

MaLoRaGW: Multi-User MIMO Transmission for LoRa

Hossein Pirayesh, Shichen Zhang, Pedram Kheirkhah Sangdeh, and Huacheng Zeng

Department of Computer Science and Engineering, Michigan State University

Email: {pirayesh, sczhang, sangdeh, hzeng}@msu.edu

ABSTRACT

LoRa has emerged as a key wireless communication technology for a gateway to provide geographically-distributed IoT devices with low-rate, long-range connections. In this paper, we present MaLoRaGW, the first-of-its-kind Multi-antenna LoRa Gateway that enables multi-user MIMO (MU-MIMO) LoRa communications in both uplink and downlink. MaLoRaGW was inspired by the success of MU-MIMO in cellular and Wi-Fi networks. The key component of MaLoRaGW is a joint baseband PHY design for uplink packet detection and downlink beamforming. Its innovation lies in three modules: spatial signal projection, accurate channel estimation, and implicit beamforming, all of which reside only in a LoRa gateway and require no modification on LoRa client devices. We have built a prototype of two-antenna MaLoRaGW on a USRP device and extensively evaluated its performance with commercial LoRa dongles in three scenarios: lab, office building, and university campus. Our experimental results show that, compared to the state-of-the-art, the two-antenna MaLoRaGW increases uplink throughput by 10% and downlink throughput by 95%.

CCS CONCEPTS

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

KEYWORDS

LoRa, chirp spread spectrum, multi-user MIMO, collision recovery, beamforming, throughput

ACM Reference Format:

Hossein Pirayesh, Shichen Zhang, Pedram Kheirkhah Sangdeh, and Huacheng Zeng. 2022. MaLoRaGW: Multi-User MIMO Transmission for LoRa. In *The 20th ACM Conference on Embedded Networked Sensor Systems (SenSys '22)*, November 6–9, 2022, Boston, MA, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3560905.3568533>

1 INTRODUCTION

Recent years have witnessed a continuous and rapid increase in the number of low-cost Internet-of-Thing (IoT) devices, most of which have a low data rate requirement but desire to prolong their battery lifetime [8, 16, 25, 26]. LoRaWAN has emerged as a key wireless communication technology to connect a large number of geographically-distributed IoT devices with a single gateway

[27, 34, 45]. LoRa devices employ chirp spread spectrum (CSS) modulation at the physical layer, which allows a LoRa receiver to decode packets at very low signal-to-noise (SNR) scenarios (e.g., -20dB [42]) and therefore permits a long communication range (e.g., 8km), as indicated by its name.

A central problem with LoRaWAN is packet collision, which fundamentally limits its packet delivery rate and network throughput in user-dense scenarios [10, 13, 18, 26]. Different collision recovery approaches have been proposed to address this problem [3, 11, 17, 22, 29, 32, 33, 36, 39]. Most of these approaches aim to decode collided LoRa packets by exploiting the unique signal features in the frequency (see, e.g., [7, 36]) and time/power (see, e.g., [28, 29, 33]) domains. While these approaches have demonstrated a significant enhancement for LoRa collision resilience, the state-of-the-art is far from satisfaction in practice. For example, mLoRa [31] and FTrack [36] can rarely decode three collided LoRa packets while maintaining its symbol error rate less than 10%. Since the potential in time and frequency domains has been well exploited, it is natural to explore the spatial domain of a LoRa gateway so that its collision resilience can be further improved by leveraging multi-antenna techniques. In addition, most existing LoRa work focuses on uplink packet detection. The downlink transmission efficiency of LoRaWAN has not been well explored. One may think that, in LoRaWAN, uplink is important while downlink is not. This was true. However, with the continuous expansion of LoRa application landscape, downlink communications have become increasingly important for some LoRa applications such as remote sensor control, massive machine operation, and smart city/building management. When LoRa devices require ACK to confirm the successful delivery of their packets (e.g., TCP connection), the downlink transmission becomes particularly useful. Multiple antennas on a LoRa gateway will make it possible to enable concurrent downlink transmission, which will significantly improve the downlink throughput and find many applications in emerging LoRaWANs.

In this paper, we present MaLoRaGW, the first-of-its-kind multi-antenna LoRa gateway that enables multi-user MIMO (MU-MIMO) LoRa communications in both uplink and downlink. MaLoRaGW was inspired by the success of MU-MIMO in cellular and Wi-Fi networks [4, 23, 40, 44, 48]. It exploits the spatial degrees of freedom (SDoF) provided by its multiple antennas for two purposes: i) enhance packet detection in uplink and ii) enable concurrent packet transmission in downlink. In uplink, MaLoRaGW projects the received multiple signal streams into different spatial subspaces, making it possible to decode a weak collided LoRa packet that cannot be decoded by a LoRa receiver with a single antenna. In downlink, MaLoRaGW performs beamforming to send multiple independent packets to different users, thereby improving downlink throughput and reducing the packet round-trip latency.

Realizing MaLoRaGW in practice is a nontrivial task. In uplink, reaping the multiplexing gain requires the MU-MIMO channel

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 '22, November 6–9, 2022, Boston, MA, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9886-2/22/11...\$15.00

<https://doi.org/10.1145/3560905.3568533>

knowledge to mitigate strong inter-user interference so as to decode weak collided packet [41]. However, estimating channel knowledge requires fine-grained frequency/timing synchronization, which in turn requires channel knowledge for signal projection [1, 43]. The synchronization and channel estimation form a death loop, which must be broken in order to decode a weak collided packet. To address this issue, we perform principal component analysis (PCA) on the received signal streams and find that the principal components (i.e., eigenvectors of the time-domain signals' covariance matrix) perform very well in the separation of weak and strong packets. We thus use the principal components of received signal streams to project them into their subspace. Based on the projected signal streams, channel coefficients are estimated for each user and the corresponding packets are decoded.

In downlink, concurrent transmission relies on beamforming to pre-mitigate inter-user interference so that each LoRa device is capable of decoding its desired packets; and the construction of beamforming filters requires downlink channels between a gateway and its users. The key challenge here is channel acquisition. An approach that is widely used in cellular and Wi-Fi networks is explicit channel feedback, i.e., each user estimates downlink channel and reports it back to base station (or access point). However, LoRa devices do not have the luxury to perform explicit channel feedback. To address this issue, MaLoRaGW adopts implicit channel estimation. That is, MaLoRaGW estimates uplink channels based on its received (collided or uncanceled) packets and performs channel calibration to infer downlink channels based on the estimated uplink channels. This approach is transparent to LoRa users, making MaLoRaGW backward compatible with off-the-shelf LoRa devices.

We have built a prototype of MaLoRaGW and evaluated its performance in realistic scenarios of three different scales: lab, office building, and university campus. It has been validated that MaLoRaGW is backward compatible with commercial off-the-shelf (COTS) LoRa devices. Extensive experiments have been conducted to evaluate the packet error rate (PER) and throughput performance of MaLoRaGW against the state-of-the-art LoRa gateways. Our experimental results show that, in uplink, the two-antenna MaLoRaGW increases the throughput by 10% and reduces the PER by 40%. In downlink, it improves the throughput by 95% while maintaining a similar PER. One may wonder why uplink throughput gain is small. This can be partially attributed to the fact that a single-antenna LoRa device is already capable of decoding collided packets thanks to its CSS modulation. Actually, in uplink, the benefit of MaLoRaGW manifests in the PER reduction (by 40%), which leads to a more reliable LoRa communication.

The contributions of this paper can be summarized as follows.

- MaLoRaGW, to the best of our knowledge, is the first one of studying MU-MIMO for LoRaWANs. It presents a novel LoRa PHY design that enhances concurrent packet detection in uplink and enables concurrent packet transmission in downlink.
- MaLoRaGW introduces new signal processing techniques to enable MU-MIMO for LoRaWAN, including PCA-based synchronization, robust channel estimation, and implicit beamforming. Through a joint uplink and downlink design, MaLoRaGW is backward compatible with COTS LoRa devices.
- MaLoRaGW has been evaluated in three different-scale scenarios with off-the-shelf LoRa devices. Experimental results show that, compared to existing LoRa gateways, a two-antenna MaLoRaGW can slightly improve the throughput in uplink but nearly double the throughput in downlink.

2 RELATED WORK

Recently, many schemes have been proposed to address the packet collision problem in LoRaWANs. Most existing schemes leverage the CSS modulation features in the frequency (e.g., [7, 36]) and time/power (e.g., [11, 28, 29, 31, 33, 39]) domains to decode collided packets. MaLoRaGW is designed based on existing work (e.g., Choir [7], FTrack [36] and PCube [35]) for its concurrent uplink transmission, and it complements existing work by enabling concurrent downlink transmission for the first time.

Frequency Feature Exploitation. In [7], Eletreby et al. proposed Choir to differentiate the collided packets from different users based on their unique frequency offsets. Choir was designed based on the observation that many hardware imperfections (e.g., time, frequency, or phase offsets) contribute to the frequency shifts of CSS chirps. In [36], Xia et al. proposed FTrack to decode collided chirps for LoRa. FTrack applies a sliding demodulation window to the received signals and traces the variations of the detected frequencies. By leveraging the timing misalignment of collided chirps, FTrack filters out the undesired chirps based on their discontinued frequencies. MaLoRaGW borrows the ideas from these works but extends them to the spatial domain. Moreover, MaLoRaGW extends concurrent transmission from uplink to downlink.

Power Feature Exploitation. In [29], Tong et al. proposed a scheme called CoLoRa to enable multi-LoRa packet reception. CoLoRa uses the time offset between collided packets as well as the peak power ratio of the demodulated chirps to differentiate collided chirps. Similarly, NScale in [28] successfully demodulates collided chirps using both normal down-chirp and non-stationary scaled down-chirp symbols. It first pairs the resultant peaks and then calculates their peak scaling factors to differentiate collided packets from different users. In [33], Wang et al. introduced a collision recovery scheme (called OCT) by leveraging the time and power offsets of collided packets. OCT first ranks the power of the detected peaks and then classifies the collided packets based on their peak power. Similar ideas have also been studied in [11, 39].

MIMO Diversity for LoRa Uplink. Thus far, spatial domain has rarely been exploited for LoRa. Pioneering work [35] presents a MIMO-based gateway design (PCube) to decode the collided packets. PCube combines LoRa signal features in time and frequency domains with the measured phase difference over different antennas to decode the collided packets. MaLoRaGW differs PCube in two aspects: i) PCube considers uplink only while MaLoRaGW mainly focuses on downlink; ii) PCube exploits spatial diversity while MaLoRaGW exploits spatial multiplexing.

New Waveform and Detector for LoRa. In [15], Li et al. introduced CurvingLoRa to improve the capacity of LoRaWANs. The key idea is to use a non-linear base chirp for the modulation of LoRa signal. While it can significantly increase the capacity of a LoRa gateway, it is not backward compatible with incumbent LoRa devices and may become more sensitive to timing offset. In [14], Li et al. proposed a demodulation scheme (called NELoRa) for LoRa




Figure 1: The spectrogram of CSS modulation: (a) the base up-chirp symbol, (b) the modulated symbol ‘1010011101’ (669 in decimal), where $f_{sym} = \frac{669}{2^{10}} \times 500 \text{ kHz} \approx 326.7 \text{ kHz}$.

8 Symbols	4.25 Symbols	8 Symbols	L Bytes	2 Bytes
Preamble	Synchronization word	PHDR	PHY payload	CRC (uplink only)

Figure 2: The frame structure for LoRa communications.

devices to improve their receivers’ sensitivity by leveraging deep neural networks. MaLoRaGW is orthogonal and complementary to this research line.

LoRa for Sensing. LoRa has recently shown a great potential to enable long-range sensing applications [38, 46, 47]. In [45], Zhang et al. proposed a LoRa-based sensing model to enable through-the-wall human activity sensing over long (e.g., 25 m) distances. In [37], Xie et al. proposed Sen-fence, a LoRa-based multi-antenna and multi-gateway system, that jointly enhances the sensing range and weakens the impact of undesired movements in the sensing field of view using virtual beamforming techniques. In [2], Chen et al. proposed WideSee, a contactless sensing scheme. It leverages a LoRa transceiver mounted on a drone to enable target human detection and localization over long ranges. MaLoRaGW does not belong to this research area but may provide insights for future LoRa sensing by leveraging multiple antennas.

3 MU-MIMO IN LORAWAN

3.1 A Primer of LoRa

LoRa is a wireless system that offers long-range and low-power wireless connectivity for IoT devices. LoRaWANs are typically configured to a star network topology, where a centralized gateway serves LoRa devices in both uplink and downlink. In North America, LoRa operates in 915 MHz frequency bands, where 64 channels of 125 kHz and 8 channels of 500 kHz are specified for upstream, and 8 channels of 500 kHz are specified for downstream. LoRa devices support data rates ranging from 980 bps to 21.9 kbps, depending on the channel bandwidth and the modulation spreading factor [5].

Chirp Spread Spectrum (CSS) Modulation. LoRa uses CSS technique for data modulation. CSS is a spreading technique that linearly sweeps the entire channel bandwidth within a symbol duration. Denote f_{sym} and f_s as the symbol frequency and the sampling frequency, respectively. Also, denote $N = 2^{SF}$ as the symbol duration, where SF is the spreading factor. LoRa supports spreading factors from 6 to 12, depending on the channel bandwidth and the required data rate. The transmitted symbol using CSS modulation can be written as [7]:

$$x[n] = c[n]e^{j2\pi\frac{f_{sym}}{f_s}n}, \quad 0 \leq n < N, \quad (1)$$




Figure 3: The transmitted LoRa frame.




Figure 4: LoRa medium access protocol.

where $c[n] = e^{j2\pi(\frac{n^2}{2N} + \frac{f_0}{f_s}n)}$ is the base up-chirp signal and f_0 is the initial base up-chirp frequency. Fig. 1 shows the spectrogram of the based chirp signal and a modulated symbol (e.g., ‘1010011101’ = 669) using CSS modulation, when SF = 10, $f_s = 500 \text{ kHz}$, and $f_0 = 0 \text{ kHz}$.

LoRa Packet Structure. Fig. 2 shows the frame structure used in LoRa communications. The frame consists of a preamble, a synchronization word, an optional PHY header and its CRC, and PHY payload. The preamble consists of 8 up-chirp symbols, followed by 2 more up-chirp symbols and 2.25 down-chirp symbols as the synchronization word. The preamble and synchronization word are used for frame extraction and synchronization purposes. The PHY header (PHDR) and PHY header CRC (PHDR_CRC) are optional. When used, they indicate the length of payload, the coding rate, and the presence of CRC for payload. The PHY payload carries MAC header and MAC data. The maximum payload size depends on the data rate, which is determined by the spreading factor and channel bandwidth. The LoRa packet is then modulated by CSS modulation and transmitted over the air, as exemplified in Fig. 3.

LoRaWAN Medium Access. Fig. 4 shows the medium access protocol for LoRa devices in a LoRaWAN. A LoRa device uses ALOHA protocol to access the medium. It can initiate uplink transmissions once it wakes up and has data to send. Upon the completion of uplink transmissions, it turns to sleep for a fixed time interval (~1 s) and opens up two time windows for data reception. The LoRa device receives the downlink information (e.g., ACK packets and gateway commands) within these two time windows and turns to sleep for the rest of the time [9].

Concurrent Transmission. While LoRaWAN already supports concurrent packet transmission in uplink [7, 22, 28, 36], it cannot support concurrent transmission in downlink. This is because the packets from a gateway will have the same features in time and frequency domains. If a gateway adds two packets together and sends them to two users, then the two users will decode the same packet (the one with a stronger power) rather than their own packets.




Figure 5: MU-MIMO transmission: (a) uplink; (b) downlink.

3.2 System Model of MU-MIMO

MU-MIMO is a key technology for OFDM-based wireless networks such as cellular [23, 30] and Wi-Fi [12, 21] networks. It has been widely deployed in real-world wireless communication systems and demonstrated a significant gain of spectrum efficiency. MU-MIMO exploits the SDoF provided by multiple antennas to separate signal streams, making it possible for a gateway (a.k.a., access point or base station) to support multiple concurrent data packet transmissions in both uplink and downlink.

Uplink Transmission. Fig. 5(a) shows uplink MU-MIMO transmission. Denote K as the number of user devices and M as the number of antennas at the gateway. Denote $\mathbf{H}_{ul} \in \mathbb{C}^{M \times K}$ as the compound uplink channel matrix. Mathematically, the received signal at the gateway, $\mathbf{y}_{gw} \in \mathbb{C}^{M \times 1}$, can be written as:

$$\mathbf{y}_{gw} = \mathbf{H}_{ul}\mathbf{x}_{ul} + \mathbf{n}, \quad (2)$$

where $\mathbf{x}_{ul} \in \mathbb{C}^{K \times 1}$ is the vector of transmitted signals from user devices, and $\mathbf{n} \in \mathbb{C}^{M \times 1}$ is the additive noise vector. To decode \mathbf{x}_{ul} , the gateway first estimates channel matrix \mathbf{H}_{ul} and then uses the estimated channel to design an equalizer for signal detection.

Downlink Transmissions: Fig. 5(b) shows downlink MU-MIMO transmission. Denote $\mathbf{H}_{dl} \in \mathbb{C}^{K \times M}$ as the compound downlink channel matrix. Denote $\mathbf{P} \in \mathbb{C}^{M \times K}$ as the beamforming matrix used by the gateway. The vector of signals received by all user device can be written as:

$$\mathbf{y}_{ud} = \mathbf{H}_{dl}\mathbf{P}\mathbf{x}_{dl} + \mathbf{n}, \quad (3)$$

where $\mathbf{x}_{dl} \in \mathbb{C}^{K \times 1}$ is the transmitted symbol vector at the gateway for K user devices, and \mathbf{n} is noise vector at user devices. In cellular and Wi-Fi networks, protocols have been specified for a gateway to obtain \mathbf{H}_{dl} . With the downlink channel knowledge, precoders such as ZF (zero-forcing) and MMSE (minimum mean square error) can be used to separate data streams for different users in the spatial domain. Particularly, the ZF precoder is constructed by letting $\mathbf{H}_{dl}\mathbf{P} = \mathbf{I}$, where \mathbf{I} is the identity matrix. Through the transmitter-side precoding, inter-user inference is pre-cancelled and each user device can decode its desired signal.

3.3 Challenges and Our Approach

MU-MIMO has received a great success in cellular and Wi-Fi networks, and it is one of key technologies for 5G/6G networks. However, while there are some prior works on studying MIMO for LoRaWAN, little research work has been done for the design of MU-MIMO schemes in LoRaWAN. This stagnation underscores the

grand challenges in the realization of MU-MIMO for LoRaWAN, which we describe as follows.


Uncoordinated Transmission in Uplink. In cellular and Wi-Fi (e.g., 802.11ax) systems, dedicated MU-MIMO protocols have been specified in the standards to coordinate user devices for concurrent uplink transmission. User devices' packets are aligned in both time and frequency when arriving at the gateway. The time and frequency alignments make it easy for a gateway to detect the collided packets. In addition, user devices in these systems carry orthogonal pilots (e.g., VHT in 802.11 [19] and demodulation reference signals in LTE [6]) in their packets. The orthogonal pilots allow a gateway to estimate the channels between itself and user devices, which play a critical role in the detection of concurrent packets. Unfortunately, LoRaWAN does not have such luxuries for MU-MIMO transmissions. It uses ALOHA protocol for medium access control, which does not support coordination among user devices [24]. The collided packets at a gateway could be fully or partially overlapping, and they may bear different carrier frequency offsets and chirp timing offsets.

Channel Acquisition in Downlink. In cellular and Wi-Fi systems, protocols have been specified for channel feedback. For example, each user involved in downlink MU-MIMO transmission will estimate its downlink channel and report the estimated channel (after quantization and compression) to an access point, which then constructs precoding vectors for downlink beamforming [20]. Such a channel sounding protocol requires cooperation from user devices. However, LoRa devices are typically of low cost, low power, and low computation. It is impractical to require LoRa devices to estimate and report their channels. In addition, one of our design objectives is to maintain backward compatibility with incumbent LoRa devices. Requiring channel feedback clearly contradicts our design objective.

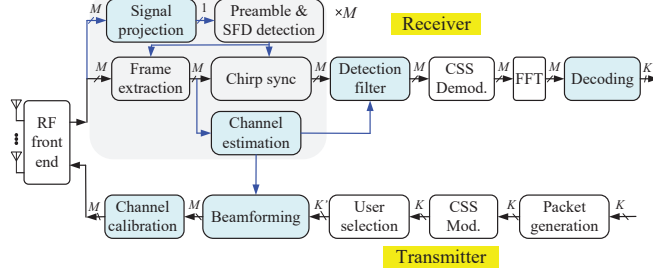
Our Approach. To address the above challenges, MaLoRaGW employs a joint design for uplink packet detection and downlink beamforming. The joint design features PCA-based signal projection, robust channel estimation, and implicit beamforming. Different from OFDM systems, LoRa has collision recovery capability even if a gateway has a single antenna [11, 22, 28, 36, 39]. Therefore, MaLoRaGW aims to enhance the collision recovery capability of a LoRa gateway by jointly exploiting signal features in spatial, frequency, and time domains. We note that, since LoRa is not a single-tone frequency channel as OFDM systems, completely mitigating inter-user interference may not be possible. However, it is still possible to support multiple packet transmission in practice, thanks to the interference resilience of CSS modulation.

4 DESIGN

MaLoRaGW is an M -antenna LoRa gateway to enhance its packet detection in uplink and enable its concurrent transmission in downlink. Fig. 6(a) shows its MAC protocol, and Fig. 6(b) shows its PHY-layer signal processing diagram. In Fig. 6(b), the colored modules are our new design, and the rest are LoRa's legacy modules. The Rx chain is to decode collided packets from multiple user devices, while the Tx chain is to deliver data packets to the same set or a subset of those user devices. MaLoRaGW can decode more than M collided packets in uplink by jointly leveraging the signal features



(a) Proposed protocol for MaLoRaGW's MU-MIMO transmission.



(b) Proposed PHY design for MaLoRaGW's MU-MIMO transmission.

Figure 6: Joint protocol and PHY design for MaLoRaGW to support both uplink and downlink MU-MIMO transmissions.

in spatial, frequency, and time domains. However, MaLoRaGW can deliver at most M concurrent packets in downlink, as the downlink transmission relies solely on the spatial degrees of freedom to separate signals.

4.1 Overview

A key component of our design is channel estimation. Prior works on LoRa collision recovery [7, 11, 29, 33, 36, 39] exploit CSS signal features in frequency, time, and/or power domains to decode packets, and thus do not require to have channel state information (CSI). In contrast, CSI plays a key role in MU-MIMO transmissions. In uplink, CSI is needed to separate strong and weak signals from their collided packets, so that weak signal can be decoded by leveraging its spatial features. In downlink, CSI is needed to construct precoding vectors for beamforming, which is the enabler of downlink concurrent packet transmission.

To estimate the CSI, some system imperfections must be corrected first. Like other wireless communication systems, LoRa receivers suffer from carrier frequency offset (CFO) and chirp timing offset (CTO) in their channel estimation. The frequency and timing offsets must be corrected prior to channel estimation. To do so, a PCA-based signal projection method is proposed for frequency and timing offset estimation. The projection is made in the spatial domain. It is effective to separate weak and strong signals from collided packets, making it possible to correct CFO/CTO for both weak and strong packets and therefore accurately estimate their channels. The estimated channels are then used for two purposes: i) enhance uplink packet detection and ii) construct precoding vectors for downlink beamforming.

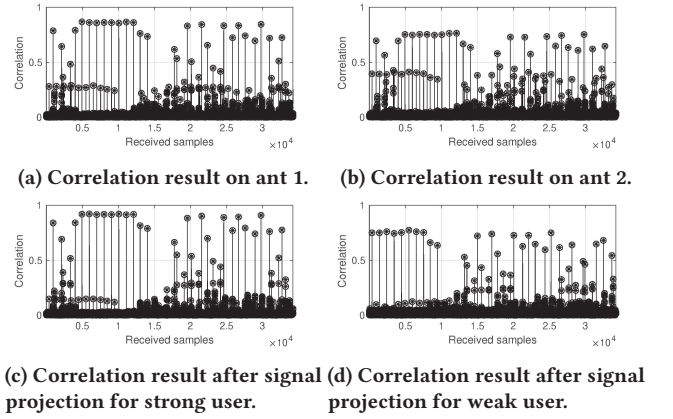


Figure 7: Correlation results before and after projection.

4.2 Preamble Detection

LoRa receivers use the cross correlation between a base up-chirp and the received signal to detect the preamble of a packet. Although cross correlation is resilient to interference, it does not work well when the interference is much stronger than the signal of interest. Consider a LoRa gateway equipped with two antennas. It receives two collided packets, which have 20 dB difference in their signal strengths. Fig. 7(a-b) shows the cross-correlation results of the received signals at the gateway's antennas. It can be seen that the cross correlation has a poor performance for the weak user as it is significantly corrupted by the strong interference. The poor performance of preamble detection always leads to a failure of packet detection. In what follows, we present our treatment to this problem.

Basic Idea. When a gateway has multiple antennas, the spatial degrees of freedom can be leveraged to improve the performance of preamble detection. Consider a gateway equipped with two antennas as shown in Fig. 8(a). It receives packets from two collided packets from two users. One user is close to gateway, while the other is far from gateway. To detect the weak signal from the distant user, a natural approach is to check the signal stream from each antenna and find the best one to perform cross correlation. However, this approach does not work well. This is because the antennas are close to each other and tend to have similar SNR, as illustrated in Fig. 8(b). To address this issue, we propose a spatial projection approach by leveraging the multiple antennas on a gateway. Fig. 8(c) illustrated the basic idea of our approach. If we can find a good spatial vector (G in Fig. 8(c)), the projected signal on G will have similar strength as the projected interference. The question to ask is how to find a good projection vector G . The challenge here is that the gateway does not have channel knowledge yet, and it must find a vector without channel knowledge.

PCA-based Spatial Projection. MaLoRaGW employs PCA to construct the projection vectors. While PCA is widely used to reduce the dimension of data, here it is used to separate strong and weak signals so that weak signal can be detected. It works as follows. MaLoRaGW first computes the covariance matrix of the received signal streams by $R_{yy} = \sum_i y_{gw}(i)y_{gw}(i)^H \in \mathbb{C}^{M \times M}$, where i is signal sample index. Then, it performs $[Q, \Lambda, Q^{-1}] = \text{svd}(R_{yy})$,




Figure 8: Illustration of spatial projection for collided packets. (a) A case of collided packet; (b) signal strengths on two antennas; (c) signal projection.

where $\text{svd}(\cdot)$ is the singular value decomposition (SVD) operator, and $\Lambda \in \mathbb{R}^{M \times M}$ is a diagonal matrix carrying the singular values of \mathbf{R}_{yy} in non-increasing order. The columns in \mathbf{Q} represent the singular vector of \mathbf{R}_{yy} . If the number of gateway's antennas is greater than that of collided packets, \mathbf{Q} can be further divided to a signaling subspace $\mathbf{Q}_s \in \mathbb{C}^{M \times K}$ and a null subspace $\mathbf{Q}_n \in \mathbb{C}^{M \times (M-K)}$, i.e., $\mathbf{Q} = [\mathbf{Q}_s | \mathbf{Q}_n]$. The principal components of received signals can be calculated as: $\mathbf{y}_p = \mathbf{Q}_s^H \mathbf{y}_{gw}$, where \mathbf{y}_p is projected signals.

The signal space includes the basis corresponding to non-zero eigenvalues, which is highly correlated with \mathbf{H}_{UL} . We use the signal space to project the received signal into the signals' direction to strengthen the desired signal SNR and reduce the interference signal. Mathematically, the projected signal for each user can be expressed as: $\mathbf{y}_p = \mathbf{q}_s^H \mathbf{y}_{gw}$, where \mathbf{q}_s is a column of \mathbf{Q}_s and it is the signal basis corresponding to the user.

To see the effectiveness of the projection, let us reconsider the example presented in Fig. 7(a-b). We compute SVD on the received signals and then perform cross correlation on the primary components of the received signals. Fig. 7(c-d) shows the cross correlation results. It can be seen that the correlation peak is significant, indicating that the weak signal can be successfully detected.

Preamble Search. After signal projection, MaLoRaGW performs cross correlation between the base up-chirp symbol and the primary components of received signals (i.e., \mathbf{y}_p). It searches for the beginning of the preamble by solving the following problem:

$$\hat{\zeta} = \arg \max_{\zeta} \sum_{l=1}^L \left| \sum_{n=0}^{N-1} y_p[n + \zeta + lN] c^{-1}[n] \right|, \quad (4)$$

where $L = 8$ is the number of base up-chirps in LoRa preamble, and $c[n]$ is the base up-chirp signal.

4.3 Channel Estimation

Consider a point-to-point LoRa communication link, where both transmitter and receiver are equipped with a single antenna. Denote h as the channel between the transmitter and the receiver. Assume that LoRa signal experiences a flat fading channel when traversing from the transmitter to the receiver. This is true in most cases as the LoRa signal has a narrow bandwidth (i.e., ≤ 500 kHz) [7]. If the transmitter and receiver are perfectly synchronized in time and frequency and noise is negligible, then the received chirp can be

written as:

$$y[n] = h \cdot c[n] e^{j2\pi \frac{f_{sym}}{f_s} n}, \quad \text{for } 0 \leq n < N, \quad (5)$$

where f_{sym} and f_s are the symbol frequency and sampling frequency, respectively.

By demodulating the received chirp and applying FFT to the demodulated signals, the channel can be estimated as:

$$\hat{h} = h[m = \frac{f_{sym}}{f_s}] = \frac{1}{N} \sum_{n=0}^{N-1} c^{-1}[n] y[n] e^{-j2\pi \frac{mn}{N}}. \quad (6)$$

However, estimating channel in LoRa is not that simple. Like other communication systems, LoRa receivers suffer from CFO and CTO, which must be estimated and corrected in order to estimate channel coefficients. In what follows, we focus on the estimation of CFO and CTO of chirps.

4.3.1 Carrier Frequency Offset (CFO). CFO is mainly caused by the frequency mismatch between transmitter and receiver. Denote Δf as the CFO between transmitter and receiver. Mathematically, the l th received chirp at the gateway in the presence of CFO can be expressed as:

$$y_l[n] = h e^{j2\pi \frac{\Delta f}{f_s} lN} c[n] e^{j2\pi \frac{f_{sym} + \Delta f}{f_s} n}, \quad (7)$$

for $0 \leq n < N$ and $0 \leq l < L$, where L here denotes the total number of chirps in one frame. Then, the channel estimation in (6) can be rewritten as:

$$\hat{h} = h[m = \frac{f_{sym} + \Delta f}{f_s}] = \frac{1}{N} e^{-j2\pi \frac{\Delta f}{f_s} lN} \sum_{n=0}^{N-1} c^{-1}[n] y_l[n] e^{-j2\pi \frac{mn}{N}}, \quad (8)$$

for $0 \leq l < L$. From (8), it can be seen that the CFO induces a phase shift to the received chirps. This phase shift changes linearly over the consecutive received chirps. Fig. 9(a) shows the measured phase of the decoded preamble chirps when the CFO between transmitter and receiver is 0kHz, 5kHz, and 10kHz. It can be seen that the unwrapped phase indeed increases linearly with the chirp index. Based on (8), we present a two-step approach to estimate the CFO.

Coarse CFO Correction. We take advantage of up-chirp and down-chirp in LoRa frame to estimate the CFO. We first measure the frequencies of preamble up-chirp symbols and synchronization down-chirp symbols in the received signal, and then use the




Figure 9: The measured phases of the decoded preamble chirps in practice.

measured symbol frequencies to infer the CFO. Nevertheless, the inference is not straightforward because the measured symbol frequencies may be caused by CFO, CTO, or both. Fortunately, the CFO and CTO have different impacts on the measured symbol frequencies. Specifically, a positive CFO will result in a positive frequency shift on both up-chirp and down-chirp symbols, while a positive CTO will cause a positive frequency shift on up-chirp symbols and a negative frequency shift on down-chirp symbols. This difference makes it possible to infer the CFO based on the measured (up-chirp and down-chirp) symbol frequencies. Based on this observation, we estimate the CFO as follows:

$$\Delta \hat{f}_{coarse} = \frac{1}{2} \left(\frac{1}{8} \sum_{l=1}^8 z_l^{up} + \frac{1}{2} \sum_{l=1}^2 z_l^{dw} \right), \quad (9)$$

where z_l^{up} is the measured frequency of the l th up-chirp symbol in preamble, and z_l^{dw} is the measured frequency of the l th down-chirp symbol in synchronization word. Once $\Delta \hat{f}_{coarse}$ is calculated, the estimated channel is updated as follows: $\hat{h}_l \leftarrow \hat{h}_l \cdot e^{j2\pi \frac{\Delta \hat{f}_{coarse}}{f_s} lN}$ for $0 \leq l < L$.

Fine CFO Correction. The CFO estimated in (9) is not accurate enough. Although the residual CFO is small, it leads to a phase shift accumulated over preamble symbols. This phase shift will degrade uplink packet detection and downlink beamforming. To estimate the residual CFO, we search a narrow frequency offset range (e.g., -10 Hz to 10 Hz) and find the frequency offset, $\Delta \hat{f}_{fine}$, that minimizes the phase variance of estimated channels over preamble symbols. Denote $\phi_l = \angle(\hat{h}_l \cdot e^{j2\pi \frac{\Delta \hat{f}}{f_s} lN})$ as the phase of the estimated channel coefficient for the l th preamble up-chirp symbol. Then, the residual CFO is calculated by:

$$\Delta \hat{f}_{fine} = \arg \min_{\Delta f} \sigma_{\Phi}^2, \quad (10)$$

where σ_{Φ}^2 is the variance of the phase vector $\Phi = [\phi_1, \phi_2, \dots, \phi_L]$.

4.3.2 Chirp Timing Offset (CTO). CTO refers to the timing misalignment between the received chirps and the applied demodulation (dechirp) window. Denote δ as CTO. Then, the demodulated chirp can be written as:

$$c^{-1}[n - \delta]y[n] = h \cdot e^{-j2\pi\phi} e^{j2\pi \left(\frac{f_s y m}{f_s} + \frac{\delta}{N} \right) n}, \quad (11)$$




Figure 10: Illustrating the decoding algorithm for MaLoRaGW in uplink.

for $\delta \leq n < N + \delta$, where $\phi = \frac{\delta^2}{2N} - \frac{f_0}{f_s} \delta$. Then, the estimated channel (6) can be re-written as:

$$\hat{h} = h[m = \frac{f_s y m}{f_s} + \frac{\delta}{N}] = \frac{1}{N} e^{j2\pi\phi} \sum_{n=0}^{N-1} c^{-1}[n - \delta]y[n] e^{-j2\pi \frac{mn}{N}}. \quad (12)$$

From (12), it can be seen that the CTO introduces an additional phase shift ϕ to the observed channel. This phase shift is constant for all chirps in preamble and can be inferred based on δ . Fig. 9 shows the measured phase of the demodulated preamble chirps, when $\delta = 0$, $\delta = 10$, and $\delta = 50$. Fortunately, as explained before, the CTO can be estimated by:

$$\hat{\delta} = \frac{N}{2f_s} \cdot \left(\frac{1}{8} \sum_{l=1}^8 z_l^{up} - \frac{1}{2} \sum_{l=1}^2 z_l^{dw} \right). \quad (13)$$

Based on the estimated CTO, the channel phase shift caused by CTO can be calculated by: $\phi = \frac{\hat{\delta}^2}{2N} - \frac{f_0}{f_s} \hat{\delta}$, which is then used to correct the estimated channel.

4.3.3 Channel Estimation for Collided Packets. While the above channel estimation was presented for a LoRa receiver when decoding an interference-free (collision-free) packet, it can be extended to the case where packets collide. When the desired preamble symbols are interfered by the symbols of other packets, MaLoRaGW traces the frequency of decoded symbols within the first 10 demodulation windows and identifies the peaks associated with the desired preamble symbols. To do so, MaLoRaGW first uses a low pass filter to mitigate the frequency components of undesired symbols. The rationale behind this operation is that, after frame synchronization, the frequencies of the desired preamble symbols are around zero. Yet, the output of low-pass filter may still include the low frequency components of interfering chirps. MaLoRaGW then uses a maximum likelihood method to search for the frequency peaks with the highest correlation value. This operation is based on the fact that the frequency shift caused by CFO and CTO is almost identical for the same user but different for different users. After identifying the frequencies of the desired preamble symbols, the above channel estimation method is applied to estimate the channel between the gateway's each antenna and the user device.

4.4 Decoding Algorithm

Our decoding algorithm comprises the following steps: spatial signal projection, chirp demodulation and FFT operation, and peak clustering. Fig. 10 shows an example of our decoding diagram. In what follows, we explain them in detail.

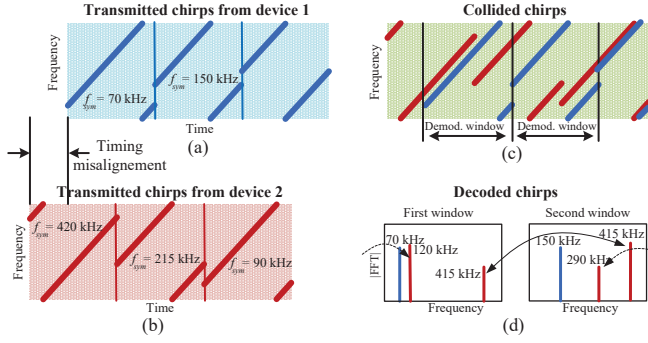


Figure 11: Illustration of demodulating collided chirps that are misaligned in time.

Spatial Signal Projection. The purpose of spatial signal projection is to alleviate the near-far effect so as to improve the signal detection accuracy. Here, we use ZF as the projection method. Denote $K' = \min(K, M)$, where K is the number of LoRa devices, and M is the number of gateway's antennas. We select K' out of K LoRa devices based on the strength of their estimated channel coefficients. Denote $\mathbf{H}_{ul} \in \mathbb{C}^{M \times K'}$ as the estimated uplink channel matrix for those selected K' LoRa devices. Then, we construct the ZF projection matrix by letting $\mathbf{U} = (\mathbf{H}_{ul}^H \mathbf{H}_{ul})^{-1} \mathbf{H}_{ul}^H$, where $(\cdot)^H$ is conjugate transpose operator, and apply it to the signal streams.

The ZF projection will reduce inter-user interference. Consider Fig. 10 for example. The gateway receives strong signals from A and B and weak signals from C and D. It constructs the ZF matrix using A's and B's channel coefficients. This ZF matrix will mitigate the signal from A for stream 1 and mitigate the signal from B for stream 2, making it easier to decode weak signals from C and D.

FFT Peak Clustering. After chirp demodulation and FFT operations, the FFT outputs will have multiple amplitude peaks in the demodulation window. These peaks may correspond to different LoRa devices and their different chirps (due to the lack of inter-user synchronization). It is critical to identify each FFT peak and find out its corresponding source device and corresponding chirp. To do so, we jointly exploit two features of each FFT peak for decoding collided packets: *Time and frequency offsets* and *channel ratio*.

i) *Time and frequency offsets:* All chirps from the same LoRa device have a unique fractional frequency offset [7], which can be used to identify an FFT peak's source device. For each of the output streams, MaLoRaGW exploits the frequency and time features of each FFT peak to disentangle collided packets. Specifically, MaLoRaGW first identifies the FFT peaks in consecutive demodulation windows and then removes those peaks with the same frequency in two consecutive demodulation windows. The rationale behind this operation is twofold. First, the chirps from different LoRa devices are very likely misaligned in time. Second, when an interfering chirp is misaligned with the demodulation window, its energy is proportionally split over two consecutive demodulation windows; however, its frequencies will be the same in the two demodulation windows. Fig. 11 illustrates the use of time and frequency features of CSS chirps to disentangle collided packets. In this example, the chirp signals from the two LoRa devices are misaligned in time. When the demodulation window is aligned with the chirps from the

first LoRa device, it misaligns with the chirps from the second LoRa device. A misaligned chirp will generate two peaks of the same frequency in two consecutive demodulation windows. Based on this observation, MaLoRaGW searches for the peaks of the same frequency in consecutive demodulation windows and removes those peaks before decoding the aligned chirps.

ii) *Channel ratio:* In addition to the fractional frequency offset, MaLoRaGW also uses the channel ratio as another feature for the classification and clustering of FFT peaks. Denote $[x_1, x_2, \dots, x_M]$ as an FFT peak observed from the M signal streams at the gateway, where M is the number of its antennas. Then, $[1, \frac{x_2}{x_1}, \dots, \frac{x_M}{x_1}]$ represents the vector of channel ratios over the gateway's M antennas, which is relatively independent of the data bit, chirp misalignment, CFO, CTO, etc. Therefore, MaLoRaGW uses this vector as another feature to classify the FFT peaks. We note that this feature is similar to that in [35], where phase difference is used for peak classification instead of channel ratio.

Decoding Collided Packets. Once the FFT peaks are classified and clustered, it is straightforward for MaLoRaGW to decode the collided packets.

4.5 Downlink Beamforming

In this subsection, we present our approach for downlink beamforming. It comprises three steps: i) select a subset of LoRa devices for downlink MU-MIMO transmission; ii) infer downlink channels between the gateway and the selected LoRa devices; iii) precode baseband signals at the gateway. In what follows, we present them in detail.

User Selection. While an M -antenna gateway can decode more than M collided uplink packets, it can send at most M concurrent packets to user devices in downlink. If more than M user devices are involved in the uplink transmission, how to select a subset of user devices for downlink MU-MIMO transmission is an open problem. This problem is not in the scope of this work and will be studied in future. Here, MaLoRaGW simply selects the $K' = \min(K, M)$ LoRa devices with strongest uplink signal strength for downlink MU-MIMO transmission.

Downlink Channel Inference. Another question we need to address is the difference between uplink and downlink channels. What we have is uplink channel (i.e., \mathbf{H}_{ul}); what we need is downlink channel (i.e., \mathbf{H}_{dl}). In what follows, we use a small case where $K' = 2$ and $M = 2$ to illustrate our approach for the inference of downlink channel. But our approach is generic and can apply to a general case where $K' \leq M$.

Referring to Fig. 12(a), at a wireless receiver, the observed/estimated channel coefficient comprises transmitter's RF response, over-the-air response, and receiver's RF response. At the gateway, the observed/estimated uplink channel \mathbf{H}_{ul} can be written as:

$$\mathbf{H}_{ul} = \begin{bmatrix} R_1^{\text{grx}} R_{11}^{\text{ota}} R_{11}^{\text{ltx}} & R_1^{\text{grx}} R_{12}^{\text{ota}} R_{21}^{\text{ltx}} \\ R_2^{\text{grx}} R_{21}^{\text{ota}} R_{11}^{\text{ltx}} & R_2^{\text{grx}} R_{22}^{\text{ota}} R_{22}^{\text{ltx}} \end{bmatrix} \\ = \begin{bmatrix} R_1^{\text{grx}} & 0 \\ 0 & R_2^{\text{grx}} \end{bmatrix} \begin{bmatrix} R_{11}^{\text{ota}} & R_{12}^{\text{ota}} \\ R_{21}^{\text{ota}} & R_{22}^{\text{ota}} \end{bmatrix} \begin{bmatrix} R_{11}^{\text{ltx}} & 0 \\ 0 & R_{22}^{\text{ltx}} \end{bmatrix}, \quad (14)$$

where R denotes response and its superscript/subscript denotes the corresponding device. Specifically, 'g' and 'l' in its superscript




Figure 12: (a) Illustrating the relation between uplink and downlink channel in a two-user MIMO case. (b) Illustrating RF calibration at gateway.

represent gateway and LoRa device, respectively; ‘tx’ and ‘rx’ represent transmitter and receiver, respectively; numbers in its subscript represent either gateway’s antenna index or LoRa device index. Superscript ‘ota’ represents the corresponding over-the-air channel response.

While the compound uplink and downlink channels are not reciprocal, their over-the-air components are identical in the coherence time. Therefore, we have

$$\begin{aligned} \mathbf{H}_{dl} &= \begin{bmatrix} R_1^{lrx} R_{11}^{ota} P_1^{gtx} & R_1^{lrx} R_{21}^{ota} P_2^{gtx} \\ R_2^{lrx} R_{12}^{ota} P_1^{gtx} & R_2^{lrx} R_{22}^{ota} P_2^{gtx} \end{bmatrix} \\ &= \begin{bmatrix} R_1^{lrx} & 0 \\ 0 & R_2^{lrx} \end{bmatrix} \begin{bmatrix} R_{11}^{ota} & R_{12}^{ota} \\ R_{21}^{ota} & R_{22}^{ota} \end{bmatrix}^T \begin{bmatrix} R_1^{gtx} & 0 \\ 0 & R_2^{gtx} \end{bmatrix}. \end{aligned} \quad (15)$$

Based (14) and (15), we have:

$$\begin{aligned} \mathbf{H}_{dl} &= \begin{bmatrix} R_1^{lrx} & 0 \\ 0 & R_2^{lrx} \end{bmatrix} \mathbf{H}_{ul}^T \begin{bmatrix} R_1^{gtx} & 0 \\ 0 & R_2^{gtx} \end{bmatrix} \\ &= \underbrace{\begin{bmatrix} R_1^{lrx} R_1^{gtx} & 0 \\ 0 & R_2^{lrx} R_2^{gtx} \end{bmatrix}}_{(a)} \underbrace{\begin{bmatrix} 1 & 0 \\ 0 & \frac{R_2^{lrx} R_1^{gtx}}{R_1^{lrx} R_2^{gtx}} \end{bmatrix}}_{(b)} \mathbf{H}_{ul}^T \underbrace{\begin{bmatrix} 1 & 0 \\ 0 & \frac{R_1^{gtx} R_2^{gtx}}{R_1^{lrx} R_2^{gtx}} \end{bmatrix}}_{(c)}. \end{aligned} \quad (16)$$

In (16), part (a) is a complex scale, which changes the signal strength for both LoRa devices. Part (b) is a diagonal matrix, which changes the signal strength difference between two devices. Therefore, both parts (a) and (b) do not change the signal beamforming directions and thus can be ignored in the beamforming process. Now, the question is how to find unknowns in part (c). To address this question, we propose an RF calibration scheme for the gateway, as shown in Fig. 12(b). The gateway first uses its first antenna to send a signal to its second antenna; denote the observed channel as $h_{1 \rightarrow 2}$. Then, it uses its second antenna to send a signal to its first antenna; denote the observed channel as $h_{2 \rightarrow 1}$. Then, we have

$$\frac{h_{2 \rightarrow 1}}{h_{1 \rightarrow 2}} = \frac{R_1^{grx} R_{11}^{ota} P_2^{gtx}}{R_2^{grx} R_{12}^{ota} P_1^{gtx}} = \frac{R_1^{grx} R_2^{gtx}}{R_1^{lrx} R_2^{gtx}}. \text{ Therefore, we have}$$

$$\tilde{\mathbf{H}}_{dl} \triangleq \mathbf{H}_{ul}^T \begin{bmatrix} 1 & 0 \\ 0 & \frac{R_1^{grx} R_2^{gtx}}{R_1^{lrx} R_2^{gtx}} \end{bmatrix} = \mathbf{H}_{ul}^T \begin{bmatrix} 1 & 0 \\ 0 & \frac{h_{2 \rightarrow 1}}{h_{1 \rightarrow 2}} \end{bmatrix}, \quad (17)$$

where \triangleq indicates the ignorance of parts (a) and (b) in (16).

Eq. (17) shows our approach for gateway’s RF calibration. It can be easily extended to a generic case. We have two remarks based on our experiments. First, the RF calibration can be done by the




Figure 13: Downlink beamforming operation for downlink MU-MIMO transmission.




Figure 14: (a) LoRaWAN gateway; (b) LoRa user devices.

gateway in a standalone mode. Second, $\frac{R_1^{grx} R_2^{gtx}}{R_1^{lrx} R_2^{gtx}}$ is stable over time and only need to calibrate at a low frequency.

Precoding for Downlink MU-MIMO Transmission. The purpose of precoding is to separate the signal streams in the over-the-air channel so that each LoRa device only receives its intended signal. Fig. 13 illustrates the precoding operation at the gateway in a toy-sized network. Recall that $\mathbf{H}_{dl} \in \mathbb{C}^{K' \times M}$ is the real downlink channel and $\tilde{\mathbf{H}}_{dl} \in \mathbb{C}^{K' \times M}$ is the inferred downlink channel. Then, the zero-forcing precoder can be computed as follows:

$$\mathbf{P} = (\tilde{\mathbf{H}}_{dl}^H \tilde{\mathbf{H}}_{dl})^{-1} \tilde{\mathbf{H}}_{dl}^H. \quad (18)$$

Denote $\mathbf{s} = [s_1, s_2, \dots, s_{K'}]^T$ as the baseband signals of data packets that the gateway wants to deliver to the K' LoRa devices. Then, the precoding operation can be expressed as $\mathbf{P}\mathbf{s}$. It can be verified that $\mathbf{H}_{dl}\mathbf{P}$ is a diagonal matrix. This means that \mathbf{P} can pre-cancel inter-user interference in downlink MU-MIMO transmission.

5 EXPERIMENTAL EVALUATION

In this section, we build a prototype of two-antenna MaLoRaGW and evaluate its performance with off-the-shelf LoRa devices in realistic wireless environments.

5.1 Implementation

Two-Antenna MaLoRaGW. We implement MaLoRaGW on a USRP X310, which has two antennas for transmission and reception. Fig. 14(a) shows the hardware for MaLoRaGW. We implemented the proposed PHY-layer design and MAC-layer design (see Fig. 6) on a laptop in C++ using GNU Radio OOT modules. The transmit power of USRP X310’s each RF channel is set to 16 dBm, and the carrier frequency is set to 900 MHz. The spreading factor is set to 10. Three channel bandwidths, 125 kHz, 250 kHz, and 500 kHz, are used to study their impacts on uplink detection and downlink beamforming. The sampling rate is 1 MSps.

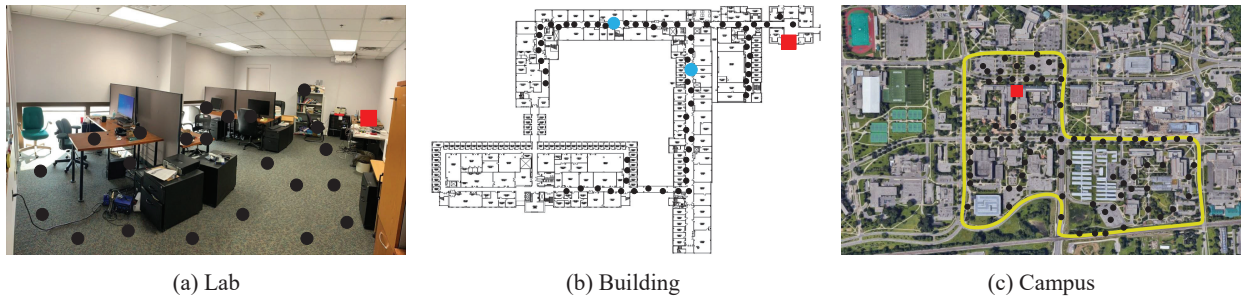


Figure 15: Experimental evaluation scenarios.

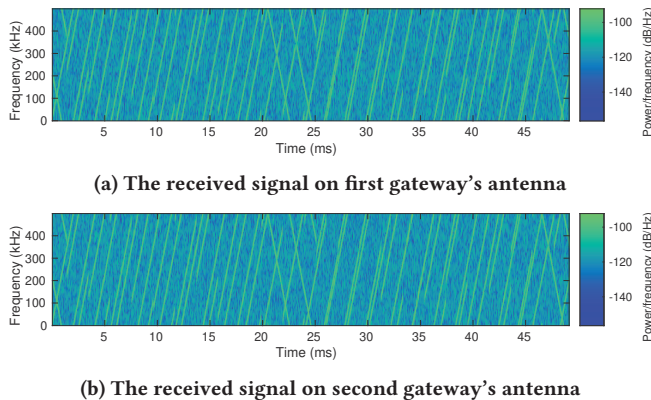


Figure 16: Spectrogram of received signals at MaLoRaGW when two LoRa devices concurrently transmit their packets.

LoRa Devices. We use HELTEC ESP32 wireless modules (developed upon LoRa SX1276 chipset) shown in Fig. 14(b) as LoRa user devices. HELTEC ESP32 has an omni-directional antenna with 3 dBi gain. The receiver sensitivity of the board is -140 dBm, and its maximum transmit power is 20 dBm. We use open-source Arduino software to drive the modules and set the communication parameters. Particularly, we configure the modules to continuously transmit their uplink packets. The modules send their packets independently without timing synchronization. We assign a unique word to each module. The word has 30 symbols, which are modulated by CSS with 10 SF. In downlink, the gateway concurrently transmits two independent packet streams to two selected LoRa modules. Each LoRa module measures its RSSI, counts the number of packets that are matched with its unique word, calculates the packet error rate (PER), and prints out the results on screen as shown in Fig. 14(b).

5.2 Experimental Setup, Metrics, and Baselines

We evaluated the performance of MaLoRaGW in three different scenarios: lab, office building, and university campus, as shown in Fig. 15. The gateway is placed at the spot marked by a red square in each scenario, while user devices are at some spots marked with solid dots. Particularly, Fig. 15(c) marks out the boundary of our test area using a yellow curve.

The following metrics will be used to evaluate the performance of MaLoRaGW: Received signal strength indicator (RSSI) in dBm,

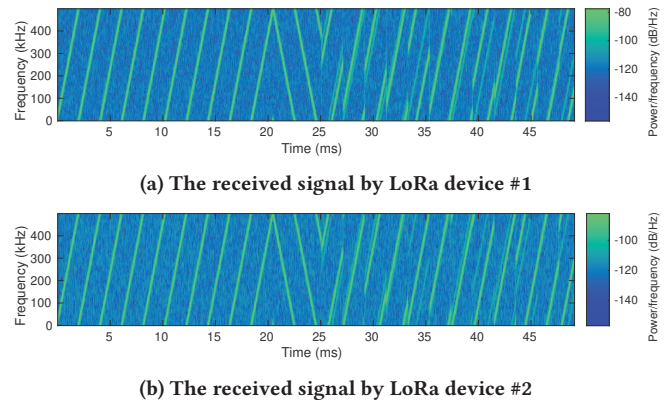


Figure 17: Spectrogram of received signals at two LoRa user devices when concurrently served by MaLoRaGW in down-link.

PER, and network throughput in kbps. Four existing LoRa gateway designs will be used as a comparison baseline to evaluate the performance of MaLoRaGW's uplink: conventional LoRa gateway, Choir [7], FTrack [36], and PCube [35]. We note that no prior work has considered concurrent downlink transmission. So we use a conventional LoRa gateway as the comparison baseline for MaLoRaGW's downlink transmission.

5.3 A Case Study

To understand how MaLoRaGW works, we consider a case study by placing MaLoRaGW and two LoRa devices in the building scenario shown in Fig. 15(b). The two LoRa devices are placed at the spots marked by blue circles. Only for this case, we use two USRP N210 devices as two LoRa devices. This is because off-the-shelf LoRa devices are a closed system and do not provide detailed information. Using N210 allows us to examine the received signals and other PHY-layer parameters.

Uplink Performance. In the uplink, the two LoRa devices concurrently transmit their packets to MaLoRaGW. Fig. 16 plots the spectrogram of the collided signals at MaLoRaGW's two antennas. It can be clearly seen that two packet collision occurs. MaLoRaGW performs PCA-based synchronization, channel estimation, and packet decoding. The estimated CFO values are 2734 Hz and -3284 Hz for two LoRa users. The measured CTO values after preamble detection and frame extractions are 5.2 μ s and 2.4 μ s for

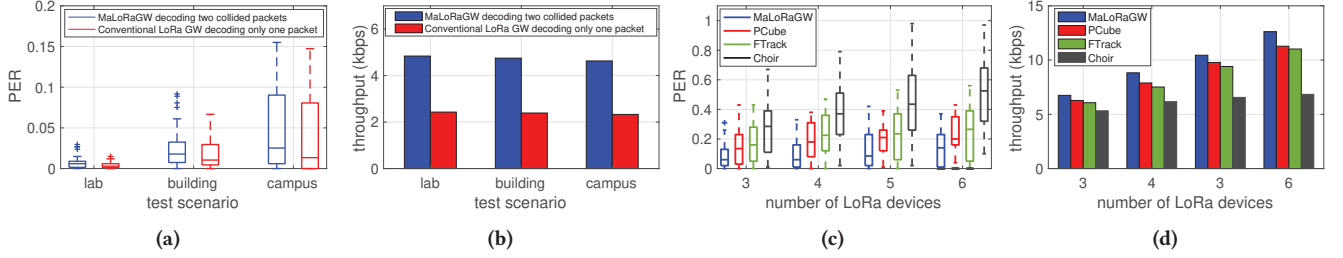


Figure 18: Uplink performance: (a) The PER performance when MaLoRaGW decodes two collided packets and a conventional LoRa gateway decodes a collision-free packet; (b) The throughput comparison between MaLoRaGW and a conventional LoRa gateway; (c) The PER performance comparison between MaLoRaGW and prior work; (d) The throughput comparison between MaLoRaGW and prior work.

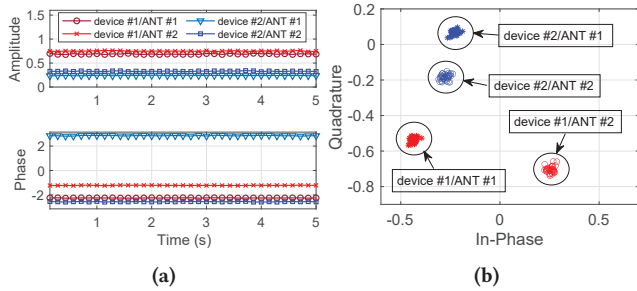


Figure 19: (a) The amplitude and phase of four channel coefficients estimated at MaLoRaGW over 40 consecutive packets; (b) The I/Q scatters of the four estimated channel coefficients.

the received signals from two devices. MaLoRaGW successfully decodes the two collided packets. The measured PER is zero for both users, and the uplink throughput is 4.9 kbps.

Downlink Performance. After decoding the uplink packets, MaLoRaGW takes advantage of the channels estimated in uplink to perform beamforming for downlink transmission. Two independent packet streams are sent to those two LoRa devices (USRPN210 in this case). Fig. 17 plots the received signals by the two LoRa devices in downlink. While the desired signal pattern is very clear in the figure, a trace of inter-user interference can be seen alongside the desired signal at both devices. The interference can be attributed to the channel estimation error, calibration error, and other circuit imperfections. As it can be seen in the figure, the interference is much weaker than the desired signal.

Accuracy of Estimated Channels. Since the channel information is critical for downlink beamforming, it would be interesting to quantify the accuracy of the channel coefficients estimated in uplink. Unfortunately, this is impossible because the ground truth of channel coefficients is unknown. To circumvent this challenge, we measure the stability of the measured channel coefficients over a short period of time. Fig. 19 presents our experimental results. It shows the four channel coefficients measured from 40 consecutive packets (spanning over 5 seconds) in both temporal and I/Q domains. We can see that the channel coefficients remain stable over time. This indirectly shows the accuracy of our channel estimation method. Another fact that shows the channel estimation accuracy is the successful packet decoding at two LoRa devices. If the estimated

channel coefficients were not accurate, the two LoRa devices would not be capable of decoding their own packets.

5.4 Uplink Performance

MaLoRaGW vs. Single-User LoRa Gateway. We compare the following two cases to quantify the performance of MaLoRaGW in uplink: i) MaLoRaGW is equipped with two antennas, and it decodes two collided packets from two uncoordinated COTS LoRa devices placed in all spots in Fig. 15. ii) A conventional LoRa gateway is equipped with one antenna. It decodes packets from only one active COTS LoRa device placed over all spots in Fig. 15. There is no collision in this case.

Fig. 18(a) presents the measured PER of MaLoRaGW and that of the single-user LoRa gateway. The mean measured PER at MaLoRaGW is 0.7% in lab, 2.6% in building, and 5.1% on campus. The mean measured PER at the single-user LoRa gateway is 0.5% in lab, 1.9% in building, and 4.4% on campus. The performance difference between MaLoRaGW and single-user LoRa gateway is 0.3% in lab, 0.7% in building, and 0.7% on campus. The experimental results indicate that a two-antenna LoRa gateway can double the number of serving users compared to a single-antenna LoRa gateway.

Fig. 18(b) plots the total uplink throughput achieved by MaLoRaGW and the single-user LoRa gateway. The average uplink throughput achieved by MaLoRaGW is 4.9 kbps in lab, 4.8 kbps in building, and 4.6 kbps on campus. In contrast, the average uplink throughput achieved by the conventional single-user gateway is 2.4 kbps in lab, 2.4 kbps in building, and 2.3 kbps on campus. The results reveal that MaLoRaGW almost doubles the uplink throughput compared to a one-antenna LoRa gateway working in non-collision mode.

MaLoRaGW vs. Existing Schemes. Above we studied the performance of MaLoRaGW when it decodes two collided packets. We increase the number of packets in collision to see how MaLoRaGW performs. To do so, we place K LoRa dongles (see Fig. 14(b)) at K different spots in Fig. 15, where $3 \leq K \leq 6$. The K LoRa dongles are configured to transmit their unique words (packets) to the gateway simultaneously and independently. We repeat the above measurement many times to cover all those marked spots in Fig. 15. We compare MaLoRaGW against Choir [7], FTrack [36] and PCube [35]. Choir and FTrack have one antenna on their gateways, while MaLoRaGW and PCube have two antennas on their gateways.




Figure 20: The PER performance of MaLoRaGW in uplink when different spreading factors are used.

Fig. 18(c) shows the measured PER of MaLoRaGW, PCube, FTrack, and Choir when the K LoRa dongles simultaneously send their packets to the gateway in the campus scenario. For MaLoRaGW, its average PER is 7.9% when $K = 3$, 9.6% when $K = 4$, 13.8% when $K = 5$, and 14.4% when $K = 6$. For PCube, its average PER is 14.3% when $K = 3$, 19.0% when $K = 4$, 19.8% when $K = 5$, and 23.1% when $K = 6$. For FTrack, its average PER is 17.1% when $K = 3$, 22.9% when $K = 4$, 23.0% when $K = 5$, and 24.8% when $K = 6$. For Choir, its average PER is 27.3% when $K = 3$, 36.9% when $K = 4$, 46.3% when $K = 5$, and 53.5% when $K = 6$. On average, MaLoRaGW reduces the average PER by 40.5% compared to PCube, 48.4% compared to FTrack, and 72.0% compared to Choir.

Fig. 18(d) plots the average uplink throughput achieved by MaLoRaGW, PCube, FTrack, and Choir when K ranges from 3 to 6. Specifically, the average uplink throughput achieved by MaLoRaGW is 6.7 kbps when $K = 3$, 8.8 kbps when $K = 4$, 10.4 kbps when $K = 5$, and 12.6 kbps when $K = 6$. The average uplink throughput achieved by PCube is 6.2 kbps when $K = 3$, 7.8 kbps when $K = 4$, 9.7 kbps when $K = 5$, and 11.2 kbps when $K = 6$. The average uplink throughput achieved by FTrack is 6.1 kbps when $K = 3$, 7.5 kbps when $K = 4$, 9.4 kbps when $K = 5$, and 11.0 kbps when $K = 6$. The average uplink throughput achieved by Choir is 5.3 kbps when $K = 3$, 6.2 kbps when $K = 4$, 6.6 kbps when $K = 5$, and 6.8 kbps when $K = 6$. On average, MaLoRaGW has 10.8% throughput gain compared to PCube, 13.5% throughput gain compared to FTrack and 53.4% throughput gain compared to Choir. The throughput gain over Choir and FTrack is from the gateway's multiple antennas, which allow us to exploit the spatial feature (channel ratio) for resolving packet collision. We believe the throughput gain over PCube is mainly from the spatial signal projection operation, which reduces the near-far effect of collided packets and thus improves the packet decoding probability.

Impact of Spreading Factor (SF). We now evaluate the PER performance of MaLoRaGW when two LoRa devices simultaneously send their packets using three different SF values, namely, 6, 8, and 10. Fig. 20 presents the PER measured at MaLoRaGW. When SF is set to 6, the average PER measured at MaLoRaGW is 1.3% in lab, 3.0% in building, and 6.3% in campus. When SF is set to 8, the average PER measured at MaLoRaGW is 1.1% in lab, 2.6% in building, and 6.2% in campus. When SF is set to 10, the average PER measured at MaLoRaGW is 0.7% in lab, 2.6% in building, and 5.1% in campus. As expected, PER decreases as SF increases. This is intuitive, as a larger SF brings a higher spreading gain for packet detection.




Figure 21: Measured RSSI at two LoRa dongles when they are concurrently served by MaLoRaGW in downlink.

5.5 Downlink Performance

Prior work does not support concurrent downlink transmission in LoRaWANs. Actually, concurrent downlink transmission has not been studied in the literature. We therefore compare MaLoRaGW with a conventional LoRa gateway in downlink, where MaLoRaGW is equipped with two antennas and the conventional LoRa gateway with one antenna. We wish to see the performance gain that can be obtained by adding one more antenna on a LoRa gateway.

In the experiments, we first trigger two LoRa dongles for concurrent uplink transmission. MaLoRaGW estimates their channels and decodes their packets. It then performs beamforming for downlink MU-MIMO transmission. Surprisingly, the channels remain coherent for a pretty long time (e.g., 5 seconds without Tx, Rx, and surrounding movement) in most cases. In the following results, the overhead of gateway's RF calibration was not considered as it only needs to be done at a low frequency.

Measured RSSI. Fig. 21 shows the measured RSSIs displayed by the two LoRa dongles' screen (see Fig. 14(b)) in downlink MU-MIMO transmission when those two LoRa dongles were placed at the locations marked in Fig. 15. The average RSSI value is -62.9 dBm in lab scenario, -107.6 dBm in building scenario, and -122.9 dBm in campus scenario. Per LoRa dongle's manual, the chip's receiver sensitivity is -140 dBm. It means that the LoRa dongles should be able to decode most of their packets. This inference is consistent with our experimental observation.

PER. In this measurement, MaLoRaGW performs downlink MU-MIMO transmission with two LoRa dongles, while the conventional LoRa gateway performs downlink transmission with a single LoRa dongle. Both gateways send 1000 packets to their dongles, and the LoRa dongles report the number of packets being successfully decoded. We then calculate their PER and throughput. Fig. 22(a) plots the PER performance of LoRa dongles when served by MaLoRaGW and the conventional LoRa gateway in those three scenarios. When served by MaLoRaGW, the average PER of LoRa dongles is 1.0% in lab scenario, 2.6% in building scenario, and 5.4% in campus scenario. In contrast, when served by the conventional LoRa gateway, the average PER at the LoRa dongle is 0.4% in lab, 1.8% in building, and 4.7% in campus scenario. Numerically, the PER difference between MaLoRaGW and a conventional LoRa gateway is less than 1% (0.6% in lab, 0.8% in building, and 0.7% in campus). This indicates that, from LoRa users' perspective, employing MU-MIMO at a gateway only slightly degrades their performance. This also indicates that employing MU-MIMO at a gateway is transparent to LoRa users.




Figure 22: Downlink performance: (a) Measured PER at LoRa dongles when they are served by MaLoRaGW (using MU-MIMO) and a conventional LoRa gateway; (b) The throughput comparison between MaLoRaGW and a conventional LoRa gateway; (c) The PER performance of MaLoRaGW when different spreading factors (SF) are used; (d) The PER performance of MaLoRaGW when different bandwidth is used.

Throughput. Fig. 22(b) shows the downlink throughput of MaLoRaGW and its conventional counterpart. The achievable network throughput in the downlink by the proposed MU-MIMO scheme. The throughput achieved by MaLoRaGW is 4.7 kbps in lab, 4.6 kbps in building, and 4.5 kbps in campus. In contrast, the throughput achieved by a conventional LoRa gateway is 2.4 kbps in lab, 2.4 kbps in building, and 2.3 kbps in campus. The measurement reveals that a two-antenna MaLoRaGW almost doubles ($1.95\times$ actually) the downlink throughput compared to a single-user LoRa gateway. This shows the potential of MU-MIMO in LoRa downlink transmission.

Impact of Spreading Factor (SF). We evaluate the PER performance of MaLoRaGW in downlink when using three different SF values, namely, 6, 8, and 10. Fig. 22(c) plots the PER measured at two LoRa dongles when served by MaLoRaGW in downlink MU-MIMO. When SF is set to 6, the average PER is 1.3% in lab, 3.2% in building, and 6.4% in campus. When SF is set to 8, the average PER is 1.2% in lab, 2.8% in building, and 6.2% in campus. When SF is set to 10, the average PER is 1.0% in lab, 2.6% in building, and 5.4% in campus. As expected, the PER at LoRa dongles decreases as the SF value increases. This is because a larger SF brings a higher spreading gain for packet detection.

Impact of Bandwidth. LoRa supports different bandwidth: 125 kHz, 250 kHz, and 500 kHz. We now conduct experiments to study the performance of MaLoRaGW with these three different bandwidths. We set SF to 10. Fig. 22(d) presents the measured PER at the two LoRa devices when served by MaLoRaGW in downlink. When bandwidth is 125 kHz, the average PER is 0.9% in lab, 2.3% in building, and 4.5% in campus. When bandwidth is 250 kHz, the average PER is 1.0% in lab, 2.6% in building, and 5.4% in campus. When bandwidth is 500 kHz, the average PER is 1.1% in lab, 2.8% in building, and 6.0% in campus. It can be seen that a smaller bandwidth offers a better PER performance. The reasons are twofold. First, a system with a smaller bandwidth will use more radio power to carry each of its bits, thereby reducing PER at its receiver. Second, a system with a smaller bandwidth tends to experience a frequency-flatter channel, which reduces the leakage of inter-user interference in downlink MU-MIMO transmission.

5.6 Limitations and Discussions

Antenna Limitation. While MaLoRaGW slightly improves the throughput of LoRa uplink transmission compared to the state-of-the-art, it significantly improves the throughput of LoRa downlink

transmission. However, our experiments of MaLoRaGW are limited to the two-antenna case. Theoretically, MaLoRaGW should work for a gateway with more (≥ 3) antennas, and it would provide higher throughput gain compared to single-antenna LoRa gateway. Validating the scalability of MaLoRaGW (over its antenna number) requires to substantiate efforts in system implementation, which will be carried out in our future work.

Hardware Cost and Energy Consumption. Compared to single-antenna LoRa gateways, MaLoRaGW has higher hardware cost and energy consumption. Fortunately, MaLoRaGW is backward compatible with commodity off-the-shelf LoRa devices and will not increase their hardware cost and energy consumption. Considering the fact that most LoRa gateways have sufficient power supplies, we believe the performance gain of MaLoRaGW outweighs the increase of its costs.

6 CONCLUSION

In this paper, we presented MaLoRaGW, the first-of-its-kind LoRa gateway that enables MU-MIMO transmission in LoRa networks. MaLoRaGW features a joint design for uplink packet detection and downlink beamforming, enabling it to concurrently serve multiple LoRa user devices in both uplink and downlink. The key component of MaLoRaGW is a joint baseband signal design for uplink packet detection and downlink beamforming, which are underpinned by three modules: spatial signal projection, accurate channel estimation, and implicit beamforming. We have evaluated a two-antenna MaLoRaGW in realistic scenarios of different scales. It has been validated that MaLoRaGW is backward compatible with COTS LoRa devices. It further demonstrates 10% throughput gain in uplink and 95% throughput gain in downlink when compared to the state-of-the-art.

ACKNOWLEDGMENTS

We sincerely thank the anonymous reviewers and our shepherd for their valuable feedback. This work was supported in part by NSF Grant CNS-2100112.

REFERENCES

- [1] Horia Vlad Balan, Ryan Rogalin, Antonios Michaloliakos, Konstantinos Psounis, and Giuseppe Caire. 2013. AirSync: Enabling distributed multiuser MIMO with full spatial multiplexing. *IEEE/ACM Transactions on Networking* 21, 6 (2013), 1681–1695.

- [2] Lili Chen, Jie Xiong, Xiaojiang Chen, Sunghoon Ivan Lee, Kai Chen, Dianhe Han, Dingyi Fang, Zhanyong Tang, and Zheng Wang. 2019. WideSee: Towards wide-area contactless wireless sensing. In *Proceedings of the 17th Conference on Embedded Networked Sensor Systems*. 258–270.
- [3] Qian Chen and Jiliang Wang. 2021. AlignTrack: Push the Limit of LoRa Collision Decoding. In *2021 IEEE 29th International Conference on Network Protocols (ICNP)*. IEEE, 1–11.
- [4] Zhe Chen, Xu Zhang, Sulei Wang, Yuedong Xu, Jie Xiong, and Xin Wang. 2017. BUSH: Empowering large-scale MU-MIMO in WLANs with hybrid beamforming. In *IEEE INFOCOM 2017-IEEE Conference on Computer Communications*. IEEE, 1–9.
- [5] LoRa Alliance Technical Committee et al. 2017. LoRaWAN 1.1 Specification. *LoRa Alliance, Standard 1*, 1 (2017).
- [6] Erik Dahlman, Stefan Parkvall, Johan Skold, and Per Beming. 2010. *3G evolution: HSPA and LTE for mobile broadband*. Academic press.
- [7] Rashad Eletreby, Diana Zhang, Swarnu Kumar, and Osman Yağan. 2017. Empowering low-power wide area networks in urban settings. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*. 309–321.
- [8] Sezana Fahmida, Venkata P Modekurthy, Mahbubur Rahman, Abusayeed Saifulah, and Marco Brocanelli. 2020. Long-lived LoRa: Prolonging the lifetime of a LoRa network. In *2020 IEEE 28th International Conference on Network Protocols (ICNP)*. IEEE, 1–12.
- [9] Amalinda Gamage, Jansen Christian Liando, Chaojie Gu, Rui Tan, and Mo Li. 2020. LMAC: Efficient carrier-sense multiple access for LoRa. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*. 1–13.
- [10] Ningning Hou, Xianjin Xia, and Yuanqing Zheng. 2021. Jamming of LoRa PHY and countermeasure. In *IEEE Conference on Computer Communications*. 1–10.
- [11] Bin Hu, Zhimeng Yin, Shuai Wang, Zhuqing Xu, and Tian He. 2020. SCLoRa: leveraging multi-dimensionality in decoding collided LoRa transmissions. In *2020 IEEE 28th International Conference on Network Protocols (ICNP)*. IEEE, 1–11.
- [12] Evgeny Khorov, Anton Kiryanov, Andrey Lyakhov, and Giuseppe Bianchi. 2018. A tutorial on IEEE 802.11ax high efficiency WLANs. *IEEE Communications Surveys & Tutorials* 21, 1 (2018), 197–216.
- [13] Chenning Li and Zhichao Cao. 2022. LoRa networking techniques for large-scale and long-term IoT: A down-to-top survey. *ACM Computing Surveys (CSUR)* 55, 3 (2022), 1–36.
- [14] Chenning Li, Hanqing Guo, Shuai Tong, Xiao Zeng, Zhichao Cao, Mi Zhang, Qiben Yan, Li Xiao, Jiliang Wang, and Yunhao Liu. 2021. NELoRa: Towards Ultra-low SNR LoRa Communication with Neural-enhanced Demodulation. In *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems*. 56–68.
- [15] Chenning Li, Xiuzhen Guo, Longfei Shangguan, Zhichao Cao, and Kyle Jamieson. 2022. CurvingLoRa to Boost LoRa Network Capacity via Concurrent Transmission. *arXiv preprint arXiv:2201.05179* (2022).
- [16] Yinghui Li, Jing Yang, and Jiliang Wang. 2020. DyLoRa: Towards energy efficient dynamic LoRa transmission control. In *IEEE Conference on Computer Communications*. 2312–2320.
- [17] Paul Marcellis, Nikolaos Kouvelas, Vijay S Rao, and Venkatesha Prasad. 2020. DaRe: Data recovery through application layer coding for LoRaWAN. *IEEE Transactions on Mobile Computing* (2020).
- [18] Konstantin Mikhaylov, Juha Petäjäjärvi, and Janne Janhunen. 2017. On LoRaWAN scalability: Empirical evaluation of susceptibility to inter-network interference. In *2017 European Conference on Networks and Communications (EuCNC)*. IEEE, 1–6.
- [19] Eldad Perahia and Robert Stacey. 2013. *Next generation wireless LANs: 802.11n and 802.11ac*. Cambridge university press.
- [20] Pedram Kheirkhah Sangdeh, Hossein Pirayesh, Aryan Mobiny, and Huacheng Zeng. 2020. LB-SciFi: Online learning-based channel feedback for MU-MIMO in wireless LANs. In *2020 IEEE 28th International Conference on Network Protocols (ICNP)*. IEEE, 1–11.
- [21] Pedram Kheirkhah Sangdeh and Huacheng Zeng. 2021. DeepMux: Deep-Learning-Based Channel Sounding and Resource Allocation for IEEE 802.11ax. *IEEE Journal on Selected Areas in Communications* 39, 8 (2021), 2333–2346.
- [22] Muhammad Osama Shahid, Millan Philipose, Krishna Chintalapudi, Suman Banerjee, and Bhuvana Krishnaswamy. 2021. Concurrent interference cancellation: decoding multi-packet collisions in LoRa. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*. 503–515.
- [23] Clayton Shepard, Hang Yu, Narendra Anand, Erran Li, Thomas Marzetta, Richard Yang, and Lin Zhong. 2012. Argos: Practical many-antenna base stations. In *Proceedings of the 18th annual international conference on Mobile computing and networking*. 53–64.
- [24] Nicolas Sornin, Miguel Luis, Thomas Eirich, Thorsten Kramp, and Olivier Hersent. 2015. LoRaWAN specification. *LoRa alliance* (2015).
- [25] Maria Stoyanova, Yannis Nikoloudakis, Spyridon Panagiotakis, Evangelos Pallis, and Evangelos K Markakis. 2020. A survey on the Internet of Things (IoT) forensics: challenges, approaches, and open issues. *IEEE Communications Surveys & Tutorials* 22, 2 (2020), 1191–1221.
- [26] Jothi Prasanna Shanmuga Sundaram, Wan Du, and Zhiwei Zhao. 2019. A survey on LoRa networking: Research problems, current solutions, and open issues. *IEEE Communications Surveys & Tutorials* 22, 1 (2019), 371–388.
- [27] Shuai Tong, Zilin Shen, Yunhao Liu, and Jiliang Wang. 2021. Combating link dynamics for reliable LoRa connection in urban settings. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*. 642–655.
- [28] Shuai Tong, Jiliang Wang, and Yunhao Liu. 2020. Combating packet collisions using non-stationary signal scaling in LPWANs. In *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*. 234–246.
- [29] Shuai Tong, Zhenqiang Xu, and Jiliang Wang. 2020. CoLoRa: Enabling Multi-Packet Reception in LoRa. In *Proceedings of IEEE International Conference on Computer Communications (INFOCOM)*. 2303–2311.
- [30] Cheng-Xiang Wang, Ji Bian, Jian Sun, Wensheng Zhang, and Minggao Zhang. 2018. A survey of 5G channel measurements and models. *IEEE Communications Surveys & Tutorials* 20, 4 (2018), 3142–3168.
- [31] Xiong Wang, Linghe Kong, Liang He, and Guihai Chen. 2019. MLoRa: A multi-packet reception protocol in LoRa networks. In *Proceedings of IEEE 27th International Conference on Network Protocols (ICNP)*. 1–11.
- [32] Xiong Wang, Linghe Kong, Zucheng Wu, Long Cheng, Chenren Xu, and Guihai Chen. 2020. SLoRa: towards secure LoRa communications with fine-grained physical layer features. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*. 258–270.
- [33] Zhe Wang, Linghe Kong, Kangjie Xu, Liang He, Kaishun Wu, and Guihai Chen. 2020. Online concurrent transmissions at LoRa gateway. In *Proceedings of IEEE International Conference on Computer Communications (INFOCOM)*. 2331–2340.
- [34] Andrew J Wixted, Peter Kinnaird, Hadi Larjani, Alan Tait, Ali Ahmadinia, and Niall Strachan. 2016. Evaluation of LoRa and LoRaWAN for wireless sensor networks. In *2016 IEEE SENSORS*. IEEE, 1–3.
- [35] Xianjin Xia, Ningning Hou, Yuanqing Zheng, and Tao Gu. 2021. PCube: scaling LoRa concurrent transmissions with reception diversities. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*. 670–683.
- [36] Xianjin Xia, Yuanqing Zheng, and Tao Gu. 2020. FTrack: Parallel decoding for LoRa transmissions. *IEEE/ACM Transactions on Networking* 28, 6 (2020), 2573–2586.
- [37] Binbin Xie and Jie Xiong. 2020. Combating interference for long range LoRa sensing. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*. 69–81.
- [38] Binbin Xie, Yuqing Yin, and Jie Xiong. 2021. Pushing the limits of long range wireless sensing with LoRa. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 3 (2021), 1–21.
- [39] Zhenqiang Xu, Pengjin Xie, and Jiliang Wang. 2021. Pyramid: Real-time LoRa collision decoding with peak tracking. In *IEEE INFOCOM 2021-IEEE Conference on Computer Communications*. IEEE, 1–9.
- [40] Qing Yang, Xiaoxiao Li, Hongyi Yao, Ji Fang, Kun Tan, Wenjun Hu, Jiansong Zhang, and Yongguang Zhang. 2013. BigStation: Enabling scalable real-time signal processing in large MU-MIMO systems. *ACM SIGCOMM Computer Communication Review* 43, 4 (2013), 399–410.
- [41] Junmei Yao, Jun Xu, Sheng Luo, Lu Wang, Chao Yang, Kaishun Wu, and Wei Lou. 2019. Comprehensive study on MIMO-related interference management in WLANs. *IEEE Communications Surveys & Tutorials* 21, 3 (2019), 2087–2110.
- [42] Yuguang Yao, Zijun Ma, and Zhichao Cao. 2019. LoSee: Long-Range Shared Bike Communication System Based on LoRaWAN Protocol. In *EWSN*. 407–412.
- [43] Huacheng Zeng, Hongxiang Li, and Qiben Yan. 2018. Uplink MU-MIMO in asynchronous wireless LANs. In *Proceedings of the Eighteenth ACM International Symposium on Mobile Ad Hoc Networking and Computing*. 21–30.
- [44] Yunze Zeng, Ioannis Pefkianakis, Kyu-Han Kim, and Prasant Mohapatra. 2017. MU-MIMO-Aware AP selection for 802.11ac networks. In *Proceedings of the 18th ACM International Symposium on Mobile Ad Hoc Networking and Computing*. 1–10.
- [45] Fusang Zhang, Zhaoxin Chang, Kai Niu, Jie Xiong, Beihong Jin, Qin Lv, and Daqing Zhang. 2020. Exploring LoRa for long-range through-wall sensing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 2 (2020), 1–27.
- [46] Fusang Zhang, Zhaoxin Chang, Jie Xiong, and Daqing Zhang. 2021. Exploring LoRa for Sensing. *GetMobile: Mobile Computing and Communications* 25, 2 (2021), 33–37.
- [47] Fusang Zhang, Zhaoxin Chang, Jie Xiong, Rong Zheng, Junqi Ma, Kai Niu, Beihong Jin, and Daqing Zhang. 2021. Unlocking the beamforming potential of LoRa for long-range multi-target respiration sensing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 2 (2021), 1–25.
- [48] Anfu Zhou, Teng Wei, Xinyu Zhang, Min Liu, and Zhongcheng Li. 2015. Signpost: Scalable MU-MIMO signaling with zero CSI feedback. In *Proceedings of the 16th ACM International Symposium on Mobile Ad Hoc Networking and Computing*. 327–336.