite phase to the base signal. For the following, we introduce how to adjust the frequency and phase of interfering signals to achieve effective interference.

7.2 Adaptive Frequency Adjusting

We propose to leverage the frequency hopping capability of commodity LoRa devices for adjusting the frequencies and phases of LoRa chirps. LoRa supports frequency hopping for enabling long-duration LoRa transmissions without violating the maximum allowable channel dwell time. The principle behind the frequency hopping scheme is that a portion of a LoRa packet can be transmitted over a hopping channel. We achieve the frequency hopping on commercial LoRa devices based on a hardware interrupt, named *ChangeChannelFhss*, which enables LoRa transmitters to select and switch to a new radio frequency during the transmission. We trigger the *ChangeChannelFhss* interrupt periodically on the payload transmission. The trigger interval should be an integer multiple of the chirp length, and the minimum interval is a single chirp.

Based on the frequency hopping, Falcon transmitters linearly vary the center frequency of chirps for the interfering packets. Therefore, for a received base signal packet with arbitrary initial frequency shift, there should be interfering chirps strictly aligned with the superposed base chirps in the frequency. For example, when the initial frequency offset between the base signal packet and the interfering packet is randomly distributed in the range of ±50 Hz, and the payload has 50 chirps modulated in $SF12$ and 125 kHz bandwidth, we linearly change the frequency of each interfering chirp with a step of 2 Hz. This promises more than 60% of base chirps superposing interference chirps with a frequency offsets smaller than $1/T$, which supports effective packet corruption. The frequency hopping also introduces a symbol-varying initial phase for each interference chirp. Denote the initial phase for the first interference chirp as φ_0. Considering the linearly varying frequency of interference chirps, the initial phase of the ith chirp is shifted by

$$\varphi_i = \varphi_0 + \sum_{k=1}^{i-1} (\Delta f \cdot kT) \tag{7}$$

where Δf is the step of frequency change, and T is the chirp length. Eq. 7 shows varying initial phases for interference chirps. Thus, given received chirps with arbitrary frequencies and initial phases, Falcon can generate interfering chirps with aligned frequencies and opposite phases, enabling effective packet corruption.

8 IMPLEMENTATION

We implement Falcon transmitters and base signal transmitters on Semtech SX1278 LoRa clients [15], and deploy an indoor testbed and a campus-scale testbed for performance evaluation. We implement Falcon clients to encode data bits using a LoRa-like forward correction code (i.e., Hamming code). As the goal of Falcon is to provide reliable emergency links, the coding rate is selected to optimize the error correction capability. For our evaluation, Falcon clients use the 4/8 Hamming code where every 4 data bits are encoded into 8 bits for error correction in transmission. For component-level evaluation, we implement a LoRa gateway on a USRP N210 software radio with a UBX daughterboard [18]. The USRP is connected to a 64-bit PC on Ubuntu 20.04 via Ethernet cables, which collects and demodulates physical samples of arrived LoRa packets. We use the collected signal samples for analyzing synchronization errors and evaluating fine-grained packet corruption performance for Falcon. It should be noted that Falcon can be totally implemented in COTS LoRa devices instead of USRP. The implementation on USRP is mainly for micro-benchmarks evaluation. For system-level evaluation, we use a commodity Raspberry-Pi based LoRa gateway with Semtech SX1301 chips [19]. Unless specified otherwise, LoRa packets in our evaluations are generated on $SF12$ with BW of $125kHz$ and transmitted at $480MHz$.

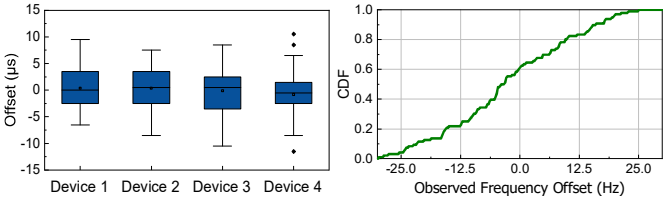

Figure 10: Accuracy for Falcon's reception based synchronization: (a) Time synchronization errors, (b) CDF of frequency synchronization errors.

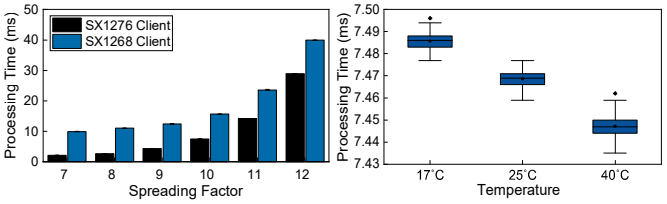

Figure 11: Time cost for Falcon reception based synchronization. Processing time for (a) various SFs and hardware platforms and (b) different environmental conditions.

9 EVALUATION

9.1 Synchronization

In this experiment, we evaluate the accuracy and time costs for synchronizing the time and frequency between Falcon clients and base signal transmitters.

Method: We deploy a testbed of five LoRa clients, placed at different locations in the campus, including buildings, roads, and parking lots, with different distances varying from several hundred meters to several kilometers to the USRP gateway. Each Falcon transmitter is informed of the coordinates of itself and the gateway after being deployed. In the implementation of this paper, we use the *RxDone* interrupt provided by LoRa SX1278 clients as a time reference for synchronization. The *RxDone* interrupt is raised after the entire LoRa packet is received. Therefore, in our implementation, we require the base signal transmitter to broadcast two consecutive reference packets for each Falcon bit transmission, the first one for detection and synchronization and the second one for interfering. In practice, a client can detect and synchronize with the ongoing transmission earlier by using other interrupts. For example, we have verified that by using the *OnPreambleDetected* interrupt for base signal detection and synchronization with the SX1268 client, the detection and synchronization can be finished at the end of the preamble. The client can interfere with the payload in the rest of the base signal packet. In such a scenario, only one packet is required for each Falcon bit. In this experiment, we use four LoRa clients as Falcon transmitters and one as the base signal transmitter. Each Falcon transmitter overhears the base signal packet, synchronizes upon the reception, and joins the transmission with a predefined frequency shift (i.e., $\Delta f = i \times 200$ kHz for the ith Falcon transmitter). The USRP gateway collects samples of each joint transmission, separates overlapped packets from their different center frequencies, and determines the synchronization accuracy for each Falcon transmitter. We record signal processing time at each Falcon transmitter under various SF configurations and with different hardware platforms for evaluating the time costs of Falcon synchronizations. We also evaluate the synchronizing time at different temperatures for showing the impacts of environmental dynamics.

Results: Observing the time plot in Fig. 10(a), we can see that the time for all four Falcon transmitters is tightly synchronized with the maximum synchronization error of $\pm 10\mu s$. The average synchronization time error for all devices is $0.4\mu s$ with a standard deviation of $3.8\mu s$. The errors in time synchronization are mainly introduced by the multi-path effects and hardware uncertainty. The remaining synchronization errors can be further resolved by the adaptive frequency adjusting mechanism in Sec. 7. It is worth noting that we do not see a significant increase in synchronization errors for clients placed distant from the receiver. This is because the Falcon transmitters can pre-estimate and compensate for propagation delays according to the positioning information. There are outliers for device 4 in Fig. 10(a) with unexpected high synchronization errors (i.e., $8 \sim 10\mu s$). We investigate the collected data and reveal that those abnormal errors are caused by intermittent interferences from neighbors of device 4. Fig. 10(b) shows the CDF of frequency errors for all Falcon transmitters. Almost all packets are received within a frequency error of ±25 Hz. These errors in frequency synchronization are mainly induced by the poor quality of low-cost Falcon crystal oscillators, whose clock drifts over time and fluctuates as the environment changes. We can use the adaptive frequency adjusting to compensate for these errors and achieve effective packet interfering for Falcon transmitters.

Fig. 11(a) shows the processing time for reception-based synchronization on Falcon clients, including packet detection, frequency and time offset correction, and RX to TX switching. The synchronization is finished within several chirps, and the processing time increases as the SF of base signals rises. LoRa chirps in high SFs have longer durations compared to those of lower SF, and thus they require more computational overhead for demodulation. Fig. 11(a) also shows a significant difference in processing time between different hardware platforms. This indicates that the processing time should be pre-estimated and compensated for each single hardware platform. Fig. 11(b) shows the processing time of the SX1278 client at $SF10$ in three environments of different temperatures. The processing time slightly varies for different temperatures as the microcontroller and clock skews are susceptible to environmental dynamics. Such impacts of environmental dynamics, however, can be resolved by extending the hopping range of Falcon's adaptive frequency adjusting. Besides, the data from temperature and humidity sensors equipped on LoRa nodes can be used for compensating the impact of environmental dynamics.

9.2 Packet Corruption

In this experiment, we evaluate the effectiveness of Falcon on packet corruption under different SNR of base signal and interference.

Method: We use two LoRa clients placed close to a USRP gateway to transmit base signal and Falcon interference signal concurrently. Those two signals are transmitted over adjacent channels with different frequencies. At the receiver, we down convert the two signals to the same base band and manually mix them up. We generate packet superpositions with different SNR settings by adjusting

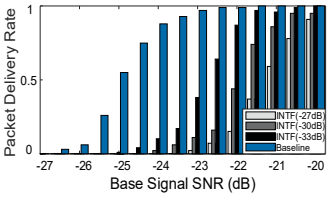

Figure 12: Packet delivery rates for Falcon's packet corruption.

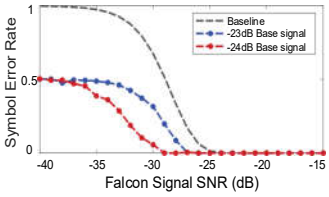

Figure 13: SER performance for the indoor platform compared with standard LoRa.

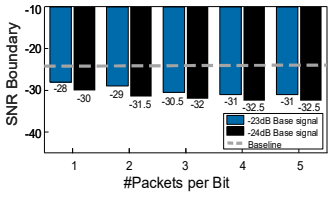

Figure 14: SNR boundaries for the indoor platform compared with standard LoRa.

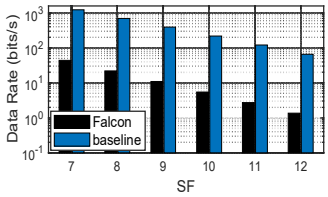

Figure 15: Data rate of networks with Falcon compared with standard LoRa.

the magnitude of each superposed signal as well as introducing additional white Gaussian noises. Then we send the superposed signals to a standard LoRa demodulator on MATLAB and show the effectiveness of Falcon packet corruption by revealing the packet delivery rates. For each SNR setting, we repeat the transmission for 1000 times. The packet delivery rates of standard LoRa without interference are set as the baseline.

Results: Fig. 12 shows the PDR of base signals at various SNRs superposed by different interferences. We focus on the SNR boundary when the packet error rate for the base signal exceeds 20%. When there is no Falcon interference, this SNR boundary is as low as −25 dB, above which more than 80% base packets are successfully decoded. When superposed with the interference of −27 dB, the SNR boundary of base signal increases by 4 dB, showing the low-SNR Falcon interference signal can cause effective packet corruptions on base signal transmissions. Even if the Falcon interference signal is as weak as −33 dB, it also introduces non-negligible interference and increases the SNR boundary by 2 dB. Falcon achieves effective packet corruption as it leverages LoRa chirps as the interference whose energy is concentrated during LoRa demodulation and thus causes significant deformation on base signal peaks. Besides, Falcon's packet corruption design which adaptively adjusts frequencies and phases also maximize the interference.

9.3 Performance for Indoor Platform

In this experiment, we evaluate the performance of Falcon for connecting low-SNR Falcon transmitters for indoor environments.

Method: We use an indoor LoRa client as a base signal transmitter to continuously transmit packets to a LoRa gateway. We program the base signal transmitter to adjust its transmitting power adaptively. Therefore, base signals received at the gateway are with a constant SNR during each Falcon transmission. In every experiment, we let the Falcon transmitter transmit a specific known sequence of symbols (i.e., a series of '0' and '1') to the gateway by selectively interfering with the base signal transmissions. We vary the transmitting power of the Falcon transmitter and evaluate its symbol error rate under different Falcon signal SNRs. We also vary the SNR of base signals for showing the impacts of base signal SNRs on Falcon transmissions. We reveal the SNR boundary of Falcon with two base signals of −23 dB and −24 dB, respectively. And we show the effectiveness of retransmissions on the SNR boundary performance. We estimate the overall data rate of the Falcon links and compare it with standard LoRa. Finally, we experiment the impacts of LoRa Forward Error Correction (FEC) on Falcon transmissions by evaluating the true positive rate and false positive rate for Falcon packet corruptions across different coding rates of base signals.

Results: Fig. 13 shows the SER of Falcon transmissions under different SNRs of Falcon signals and base signals. We set the SER of standard LoRa demodulation for SF12 as the baseline. We define symbol errors for Falcon transmissions as packet corruption failures (true negatives) or base packet loss with no interference (false positives). False positives are rare as most base packets are successfully delivered when they are transmitted beyond the SNR boundary (i.e., −23 dB and −24 dB for this experiment). For the Falcon signal received with a high SNR, it can efficiently corrupt the concurrently transmitted base signal, leading to a high success rate of Falcon transmissions. As the SNR of the Falcon signal decreases, the interference of the Falcon signal fails to corrupt base signals, causing an increase in Falcon transmission errors. Falcon transmissions assisted by weak base signals (i.e., −24 dB) perform much better than those assisted by strong base signals (i.e., −23 dB). This is because base signals with low SNRs tend to be more vulnerable for interference. Fig. 14 shows the SNR boundary for Falcon using multiple retransmissions representing each Falcon bit. The SNR boundary is the lowest SNR when the error rate of the transmitted messages is below 20%. We use the SNR boundary for standard LoRa at SF12 as the baseline. Cooperation with retransmissions significantly enhances the performance of Falcon transmission, as the retransmissions can improve the tolerance for true negatives and false positives for link dynamics. When using five base packets representing one Falcon bit, the SNR boundary for Falcon can exceed −32.5 dB, about 7.5 dB improvement compared with the baseline approach. Fig. 15 shows the data rate of a Falcon link at different SF configurations. Falcon provides reliable emergency links to disconnected LoRa clients, which can transmit tens to hundreds of data bits every minute. We compare the data rate of Falcon with that of standard LoRa at $SF12$. Note that the SNR requirement for standard LoRa is much higher than that of Falcon. When the SNR of received signals falls below the lowest threshold of standard LoRa, the data rate of baseline decreases to zero. For the scenario of emergency communications that Falcon aims to provide, the data rate shown in Fig. 15 will be sufficient.

We finally exploit the impact of LoRa FEC on Falcon transmissions. LoRa has four different coding rates for FEC from CR4/5 to CR4/8, with 1 to 4 additional bits of redundancy added to every four data bits. Fig. 16(a) shows the true positive rate for Falcon under four different coding rates. The SNR of the base signal is −24 dB and the SNR of the Falcon signal is −27 dB. It is interesting to see that the true positive rate for Falcon decreases as the coding rate of base packets increases. It is because that the LoRa receiver corrects some of the base signal errors caused by Falcon interference with a high FEC coding rate. We resolve this problem by symbol retransmission

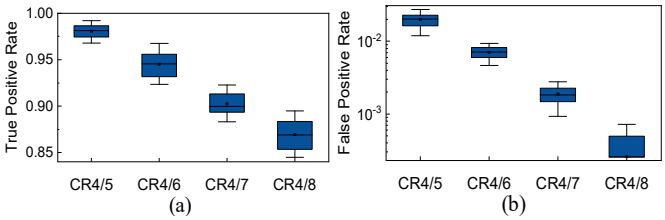

Figure 16: Falcon transmissions under different configurations of LoRa forward error correction (FEC) schemes: (a) True positive rates and (b) False positive rates.

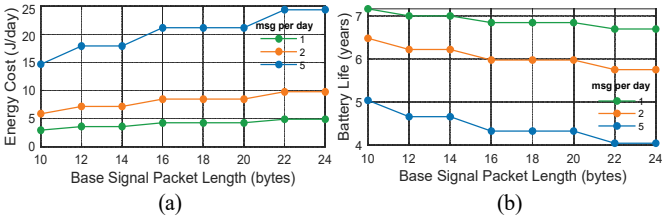

Figure 17: Energy efficiency for Falcon base signal transmitters: (a) energy cost for transmitting base signal packets and (b) battery life for base signal transmitters.

Status	Current
MCU + Radio sleep	$9.4mA$
MCU + TX	$43 \sim 129mA$
MCU + RX	$21.3mA$
MCU sleep + Radio sleep	$331nA$

Table 2: Current for a typical LoRa client device with different operations. The device is powered at $3.6V$.

or shutting down the FEC in practice. The false positive rate of Falcon also decreases with the CR increasing (see Fig. 16(b)), as base signal errors from channel noise are corrected by FEC.

9.4 Energy Cost

We evaluate the energy cost of base signal transmitters on assisting the Falcon transmission.

Method: We study the energy profile of a typical LoRa client with the Monsoon HV Power Monitor [20]. We estimate the current of a LoRa client at different statuses, such as MCU computing, transmitting, receiving, and sleeping. Based on the estimation, we calculate the energy cost for base signal transmitters by multiplying the currents with input voltages (i.e., $3.6V$ for our evaluation) and accumulating the multiplication results over time. We use short messages for Falcon emergency links, where each message consists of 10 data bits for reporting Falcon client's status. We evaluate the energy costs for transmitting base signals at various packet lengths and duty cycles. We further simulate the battery life of the base signal transmitters given the battery capacity of 2200 mAh.

Results: The current profile of the LoRa client is shown in Table 2. We can see that radio operations, including transmitting and receiving, consume the highest amount of the energy. The LoRa client can adjust their TX power for transmissions, and thus the estimated current varies from 43 to 129 mA on different TX settings. Fig 17(a) shows the energy cost of base signal transmitters for assisting Falcon transmissions. The energy cost is related to the base signal packet length and the duty cycle. A larger packet length and a higher duty cycle can lead to a more reliable link but at the expense of higher energy cost. We further translate the energy cost to battery life. When there is no Falcon transmission, the lifetime of the LoRa client can be 8 years. Fig 17(b) gives the battery life of the base signal transmitters, which decreases when there are more Falcon transmission requirements and longer base signal packets. In real-world deployments, the base signal transmitter can be any of the LoRa devices (e.g., battery-powered clients, wall-plugged clients, and gateways) as long as the device can emit base signal packets in the form of chirps. Thus, we also suggest that a wall-plugged gateway can perform as a base signal transmitter to avoid the power consumption issue.

9.5 Coverage

We use trace-driven simulations to show the advantages of Falcon in extending the coverage area.

Method: Given a specific receiver (i.e., gateway), we estimate its coverage by calculating the receiving power from any possible transmitter in the area of campus. We estimate the signal power with a land-cover map based path-loss model described in SateLoc [21]. Further, we calibrate the model using more than 10000 points collected in the campus using GPS-equipped LoRa nodes. Specifically, we first generate land-cover maps according to the satellite images, where the generated maps are composed of pixels of buildings, forests, waters, and fields. Then, we transform the land-cover map to signal strength map by corresponding pixels to degradation values. We finally produce the path loss model based on the signal strength map. The SNR boundaries at the receiver are taken from our evaluations with indoor deployments. We compare the SNR and coverage performance of Falcon with Choir [12].

Results: Fig 18(a) shows the coverage of normal LoRa. Due to the blockage of buildings and trees, link attenuation to different directions varies, showing directional isomerism. Fig 18(a) also indicates that not all LoRa devices can experience the promised communication range. Some devices are disconnected even if they are deployed only hundreds of meters away from the gateway. This is also validated in our real campus measurement. Fig 18(b) shows the coverage of Falcon . Falcon provides emergency links to disconnected clients in normal LoRa. Thus, the coverage area extends significantly compared with that of normal LoRa. Fig 18(c) shows the coverage of Choir with two collaborated transmitters (grey area) and five collaborated transmitters (red area). As the number of collaborated transmitters increases, the coverage range of Choir extends. Note here we assume different Choir nodes transmit identical data in order to combine their signals to extend coverage. This limits its usage in real-world systems.

9.6 Deployment of A Campus-Scale Testbed

We deploy a campus-scale LoRa testbed to show the performance of Falcon.

Method: As illustrated in Fig. 20, the LoRa testbed consists of 10 LoRa clients, each deployed at a different distance to the gateway. To examine the performance of Falcon, most deployment positions are with severe blockages, and thus clients at those positions cannot reach the gateway for most of the time. We make the LoRa clients first transmit with standard LoRa modulation and then with Falcon. Then, we evaluate the SER performance for both transmissions.

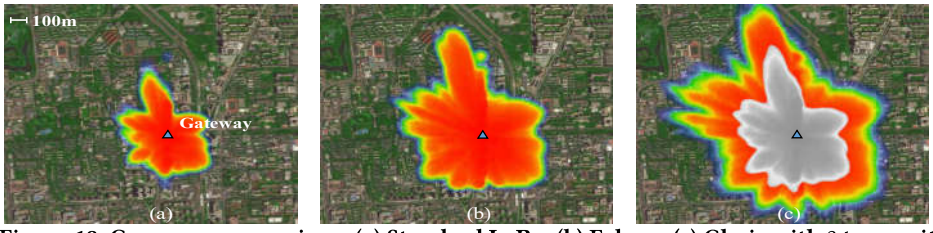

Figure 18: Coverage comparison: (a) Standard LoRa, (b) Falcon, (c) Choir with 2 transmitters (grey area) and 5 transmitters (red area) for the ideal case with identical data.

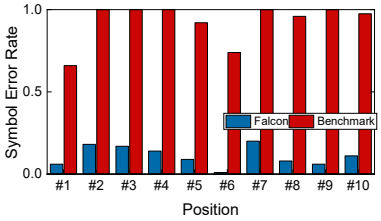

Figure 19: Performance for campus-scale testbed.

Figure 20: Campus-scale LoRa testbed deployment.

Results: The SERs for LoRa clients at different deployment positions are shown in Fig. 19. For all of the ten positions, Falcon achieves a much lower SER for data transmission than the standard LoRa. The SER for standard LoRa at all positions is higher than 60%, which has no chance of being corrected even with the highest coding rate of LoRa FEC (i.e., CR4/8). The SER for Falcon transmissions is much lower, and most errors can be resolved by leveraging more retransmissions for each Falcon bit. This proves that Falcon can provide reliable communication links to LoRa clients when the standard LoRa communication cannot reach.

10 DISCUSSION

In this section, we briefly discuss some issues in design of Falcon.

• **System Deployment.** Falcon can be deployed in operational LoRa networks using existing devices with no hardware modification. Falcon can select base signal transmitters from LoRa devices that have adequate power budgets and are currently in idle state. For most LoRa applications, devices work at a low duty-cycle. Therefore, it will be easy for Falcon to find appropriate devices as base signal transmitters without interfering the original application. Besides, dedicate base signal transmitters deployed by network operators will further improve the Falcon transmissions. All clients connected to the same gateway can share the same base signal transmitter, and thus the deployment costs will be low. When Falcon clients are transmitting, regular LoRa transmissions at the same channel may be interfered. Such negative channel impacts, however, can be resolved through retransmissions, as Falcon only takes a small portion of channel occupancy for providing emergency links.

• **Noise Resistance.** Falcon achieves noise resistance by taking advantage of the LoRa modulation. Base signals and low-SNR Falcon signals are modulated in chirps. During the demodulation, the chirp energy can be efficiently concentrated, leading to effective interference resistance. Unexpected noise, which is not chirps, cannot concentrate its energy in demodulation and has a limited impact on base signal decoding. The redundant coding mechanism also correct code errors caused by unexpected noises. On the one hand, by carefully tuning the transmitting power of the base signal transmitter, base signals can be resistant to channel noise. On the other hand, by carefully designing the interfering signal, base signals can be successfully interfered with by Falcon transmissions.

• **Environment Dynamics.** Falcon communication is impacted by the environmental dynamics as the low power microcontrollers and clock skews are susceptible to environmental conditions. Falcon resolves the environmental dynamics by treating their impacts as random frequency drifts and compensates those drifts by adaptive frequency adjusting as illustrated in Sec. 7.

• **Security.** Falcon can address the malicious interference by (1) keeping base signal packets being secret between Falcon transmitters and base signal transmitters, and (2) monitoring the SNR of base signals at the gateway to detect malicious interference of random sequences or noise. Meanwhile, by encoding the transmitted data, the bits impacted by malicious interference can also be identified.

• **Coverage vs. Throughput and Energy Consumption.** Falcon makes a tradeoff on network throughput and energy consumption to enable a longer range connectivity. For each Falcon transmission, it requires the base transmitter to emit base signals, which increases the energy consumption for base signal transmitters as well as compromises the network throughput. To minimize the impact on original LoRa transmissions, Falcon selects LoRa devices as base signal transmitters that have adequate power budget and are currently idle. This also balances the lifetime and accessability for all devices in the network, which improves the system reliability and reduces the maintenance costs. Normally, the number of nodes disconnected from the gateway should be very low. Falcon is supposed to support those nodes and should introduce a small overhead to the entire network.

• **Falcon vs. Multi-hop/Mesh Protocols.** Multi-hop/mesh protocols can extend network coverage by signal forwarding with relays. Applying those protocols in LoRa, however, cannot achieve the same goal as Falcon. For a multi-hop network, the end-devices need to use the multi-hop network protocol all the time, whereas a node resorts to Falcon only when its SNR is too low to reach the gateway. Note that the number of clients disconnected from the gateway should be very low compared with that of normal LoRa clients. Most of the clients should work in the normal LoRa mode without any change. Multi-hop/mesh protocols introduce extra overhead especially for LPWANs as they have to change the behavior of all clients (e.g., building routing information) in the network. Further, in Falcon , all clients in the same network can share the same base signal transmitter, which significantly reduces its

deployment costs compared to relay-based approaches. Finally, the multi-hop/mesh network requires a dense deployment of clients. A client that cannot reach any neighboring relay will be inaccessible in the mesh/relaying architecture. Such kind of clients can still successfully transmit in Falcon, as long as the downlink message from gateways or base signal transmitters can be overheard by the client. Such asymmetric links are easy to achieve as gateways or base signal transmitters can use a higher TX power.

• **Falcon vs. Backscatter.** Falcon is conceptually akin to wireless backscatter in some aspects. It is, however, different from backscatter from the following aspects: (1) Falcon aims to provide emergency links to disconnected low-SNR LoRa clients, while backscatter aims to enable ubiquitous connections to battery-free tags with as little energy consumption as possible. (2) Falcon achieves a coverage extension compared with regular LoRa, while backscatter still has the communication range limited by the original signal SNR (e.g., up to 1.1km for PLoRa [22]). Falcon modulates information by interfering with the signal, and the receiver does not need to decode the signal; while typical backscatter systems modulate information in the amplitude/phase of the backscatter signal, and the receiver needs to decode the signal. (3) Falcon is designed and implemented with commodity LoRa clients, while most backscatter systems use specially designed backscatter tags and specialized receivers that cannot be implemented on commodity LoRa devices.

11 RELATED WORKS

LoRa transmission optimization. There have been much research effort on LoRa transmission optimization, including ways to mitigate collisions [12, 23–29], improve parameter allocations [11, 30–32], build up enhanced coding mechanisms [33–39], and design a variety of channel sensing and estimation solutions [14, 21, 40–43]. Extensive works focus on mitigating LoRa collisions, which provide multi-user Gaussian channels in LoRa by disentangling concurrent transmissions. They resolve multiple challenges of high concurrency [12, 23], real time processing [24–26], bad channel conditions [27–29] and so on. Other works improve LoRa transmissions by selecting optimum parameters [11, 14, 30–32] or adapting specially designed PHY coding mechanisms [34, 36–39].

Communication range extension. Extending communication range is a traditional problem in wireless. Some works aim to extend communication ranges by using MIMO on wireless devices [44–48]. They synchronize signal phases from multiple antennas, enabling constructive superposition of the transmitted signal and thus improving the signal power at the receiver. ONPC [49] extends the range of standard WiFi by piggybacking data bits on noises and interference power levels. However, these solutions have high demands on hardware overheads or rely on specific signal patterns, and thus are not appropriate for LoRa devices.

Closely related to our system are recent solutions on extending the communication range of low power LoRa nodes [8–12, 40]. Choir [12] exploits coordinating transmissions from a team of LoRa sensors which are individually too weak to reach the base station. The key idea is that sensing data from spatially adjacent sensors are closely related and thus the data packets for those sensing data tend to superpose constructively. Charm [8], OPR [10] and Nephalai [9]exploit spatial diversity gains by gathering signals from 3-8 spatial distributed gateways. Chime [11] identifies an optimal frequency of LoRa operation by analyzing wireless channels based on a single heartbeat packet reception at 3-6 gateways. While these systems improve range to commercial LoRa devices, our system is different from those systems as it provides reliable communication links to low SNR LoRa clients without collaboration of multiple gateways or assistance from cloud servers.

Backscatter. Backscatter provides near-zero-power communications by passively reflecting radio signals in the environment [22, 50–55]. Recent backscatter systems focus on resolving different practical challenges, such as communication range [22, 55], tag concurrency [54], and specialized hardware [50–53]. This work is inspired by WiTAG [53] that backscatter to Wi-Fi APs by selectively interfering with the on-transmitting Wi-Fi frames. However, the goals of the two systems are different: Falcon provides emergency links to LoRa clients with extremely low SNR, while WiTAG is to provide low-power short-range connections to backscatter tags.

Multi-hop/mesh protocols. Multi-hop/mesh protocols have been extensively studied in literatures [56–60]. Most protocols are designed for improving reliability, saving energy, and avoiding collisions in sensor networks. These protocols, however, are not well suitable for LPWANs where end nodes are sparsely deployed and communicate at an extremely low data rate. Besides, mesh protocols introduce non-negligible extra overhead for LPWANs as they have to maintain route tables for each node in the network and exchange messages between neighbors frequently.

12 CONCLUSION

This paper proposes Falcon, a physical layer mechanism to provide reliable communication links for disconnected LoRa devices. Falcon utilizes the finding that LoRa signals too weak to be decoded can still introduce significant interference on other LoRa transmissions. Thus, Falcon achieves communication by selectively interfering with other concurrent LoRa transmissions. We propose several novel techniques to address practical challenges in Falcon design. We propose a low-power channel activity detection method to efficiently detect concurrent LoRa transmissions for selectively interfering. We corrupt so-called interference-resilient LoRa packets by presenting a reception-based synchronization algorithm and an adaptive frequency adjusting strategy to maximize the interference. We implement Falcon totally using commercial off-the-shelf LoRa devices and thoroughly evaluate its performance in both an indoor testbed and an outdoor campus-scale testbed. The evaluation results show that Falcon can provide reliable emergency links for disconnected LoRa devices and achieves the SNR boundary upto 7.5dB lower compared to the standard LoRa.

ACKNOWLEDGMENT

We thank the anonymous shepherd and reviewers for their insightful comments to improve the quality of our work. We also thank Zhenqiang Xu from Tsinghua University for joining the discussion of this paper. This work is in part supported by NSFC No. 61932013 and Tsinghua-Foshan Innovation Special Fund (TFISF). Jiliang Wang is the corresponding author.

REFERENCES

[1] Yiwei Ma and Jianliang Chen. Toward intelligent agriculture service platform with lora-based wireless sensor network. In *Proceedings of IEEE ICASI*, Tokyo, Japan, April 13-17, 2018.

[2] Jithu G. Panicker, Mohamed Azman, and R. Kashyap. A lora wireless mesh network for wide-area animal tracking. In *Proceedings of IEEE ICECCT*, Tamil Nadu, India, February 20-22, 2019.

[3] Lili Chen, Jie Xiong, Xiaojiang Chen, Sunghoon Ivan Lee, Kai Chen, Dianhe Han, Dingyi Fang, Zhanyong Tang, and Zheng Wang. Widesee: Towards wide-area contactless wireless sensing. In *Proceedings of ACM SenSys*, New York, NY, USA, November 10-13, 2019.

[4] R. Jedermann, M. Borysov, N. Hartgenbusch, S. Jaeger, M. Sellwig, and W. Lang. Testing lora for food applications - example application for airflow measurements inside cooled warehouses with apples. *Procedia Manufacturing*, 24(5):284–289, February 2018.

[5] J. P. Shanmuga Sundaram, W. Du, and Z. Zhao. A survey on lora networking: Research problems, current solutions, and open issues. *IEEE Communications Surveys & Tutorials*, 22(1):371–388, October 2019.

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

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

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

[9] Jun Liu, Weitao Xu, Sanjay Jha, and Wen Hu. Nephalai: Towards lpwan c-ran with physical layer compression. In *Proceedings of ACM MobiCom*, Online, September 21-25, 2020.

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

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

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

[13] Tallal Elshabrawy and Joerg Robert. Closed-form approximation of lora modulation ber performance. *IEEE Communications Letters*, 22(9):1778–1781, June 2018.

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

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

[16] Bernat Carbonés Fargas and Martin Nordal Petersen. Gps-free geolocation using lora in low-power wans. In *Proceedings of IEEE GIoTS*, 2017.

[17] Nico Podevijn, David Plets, Jens Trogh, Luc Martens, Pieter Suanet, Kim Hendrikse, and Wout Joseph. Tdoa-based outdoor positioning with tracking algorithm in a public lora network. *Wireless Communications and Mobile Computing*, 2018.

[18] USRP Ettus. N210 datasheet. *Available: https://www.ettus.com/*.

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

[20] Monsoon Solutions Inc. High voltage power monitor. *Available: https://www.msoon.com/high-voltage-power-monitor*.

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

[22] 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 *Proceedings of ACM SIGCOMM*, Budapest, Hungary, August 20-25, 2018.

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

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

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

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

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

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

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

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

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

[32] Liu Li, Yao Yuguang, Cao Zhichao, and Zhang Mi. Deeplora: Learning accurate path loss model for long distance links in lpwan. In *Proceedings of IEEE INFOCOM*, Online, May 10-13, 2021.

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

[34] Akshay Gadre, Fan Yi, Anthony Rowe, Bob Iannucci, and Swarun Kumar. Quick (and dirty) aggregate queries on low-power wans. In *Proceedings of ACM/IEEE IPSN*, Online, April 21-24, 2020.

[35] Xianjin Xia, Yuanqing Zheng, and Tao Gu. Litenap: Downclocking lora reception. In *Proceedings of IEEE INFOCOM*, Online, July 6-9, 2020.

[36] Ruofeng Liu, Zhimeng Yin, Wenchao Jiang, and Tian He. Xfi: Cross-technology iot data collection via commodity wifi. In *Proceedings of IEEE ICNP*, Online, October 13-16, 2020.

[37] Junyang Shi, Di Mu, and Mo Sha. Lorabee: Cross-technology communication from lora to zigbee via payload encoding. In *Proceedings of IEEE ICNP*, Chicago, Illinois, USA, October 7-10, 2019.

[38] Ningning Hou and Yuanqing Zheng. Cloaklora : A covert channel over lora phy. In *Proceedings of IEEE ICNP*, Online, October 13-16, 2020.

[39] Zhijun Li and Yongrui Chen. Achieving universal low-power wide-area networks on existing wireless devices. In *Proceedings of IEEE ICNP*, Chicago, Illinois, USA, October 7-10, 2019.

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

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

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

[43] Chaojie Gu, Linshan Jiang, Rui Tan, Mo Li, and Jun Huang. Attack-aware synchronization-free data timestamping in lorawan. *ACM Transactions on Sensor Networks*, In press, 2021.

[44] Ezzeldin Hamed, Hariharan Rahul, Mohammed A. Abdelghany, and Dina Katabi. Real-time distributed mimo systems. In *SIGCOMM*, Florianopolis, Brazil, August 22-26, 2016.

[45] Ezzeldin Hamed, Hariharan Rahul, and Bahar Partov. Chorus: truly distributed distributed-mimo. In *SIGCOMM*, Budapest, Hungary, August 20-25, 2018.

[46] Se-Yeon Jeon, Min-Ho Ka, Seungha Shin, Munsung Kim, Seok Nyeon Kim, Sunwook Kim, Jeongbae Kim, Aulia Dewantari, Jaeheung Kim, and Hansup Chung. W-band mimo fmcw radar system with simultaneous transmission of orthogonal waveforms for high-resolution imaging. *IEEE Transactions on Microwave Theory and Techniques*, 66(11):5051–5064, September 2018.

[47] Arijit Roy, Debasish Deb, Harshal B. Nemade, and Ratnajit Bhattacharjee. Design of discrete frequency-coding waveforms using phase-coded linear chirp for multiuser and mimo radar systems. Karnataka, India, February 20-23, 2019.

[48] Diep N Nguyen and Marwan Krunz. A cooperative mimo framework for wireless sensor networks. *ACM Transactions on Sensor Networks*, 10(3):1–28, 2014.

[49] Philip Lundrigan, Neal Patwari, and Sneha K Kasera. On-off noise power communication. In *Proceedings of ACM Mobicom*, Los Cabos, Mexico, October 21-25, 2019.

[50] Vincent Liu, Aaron Parks, Vamsi Talla, Shyamnath Gollakota, David Wetherall, and Joshua R. Smith. Ambient backscatter: Wireless communication out of thin air. In *Proceedings of ACM SIGCOMM*, Hong Kong, China, 2013.

[51] Anran Wang, Vikram Iyer, Vamsi Talla, Joshua R. Smith, and Shyamnath Gollakota. FM backscatter: Enabling connected cities and smart fabrics. In *Proceedings of USENIX NSDI*, Boston, MA, USA, 2017.

[52] Pengyu Zhang, Dinesh Bharadia, Kiran Joshi, and Sachin Katti. HitchHike: Practical backscatter using commodity WiFi. In *Proceedings of ACM SenSys*, Stanford, CA, USA, 2016.

[53] Ali Abedi, Farzan Dehbashi, Mohammad Hossein Mazaheri, Omid Abari, and Tim Brecht. Witag: Seamless wifi backscatter communication. In *Proceedings of*

ACM SIGCOMM, Virtual Event, USA, 2020.

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

[55] Vamsi Talla, Mehrdad Hessar, Bryce Kellogg, Ali Najafi, Joshua R Smith, and Shyamnath Gollakota. Lora backscatter: Enabling the vision of ubiquitous connectivity. In *Proceedings of ACM Ubicomp*, Hawaii, USA, September 11-15, 2017.

[56] Szymon Chachulski, Michael Jennings, Sachin Katti, and Dina Katabi. Trading structure for randomness in wireless opportunistic routing. In *Proceedings of ACM SIGCOMM*, ACM New York, NY, USA, 2007.

[57] Omprakash Gnawali, Rodrigo Fonseca, Kyle Jamieson, David Moss, and Philip Levis. Collection tree protocol. In *Proceedings of ACM SenSys*, Berkeley, California, USA, 2009.

[58] Michael Buettner, Gary V Yee, Eric Anderson, and Richard Han. X-mac: a short preamble mac protocol for duty-cycled wireless sensor networks. In *Proceedings of ACM SenSys*, pages 307–320, 2006.

[59] Gang Zhou, Tian He, Sudha Krishnamurthy, and John A. Stankovic. Models and solutions for radio irregularity in wireless sensor networks. *ACM Transactions on Sensor Networks*, 2(2):221–262, 2006.

[60] 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.