Are LoRa Logical Channels Really Orthogonal? Practically Orthogonalizing Massive Logical Channels

Shiming Yu[1], Ziyue Zhang[1], Xianjin Xia[1], Yuanqing Zheng[1], Jiliang Wang[2]
[1] The Hong Kong Polytechnic University, Hong Kong SAR, China
[2] Tsinghua University, Beijing, China
{shiming.yu,ziyue.zhang}@connect.polyu.hk,{xianjin.xia,yuanqing.zheng}@polyu.edu.hk,jiliangwang@tsinghua.edu.cn

ABSTRACT

LoRaWANs are envisioned to connect billions of IoT devices through thousands of physically overlapping yet logically orthogonal channels (termed logical channels). These logical channels hold significant potential for enabling highly concurrent scalable IoT connectivity. Large-scale deployments however face strong interference between logical channels. This practical issue has been largely overlooked by existing works but becomes increasingly prominent as LoRaWAN scales up. To address this issue, we introduce Canas, an innovative gateway design that is poised to orthogonalize the logical channels by eliminating mutual interference. To this end, Canas develops a series of novel solutions to accurately extract the meta-information of individual ultra-weak LoRa signals from the received overlapping channels. The meta-information is then leveraged to accurately reconstruct and subtract the LoRa signals over thousands of logical channels iteratively. Real-world evaluations demonstrate that Canas can enhance concurrent transmissions across overlapping logical channels by 2.3× compared to the best known related works.

CCS CONCEPTS

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

KEYWORDS

Internet of Things, LPWAN, LoRa, Logical Channel

ACM Reference Format:
Shiming Yu[1], Ziyue Zhang[1], Xianjin Xia[1], Yuanqing Zheng[1], Jiliang Wang[2] . 2025. Are LoRa Logical Channels Really Orthogonal? Practically Orthogonalizing Massive Logical Channels. In *The 23rd Annual International Conference on Mobile Systems, Applications and Services (MobiSys '25), June 23–27, 2025, Anaheim, CA, USA.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3711875.3729125

1 INTRODUCTION

Recent years have witnessed Low-Power Wide-Area Networks (LPWANs) as a promising wireless technology to enable ubiquitous connectivity for wireless sensing devices and empower various mobile services and applications. As a leading LPWAN technology, LoRa stands out with its extensive coverage and energy-efficient communication capabilities, making it suitable for Internet of Things (IoT) applications [4, 5, 16, 18, 24, 26, 28–30, 38, 44, 45, 59–61, 81, 84] such as smart agriculture [8, 49, 83], smart metering [9, 32], and satellite [7, 15, 46]. According to industry reports [53], the LoRaWAN ecosystem has already connected over 350 million LoRa nodes and 6.9 million gateways globally.

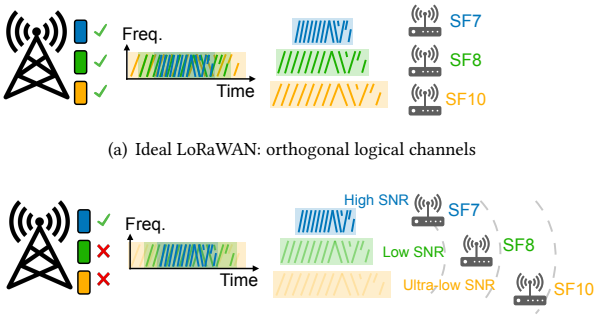

Figure 1: The orthogonality between logical channels can be broken because of *inter-logical channel interference*. The research problem of how to resolve the inter-logical channel interference remains unexplored.

LoRa networks can service vast physical areas at the kilometer level, covering massive end nodes with link SNRs ranging from −15 dB to 20 dB [66]. To fully harness the spectrum for massive IoT communications, LoRa adopts diverse and flexible channel configurations. The LoRa spectrum can be divided into many physically separated channels (*e.g.*, 208 channels in US). Within each channel, LoRa supports overlapping *logical channels* by modulating packets with different spreading factors (*e.g.*, SF6~12). Owing to the SF-selective LoRa demodulation, signals with different SFs can be separated by the receiver even if they physically overlap with each other. Therefore, overlapping logical channels are traditionally deemed "orthogonal" in LoRa networks. Such orthogonality is pivotal to greatly facilitate wireless access without channel activity detection, enable flexible rate adaptation of individual nodes, and allow operators to scale LoRa networks.

However, large-scale LoRa networks encounter practical challenges due to the *imperfect orthogonality between logical channels*. Our measurements (§3) discover an unexpected finding: concurrent transmissions across logical channels presumed orthogonal (which dominate the concurrent traffic in large-scale LoRa deployments), could suffer 61% packet loss due to the imperfect orthogonality.

MobiSys '25, June 23–27, 2025, Anaheim, CA, USA

ACM ISBN 979-8-4007-1453-5/25/06
https://doi.org/10.1145/3711875.3729125

Many ultra-weak packets from distant LoRa nodes are lost due to interference from concurrent logical channels, leading to serious starvation and power depletion of weak nodes. Although much research effort has been dedicated to supporting concurrent transmissions, the previous research mainly focused on resolving explicit collisions for single-channel [56, 65, 67, 76] or cross-channels [87, 88] scenarios. The interference among logical channels has been largely overlooked and under-explored. Given that numerous LoRa nodes often use different logical channels for concurrent transmissions [54], the lack of reliable orthogonality between logical channels presents a significant scalability challenge for LoRa networks. In this paper, we ask the following research question: *Is it possible to resolve inter-logical channel interference to practically orthogonalize massive logical channels for concurrent transmission?*

To answer this question, we first investigate the existing strategies to receive concurrent logical channels in LoRaWAN. Existing works primarily rely on the SF-specific LoRa demodulation to separate overlapping logical channels. When a LoRa receiver targets at a specific logical channel, the packets over that logical channel will benefit from a distinct demodulation gain (*e.g.*, 3∼17 dB [51]), while the signals over adjacent logical channels will be suppressed. Nevertheless, the signals from other logical channels, though not demodulated, could physically overlap with the target channel and increase the demodulation noise floor (*i.e.*, termed *logical channel interference*[1]). The problem becomes more challenging in large-scale LoRa deployments. Because of the high concurrency and near-far effect, we notice the increased noise floor can affect and even overwhelm the SF-selective demodulation gain, thereby breaking the logical channel orthogonality.

In this paper, we present *Canas*, the first gateway design that cancels the mutual interference between logical channels for concurrent LoRa transmissions. At the core of *Canas* is its novel logical channel interference cancellation technique that suppresses the signal of non-targeted logical channels down to the noise floor. We achieve this by precisely reconstructing the raw signal of non-target logical channels and removing them from the Rx signal before demodulation of target packets. By doing so, *Canas* enhances the orthogonality of logical channels and better supports LoRa communications over massive logical channels in large-scale LoRa deployments.

However, turning this basic idea into a practical system faces substantial challenges. Unlike traditional gateways that directly demodulate a target logical channel, *Canas* aims to eliminate other interfering logical channels before demodulation. Since the concurrent logical channels overlap, it is non-trivial to extract the signal of an individual logical channel, let alone subtract it from the Rx signal. Though previous works [21] have demonstrated the effectiveness of strong signal reconstruction and cancellation for other wireless technologies (*e.g.*, Wi-Fi), these solutions cannot be applied in LoRa due to two practical challenges: Firstly, the traditional methods (*e.g.*, Zig-zag decoding [21]) require a clean strong signal chunk to bootstrap signal reconstruction and cancellation process. In contrast, under concurrent logical channels, the concurrent signals can overlap and distort the amplitude and phase of each logical channel, making it challenging to obtain the meta-information of a specific ultra-weak logical channel without being influenced by others. Moreover, ultra-weak LoRa packets are also subject to various hardware imperfections caused by their unique long airtime, low-cost Tx hardware modules, and unreliable power supply. These imperfections increase the randomness and unpredictability of interference signals, which pose significant obstacles to accurate interference signal reconstruction and cancellation. Note that existing techniques, such as frequency offset cancellation methods, are primarily designed for packet reception rather than interference signal reconstruction and thus cannot be directly applied to tackle the obstacles. That is because packet reception only needs to figure out transmitted bits, while interference cancellation requires precise reconstruction of physical samples corresponding to their waveforms. Therefore, achieving logical channel interference cancellation over massive LoRa logical channels remains challenging.

Canas presents a novel design to practically orthogonalize massive logical channels by precisely reconstructing and canceling the interference of ultra-weak LoRa signals. We find that orthogonality between logical channels is compromised when the noise floor of a weak logical channel is elevated by stronger ones. Conversely, the strong logical channels are less affected by the weaker ones. Based on this observation, it is feasible to clearly demodulate the relatively strong logical channel and translate the payload bits to its raw signal. To achieve accurate raw signal reconstruction, *Canas* employs novel techniques to capture various signal offsets of the target logical channel without being affected by others. These meta-information capturing the impact of various factors (*e.g.*, hardware imperfections, air-channel, *etc*) are then emulated to faithfully reconstruct the raw interference signals. Subsequently, we can subtract and cancel the strong interference from the Rx signal. This process can be iterated to receive all concurrent logical channels, progressively from the strongest to the weakest.

We implement and evaluate *Canas* with commercial off-the-shelf (COTS) LoRa devices and Software Defined Radio (SDR). Results demonstrate that *Canas* significantly outperforms existing techniques adapted to cancel interference among logical channels. Specifically, *Canas* achieves 2.3× concurrency gain on packet reception across logical channels than the best known baselines by enhancing the orthogonality between logical channels. *Canas* can be seamlessly integrated with the existing LoRa gateways with software modification. In summary, our work makes the following contributions:

- For the first time, we draw attention to the problem of inter-logical channel interference in large-scale LoRa deployments, emphasizing its uniqueness and how it differs from previously studied problems (§2).
- We conduct a comprehensive analysis of the inter-logical channel interference problem, supported by extensive measurements, to illustrate its implications, practical impacts, and root causes (§3).
- We present *Canas* to effectively orthogonalize massive LoRa logical channels. *Canas* employs innovative techniques to overcome the unique challenges involved in precise reconstruction and iterative cancellation of ultra-weak LoRa signals (§4).

[1] used interchangeably with "inter-logical channel interference" in this paper

System Name	Concurrent Transmission	Target Problem
FTrack [76]	Same <BW, SF> Same Chirp Slope (non-orthogonal)	Single-channel Collision
NScale [65]		
CIC [56]		
Mc-LoRa [87]	Different <BW, SF> Same Chirp Slope (non-orthogonal)	Cross-channel Collision
XGate [89]	Different <BW, SF> Different Chirp Slope ("orthogonal")	Scale to More Logical Channels
Canas	Different <BW, SF> Different Chirp Slope ("orthogonal")	Inter-Logical Channel Interference

Table 1: Comparisons between *Canas* and SOTA. *Canas* addresses related but different research problem. Note that XGate [89] focuses on packet detection and decoding over massive logical channels, and does not address the problem of inter-logical channel interference.

- We implement and evaluate *Canas* on an outdoor testbed. Our evaluations demonstrate significant improvements in concurrent transmission across logical channels compared to the best baselines, as well as seamless integration and enhancement of packet reception in existing systems under practical near-far effects (§5).

2 LIMITATIONS OF EXISTING SOLUTIONS

Existing solutions for LoRa concurrent transmissions mainly focus on addressing collisions within a single logical channel or cross-channel collisions. To the best of our knowledge, this paper is the first to systematically investigate the inter-logical channel interference as summarized in Table 1.

Packet Reception over Multiple Non-overlapping Channels. To facilitate concurrent packet reception, commercial off-the-shelf (COTS) gateway chips (*e.g.*, SX1301 [52]) divide the 1.6 MHz bandwidth spectrum into eight independent, non-overlapping physical channels, along with an additional configurable high-rate channel. After frequency shifting and filtering through the hardware Rx chains, gateways employ SF-specific LoRa detectors and decoders to receive concurrent logical channels.

Resolving Single-channel Collisions. Most recent works primarily focus on resolving collisions within one single logical channel (see Figure 2(a)) to support concurrent transmissions [14, 36, 56, 70, 76, 78, 80]. The primary approach involves separating collided symbols within each decoding window and then decoding them individually using standard LoRa techniques. State-of-the-art separation methods exploit hardware imperfections [14], time misalignment [76], non-stationary features [65], reception diversity [75], and sub-symbol features [56]. However, when faced with inter-logical channel interference from overlapping logical channels, LoRa packet reception encounters the unique challenge of an increased noise floor, rendering these symbol separation techniques ineffective (as detailed in §3.2).

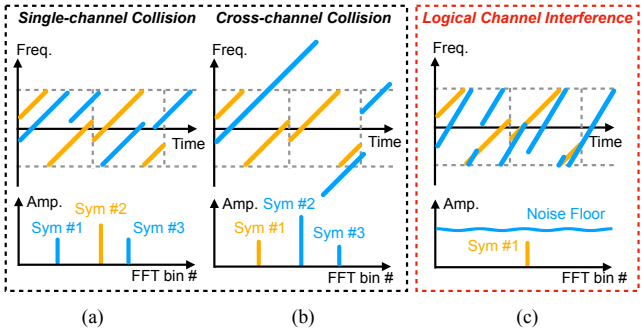

Figure 2: Existing works deal with explicit collisions under single-channel [56] and cross-channel [87] scenarios while *Canas* solves the increased noise floor from inter-logical channel interference.

Decoding Cross-channel Collisions. A few studies [87, 88] have noticed that non-orthogonal logical channels with different bandwidths can also cause collisions in LoRa networks (see Figure 2(b)). To enable concurrent transmission under cross-channel collisions, Mc-LoRa [87] employs signal viper techniques to diversify time-frequency domain features and separate collisions. SD-LoRa [88] further enhances performance by utilizing self-dechirp operations. Similar to single-channel collision recovery methods, these techniques are ineffective under inter-logical channel interference due to the increased noise floor, caused by imperfect orthogonality among logical channels.

Scaling to More Logical Channels. Recent work [89] highlights the disparity between the numerous available logical channels in the LoRa spectrum and the limited number covered by existing gateways. To enable flexible scaling, XGate [89] introduces an auto-configured LoRa gateway capable of covering all available logical channels in the Rx spectrum. Although XGate achieves hundreds of concurrent transmissions using the same spectrum as COTS gateways, it still suffers from inter-logical channel interference, leading to significant packet loss when all "orthogonal" logical channels are used concurrently.

Successive Interference Cancellation. Previous studies [3, 10, 43, 55] have applied successive interference cancellation (SIC) in Code Division Multiple Access (CDMA) systems to support concurrent transmissions. However, under inter-logical channel interference, traditional SIC methods fail because LoRa experiences significant signal fluctuation and instability due to longer packet airtime as well as frequency drift due to low-cost hardware and clocks. Additionally, SIC requires high signal quality, whereas concurrent LoRa packets often have diverse SNR conditions and poor signal quality. We draw strength from these pioneering works and address the unique technical challenges involved in accurate reconstruction and cancellation of ultra-weak LoRa signals affected by various practical factors (*e.g.*, hardware imperfection of lightweight LoRa devices, frequency offset and jitters, large disparity in signal strength and quality).

Difference with Existing Works. Existing works mainly focused on the collision recovery for concurrent transmission, addressing both single-channel [14, 56, 76] and cross-channel [87, 88]

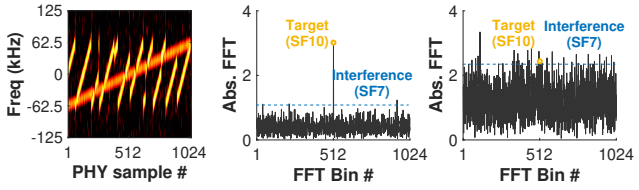

(a) Raw signal (b) ΔSNR = 3 dB (c) ΔSNR = 10 dB

Figure 3: Demodulation under logical channel interference: (a) Spectra of concurrent logical channels; (b) Demodulation of SF10 logical channel; (c) Orthogonality is broken by the increased noise floor.

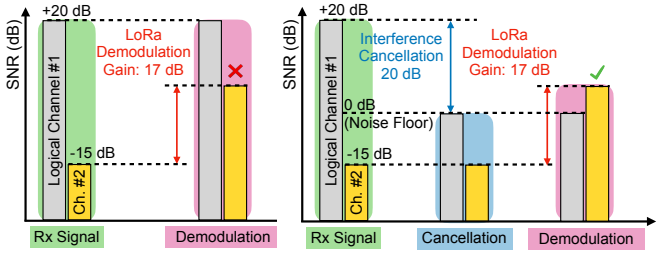

(a) Standard LoRa (b) *Canas*

Figure 4: Illustration and requirements of logical channel interference cancellation problem.

scenarios where explicit collisions occur during demodulation. In contrast, our work focuses on the concurrent transmission over multiple logical channels (see Figure 2(c)). These channels introduce strong mutual interference instead of explicit collision in large-scale deployments. Our extensive measurements (§3) have revealed that this issue poses a significant bottleneck for reliable packet delivery from distant nodes in large-scale LoRa deployments, a problem that has been largely overlooked until now. Unlike the widely studied collision recovery problem [14, 20, 75, 87], this study is, to the best of our knowledge, the first to identify and tackle the problem of inter-logical channel interference.

3 PROBLEM AND MOTIVATION

3.1 LoRa Logical Channels

LoRa leverages chirp spread spectrum (CSS) modulation as the physical layer technique to enable its long-range communication capabilities while maintaining low power consumption [33, 40, 48, 62, 82, 85]. CSS uses chirps to modulate symbols. The duration of a chirp is determined by two key LoRa parameters: Spreading Factor (SF) and Bandwidth (BW). Specifically, the chirp duration T is calculated as $T = \frac{2^{SF}}{BW}$. A *base up-chirp*, whose frequency linearly increases from $-\frac{BW}{2}$ to $\frac{BW}{2}$ over the chirp duration, can be represented as:

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

where the chirp slope $k = \frac{BW}{T}$ denotes the rate of frequency change over time.

LoRa modulates symbols by changing the initial frequency of the base up-chirp. A modulated symbol $S(f_{sym}, k, t)$ can be represented as:

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

The frequency shift is in a cyclic manner, where the frequencies higher than $\frac{BW}{2}$ will align to $-\frac{BW}{2}$. To demodulate the received symbol, the LoRa receiver performs *dechirp* by multiplying $S(f_{sym}, k, t)$ with the conjugate of the base up-chirp, denoted as $C^{-1}(k,t)$. The operation is represented as:

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

The *dechirp* operation removes the CSS modulation and converts the LoRa symbol into a single tone. Importantly, the energy of the symbol, which was initially spread across the chirp bandwidth, is concentrated at the tone frequency after *dechirp*. Subsequently, the Fast Fourier Transform (FFT) consolidates the signal energy into a peak at f_{sym} in the frequency domain, identifying the symbol. Essentially, the FFT not only identifies the frequency but also gathers the entire chirp energy into a single FFT bin, resulting in a significant *SNR gain after dechirping*. This SNR gain raises the LoRa symbol above the surrounding noise, enabling long-distance and below-the-noise communications.

3.2 Understanding Logical Channel Interference

Imperfect orthogonality among logical channels. LoRa aims to build orthogonal logical channels by utilizing different spreading factors (SF), which controls the chirp slope k and duration T. When receiving a signal on a specific target logical channel with chirp slope k, the signal originating from another logical channel with chirp slope k' can be represented as follows after the dechirping process:

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

Since $k' \neq k$, we observe in the above equation that the signal frequency after dechirp continues to vary over time, rather than exhibiting as a constant frequency. Consequently, the energy of the other interfering chirp will spread across the spectrum instead of being concentrated into a single peak after performing FFT. Therefore, LoRa could potentially enable concurrent transmissions over these logical channels (whose SFs differ, *i.e.*, $k' \neq k$), as if they were orthogonal.

In practice, however, the energy of the other interfering chirp (SF=k') spreads across rather than disappears in the demodulation window. As a result, the spread energy inevitably brings up the noise floor. Thus, the increased noise floor caused by a nearby interfering node could sometimes completely overwhelm the ultra-weak signals from a remote target node even when the target signals (SF=k) can be concentrated into a single peak.

As a matter of fact, practical deployments and measurements reveal that logical channels are seldom completely orthogonal. Figure 3(a) illustrates the spectrum of two concurrent logical channels with different SFs (*e.g.*, SF7 and SF10). When the receiver locks on SF10, the target logical channel is clearly distinguishable from the Rx signal, benefiting from the processing gain after dechirping

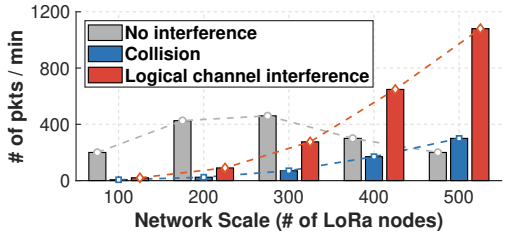

Figure 5: Packets that experience logical channel interference dominate the traffic in large-scale LoRa deployments.

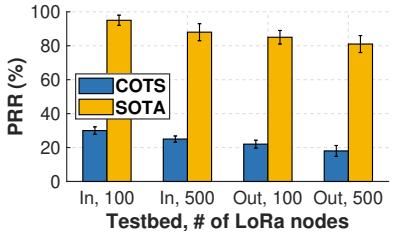

Figure 6: Packets under collisions can be recovered by SOTA solution.

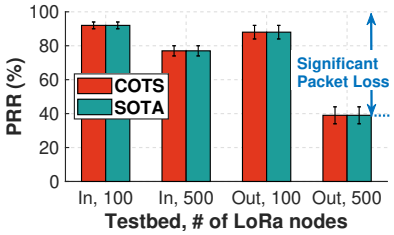

Figure 7: Packets under logical channel interference suffer from significant packet loss.

(as shown in Figure 3(b)). However, the signals from other logical channels do not vanish but remain present within the demodulation window, dramatically increasing the noise floor. As the SNR difference increases, the orthogonality between logical channels weakens and can even be broken, as shown in Figure 3(c). Such SNR differences (*e.g.*, >10 dB) are frequently encountered due to various factors (*e.g.*, near-far effect, adaptive transmission power, diverse antenna gains) in LoRa networks that cover vast physical areas at the kilometer scale [66]. Thus, the logical channel interference could severely affect the packet reception of ultra-weak nodes.

Limitations of current logical channel reception. The root cause for logical channel interference is that standard LoRa solely relies on the demodulation SNR gain to separate different logical channels. As shown in Figure 4(a), LoRa demodulation can provide 3~17 dB SNR gain [51] for specific logical channel. However, a large-scale LoRa network covers massive nodes with varying distances from meters to kilometers and SNRs from −15 dB to 20 dB [54, 66]. Given that different SFs typically serve different SNR conditions, the power gap between logical channels can be as large as 35 dB, which easily exceeds the demodulation gain and destroys the orthogonality.

Requirement on orthogonalizing logical channels. In this paper, we define orthogonality breaks as situations where a weak logical channel cannot be demodulated due to interference from stronger logical channels. The general strategy of orthogonalizing the logical channels is to cancel the interference of strong logical channels on the weak ones. As shown in Figure 4(a), a gateway's Rx signal consists of concurrent logical channels from multiple nodes, where the power of strong logical channels can be orders of magnitude higher than the weak ones [66]. To achieve reliable concurrent transmission across logical channels, it is crucial to minimize and cancel the interference of strong logical channels, *i.e.*, *logical channel interference cancellation.*

Note that LoRa demodulation provides each logical channel with a unique SNR gain, allowing it to stand out from the noise floor and enable long-range below-the-noise communication. However, the presence of strong logical channels raises the noise floor of weaker ones, thus compromising the orthogonality between the concurrent logical channels. Therefore, to effectively cancel the interference caused by strong logical channels, it is imperative to suppress them below the noise floor. Next, we can exploit the SNR gain of LoRa demodulation to efficiently cancel the residual interference (see Figure 4(b)).

Orthogonality breaks as networks scale. We conduct measurement studies on the indoor and outdoor testbeds (see Figure 11) to quantitatively investigate the impact of increased networks on the orthogonality requirement. Our setup utilizes both COTS and USRP gateways to capture measurements and record data traces. The testbed comprises 50 COTS LoRa nodes distributed across various ranges and SNRs. We increase the transmit duty cycle to practically emulate a larger number of nodes. The nodes adopt an ALOHA-based MAC protocol.

Figure 5 illustrates the variation among three types of packets (*i.e.*, no interference, collision, and logical channel interference). As the network scales, an increasing number of packets experience collisions or logical channel interference. Notably, packets experiencing logical channel interference (*i.e.*, transmitted across different logical channels) dominate, significantly outweighing those experiencing collisions (*i.e.*, transmitted within a single logical channel). While the SOTA solution of collision recovery (*e.g.*, CIC [56]) can recover some packets affected by collisions (see Figure 6), the packet reception ratio for packets affected by logical channel interference—which dominate the traffic—still drops significantly to 39% in the outdoor scenario (see Figure 7). Such logical channel interference cannot be resolved by the collision recovery solution. The reason is that collision recovery techniques (*e.g.*, CIC [56]) are designed for resolving explicit collision (see Figure 2(a)) and cannot handle the increased noise floor from logical channel interference (see Figure 2(c)).

3.3 Objective

Logical channels have demonstrated significant advantages in enabling high concurrency [89, 90], efficient channel access [19, 86], and adaptive data rate [31] in LoRa networks. The LoRa nodes usually switch between logical channels to adapt to diverse link quality and data rate requirements [1, 39, 41, 72]. The use of concurrent logical channels is widely applied in LoRa and brings great opportunities to connect millions of end nodes with high spectrum efficiency [66, 89]. However, the limitation of imperfect orthogonality practically hinders the performance of connecting massive LoRa nodes in large-scale deployments, necessitating novel solutions. Conventional LoRa gateways solely rely on LoRa demodulation to separate logical channels. Unfortunately, this approach is inadequate in large-scale deployments with high concurrency and significant SNR disparity among nodes. To unleash the potential of massive logical channels, we propose a new LoRa gateway design -

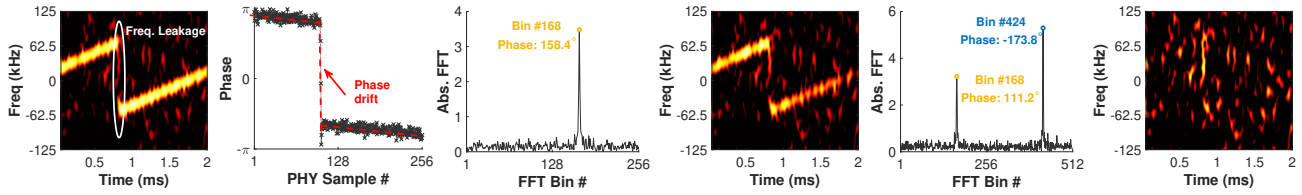

(a) A received symbol (b) Phase drift (c) Standard dechirp (d) w/o calibration (e) Up-sample dechirp (f) w calibration

Figure 8: Illustration of frequency leakage and impacts on interference cancellation: (a) LoRa symbol experiences frequency leakage at chirp edges; (b) Random phase drift; (c) Standard dechirp: two chirp segments are aliased; (d) Cancellation fails due to phase misalignment between the local chirp and Rx signal; (e) Up-sample dechirp: two chirp segments are separated; (f) Cancellation succeeds after calibrating frequency leakage in the local chirp signal.

Canas, to orthogonalize logical channels by canceling inter-logical channel interference. *Canas* only requires updates on gateways, without any modifications to the deployed COTS LoRa nodes.

4 SYSTEM DESIGN

In this section, we introduce *Canas*, a novel gateway design that eliminates logical channel interference to support reliable concurrent transmission across logical channels. We first present the key technical components for signal reconstruction and then give the system architecture.

4.1 Reconstruction of Interfering Logical Channel

Without loss of generality, suppose N nodes communicate concurrently on overlapping logical channels. The Rx signal of concurrent logical channels can be denoted as $y(t) = \sum_{i=1}^{N} S_i(t - t_i) + n(t)$, where $S_i(t)$ denotes the raw signal of i^{th} logical channel, t_i indicates the time shift of raw signal, and $n(t)$ is the Rx noises.

Assuming that $S_1(t)$ has a high signal strength, it interferes with other logical channels. Accurately subtracting the interfering logical channel from the Rx signal requires a precise replica of $S_1(t)$. However, extracting $S_1(t)$ directly from the received signal $y(t)$ is ineffective due to the presence of overlapping concurrent logical channels. These overlapping signals introduce distortions to the amplitude and phase characteristics of $S_1(t)$, making it challenging to obtain an accurate replica of the raw signal.

To address this practical issue, the basic idea behind *Canas* signal reconstruction is as follows: we leverage the fact that the interfering logical channel $S_1(t)$ is typically strong and less interfered with by other weaker logical channels. This allows us to demodulate $S_1(t)$ and extract its payload data explicitly. Instead of directly estimating $S_1(t)$ from the Rx signal $y(t)$, *Canas* generates a local signal replica of the interfering logical channel $S_1'(t)$ using the received payload data and reconstructs $\hat{S_1}(t)$ based on this signal replica.

Nevertheless, reconstructing a LoRa packet is considerably more complex than traditional wireless signals [10, 21]. The reason is that LoRa utilizes CSS modulation that has a much longer symbol duration, making it more vulnerable to various sophisticated signal offsets. Since the LoRa packets are transmitted by low-cost hardware, the received signals experience significant hardware imperfections, causing them to deviate substantially from the ideal signals. As a result, the locally generated signal $S_1'(t)$ is not yet an effective signal replica of the interfering logical channel $S_1(t)$. *Canas* further calibrates the various frequency and phase offsets introduced by low-cost RF components to achieve precise signal reconstruction.

Phase Drift. Since LoRa modulates data by cyclically shifting the initial frequency of the chirp signal, a typical LoRa symbol consists of two chirp segments, as illustrated in Figure 8(a). LoRa radios are susceptible to frequency leakage [75, 77] when the frequency transitions from one chirp to another (*e.g.*, at the chirp edges). This frequency leakage often introduces phase drifts to the transmitted symbols, resulting in inner-symbol phase variance (see Figure 8(b)).

A naive reconstruction method is to directly build on the locally generated signal. We use the demodulation result in Figure 8(c) to reconstruct a local chirp and subtract it from the received signal (*i.e.*, the symbol in Figure 8(a)). The outcome is displayed in Figure 8(d)). We observe that the residual signal still retains high energy, indicating that the cancellation has failed. This failure occurs because the local signal and the received signal, despite carrying the same symbol, are not phase-aligned. To effectively reconstruct the received signal and further subtract it from the Rx signal, it is essential to ensure their phase consistency.

Therefore, the key problem is to precisely capture and emulate the phase drift when reconstructing the local signal $S_1'(t)$. This task is non-trivial because the phase drifts introduced by frequency leakage are typically random and unpredictable. Moreover, the phase drift information is distorted during LoRa demodulation since the two chirp segments overlap in the same FFT bin due to spectrum aliasing (see Figure 8(c)).

To address this practical issue, we exploit the observation that the phase drift at the chirp edges does not affect the inner phase variation within the two chirp segments. it is possible to estimate the phase drift by measuring the phase difference between these segments, provided they can be accurately separated. Our key insight is that the distortion of phase drift information is caused by spectrum aliasing [79], resulting from the limited sampling rate. By up-sampling the LoRa symbol to a higher sampling rate (*e.g.*, 2× bandwidth), we can separate the two chirp segments after dechirping and FFT, as depicted in Figure 8(e). From the two peaks, we can explicitly extract the phase difference between the two chirp

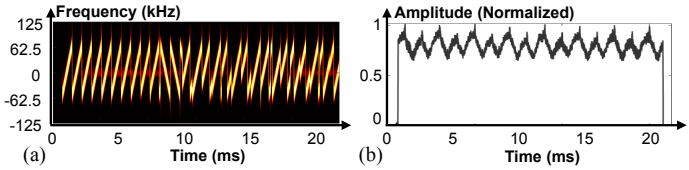

Figure 9: Dynamic channel fluctuation: (a) spectra and (b) amplitude of a received LoRa packet.

segments, which exactly indicates the phase drift at the edge. Subsequently, we calibrate the phase drift during the local chirp reconstruction process and subtract it from the received signal. As shown in Figure 8(f), the received symbol is successfully suppressed to a noise level, demonstrating effective signal reconstruction.

Carrier Frequency Offset. The oscillator frequency of the gateway and the LoRa node cannot be perfectly matched. As a result, there is usually a carrier frequency offset (CFO) in the received signal, causing linear phase rotations represented as $e^{j2\pi\Delta f t}S'_1(t)$. The CFO has been measured to be approximately 10 kHz [58]. Given the narrow bandwidth of LoRa (e.g., 125 kHz), if left uncorrected, the CFO can significantly impact signal reconstruction.

Only using the preamble to estimate CFO requires a fine-grained sliding window at the sample level, incurring significant computational overhead. To achieve lightweight CFO estimation during the signal reconstruction process, we exploit the packet structure of LoRa to extract the CFO information. The key idea is that LoRa packets contain not only base up-chirps as preamble but also base down-chirps as Start Frame Delimiter (SFD) [25, 71, 75], where different parts of the same packet experience the same CFO. Moreover, the CFO produces opposite frequency effects on the preamble and SFD parts of the packet. By utilizing the up-chirp preamble ($C_{pr}(t)$) and the down-chirp SFD ($C^{-1}_{sfd}(t)$), we can cancel out the CFO effects as follows:

$$he^{j2\pi\Delta f t}C_{pr}(t)\cdot he^{j2\pi\Delta f t}C^{-1}_{sfd}(t) = h^2e^{j2\pi(2\Delta f)t} \quad (5)$$

Here, h represents the impact of the wireless channel. By performing FFT on the resulting signal in Equation 5, the peak location indicates the frequency of $2\Delta f$. This information allows us to estimate and calibrate the CFO during signal reconstruction.

4.2 Enhancing Reconstruction Granularity

Although calibrating the frequency leakage and CFO makes the local signal replica closely resemble the interfering logical channel in the Rx signal, it remains unsatisfactory to directly subtract this replica from the Rx signal. The reason is that the received LoRa packets are subject to dynamic channel fluctuation in addition to various hardware imperfections. Figure 9 illustrates the spectrum and signal amplitude of a received LoRa packet. The long air-time of a LoRa packet (*e.g.*, hundreds of ms) and the inherent instability of the Tx hardware (e.g., Low-cost LoRa node) result in noticeable channel fluctuations in the received raw signal, as depicted in Figure 9(b).

A naive solution is to apply the same amplitude fluctuation on the reconstructed signal replica. However, it is challenging to obtain precise amplitude fluctuation of individual logical channels from

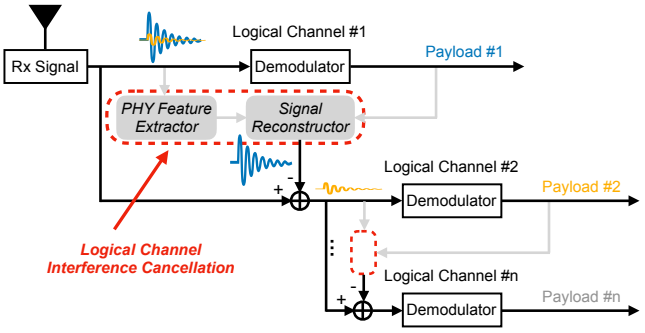

Figure 10: Workflow of *Canas*: *Canas* iteratively demodulates, reconstructs, and removes individual logical channels from the Rx signal.

the Rx signal as it contains multiple concurrent logical channels. Additionally, the dynamic channel fluctuations introduce not only amplitude variations but also phase variations. We address this problem to reconstruct the signals of individual logical channels with finer granularity.

To address the time-varying nature of practical LoRa channels, *Canas* introduces a sophisticated channel estimation technique to facilitate precise signal reconstruction. capitalizes on the observation that each chirp present in a LoRa packet can serve as a basis for channel estimation. By adopting a per-chirp channel estimation strategy, *Canas* enhances the granularity of signal restoration. The received signal $y(t)$ can be represented as the combination of one logical channel to be reconstructed and other logical channels:

$$y(t) = S_1(t) + \sum_{i=2}^{N} S_i(t) = h_1^j S_1'^j(t) + \sum_{i=2}^{N} h_i^j S_i(t) \quad (6)$$

Initially, we cross-correlate the received signal $y(t)$ with a locally generated raw signal $S'_1(t)$ to pinpoint the signal start of the interfering logical channel. Let $y^j(t)$ denote the j^{th} chirp of $y(t)$, we next employ conjugate multiplication and FFT to derive the chirp-level channel response h_1:

$$\mathcal{F}[y^j(t)\cdot[S_1'^j(t)]^*] \approx \mathcal{F}[h_1^j S_1'^j(t)\cdot[S_1'^j(t)]^*] = h_1^j \quad (7)$$

where $(\cdot)^*$ indicates the complex conjugate operation. It is worth noting that conjugate multiplication (*i.e.*, Equation 7) serves a dual purpose. Firstly, it eliminates the common chirp signal in $S_1^j(t)$ and $S_1'^j(t)$. Secondly, the signals of other concurrent logical channels (*i.e.*, $\sum_{i=2}^{N} h_i^j S_i(t)$) are spread across the spectrum and diminish towards the noise floor. Consequently, the conjugate multiplication results in a tone frequency corresponding to $S_1^j(t)$, representing the signal deviation from $S_1'^j(t)$ to $S_1^j(t)$. By employing FFT to aggregate the signal power at this tone frequency into a distinct peak, the chirp-level channel estimation $\hat{h_1^j}$ can be derived. We then iterate through all chirps in $S_1(t)$ and $S'_1(t)$, the collected fine-grained channel data $\hat{h_1} = \{\hat{h_1^j}\}$ can be utilized for accurate reconstruction of $\hat{S_1}(t)$. *Canas* not only tracks the dynamic nature of practical LoRa packet channels but also mitigates the influence of other concurrent logical channels during channel estimation.

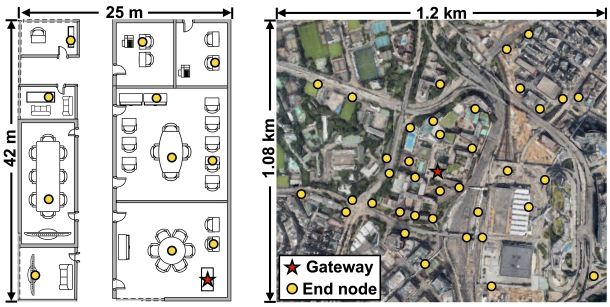

(a) Indoor testbed (b) Outdoor testbed

Figure 11: Testbed setting of *Canas*.

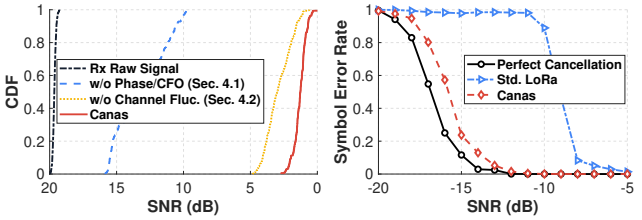

(a) Cancellation performance (b) Weak packet reception

Figure 12: Performance of logical channel interference cancellation: (a) Signal strength of strong logical channel. (b) Packet reception of weak logical channel.

Subsequently, we align $\hat{S}_1(t)$ with $S_1(t)$ and then subtract it from $y(t)$ to eliminate the interfering logical channel.

It is worth noting that the channel response estimation is robust even with multiple interfering signals of similar strength. Note that interfering logical channels use different chirp slopes, and *Canas* uses conjugate multiplication to extract the channel response $\hat{h}_1^j$. The conjugate multiplication (*i.e.*, Equation 7) spreads other interfering signals across the spectrum while $S_1(t)$ accumulates to a peak. This process provides a significant SNR gain, ensuring a sufficiently accurate $\hat{h}_1^j$ estimation. The only exception occurs when an interfering signal $S_2(t)$ is significantly stronger than $S_1(t)$ (*e.g.*, >5 dB). In such cases, *Canas* would initiate cancellation starting with $S_2(t)$.

4.3 *Canas*: Put All Together

Figure 10 shows the general workflow of concurrent logical channel reception in *Canas*.

Logical channel interference cancellation. *Canas* generally applies signal reconstruction to cancel the logical channel interference. Initially, *Canas* demodulates the strongest logical channel, which experiences minimal interference from weaker logical channels. We use the received payload data to locally generate an ideal raw signal. Furthermore, *Canas* extracts diverse signal offsets present within the received LoRa packet, serving as physical layer fingerprints. These offsets are added to the ideal local raw signal to facilitate more precise signal reconstruction. Additionally, *Canas* utilizes the local raw signal to track channel variations in practical received packets, enabling fine-grained signal reconstruction. Finally, *Canas* by subtracting the reconstructed logical channel from the Rx signal, *Canas* effectively suppresses its signal strength to the level of Rx noises. Consequently, the remaining logical channels are free from interference caused by the strong logical channel.

Iteration to receive concurrent logical channels. *Canas* adopts an iterative approach to receive more concurrent logical channels. Since concurrent logical channels may exhibit different signal strengths, they are affected by other logical channels to different degrees. To ensure successful logical channel interference cancellation, *Canas* requires prior demodulation of payload data. We begin by using cross-correlation to detect the preamble of the LoRa packet across all logical channels. This process allows us to determine the packet arrival timing and measure the signal's SNR by analyzing the peak level. Next, we sort the detected packets according to their SNRs and proceed to demodulate the strongest signal. After demodulating the strongest signal, we remove it and move on to the next logical channel. We continue this iterative process to receive concurrent logical channels from the strongest channel to the weakest. The weak logical channels will be demodulated in the later rounds if they are interfered with by other strong packets, even if they arrive earlier. However, a practical challenge arises when the signal strengths of weak LoRa packets fall below the noise floor and become overshadowed by stronger packets, making direct measurement from the Rx signal difficult. To address this, measurements for weak packets can be obtained in subsequent cancellation rounds after removing the signals of strong logical channels.

Integration with collision recovery. *Canas* can seamlessly integrate with SOTA collision recovery techniques to deal with hybrid interference (*i.e.*, coexistence of collision and logical channel interference). The integrated approach still follows the iterations from the strongest to the weakest signals. Within each iteration focusing on one logical channel, if collisions are detected, we can replace the standard demodulator (see Figure 10) with SOTA decoder (*e.g.*, CIC [56]) to handle collisions. Then, Canas reconstructs and cancels each colliding packet in the logical channel.

5 EVALUATION

5.1 Methodology

Implementation. We implement *Canas* based on the software-defined radio platform (USRP N210) and `gr-lora` project [22]. The USRP is employed to receive signals from LoRa nodes, and the received samples are forwarded to a workstation running *Canas* for signal processing. We use COTS LoRa nodes with Semtech SX1276 radio [51] as transmitters and employ Arduino Uno boards for node configuration. We deploy a testbed including 50 LoRa nodes and 2 gateways. We conduct experiments on our campus, spanning an area of 1.08 km × 1.2 km that represents a typical urban environment. The gateways are installed on the rooftop of a 20-storey building, and all nodes operate in the 915 MHz ISM band.

Experiment Setup. We collect data traces from >200 links for over two months in our testbed. The traces encompass diverse channel conditions in typical urban settings, including indoor and outdoor, low and high SNRs (see Figure 11). All evaluations with ≤50 nodes are conducted via real-world experiments in the testbed. Additionally, we also perform large-scale trace-driven evaluations

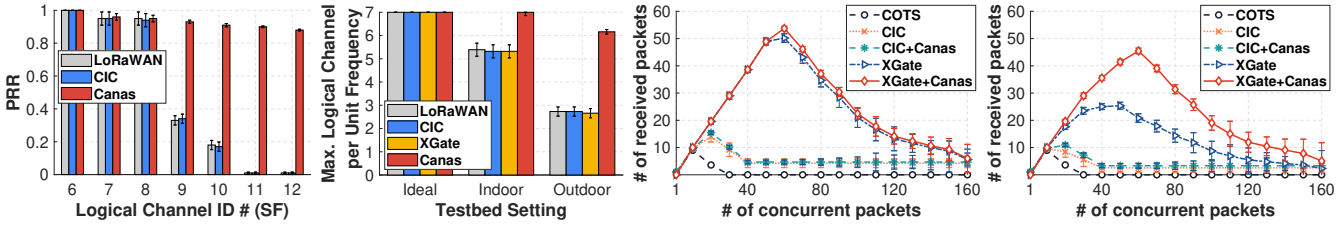

(a) Individual logical channel (b) Logical channel concurrency (c) Overall Concurrency (Indoor) (d) Overall Concurrency (Outdoor)

Figure 13: Comparision with state-of-the-art: (a) packet reception ratio of each logical channel under concurrent transmission; (b) the maximum number of concurrent logical channels supported on a unit frequency; (c) indoor and (d) outdoor performance of concurrent packet reception in a 1.6 MHz spectrum.

using collected data traces. Our evaluation primarily focuses on addressing the following questions: (1) How does *Canas* perform in logical channel interference cancellation? (2) How many concurrent logical channels can be supported by *Canas*? and (3) How does *Canas* perform in supporting practical IoT applications?

Baselines. We compare *Canas* with three baselines on concurrent transmission for LoRa. (1) *LoRaWAN* directly demodulates individual logical channels from the Rx signal; (2) *CIC* adopts collision recovery technique to receive concurrent packets within a single logical channel; (3) *XGate* implements auto-configuration gateway to cover all available logical channels in the Rx spectrum.

5.2 Logical Channel Interference Cancellation Performance

Cancellation effectiveness. We first evaluate *Canas* in logical channel interference cancellation. We set up two LoRa nodes that transmit concurrently at low and high SNRs (*i.e.*, −15 dB and 20 dB) using SF10 and SF7, respectively. Figure 12(a) illustrates the SNR of SF7 packets over 15 seconds. We observe that the strong logical channel (*i.e.*, SF7) is originally received with high SNR (*i.e.*, around 20 dB) in the Rx signal, which significantly interferes with other logical channels. *Canas* effectively reduces its power to the noise floor, resulting in a median SNR of 1.3 dB after cancellation. We also conduct ablation studies to evaluate the novel techniques in *Canas*. The results demonstrate the crucial role of calibrating the phase drift, frequency offset, and channel fluctuation in achieving precise signal reconstruction and enabling effective logical channel interference cancellation. The median SNR are 12.9 dB and 3.1 dB without canceling phase drift, frequency offset, and calibrating channel fluctuation, respectively (see Figure 12(a)).

Ultra-weak packet reception under interference. We next investigate the Rx sensitivity of *Canas* after logical channel interference cancellation. We fix the SNR of the interferer and measure the symbol error rate of the transmitter at different SNRs. Specifically, we measure the received packets of the SF10 node under different SNR conditions and analyze >100 packets for each SNR condition. The experiments are under interference from a concurrently transmitting SF7 node at the SNR of 10 dB. For comparison, we also disable the SF7 node (*i.e.*, the interferer) to represent perfect cancellation. As shown in Figure 12(b), the perfect cancellation performs best as it has no interference from other logical channels. LoRaWAN [2], which directly demodulates the weak logical channel, suffers from significant SNR loss due to interference from strong logical channels. In contrast, *Canas* can receive weak packets at ultra-low SNRs (*e.g.*, −14.8 dB for SF10) by reconstructing and subtracting the strong logical channel. Compared to the perfect cancellation, the interference from the strong logical channel only leads to 0.9 dB SNR loss on Rx sensitivity of *Canas* while achieving a symbol error rate of <20 %.

5.3 Comparison with state-of-the-art

Reception of individual logical channel. This experiment delves into each logical channel to evaluate the performance of *Canas* in supporting concurrent logical channels. We configure seven LoRa nodes to transmit concurrently using different logical channels (SF 6∼12) on the same central frequency. These nodes are evenly distributed with SNRs ranging from −15 dB to 10 dB in the outdoor testbed area. Nodes select logical channels based on their SNRs, where large (low) SFs are used for low (high) SNRs. We compare *Canas* against LoRaWAN [2] and CIC [56]. LoRaWAN employs a standard LoRa decoder for receiving each logical channel, while CIC utilizes a collision recovery decoder. *Canas* receives concurrent logical channels with iterative interference cancellation.

Figure 13(a) presents the packet reception ratio (PRR) of each logical channel. We observe that LoRaWAN can only reliably receive the low-SF logical channels, and the PRR drops significantly for the high-SF ones. The reason is that the high-SF logical channels are utilized for low SNR links and are susceptible to logical channel interference. The signals from strong logical channels (SF 6∼8) can raise the Rx noise floor for the weaker ones, leading to disrupted demodulation. CIC cannot improve the reception of concurrent logical channels as it focuses on resolving collisions within a single logical channel. In contrast, *Canas* demonstrates a substantial enhancement in packet reception by effectively canceling the mutual interference between logical channels. Even the weakest logical channel (SF12) achieves an average PRR of >80%.

Logical channel concurrency. This experiment compares *Canas* with the current leading strategies for concurrent transmissions in LoRa. We consider two state-of-the-art (SOTA) strategies, *i.e.*, CIC [56] and XGate [89]. To ensure a fair comparison, we focus on the maximum number of concurrent logical channels on the same physical channel (*i.e.*, overlap on the same central frequency). We configure seven LoRa nodes to transmit concurrently using different logical channels (SF 6∼12) on the same central frequency. Both CIC and XGate are implemented on USRP devices, as they require access to the PHY raw signal. We test three scenarios in our

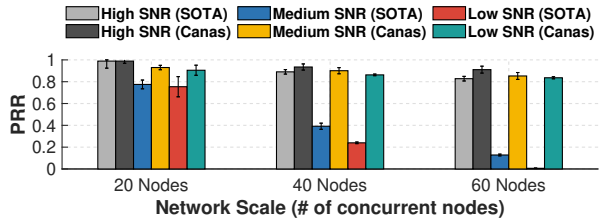

Figure 14: Near-far effect at different network scales.

testbed: (1) Ideal, where nodes transmit independently; (2) Indoor; and (3) Outdoor, representing different environmental conditions (see Figure 11).

Figure 13(b) displays the maximum number of concurrent logical channels per unit frequency received by four strategies. When transmitting independently, all logical channels can be received, as expected. However, when transmitting concurrently, LoRaWAN experiences packet loss due to logical channel interference, resulting in an average concurrency of 5.3 and 2.7 for indoor and outdoor deployments, respectively. The outdoor performance is particularly impacted due to the more pronounced near-far effect and larger SNR variance. CIC improves concurrency by resolving collisions within a single logical channel, yet it fails to support concurrent logical channels. XGate enhances concurrency by automatically configuring and receiving all available logical channels, but it does not address the mutual interference between overlapping logical channels. In comparison, *Canas* outperforms the other strategies, achieving an average concurrency of 7 and 6.1 for indoor and outdoor deployments, respectively (*i.e.*, 1.3× and 2.3× higher logical channel concurrency than SOTA for indoor and outdoor testbed). This superior performance is attributed to *Canas*'s ability to eliminate logical channel interference through iterative signal reconstruction and subtraction.

Integration with SOTAs. This experiment evaluates the performance of *Canas* in enhancing concurrent transmission alongside existing SOTAs. We compare *Canas* with CIC [56] and XGate [89] within the same 1.6 MHz spectrum and 9 Rx chains as the COTS gateway [52]. We control a varying number of LoRa nodes to transmit concurrently on different LoRa channels, each with a 125 kHz bandwidth. By varying the spreading factor (SF) from 7 to 12, we support 54 LoRa logical channels. We utilize a beacon to coordinate the simultaneous transmission of 50 nodes, aiming to assess system performance under extreme concurrency. For experiments involving more than 50 nodes, we aggregate multiple signal traces with randomly assigned time offsets to simulate real-world network traffic. These synthesized traces are then replayed for evaluation. Importantly, we also integrate SOTAs with *Canas* to evaluate the additional benefits that *Canas*'s logical channel interference cancellation technique can bring to existing LoRa systems. Specifically, we implement *CIC+Canas* by replacing the standard LoRa demodulator in *Canas* with CIC (see Figure 10). We implement *XGate+Canas* by replacing the demodulators and decoders in XGate with Canas. The evaluations are conducted on both indoor and outdoor testbeds, representing the signals with similar and diverse SNRs, respectively. The outdoor testbed experiences more significant near-far effect.

Figure 13(c,d) presents the number of received packets for these strategies under different levels of concurrent transmissions. As expected, network concurrency initially increases, reaches a peak, and then declines. CIC achieves 14 and 10 concurrent transmissions for indoor and outdoor testbeds, respectively. *Canas* can work in conjunction with CIC to improve the average concurrency to 15.5 and 11. Their gains are complementary, as they address different types of packet loss. For XGate, it scales concurrent transfers to more logical channels, which also results in more inter-logical channel interference than CIC. We observe that among the 54 concurrent logical channels, XGate in an indoor deployment (*i.e.*, high signal quality) achieves an average of 51 concurrent transmissions with minimal packet loss (see Figure 13(c)). However, in practical outdoor deployments, the performance of XGate drops significantly due to more severe inter-logical channel interference. We note that the maximum concurrency of XGate drops to 25 with considerable packet loss. By augmenting with *Canas*, many of the lost ultra-weak packets can be recovered, resulting in an overall concurrency of 45, boosting the XGate by 80% in the outdoor testbed through seamless integration. These results demonstrate that *Canas* complements SOTAs and significantly enhances LoRa systems under near-far effects in outdoor deployments.

Near-far effect at scale. This experiment evaluates the performance gains of *Canas* under the near-far effect. We test three network scales with 20, 40, and 60 nodes, all evenly distributed within the testbed area and connecting to the gateway with SNRs ranging from −10 dB to 10 dB. The gateways are configured within the same 1.6 MHz spectrum and 9 Rx chains as the COTS gateway [52]. The SOTA combines XGate [89] and CIC [56]. LoRa nodes select logical channels based on their SNRs, with larger (lower) SFs used by nodes with lower (higher) SNRs. To investigate the near-far effects, we control all nodes to transmit concurrently and measure the packet reception performance across different SNR regimes.

Figure 14 illustrates the packet reception ratio (PRR) for the three SNR regimes as received by SOTA and *Canas*. Most of the received packets experience imperfect orthogonality and suffer from varying degrees of PRR loss. We observe that the PRRs for high-SNR packets are less affected as the network scales, remaining above 80 % even when 60 nodes transmit concurrently. In contrast, the PRRs for medium- and low-SNR packets drop dramatically, as weaker packets suffer from significantly more logical channel interference than stronger packets, which SOTAs cannot resolve. *Canas* addresses this by iteratively reconstructing and removing strong logical channels, thereby canceling the interference and improving the PRRs for medium- and low-SNR packets to above 80 %. These results demonstrate that *Canas* significantly enhances the reception of ultra-weak packets under practical near-far effects in large-scale deployments.

Massive packet reception performance. This experiment assesses *Canas* for supporting massive LoRa nodes in large-scale deployments using trace-driven emulations. Specifically, we aim to deploy LoRa gateways to connect up to 1,500 IoT sensors. To ensure a fair comparison, we adopt the same channel plan as the COTS gateway [52], which includes 54 logical channels (BW125 kHz, SF7∼12) distributed over 9 central frequencies. Each IoT sensor transmits a 20-byte message every 30 minutes, with a duty cycle of ≤1%. To investigate communications for thousands of sensors, we employ a

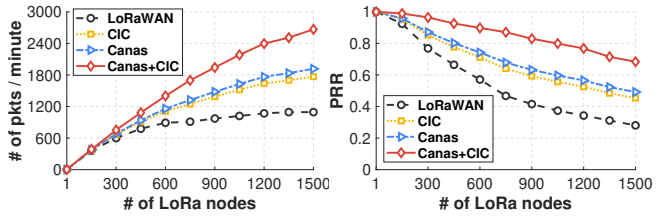

(a) Throughput (b) Packet reception ratio

Figure 15: Performance of supporting massive LoRa nodes in large-scale deployments.

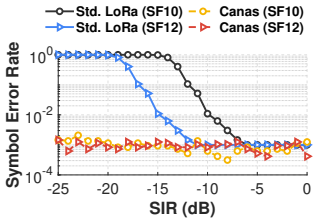

Figure 16: Impact of SIR.

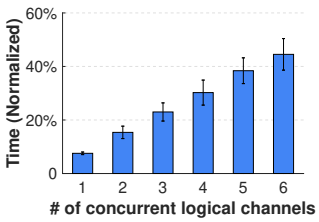

Figure 17: Time overhead.

trace-driven approach. Specifically, we capture LoRa packet profiles from over 200 sites in the outdoor testbed, with SNR ranging from −20 dB to 10 dB. We use the collected data traces to synthesize the received packet signals with randomly selected link profiles. LoRa nodes freely select from the available logical channels to transmit messages using an ALOHA-based MAC. For benchmarking, we compare *Canas* with two baselines: (1) LoRaWAN, which utilizes a COTS gateway [52] for reception; and (2) CIC, which implements collision recovery [56] upon LoRaWAN.

Figure 15 compares the receive packet throughput and packet reception ratio (PRR) of *Canas* and baselines under different numbers of active LoRa nodes. We see that the packet throughput of LoRaWAN first increases and then becomes saturated. The reason is twofold. Firstly, packets using the same logical channel may arrive simultaneously, leading to collisions. Secondly, packets from strong logical channels can interfere with weaker ones, resulting in significant packet loss for high-SF logical channels. As shown in Figure 15(a), CIC improves the packet rate through collision recovery techniques. Meanwhile, *Canas* achieves higher improvement by canceling the logical channel interference. Importantly, the benefit provided by *Canas* is orthogonal to traditional collision recovery methods [56, 67, 76]. *Canas* can jointly work with CIC to support >800 LoRa nodes with PRR>85% (see Figure 15(b)). Importantly, the concurrency improvement of *Canas+CIC* compared to pure CIC is higher than Figure 13(c,d). This is because *Canas+CIC* experiences a performance decrease when handling hybrid interference. In this experiment, which uses a real-world duty cycle, the concurrency level is not as pronounced as in Figure 13(c,d), where strict concurrency is evaluated.

5.4 Microbenchmarks

Impact of SIR. To assess the impact of the signal-to-interference ratio (SIR) on packet reception, we conduct experiments that quantitatively measure the interference caused by strong logical channels on weaker ones. We fix the SNR of the transmitter and change the SNR of the interferer, aiming to measure the symbol error rate of the transmitter at different SIRs. We deploy a LoRa node as the transmitter, which transmits packets with a fixed link SNR of −10 dB. Simultaneously, another LoRa node was deployed as the interferer, transmitting packets with varying link SNRs ranging from −10 dB to 15 dB. The transmitter utilizes SF10 and SF12, while the interferer uses SF7.

Figure 16 illustrates the decoding errors observed at the transmitter node under different SIR conditions. Generally, the symbol error rate (SER) increases as the interference becomes stronger (*i.e.*, as the SIR decreases). We observe that the SIR thresholds (with SER <10%) for SF10 and SF12 are −17 dB and −12 dB, respectively. High-SF logical channels demonstrate greater resilience to logical channel interference. However, such SIR conditions between logical channels frequently occur in large-scale LoRa deployments characterized by significant SNR variances and concurrent transmissions. Compared to LoRaWAN, *Canas* significantly reduces the SIR during LoRa demodulation, thereby enhancing the reception performance of weaker logical channels.

Time overhead. In this experiment, we examine the time overhead of logical channel interference cancellation in *Canas*. The algorithm is executed on a laptop, and we measure the runtime for canceling different numbers of logical channels. We normalize the overhead to the ratio of packet duration to assess the real-time performance. The experimental results demonstrate that the overhead increases linearly, as depicted in Figure 17. Importantly, even when canceling six logical channels, the overhead remains below 50% of packet duration. This duration is much shorter than the airtime of LoRa packets and can be effortlessly handled by the processors in the gateway.

6 RELATED WORK

Weak packet reception. Extending the coverage area and receiving ultra-weak packets is a core challenge for LoRa and has been extensively studied [6, 13, 17, 23, 34, 35, 37, 42, 47, 57, 63, 68, 69, 73]. Charm [11] leverages the spatial diversity of multiple gateways to enhance LoRaWAN coverage. MALoRa [27] combines weak signals from multi-antenna gateways constructively to boost LoRa packet reception. Falcon [64] extends LoRa coverage by allowing weak links to selectively interfere with strong transmissions. XCopy [74] enhances weak link reception by coherently combining retransmitted packets. Previous research mainly focused on weak packet reception of a single node. In contrast, our work aims to improve the reception of weak packets in large-scale deployments. In practical scenarios with high concurrency and the near-far effect, weak logical channels experience significant interference from stronger ones. This problem is distinct and more complex than those addressed in earlier studies.

Tackling hardware and channel uncertainty. LoRa packets experience significant signal offsets, such as CFO, phase drift, and amplitude fluctuation, primarily due to the long packet air-time and low-cost hardware. This hardware-induced uncertainty differs among devices and transmission parameters, making it challenging to predict. Previous offset compensation methods aim to enhance demodulation. In contrast, *Canas* aims to reconstruct the raw signal

of individual logical channels with special consideration for replaying these offsets. PCube [75] utilizes synchronized antenna pairs to calibrate the CFO before demodulation, while *Canas* leverages the inherent conjugate chirp in the LoRa packet to estimate CFO and further replay it to reconstruct a local signal. LoRaTrimmer [12] uses probabilistic model to mitigate phase drift but does not precisely capture it, while *Canas* does and further replays on local signal for precise signal reconstruction. GLoRiPHY [50] extracts preamble channel information for payload denoising. This method can decode packets, but the granularity is insufficient for precise reconstruction (see Figure 12(a)). Canas presents new methods to capture and replay the signal fluctuation for precise signal reconstruction.

Interference signal cancellation. Previous studies have investigated interference signal cancellation techniques for successive interference cancellation (SIC) [3, 10, 43, 55]. While both *Canas* and traditional SIC-based methods involve reconstructing and canceling signals from the strongest to the weakest, they differ fundamentally in several key aspects: (1) Research Problem: Canas focuses on addressing inter-logical channel interference in LoRa networks, which arises due to imperfect orthogonality between logical channels. Traditional SIC, on the other hand, deals with multi-user interference signals within the same channel, where differences in signal strength are the primary concern. Canas, however, handles concurrent logical channels where interfering signals vary not only in signal strength but also in modulation parameters and symbol/packet durations. (2) Unique Challenges: Canas introduces new techniques to address these issues. Precise signal reconstruction in LoRa is challenged by factors such as weak signal strength, hardware imperfections, and long symbol/packet durations. Canas develops innovative methods to extract hardware imperfections and track the air-channel with fine granularity.

7 DISCUSSION AND FUTURE WORK

Benefits and overheads of *Canas*. *Canas* greatly empowers the reception reliability of weak packets (*i.e.*, from distant LoRa nodes) in large-scale deployments. The evaluation of corresponding energy efficiency gains will be carried out in future work. *Canas* can integrate with existing LoRa gateways without requiring any modifications on the deployed COTS LoRa nodes. To cancel the logical channel interference, *Canas* introduces higher computational overheads at the gateway but can be processed in real time.

Interaction with PHY error correction. LoRa's physical layer utilizes a forward error correction mechanism that employs a coding rate to introduce redundancy into the transmitted data, enabling it to correct a certain percentage of symbol errors (*e.g.*, <20% symbol error for a coding rate of 4/5). This error correction can handle limited instances of logical channel interference, such as when a strong SF7 packet interferes with only a few SF12 symbols due to the significant difference in symbol duration. However, it cannot fully resolve the issue. Firstly, SF12 packets have a long duration and can be disrupted by multiple low-SF packets, leading to accumulated symbol errors that exceed the error correction's capacity. Secondly, neighboring SFs can interfere with each other due to similar packet airtime, resulting in symbol errors that surpass the 20% threshold. To effectively address logical channel interference, the error correction mechanism should be complemented by *Canas*. For instance, *Canas* can reduce the symbol error rate of weak packets from over 30% to below 10%, allowing the error correction algorithm to resolve the residual errors.

Handling hybrid interference. Both *Canas* and SOTA collision recovery methods, such as CIC [56], address interference scenarios where multiple nodes transmit simultaneously on the same channel frequency. SOTA methods can resolve traditional collision issues (same BW, same SF), but fail to handle logical channel interference (same SF, different SF). A further challenge is hybrid interference, where traditional packet collisions occur alongside logical channel interference. Preliminary results suggest that Canas+CIC can partially address this issue. However, the performance of both Canas and CIC diminishes under conditions of extreme hybrid concurrency. The combined performance improvement of Canas+CIC is less than the optimal sum of their individual gains, as illustrated in Figure 13(c,d). Our future research will focus on developing unified and comprehensive strategies to tackle hybrid interference signals.

8 CONCLUSION

This paper introduces *Canas*, a novel design of LoRa gateway aimed at practically orthogonalizing massive LoRa logical channels. Unlike traditional LoRa gateways that directly demodulate each logical channel, *Canas* incorporates a new logical channel interference cancellation technique, which reconstructs the raw signal of individual logical channels from the superimposed Rx signal and iteratively eliminates their interference. This breakthrough addresses the imperfect orthogonality and strong mutual interference between the long-believed "orthogonal" logical channels in large-scale LoRaWAN deployments, thereby unlocking the full potential of utilizing massive logical channels for achieving high concurrency, flexible channel access, adaptive data rates, and more.

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, 15231424, 15211924, and 15206123), National Natural Science Foundation of China (U22A2031, No. 61932013), 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] 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 Transactions on Internet of Things* 4, 1 (2023), 1–25.

[2] LoRa Alliance. 2022. *Lorawan Specification.* "https://lora-alliance.org/about-lorawan".

[3] Jeffrey G Andrews and Teresa H-Y Meng. 2004. Performance of multicarrier CDMA with successive interference cancellation in a multipath fading channel. *IEEE Transactions on Communications* 52, 5 (2004), 811–822.

[4] Atul Bansal, Akshay Gadre, Vaibhav Singh, Anthony Rowe, Bob Iannucci, and Swarun Kumar. 2021. Owll: Accurate lora localization using the tv whitespaces. In *Proceedings of the 20th International Conference on Information Processing in Sensor Networks (co-located with CPS-IoT Week 2021).* 148–162.

[5] Jiani Cao, Jiesong Chen, Chengdong Lin, Yang Liu, Kun Wang, and Zhenjiang Li. 2024. Practical Gaze Tracking on Any Surface with Your Phone. *IEEE Transactions on Mobile Computing* (2024).

[6] Qianwu Chen, Xianjin Xia, Zhigang Li, and Yuanqing Zheng. 2021. Collaborative Transmission over Intermediate Links in Duty-Cycle WSNs. In *2021 IEEE 27th*

International Conference on Parallel and Distributed Systems (ICPADS). IEEE, 843–850.

[7] Ying Chen, Jintao Hu, Jie Zhao, and Geyong Min. 2024. Qos-aware computation offloading in leo satellite edge computing for iot: A game-theoretical approach. *Chinese Journal of Electronics* 33, 4 (2024), 875–885.

[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] Yao Cheng, Hendra Saputra, Leng Meng Goh, and Yongdong Wu. 2018. Secure smart metering based on LoRa technology. In *2018 IEEE 4th International Conference on Identity, Security, and Behavior Analysis (ISBA)*. 1–8.

[10] GS Deepthy and RJ Susan. 2012. Analysis of successive interference cancellation in CDMA systems. In *2012 Second International Conference on Advanced Computing & Communication Technologies*. IEEE, 481–485.

[11] Adwait Dongare, Revathy Narayanan, Akshay Gadre, Anh Luong, Artur Balanuta, Swarun Kumar, Bob Iannucci, and Anthony Rowe. 2018. Charm: exploiting geographical diversity through coherent combining in low-power wide-area networks. In *ACM/IEEE IPSN 2018*. IEEE, 60–71.

[12] Jialuo Du, Yunhao Liu, Yidong Ren, Li Liu, and Zhichao Cao. 2024. LoRaTrimmer: Optimal Energy Condensation with Chirp Trimming for LoRa Weak Signal Decoding. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*.

[13] 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.

[14] Rashad Eletreby, Diana Zhang, Swarun Kumar, and Osman Yağan. 2017. Empowering low-power wide area networks in urban settings. In *ACM SIGCOMM 2017*. 309–321.

[15] Akshay Gadre, Zachary Machester, and Swarun Kumar. 2024. Adapting LoRa Ground Stations for Low-latency Imaging and Inference from LoRa-enabled CubeSats. *ACM Transactions on Sensor Networks* 20, 5 (2024), 1–30.

[16] Akshay Gadre, Revathy Narayanan, and Swarun Kumar. 2018. Maintaining UAV stability using low-power WANs. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*. 738–740.

[17] Akshay Gadre, Revathy Narayanan, Anh Luong, Anthony Rowe, Bob Iannucci, and Swarun Kumar. 2020. Frequency Configuration for Low-Power Wide-Area Networks in a Heartbeat.. In *USENIX NSDI*. 339–352.

[18] Akshay Gadre, Fan Yi, Anthony Rowe, Bob Iannucci, and Swarun Kumar. 2020. Quick (and dirty) aggregate queries on low-power WANs. In *2020 19th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, 277–288.

[19] Amalinda Gamage, Jansen Liando, Chaojie Gu, Rui Tan, Mo Li, and Olivier Seller. 2023. LMAC: Efficient carrier-sense multiple access for LoRa. *ACM Transactions on Sensor Networks* 19, 2 (2023), 1–27.

[20] Orestis Georgiou and Usman Raza. 2017. Low power wide area network analysis: Can LoRa scale? *IEEE Wireless Communications Letters* 6, 2 (2017), 162–165.

[21] Shyamnath Gollakota and Dina Katabi. 2008. Zigzag decoding: Combating hidden terminals in wireless networks. In *ACM SIGCOMM 2008*. 159–170.

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

[23] Ningning Hou, Yifeng Wang, Xianjin Xia, Shiming Yu, Yuanqing Zheng, and Tao Gu. 2025. MoLoRa: Intelligent Mobile Antenna System for Enhanced LoRa Reception in Urban Environments. In *ACM SenSys*.

[24] 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).

[25] Ningning Hou, Xianjin Xia, and Yuanqing Zheng. 2021. Jamming of lora phy and countermeasure. In *IEEE INFOCOM 2021*. IEEE, 1–10.

[26] Ningning Hou, Xianjin Xia, and Yuanqing Zheng. 2022. Cloaklora: A covert channel over lora phy. *IEEE/ACM Transactions on Networking* 31, 3 (2022), 1159–1172.

[27] 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.

[28] Jinyan Jiang, Jiliang Wang, Yijie Chen, Shuai Tong, Pengjin Xie, Yihao Liu, and Yunhao Liu. 2024. Willow: Practical WiFi Backscatter Localization with Parallel Tags. In *Proceedings of the 22nd Annual International Conference on Mobile Systems, Applications and Services*. 265–277.

[29] Jinyan Jiang, Jiliang Wang, Yihao Liu, Yijie Chen, and Yunhao Liu. 2024. WiCloak: Protect Location Privacy of WiFi Devices. In *ACM/IEEE IPSN*. 101–112.

[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] 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.

[32] Preti Kumari, Rahul Mishra, Hari Prabhat Gupta, Tanima Dutta, and Sajal K Das. 2021. An energy efficient smart metering system using edge computing in LoRa network. *IEEE Transactions on Sustainable Computing* 7, 4 (2021), 786–798.

[33] 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.

[34] 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.

[35] Chenning Li, Hanqing Guo, Shuai Tong, Xiao Zeng, Zhichao Cao, Mi Zhang, Qiben Yan, Li Xiao, Jiliang Wang, and Yunhao Liu. 2021. NELoRa: Towards ultra-low SNR LoRa communication with neural-enhanced demodulation. In *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems*. 56–68.

[36] Chenning Li, Xiuzhen Guo, Longfei Shangguan, Zhichao Cao, and Kyle Jamieson. 2022. CurvingLoRa to Boost LoRa Network Throughput via Concurrent Transmission. In *USENIX NSDI 2022*. USENIX Association, Renton, WA, 879–895.

[37] 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 *ACM MobiSys 2024*. 479–491.

[38] Ruonan Li, Ziyue Zhang, Xianjin Xia, Ningning Hou, Wenchang Chai, Shiming Yu, Yuanqing Zheng, and Tao Gu. 2025. From Interference Mitigation to Toleration: Pathway to Practical Spatial Reuse in LPWANs. In *Proceedings of the 31st Annual International Conference on Mobile Computing and Networking*.

[39] Ruinan Li, Xiaolong Zheng, Yuting Wang, Liang Liu, and Huadong Ma. 2022. Polarscheduler: Dynamic transmission control for floating lora networks. In *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*. IEEE, 550–559.

[40] 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.

[41] 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 Transactions on Internet of Things* 4, 4 (2023), 1–25.

[42] Qianhe Meng, Han Wang, Chong Zhang, Yihang Song, Songfan Li, Li Lu, and Hongzi Zhu. 2024. Processor-Sharing Internet of Things Architecture for Large-scale Deployment. In *Proceedings of the 22nd ACM Conference on Embedded Networked Sensor Systems*. 211–224.

[43] Pulin Patel and Jack Holtzman. 1994. Analysis of a simple successive interference cancellation scheme in a DS/CDMA system. *IEEE journal on selected areas in communications* 12, 5 (1994), 796–807.

[44] Yuben Qu, Zhenhua Wei, Zhen Qin, Tao Wu, Jinghao Ma, Haipeng Dai, and Chao Dong. 2024. Collaborative Service Provisioning for UAV-Assisted Mobile Edge Computing. *Chinese Journal of Electronics* 33, 6 (2024), 1504–1514.

[45] 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.

[46] 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*. 755–769.

[47] Yidong Ren, Gen Li, Yimeng Liu, Younsuk Dong, and Zhichao Cao. 2025. AeroEcho: Towards Agricultural Low-power Wide-area Backscatter with Aerial Excitation Source. In *IEEE INFOCOM*.

[48] 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)*. IEEE, 1–12.

[49] 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.

[50] Kanav Sabharwal, Soundarya Ramesh, Jingxian Wang, Dinil Mon Divakaran, and Mun Choon Chan. 2024. Enhancing LoRa Reception with Generative Models: Channel-Aware Denoising of LoRaPHY Signals. In *Proceedings of the 22nd ACM Conference on Embedded Networked Sensor Systems*. 507–520.

[51] Semtech. 2022. *SX1276*. "https://www.semtech.com/products\/wireless-rf/lora-connect/sx1276".

[52] Semtech. 2022. *SX1301*. "https://www.semtech.com/products\/wireless-rf/lora-core/sx1301".

[53] Semtech. 2024. *LoRa*. "https://www.semtech.com/lora".

[54] Semtech. 2024. *Understanding ADR*. "https://lora-developers.semtech.com/documentation/tech-papers-and-guides/understanding-adr".

[55] Souvik Sen, Naveen Santhapuri, Romit Roy Choudhury, and Srihari Nelakuditi. 2012. Successive interference cancellation: Carving out MAC layer opportunities. *IEEE Transactions on Mobile Computing* 12, 2 (2012), 346–357.

[56] Muhammad Osama Shahid, Millan Philipose, Krishna Chintalapudi, Suman Banerjee, and Bhuvana Krishnaswamy. 2021. Concurrent Interference Cancellation: Decoding Multi-Packet Collisions in LoRa. In *ACM SIGCOMM 2021*. 503–515.

[57] Chenglong Shao and Osamu Muta. 2024. TONARI: Reactive Detection of Close Physical Contact Using Unlicensed LPWAN Signals. *ACM Transactions on Internet of Things* 5, 2 (2024), 1–30.

[58] Guanxiong Shen, Junqing Zhang, Alan Marshall, Linning Peng, and Xianbin Wang. 2021. Radio frequency fingerprint identification for LoRa using spectrogram and CNN. In *IEEE INFOCOM 2021-IEEE Conference on Computer Communications*. IEEE, 1–10.

[59] 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 *ACM MobiSys*. 398–411.

[60] Leming Shen, Qiang Yang, Xinyu Huang, Zijing Ma, and Yuanqing Zheng. 2025. GPIoT: Tailoring Small Language Models for IoT Program Synthesis and Development. In *ACM SenSys*.

[61] Leming Shen, Qiang Yang, Yuanqing Zheng, and Mo Li. 2025. AutoIOT: LLM-Driven Automated Natural Language Programming for AIoT Applications. In *Proceedings of the 31st Annual International Conference on Mobile Computing and Networking*.

[62] 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.

[63] 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.

[64] Shuai Tong, Zilin Shen, Yunhao Liu, and Jiliang Wang. 2021. Combating Link Dynamics for Reliable Lora Connection in Urban Settings. In *ACM MobiCom 2021*. New York, NY, USA, 642–655.

[65] Shuai Tong, Jiliang Wang, and Yunhao Liu. 2020. Combating packet collisions using non-stationary signal scaling in LPWANs. In *ACM MobiSys 2020*. ACM, Toronto Ontario Canada, 234–246.

[66] Shuai Tong, Jiliang Wang, Jing Yang, Yunhao Liu, and Jun Zhang. 2023. City-wide LoRa Network Deployment and Operation: Measurements, Analysis, and Implications. In *ACM SenSys 2023*. 362–375.

[67] Shuai Tong, Zhenqiang Xu, and Jiliang Wang. 2020. CoLoRa: Enabling Multi-Packet Reception in LoRa. In *IEEE INFOCOM 2020*. 2303–2311. ISSN: 2641-9874.

[68] 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 Transactions on Internet of Things* 4, 2 (2023), 1–36.

[69] Han Wang, Yihang Song, Qianhe Meng, Zetao Gao, Chong Zhang, and Li Lu. 2024. Sisyphus: Redefining Low Power for LoRa Receiver. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*. 1177–1191.

[70] Jiliang Wang, Shuai Tong, Zhenqiang Xu, and Pengjin Xie. 2024. Real-Time Concurrent LoRa Transmissions Based on Peak Tracking. *IEEE Transactions on Mobile Computing* (2024).

[71] Xiong Wang, Linghe Kong, Zucheng Wu, Long Cheng, Chenren Xu, and Guihai Chen. 2020. SLoRa: Towards secure LoRa communications with fine-grained physical layer features. In *ACM SenSys 2020*. 258–270.

[72] Yuting Wang, Xiaolong Zheng, Liang Liu, and Huadong Ma. 2022. PolarTracker: Attitude-aware channel access for floating low power wide area networks. *IEEE/ACM Transactions on Networking* 30, 4 (2022), 1807–1821.

[73] 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.

[74] Xianjin Xia, Qianwu Chen, Ningning Hou, Yuanqing Zheng, and Mo Li. 2023. XCopy: Boosting Weak Links for Reliable LoRa Communication. In *ACM MobiCom 2023*. Association for Computing Machinery, 1–15.

[75] Xianjin Xia, Ningning Hou, Yuanqing Zheng, and Tao Gu. 2021. PCube: scaling LoRa concurrent transmissions with reception diversities. In *ACM MobiCom 2021*. Association for Computing Machinery, 670–683.

[76] Xianjin Xia, Yuanqing Zheng, and Tao Gu. 2019. FTrack: parallel decoding for LoRa transmissions. In *ACM SenSys 2019*. 192–204.

[77] Xianjin Xia, Yuanqing Zheng, and Tao Gu. 2021. Litenap: Downclocking lora reception. *IEEE/ACM Transactions on Networking* 29, 6 (2021), 2632–2645.

[78] Zhenqiang Xu, Shuai Tong, Pengjin Xie, and Jiliang Wang. 2020. FlipLoRa: Resolving collisions with up-down quasi-orthogonality. In *IEEE SECON*. 1–9.

[79] 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.

[80] Zhenqiang Xu, Pengjin Xie, and Jiliang Wang. 2021. Pyramid: Real-time lora collision decoding with peak tracking. In *IEEE INFOCOM 2021*. IEEE, 1–9.

[81] Huanqi Yang, Zehua Sun, Hongbo Liu, Xianjin Xia, Yu Zhang, Tao Gu, Gerhard Hancke, and Weitao Xu. 2023. ChirpKey: a chirp-level information-based key generation scheme for LoRa networks via perturbed compressed sensing. In *IEEE INFOCOM*. IEEE, 1–10.

[82] Kang Yang, Yuning Chen, Xuanren Chen, and Wan Du. 2023. Link quality modeling for lora networks in orchards. In *IPSN 2023*. 27–39.

[83] 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 Annual International Conference on Mobile Systems, Applications and Services*. 304–317.

[84] Kang Yang and Wan Du. 2024. A Low-Density Parity-Check Coding Scheme for LoRa Networking. *ACM Transactions on Sensor Networks* (2024).

[85] Mingran Yang, Junbo Zhang, Akshay Gadre, Zaoxing Liu, Swarun Kumar, and Vyas Sekar. 2020. Joltik: enabling energy-efficient" future-proof" analytics on low-power wide-area networks. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*. 1–14.

[86] 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, London, United Kingdom, 540–549.

[87] Fu Yu, Xiaolong Zheng, Liang Liu, and Huadong Ma. 2023. Enabling concurrency for non-orthogonal lora channels. In *ACM MobiCom 2023*. 1–15.

[88] 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.

[89] 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.

[90] Shiming Yu, Xianjin Xia, Ziyue Zhang, Ningning Hou, and Yuanqing Zheng. 2024. FDLoRa: Tackling Downlink-Uplink Asymmetry with Full-duplex LoRa Gateways. In *Proceedings of the 22nd ACM Conference on Embedded Networked Sensor Systems*. 281–294.