# FDLoRa: Tackling Downlink-Uplink Asymmetry with Full-duplex LoRa Gateways

Shiming Yu[1], Xianjin Xia[1], Ziyue Zhang[1], Ningning Hou[2], Yuanqing Zheng[1]
[1] The Hong Kong Polytechnic University, Hong Kong SAR, China
[2] Macquarie University, Sydney, Australia
shiming.yu@connect.polyu.hk,xianjin.xia@polyu.edu.hk,ziyue.zhang@connect.polyu.hk,
ningning.hou@mq.edu.au,yuanqing.zheng@polyu.edu.hk

## ABSTRACT

Unlike traditional data collection applications (*e.g.*, environment monitoring) that are dominated by uplink transmissions, the newly emerging applications (*e.g.*, device actuation, firmware update, packet reception acknowledgement) also pose ever-increasing demands on downlink transmission capabilities. However, current LoRaWAN falls short in supporting such applications primarily due to downlink-uplink asymmetry. While the uplink can concurrently receive multiple packets, downlink transmission is limited to a single logical channel at a time, which fundamentally hinders the deployment of downlink-hungry applications. To tackle this practical challenge, *FDLoRa* develops the first-of-its-kind in-band full-duplex LoRa gateway design with novel solutions to mitigate the impact of self-interference (*i.e.*, strong downlink interference to ultra-weak uplink reception), which unleashes the full spectrum for in-band downlink transmissions without compromising the reception of weak uplink packets. Built upon the full-duplex gateways, *FDLoRa* introduces a new downlink framework to support concurrent downlink transmissions over multiple logical channels of available gateways. Evaluation results demonstrate that *FDLoRa* boosts downlink capacity by 5.7× compared to LoRaWAN on a three-gateway testbed and achieves 2.58× higher downlink concurrency per gateway than the state-of-the-art.

## CCS CONCEPTS

• **Computer systems organization** → **Embedded systems**; • **Networks** → **Network protocol design**.

## KEYWORDS

Internet of Things, LPWAN, LoRa, Full Duplex, Logical Channel

**ACM Reference Format:**
Shiming Yu[1], Xianjin Xia[1], Ziyue Zhang[1], Ningning Hou[2], Yuanqing Zheng[1] . 2024. FDLoRa: Tackling Downlink-Uplink Asymmetry with Full-duplex LoRa Gateways. In *The 22nd ACM Conference on Embedded Networked Sensor Systems (SenSys '24), November 4–7, 2024, Hangzhou, China.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3666025.3699338


*SenSys '24, November 4–7, 2024, Hangzhou, China*

ACM ISBN 979-8-4007-0697-4/24/11
https://doi.org/10.1145/3666025.3699338

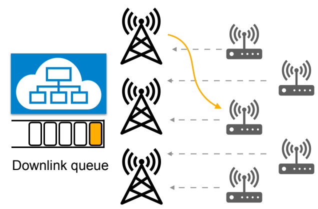


(a) Traditional LoRa downlink

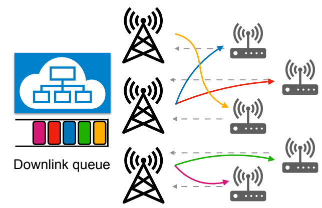


(b) Concurrent downlink framework

**Figure 1: Downlink-uplink asymmetry: (a) LoRaWAN receives uplink concurrently but transmits downlink sequentially; (b) *FDLoRa* scales concurrent downlink to massive logical channels of available gateways.**

## 1 INTRODUCTION

Recent years have witnessed Low-Power Wide-Area Networks (LPWANs) as promising technologies to enable ubiquitous connectivity of massive networked sensors to cyberspace, underpinning a variety of embedded AI and sensing systems and their applications [2, 8, 10, 11, 15, 18–21, 23, 24, 29, 36, 39, 40, 43, 44, 46, 48, 49, 51, 59–61, 71–73, 77, 80, 85, 91–94, 96, 97, 99]. As a leading LPWAN technology, LoRa features extensive coverage and energy-efficient communication, connecting over 300 million LoRa nodes and 5.9 million gateways worldwide [54].

These deployed nodes are supporting various IoT applications, where both uplink and downlink transmissions are indispensable to read and control sensors, enable essential systems and networking functionalities, ensure communication reliability, and optimize overall performance. For example, smart lamppost systems [16] necessitate both uplink and downlink transmissions for status acquisition and device actuation (*e.g.*, turn on/off lights), and firmware update services [62, 64] need both uplink and downlink for device authentication and firmware transfer. To support reliable data transfer [22, 28] and data rate adaptation of LoRa nodes [33, 55], gateways need to acknowledge successful uplink transmissions to the nodes via downlink.

In supporting the above applications and systems operations, current LoRa networks face practical challenges due to the *asymmetry between downlink and uplink transmissions*. Although recent research works have made remarkable strides [58, 84] in improving the uplink performance (*e.g.*, concurrent uplink transmissions [58, 83, 84], weak packet reception [25, 37]), the problem of poor downlink performance largely remains under-explored. While commodity LoRa gateways are typically equipped with multiple Rx chains (*e.g.*, nine channels [53]) to handle concurrent uplink packets, they have only one Tx chain for downlink transmission. As a result, while multiple uplink packets can be received concurrently, their corresponding downlink packets have to be queued and processed one by one. As systems scale in supporting ever-increasing numbers of IoT devices, more and more downlink packets are being accumulated at network servers, leading to increased latency and excessive congestion. The asymmetry between downlink and uplink transmissions could also cause the power depletion of sensors in excessive packet retransmissions and even potential network collapse.

Current LoRaWAN transmits downlink packets sequentially (as illustrated in Figure 1(a)), creating a bottleneck that hampers overall network performance due to limited downlink capacity. In this paper, we explore whether it is possible to improve downlink concurrency within the LoRaWAN architecture. We notice that LoRaWAN adopts massive logical channels [95, 98] for both uplink and downlink transmissions. The downlink packets queued at the network server could be assigned to orthogonal logical channels and transmitted simultaneously without causing much interference (Figure 1(b)). Rather than transmitting downlinks sequentially through a single logical channel as the existing LoRaWAN, we aim to scale concurrent downlink transmissions to massive logical channels. By enabling downlink concurrency, we aim to fundamentally break the barrier of the current downlink transmission paradigm and fill the gap caused by the asymmetry between the downlink and uplink transmissions to better support various IoT applications with diverse communication demands.

However, implementing the above idea into a practical system entails substantial technical challenges. Firstly, unlike traditional LoRaWAN which transmits only one downlink packet via a single gateway, we aim to support concurrent downlink transmissions over multiple logical channels through available gateways. However, current LoRa gateways [68, 69] face a dilemma: allocating more gateways for downlink transmissions would miss uplink packet receptions. Previous full-duplex radio designs [6, 63] cannot support full-duplex LoRa communication because of high self-interference of gateways to the ultra-weak uplink packets over long ranges. Moreover, concurrent downlink over logical channels would divide the transmit power and increase the demodulation noise floor, potentially compromising the reliability of downlink packet reception at low-cost COTS end nodes. Scaling to more gateways even requires a new scheduling strategy and faces inter-gateway interference. Therefore, it is crucial, albeit challenging, to enhance downlink performance without compromising reliability and uplink reception.

To address the practical challenges, this paper presents *FDLoRa* which supports in-band full-duplex communications for LoRa gateways. We find that uplink and downlink transmissions in LoRaWAN can be configured to adopt orthogonal modulations that cause minimum self-interference. By leveraging this orthogonality, we develop self-interference cancellation methods to further reduce the impact of downlink transmissions to the receptions of ultra-weak uplink packets. By doing so, *FDLoRa* can make the full use of wireless spectrum by simultaneously supporting downlink transmissions and uplink receptions. Unlike traditional full-duplex designs [6, 63], *FDLoRa* leverages the unique LoRa modulation and develops novel self-interference cancellation algorithms to further eliminate the self-interference below the noise floor. To unlock the potential of full-duplex gateways, we develop a new downlink framework that transmits downlink packets on multiple logical channels simultaneously and utilizes all available gateways. Based on the global information of instant link qualities and network dynamics, *FDLoRa* predicts the physical layer performance of concurrent downlinks to ensure communication reliability. The prediction results can further guide packet configurations (*e.g.*, channel selection, data rate adaptation, power control) and downlink scheduling across gateways.

We implement and evaluate *FDLoRa* using commercial off-the-shelf (COTS) LoRa devices and Software Defined Radio (SDR). Results show that *FDLoRa* significantly outperforms existing LoRaWAN downlink. Specifically, *FDLoRa* supports up to 5.16× downlink concurrency per gateway and scales the network downlink capacity by 5.7× than current LoRaWAN. *FDLoRa* can be seamlessly integrated with existing LoRaWAN architecture and hardware with software modification on the network server.

In summary, this paper makes the following contributions: (1) *FDLoRa* presents the-first-of-its-kind gateway design that enables in-band simultaneous downlink transmission and uplink reception, fully unlocking the spectrum and gateways for concurrent downlink transmission without compromising uplink; (2) *FDLoRa* introduces a new downlink framework to scale concurrent downlink transmissions to massive logical channels through available gateways; (3) We propose novel techniques to cancel the self-interference in in-band full-duplex LoRa communication and enable concurrent downlink to massive end nodes over logical channels and gateways; (4) We implement and evaluate *FDLoRa* to validate its effectiveness. Our evaluations demonstrate significant downlink scalability gains than LoRaWAN and state-of-the-art.

## 2 BACKGROUND AND MOTIVATION

In this section, we introduce the background of LoRa (§2.1) and analyze the root cause of downlink-uplink asymmetry in LoRa networks and its impact on network functions and applications (§2.2). Finally, we present our motivation of full-duplex gateways and concurrent downlink transmissions over massive logical channels (§2.3).

### 2.1 LoRa Background

**LoRa PHY.** LoRa employs chirp spread spectrum (CSS) modulation to enable long-range communication while maintaining low power consumption [34, 41, 50, 65, 66, 88, 90]. CSS utilizes chirps to modulate symbols. The duration of a chirp is defined by Spreading Factor (SF) and bandwidth (BW) parameters, *i.e.*, $T = \frac{2^{SF}}{BW}$. As shown in Figure 2(a), a *base up-chirp* whose frequency increases from $-\frac{BW}{2}$ to $\frac{BW}{2}$ can be represented as follows:

$$C(k,t) = e^{j2\pi(\frac{1}{2}kt - \frac{BW}{2})t} \tag{1}$$

where the chirp slope $k = \frac{BW}{T}$ denotes the frequency changing rate.

LoRa modulates a symbol by varying the initial frequency of an up-chirp. As shown in Figure 2(b), we denote a modulated symbol

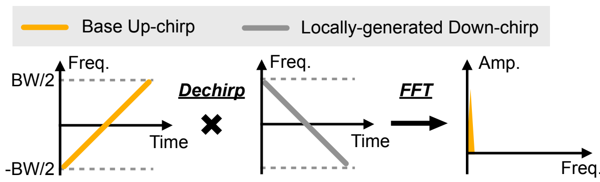

(a) Base Up-chirp

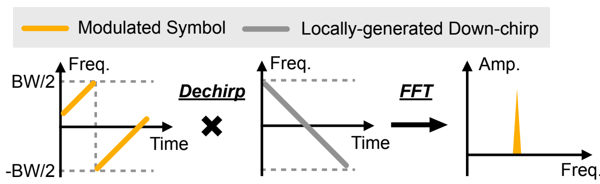

(b) Modulated Symbol

**Figure 2: Chirp Spread Spectrum of LoRa.**

$S(f_{sym}, k, t)$ as follows:

$$S(f_{sym}, k, t) = C(k, t)\ e^{j2\pi f_{sym} t} \tag{2}$$

To demodulate the symbol, a LoRa receiver performs *dechirp* by multiplying $S(f_{sym}, k, t)$ with a *down-chirp* represented as $C^{-1}(k, t)$, which is the conjugate of a base up-chirp. The operation is represented as follows:

$$S(f_{sym}, k, t)\ C^{-1}(k, t) = e^{j2\pi f_{sym} t} \tag{3}$$

The *dechirp* operation eliminates the CSS modulation and transforms the LoRa symbol into a single tone. Notably, the symbol energy that initially spread across the chirp bandwidth is migrated at the tone frequency after *dechirp*. Finally, Fast Fourier Transform (FFT) aggregates the signal energy into a peak at $f_{sym}$ on the frequency domain, which indicates the symbol. Essentially, FFT not only extracts the frequency itself but also accumulates the whole chirp energy into a single FFT bin, which yields a substantial *SNR gain after dechirping*. The SNR gain elevates the LoRa symbol above the ambient noise, facilitating long-distance and below-the-noise communications.

A LoRa packet initiates with a preamble, continues with two up-chirps serving as synchronization words and 2.25 down-chirps as a Start Frame Delimiter (SFD), and ends with a packet payload. The preamble, synchronization words, and SFD serve as the header of the LoRa packet and have a fixed signal pattern. In contrast, the payload carries modulated symbols with different initial frequencies.

**LoRa Channels.** LoRa allows nodes to select and configure channel frequency and bandwidth (BW) before packet transmissions. The channel frequency and bandwidth uniquely define a *physical channel*. Additionally, multiple LoRa nodes can simultaneously transmit in the same physical channel but with different spreading factors (SFs), creating multiple orthogonal *logical channels* [98]. Different logical channels are orthogonal because they use different chirp slopes. When receiving over the logical channel with chirp slope $k$, the interference from another logical channel

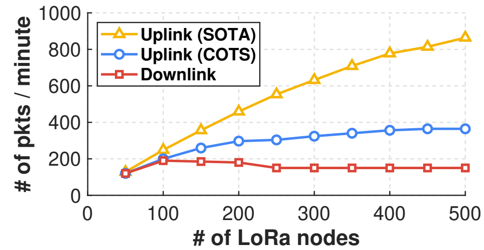


**Figure 3: Downlink-uplink asymmetry bottlenecks effective network capacity. SOTA indicates multi-channel reception [53] with collision recovery [58].**

with chirp slope $k'$ after dechirp can be represented as follows:

$$S(f_{sym}, k', t)\ C^{-1}(k, t) = e^{j2\pi[\frac{1}{2}(k'-k)t + f_{sym}]t} \tag{4}$$

In the above equation, we see the signal frequency after dechirp still changes over time, since $k' \neq k$. As a result, the chirp energy will spread over the spectrum rather than concentrating into a single peak after FFT (if $k' = k$). Hence, chirps with different slopes can be largely considered orthogonal after dechirping.

**LoRa Network Architecture.** A LoRa network typically consists of massive end nodes, multiple gateways, and a network server. Gateways are equipped with multiple Rx chains to receive uplink packets simultaneously, and one Tx chain for downlink transmission. Generally, gateways forward received packets to the network server over the backhaul. The network server implements and optimizes various network functionalities such as uplink consolidation, downlink scheduling, and network resource optimization, with a global view of network status [56].

## 2.2 Understanding Downlink Bottleneck

**Need for robust LoRa downlink.** LoRa and other LPWAN technologies are designed for IoT communications, where the uplink is dominant, and the downlink requirements are typically minimal (*e.g.*, <10 downlink packets per day [27]). However, recent booming IoT applications pose an increasing demand on downlinks. For example, emerging applications like smart lamppost control systems [16] require frequent status acquisition and device control through downlink communication. Even in uplink-dominant applications such as smart agriculture and environmental monitoring, practical systems still rely on various downlink services such as data rate adaptation, uplink acknowledgment, and firmware updates to facilitate essential network operations. Considering the massive IoT devices deployed worldwide [4], efficient and reliable downlink communication becomes increasingly important.

**Downlink-uplink asymmetry.** Current LoRaWAN faces significant downlink-uplink asymmetry. As shown in Figure 3, our measurements[1] reveal that the downlink remains highly constrained (*i.e.*, <200 packets per minute) as the number of LoRa nodes increases. In contrast, the uplink capacity continues to grow with the support of multi-channel gateways with concurrent transmission

[1]We conduct measurements on a three-gateway testbed with fifty nodes and emulate more nodes by increasing the transmit duty cycle. Nodes randomly select uplink/downlink channel for each transmission.

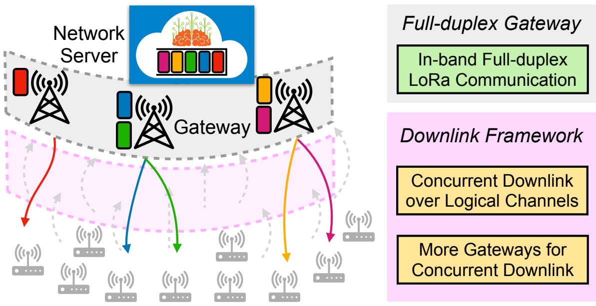


**Figure 4: System overview: *FDLoRa* develops full-duplex gateways and a new downlink framework to tackle downlink-uplink asymmetry.**

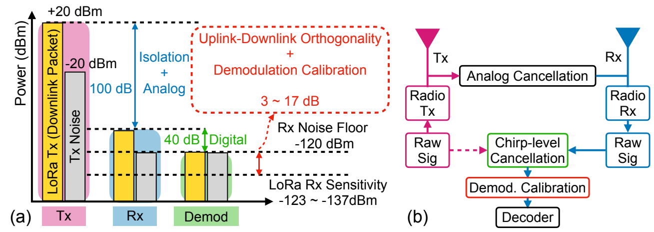


**Figure 5: Full-duplex LoRa gateway: (a) Requirements of self-interference cancellation; (b) full-duplex cancellation architecture of *FDLoRa*.**

capabilities [58, 84]. As a result, we observe high latency and excessive congestion because the downlink capacity cannot catch up with the relatively high uplink transmissions. Moreover, as more nodes keep retransmitting the same unacknowledged packets in vain of downlink feedback, the downlink-uplink asymmetry leads to ineffective network transmissions and power depletion of nodes.

**Limitations of current downlink.** The main reason for downlink-uplink asymmetry lies in three folds. Firstly, as most gateways [68, 69] operate in the half-duplex mode, downlink transmissions can conflict with uplink. Recent work has made tremendous progresses but only enables *out-of-band* full-duplex, where the uplink and downlink spectra need to be separated [3, 70]. The *out-of-band* full-duplex limits their applicability to specific LoRaWAN regional specifications, such as AU915 and US915 [3], and results in doubled spectrum usage. Secondly, gateways suffer from high asymmetry in uplink and downlink concurrency. Commodity gateways employ multiple Rx chains (*e.g.*, nine channels [53]) to receive concurrent uplink reception, while only one Tx chain is reserved for downlink transmission. Lastly, the gateway utilization is also asymmetric. LoRa networks typically deploy multiple gateways to collect uplink transmissions simultaneously. In contrast, current LoRaWAN only selects one network server and gateway at a time for downlink transmission [1, 57].

## 2.3 Motivation

The downlink-uplink asymmetry limits network operations in embedded AI and sensing systems, calling for novel solutions. We observe that conventional LoRaWANs typically employ half-duplex gateways and transmit downlink messages one by one through a single gateway. This design paradigm fundamentally limits the downlink capacity. To deal with resource conflict, we propose *a new in-band full-duplex LoRa gateway design – FDLoRa*, to make the full use of available network resources without compromising uplink reception. In addition, we aim to scale the downlink transmissions across massive logical channels and enable concurrent downlink transmissions. We note that *FDLoRa* can be seamlessly integrated with existing LoRaWANs without any hardware or software modification to deployed COTS LoRa nodes. Compared with half-duplex and out-of-band full-duplex designs, *FDLoRa* can offer broader feasibility for various channel divisions, better spectrum utilization efficiency, and hold greater promise for representative applications such as real-time data transmission, rapid response to safety commands, and efficient management of high-density networks [42].

# 3 SYSTEM OVERVIEW

To tackle the downlink-uplink asymmetry of LoRaWAN, our work *FDLoRa* enables in-band full-duplex gateways with concurrent downlink transmission capabilities. The workflow of *FDLoRa* is shown in Figure 4.

To resolve the resource conflicts between uplink and downlink, *FDLoRa* develops in-band full-duplex LoRa gateways. *FDLoRa* addresses unique challenges involved in achieving in-band full-duplex LoRa communication without compromising either uplink or downlink performance.

Building upon the full-duplex gateways, *FDLoRa* introduces a novel concurrent downlink transmission framework to substantially improve the scalability of downlink transmissions. *FDLoRa* leverages massive LoRa logical channels to enable one gateway to concurrently send multiple downlink packets that can be readily received by deployed COTS LoRa nodes. *FDLoRa* also utilizes all available gateways for downlink transmission. By filling the gap between downlink and uplink (see Figure 3), *FDLoRa* is envisioned to meet the practical diverse demands of current and emerging IoT applications [16, 62, 64].

# 4 IN-BAND FULL-DUPLEX GATEWAY

In this section, we introduce the design of *FDLoRa* in-band full-duplex gateway, which enables simultaneous downlink transmission and uplink reception. We first identify the requirements of an in-band full-duplex LoRa gateway and introduce the technical challenges (§4.1). Next, we present our solutions (§4.2).

## 4.1 Requirements of in-band full-duplex LoRa gateway

The key challenge of developing an in-band full-duplex LoRa gateway is to cancel the strong *self-interference (SI)* to uplink packets. As illustrated in Figure 5, a gateway's downlink transmission consists of the modulated signals of a downlink packet and the hardware noise of the transmitter, both of which can be orders of magnitude stronger than the incoming uplink packets received over long communication distances. The SI can be times stronger in an in-band full-duplex radio than that of an out-of-band full-duplex radio (*e.g.*,

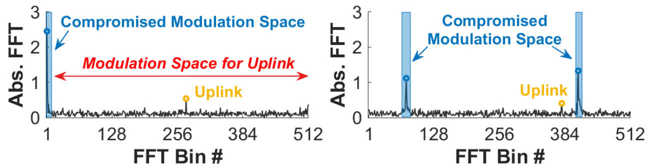


(a) Base chirp (preamble) (b) Non-base chirp (payload)

**Figure 6: (a) Suppressing interference bins harms the effective modulation space for uplink. The uplink symbol is submerged if it lies close to interference bins (represented as blue zones). (b) The problem is more severe under payload self-interference.**

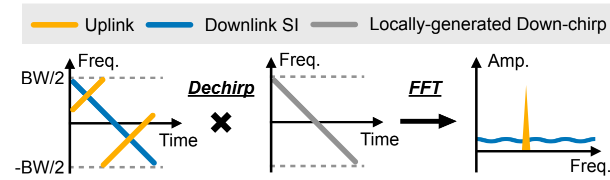


**Figure 7: Uplink-downlink orthogonality: downlink self-interference can be separated with conjugate modulation.**

WisGate RAK7285 [70]) because SI degrades fast across channels. However, an out-of-band full-duplex radio occupies extra channels that could otherwise be used for more LoRa transmissions.

To enable in-band full-duplex communication without compromising Rx sensitivity, it is crucial to minimize and cancel the in-band self-interference from the downlink transmissions (both downlink packet and Tx noise), which is known as *Self-Interference Cancellation (SIC)*. Existing works [6, 7, 32] typically adopt a three-stage solution on both analog and digital domains: (1) Isolation between Tx-Rx antennas; (2) Analog cancellation; and (3) Digital cancellation. As digital cancellation alone cannot completely remove SI and Tx noises [6], we need to preliminarily suppress SI and Tx noises through analog components before digital processing. To achieve this, the Tx and Rx antennas of a full-duplex radio are physically separated apart [9] and wire-connected with a specifically designed analog circuit for SIC [63]. Such antenna isolation and analog cancellation techniques suppress SI and Tx noises by >100 dB [63] (see Figure 5(a)). Subsequently, digital cancellation leverages the prior knowledge of the Tx signal to remove the residual SI signals in the digital domain, which can further reduce interference by around 40 dB and suppress SI to the same power level of Rx noises.

**Limitations of state-of-the-art (SOTA) full-duplex radios.** Given the high sensitivity of a LoRa radio (*e.g.*, down to -137 dBm [52]), the self-interference from a gateway's downlink can still trigger packet detection and reception of the gateway even with the SI being suppressed close to the noise floor (*e.g.*, using a traditional SIC technique). It imposes extra requirements on the SIC of a full-duplex gateway to further suppress SI signals below the Rx sensitivity of a LoRa radio (*e.g.*, usually tens of dB lower than the noise floor [52]).

A recent work [63] incorporates new analog and digital SIC techniques that take account of LoRa signal features and use optimized circuits together with digital SIC to achieve a remarkable 117.7 dB SIC gain. The full-duplex gateway design is proven effective in mitigating the self-interference of a specific downlink busy signal consisting of the same base up-chirps, thereby reserving uplinks for end nodes and avoiding packet collisions [63]. However, the previous work [63] cannot be applied to the SIC of in-band full-duplex LoRa communications in our scenario where the downlink carries LoRa data packets rather than pre-defined busy signals for media access control purposes as in [63].

In-band full-duplex LoRa communication faces two unique challenges that have yet to be addressed. Firstly, a normal downlink packet consists of not only base chirps in the preamble but also non-base chirps in the payload part. The existing LoRa SIC method [63] suppresses interference during demodulation by removing certain frequency bins, which would reduce the effective space for payload data modulation (see Figure 6(a)). Such a method is particularly difficult for the non-base chirps of LoRa payload because non-base chirps usually suffer more severe frequency misalignment and side lobe issues than base chirps [23]. Given that downlink SI is much stronger than uplink signals, even the side-lobe bins of SI can overwhelm the target bin of a weak uplink. We have to remove multiple surrounding bins to effectively suppress interference, which can drastically reduce the modulation space of payload data (see Figure 6(b)). On the other hand, existing SIC methods [6, 63] require estimating channels from a preamble and applying channel parameters to the entire packet for digital cancellation. However, a LoRa packet has a long air-time (*e.g.*, hundreds of ms), while a preamble constitutes only a small portion of the packet (*e.g.*, <20 %). Channel parameters may change across a LoRa packet as the channel coherence time lasts only tens of ms (*i.e.*, shorter than a LoRa packet) [78]. The channel estimated from a preamble may become outdated when applied to a payload and fail the digital cancellation. It is important yet challenging to acquire channel information from a packet at fine degrees to enable full-duplex LoRa communications.

## 4.2 SIC for Full-duplex Communication

*FDLoRa* implements an in-band full-duplex LoRa gateway by systematically mitigating the self-interference (SI) from downlink to uplink using both analog and digital SIC techniques. *FDLoRa* uses Tx/Rx antenna isolation and an analog SIC circuit to suppress strong downlink SI to an acceptable power level, allowing normal operations of Rx radio (*e.g.*, ADC) for sampling a weak uplink packet [63]. Beyond the analog SIC, *FDLoRa* presents a new digital SIC solution that can effectively remove the SI from normal downlink LoRa packets to achieve in-band full-duplex communication. Unlike the existing digital SIC (*i.e.*, bin suppression [63]) that may damage the modulation space of uplink packets, *FDLoRa* cancels the SI from downlink without compromising (de)modulation performance or Rx sensitivity for uplink.

The basic idea of *FDLoRa* SIC solution can be described as follows: *FDLoRa* leverages the fact that the uplink and downlink in LoRaWAN can be configured with conjugated modulation schemes, *e.g.*, using normal up-chirps as the modulation basis for uplink but down-chirps (*i.e.*, the conjugate of up-chirps) for downlink. While conjugated modulation schemes are used for uplink and downlink, the SI signals from downlink (*i.e.*, down-chirp) now become orthogonal to the uplink signals (*i.e.*, up-chirp). Such signal orthogonality

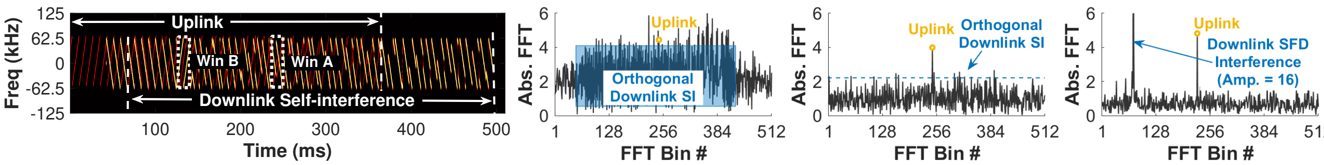


(a) Raw signal after analog SIC (b) Win A (before digital SIC) (c) Win A (after digital SIC) (d) Win B (after digital SIC)

**Figure 8: SIC for in-band full-duplex LoRa communication: (a) Spectra of Rx signal: downlink SI is 15 dB stronger than uplink; (b) Demodulation of an uplink symbol: orthogonal downlink SI overwhelms the uplink; (c) Digital cancellation suppresses downlink SI to noise floor; (d) Downlink SFD still interferes with uplink symbols.**

allows *FDLoRa* to adopt a simple yet effective signal processing method (*e.g.*, dechirping) to efficiently cancel the downlink SI. For instance, as the signal received by the Rx antenna is a superimposition of down-chirps (*i.e.*, downlink SI) and up-chirps (*i.e.*, uplink signal), if we apply the demodulation method for up-chirps (*i.e.*, dechirping followed by FFT) to the superimposed signal as illustrated in Figure 7, down-chirps will not get demodulated because of the orthogonality between up-chirp and down-chirp. It can logically cancel the downlink SI during demodulation and only give symbols of the uplink packet as demodulation results. The (de)modulation performance for uplink will not be compromised when performing SIC.

However, it remains non-trivial to directly apply orthogonal uplink and downlink to an in-band full-duplex radio in practice to achieve digital domain SIC. The reason is that the downlink SI signals received by digital components of the radio, though being largely suppressed by analog SIC, may still be tens of dB stronger than uplink signals, and this large power disparity between uplink and downlink can destroy their signal orthogonality. For example, Figure 8(a) shows the spectrogram of signals received by our full-duplex radio after analog SIC but before digital SIC. We see brighter downlink SI signals than uplink signals (*i.e.*, around 15 dB difference in signal power). When we apply the demodulation method of uplink (*i.e.*, de-chirp with a standard down-chirp, then perform FFT) to the compound signals, we expect the dechirping operation spread the signal power of down-chirps (*i.e.*, downlink SI) but concentrate the energy of up-chirps into FFT bins that give us the demodulated symbols of the uplink. However, Figure 8(b) shows that the energy of the strong downlink SI after spreading into multiple bins is still comparable with (or even higher than) the accumulated energy of the uplink and interferes with the demodulation of the uplink. It requires another digital cancellation procedure before demodulation that can further deduce downlink SI power to protect signal orthogonality between uplink and downlink for correction demodulation of an uplink.

To deal with the practical issues above, *FDLoRa* performs digital SIC in two steps: *FDLoRa* first employs a digital cancellation technique to remove downlink SI from the received compound signals of uplink and downlink, aiming at suppressing SI to the same level of Rx noises. After that, *FDLoRa* exploits the signal orthogonality between uplink and downlink to logically remove any residual downlink SI via a normal demodulation operation of the uplink. The method details are provided in the following.

**Fine-grained digital cancellation.** Digital cancellation uses a digital domain signal processing method to remove the signals of downlink SI from the received raw signals of an in-band full-duplex radio. Without loss of generality, we denote the raw signals ($y(t)$) digitized by a full-duplex radio as below.

$$y(t) = SI(t) + U(t) + n(t) \tag{5}$$

where $SI(t)$ denotes the downlink self-interference signal, $U(t)$ is the signal of an uplink packet, and $n(t)$ denotes Rx noises. The goal of digital cancellation is to remove $SI(t)$ from $y(t)$. Note that $SI(t)$ is a copy of the signal of a downlink packet that is sent out by the Tx antenna, traversing through the air and re-entering the Rx antenna of the same radio. In practice, a full-duplex radio has prior knowledge about the transmitted downlink signals (*i.e.*, denoted by $x(t)$), and thus we can model $SI(t)$ as $SI(t) = h_{tx-rx} \cdot x(t)$, where $h_{tx-rx}$ is the channel from Tx antenna to Rx antenna. To perform effective digital cancellation, *FDLoRa* takes a three-step procedure: (1) use the known signal of downlink SI (*i.e.*, $x(t)$) to estimate channel $h_{tx-rx}$ from $y(t)$; (2) use the estimated channel (*i.e.*, $\hat{h}_{tx-rx}$) to rebuild a received downlink SI as $\hat{SI}(t) = \hat{h}_{tx-rx} \cdot x(t)$; and (3) subtract $\hat{SI}(t)$ from $y(t)$.

To handle potential channel changes over a long duration of a LoRa packet (*e.g.*, hundreds of ms), *FDLoRa* proposes a fine-grained channel estimation method to estimate the channel of downlink SI (*i.e.*, $h_{tx-rx}$) on a per-chirp basis. In particular, as the Tx and Rx of a full-duplex radio share the same clock, the transmitted downlink signals and received SI do not have carrier frequency offsets (CFOs). We can correlate $y(t)$ with the known downlink signal ($x(t)$) to accurately locate the start of downlink SI in received signals. Then, we estimate $h_{tx-rx}$, reconstruct and cancel SI signals chirp-by-chirp. Recall that $SI(t)$ is orthogonal with but stronger than uplink signals $U(t)$ in $y(t)$. We can multiply the $i^{th}$ chirp of $SI(t)$ with the conjugate of the corresponding chirp from $x(t)$ to estimate channel $h_{tx-rx}$ as follows:

$$\begin{aligned} y^i(t) \cdot [x^i(t)]^* &= SI^i(t) \cdot [x^i(t)]^* + U^i(t) \cdot [x^i(t)]^* \\ &\approx h^i_{tx-rx} x^i(t) \cdot [x^i(t)]^* = h^i_{tx-rx} \end{aligned} \tag{6}$$

where $(\cdot)^i$ denotes the $i^{th}$ chirp and $(\cdot)^*$ denotes a complex conjugate. Equation 6 can remove the common down-chirp signal shared by $SI^i(t)$ and $x^i(t)$, spread the uplink chirp $U^i(t)$ to noise levels and produce a tone-frequency of the SI. We can apply FFT to accumulate signal power of the tone-frequency to a high peak from which $\hat{h}^i_{tx-rx}$ can be extracted. After traversing the chirps in $SI(t)$

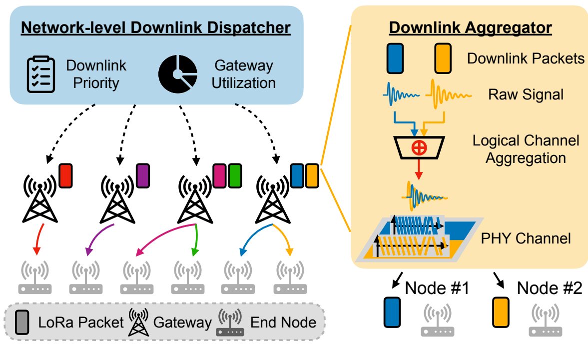


Figure 9: The downlink framework of *FDLoRa*.

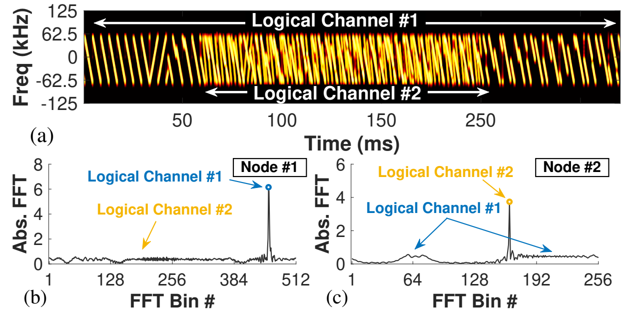


**Figure 10: Concurrent downlink over logical channels: (a) Spectra of aggregated logical channels; (b,c) LoRa node #1 and node #2 can extract and receive over different logical channels from the aggregated PHY signal.**

and $x(t)$, we can use the extracted fine-grained channel information $\hat{h}_{tx-rx} = \{\hat{h}^i_{tx-rx}\}$ to rebuild $\hat{SI}(t)$. Next, we align $\hat{SI}(t)$ with $SI(t)$ and finally subtract it from $y(t)$ for digital cancellation of the SI signals.

Figure 8(c) shows the demodulation results of the same signal from Figure 8(b) after a fine-grained digital cancellation. Comparing the two figures, we can see that the strong downlink SI shown in Figure 8(b) has been successfully suppressed to a noise level.

## 5 *FDLORA* DOWNLINK FRAMEWORK

**Handling downlink SFD interference.** With the above digital cancellation method, we have effectively suppressed the power strength of downlink SI. The residual downlink SI, if any, can next be logically mitigated during the demodulation of the uplink by leveraging signal orthogonality between the uplink and the downlink. However, as a practical downlink LoRa packet also contains 2.25 up-chirps as the Start Frame Delimiter (SFD) of the packet, the SFD chirps are not orthogonal to the data chirps of uplink packets and may interfere with the uplink demodulation when the downlink SI is not completely removed, as shown in Figure 8(d). Our study reveals that although such interference from downlink SFD only impacts the demodulation of 2∼4 symbols of an uplink packet, it can lead to 8% more packet loss on uplink reception (see Figure 15).

*FDLoRa* handles downlink SFD interference from the residual SI signals through two techniques. Firstly, as the transmitted downlink signals are known to a full-duplex radio, *FDLoRa* can precisely detect the start of downlink SI in received signals and track the location of downlink SFD based on the prior knowledge about the frame structure of a downlink packet. With the located downlink SFD information, *FDLoRa* can accurately predict the FFT bins of a downlink SFD chirp in an affected demodulation window of the uplink as shown in Figure 8(d), and thus interpret those bins as SFD interference and extract other FFT peaks as the demodulated symbol of uplink. Secondly, we observe that symbols from the same uplink packet usually have the same FFT amplitude and differ from the amplitude of downlink SFD chirps. As downlink SFD only affects 2∼4 uplink symbols, *FDLoRa* can extract the amplitude of uplink symbols from the demodulation windows not affected by downlink SFD and use the amplitude as a feature to differentiate uplink symbols from downlink SFD in those affected windows. *FDLoRa* jointly employs the two techniques to reliably demodulate symbols of an uplink without compromising Rx sensitivity.

To unlock the potential of full-duplex LoRa gateways, *FDLoRa* introduces a novel downlink framework that utilizes multiple logical channels of all available gateways for concurrent downlink transmissions. We aim to address the limitations of the current LoRa downlink, *i.e.*, each gateway only supports a single downlink channel, and the network server only selects a single gateway for downlink transmission.

To breakthrough the bottleneck, as shown in Figure 9, our downlink framework adopts a two-tier structure to scale concurrent downlink transmissions over multiple logical channels (§5.1) through multiple gateways (§5.2). First, *FDLoRa* includes a *downlink aggregator* that empowers a gateway to serve multiple nodes by transmitting concurrent downlink packets over multiple logical channels. Second, *FDLoRa* employs a *network-level downlink dispatcher*, which coordinates and dispatches concurrent downlink packets through multiple gateways. In the following, we present the design of *FDLoRa* downlink framework in detail.

### 5.1 Concurrent Downlink over Multiple Logical Channels

**Logical channel aggregation.** *FDLoRa* leverages the orthogonality of LoRa logical channels to enable concurrent downlink transmissions from one gateway to multiple nodes. *FDLoRa* develops novel techniques to aggregate multiple logical channels in a unified Tx raw signal while maintaining compatibility with COTS LoRa nodes. This allows *FDLoRa* to send concurrent downlink packets over multiple logical channels to LoRa nodes without any hardware or software extensions to the nodes. As illustrated in Figure 9, *FDLoRa* takes the concurrent downlink packets from the network server as input, and then generates a unified raw signal containing multiple downlink packets. A unique challenge here is that the concurrent downlink packets have diverse demands for delay and power, as they target different nodes with varying Rx windows and link qualities. To address this, *FDLoRa* assigns each packet with unique PHY features before aggregation to meet the delay and power requirements. Formally, we represent the aggregated Tx raw signal as follows:

$$x(t) = \sum_{i=1}^{N} \alpha_i \cdot P_i(t - t_i) \tag{7}$$

where $P_i(t)$ denotes the raw signal of the $i^{th}$ downlink packet, $\alpha_i$ and $t_i$ represent the transmit power and time delay of $P_i(t)$, respectively.

**Reception at COTS LoRa nodes.** Concurrent downlink reception of the aggregated packets over logical channels requires no modification to COTS LoRa nodes. Figure 10(a) illustrates the aggregated downlink packets over two concurrent logical channels targeting at two COTS nodes configured with SF9 and SF8, respectively. The received signal of node #1 can be represented as:

$$y_1(t) = h_1\alpha_1 \underbrace{S(f_{sym1}, k_1, t)}_{\text{Signal of logical ch.\#1}} + h_1\alpha_2 \underbrace{S(f_{sym2}, k_2, t)}_{\text{Signal of logical ch.\#2}} \tag{8}$$

where $0 \le t \le T$ and $h_1$ denotes the channel from the gateway to node #1. Node #1 monitors incoming chirps with SF9 (*i.e.*, slope $k = k_1$) to extract the target chirps from the aggregated downlink signal, with the demodulation results as below (according to Equation 3 and 4).

$$y_1^{demod}(t) = h_1\alpha_1 \underbrace{e^{j2\pi f_{sym1} t}}_{\text{Single tone}} + h_1\alpha_2 \underbrace{e^{j2\pi[\frac{1}{2}(k_2-k_1)t + f_{sym2}]t}}_{\text{Time-varying frequency}} \tag{9}$$

After performing FFT, the signal from logical channel #1 (single tone) is accumulated to a distinct energy peak. In contrast, the signal from logical channel #2 (Time-varying frequency) disperses across the spectrum (see Figure 10(b)). We can observe similar demodulation results for node #2 on logical channel #2 in Figure 10(c). The experiment results show that by leveraging the logical channel orthogonality, COTS LoRa nodes can extract and receive their individual downlink packets from the aggregated downlink signal without any hardware or software modification.

**Downlink quality prediction and power allocation.** While the LoRa nodes can receive their individual downlink packets from the aggregated signal, this approach entails practical challenges. Firstly, transmitting multiple packets would inevitably divide the transmit power of a gateway, potentially affecting the communication range and reliability. Second, signals from concurrent logical channels can spread across the spectrum rather than being perfectly canceled after dechirping as illustrated in Figure 10(b, c). The increased noise floor can negatively affect the reception of the target packet for each LoRa node.

To ensure downlink reliability, we adopt a novel downlink quality prediction method for concurrent logical channels. *FDLoRa* rehearses the aggregated physical layer signal arriving at each target node based on their link quality estimation. Next, it predicts the downlink demodulation performance of LoRa nodes and computes the minimum transmit power required for each downlink packet (*i.e.*, $\{\alpha_i^{min}\}$).

Based on the downlink quality prediction, our goal is to maximize the number of concurrent downlink transmissions while meeting the transmit power requirements as presented in (10) below. $\lambda_i$ equals 1 if the packet $P_i$ can be transmitted, else it equals 0. $A$ denotes the maximum transmit power of the gateway.

$$\begin{aligned} \max \quad & \sum_{i=1}^{N} \lambda_i \\ s.t. \quad & \sum_{i=1}^{N} \lambda_i \alpha_i^{min} \le A, \quad \lambda_i \in \{0, 1\} \end{aligned} \tag{10}$$

This problem can be formulated as a *Knapsack Problem* [31], which is known to be NP-complete. In practice, there are at most 6 logical channels (*i.e.*, SF 7~12 [3]). The maximum Tx power of the LoRa gateway is 20 dBm [69], and we set the minimum power unit as 1 dBm. In practice, we can efficiently solve this problem in $O(NA)$, where $N \le 6$, $A = 20$ dBm, with a *Dynamic Programming* algorithm [81].

## 5.2 Coordinating All Available Gateways

Beyond concurrent downlink transmission of a single gateway, *FDLoRa* also utilizes all available gateways to further enhance downlink performance. However, scaling to more gateways is a non-trivial task. Firstly, dispatching concurrent downlink transmissions to multiple gateways is more complex. Traditional LoRaWAN selects one gateway with the highest SNR for downlink transmission [1]. This approach cannot be applied to schedule more concurrent downlink transmissions of multiple gateways. Moreover, utilizing multiple gateways simultaneously for downlink communications could potentially cause *inter-gateway interference.*

**Dispatching concurrent downlink packets.** *FDLoRa* introduces a *network-level downlink dispatcher* to optimize gateway utilization and maximize network-level downlink concurrency. Measurement studies [75] have shown that due to the long communication range most LoRa nodes can be simultaneously covered by multiple gateways, while a few LoRa nodes can only connect to a single gateway due to signal fading and blockage. To address practical LoRa deployments, our dispatching algorithm is designed as follows. Firstly, we identify the downlink packets whose target nodes can only connect to one gateway, and assign them to the corresponding gateways. Secondly, we arrange the other downlink packets that have multiple optional gateways. We allocate the downlink packets with orthogonal logical channels to the same gateway to maximize the network-level downlink concurrency. Finally, since the second step may lead to unbalanced allocation, we revisit the gateway availability and balance the workload if there are idle gateways. The algorithm traverses the downlink packets two times and can be completed in $O(N)$, which is evaluated in Figure 16(b).

**Resolving inter-gateway interference.** Utilizing multiple gateways simultaneously faces inter-gateway interference problem. A gateway can potentially receive the downlink signal transmitted by other gateways. The uplink packet reception can be influenced if the interference from other gateways unfortunately overlaps with its own uplink. Analog cancellation circuits can help, but it is not feasible as gateways are geographically separated. To mitigate the impact of inter-gateway interference in concurrent downlink transmissions, *FDLoRa* identifies downlink packets that could potentially cause inter-gateway interference (*i.e.*, using the same logical channel) and assigns them to gateways that are as far apart as possible, thereby exploiting the wireless signal attenuation between gateways. As illustrated in Figure 5(a), the wireless signal attenuation

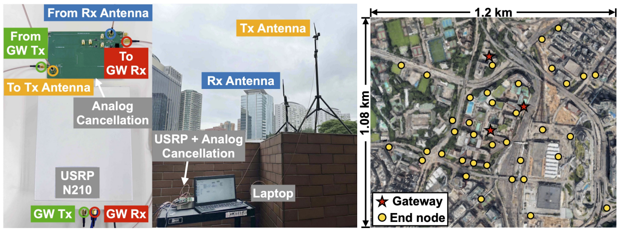


(a) Gateway implementation (b) Campus-scale testbed

**Figure 11: (a) USRP gateway connects with analog cancellation circuit; (b) Deployment map of our testbed.**

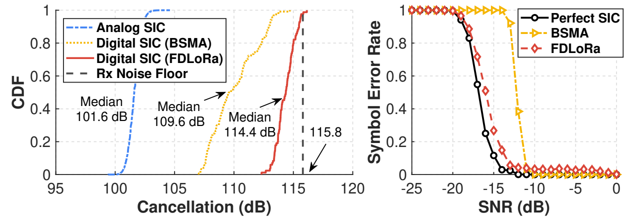


(a) Self-interference cancellation (b) Weak uplink reception

**Figure 12: Performance of full-duplex communication.**

acts as an analog SIC in the air to suppress the inter-gateway interference. *FDLoRa* then applies the digital methods described in §4.2 to further cancel the residual interference signal. Note that such a digital cancellation method only requires the data sharing of concurrent downlink packets and their meta-information (but not raw PHY samples) among neighbor gateways which can be easily handled by the backhaul network between gateways. The requirement of signal attenuation between two gateways is to attenuate the Tx noise of one gateway to the Rx noise floor of the other gateway (see Figure 5(a)). Otherwise, the conflicted packets should be sent at different times.

# 6 EVALUATION

## 6.1 Methodology

**Implementation.** We implement *FDLoRa* on software-defined radios (USRP N210) based on the `gr-lora` project [17]. Each full-duplex LoRa gateway is integrated with an analog cancellation circuit [63] between USRP and Tx/Rx antennas as shown in Figure 11(a). The Tx/Rx antennas are vertically separated for antenna isolation. We use COTS LoRa nodes with Semtech SX1276 radio [52] as transceivers and Arduino Uno boards for node configuration. We build a testbed consisting of 50 LoRa nodes and 3 gateways, where each gateway is connected to a laptop running *FDLoRa*. The gateways are connected to a network server for gateway coordination. We conduct experiments on our campus, covering an area of 1.08 km × 1.2 km in a typical urban environment (see Figure 11(b)). All nodes operate in the 915 MHz ISM band. We use the Arduino-LMIC library [26] to configure the LoRa nodes to apply down-chirp as the modulation basis for downlink packets. The LoRaWAN benchmarks are implemented with RAK7268 COTS LoRa gateways [69].

**Experiment Setup.** We collect data traces for over three months from >200 links in our testbed. The traces cover various channel conditions in urban settings, including indoor and outdoor, low SNR, and high SNR. All evaluations with ≤50 nodes and ≤3 gateways are conducted via real-world experiments in the testbed. In addition, we also perform large-scale trace-driven evaluations based on the data traces collected in the testbed. Our evaluation primarily focuses on the following questions: (1) How does *FDLoRa* perform in full-duplex LoRa communication? (2) How many concurrent downlink transmissions can be supported by *FDLoRa*? and (3) How does *FDLoRa* perform in supporting practical IoT applications?

**Baselines.** We compare *FDLoRa* with four baselines on both full-duplex gateway and downlink transmission. (1) *LoRaWAN* [4] transmits one downlink packet through a single gateway; (2) *MaLoRaGW* [47] adopts two-antenna gateway to transmit concurrent downlinks by beamforming; (3) *HyLink* [82] modulates multiple chirps in the same packet to enhance link throughput; and (4) *BSMA* [63] applies digital cancellation methods designed for full-duplex busy signal transmission.

## 6.2 Basic Performance

**Full-duplex LoRa communication.** We first evaluate *FDLoRa* in supporting full-duplex LoRa communication. We set up one gateway to transmit downlink packets and receive uplink simultaneously. Figure 12(a) presents the full-duplex SIC performance over 20 seconds. We see that antenna isolation and analog cancellation combined provide 101.6 dB SIC. The digital cancellation of BSMA [63] provides additional 8 dB for downlink packets, which is insufficient for receiving weak uplink. *FDLoRa* utilizes a fine-grained digital cancellation method that achieves 12.8 dB and thereby takes the downlink SI to the noise floor.

We further examine the uplink reception sensitivity of *FDLoRa* (below the noise floor) under full-duplex downlink transmission. We set up one LoRa node at SF10 under different SNR conditions and analyzed >100 packets for each SNR condition. The gateway receives uplink while transmitting downlink packets. We also test the uplink-only scenario as the perfect SIC for comparison. As illustrated in Figure 12(b), the perfect SIC performs best as it has no self-interference from downlink. The bin suppression method [63] does not work well because payload symbols of SI have strong side lobe bins that overpower the weak uplink. *FDLoRa* can receive the weak uplink at ultra-low SNR (*e.g.*, −14.5 dB for SF10) by leveraging the signal orthogonality and calibrating the SFD interference. Compared to the uplink-only gateway, the full-duplex communication in *FDLoRa* only leads to 1.1 dB SNR loss on uplink reception while achieving <20% symbol error rate.

**Gains on downlink concurrency.** We next evaluate the performance of *FDLoRa* in supporting concurrent downlink transmissions. We control a varying number of LoRa nodes to receive concurrently on different logical channels. We compare *FDLoRa* against LoRaWAN and MaLoRaGW [47]. LoRaWAN uses three COTS gateways, and MaLoRaGW adopts a two-antenna gateway for downlink transmission, respectively. Both baselines are coordinated by a network server. The nodes use SF9, as the baselines can only transmit downlink packets on a single logical channel. *FDLoRa* contains

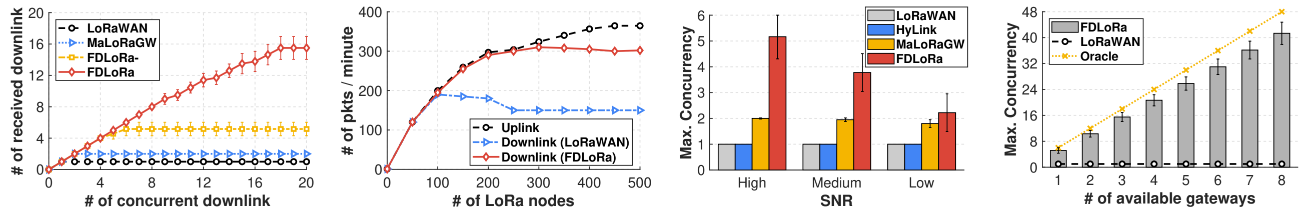


(a) Concurrent downlink transmission (b) Large-scale network (c) Comparison with SOTA (d) Scaling with more gateways

**Figure 13: Performance of downlink scalability: (a) different number of concurrent transmissions; (b) downlink packet rate in large-scale networks; (c) comparison with state-of-the-art; (d) concurrency with more gateways.**

three full-duplex gateways. To fully evaluate the gain of full-duplex communication for unlocking the downlink spectrum and available gateways, we implement a version of *FDLoRa* without full-duplex techniques, which utilizes only one gateway for downlink transmission, represented as *FDLoRa-* in the evaluation. The nodes use SF7∼12, as *FDLoRa* can transmit concurrent downlink packets on multiple logical channels.

Figure 13(a) illustrates the number of received downlink packets under different number of concurrent downlink packets. The number of received downlinks first increases and then reaches the maximum downlink concurrency and stops growing. As expected, LoRaWAN supports at most one downlink packet at a time. MaLoRaGW supports two concurrent downlinks but is still limited by the half-duplex design. In contrast, *FDLoRa-* can transmit an average of >5 concurrent packets over multiple logical channels through a single gateway. The maximum concurrency increases to >15 when unlocking all gateways for downlink with full-duplex communications of *FDLoRa*. The downlink concurrency gain of *FDLoRa* is mainly because of its full-duplex communication capability. By breaking the conflict between uplink and downlink transmissions, it can fully utilize the available gateways and logical channels to enhance downlink concurrency.

**Downlink in large-scale networks.** This experiment uses ACK as a representative downlink application to evaluate *FDLoRa* in large-scale networks. We deploy 50 LoRa nodes in our testbed area to compare *FDLoRa* against LoRaWAN [4]. Both strategies consist of three gateways to cover the LoRa nodes. Each node transmits a *ConfirmedDataUp* message that requires downlink ACK with a duty-cycle ≤1%. We modify the duty-cycle of LoRa nodes to emulate a larger number of nodes (*e.g.*, one node with 5% duty-cycle represents 5 nodes with 1% duty-cycle). The nodes randomly select the channel for each transmission. For a fair comparison, we use the standard LoRa decoder for uplink reception.

Figure 13(b) shows the packet rate for both uplink and downlink. We observe that although the uplink packet rate grows continuously as the number of nodes increases, the downlink packet rate of standard LoRaWAN stops at 100 LoRa nodes and is limited to <200 packets per minute. This significant asymmetry between the uplink and downlink prevents more nodes from receiving effective downlink services. In contrast, *FDLoRa* fills this gap by unlocking downlink logical channels of available gateways. Results show that *FDLoRa* improves the downlink packet rate to 300 packets per minute. It is only when the number of LoRa nodes exceeds 300 that the downlink begins to lag behind the uplink.

**Comparison with state-of-the-art.** This experiment evaluates *FDLoRa* against current leading LoRa downlink methods. We examine two state-of-the-art (SOTA) strategies (*i.e.*, HyLink [82] and MaLoRaGW [47]). To ensure a fair comparison, we focus on the maximum number of concurrent downlink transmissions each gateway can support. For HyLink [82], we implement both the transmitter and receiver with USRPs as it has to access the PHY raw signal. For MaLoRaGW [47], we employ two synchronized USRPs to facilitate 2×2 MIMO beamforming for the transmitter, while the reception is handled by COTS LoRa nodes. For *FDLoRa*, we aggregate multiple downlink logical channels and transmit with USRP, with COTS LoRa nodes serving as receivers. A standard LoRaWAN setup is also implemented for baseline comparison. The baselines use SF9 since they can only transmit downlink on a single logical channel. *FDLoRa* uses SF7∼12. We test three SNR scenarios: high (>5 dB), medium (-5∼5 dB), and low (<-5 dB), with the average communication range of <30 m, ∼100 m and >200 m.

Figure 13(c) displays the maximum downlink concurrency per gateway achieved by four strategies at different SNR conditions. As expected, LoRaWAN is limited to a single downlink transmission. HyLink enhances link throughput by encoding parallel chirps within a single packet payload, yet it fails to accommodate multiple clients simultaneously. MaLoRaGW achieves up to 2× concurrent transmissions at high SNR, with a slight reduction as SNR worsens. The reason is that the concurrency of multi-user MIMO is limited by the number of Tx antennas and the precision of channel estimation. *FDLoRa* outperforms the other strategies, achieving on average 5× and 2.2× concurrent transmissions at high and low SNR conditions, respectively. This superior performance is because *FDLoRa* exploits logical channel orthogonality of LoRa to enhance downlink capabilities, which is feasible with the gateway's single Tx chain and antenna. We notice that the downlink concurrency drops as the communication range increases. The reason is that low-SNR links require higher transmit power, but the budget is restricted by the maximum transmit power. Since the low-SF packets require higher transmit power to reach the same communication range as the high-SF ones, *FDLoRa* prioritizes the high-SF logical channels (*e.g.*, SF10∼12) to achieve higher concurrency for low-SNR links.

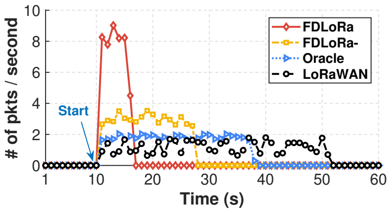


**Figure 14: Practical application with bursty downlink traffic.**

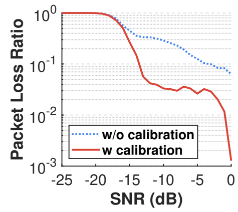


**Figure 15: Downlink SFD interference.**

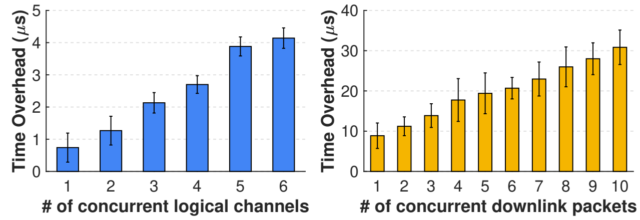


(a) Power allocation (b) Dispatch concurrent downlink

**Figure 16: Time overhead of downlink framework.**

**Scaling with more gateways.** This experiment assesses the maximum number of concurrent downlink transmissions that *FDLoRa* can support with an increased number of gateways. We increment the number of gateways from 1 to 8 and configure various LoRa nodes to receive on different logical channels. We employ trace-driven evaluation to conduct experiments involving more than three gateways. To achieve this, we aggregate multiple downlink packet traces with random time offsets to simulate real-world network traffic. These synthesized traces are then replayed to assess the performance of *FDLoRa* in scenarios with over three gateways. We compare *FDLoRa* with two benchmarks: (1) LoRaWAN, which transmits one downlink packet through a single gateway, and (2) Oracle, which explicitly receives all the downlink logical channels.

We measure the maximum number of concurrent downlink transmissions and present the results in Figure 13(d). The maximum concurrency of LoRaWAN is limited to 1 due to the resource conflict of half-duplex gateways, which restricts the allocation of additional gateways for downlink transmissions. However, the number of available downlink logical channels (*i.e.*, Oracle) continues to increase. *FDLoRa* closely aligns with Oracle, *i.e.*, within 86%. Specifically, when the number of available gateways increases to 8, *FDLoRa* supports 41 concurrent downlink transmissions. The results indicate that *FDLoRa* can effectively utilize all available gateways to enhance downlink transmission.

**Practical application with bursty downlink traffic.** This experiment evaluates *FDLoRa* within a smart lamppost system [16] that experiences bursty downlink traffic. The objective is to switch on/off 50 lampposts in the least amount of time. We deploy 50 LoRa nodes for device control and transmit commands via three gateways. We compare *FDLoRa* against two benchmarks: (1) LoRaWAN: COTS gateways with a LoRaWAN network server. (2) Oracle: LoRaWAN with an Oracle scheduler that explicitly schedules downlinks to achieve optimal time slot utilization.

We record the downlink packet rate obtained in each time slot and display the results in Figure 14. The transmission starts at the 10-second mark. We see that LoRaWAN requires >40 seconds to complete the task. Even when enhanced with an Oracle scheduler for optimal time slot utilization, LoRaWAN still requires >27 seconds, with a limited downlink packet rate of <1.9 packets/s. In contrast, *FDLoRa-* improves the rate to 3 packets/s by enabling concurrent logical channels through a single gateway. The full-duplex communication of *FDLoRa* further boosts the rate to approximately 8 packets/s by unlocking multiple gateways for concurrent transmission, which is 5.7× higher than Oracle LoRaWAN. *FDLoRa-* and *FDLoRa* take <17 and <6 seconds to accomplish this task, respectively.

## 6.3 Microbenchmarks

**Impact of downlink SFD interference.** This experiment evaluates the impact of downlink SFD on weak uplink reception. Specifically, we deploy a LoRa node at different SNR conditions and measure the packet loss ratio of *FDLoRa* without downlink SFD calibration. The packets are considered lost when the symbol error rate is more than 20%. We also implement a full version of *FDLoRa* for comparison.

Figure 15 shows the packet loss ratio under different SNR conditions. We see that the downlink SFD (including 2.25 interference chirps) leads to about 8 % more packet loss. The reason is that the strong SFD interference introduces more symbol errors for weak uplink reception. *FDLoRa* effectively mitigates this issue by calibrating downlink SFD interference during demodulation.

**Time overhead.** This experiment investigates the time overhead of *FDLoRa* for logical channel power allocation and dispatching concurrent downlink packets. We first test the power allocation algorithm for concurrent logical channels with a time complexity of $O(NA)$. This algorithm is executed on a laptop. We measure the runtime under different number of concurrent logical channels, specifically for $N = 1 \sim 6$. As shown in Figure 16(a), the experimental overhead observed is slightly higher than linear $O(N)$. Next, we evaluate the network-level downlink dispatcher. We also measure the runtime on a laptop under different number of concurrent downlink packets. The results demonstrate a linear $O(N)$ overhead as depicted in Figure 16(b). Notably, the overheads of the two algorithms under normal downlink traffic are below 5 and 40 $\mu$s, respectively. This duration is much shorter than the airtime of LoRa packets (*e.g.*, hundreds of ms) and can be easily handled by the gateway/server processors.

# 7 RELATED WORK

Table 1 contrasts *FDLoRa* with the state-of-the-art.

**Full-duplex radio.** Self-interference cancellation for full-duplex radios employs a combination of antenna isolation, analog cancellation, and digital cancellation [5–7, 12, 14, 32, 67]. Many studies [5, 6] focused on full-duplex communication for WiFi systems. Regarding full-duplex LoRa gateways, BSMA [63] has developed analog cancellation circuits and a bin suppression technique to enable full-duplex downlink busy signal transmission for channel access control. However, the transmit signal in [63] is a pre-defined busy signal consisting of only base up-chirps. Its digital domain solution faces limitations when applied to standard downlink packets (*e.g.*, data transfer, ACK). To support full-duplex uplink and downlink communications, *FDLoRa* borrows from the analog cancellation

| System Name | Full Duplex | # of gateway | Concur. per gateway | Method. |
|---|---|---|---|---|
| LoRaWAN [4] | ✗ | 1 | 1 | Single logical ch. & single gateway |
| HyLink [82] | ✗ | 1 | 1 | Parallel-CSS |
| MaLoRaGW [47] | ✗ | 1 | 2 | MU-MIMO (2 Tx) |
| BSMA [63] | Busy Signal | 1 | 1 | Media access control (for uplink) |
| ***FDLoRa*** | Uplink & Downlink | All | **5.16** | Full-duplex comm. & Multi. logical ch. & Multi. gateways |

**Table 1: Comparisons between *FDLoRa* and SoTA.**

circuits [63] and introduces innovative digital SIC methods that outperform existing approaches.

**LoRa downlink.** The newly emerging IoT applications are calling for robust LoRa downlink to support efficient device management, real-time response, and high throughput data transfer. LoRaWAN [4] sequentially transmits the downlink packets and thus suffers from significant downlink-uplink asymmetry. A recent work [1] proposes a downlink gateway selection strategy but is still restricted by the sequential transmit methodology. HyLink [82] enhances link throughput by modulating parallel chirps but can only serve a single client. MaLoRaGW [47] requires a multi-antenna gateway to support concurrent downlink via beamforming. However, existing works either limit to one client [82] or narrowly focus on concurrency within a single logical channel [47]. Built on top of the full-duplex gateway, *FDLoRa* unlocks the available spectrum and resources for downlink transmission, enabling concurrent downlink transmission through multiple logical channels per gateway and across more gateways.

**Concurrent LoRa transmission.** Many previous works have explored concurrent LoRa transmissions using parallel decoding [13, 30, 35, 38, 45, 58, 74, 76, 79, 83, 84, 86, 89]. Choir [13] utilizes the hardware imperfection of LoRa devices to classify symbols based on their unique frequency offsets. FTrack [84] exploits the time misalignment of symbol edges to separate collisions. FlipLoRa [87] employs the orthogonality between up- and down-chirps to resolve collisions. CIC [58] leverages the feature of sub-symbols to decode from the multi-packet collision. These works aim to improve uplink concurrency. Unlike all these works, *FDLoRa* develops a new in-band full-duplex LoRa gateways and supports concurrent downlink transmissions to tackle downlink-uplink asymmetry.

## 8 LIMITATIONS AND FUTURE WORK

**Overheads of *FDLoRa*.** The key advantage of *FDLoRa* is that it fully unlocks the spectrum and gateway utilization for downlink with full-duplex gateways. By leveraging uplink-downlink orthogonality, *FDLoRa* implements a novel digital SIC method that outperforms prior works. Importantly, the (de)modulation of downlink packets based on down-chirp is supported by COTS LoRa nodes [26, 52] and gateways [69], allowing *FDLoRa* to operate with COTS LoRa end devices without requiring any modifications. Despite these benefits, *FDLoRa* introduces higher computational and energy overheads at the gateway and server sides, associated with additional analog circuits, digital SIC methods, and the management of downlink packets and logical channels. A comprehensive evaluation of these overheads will be carried out in future work. Additionally, *FDLoRa* experiences slight SNR losses ($\sim 1.1$ dB) as a result of full-duplex communication. Nonetheless, we believe these trade-offs are justified by the significant improvements in downlink scalability and overall system performance.

**Achieving true downlink-uplink symmetry.** The goal of *FDLoRa* is to address the downlink-uplink asymmetry in LoRaWAN. *FDLoRa* mitigates this gap by scaling downlink communications from one to multiple logical channels. However, the gap is not entirely closed due to the hardware constraints of commodity LoRa gateways, which utilize 9 Rx chains for uplink but only one Tx chain for downlink. This physical asymmetry presents fundamental challenges for achieving true downlink-uplink symmetry in LoRaWAN. We will explore additional possibilities to enhance downlink communications for newly emerging IoT applications using LoRa.

**Further enhancing downlink communication.** Building on *FDLoRa*, our future work will focus on developing software-based gateway solutions and new downlink protocols. We will also conduct more comprehensive evaluations of *FDLoRa*, including fine-grained and quantitative system overhead analysis, real-time performance enhancements, multi-gateway deployment issues, *etc.*

## 9 CONCLUSION

This paper presents *FDLoRa*, a collaborative design of LoRa gateway and downlink framework that revolutionizes the sequential downlink paradigm in existing LoRaWAN. Unlike traditional LoRaWAN that utilizes half-duplex gateways, *FDLoRa* presents a new in-band full-duplex gateway design. To harness the capabilities of full-duplex gateways, *FDLoRa* develops a novel downlink framework that scales concurrent downlink packets to massive logical channels of all available gateways. This breakthrough addresses the downlink-uplink asymmetry in current LoRaWAN and accommodates the emerging demands of new applications on increasing downlink transmissions.

## ACKNOWLEDGMENTS

We sincerely thank the anonymous shepherd and reviewers for their comments and feedback. This work is supported in part by HK GRF (Grant No. 15218022, 15211924 and 15206123), NSFC (Grant No. 62102336), and the Innovation Capability Support Program of Shaanxi (No. 2023-CX-TD-08) Shaanxi Qinchuangyuan "scientists+engineers" team (No. 2023KXJ-040). Xianjin Xia and Yuanqing Zheng are the corresponding authors.

## REFERENCES

[1] Samira Abboud, Nancy El Rachkidy, Alexandre Guitton, and Haidar Safa. 2019. Gateway selection for downlink communication in LoRaWAN. In *2019 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 1–6.
[2] Absar-Ul-Haque Ahmar, Emekcan Aras, Thien Duc Nguyen, Sam Michiels, Wouter Joosen, and Danny Hughes. 2023. Design of a Robust MAC Protocol for LoRa. *ACM Trans. Internet Things* 4, 1, Article 3 (feb 2023), 25 pages. https://doi.org/10.1145/3557048
[3] LoRa Alliance. 2021. *RP2-1.0.3 LoRaWAN Regional Parameters*. "https://lora-alliance.org/resource_hub/rp2-1-0-3-lorawan-regional-parameters".
[4] LoRa Alliance. 2022. *Lorawan Specification*. "https://lora-alliance.org/about-lorawan".
[5] Dinesh Bharadia and Sachin Katti. 2014. Full duplex {MIMO} radios. In *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*. 359–372.
[6] Dinesh Bharadia, Emily McMilin, and Sachin Katti. 2013. Full duplex radios. In *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*. 375–386.

[7] Tingjun Chen, Mahmood Baraani Dastjerdi, Jin Zhou, Harish Krishnaswamy, and Gil Zussman. 2019. Wideband full-duplex wireless via frequency-domain equalization: Design and experimentation. In *The 25th annual international conference on mobile computing and networking*. 1–16.

[8] Yuning Chen, Kang Yang, Zhiyu An, Brady Holder, Luke Paloutzian, Khaled Bali, and Wan Du. 2024. MARLP: Time-series Forecasting Control for Agricultural Managed Aquifer Recharge. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'24)*.

[9] COMMSCOPE. 2024. *Antenna Vertical Isolation Calculator.* "https://calc.commscope.com/qvisolation.aspx".

[10] Kaiyan Cui, Qiang Yang, Leming Shen, Yuanqing Zheng, Fu Xiao, and Jinsong Han. 2024. Towards ISAC-Empowered mmWave Radars by Capturing Modulated Vibrations. *IEEE Transactions on Mobile Computing* (2024).

[11] Jialuo Du, Yidong Ren, Zhui Zhu, Chenning Li, Zhichao Cao, Qiang Ma, and Yunhao Liu. 2023. SRLoRa: Neural-enhanced LoRa Weak Signal Decoding with Multi-gateway Super Resolution. In *Proceedings of the Twenty-fourth International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing*. 270–279.

[12] Melissa Duarte, Ashutosh Sabharwal, Vaneet Aggarwal, Rittwik Jana, Kadangode K Ramakrishnan, Christopher W Rice, and NK Shankaranarayanan. 2013. Design and characterization of a full-duplex multiantenna system for WiFi networks. *IEEE Transactions on Vehicular Technology* 63, 3 (2013), 1160–1177.

[13] Rashad Eletreby, Diana Zhang, Swarun 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 (SIGCOMM '17)*. 309–321.

[14] Evan Everett, Achaleshwar Sahai, and Ashutosh Sabharwal. 2014. Passive self-interference suppression for full-duplex infrastructure nodes. *IEEE Transactions on Wireless Communications* 13, 2 (2014), 680–694.

[15] 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*. Association for Computing Machinery, New York, NY, USA, Article 43, 13 pages.

[16] Anita Gehlot, Sultan S Alshamrani, Rajesh Singh, Mamoon Rashid, Shaik Vaseem Akram, Ahmed Saeed AlGhamdi, and Fahad R Albogamy. 2021. Internet of things and long-range-based smart lampposts for illuminating smart cities. *Sustainability* 13, 11 (2021), 6398.

[17] Gr-LoRa GitHub community. 2021. *gr-lora projects.* "https://github.com/rpp0/gr-lora".

[18] Chaojie Gu, Linshan Jiang, Rui Tan, Mo Li, and Jun Huang. 2021. Attack-aware synchronization-free data timestamping in lorawan. *ACM Transactions on Sensor Networks (TOSN)* 18, 1 (2021), 1–31.

[19] Dongfang Guo, Chaojie Gu, Linshan Jiang, Wenjie Luo, and Rui Tan. 2022. Illoc: In-hall localization with standard lorawan uplink frames. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 1 (2022), 1–26.

[20] Xiuzhen Guo, Longfei Shangguan, Yuan He, Nan Jing, Jiacheng Zhang, Haotian Jiang, and Yunhao Liu. 2022. Saiyan: Design and implementation of a low-power demodulator for LoRa backscatter systems. In *Proc. USENIX NSDI*. 437–451.

[21] Xiuzhen Guo, Longfei Shangguan, Yuan He, Jia Zhang, Haotian Jiang, Awais Ahmad Siddiqi, and Yunhao Liu. 2020. Aloba: rethinking ON-OFF keying modulation for ambient LoRa backscatter. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*. 192–204.

[22] Yohei Hasegawa and Kazuya Suzuki. 2019. A multi-user ack-aggregation method for large-scale reliable lorawan service. In *ICC 2019-2019 IEEE International Conference on Communications (ICC)*. IEEE, 1–7.

[23] Ningning Hou, Xianjin Xia, Yifeng Wang, and Yuanqing Zheng. 2024. One shot for all: Quick and accurate data aggregation for LPWANs. *IEEE/ACM Transactions on Networking* (2024).

[24] Ningning Hou, Xianjin Xia, and Yuanqing Zheng. 2022. Cloaklora: A covert channel over lora phy. *IEEE/ACM Transactions on Networking* (2022).

[25] Ningning Hou, Xianjin Xia, and Yuanqing Zheng. 2023. Don't Miss Weak Packets: Boosting LoRa Reception with Antenna Diversities. *ACM Trans. Sen. Netw.* 19, 2, Article 41 (feb 2023), 25 pages. https://doi.org/10.1145/3563698

[26] IBM. 2022. *MCCI LoRaWAN LMIC Library.* "https://github.com/mcci-catena/arduino-lmic".

[27] The Things Industries. 2024. *Duty Cycle.* "https://www.thethingsnetwork.org/docs/lorawan/duty-cycle".

[28] Wha Sook Jeon and Dong Geun Jeong. 2020. Adaptive uplink rate control for confirmed class A transmission in LoRa networks. *IEEE Internet of Things Journal* 7, 10 (2020), 10361–10374.

[29] Jinyan Jiang, Jiliang Wang, Yijie Chen, Yihao Liu, and Yunhao Liu. 2023. Locra: Enable practical long-range backscatter localization for low-cost tags. In *Proceedings of the 21st Annual International Conference on Mobile Systems, Applications and Services*. 317–329.

[30] Jinyan Jiang, Zhenqiang Xu, Fan Dang, and Jiliang Wang. 2021. Long-range ambient LoRa backscatter with parallel decoding. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*. 684–696.

[31] Hans Kellerer, Ulrich Pferschy, and David Pisinger. 2004. Knapsack problems. https://api.semanticscholar.org/CorpusID:28836720

[32] Dani Korpi, Joose Tamminen, Matias Turunen, Timo Huusari, Yang-Seok Choi, Lauri Anttila, Shilpa Talwar, and Mikko Valkama. 2016. Full-duplex mobile device: Pushing the limits. *IEEE Communications Magazine* 54, 9 (2016), 80–87.

[33] Rachel Kufakunesu, Gerhard P Hancke, and Adnan M Abu-Mahfouz. 2020. A survey on adaptive data rate optimization in lorawan: Recent solutions and major challenges. *Sensors* 20, 18 (2020), 5044.

[34] 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.

[35] Chenning Li, Zhichao Cao, and Li Xiao. 2022. Curvealoha: Non-linear chirps enabled high throughput random channel access for lora. In *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*. IEEE, 520–529.

[36] Chenning Li, Hanqing Guo, Shuai Tong, Zhichao Cao, Mi Zhang, Qiben Yang, Li Xiao, Jiliang Wang, and Yunhao Liu. 2022. Nelora: Neural-enhanced demodulation for low-power wans. *GetMobile: Mobile Computing and Communications* 26, 3 (2022), 34–38.

[37] 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* (Coimbra, Portugal) *(SenSys '21)*. Association for Computing Machinery, New York, NY, USA, 56–68. https://doi.org/10.1145/3485730.3485928

[38] Chenning Li, Xiuzhen Guo, Longfei Shangguan, Zhichao Cao, and Kyle Jamieson. 2022. CurvingLoRa to Boost LoRa Network Throughput via Concurrent Transmission. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*. USENIX Association, Renton, WA, 879–895.

[39] Chenning Li, Yidong Ren, Shuai Tong, Shakhrul Iman Siam, Mi Zhang, Jiliang Wang, Yunhao Liu, and Zhichao Cao. 2024. ChirpTransformer: Versatile LoRa Encoding for Low-power Wide-area IoT. In *Proceedings of the 22nd Annual International Conference on Mobile Systems, Applications and Services*. 479–491.

[40] Yinghui Li, Jing Yang, and Jiliang Wang. 2020. DyLoRa: Towards Energy Efficient Dynamic LoRa Transmission Control. In *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications*. 2312–2320. https://doi.org/10.1109/INFOCOM41043.2020.9155407

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

[42] RAKwireless Technology Limited. 2024. *Revolutionizing High-Density Connectivity with the WisGate Edge Ultra, a Full-Duplex Gateway for LoRaWAN.* "https://news.rakwireless.com/high-density-connectivity-with-the-wisgate-edge-ultra-a-full-duplex-lorawan-gateway/".

[43] Li Liu, Yuguang Yao, Zhichao Cao, and Mi Zhang. 2021. Deeplora: Learning accurate path loss model for long distance links in lpwan. In *IEEE INFOCOM 2021-IEEE Conference on Computer Communications*. IEEE, 1–10.

[44] Junzhou Luo, Zhuqing Xu, Jingkai Lin, Ciyuan Chen, and Runqun Xiong. 2023. CH-MAC: Achieving Low-latency Reliable Communication via Coding and Hopping in LPWAN. *ACM Trans. Internet Things* 4, 4, Article 24 (nov 2023), 25 pages. https://doi.org/10.1145/3617505

[45] Manan Mishra, Daniel Koch, Muhammad Osama Shahid, Bhuvana Krishnaswamy, Krishna Chintalapudi, and Suman Banerjee. 2023. {OpenLoRa}: Validating {LoRa} Implementations through an Extensible and Open-sourced Framework. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*. 1165–1183.

[46] Yao Peng, Longfei Shangguan, Yue Hu, Yujie Qian, Xianshang Lin, Xiaojiang Chen, Dingyi Fang, and Kyle Jamieson. 2018. PLoRa: A passive long-range data network from ambient LoRa transmissions. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*. 147–160.

[47] Hossein Pirayesh, Shichen Zhang, Pedram Kheirkhah Sangdeh, and Huacheng Zeng. 2022. MaLoRaGW: Multi-User MIMO Transmission for LoRa. In *Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems*. 179–192.

[48] Yidong Ren, Puyu Cai, Jinyan Jiang, Jialuo Du, and Zhichao Cao. 2023. Prism: High-throughput LoRa backscatter with non-linear chirps. In *IEEE INFOCOM 2023-IEEE Conference on Computer Communications*. IEEE, 1–10.

[49] Yidong Ren, Amalinda Gamage, Li Liu, Mo Li, Shigang Chen, Younsuk Dong, and Zhichao Cao. 2024. SateRIoT: High-performance Ground-Space Networking for Rural IoT. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*.

[50] Yidong Ren, Li Liu, Chenning Li, Zhichao Cao, and Shigang Chen. 2022. Is LoRaWAN Really Wide? Fine-grained LoRa Link-level Measurement in An Urban Environment. In *2022 IEEE 30th International Conference on Network Protocols (ICNP)*. 1–12. https://doi.org/10.1109/ICNP55882.2022.9940375

[51] Yidong Ren, Wei Sun, Jialuo Du, Huaili Zeng, Younsuk Dong, Mi Zhang, Shigang Chen, Yunhao Liu, Tianxing Li, and Zhichao Cao. 2024. Demeter: Reliable Cross-soil LPWAN with Low-cost Signal Polarization Alignment. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*. 230–245.

[52] Semtech. 2022. *SX1276.* "https://www.semtech.com/products/wireless-rf/lora-connect/sx1276".
[53] Semtech. 2022. *SX1301.* "https://www.semtech.com/products/wireless-rf/lora-core/sx1301".
[54] Semtech. 2023. *LoRa.* "https://www.semtech.com/lora".
[55] Semtech. 2024. *Understanding ADR.* "https://lora-developers.semtech.com/documentation/tech-papers-and-guides/understanding-adr".
[56] Semtech. 2024. *Understanding LoRaWAN Architecture.* "https://blog.semtech.com/understanding-the-lorawan-architecture".
[57] Amazon Web Services. 2024. *Queue downlink messages to send to LoRaWAN devices.* "https://docs.aws.amazon.com/iot/latest/developerguide/connect-iot-lorawan-downlink-queue.html".
[58] Muhammad Osama Shahid, Millan Philipose, Krishna Chintalapudi, Suman Banerjee, and Bhuvana Krishnaswamy. 2021. Concurrent Interference Cancellation: Decoding Multi-Packet Collisions in LoRa *(SIGCOMM '21)*. Association for Computing Machinery, New York, NY, USA.
[59] Chenglong Shao and Osamu Muta. 2024. TONARI: Reactive Detection of Close Physical Contact Using Unlicensed LPWAN Signals. *ACM Trans. Internet Things* 5, 2, Article 13 (apr 2024), 30 pages. https://doi.org/10.1145/3648572
[60] Leming Shen, Qiang Yang, Kaiyan Cui, Yuanqing Zheng, Xiao-Yong Wei, Jianwei Liu, and Jinsong Han. 2024. FedConv: A Learning-on-Model Paradigm for Heterogeneous Federated Clients. In *Proceedings of the 22nd Annual International Conference on Mobile Systems, Applications and Services.* 398–411.
[61] Yihang Song, Li Lu, Jiliang Wang, Chong Zhang, Hui Zheng, Shen Yang, Jingsong Han, and Jian Li. 2023. $\mu$Mote: Enabling Passive Chirp De-spreading and $\mu$W-level Long-Range Downlink for Backscatter Devices. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*. USENIX Association.
[62] Nicolas Sornin. 2020. LoRaWAN®: Firmware Updates Over-the-Air. *Semtech Technical Paper* 200 (2020), 93012–8790.
[63] Raghav Subbaraman, Yeswanth Guntupalli, Shruti Jain, Rohit Kumar, Krishna Chintalapudi, and Dinesh Bharadia. 2022. BSMA: scalable LoRa networks using full duplex gateways. In *Proceedings of the 28th Annual International Conference on Mobile Computing And Networking* (Sydney, NSW, Australia) *(MobiCom '22)*. Association for Computing Machinery, New York, NY, USA, 676–689. https://doi.org/10.1145/3495243.3560544
[64] Zehua Sun, Tao Ni, Huanqi Yang, Kai Liu, Yu Zhang, Tao Gu, and Weitao Xu. 2023. FLoRa: Energy-Efficient, Reliable, and Beamforming-Assisted Over-The-Air Firmware Update in LoRa Networks. In *Proceedings of the 22nd International Conference on Information Processing in Sensor Networks.* 14–26.
[65] Zehua Sun, Huanqi Yang, Kai Liu, Zhimeng Yin, Zhenjiang Li, and Weitao Xu. 2022. Recent advances in LoRa: A comprehensive survey. *ACM Transactions on Sensor Networks* 18, 4 (2022), 1–44.
[66] 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.
[67] Joose Tamminen, Matias Turunen, Dani Korpi, Timo Huusari, Yang-Seok Choi, Shilpa Talwar, and Mikko Valkama. 2016. Digitally-controlled RF self-interference canceller for full-duplex radios. In *2016 24th European Signal Processing Conference (EUSIPCO)*. IEEE, 783–787.
[68] Dragino Technology. 2024. *lps8n.* "https://www.dragino.com/products/lora-lorawan-gateway/item/200-lps8n.html".
[69] RAKwireless Technology. 2024. *RAK7268.* "https://docs.rakwireless.com/Product-Categories/WisGate/RAK7268/Overview".
[70] RAKwireless Technology. 2024. *RAK7285.* "https://docs.rakwireless.com/Product-Categories/WisGate/RAK7285/Overview".
[71] Xu Tian, Joseph Sarkis, Wei Chen, Yong Geng, Haozhi Pan, Zuoxi Liu, and Sergio Ulgiati. 2024. Greening the Belt and Road Initiative: Evidence from emergy evaluation of China's provincial trade with ASEAN countries. *Fundamental Research* 4, 2 (2024), 379–393. https://doi.org/10.1016/j.fmre.2022.11.007
[72] Shuai Tong, Yangliang He, Yunhao Liu, and Jiliang Wang. 2022. De-spreading over the air: long-range ctc for diverse receivers with lora. In *Proceedings of the 28th Annual International Conference on Mobile Computing and Networking.* 42–54.
[73] 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* (New Orleans, Louisiana) *(MobiCom '21)*. Association for Computing Machinery, New York, NY, USA, 642–655. https://doi.org/10.1145/3447993.3483250
[74] 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.* ACM, Toronto Ontario Canada, 234–246.
[75] Shuai Tong, Jiliang Wang, Jing Yang, Yunhao Liu, and Jun Zhang. 2023. City-wide LoRa Network Deployment and Operation: Measurements, Analysis, and Implications. In *Proceedings of the 21st ACM Conference on Embedded Networked Sensor Systems.* 362–375.
[76] Shuai Tong, Zhenqiang Xu, and Jiliang Wang. 2020. CoLoRa: Enabling Multi-Packet Reception in LoRa. In *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications.* 2303–2311. ISSN: 2641-9874.
[77] Roman Trüb, Reto Da Forno, Andreas Biri, Jan Beutel, and Lothar Thiele. 2023. LSR: Energy-Efficient Multi-Modulation Communication for Inhomogeneous Wireless IoT Networks. *ACM Trans. Internet Things* 4, 2, Article 10 (apr 2023), 36 pages. https://doi.org/10.1145/3579366
[78] Vutha Va and Robert W Heath. 2015. Basic relationship between channel coherence time and beamwidth in vehicular channels. In *2015 IEEE 82nd Vehicular Technology Conference (VTC2015-Fall)*. IEEE, 1–5.
[79] Xiong Wang, Linghe Kong, Liang He, and Guihai Chen. 2019. mlora: A multi-packet reception protocol in lora networks. In *2019 IEEE 27th International Conference on Network Protocols (ICNP)*. IEEE, 1–11.
[80] Zhe Wang, Linghe Kong, Longfei Shangguan, Liang He, Kangjie Xu, Yifeng Cao, Hui Yu, Qiao Xiang, Jiadi Yu, Teng Ma, et al. 2023. LigBee: Symbol-Level Cross-Technology Communication from LoRa to ZigBee. In *IEEE INFOCOM 2023-IEEE Conference on Computer Communications.* IEEE, 1–10.
[81] Wikipedia. 2024. Knapsack problem — Wikipedia, The Free Encyclopedia. http://en.wikipedia.org/w/index.php?title=Knapsack%20problem&oldid=1223595758. [Online; accessed 28-May-2024].
[82] Xianjin Xia, Qianwu Chen, Ningning Hou, and Yuanqing Zheng. 2022. Hylink: Towards high throughput lpwans with lora compatible communication. In *Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems.* 578–591.
[83] 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.
[84] Xianjin Xia, Yuanqing Zheng, and Tao Gu. 2019. FTrack: parallel decoding for LoRa transmissions. In *Proceedings of the 17th Conference on Embedded Networked Sensor Systems.* ACM, New York New York, 192–204.
[85] Xianjin Xia, Yuanqing Zheng, and Tao Gu. 2021. Litenap: Downclocking lora reception. *IEEE/ACM Transactions on Networking* 29, 6 (2021), 2632–2645.
[86] Pengjin Xie, Yinghui Li, Zhenqiang Xu, Qian Chen, Yunhao Liu, and Jiliang Wang. 2023. Push the Limit of LPWANs with Concurrent Transmissions. In *IEEE INFOCOM 2023-IEEE Conference on Computer Communications.* IEEE, 1–10.
[87] Zhenqiang Xu, Shuai Tong, Pengjin Xie, and Jiliang Wang. 2020. FlipLoRa: Resolving collisions with up-down quasi-orthogonality. In *2020 17th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*. IEEE, 1–9.
[88] Zhenqiang Xu, Shuai Tong, Pengjin Xie, and Jiliang Wang. 2023. From demodulation to decoding: Toward complete LoRa PHY understanding and implementation. *ACM Transactions on Sensor Networks* 18, 4 (2023), 1–27.
[89] 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.
[90] Kang Yang, Yuning Chen, Xuanren Chen, and Wan Du. 2023. Link Quality Modeling for LoRa Networks in Orchards. In *Proceedings of the 22nd International Conference on Information Processing in Sensor Networks* (San Antonio, TX, USA) *(IPSN '23)*. Association for Computing Machinery, New York, NY, USA, 27–39. https://doi.org/10.1145/3583120.3586969
[91] Kang Yang, Yuning Chen, and Wan Du. 2024. OrchLoc: In-orchard localization via a single LoRa gateway and generative diffusion model-based fingerprinting. In *Proceedings of the 22nd ACM International Conference on Mobile Systems, Applications, and Services (MobiSys' 24)*.
[92] Kang Yang and Wan Du. 2022. LLDPC: A low-density parity-check coding scheme for LoRa networks. In *Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems.* 193–206.
[93] Kang Yang and Wan Du. 2024. A Low-Density Parity-Check Coding Scheme for LoRa Networking. *ACM Transactions on Sensor Networks* (2024).
[94] Qiang Yang and Yuanqing Zheng. 2023. Aquahelper: Underwater sos transmission and detection in swimming pools. In *Proceedings of the 21st ACM Conference on Embedded Networked Sensor Systems.* 294–307.
[95] Fu Yu, Xiaolong Zheng, Liang Liu, and Huadong Ma. 2022. LoRadar: An Efficient LoRa Channel Occupancy Acquirer based on Cross-channel Scanning. In *IEEE INFOCOM 2022 - IEEE Conference on Computer Communications.* IEEE, London, United Kingdom, 540–549.
[96] Fu Yu, Xiaolong Zheng, Liang Liu, and Huadong Ma. 2023. Enabling concurrency for non-orthogonal lora channels. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking.* 1–15.
[97] Fu Yu, Xiaolong Zheng, Yuhao Ma, Liang Liu, and Huadong Ma. 2024. Resolve Cross-Channel Interference for LoRa. In *2024 IEEE 44th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 1027–1038.
[98] Shiming Yu, Xianjin Xia, Ningning Hou, Yuanqing Zheng, and Tao Gu. 2024. Revolutionizing LoRa Gateway with XGate: Scalable Concurrent Transmission across Massive Logical Channels. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking.* 482–496.
[99] Zhiwei Zhao, Weifeng Gao, Wan Du, Geyong Min, Wenliang Mao, and Mukesh Singhal. 2022. Towards energy-fairness in LoRa networks. *IEEE Transactions on Mobile Computing* (2022).