

FTrack: Parallel Decoding for LoRa Transmissions

Xianjin Xia
The Hong Kong Polytechnic
University, Hong Kong, China
csxxia@comp.polyu.edu.hk

Yuanqing Zheng
The Hong Kong Polytechnic
University, Hong Kong, China
csyqzheng@comp.polyu.edu.hk

Tao Gu
RMIT University,
Melbourne, Australia
tao.gu@rmit.edu.au

ABSTRACT

LoRa has emerged as a promising Low-Power Wide Area Network (LP-WAN) technology to connect a huge number of Internet-of-Things (IoT) devices. The dense deployment and an increasing number of IoT devices lead to intense collisions due to uncoordinated transmissions. However, the current MAC/PHY design of LoRaWAN fails to recover collisions, resulting in degraded performance as the system scales. This paper presents FTrack, a novel communication paradigm that enables demodulation of collided LoRa transmissions. FTrack resolves LoRa collisions at the physical layer and thereby supports parallel decoding for LoRa transmissions. We propose a novel technique to separate collided transmissions by jointly considering both the time domain and the frequency domain features. The proposed technique is motivated from two key observations: (1) the symbol edges of the same frame exhibit periodic patterns, while the symbol edges of different frames are usually misaligned in time; (2) the frequency of LoRa signal increases continuously in between the edges of symbol, yet exhibits sudden changes at the symbol edges. We detect the continuity of signal frequency to remove interference and further exploit the time-domain information of symbol edges to recover symbols of all collided frames. We implement FTrack on a low-cost software defined radio. Our testbed evaluations show that FTrack demodulates collided LoRa frames with low symbol error rates in diverse SNR conditions. It increases the throughput of LoRaWAN in real usage scenarios by up to 3 times.

CCS CONCEPTS

• **Networks** → **Network protocol design**; **Link-layer protocols**; • **Computer systems organization** → *Embedded systems*;

KEYWORDS

Internet of Things, LoRaWAN, collision, parallel decoding

ACM Reference Format:

Xianjin Xia, Yuanqing Zheng, and Tao Gu. 2019. FTrack: Parallel Decoding for LoRa Transmissions. In *The 17th ACM Conference on Embedded Networked Sensor Systems (SenSys '19)*, November 10–13, 2019, New York, NY, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3356250.3360024>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SenSys '19, November 10–13, 2019, New York, NY, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6950-3/19/11...\$15.00

<https://doi.org/10.1145/3356250.3360024>

1 INTRODUCTION

Recently, LoRa has emerged as a promising technology for Low-Power Wide Area Networks (LP-WANs). Among many LPWAN technologies (e.g., SigFox [47], NB-IoT [41]), LoRa technology [3] has attracted wide attention due to its low cost, long communication range, and supports from IoT industry as well as open-source and research community. LoRaWAN is an open-standard networking layer governed by the LoRa Alliance [3], which has about 400 member companies including Tencent, IBM, Cisco, Semtech, *etc.* LoRa employs a variant of Chirp Spread Spectrum (CSS) modulation to support several kilometers of wireless transmissions at very low power consumption. The PHY layer of LoRa is robust against interference, noise, multi-path and Doppler effects. Such characteristics make LoRaWAN a promising communication technology for IoT innovations such as smart city, health care, environment monitoring, warehouse management, *etc.*

When a LoRa node receives a packet, it first detects the preamble, then searches for chirp boundaries (*i.e.*, symbol edges) and locates the chirp signal of each symbol to decode. The standard demodulation scheme dechirps the incoming signals by multiplying with a down chirp, then performs FFT on the multiplication, which produces a single FFT peak indicating the symbol. The PHY technique of LoRa (*i.e.*, CSS) is inherently robust to interference and noise in the ISM band (*e.g.*, WiFi, RFID, *etc.*).

Although LoRa uses several PHY techniques for parallel transmissions of multiple LoRa nodes, a high-end LoRa gateway can only support up to 8 LoRa nodes to transmit at the same time [32]. It fails to meet the need of many IoT applications which require dense deployment of LoRa devices in practice. Constrained by the hardware capability and power supply, LoRa nodes typically adopt a simple aloha-based MAC for collision avoidance. As a result, collisions may occur among LoRa nodes with the same PHY configuration. One may configure LoRa nodes with different radio parameters (*e.g.*, channel, spreading factor, *etc.*) to mitigate collision, but it requires cooperation among different operators and service providers. A more effective method is to enable parallel decoding for LoRa transmissions without any extension to COTS LoRa nodes or any coordination among the users. Ideally, a LoRa node should be able to join on-going communications in parallel with other nodes without specific coordination.

To enable parallel decoding for LoRa transmissions, prior work (*e.g.*, Choir [15]) exploits the frequency offset introduced by the imperfection and diversity of LoRa nodes to separate collided frames. The key idea is to classify the collided frames according to the distinct frequency offset of each LoRa node (*e.g.*, carrier frequency). However, due to interference and noise, it is difficult to accurately extract the hardware-induced offset of carrier frequencies in practice. Moreover, the carrier frequencies of LoRa nodes inevitably

resemble each other as the number of LoRa nodes increases. Consequently, Choir may incorrectly classify collided frames and result in decoding errors.

In this paper, we present FTrack, a parallel decoding scheme for LoRa that resolves collision by jointly considering both time and frequency domain features. We exploit an observation that, as the airtime of a low-rate LoRa packet is much longer than the packets in other wireless technologies (*e.g.*, WiFi and RFID), LoRa packets are more likely to be misaligned and separated in the time domain. Due to the periodicity of LoRa symbols, the symbols from the same frame share the same pattern of edge misalignment in time. We leverage such time-domain information to separate collided LoRa symbols. In addition, we also leverage the fact that the chirp of a LoRa symbol starts from an initial frequency and then increases continuously along a fixed ‘track’ to sweep through the whole LoRa bandwidth. The different arrival time of colliding frames would result in timing offset and a corresponding shift in frequency between collided symbols. Thus, chirps of collided symbols are also likely to exhibit as different tracks in the frequency domain. We separate the frequency tracks of individual chirps to recover collisions of all frames transmitted in parallel.

Intuitively, the decoding process consists of the following key steps. FTrack first detects the edges of LoRa symbols in the time domain and groups the symbols according to different LoRa nodes. Then, it iteratively demodulates and decodes the symbols of each individual LoRa node. FTrack is orthogonal to existing parallel decoding techniques for LoRa transmissions. As long as the transmissions are interleaved with misaligned symbol edges, LoRa nodes should be able to transmit concurrently even with the same parameter configuration in the same channel.

Turning the idea into reality, however, entails tremendous challenges. First, it is non-trivial to detect the symbol edges of each individual transmission from collided signals and group the symbol edges according to each LoRa node. To address this problem, we leverage the fact that the preamble of a LoRa node is known in advance, which allows us to extract the symbol edges from collided signals by correlating the known preamble with the collided signals. The correlation peaks indicate the symbol edges in the preamble. The extracted symbol edges, however, are interleaved in the time domain due to collisions. As a LoRa preamble consists of repeated base chirps, the signal frequency would increase continuously across boundaries of all chirps associated with the preamble, resulting in a long track of continuous frequency. As such, we filter out any interfering chirps not belonging to the long frequency track to extract a pure collision-free preamble. We detect symbol edges from the extracted preamble and use that timing information to accurately pinpoint the symbol edges in its payload. Between the symbol edges corresponding to the same LoRa node, the frequency of a LoRa chirp should continuously increase, while the frequencies of other coexisting chirps may exhibit sudden changes. We leverage such a fact to filter out coexisting chirps and resolve collisions. We iterate the above process until all chirps are correctly associated with their corresponding nodes.

We implement FTrack using software defined radios (SDRs). To reduce the deployment cost, we employ low-cost SDRs to collect PHY samples and run the proposed parallel decoding scheme. We build a testbed of 20 LoRa nodes to evaluate FTrack with a variety of

transmission configurations. Results show that FTrack decodes collided frames with low symbol error rates in diverse SNRs. Compared with the state-of-the-art schemes, FTrack improves the throughput of LoRaWAN networks in real usage scenarios by up to 3 times.

2 BACKGROUND AND MOTIVATION

LoRa adopts Chirp Spread Spectrum (CSS) as the PHY modulation scheme. Symbols are modulated as up-chirp signals whose frequencies increase linearly with time over a predefined bandwidth. LoRa varies the initial frequency of an up-chirp to modulate different data. Such a procedure can be represented as follows.

$$S(t, f_{sym}) = e^{j2\pi(\frac{k}{2}t+f_0)t} \cdot e^{j2\pi f_{sym}t} = C(t) \cdot e^{j2\pi f_{sym}t}, \quad (1)$$

where f_{sym} denotes the initial frequency of the up-chirp (*i.e.*, encoded symbol). $C(t) = e^{j2\pi(\frac{k}{2}t+f_0)t}$ represents the raw chirp signal (termed *base chirp*); f_0 and k denote the starting frequency and frequency increasing rate of the chirp, respectively.

A LoRa receiver can demodulate an incoming chirp as follows. It first multiplies the received signal with the *conjugate of the base chirp* denoted as $C^{-1}(t)$, performs a Fast Fourier Transform (FFT) on the multiplication, and searches for power peak in the FFT bins to demodulate symbol. The procedure can be represented as follows

$$S(t, f_{sym}) \cdot C^{-1}(t) = e^{j2\pi f_{sym}t} \quad (2)$$

The FFT of $e^{j2\pi f_{sym}t}$ produces one peak in the FFT bins, that indicates the frequency component of f_{sym} .

When multiple LoRa nodes transmit concurrently, their signals add up at the receiver. Figure 1 presents the signal received with USRP N210 when two LoRa nodes transmit simultaneously. In the figure, we observe multiple chirps overlapping in time, each of which corresponds to the symbol of one transmitter. As a standard LoRa demodulator (Eq.(2)) searches for the maximum in FFT results, it cannot reliably decode the collided signals. We aim to support parallel transmissions of multiple LoRa nodes and resolve collisions.

A recent work, Choir [15], exploits the frequency offset (*i.e.*, offset of carrier frequency) introduced by the hardware imperfection and diversities of LoRa nodes to separate collided frames. However, it is difficult to extract the frequency offset of LoRa hardware in the presence of low SNR and inter-packet collisions. More importantly, as the number of LoRa nodes increases, the carrier frequencies of nodes are likely to resemble each other. Solely relying on hardware-induced frequency offsets may fail to differentiate LoRa nodes in practice.

In this paper, we leverage the fact that the concurrent transmissions of LoRa packets are likely to be misaligned in time. Figure 1 illustrates the time misalignment of two collided LoRa transmissions. A receiver can increase the sampling rate to better differentiate the misaligned chirps in the time domain. Besides, as the frequencies of LoRa chirps vary linearly with time, the time misalignment of chirps leads to the corresponding frequency offsets between the chirps. Thus, we can separate collided chirps in both time and frequency domains. We configure COTS LoRa nodes to transmit with the default parameters (*i.e.*, Spreading Factor=8, bandwidth=250kHz) and use a low-cost receive-only SDR (*e.g.*, RTL-SDR dongle) to sample at a high rate (*e.g.*, 1MS/s). The receiver chain will decimate the






Figure 1: Collision of two LoRa transmissions. Red/blue dashed lines indicate the symbol edges of two LoRa transmissions.



(a) Correlation detection on the preamble of the first transmission



(b) Correlation detection on the preamble of the second transmission

Figure 2: Edge detection on (a) collision-free and (b) colliding preambles, i.e., the preambles of the first and second transmissions in Figure 1. The correlation peaks indicate the boundaries of base chirps, i.e., symbol edges.

received LoRa packets and produce 256 PHY samples for each symbol. As the data rate of LoRa node is low, the high sampling rates of SDRs provide sufficient time resolution to separate misaligned chirps.

3 DESIGN

In this section, we present the design of parallel decoding for LoRa transmissions and describe the key design considerations.

3.1 Detecting Symbol Edges

A LoRa receiver must accurately locate the boundaries of a chirp (i.e., symbol edges) to demodulate the chirp. For example, Figure 3(a) illustrates three base chirps that start with the same initial frequency f_0 . If the chirp boundary t_0 is correctly located, the initial frequency can be correctly measured as f_0 and the base chirp can be demodulated. However, if the boundaries are located mistakenly with an offset of Δt , where $t_1 = t_0 + \Delta t$, the measured initial frequency would become $f_1 = f_0 + \Delta f$, leading to demodulation errors. The increasing rate of chirp frequency $k = \Delta f / \Delta t$ depends on the spreading factor of LoRa modulation which remains constant during the transmission of a packet.

Intuitively, we detect chirp boundaries and calibrate the frequency offset introduced by timing misalignment by leveraging




Figure 3: Illustration of (a) preamble chirps, and (b) the corresponding frequency tracks, i.e., normalize instantaneous frequency into the meta initial frequency of chirps.

our prior knowledge of LoRa preamble which is a standard method in LoRa demodulation. We calculate the correlation of received signals with a base chirp to detect symbol edges. When we apply the method on a non-collided preamble (e.g., the first preamble in Figure 1), symbol edges can be well identified by the correlation peaks as shown in Figure 2(a). However, when collision happens (e.g., the second preamble in Figure 1), the concurrent LoRa transmissions may introduce interference as shown in Figure 2(b). As such, a key problem arises: *how shall we detect symbol edges in the presence of collisions?*

To solve the problem, we extract pure preamble from collisions and then detect symbol edges with the extracted collision-free preamble. The method is inspired by the following observation. The frequency of a LoRa chirp will start from an initial value (named *meta initial frequency*) and increase linearly with time. All chirps in the preamble have the same meta initial frequency (i.e., f_0 in Figure 3(a)), yet those from the payload of LoRa frame do not. We exploit this fact to detect LoRa preamble.

Ideally, the meta initial frequency of a LoRa chirp (i.e., f_0) should be measured at the arrival time of the chirp (i.e., chirp boundary t_0). Without knowing the exact boundary of the chirp, we may obtain different initial frequency (e.g., f_1) when measuring at different locations (e.g., t_1) of the chirp, as illustrated in Figure 3(a). Indeed, f_1 corresponds to the instantaneous frequency of the chirp at the measuring point, but not the meta initial frequency (i.e., f_0). To deduce f_0 , we leverage the relationship between chirp frequency and time to *remove the frequency deviation (i.e., Δf) induced by time offset Δt* , as illustrated in Figure 3(b). Doing so, we normalize the time-varying frequency of chirps (e.g., f_1) into horizontal frequency lines, termed as *frequency tracks*, that indicate the meta initial frequency of chirps, i.e., f_0 . With such a normalization process, we can measure the meta initial frequency at any time of a chirp.

We employ a sliding window on the received signals shown in Figure 1 to extract the normalized initial frequency of chirps. As shown in Figure 4(a), the normalized initial frequency of each LoRa

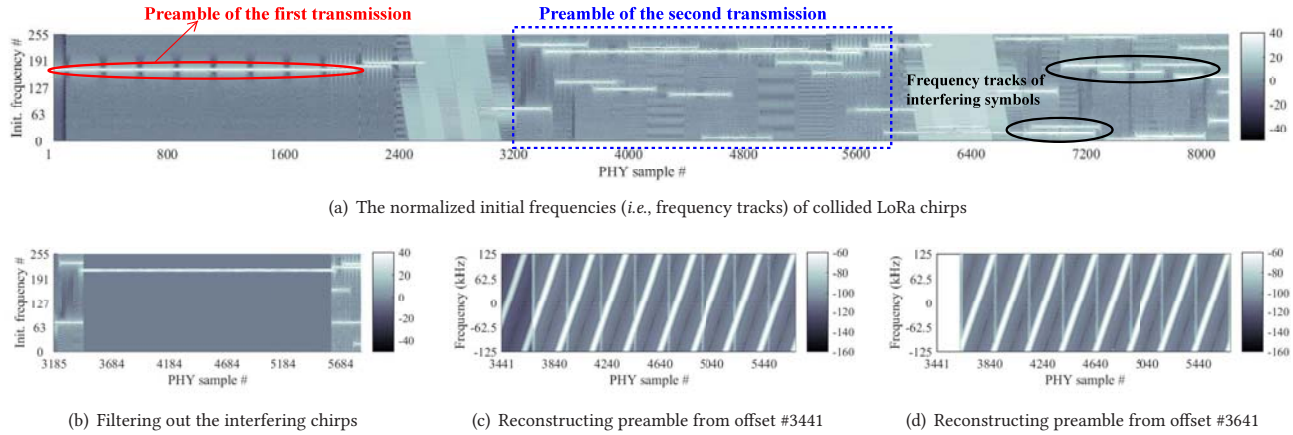


Figure 4: Extraction of collision-free preambles: (a) normalize LoRa chirps into meta initial frequency and transform colliding symbols into different frequency tracks, (b) remove interfering symbols, and (c-d) use inverse operations (i.e., the inverse of frequency normalization, followed by an inverse FFT) to reconstruct the second preamble.

chirp corresponds to a horizontal high-power line (i.e., frequency track). The long frequency tracks correspond to LoRa preamble. Since the base chirps in preamble have identical meta initial frequency, their frequency tracks form a single long track as illustrated in Figure 4(a). We exploit such a property to detect LoRa preambles.

From Figure 4(a), we also observe that the preambles of collided frames are separated into different frequency tracks. The location of frequency track is determined by both the meta initial frequency and the arrival time of a LoRa chirp (see §3.3 for detail). Therefore, although the chirps of two preambles have the same meta initial frequency, collided preambles will be separated into different tracks as long as their chirp boundaries are interleaved in time. We exploit the property to extract LoRa preamble from collisions. Specifically, we filter out interference by setting the frequency tracks of interfering chirps to zero while only keeping the long frequency track as shown in Figure 4(b). As the extracted frequency track indicates the normalized initial frequency of base chirps in preamble, we employ an inverse operation to reconstruct a pure, collision-free preamble, as shown in Figure 4(c). We can reconstruct the same preamble using a segment of the frequency track starting from different offsets, as shown in Figure 4(c) and (d). We then perform correlation detection on the reconstructed preamble to detect symbol edges. In this way, we can eliminate the impact of collision and detect the symbol edges of concurrent transmissions.

3.2 Demodulating Symbols

To demodulate symbols in the payload of a LoRa frame, we need to first locate the chirp of each symbol. We leverage the timing information of symbol edges detected from a LoRa preamble to deduce symbol edges in the payload. We refer to the signals located between two symbol edges as a *demodulation window*. Ideally, when there is no collision, there will be only one complete chirp (i.e., *target symbol*) in a demodulation window, as illustrated in Figure 5(a). We can demodulate the chirp using the standard procedure presented in §2, which will produce a single FFT peak as illustrated in




Figure 5: Symbol demodulation: (a) locate a target chirp within a demodulation window, (b) a target chirp and an interfering chirp within the same demodulation window due to collisions, (c) filter out the interfering chirp by detecting frequency continuity.




Figure 6: Demodulation with interference: the target chirp can be demodulated incorrectly due to interference.

Figure 6(a). However, when collision happens, chirps of interfering transmissions would also appear in the demodulation window, as shown in Figure 5(b). Directly applying the standard algorithm to demodulate the collided chirps will produce multiple FFT peaks, as shown in Figure 6(b), where the highest peak may not correspond to the target symbol. Therefore, how to correctly demodulate the target symbol in the presence of collisions becomes a key issue.

Our solution separates the target symbol from collisions by jointly exploiting the frequency-domain and time-domain features of concurrent LoRa transmissions. We find that the frequency of the target chirp changes continuously in the demodulation window, as there is only one chirp of the target symbol in the window. Whereas




Figure 7: Successive interfering chirps carry identical data. We filter interference by leveraging timing information.

for interfering symbols, as their symbol edges are misaligned with the demodulation window of the target symbol, the window spans across two chirps. Since the second chirp may change its initial frequency to modulate a different symbol, it will result in a sudden change of chirp frequency. Based on this observation, we examine the continuity of frequency tracks to filter out interfering symbols in each demodulation window. As illustrated in Figure 5(c), the normalized initial frequency of interfering symbols will switch to a different track. In contrast, the target symbol produces a single complete frequency track that spans over the whole window. A complete frequency track within the demodulation window indicates the initial frequency (i.e., encoded data) of the target symbol.

In case that two successive chirps of an interfering node carry the same data, the corresponding frequency tracks exhibit no sudden changes within a demodulation window, as shown in Figure 7. To address this problem, we may look into a longer window (e.g., two or more symbols). As illustrated in Figure 7, for the target symbol, the frequency track’s starting and ending points are located at the target symbol edges. In contrast, the starting and ending points of the interfering chirps are not located at the symbol edges. We can filter out the interference by checking the starting and ending points of frequency tracks against the target symbol edges.

3.3 Frequency Track in Practice

Frequency track plays important roles in collision resolving: Firstly, it separates colliding symbols into different tracks, indicating different meta initial frequencies. We exploit this property for collision recovery. Secondly, we can directly demodulate symbols from frequency tracks because data are encoded by the meta initial frequency of LoRa chirps (§2). In this subsection, we describe how to extract frequency tracks from received signals.

The signal of a LoRa chirp is represented by $S(t, f_{sym})$ in Eq. (1). Based on Eq. (1), we can formally represent the frequency track of chirp $S(t, f_{sym})$ as follows

$$F_{track}(f_{sym}) = e^{j2\pi f_{sym} t}, \quad (3)$$

where f_{sym} is the meta initial frequency of the chirp. To extract frequency track from the received signal of a LoRa chirp, one approach is to measure the instantaneous frequency of the chirp and subtract a frequency deviation Δf from the instantaneous frequency, as illustrated in Figure 3(b). In order to measure the instantaneous frequency, we can use the short-time Fourier transform (STFT) to analyze the changing spectra of signal over time (i.e., spectrogram of signal). However, such a method only provides a coarse-grained estimate of chirp frequency, which cannot precisely identify the frequency track (i.e., meta initial frequency) of LoRa chirp.

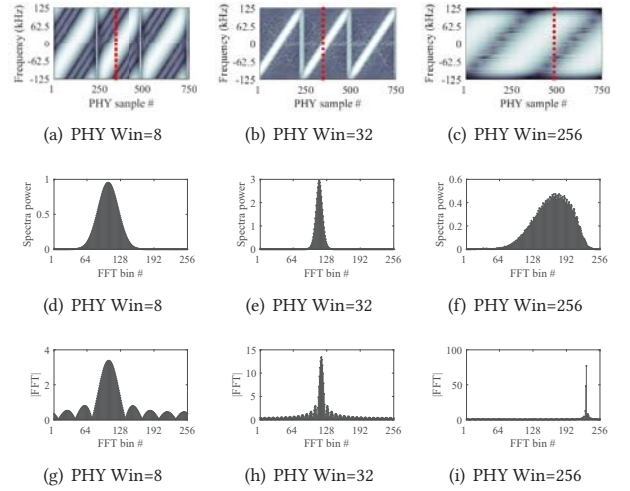


Figure 8: Frequency estimation with different PHY Wins (the FFT Win is 256). (a-c): Spectrogram of preamble chirps; (d-f): Width of the FFT peak of original signal; (g-i): Width of the FFT peak of dechirped signal.

As STFT takes PHY samples from received signal with a PHY Win to analyze the frequency components, the size of PHY Win affects the resolution of frequency estimation. Figures 8(a-c) present the resulting spectrograms of the same signal with different configurations of PHY Win. Comparing Figure 8(b) and (c), we see that the measured frequency of chirps become narrower with a smaller PHY Win of 32, because a smaller PHY Win would contain PHY samples with less time-varying frequencies. Yet, the chirps do not become thinner as PHY Win further decreases from 32 to 8 (see Figure 8(a) and (b)). The reason is that, when PHY Win is too short, STFT must pad a number of zeros to the PHY samples before FFT analysis, which will produce a wide main-lobe and many sinc side-lobes between two main chirps, as shown in Figure 8(a). As displayed in Figures 8(d-f), even with the finest configuration for spectrogram (i.e., PHY Win=32), the width of obtained chirp frequency would still span tens of FFT bins, which does not provide sufficient resolution to precisely locate the frequency track of the chirp.

We aim to extract narrow frequency tracks that can precisely indicate the meta initial frequencies of LoRa chirps. As we mentioned, due to the CSS modulation scheme of LoRa, signals in a period of Δt would sweep a frequency width of Δf (see Figure 3(a)). The time-varying frequencies of chirp within PHY Win prevent spectrogram from producing narrow frequency tracks. To solve the problem, we dechirp the signal in PHY Win by multiplying with a down chirp (i.e., $C^{-1}(t)$, see Eq. (2)). Figures 8(g-i) present the FFT of the signal samples in PHY Win after being dechirped. The resulting FFT peaks become narrower and can precisely pinpoint the frequency of corresponding chirps shown in Figure 8(d-f). Besides, the width of FFT peak becomes thinner as PHY Win increases, because the window includes more PHY samples with the same meta initial frequency. In particular, when PHY Win=256, i.e., containing the samples of a whole chirp, the resulting peak width can reach the resolution of one FFT bin, as shown in Figure 8(i).




Figure 9: Illustration of chirp signals in a *PHY Win*.

Our approach of frequency track extraction involves three key steps: Firstly, we dechirp the signal samples in *PHY Win* and perform FFT on dechirped signals to measure the instantaneous frequency of chirps at a given time offset. Secondly, we normalize the instantaneous frequency measured at different offsets by subtracting the corresponding frequency deviation. Finally, we move *PHY Win* to the next PHY sample of received signals and repeat the above two steps. We describe the approach in details in the following.

Suppose that n LoRa nodes transmit simultaneously. Let $X(t)$ denote the signal received by a LoRa device and $x_i(t)$ denote the signal from node i , i.e., $X(t) = \sum_{i=1}^n x_i(t)$. We focus on the signal of one transmitter (i.e., $x_i(t)$ of node i) in the following. In particular, we configure the length of *PHY Win* to T , i.e., the duration of a LoRa chirp. Let t_w denote the offset position of *PHY Win* in $X(t)$. A *PHY Win* generally spans two LoRa symbols as illustrated in Figure 9, where the starting edges of the two symbols are represented by t_1 and $t_1 + T$, respectively. We denote the offset between t_1 and *PHY Win* by Δt , i.e., $\Delta t = t_w - t_1$. Specifically, $0 \leq \Delta t < T$ (note that when $\Delta t = 0$, only the first symbol is included in *PHY Win*, which is a special case). Let t denote the relative time of chirp signals within *PHY Win, i.e., $0 \leq t < T$. As illustrated in Figure 9, signals in *PHY Win* are composed of two segments from the first and the second symbols, respectively. Based on Eq. (1), we can denote the first segment of signal in *PHY Win* as*

$$x_i(t) = h_i \cdot S(t + \Delta t, f_1), \quad t \in [0, T - \Delta t] \quad (4)$$

where f_1 represents the initial frequency of the first chirp. h_i is a complex value representing amplitude and phase changes on the wireless channel between the LoRa transmitter and the receiver.

According to the CSS modulation, a time shift of Δt in chirp signals would result in a frequency offset of $\Delta f = \frac{k}{2}\Delta t$ as illustrated in Figure 3(a). Therefore, Eq. (4) can be transformed as follows

$$x_i(t) = h_i \cdot S(t, f_1 + \frac{k}{2}\Delta t), \quad t \in [0, T - \Delta t] \quad (5)$$

Similarly, the second segment of signal in *PHY Win* can be represented by

$$\begin{aligned} x_i(t) &= h_i \cdot S(t - (T - \Delta t), f_2) \\ \Leftrightarrow x_i(t) &= h_i \cdot S(t, f_2 - \frac{k}{2}(T - \Delta t)), \quad t \in [T - \Delta t, T] \end{aligned} \quad (6)$$

where f_2 represents the initial frequency of the second symbol chirp. For conciseness, we focus on the first segment in the following. Similar results can be obtained for the second segment.

We dechirp $x_i(t)$ in Eq. (5) by multiplying a standard down chirp (i.e., $C^{-1}(t)$) as follows

$$x_i(t) \cdot C^{-1}(t) = h_i \cdot \underbrace{e^{j2\pi(\frac{k}{2}\Delta t)t}}_{\text{Offset of } PHY \text{ Win}} \cdot \underbrace{e^{j2\pi f_1 t}}_{\text{Initial freq.}}, \quad t \in [0, T - \Delta t] \quad (7)$$

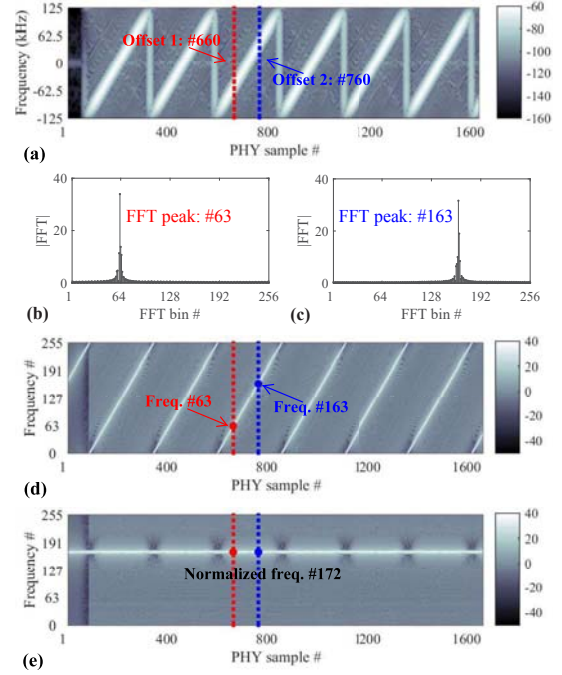


Figure 10: Steps of extracting frequency tracks: (a) spectrogram of received signals (the preamble part), (b-c) dechirp the signal in *PHY Win* (FFT of dechirped signals at #660 and #760), (d) slide *PHY Win* across all samples (the frequency of dechirped signal vs. offsets of *PHY Win*), (e) remove the frequency deviation induced by t_w (i.e., the time offset of *PHY Win*) to produce the normalized frequency of LoRa chirps, i.e., frequency tracks.

The FFT of Eq. (7) produces a peak at $f_1 + \frac{k}{2}\Delta t$, which corresponds to the frequency of the first chirp with time shift Δt , i.e., instantaneous frequency of the chirp in *PHY Win*. Figure 10(d) presents the FFT of dechirped signals when *PHY Win* slides across the PHY samples of chirps shown in Figure 10(a). It produces thin bright lines that precisely indicate the instantaneous frequency of chirps.

Lastly, we normalize the instantaneous frequency of chirps to extract frequency tracks. We need to subtract $\Delta f = \frac{k}{2}\Delta t$ from the resulting frequency of the dechirped signal in Eq. (7), i.e., $f_1 + \frac{k}{2}\Delta t$. Note that $\Delta t = t_w - t_1$, where t_1 denotes the starting edge of symbol chirp f_1 . As t_1 is not known, we cannot directly remove Δt from Eq. (7). Instead, we subtract t_w , i.e., the offset of *PHY Win* from the second term of Eq. (7) as follows

$$\begin{aligned} F_{norm}(f_1, t_1) &= h_i \cdot e^{j2\pi(-\frac{k}{2}t_1)t} \cdot e^{j2\pi f_1 t} \\ &= h_i \cdot F_{edge}(t_1) \cdot F_{track}(f_1) \end{aligned} \quad (8)$$

where $F_{norm}(f_1, t_1)$ represents the normalized frequency of the first chirp in *PHY Win*, where f_1 and t_1 are the meta initial frequency and arrival time of the chirp, respectively. $F_{track}(f_1)$ denotes the ideal frequency track of the chirp, as defined in Eq. (3). $F_{edge}(t_1)$ denotes the frequency shift induced by t_1 , i.e., the arrival time (symbol edge) of the first chirp. Since t_1 is invariant as *PHY Win* slides across the signal samples, $F_{edge}(t_1)$ is a constant.

Figure 10(e) presents the normalized frequency of the dechirped signals shown in Figure 10(d). Note that, with Eq. (8), we actually normalize the instantaneous frequency of chirp signals into a different frequency track (*i.e.*, $F_{norm}(f_1, t_1)$) that shifts from the ideal frequency track (*i.e.*, $F_{track}(f_1)$) with a specific offset $F_{edge}(t_1)$. Since $F_{edge}(t_1)$, *i.e.*, the frequency offset between $F_{track}(f_1)$ and $F_{norm}(f_1, t_1)$, remains constant across all symbols in the received signal, we can employ $F_{norm}(f_1, t_1)$ to separate collided symbols in practice. Moreover, as $F_{norm}(f_1, t_1)$ is determined by both the meta initial frequency (f_1) and arrival time (t_1) of LoRa chirps, collided symbols would be separated into different frequency tracks as long as their edges (*i.e.*, t_1) are interleaved in time.

3.4 FTrack: Put It Together

Figure 11 illustrates the general workflow of *FTrack*. Specifically, it involves the following key operations:

Frequency track extraction. *FTrack* first dechirps the received signal with a sliding window. The collided symbols in each window are normalized into the initial frequency of corresponding LoRa chirps (*i.e.*, frequency track). The collided signals are thus converted into a sequence of frequency tracks for further processing.

Frame identification. *FTrack* extracts long frequency tracks to separate preamble from collisions. It detects symbol edges from the preamble and searches for sync words to identify LoRa packets. The edges of payload symbols can then be deduced from the symbol edges of preamble based on the frame structure of LoRa packets.

Symbol demodulation. *FTrack* employs the detected symbol edges of a specific preamble to locate the demodulation window of each payload symbol associating with the preamble. *FTrack* checks the continuity of frequency tracks, as well as the timing information, in each window to filter interference and demodulate symbol from the resulting frequency track of target based on Eq. (8).

Iteration to decode parallel transmissions. *FTrack* iterates to detect more preambles and demodulate concurrent LoRa transmissions. As collided transmissions may have misaligned symbol edges, they would have different sequences of payload symbol. For example, when we apply *FTrack* to demodulate the collided LoRa signals shown in Figure 1, we can detect two LoRa preambles (exhibiting as long frequency tracks), as shown in Figure 12. Thereafter, we separate the collided symbols into two concurrent transmissions (*i.e.*, frequency track classifying) and demodulate symbols of each transmission. *FTrack* removes a frequency track after demodulating symbol corresponding to the frequency track, and outputs a sequence of demodulated symbols for each LoRa transmission, as shown in Figure 12. A conventional LoRa decoder can be applied to decode data from such symbols.

3.5 Discussion

Computation overhead. The overheads of *FTrack* mainly come from frequency track extraction and frame identification. The frequency track extraction procedure needs to slide a window across each PHY sample and perform FFT. The computational complexity is $O(N_s n_{fft} \log(n_{fft}))$, where n_{fft} is the FFT window size and N_s is the number of PHY samples. To reduce the overhead, one can set a step τ for the sliding window, *i.e.*, slide the window across τ samples each time. The resulting computation complexity would be




Figure 11: The workflow of *FTrack* system. Long track and short track represent the frequency tracks of LoRa preamble and payload symbol, respectively.

$O(\frac{N_s}{\tau} n_{fft} \log(n_{fft}))$. The procedure of frame identification needs to reconstruct preamble from collided signals using inverse FFT. The overhead is $N_p \cdot C_p n_{fft} \log(n_{fft})$, where N_p denotes the number of detected preambles and C_p denotes the number of chirps in the preamble. As the maximum of C_p is a constant, the computation complexity can be denoted as $O(N_p n_{fft} \log(n_{fft}))$. As *FTrack* detects the correlation of preamble with a base chirp to detect symbol edges, the complexity is $O(N_p n_{sym}^2)$, where n_{sym} denotes the number of PHY samples corresponding to the base chirp. Finally, the overhead of iterative symbol demodulation is $O(N_f)$, where N_f denotes the number of detected frequency tracks.

Demodulating capacity and the gain. The demodulation capacity of *FTrack* is not unbounded as the number of colliding nodes increases. Basically, *FTrack* exploits the timing misalignment of symbol edges to separate collisions. If the edges of collided symbols are aligned or closely-located, *FTrack* may not be able to separate them in time. Besides, the configurations of LoRa transmission, such as spreading factor (SF) and bandwidth (BW), also affect the demodulation performance. Nevertheless, *FTrack* can separate collisions in frequency and time domains in most cases. Assisted by the coding/decoding schemes of LoRa, it can handle symbol errors with error correction code at upper layers. In fact, current LoRaWAN supports concurrent transmission of packets using orthogonal spreading factors which require careful configuration and coordinations of LoRa nodes. However, due to lack of coordination, some nodes may still collide if they choose the same parameter settings especially in densely-populated LoRa networks (*e.g.*, an urban or warehouse scenario). As such, the proposed collision recovery approach can be complementary to current LoRa network.

Accounting for SNR variations and the near-far problem. In practice, the power strength of received signal would impact the extraction of frequency tracks. As SNR varies across diverse transmitters and channel conditions, *FTrack* may fail to detect the frequency power of weak transmitter (*i.e.*, false negative error) or mistakenly detect the power leaking from the main frequency of strong transmitter as frequency track (*i.e.*, false positive error). Both can lead to incorrect demodulation results. *FTrack* solves the problem by selecting thresholds for frequency track detection dynamically based on the signal strength of different transmitters. Moreover, to handle the case that strong receptions of a nearby transmitter overwhelm the weak signals of far-away transmitters (*i.e.*, the near-far problem [18]), we can employ a method similar to ZigZag [17] and successive interference cancellation [19] to extract signals of comparable power strength. We apply *FTrack* on each

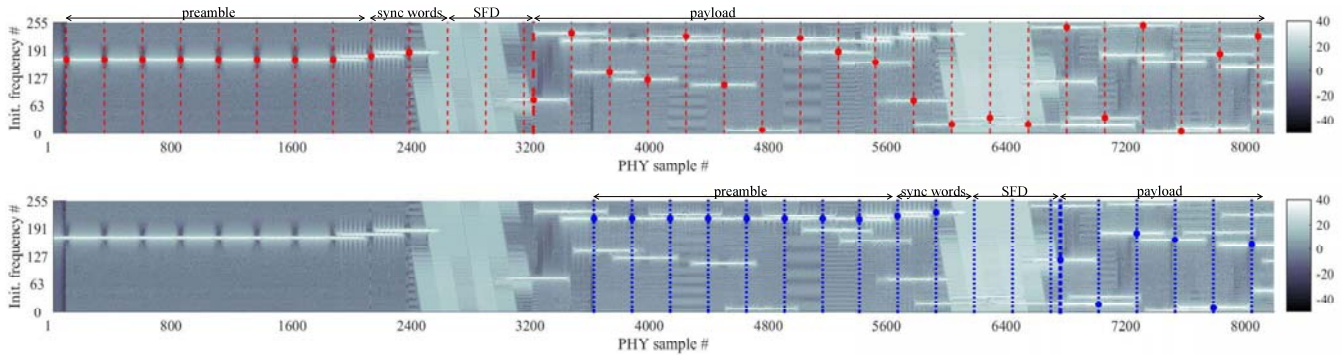


Figure 12: Separating the collided signals of Figure 1 into two LoRa frames: the frame structure and demodulated symbols of (a) the first and (b) second transmission.




Figure 13: Experiment equipment.

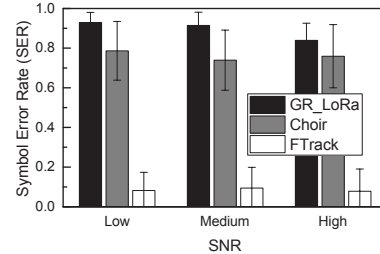


Figure 14: Demodulation with two-node collisions.

set of the divided signals that have similar power levels, and use adaptive power thresholds to detect frequency tracks.

Combating the absence of LoRa configurations. A LoRa packet transmitted at given bandwidth (BW) and spreading factor (SF) can only be demodulated with a unique chirp $C^{-1}(t)$ associated with the same configurations. Generally, BW and SF are chosen by a transmitter. A receiver may not know the configurations of received frames. With FTrack, we can detect BW and SF from a short segment of the received signals (*e.g.*, ≤ 5 chirps). We exploit the fact that chirps are transformed into horizontal frequency tracks only if down-chirp $C^{-1}(t)$ is produced with the correct BW and SF. One approach is to detect the existence of frequency tracks by trying all BW-SF combinations. The search space is (10 BWs \times 7 SFs) for the Semtech SX1276 LoRa transmitter [46]. To accelerate the process, we can process the PHY samples in parallel with multiple threads, each of which is configured by a particular pair of BW and SF. Specifically, if two packets transmitted with different configurations collide, due to the orthogonality of transmission, they will be demodulated separately by different instances of FTrack.

4 EVALUATION

We implement and evaluate the system using software radio base stations and commodity LoRa devices. For performance evaluation, we use both high-end software defined radio (*i.e.*, USRP N210) and low-cost receive-only software defined radio (*i.e.*, RTL-SDR dongle) as shown in Figure 13(a). We develop our own LoRa demodulator based on the GNU Radio library, and implement FTrack in MATLAB

to process PHY samples. If not otherwise specified, we employ RTL-SDR dongles to receive PHY samples at the 900 MHz bands. The USRP N210 is only used for performance evaluations.

The LoRa nodes are composed of Dragino LoRa shields [13], which consist of HopeRF’s RFM96W transceiver module embedded with the Semtech SX1276 chip, as shown in Figure 13(b). We connect the LoRa shield to Arduino Uno motherboard to control SX1276 chip in packet transmission and reception. The SX1276 chip operates at 915MHz with the bandwidth of 250kHz or 125kHz, depending on the configuration of LoRa parameters. We note that radio broadcasting, RFID and P-GSM-900 may also work in the 900MHz frequency band. The measured noise floor is around -90dBm in our experiments. By default, we set the spreading factor (SF), coding rate (CR) and bandwidth (BW) of LoRa communication as 8, 4/5 and 250kHz, respectively. The sampling rate of RTL-SDR dongle is 1 MS/s.

In the following, we first evaluate the performance of FTrack in two-node collisions and consider various factors (*e.g.*, timing offsets, SNR conditions, LoRa parameters, *etc.*). We examine the capability of FTrack on parallel decoding in §4.2. We evaluate the performance of FTrack in real usage scenarios in §4.3.

4.1 Parallel Demodulation

4.1.1 Basic performance. In the following experiment, we use an RTL-SDR dongle as a receiver and three LoRa nodes as transmitters in an indoor environment. To experiment with varied colliding time, we configure one transmitter to send beacons every 2 seconds. Upon receiving such beacons, the other two LoRa nodes reply a

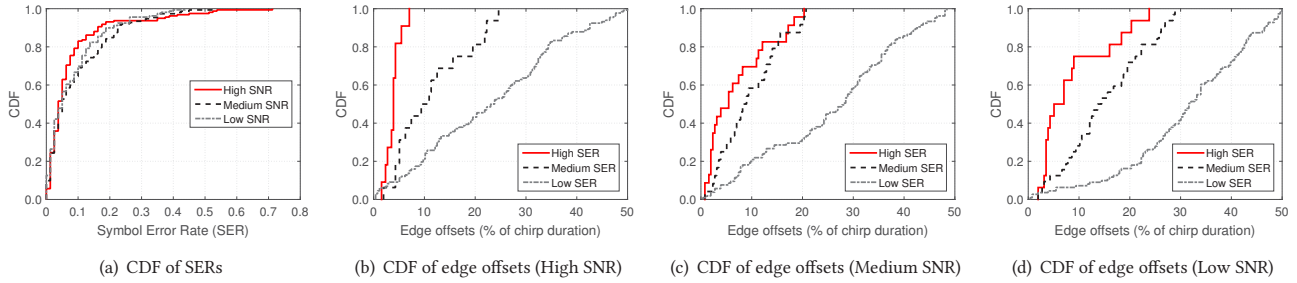


Figure 15: In-depth study of FTrack performance: (a) CDF of the Symbol Error Rates (SERs); (b-d) Relationships between SERs and the normalized edge offsets of colliding symbols: FTrack produces high SERs when colliding symbols are closely-aligned in time (*i.e.*, with short edge offsets).

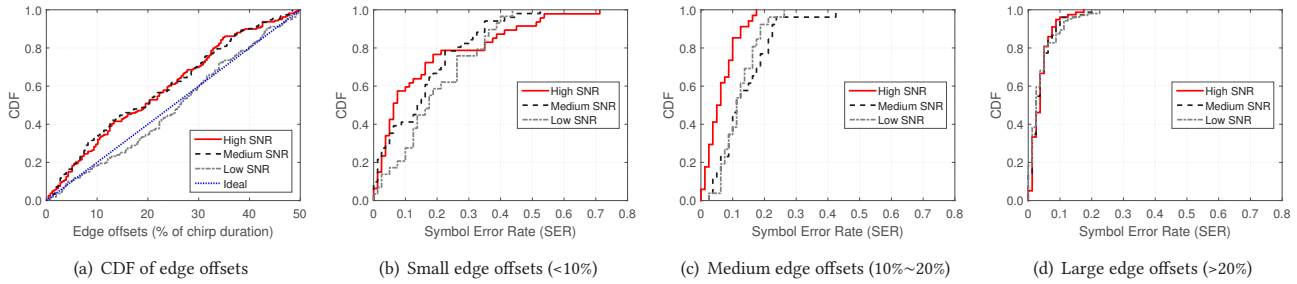


Figure 16: Relationship between the demodulation performance of FTrack and edge offsets of colliding symbols: (a) CDF of edge offsets in the collected frames, (b-d) Demodulation performance when edge offsets are small (<10%), medium (10%~20%) and large (>20%), respectively.

30-Byte data frame, which consists of 61 payload symbols and lasts for about 80ms when SF=8, BW=250kHz. Specifically, we configure the two LoRa nodes to delay for a random period of time before transmission with a maximum delay of 20ms (*i.e.*, the air time of 20 symbols). The experiment setting can result in collisions at different parts of packets (*e.g.*, preambles, sync words and payloads) as reported in [39]. We vary the transmission power of LoRa nodes to evaluate demodulation performance across three SNR regimes: low (<5dB), medium (5-20 dB) and high (>20 dB). For each SNR regime, we collect 500 colliding frames and repeat the experiment 5 times. We compare FTrack against GR_LoRa (*i.e.*, a standard LoRa demodulation scheme) and Choir [15] which represents the state-of-the-art on LoRa collision recovery.

Figure 14 compares the average SERs of the three demodulation schemes when two LoRa nodes transmit concurrently. According to the experiment results, FTrack performs the best and GR_LoRa is the worst (SER > 80%) because GR_LoRa is not capable of recovering any collisions. Our experimental study reveals that the hardware frequency offset (*i.e.*, the fractional part of chirp frequency) extracted by Choir is not reliable to classify colliding symbols in practice. The main reason is that along with the hardware imperfection, various influencing factors (*e.g.*, phase jitters, noise) may cause variations in the measurement of fractional part of carrier frequency. We also notice that the time offset between a detection window and a symbol edge also influences the measurement of fractional part of

carrier frequency. As a result, Choir suffers high symbol error rates (70%~80%) in our experiments. In contrast, FTrack can leverage the timing misalignment to separate colliding symbols. It yields <10% symbol error rates across all SNR regimes. The low symbol error rates of FTrack can be corrected by standard error correction schemes (*e.g.*, Hamming code) adopted by current LoRa nodes.

We present the CDF of SERs of FTrack in Figure 15(a). As we can see, FTrack achieves better performance in the high SNR condition than in the low SNR condition. More importantly, the performance of FTrack does not degrade dramatically even in the low SNR condition. That is because FTrack can still leverage the long transmission time of a LoRa packet to boost the signal strength. In particular, 80% of the symbol error rates are below 10% when SNR is high. Even in the medium and low SNR conditions, nearly 70% of the collided frames are demodulated with <10% SERs.

4.1.2 Impact of time offsets and SNR. We analyze the major factors that may influence the performance of FTrack. Figures 15(b-d) examine the relationships between SERs and the timing offsets of collided symbols. In each SNR regime, we divide the demodulation results into three groups for analysis: low SER (<0.1), medium SER (0.1~0.2) and high SER (>0.2). We observe that in the high-SNR condition, 100% of high SERs and 50% of medium SERs appear when the edges of colliding symbols are closer than 10% of a chirp duration, as shown in Figure 15(b). Similar results are observed in

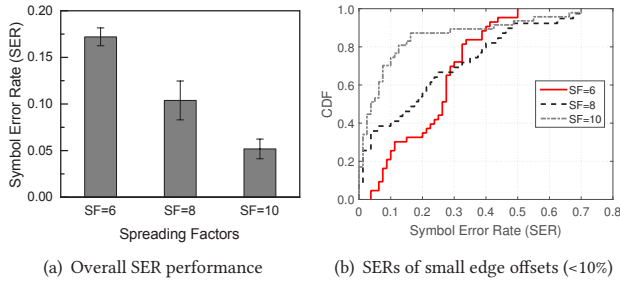


Figure 17: Impact of the Spreading Factor (SF) of LoRa PHY modulation, BW=250kHz: (a) SERs under different SFs, (b) Demodulation results when edge offsets are small (<10%).

the medium and low SNR cases, as shown in Figures 15(c,d). That is because once the symbols are not separated in time with sufficient margin, FTrack may not be able to separate such edges, leading to incorrect symbol grouping and demodulating results.

In practice, LoRa nodes may transmit packets in random time slots. Figure 16(a) shows the edge offsets of collided symbols in our collected traces, where edge offsets are normalized in percentages of a chirp duration. As LoRa nodes transmit at random time, the edge offsets of colliding symbols distribute uniformly across the whole range of one chirp duration (*i.e.*, length of a LoRa symbol). As we mentioned, FTrack may fail to demodulate concurrent transmissions when the edge offsets are short (*e.g.*, <10% chirp duration), while FTrack can successfully recover collisions in most other cases.

We examine the symbol error rates of demodulation when the edge offsets of colliding symbols are small (<10%), medium (10%~20%) and large (>20%), respectively. By comparing Figures 16(b-d), we see that FTrack achieves lower SERs as the edge offsets of interfering symbols increases. For instance, as shown in Figure 16(b), when the edge offsets are shorter than 10% chirp duration, nearly 40% of the results of high-SNR (60% of medium-SNR and 80% of low-SNR) have high or medium symbol error rates (*i.e.*, SER>0.1, collisions may not be recovered). This number decreases to lower than 5% when the edge offsets increases to 20%, as shown in Figure 16(d). In this case, 80% of the results have SERs<0.05.

In addition, by increasing the SNR of received signals, FTrack can produce better demodulation results. For example, as shown in Figures 15(b) and (d), when the SNR is low, 82% of low SERs are produced in the case that symbol edges are apart farther than 20% of the symbol length; whereas the results of low SER appear uniformly in all edge offset occasions when SNR becomes high. Therefore, as the SNR increases, FTrack may demodulate with low SERs even when the edges of colliding symbols are closely located. This implies that we can increase the transmission power of LoRa nodes for better collision recovery performance.

4.1.3 Impact of LoRa packet configuration. The demodulation performance of FTrack can be influenced by the symbol duration of LoRa packets. Some parameters configure the symbol duration of LoRa (*e.g.*, Spreading Factor (SF) and Bandwidth (BW)). In the following, we investigate the impact of SF and BW on collision recovery performance. Unless otherwise specified, we adopt the same experimental settings as in §4.1.1. We only present the evaluation

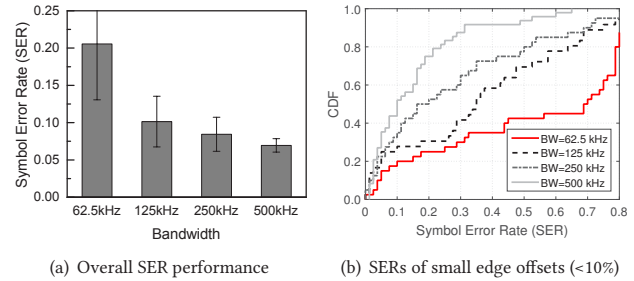


Figure 18: Impact of LoRa Bandwidth (BW), SF=8: (a) SERs under different BW settings, (b) Demodulation results when edge offsets are small (<10%).

results of high-SNR. The experiment results exhibit similar trends in both medium-SNR and low-SNR regimes (not presented).

We set the Bandwidth of LoRa to 250kHz and study the impact of Spreading Factor on demodulation performance. Figure 17(a) presents the overall symbol error rates of FTrack with varied SFs (*e.g.*, SF=6, 8, and 10). We see that FTrack achieves lower SERs with a larger SF. The average SER decreases from 0.17 to 0.05 as SF increases from 6 to 10. In particular, a larger SF facilitates collision recovery especially when the edge offsets of interfering symbols are small, as shown in Figure 17(b). When SF=6, only 20% of the collisions with small edge offset can be correctly recovered (*i.e.*, SER<0.1). This number increases to 40% for SF=8 and around 70% for SF=10. That is because when SF increases, a symbol takes a longer air time [32], which helps FTrack separate collisions in time.

Next, Figure 18 evaluates the demodulation performance of FTrack with varied BWs. The Spreading Factor is fixed to 8. The results show that FTrack performs better with a larger BW. The average SER decreases from 0.21 to 0.07 as the LoRa bandwidth increases from 62.5kHz to 500kHz. That is because when the LoRa bandwidth increases, the frequency gap between the symbols coexisting within a demodulation window also increases, which helps the differentiation of symbols in the frequency domain. As shown in Figure 18(b), when the edge offsets of colliding symbols are small, even though it is hard to separate symbols in time, FTrack recovers more collisions with larger BWs. For instance, less than 20% collisions can be correctly recovered (*i.e.*, SER<0.1) when BW=62.5kHz. The number increases to about 45% when BW=500kHz. However, we note that a larger BW comes at the cost of an increased number of PHY samples and higher computational overhead to process the PHY samples.

4.2 FTrack Capability

Impact of a varied number of LoRa nodes. We examine the capability of FTrack on collision recovery with an increasing number of colliding nodes. In this experiment, we use the same configuration as in §4.1.1 and increase the number of concurrent transmissions from 1 node to 10 nodes. We measure the PHY-layer symbol error rate (SER) and the ratio of packets being correctly decoded (*i.e.*, packet reception ratio, PRR). As shown in Figure 19, the demodulation errors of FTrack increase as more nodes transmit in parallel. For instance, the average SER is 7% when two nodes collide. Yet it

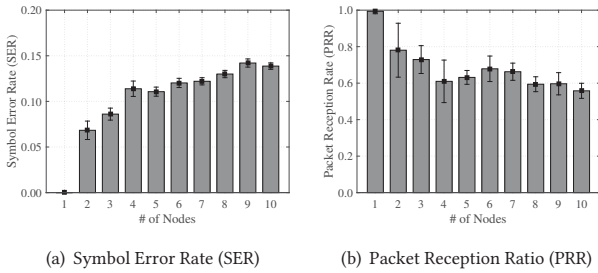


Figure 19: Performance of FTrack on decoding a varied number of parallel transmissions.

| # of transmissions | N=1 | N=2 | N=4 | N=8 |
|--------------------|---------|---------|---------|---------|
| Choir | 0.246 s | 0.328 s | 0.443 s | 0.475 s |
| FTrack | 0.139 s | 0.153 s | 0.205 s | 0.373 s |

Table 1: Comparison of computational time (in second).

increases to 14% in the case of 10-node collisions. When two nodes transmit concurrently, nearly 80% of the packets can be successfully recovered. The packet reception ratio, however, decreases as more nodes collide. When 10 nodes transmit in parallel, although we cannot recover all packets, 58% of the packets can still be correctly decoded leading to a considerable throughput gain.

Computational overhead. Table 1 compares the computational time of Choir and FTrack when running on a PC with the Intel Core i5 CPU processor. In general, the computational time of FTrack is shorter than that of Choir. The decoding time increases proportionally with the number of collided packets. As the data rates of LoRa are much lower than other wireless technologies (e.g., WiFi, RFID), it leaves sufficient time for LoRa base stations to decode collisions. In the future, we plan to optimize and adopt more efficient software and hardware designs (e.g., multi-thread, FPGA, ASIC) to accelerate FTrack so as to meet the real-time processing requirement.

4.3 Performance in Real Network

In the following, we evaluate how FTrack performs in practice. We deploy the testbed LoRaWAN network within our office building, as illustrated in Figure 20. Each node senses the environments (e.g., light, temperature, humidity, etc.) and randomly selects a 200ms-slot to transmit data to a base station. The payload is 30 Bytes, corresponding to about 80 PHY symbols (payloads + headers) and 80ms air time when SF=8, BW=250kHz. The duty cycle of LoRa nodes are set to 10% for evaluation purpose. To reduce power consumption, LoRa node can adopt a much lower duty cycle (e.g., <2%) in real usage scenarios that results in lower traffic rates and thus less severe collisions. In particular, the ordinary LoRa nodes in our testbed can transmit in 8 uplink slots during each data collection cycle (the rest two slots are reserved by the LoRaWAN testbed to send control messages). We increase the number of nodes from 1 to 20 to compare the scalability of different approaches. We compare FTrack with GR_LoRa [27] and Choir [15]. We also compare the performance against an Oracle scheme that is assumed to optimally schedule the LoRa nodes such that no collision could happen.

Figure 21(a) shows the network throughput of four approaches. We see that the throughput of all approaches increase as more nodes join the network, when the network size is small (e.g., # nodes ≤ 4).



Figure 20: Layout of a LoRaWAN testbed.

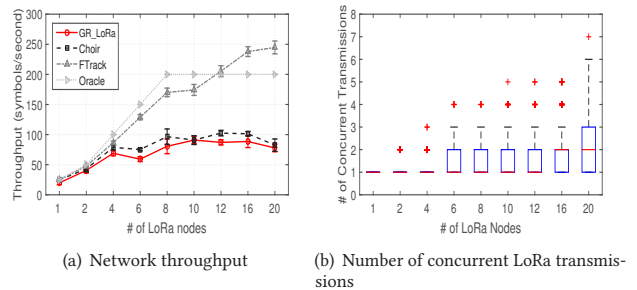


Figure 21: Performance of decoding concurrent transmissions in real-world LoRaWAN network (SF=8, BW=250kHz).

However, the throughput of GR_LoRa and Choir saturate (around 80 symbols/sec) rapidly when the network size increases to 4 nodes. That is because as more nodes transmit concurrently, GR_LoRa and Choir cannot recover collisions among the nodes. Benefiting from a perfect transmission schedule, Oracle yields the highest throughput that increases linearly with network size as the number of nodes increases from 1 to 8. However, the Oracle reaches its capacity limit (200 symbols/sec) when network size is 8. In contrast, the throughput of FTrack further increases to 250 symbols/second as network size increases to 20 nodes, which is about 3 times the throughput of Choir and GR_LoRa. The performance gain stems from FTrack’s capability of decoding concurrent LoRa transmissions.

Figure 21(b) presents the number of concurrent transmissions with a varied number of LoRa nodes working in the low-duty cycle mode. We see that when the network size exceeds 4, nodes collide more frequently even if they work in the low-duty cycle mode. As such, the performance of GR_LoRa and Choir starts to degrade when the number of nodes exceeds 4, as shown in Figure 21(a). As the network size further increases, we observe more collisions: A median of 2 and a maximum of 5 nodes transmit concurrently at a slot when the network size exceeds 16. In this case, FTrack achieves the highest network throughput as it is able to recover collisions.

5 RELATED WORK

A variety of LP-WAN technologies [7, 29] have been proposed to enable the fast-growing IoT applications. SigFox [47] uses Ultra

Narrow Band (UNB) technology combined with DBPSK and GFSK modulation to support low-power long-range communication in the ISM band. NB-IoT [41] and LTE-M [28] are introduced by 3GPP. They use a subset of the LTE standard, operate at licensed spectrum yet limit the bandwidth to a single narrow-band of 200kHz. LoRa [3, 7], like SigFox, works at license-free frequency band. It employs Chirp Spread Spectrum (CSS) modulation to transmit data, which is resilient to interference, multi-path fading and Doppler effects. Such characteristics make LoRa a competitive candidate for long-range low-cost IoT networks [5, 12]. We focus on LoRa in this paper and refer the readers to [48–50] for detailed comparisons of existing LP-WAN technologies.

There are some pioneer researches on LoRa and LoRaWAN [21]. Early efforts have been devoted to the measurement study [8, 11, 32] and performance analysis, such as transmission air time [30, 34], power consumption [10, 16, 32], coverage [36, 37], *etc.* Based on these studies, some improvement schemes [1, 33, 42] are introduced for better performance. Although the LoRa PHY is proprietary, authors of [27] and [44] employ reverse engineering to study the encoding and decoding schemes of LoRa.

With respect to the limited capability of hardware, LoRaWAN [3] adopts a simple ALOHA-based MAC for access control. The limits of LoRaWAN MAC have been analytically concluded in [2, 6]. [9] studies LoRa collisions via simulation. [20] employs commodity devices to empirically study the characteristic of LoRa collisions. More recently, [39] presents an in-depth investigation of LoRa collisions within actual running networks. To avoid collisions, some researchers [38, 43] proposed new MAC designs that incorporate advanced scheduling schemes on top of LoRa. However, these schemes would add higher complexity yet produce lower efficiency in terms of network capacity and throughput.

The work most related to ours is Choir [15], which exploits the frequency offsets introduced by LoRa hardware to separate collisions. However, in practice, the extracted frequency offset is not reliable to classify colliding symbols due to various influencing factors (*e.g.*, phase jitters, time offset).

Our work is inspired by the previous works of collision recovery and parallel decoding in various wireless systems (*e.g.*, WiFi, RFID). ZigZag [17] decodes collisions by exploiting the fact that the time offset of collided transmissions produces some interference-free chunks. It extracts the interference-free chunks and subtracts from collisions to separate each individual transmission. BiGroup [35] examines the collision states of concurrent transmissions of RFID tags, and iteratively detects the state transitions of collided signals to decode tag transmissions. LF-Backscatter [23] employs powerful RFID readers to detect the interleaving signal edges with high sampling rates and separate collided tag signals. More recently, Flip-Tracer [26] and Hubble [25] support parallel decoding of backscatter communications by leveraging both PHY and time domain information. NetScatter [22] proposes a new coding scheme that combines On-Off Keying and CSS to support concurrent transmissions.

In addition to the collision recovery approach, some works aim to *avoid collisions*, including MIMO [4, 40, 53], TDMA [45, 52], collision-recovery methods [24, 51], and constructive interference [14, 31]. LoRa also supports orthogonal communications by transmitting with different channels, bandwidths and spreading factors. However, even with orthogonal parameter settings, transmission

collisions may still happen as the number of devices further increases, *e.g.*, in scenarios like urban and warehouse [20, 39]. Our work is complementary to such collision avoidance schemes.

6 LIMITATIONS AND FUTURE WORK

FTrack exploits the timing misalignment of collided frames to enable parallel decoding. The evaluation results indicate that the performance of FTrack is affected by the edge offsets of collided symbols. When the symbol edges of two colliding LoRa nodes are well-aligned (*e.g.*, <10% of symbol duration), the performance of FTrack will degrade. In practice, some LoRaWAN MAC protocols schedule LoRa transmissions with CSMA or random slotting mechanisms, which can lead to collisions in synchronized time slots. As a result, the symbol edges of collided frames can be aligned, which could potentially affect the capability of FTrack in resolving collisions. To address this problem, we can modify the MAC protocols in practice to intentionally mis-align collided transmissions. Moreover, we can properly select the parameters of LoRa communication (*e.g.*, spreading factor, bandwidth, transmission power, *etc.*) to maximize the capability of FTrack in supporting more parallel transmissions. We leave such a exploration for future work.

7 CONCLUSION

This paper presents FTrack, a practical strategy that resolves LoRa collisions in both time and frequency domains. FTrack jointly exploits the distinct frequency tracks and misaligned edges of LoRa symbols to separate collisions. It enables a novel communication paradigm that allows LoRa node to join on-going communications in parallel without specific coordination. We implement FTrack on a low-cost SDR platform and deploy an indoor testbed to evaluate the performance of FTrack in a variety of network settings. Results show that FTrack recovers collided LoRa frames with low symbol error rates in diverse SNR conditions. It can boost the throughput of real-world LoRaWAN by up to 3 times. The parallel decoding capability of FTrack can benefit the deployment of large-scale LoRaWAN in densely-populated IoT scenarios.

8 ACKNOWLEDGEMENTS

We thank anonymous shepherd and reviewers for helpful comments. This work is supported in part by the National Nature Science Foundation of China under grant 61702437 and Hong Kong GRF under grant PolyU 152165/19E, and in part by the Australian Research Council (ARC) Discovery Project Grants DP190101888 and DP180103932. Yuanqing Zheng is the corresponding author.

REFERENCES

- [1] Khaled Abdelfadeel, Victor Cionca, and Dirk Pesch. 2018. Poster: A Fair Adaptive Data Rate Algorithm for LoRaWAN. In *Proceedings of the 2018 International Conference on Embedded Wireless Systems and Networks (EWSN'18)*. 169–170.
- [2] Ferran Adelantado, Xavier Vilajosana, Pere Tuset-Peiro, Borja Martinez, Joan Melia-Segui, and Thomas Watteyne. 2017. Understanding the Limits of LoRaWAN. *IEEE Communications Magazine* 55, 9 (Sept. 2017), 34–40.
- [3] LoRa Alliance. 2019. LoRaWAN for Developer. Retrieved Jan 10, 2019 from <https://lora-alliance.org/lorawan-for-developers>
- [4] Narendra Anand, Ryan E. Guerra, and Edward W. Knightly. 2014. The Case for UHF-band MU-MIMO. In *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking (MobiCom '14)*. 29–40.
- [5] Aloïs Augustin, Jiazi Yi, Thomas Clausen, and William Mark Townsley. 2016. A Study of LoRa: Long Range & Low Power Networks for the Internet of Things. *Sensors* 16, 9 (Sept. 2016), 1446.

- [6] Dmitry Bankov, Evgeny Khorov, and Andrey Lyakhov. 2016. On the Limits of LoRaWAN Channel Access. In *Proceedings of the 2016 International Conference on Engineering and Telecommunication (EnT)*. 10–14.
- [7] Jean-Paul Bardyn, Thierry Melly, Olivier Seller, and Nicolas Sornin. 2016. IoT: The era of LPWAN is starting now. In *Proceedings of the ESSCIRC Conference 2016: 42nd European Solid-State Circuits Conference*. 25–30.
- [8] Norbert Blenn and Fernando A. Kuipers. 2017. LoRaWAN in the Wild: Measurements from The Things Network. *CoRR(arXiv preprint) abs/1706.03086* (2017).
- [9] Martin C. Bor, Utz Roedig, Thimo Voigt, and Juan M. Alonso. 2016. Do LoRa Low-Power Wide-Area Networks Scale?. In *Proceedings of the 19th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM'16)*. 59–67.
- [10] Taoufik Bouguera, Jean-François Diouris, Jean-Jacques Chaillout, Randa Jaouadi, and Guillaume Andrieux. 2018. Energy Consumption Model for Sensor Nodes Based on LoRa and LoRaWAN. *Sensors* 18, 7 (Jul 2018), 2104.
- [11] Anders Carlsson, Ievgeniia Kuzminykh, Robin Franksson, and Alexander Liljgren. 2018. Measuring a LoRa Network: Performance, Possibilities and Limitations. In *Proceedings of the 18th International Conference, NEW2AN 2018, and 11th Conference, ruSMART 2018*. 116–128.
- [12] Jonathan de Carvalho Silva, Joel J.P.C. Rodrigues, Antonio M. Alberti, Petar Solic, and Andre L. L. Aquino. 2017. LoRaWAN—A low power WAN protocol for Internet of Things: A review and opportunities. In *Proceedings of the 2017 2nd International Multidisciplinary Conference on Computer and Energy Science (SpliTech)*. 1–6.
- [13] Dragino. 2019. LoRa Shield for Arduino. Retrieved Mar 25, 2019 from <http://www.dragino.com/products/module/item/102-lora-shield.html>
- [14] Wan Du, Jansen Christian Liando, Huanle Zhang, and Mo Li. 2015. When pipelines meet fountain: Fast data dissemination in wireless sensor networks. In *Proceedings of the ACM Conference on Embedded Networked Sensor Systems*. 365–378.
- [15] R. Eletreby, D. Zhang, S. Kumar, and O. Yagan. 2017. Empowering Low-Power Wide Area Networks in Urban Settings. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication (SIGCOMM'17)*. 309–321.
- [16] Weifeng Gao, Wan Du, Zhiwei Zhao, Geyong Min, and Mukesh Singhal. 2019. Towards Energy-Fairness in LoRa Networks. In *Proceedings of the 39th IEEE International Conference on Distributed Computing Systems (ICDCS)*.
- [17] Shyamath Gollakota and Dina Katabi. 2008. Zigzag decoding: combating hidden terminals in wireless networks. In *Proceedings of the ACM SIGCOMM 2008 conference on Data communication (SIGCOMM'08)*. 159–170.
- [18] D.J. Goodman and A.A.M. Saleh. 1987. The near/far effect in local ALOHA radio communications. *IEEE Transactions on Vehicular Technology* 36, 1 (Feb. 1987), 19–27.
- [19] Daniel Halperin, Thomas Anderson, and David Wetherall. 2008. Taking the Sting out of Carrier Sense: Interference Cancellation for Wireless LANs. In *Proceedings of the 14th ACM International Conference on Mobile Computing and Networking (MobiCom'08)*. 339–350.
- [20] Jetmir Haxhibeqiri, Floris Van den Abele, Ingrid Moerman, and Jeroen Hoebeke. 2017. LoRa Scalability: A Simulation Model Based on Interference Measurements. *Sensors* 17, 6 (Mar 2017), 1193.
- [21] Jetmir Haxhibeqiri, Eli De Poorter, Ingrid Moerman, and Jeroen Hoebeke. 2018. A Survey of LoRaWAN for IoT: From Technology to Application. *Sensors* 18, 11 (Nov. 2018), 3995.
- [22] Mehrdad Hesar, Ali Najafi, and Shyamath Gollakota. 2019. NetScatter: Enabling Large-Scale Backscatter Networks. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI'19)*. 271–284.
- [23] Pan Hu, Pengyu Zhang, and Deepak Ganesan. 2015. Laissez-Faire: Fully Asymmetric Backscatter Communication. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication (SIGCOMM'15)*. 255–267.
- [24] Kyle Jamieson and Hari Balakrishnan. 2007. PPR: Partial Packet Recovery for Wireless Networks. In *Proceedings of the 2007 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM'07)*. 409–420.
- [25] Meng Jin, Yuan He, Xin Meng, Dingyi Fang, and Xiaojiang Chen. 2018. Parallel Backscatter in the Wild: When Burstiness and Randomness Play with You. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking (MobiCom '18)*. 471–485.
- [26] Meng Jin, Yuan He, Xin Meng, Yilun Zheng, Dingyi Fang, and Xiaojiang Chen. 2017. FlipTracer: Practical Parallel Decoding for Backscatter Communication. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking (MobiCom '17)*. 275–287.
- [27] Matthew Knight and Balint Seeber. 2016. Decoding LoRa : Realizing a Modern LPWAN with SDR. In *Proceedings of the 6th GNU Radio Conference*. 1–5.
- [28] Mads Lauridsen, Istvan Z. Kovacs, Preben Mogensen, Mads Sorensen, and Steffen Holst. 2016. Coverage and Capacity Analysis of LTE-M and NB-IoT in a Rural Area. In *Proceedings of IEEE 84th Vehicular Technology Conference (VTC-Fall 2016)*. 1–5.
- [29] Alexandru Lavric and Adrian Ioan Petrariu. 2018. LoRaWAN communication protocol: The new era of IoT. In *Proceedings of the 2018 International Conference on Development and Application Systems (DAS)*. 74–77.
- [30] Alexandru Lavric and Valentin Popa. 2017. A LoRaWAN: Long range wide area networks study. In *Proceedings of the 2017 International Conference on Electromechanical and Power Systems (SIELMEN)*. 417–420.
- [31] Zhenjiang Li, Wan Du, Yuanqing Zheng, Mo Li, and Dapeng Wu. 2017. From rateless to hopless. *IEEE/ACM transactions on networking* 25, 1 (2017), 69–82.
- [32] J.C. Liando, A. Gamage, A.W. Tengourtius, and M. Li. 2019. Known and Unknown Facts of LoRa: Experiences from a Large Scale Measurement Study. *ACM Transactions on Sensor Networks (TOSN)* 15, 2 (Feb. 2019), 16:1–16:35.
- [33] Jin-Taek Lim and Youngnam Han. 2018. Spreading Factor Allocation for Massive Connectivity in LoRa Systems. *IEEE Communications Letters* 22, 4 (Apr 2018), 800–803.
- [34] Umer Noreen, Ahdhane Bounceur, and Laurent Clavier. 2017. A study of LoRa low power and wide area network technology. In *Proceedings of the 2017 International Conference on Advanced Technologies for Signal and Image Processing (ATSIP)*. 1–6.
- [35] Jiajue Ou, Mo Li, and Yuanqing Zheng. 2015. Come and Be Served: Parallel Decoding for COTS RFID Tags. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking (MobiCom'15)*. 500–511.
- [36] Juha Petajarvi, Konstantin Mikhaylov, Antti Roivainen, Tuomo Hanninen, and Marko Pettissalo. 2015. On the coverage of LPWANs: range evaluation and channel attenuation model for LoRa technology. In *Proceedings of the 2015 14th International Conference on ITS Telecommunications (ITST)*. 55–59.
- [37] Juha PetÄdjÄdrvi, Konstantin Mikhaylov, Marko Pettissalo, Janne Janhunen, and Jari Inatti. 2017. Performance of a low-power wide-area network based on LoRa technology: Doppler robustness, scalability, and coverage. *International Journal of Distributed Sensor Networks* 13, 3 (Mar 2017), 1–16.
- [38] Congduc Pham. 2018. Investigating and experimenting CSMA channel access mechanisms for LoRa IoT networks. In *Proceedings of the 2018 IEEE Wireless Communications and Networking Conference (WCNC)*. 1–6.
- [39] Andri Rahmadhani and Fernando Kuipers. 2018. When LoRaWAN Frames Collide. In *Proceedings of the 12th International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization (WiNTECH'18)*. 89–97.
- [40] Hariharan Shankar Rahul, Swaran Kumar, and Dina Katabi. 2012. JMB: Scaling Wireless Capacity with User Demands. In *Proceedings of the ACM SIGCOMM 2012 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM '12)*. 235–246.
- [41] Rapeepat Ratasuk, Benny Vejlgaard, Nitin Mangalvedhe, and Amitava Ghosh. 2016. NB-IoT system for M2M communication. In *Proceedings of IEEE Wireless Communications and Networking Conference (WCNC 2016)*. 1–5.
- [42] Brecht Reynders, Wannas Meert, and Sofie Pollin. 2017. Power and spreading factor control in low power wide area networks. In *Proceedings of the 2017 IEEE International Conference on Communications (ICC)*. 1–6.
- [43] Brecht Reynders, Qing Wang, Pere Tuset-Peiro, Xavier Vilajosana, and Sofie Pollin. 2018. Improving Reliability and Scalability of LoRaWANs Through Lightweight Scheduling. *IEEE Internet of Things Journal* 5, 3 (Jun 2018), 1830–1842.
- [44] Pieter Robyns, Peter Quax, Wim Lamotte, and William Thenaers. 2018. A Multi-Channel Software Decoder for the LoRa Modulation Scheme. In *Proceedings of the 3rd International Conference on Internet of Things, Big Data and Security (IoTBSDS)*. 1–11.
- [45] Anthony Rowe, Dhiraj Goel, and Raguathan Rajkumar. 2007. FireFly Mosaic: A Vision-Enabled Wireless Sensor Networking System. In *Proceedings of the 28th IEEE International Real-Time Systems Symposium (RTSS 2007)*. 459–468.
- [46] Semtech. 2019. Semtech SX1276: 137MHz to 1020MHz Long Range Low Power Transceiver. Retrieved Apr 5, 2019 from <https://www.semtech.com/products/wireless-rl/lora-transceivers/sx1276>
- [47] SigFox. 2019. SigFox Overview. Retrieved Jan 10, 2019 from <https://www.sigfox.com/en/sigfox-iot-technology-overview>
- [48] Rashmi Sharan Sinha, Yiqiao Wei, and Seung-Hoon Hwang. 2017. A survey on LPWA technology: LoRa and NB-IoT. *ICT Express* 3, 1 (March 2017), 14–21.
- [49] Jothi Prasanna Shanmuga Sundaram, Wan Du, and Zhiwei Zhao. 2019. A Survey on LoRa Networking: Research Problems, Current Solutions and Open Issues. *arXiv preprint arXiv:1908.10195* (2019).
- [50] Benny Vejlgaard, Mads Lauridsen, Huan Nguyen, Istvan Z. Kovacs, Preben Mogensen, and Mads Sorensen. 2017. Coverage and Capacity Analysis of Sigfox, LoRa, GPRS, and NB-IoT. In *Proceedings of the 2017 IEEE 85th Vehicular Technology Conference (VTC Spring)*. 1–5.
- [51] Yafeng Wu, Gang Zhou, and John A. Stankovic. 2010. ACR: Active Collision Recovery in Dense Wireless Sensor Networks. In *Proceedings of the IEEE INFOCOM 2010*. 1–9.
- [52] Xianjin Xia, Shining Li, Yu Zhang, Lin Li, Tao Gu, and et al. 2017. Surviving screen-off battery through out-of-band Wi-Fi coordination. In *Proceedings of the IEEE INFOCOM 2017 - IEEE Conference on Computer Communications*. 1–9.
- [53] Yaxiong Xie, Jie Xiong, Mo Li, and Kyle Jamieson. 2019. mD-Track: Leveraging Multi-Dimensionality for Passive Indoor Wi-Fi Tracking. In *Proceedings of MobiCom'19*. 8:1–8:16.