Resolving Inter-Logical Channel Interference for Large-scale LoRa Deployments

Shiming Yu, *Student Member, IEEE*, Ziyue Zhang, *Student Member, IEEE*, Xianjin Xia, *Member, IEEE*, Yuanqing Zheng, *Senior Member, IEEE*, and Jiliang Wang, *Senior Member, IEEE*

***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.**

***Index Terms*—Internet of Things, LPWAN, LoRa**

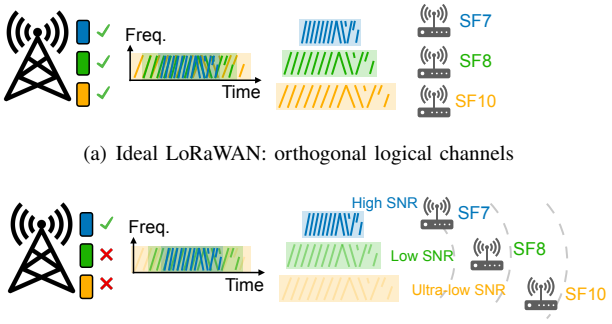

(a) Ideal LoRaWAN: orthogonal logical channels

(b) Practical LoRaWAN: inter-logical channel interference

Fig. 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.

I. 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 [1]–[12] such as smart agriculture [13]–[15], smart metering [16], and satellite [17]. According to industry reports [18], the LoRaWAN ecosystem has already connected over 350 million LoRa nodes and 6.9 million gateways globally.

LoRa networks can service vast physical areas at the kilometer level, covering massive end nodes with link SNRs ranging from $-15\,\mathrm{dB}$ to $20\,\mathrm{dB}$ [19]. 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 (§III) 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. 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 [20]–[23] or cross-channels [24] 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

Manuscript received 8 April 2025; revised; revised 19 August 2025; accepted 9 September 2025; approved by IEEE TRANSACTIONS ON MOBILE COMPUTING Editor xxx. Date of publication xxx; date of current version xxx; This work was supported in part by the Hong Kong General Research Fund (GRF) under Grant 15218022, 15231424, 15211924, 15216325, in part by the National Nature Science Foundation of China (U22A2031, 61932013). Xianjin Xia and Yuanqing Zheng are the corresponding authors.

Shiming Yu, Ziyue Zhang, Xianjin Xia, Yuanqing Zheng are with the Department of Computing, The Hong Kong Polytechnic University, Hong Kong, China (email: shiming.yu@connect.polyu.hk; ziyue.zhang@connect.polyu.hk; xianjin.xia@polyu.edu.hk; yuanqing.zheng@polyu.edu.hk)

Jiliang Wang is with the School of Software, BNRist, Tsinghua University, Beijing 100190, China (email: jiliangwang@tsinghua.edu.cn)

Digital Object Identifier xxxx

transmissions [25], 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.*, $3{\sim}17\,\mathrm{dB}$ [26]), 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 support LoRa communications over massive logical channels in large-scale LoRa deployments.

Turning this basic idea into a practical system presents significant challenges. Unlike traditional gateways that directly demodulate a target logical channel, *Canas* seeks to eliminate interfering logical channels before demodulation. The overlapping nature of concurrent logical channels makes it difficult to isolate the signal of an individual channel, let alone subtract it from the received signal. Previous works [27] have shown the effectiveness of strong signal reconstruction and cancellation in other wireless technologies like Wi-Fi. However, these solutions are not applicable to LoRa due to two main challenges: Firstly, traditional methods like Zigzag decoding [27] require a clean, strong signal segment to initiate the signal reconstruction and cancellation process. In the presence of concurrent logical channels, overlapping signals can distort the amplitude and phase of each channel, making it difficult to obtain the meta-information of a specific ultra-weak logical channel without interference from others. Additionally, ultra-weak LoRa packets are affected by various hardware imperfections due to their long airtime, low-cost transmission hardware, and unreliable power supply. These imperfections increase the randomness and unpredictability of interference signals, posing significant obstacles to accurate interference signal reconstruction and cancellation. Existing techniques, such as frequency offset cancellation methods, are primarily designed for packet reception rather than interference signal reconstruction and cannot be directly applied to address these challenges. Packet reception focuses on determining transmitted bits, whereas interference cancellation requires precise reconstruction of physical samples corresponding to their waveforms. Moreover, the need to analyze the presence of logical channel interference adds to the challenge of algorithm efficiency. Traditional methods that repeatedly correlate the received signal with different spreading factor-modulated preambles can introduce significant detection latency and redundant computation. Therefore, achieving logical channel interference cancellation over numerous LoRa logical channels remains a challenging task.

Canas introduces an innovative approach to effectively orthogonalize massive logical channels by accurately reconstructing and canceling the interference from ultra-weak LoRa signals. We observe that the orthogonality between logical channels is disrupted when the noise floor of a weak logical channel is raised by stronger ones, while stronger channels are less impacted by weaker ones. Leveraging this insight, it becomes feasible to clearly demodulate the relatively strong logical channel and convert its payload bits back into the raw signal. To ensure precise raw signal reconstruction, *Canas* utilizes novel techniques to capture various signal offsets of the target logical channel without interference from others. This meta-information, which accounts for factors such as hardware imperfections and air-channel effects, is then used to accurately emulate and reconstruct the raw interference signals. Once reconstructed, the strong interference can be subtracted and canceled from the received signal. This process is iteratively applied to decode all concurrent logical channels, starting from the strongest and progressing to the weakest. To minimize packet detection latency across concurrent logical channels, *Canas* introduces an innovative logical channel analyzer capable of detecting all concurrent logical channels in a single round.

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 $2.3\times$ 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 (§II).
- 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 (§III).
- We present *Canas* to effectively orthogonalize massive LoRa logical channels. *Canas* employs innovative tech-

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

System Name	Concurrent Transmission	Target Problem
FTrack [20]	Same $< BW, SF >$ Same Chirp Slope (non-orthogonal)	Single-channel Collision
NScale [21]		
CIC [23]		
Mc-LoRa [24]	Different $< BW, SF >$ Same Chirp Slope (non-orthogonal)	Cross-channel Collision
XGate [28]	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 I
COMPARISONS BETWEEN *CANAS* AND SOTA.

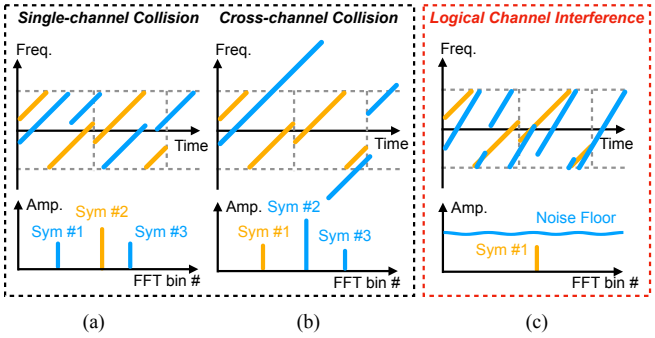

Fig. 2. Existing works deal with explicit collisions under single-channel [23] and cross-channel [24] scenarios while *Canas* solves the increased noise floor from inter-logical channel interference.

niques to overcome the unique challenges involved in precise reconstruction, iterative cancellation of ultra-weak LoRa signals, and efficient logical channel interference analysis (§IV).

- 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 (§V).

II. Related Work

Concurrent transmission. Existing solutions for LoRa concurrent transmissions mainly focus on addressing collisions within a single logical channel [20], [23], [29]–[32] or cross-channel collisions [24], [33]. The primary approach involves separating collided symbols within each decoding window and then decoding them individually using standard LoRa techniques. More detailed analysis is introduced in [34]. However, when faced with inter-logical channel interference from overlapping logical channels (see Figure 2, LoRa packet reception encounters the unique challenge of an increased noise floor, rendering these symbol separation techniques ineffective (as detailed in §III-B) To the best of our knowledge, this paper is the first to systematically investigate the inter-logical channel interference as summarized in Table I.

Succesive interference cancellation. Prior research has explored interference signal cancellation techniques, particularly in the context of successive interference cancellation (SIC) [35]–[38]. While both *Canas* and traditional SIC-based approaches involve reconstructing and canceling signals in order from the strongest to the weakest, they differ fundamentally in several key aspects: (1) Research Problem: Canas specifically targets inter-logical channel interference in LoRa networks, which stems from the imperfect orthogonality between logical channels. In contrast, traditional SIC methods are designed to address multi-user interference within the same channel, primarily focusing on differences in signal strength. Canas must contend with concurrent logical channels where interfering signals differ not only in power but also in modulation parameters and symbol or packet durations. (2) Unique Challenges: Canas introduces novel techniques to overcome these unique challenges. Accurate signal reconstruction in LoRa is complicated by factors such as weak signal strength, hardware imperfections, and extended symbol or packet durations. To address these, Canas implements advanced methods for extracting hardware imperfections and tracking the air-channel with high precision.

Weak packet reception. Significant research efforts have been dedicated to improving the reception of weak LoRa packets [39], [40]. Falcon [41] extends LoRa coverage by allowing weak links to selectively interfere with strong transmissions. MALoRa [42] enhances packet reception by constructively combining weak signals from multi-antenna gateways. Ostinato [43] uses payload collaboration to improve the reception of ultra-low SNR packets. XCopy [44] boosts weak link reception by coherently combining retransmitted packets. While previous research has primarily focused on the reception of weak packets from a single node, *Canas* targets the improvement of weak packet reception in large-scale deployments. In real-world scenarios characterized by high concurrency and the near-far effect, weak logical channels face substantial interference from stronger ones. This issue is distinct and more complex than those tackled in earlier studies.

Logical channel. Many studies have explored the use of logical channels in LoRaWAN to enhance scalability [45]–[47]. CurvingLoRa [48] employs non-linear chirps to create additional logical channels, supporting higher concurrency. OrthoRa [49] divides a full LoRa chirp into scattered sub-chirps to accommodate more logical channels. ChirpTransformer [50] utilizes hardware interrupts to implement chirp hopping for versatile logical channel encoding. MAC layer methods [51], [52] coordinate LoRa nodes to switch between logical channels, improving network performance. While previous studies have primarily focused on creating more logical channels and using them flexibly to enhance network performance, they have not adequately addressed the interference between logical channels and its negative impacts. In contrast, *Canas* identifies and resolves a new research challenge, thereby enhancing the ability of existing approaches that leverage numerous logical channels to achieve scalable LoRaWAN.

Link throughput enhancement. Various efforts have been

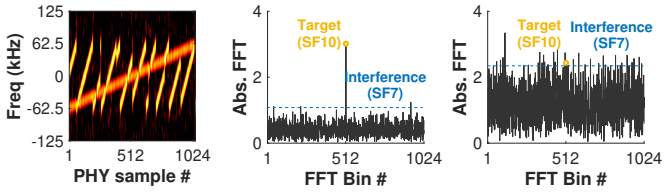

(a) Raw signal (b) ΔSNR = 3 dB (c) ΔSNR = 10 dB

Fig. 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.

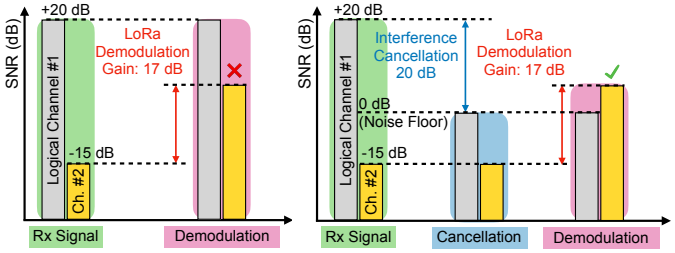

(a) Standard LoRa (b) *Canas*

Fig. 4. Illustration and requirements of logical channel interference cancellation problem.

made to improve the link throughput of LoRaWAN [3]–[5], [50], [53], [54]. HyLink [55] enhances throughput by enabling concurrent chirp (de)modulation within the same packet. MaLoRaGW [56] applies MU-MIMO techniques to LoRa, boosting both uplink and downlink throughput. FDLoRa [53] uses full-duplex gateways to improve scalable downlink communications. In comparison, *Canas* primarily focuses on enhancing link throughput for large-scale outdoor deployments, addressing practical challenges of the near-far effect. *Canas* can complement these existing approaches to further push the throughput limits of LoRaWAN.

III. Problem and Motivation

A. 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 [19], [57]–[60]. 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.

B. 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 (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

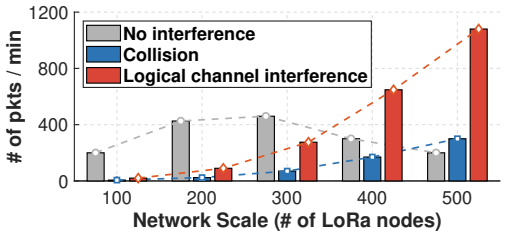

Fig. 5. Packets that experience logical channel interference dominate the traffic in large-scale LoRa deployments.

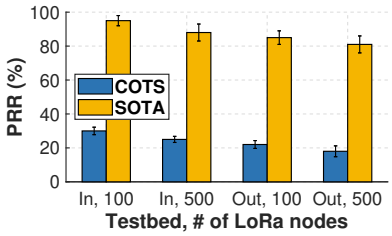

Fig. 6. Packets under collisions can be recovered by SOTA solution.

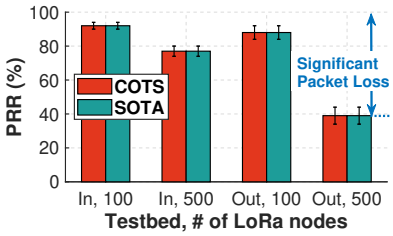

Fig. 7. Packets under logical channel interference suffer from significant packet loss.

and can even be broken, as shown in Figure 3(c). Such SNR differences (*e.g.*, $>10\,\mathrm{dB}$) 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 [19]. 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 $3{\sim}17$ dB SNR gain [26] for specific logical channel. However, a large-scale LoRa network covers massive nodes with varying distances from meters to kilometers and SNRs from $-15\,\mathrm{dB}$ to $20\,\mathrm{dB}$ [19], [25]. Given that different SFs typically serve different SNR conditions, the power gap between logical channels can be as large as $35\,\mathrm{dB}$, which easily exceeds the demodulation gain and destroys the orthogonality.

Requirement on orthogonalizing 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 [19]. 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 12) 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 [23]) 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.

C. Objective

Logical channels have demonstrated significant advantages in enabling high concurrency [28], [53], efficient channel access [51], [61], and adaptive data rate [62] in LoRa networks. 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 - *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.

IV. 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. Firstly, we present the key technical components for signal reconstruction. Then, we introduce an efficient logical channel analyzer to accelerate the detection process. Finally, we give the system architecture.

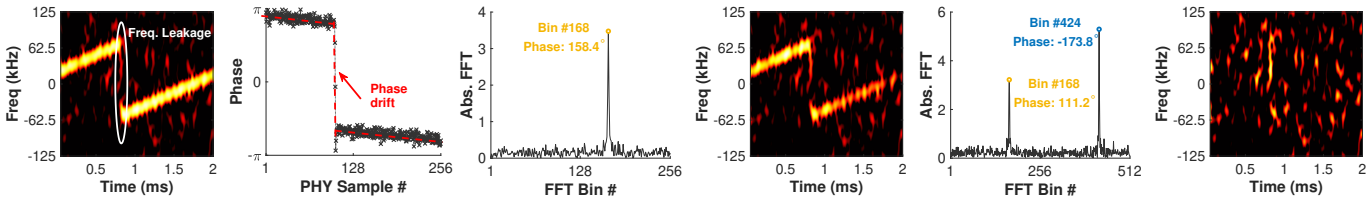

(a) A received symbol (b) Phase drift (c) Standard dechirp (d) w/o calibration (e) Up-sample dechirp (f) w calibration

Fig. 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.

A. *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 [27], [35]. 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 [31], [63] 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)).

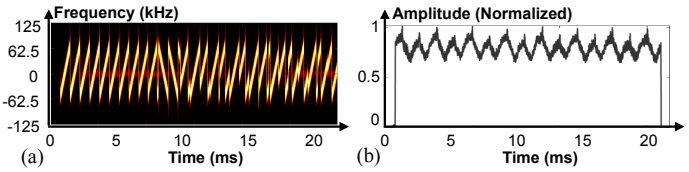

Fig. 9. Dynamic channel fluctuation: (a) spectra and (b) amplitude of a received LoRa packet.

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 [64], resulting from the limited sampling rate. By up-sampling the LoRa symbol to a higher sampling rate (*e.g.*, $2\times$ 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 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

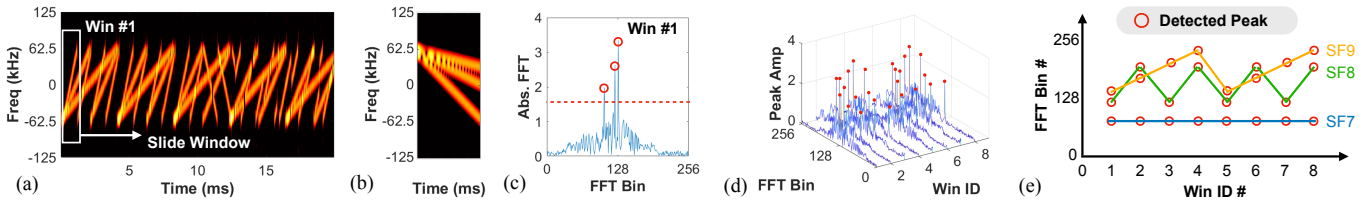

Fig. 10. Illustration of efficient logical channel analyzer: (a) spectrum of three concurrent logical channels (SF7, SF8, and SF9); (b) aggregated downchirps for detection acceleration; (c) result of Win #1: three packets are captured simultaneously via one detection; (d) result across consecutive windows; (e) infer SF based on periodic peak pattern.

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 [65]. Given the narrow bandwidth of LoRa (e.g., 125 kHz), if left uncorrected, the CFO can significantly impact signal reconstruction.

To estimate and calibrate the CFO 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) [31], [66], 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.

B. 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 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. 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

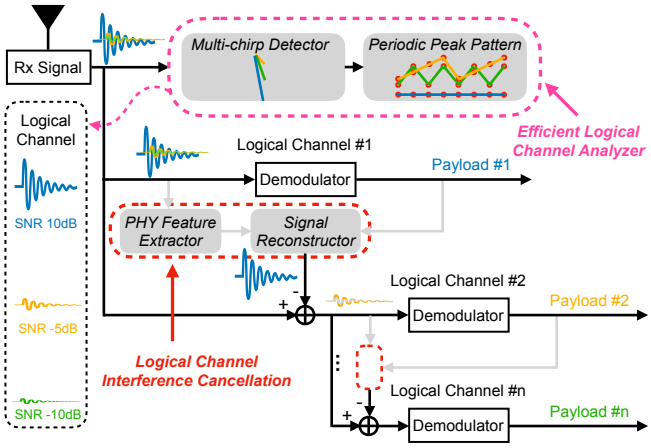

Fig. 11. System architecture of *Canas*: *Canas* firstly analyze the logical channel activities and then iteratively demodulates, reconstructs, and removes individual logical channels from the Rx signal.

$\hat{h_1^j}$ 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 $S_2(t)$.

C. Efficient Logical Channel Analyzer

Canas employs an iterative method to receive concurrent logical channels, prioritizing them from strongest to weakest. However, a challenge arises in initiating the iterative cancellation process. How can we efficiently identify activities across multiple concurrent logical channels and accurately assess their SNRs? The conventional LoRa gateway [67] uses cross-correlation to scan all SFs for detecting concurrent logical channels. This approach has two main drawbacks: Firstly, performing multiple rounds of correlation results in significant computational overhead. Secondly, the varying SNR gains introduced by different SFs after correlation complicate the fair comparison and further sorting of concurrent logical channels based on their SNRs.

To address these challenges, *Canas* introduces an efficient logical channel analyzer with enhanced detection capabilities to minimize latency. Figure 10(a) illustrates the input signal comprising three concurrent logical channels. Unlike traditional receivers that rely on a single local chirp for detection (*i.e.*, $C^{-1}(k,t)$ in Equation 3), *Canas* employs a synthetic signal composed of multiple down-chirps, as depicted in Figure 10(b). This approach allows *Canas* to detect multiple logical channels simultaneously with just one round of calculation (see Figure 10(c)). This method also ensures a consistent SNR gain across all logical channels by aggregating chirp signals of the same length. A potential question arises regarding how to determine the specific SF, given that all peaks cluster within the same window without distinct features. To resolve this ambiguity, we leverage the fact that chirps with different SFs have varying symbol durations. *Canas* then slides the detection window and captures all detected peaks into a matrix (see Figure 10(d)). Since the chirp durations are $1\times$, $2\times$, and $4\times$ the window length, respectively, we can identify a peak pattern with consistent periodicity in the matrix. As shown in Figure 10(e), this extracted pattern periodicity uniquely corresponds to the specific SF in the original signal. In practice, LoRa nodes typically use SF7 to SF12. With this accelerated method, *Canas* can divide the concurrent logical channels into subgroups and detect them simultaneously.

One may be concerned that the aggregate downchirp will also raise the noise floor as shown in Figure 10(c), introducing side effects for the concurrent packet detection. However, the method is still effective in practical experiments. The rationale is that the signal that leads to the logical channel interference problem tends to be a high SNR (*i.e.*, above the noise floor) as illustrated in Figure 4. The original signal power plus the demodulation gain will not be submerged by the slightly increased noise floor.

D. Canas: Put All Together

Figure 11 shows the general workflow of concurrent logical channel reception in *Canas*.

Logical channel detection and interference analysis. *Canas* firstly detect the channel activities on all logical channels and identify if there is logical channel interference. Specifically, we divide the logical channels into two subgroups (*i.e.*, SF7$\sim$9 and SF10$\sim$12) and apply the efficient logical channel analyzer. This method not only accelerates the detection process but also 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 bootstrap the logical channel interference cancellation process.

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. The received packets are initially sorted based on their SNRs, and the logical channels are demodulated and canceled from the strongest channel to the weakest. The weak

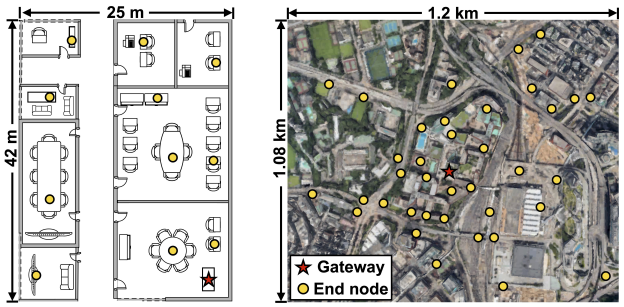

(a) Indoor testbed (b) Outdoor testbed

Fig. 12. Testbed setting of *Canas*.

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 11) with SOTA decoder (*e.g.*, CIC [23]) to handle collisions. Then, Canas reconstructs and cancels each colliding packet in the logical channel.

V. Evaluation

A. Methodology

Implementation. We implement *Canas* based on the software-defined radio platform (USRP N210) and `gr-lora` project [68]. 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 [26] 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 $1.08\,\mathrm{km} \times 1.2\,\mathrm{km}$ that represents a typical urban environment. The gateways are installed on the rooftop of a 20-storey building, and all nodes operate in the $915\,\mathrm{MHz}$ 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 12). All evaluations with $\leq$50 nodes are conducted via real-world experiments in the testbed. Additionally, we also perform large-scale trace-driven evaluations using collected data traces. Our evaluation. 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?

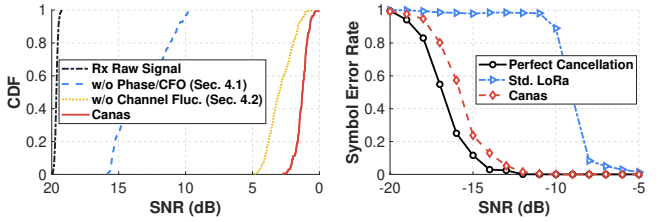

(a) Cancellation performance (b) Weak packet reception

Fig. 13. Performance of logical channel interference cancellation: (a) Signal strength of strong logical channel. (b) Packet reception of weak logical channel.

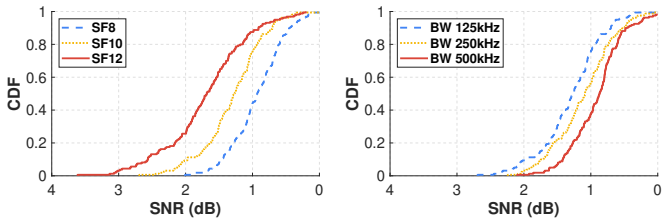

(a) SNR vs. SF (b) SNR vs. BW

Fig. 14. Performance of logical channel interference cancellation of different (a) spreading factor and (b) bandwidth.

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.

B. 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.*, $-15\,\mathrm{dB}$ and $20\,\mathrm{dB}$) using SF10 and SF7, respectively. Figure 13(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 $20\,\mathrm{dB}$) 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 $1.3\,\mathrm{dB}$ 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 $12.9\,\mathrm{dB}$ and $3.1\,\mathrm{dB}$ without canceling phase drift, frequency offset, and calibrating channel fluctuation, respectively (see Figure 13(a)).

Ultra-weak packet reception under interference. We next investigate the Rx sensitivity of *Canas* after logical channel interference cancellation. Specifically, we measure the received packets of the SF10 node under different SNR

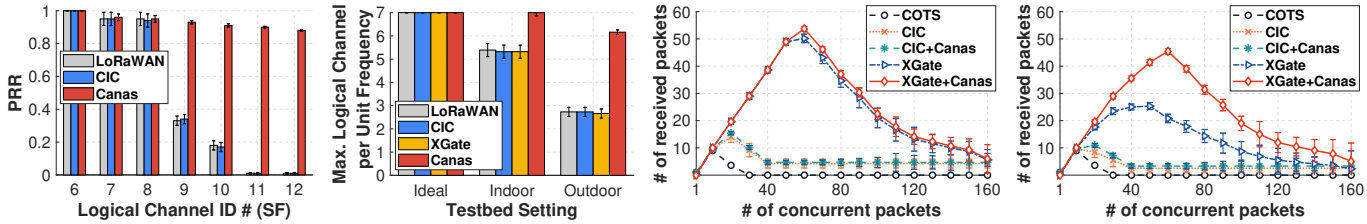

(a) Individual logical channel (b) Logical channel concurrency (c) Overall Concurrency (Indoor) (d) Overall Concurrency (Outdoor)

Fig. 15. 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.

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 13(b), the perfect cancellation performs best as it has no interference from other logical channels. LoRaWAN [69], 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 %.

Impact of bandwidth and spreading factor. This experiment examines how different transmission parameters influence the effectiveness of logical channel interference cancellation. We configure a strong LoRa node to transmit using various bandwidths and spreading factors, alongside a weak LoRa node. For each parameter combination, we analyze the SNR after interference cancellation across more than 100 packets. The initial SNR of the strong node (without cancellation) is measured at 20 dB, while the weak node's SNR is −15 dB. First, we fix the bandwidth at 125 kHz and vary the spreading factor of the strong node, with the weak node set to SF11. As shown in Figure 14(a), the median SNR after cancellation is 0.9 dB, 1.3 dB, and 1.7 dB for SF8, SF10, and SF12, respectively. We observe that cancellation performance degrades (reflected by higher post-cancellation SNR) as the spreading factor increases. This is because higher spreading factors lead to longer symbol durations, which introduce greater uncertainty in signal reconstruction and cancellation. Next, we fix the spreading factor of the strong node at SF10 and vary its bandwidth. The weak node is set to SF11 and its bandwidth is matched to that of the strong node to ensure both use the same spectrum. As illustrated in Figure 14(b), the median SNR after cancellation is 1.3 dB, 1.1 dB, and 0.9 dB for bandwidths of 125 kHz, 250 kHz, and 500 kHz, respectively. This improvement is attributed to the reduction in symbol duration as bandwidth increases.

C. 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 [69] and CIC [23]. 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 15(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 [23] and XGate [28]. 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 testbed: (1) Ideal, where nodes transmit independently; (2) Indoor; and (3) Outdoor, representing different environmental conditions (see Figure 12).

Figure 15(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

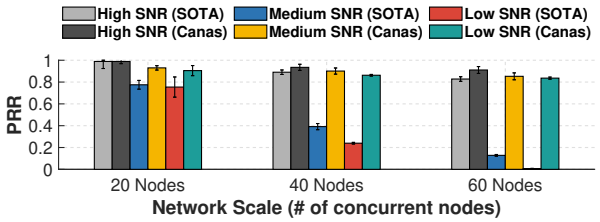

Fig. 16. Near-far effect at different network scales.

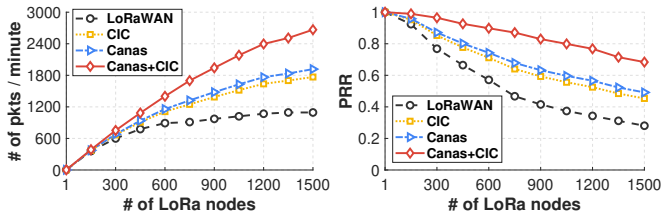

(a) Throughput (b) Packet reception ratio

Fig. 17. Performance of supporting massive LoRa nodes in large-scale deployments.

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 [23] and XGate [28] within the same 1.6 MHz spectrum and 9 Rx chains as the COTS gateway [67]. 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. 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.

Figure 15(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 15(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 [67]. The SOTA combines XGate [28] and CIC [23]. 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 16 illustrates the packet reception ratio (PRR) for the three SNR regimes as received by SOTA and *Canas*. 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 [67], 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 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

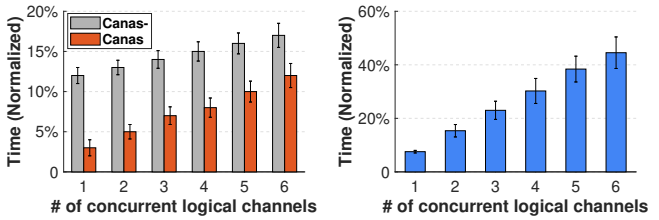

(a) Packet Detection (b) Interference Cancellation

Fig. 18. Time overhead of (a) logical channel detection; and (b) logical channel interference cancellation. The time is normalized the the percentage of the packet duration to showcase the real-time performance.

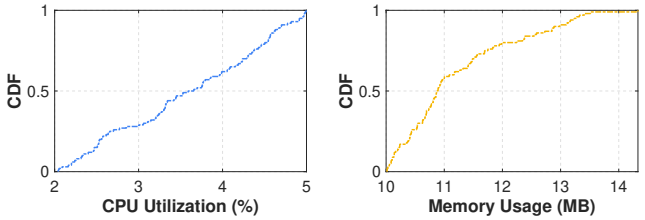

(a) CPU (b) Memory

Fig. 19. CPU utilization and memory usage when running *Canas* on a laptop.

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 [67] for reception; and (2) CIC, which implements collision recovery [23] upon LoRaWAN.

Figure 17 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 17(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 [20], [22], [23]. *Canas* can jointly work with CIC to support >800 LoRa nodes with PRR>85% (see Figure 17(b)).

D. Microbenchmarks

Time overhead of logical channel analyzer. *Canas* introduces an efficient logical channel analyzer designed to speed up the detection of concurrent logical channels. This experiment evaluates its time overhead and compares the performance gain against the traditional detection method. We deployed six LoRa nodes to transmit simultaneously on different logical channels. The traditional method, referred to as *Canas*-, involves iteratively correlating the received signal with six distinct base up-chirps (from SF7 to SF12) to detect the logical channels. In contrast, *Canas* groups the logical channels into two subgroups (SF7 to SF9 and SF10 to SF12) and uses an aggregated down-chirp for detection. Figure 18(a) illustrates the time overhead for different numbers of concurrent nodes. We normalized the overhead to the packet duration ratio to evaluate real-time performance. The results show that *Canas* reduces the overhead to less than 14% of the packet duration under extreme concurrency. The traditional method, *Canas*-, incurs high overhead because it requires six rounds of cross-correlation. In contrast, *Canas* accelerates the process by detecting multiple logical channels in a single round, with only a slight additional overhead needed to determine the peak pattern.

Time overhead of interference cancellation. 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 18(b). Importantly, even when canceling six logical channels, the overhead remains below 50% of packet duration. The latency per iteration is around 8% of the packet duration, which is much shorter than the airtime of LoRa packets and can be effortlessly handled by the gateway processors. We also evaluate the CPU and memory utilization when executing *Canas* on a laptop. We cut down all the other applications and only leave *Canas* for faithful measurement. As shown in Figure 19(a,b), the average CPU utilization and memory usage are 3.5% and 12 MB, respectively.

VI. Discussion

Interaction with PHY Error Correction. LoRa's physical layer employs forward error correction (FEC) by using a coding rate that adds redundancy to transmitted data, enabling the correction of a certain proportion of symbol errors (e.g., less than 20% symbol errors for a coding rate of 4/5). This FEC mechanism can tolerate limited logical channel interference, such as when a strong SF7 packet overlaps with only a few symbols of an SF12 packet due to their differing symbol durations. However, FEC alone is insufficient to fully address the problem. First, SF12 packets have long airtimes and can be disrupted by multiple low-SF packets, causing cumulative symbol errors that exceed the correction capability. Second, adjacent SFs with similar airtimes can interfere with each other, resulting in symbol error rates that surpass the 20% correction threshold. To effectively mitigate logical channel interference, FEC should be complemented by *Canas*. For example, *Canas* can reduce the symbol error rate of weak packets from over 30% to below 10%, allowing the FEC algorithm to correct the remaining errors.

Handling Hybrid Interference. Both *Canas* and state-of-the-art (SOTA) collision recovery techniques, such as CIC, are designed to address interference when multiple nodes transmit simultaneously on the same channel frequency. While SOTA methods can resolve conventional collisions (same BW and SF), they are ineffective against logical channel interference (same BW, different SF). An additional challenge arises with

hybrid interference, where traditional packet collisions and logical channel interference occur together. Preliminary results indicate that combining Canas with CIC can partially mitigate this issue. However, the effectiveness of both Canas and CIC decreases under conditions of high hybrid concurrency. The combined performance improvement of Canas+CIC is less than the sum of their individual benefits, as shown in figure 15(c,d). Future work will focus on developing unified and robust approaches to comprehensively address hybrid interference scenarios.

Uniqueness of Iterative Cancellation in LoRa. LoRa possesses unique characteristics that allow it to overcome several common challenges faced by traditional iterative cancellation schemes. (1) Similar signal strength: LoRa provides a distinct SNR gain after demodulation, making its channel estimation method robust even under close-SNR conditions, as described in Equation 7. (2) Overlapping symbol boundaries: *Canas* addresses logical channel interference, where concurrent packets are modulated with different spreading factors and symbol durations. Overlapping symbol boundaries do not cause ambiguity, since each packet uses a different base signal for demodulation. (3) Doppler frequency shift: Doppler frequency shift poses minimal issues for LoRa, as it operates in the sub-GHz band. The resulting frequency drift is much smaller than the guard bin bandwidth in the frequency domain, even when nodes are deployed on high-mobility platforms such as drones.

Overheads of *Canas*. *Canas* significantly enhances the reception reliability of weak packets (*i.e.*, those from distant LoRa nodes) in large-scale deployments. However, this improvement comes with increased computational overhead due to the iterative cancellation algorithms. While *Canas* imposes higher processing overheads at the gateway, it is capable of operating in real time, as demonstrated in Figure 18. Additionally, *Canas* requires access to raw physical layer signals for processing, a feature not currently supported by commodity LoRa gateways. Nonetheless, *Canas* can be seamlessly integrated with existing LoRa nodes without necessitating any hardware modifications. Future work would include more real-time optimization to reduce the memory usage and implement *Canas* on embedded processors.

VII. 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.

References

[1] K. Yang and W. Du, "A low-density parity-check coding scheme for lora networking," *ACM Transactions on Sensor Networks*, 2024.

[2] Y. Ren, P. Cai, J. Jiang, J. Du, and Z. Cao, "Prism: High-throughput lora backscatter with non-linear chirps," in *IEEE INFOCOM*, 2023, pp. 1–10.

[3] J. Jiang, Z. Xu, F. Dang, and J. Wang, "Long-range ambient lora backscatter with parallel decoding," in *ACM MobiCom*, 2021, pp. 684–696.

[4] N. Hou, X. Xia, and Y. Zheng, "Cloaklora: A covert channel over lora phy," *IEEE/ACM Transactions on Networking*, 2022.

[5] N. Hou, X. Xia, Y. Wang, and Y. Zheng, "One shot for all: Quick and accurate data aggregation for lpwans," *IEEE/ACM Transactions on Networking*, 2024.

[6] L. Shen, Q. Yang, K. Cui, Y. Zheng, X.-Y. Wei, J. Liu, and J. Han, "Fedconv: A learning-on-model paradigm for heterogeneous federated clients," in *ACM MobiSys*, 2024, pp. 398–411.

[7] K. Wang, Z. Zhou, and Z. Li, "Latte: Layer algorithm-aware training time estimation for heterogeneous federated learning," in *ACM MobiCom*, 2024, pp. 1470–1484.

[8] K. Cui, Q. Yang, L. Shen, Y. Zheng, F. Xiao, and J. Han, "Towards isac-empowered mmwave radars by capturing modulated vibrations," *IEEE Transactions on Mobile Computing*, 2024.

[9] J. Cao, J. Chen, C. Lin, Y. Liu, K. Wang, and Z. Li, "Practical gaze tracking on any surface with your phone," *IEEE Transactions on Mobile Computing*, 2024.

[10] S. Ji, X. Zhang, Y. Zheng, and M. Li, "Construct 3d hand skeleton with commercial wifi," in *ACM SenSys*, 2023, pp. 322–334.

[11] Q. Yang and Y. Zheng, "Aquahelper: Underwater sos transmission and detection in swimming pools," in *ACM SenSys*, 2023, pp. 294–307.

[12] K. Wang, J. Cao, Z. Zhou, and Z. Li, "Swapnet: Efficient swapping for dnn inference on edge ai devices beyond the memory budget," *IEEE Transactions on Mobile Computing*, vol. 23, no. 9, pp. 8935–8950, 2024.

[13] K. Yang, Y. Chen, and W. Du, "Orchloc: In-orchard localization via a single lora gateway and generative diffusion model-based fingerprinting," in *ACM MobiSys*, 2024.

[14] Y. Ren, W. Sun, J. Du, H. Zeng, Y. Dong, M. Zhang, S. Chen, Y. Liu, T. Li, and Z. Cao, "Demeter: Reliable cross-soil lpwan with low-cost signal polarization alignment," in *ACM MobiCom*, 2024, pp. 230–245.

[15] Y. Chen, K. Yang, Z. An, B. Holder, L. Paloutzian, K. Bali, and W. Du, "Marlp: Time-series forecasting control for agricultural managed aquifer recharge," in *ACM KDD*, 2024.

[16] Y. Cheng, H. Saputra, L. M. Goh, and Y. Wu, "Secure smart metering based on lora technology," in *2018 IEEE 4th International Conference on Identity, Security, and Behavior Analysis (ISBA)*, 2018, pp. 1–8.

[17] Y. Ren, A. Gamage, L. Liu, M. Li, S. Chen, Y. Dong, and Z. Cao, "Sateriot: High-performance ground-space networking for rural iot," in *ACM MobiCom*, 2024.

[18] Semtech. (2024, Nov.) Lora. "https://www.semtech.com/lora".

[19] S. Tong, J. Wang, J. Yang, Y. Liu, and J. Zhang, "Citywide lora network deployment and operation: Measurements, analysis, and implications," in *ACM SenSys*, 2023, pp. 362–375.

[20] X. Xia, Y. Zheng, and T. Gu, "FTrack: parallel decoding for LoRa transmissions," in *ACM SenSys*. ACM, Nov. 2019, pp. 192–204.

[21] S. Tong, J. Wang, and Y. Liu, "Combating packet collisions using non-stationary signal scaling in LPWANs," in *ACM MobiSys*. Toronto Ontario Canada: ACM, Jun. 2020.

[22] S. Tong, Z. Xu, and J. Wang, "CoLoRa: Enabling Multi-Packet Reception in LoRa," in *IEEE INFOCOM*, Jul. 2020.

[23] M. O. Shahid, M. Philipose, K. Chintalapudi, S. Banerjee, and B. Krishnaswamy, "Concurrent interference cancellation: Decoding multi-packet collisions in lora," in *ACM SIGCOMM*, 2021.

[24] F. Yu, X. Zheng, L. Liu, and H. Ma, "Enabling concurrency for non-orthogonal lora channels," in *ACM MobiCom*, 2023, pp. 1–15.

[25] Semtech. (2024, Mar.) Understanding adr. "https://lora-developers.semtech.com/documentation/tech-papers-and-guides/understanding-adr".

[26] ——. (2022, Dec.) Sx1276. "https://www.semtech.com/products/wireless-rf/lora-connect/sx1276".

[27] S. Gollakota and D. Katabi, "Zigzag decoding: Combating hidden terminals in wireless networks," in *ACM SIGCOMM*, 2008, pp. 159–170.

[28] S. Yu, X. Xia, N. Hou, Y. Zheng, and T. Gu, "Revolutionizing lora gateway with xgate: Scalable concurrent transmission across massive logical channels," in *ACM MobiCom*, 2024, pp. 482–496.

[29] R. Eletreby, D. Zhang, S. Kumar, and O. Yağan, "Empowering low-power wide area networks in urban settings," in *ACM SIGCOMM*, 2017, pp. 309–321.
[30] Z. Xu, S. Tong, P. Xie, and J. Wang, "Fliplora: Resolving collisions with up-down quasi-orthogonality," in *IEEE SECON*, 2020, pp. 1–9.
[31] X. Xia, N. Hou, Y. Zheng, and T. Gu, "Pcube: scaling lora concurrent transmissions with reception diversities," in *ACM MobiCom*, 2021, pp. 670–683.
[32] Z. Xu, P. Xie, and J. Wang, "Pyramid: Real-time lora collision decoding with peak tracking," in *IEEE INFOCOM*. IEEE, 2021, pp. 1–9.
[33] F. Yu, X. Zheng, Y. Ma, L. Liu, and H. Ma, "Resolve cross-channel interference for lora," in *IEEE ICDCS*, 2024, pp. 1027–1038.
[34] S. Yu, Z. Zhang, X. Xia, Y. Zheng, and J. Wang, "Are lora logical channels really orthogonal? practically orthogonalizing massive logical channels," in *ACM MobiSys*, 2025.
[35] G. Deepthy and R. Susan, "Analysis of successive interference cancellation in cdma systems," in *IEEE ICACCTech*. IEEE, 2012, pp. 481–485.
[36] S. Sen, N. Santhapuri, R. R. Choudhury, and S. Nelakuditi, "Successive interference cancellation: Carving out mac layer opportunities," *IEEE Transactions on Mobile Computing*, vol. 12, no. 2, pp. 346–357, 2012.
[37] P. Patel and J. Holtzman, "Analysis of a simple successive interference cancellation scheme in a ds/cdma system," *IEEE journal on selected areas in communications*, vol. 12, no. 5, pp. 796–807, 1994.
[38] J. G. Andrews and T. H.-Y. Meng, "Performance of multicarrier cdma with successive interference cancellation in a multipath fading channel," *IEEE Transactions on Communications*, vol. 52, no. 5, pp. 811–822, 2004.
[39] N. Hou, Y. Wang, X. Xia, S. Yu, Y. Zheng, and T. Gu, "Molora: Intelligent mobile antenna system for enhanced lora reception in urban environments," in *ACM SenSys*, 2025, pp. 424–436.
[40] R. Li, Z. Zhang, X. Xia, N. Hou, W. Chai, S. Yu, Y. Zheng, and T. Gu, "From interference mitigation to toleration: Pathway to practical spatial reuse in lpwans," *ACM MobiCom*, 2025.
[41] S. Tong, Z. Shen, Y. Liu, and J. Wang, "Combating link dynamics for reliable lora connection in urban settings," in *ACM MobiCom*, 2021.
[42] N. Hou, X. Xia, and Y. Zheng, "Don't miss weak packets: Boosting lora reception with antenna diversities," *ACM Trans. Sen. Netw.*, vol. 19, no. 2, feb 2023.
[43] Z. Xu, P. Xie, J. Wang, and Y. Liu, "Ostinato: Combating lora weak links in real deployments," in *IEEE ICNP*, 2022, pp. 1–11.
[44] X. Xia, Q. Chen, N. Hou, Y. Zheng, and M. Li, "Xcopy: Boosting weak links for reliable lora communication," in *ACM MobiCom*, 2023, pp. 1–15.
[45] R. Trüb, R. Da Forno, A. Biri, J. Beutel, and L. Thiele, "Lsr: Energy-efficient multi-modulation communication for inhomogeneous wireless iot networks," *ACM Trans. Internet Things*, vol. 4, no. 2, apr 2023.
[46] C. Shao and O. Muta, "Tonari: Reactive detection of close physical contact using unlicensed lpwan signals," *ACM Trans. Internet Things*, vol. 5, no. 2, apr 2024.
[47] A.-U.-H. Ahmar, E. Aras, T. D. Nguyen, S. Michiels, W. Joosen, and D. Hughes, "Design of a robust mac protocol for lora," *ACM Transactions on Internet of Things*, vol. 4, no. 1, pp. 1–25, 2023.
[48] C. Li, X. Guo, L. Shangguan, Z. Cao, and K. Jamieson, "CurvingLoRa to boost LoRa network throughput via concurrent transmission," in *USENIX NSDI*. USENIX Association, 2022, pp. 879–895.
[49] P. Xie, Y. Li, Z. Xu, Q. Chen, Y. Liu, and J. Wang, "Push the limit of lpwans with concurrent transmissions," in *IEEE INFOCOM*. IEEE, 2023, pp. 1–10.
[50] C. Li, Y. Ren, S. Tong, S. I. Siam, M. Zhang, J. Wang, Y. Liu, and Z. Cao, "Chirptransformer: Versatile lora encoding for low-power wide-area iot," in *ACM MobiSys*, 2024, pp. 479–491.
[51] A. Gamage, J. C. Liando, C. Gu, R. Tan, and M. Li, "Lmac: efficient carrier-sense multiple access for lora," in *ACM MobiCom*, 2020.
[52] J. Luo, Z. Xu, J. Lin, C. Chen, and R. Xiong, "Ch-mac: Achieving low-latency reliable communication via coding and hopping in lpwan," *ACM Transactions on Internet of Things*, vol. 4, no. 4, pp. 1–25, 2023.
[53] S. Yu, X. Xia, Z. Zhang, N. Hou, and Y. Zheng, "Fdlora: Tackling downlink-uplink asymmetry with full-duplex lora gateways," in *ACM SenSys*, 2024, pp. 281–294.
[54] ——, "Fdlora: Scaling downlink concurrent transmissions with full-duplex lora gateways," *IEEE Transactions on Mobile Computing*, 2025.
[55] X. Xia, Q. Chen, N. Hou, and Y. Zheng, "Hylink: Towards high throughput lpwans with lora compatible communication," in *ACM SenSys*, 2022, pp. 578–591.
[56] H. Pirayesh, S. Zhang, P. K. Sangdeh, and H. Zeng, "Maloragw: Multi-user mimo transmission for lora," in *ACM SenSys*, 2022, pp. 179–192.
[57] C. Li and Z. Cao, "Lora networking techniques for large-scale and long-term iot: A down-to-top survey," *ACM Computing Surveys (CSUR)*, vol. 55, no. 3, pp. 1–36, 2022.
[58] J. C. Liando, A. Gamage, A. W. Tengourtius, and M. Li, "Known and unknown facts of lora: Experiences from a large-scale measurement study," *ACM Transactions on Sensor Networks (TOSN)*, vol. 15, no. 2, pp. 1–35, 2019.
[59] Y. Ren, L. Liu, C. Li, Z. Cao, and S. Chen, "Is lorawan really wide? fine-grained lora link-level measurement in an urban environment," in *IEEE ICNP*, 2022, pp. 1–12.
[60] K. Yang, Y. Chen, X. Chen, and W. Du, "Link quality modeling for lora networks in orchards," in *IEEE/ACM IPSN*, 2023, p. 27–39.
[61] F. Yu, X. Zheng, L. Liu, and H. Ma, "LoRadar: An Efficient LoRa Channel Occupancy Acquirer based on Cross-channel Scanning," in *IEEE INFOCOM*, May 2022, pp. 540–549.
[62] R. Kufakunesu, G. P. Hancke, and A. M. Abu-Mahfouz, "A survey on adaptive data rate optimization in lorawan: Recent solutions and major challenges," *Sensors*, vol. 20, no. 18, p. 5044, 2020.
[63] X. Xia, Y. Zheng, and T. Gu, "Litenap: Downclocking lora reception," *IEEE/ACM Transactions on Networking*, vol. 29, no. 6, pp. 2632–2645, 2021.
[64] Z. Xu, S. Tong, P. Xie, and J. Wang, "From demodulation to decoding: Toward complete lora phy understanding and implementation," *ACM Trans. Sen. Netw.*, vol. 18, no. 4, jan 2023.
[65] G. Shen, J. Zhang, A. Marshall, L. Peng, and X. Wang, "Radio frequency fingerprint identification for lora using spectrogram and cnn," in *IEEE INFOCOM*. IEEE, 2021, pp. 1–10.
[66] N. Hou, X. Xia, and Y. Zheng, "Jamming of lora phy and countermeasure," in *IEEE INFOCOM*. IEEE, 2021, pp. 1–10.
[67] Semtech. (2022, Oct.) Sx1301. "https://www.semtech.com/products/wireless-rf/lora-core/sx1301".
[68] Gr-LoRa GitHub community. (2021, Jul) gr-lora projects. "https://github.com/rpp0/gr-lora".
[69] L. Alliance. (2022, Oct.) Lorawan specification. "https://lora-alliance.org/about-lorawan".

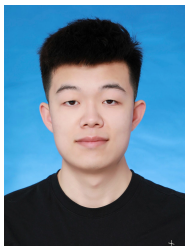

Shiming Yu (Student Member, IEEE) received the B.E. degree from University of Electronic Science and Technology of China in 2022. He is currently a PhD candidate in the Department of Computing, The Hong Kong Polytechnic University. His research interests include low-power wide-area networks, wireless networking, mobile computing.

Ziyue Zhang (Student Member, IEEE) received the B.E. degree from University of Electronic Science and Technology of China in 2023. He is currently a PhD student in the Department of Computing, The Hong Kong Polytechnic University. His research interests include low-power-wide-area network, satellite networks, mobile computing.

Xianjin Xia (Member, IEEE) received the B.S, M.Sc and PhD degrees in Computer Science from Northwestern Polytechnical University, Xi'an, China, in 2010, 2013 and 2018, respectively. He is currently a Research Assistant Professor (RAP) in the Department of Computing, The Hong Kong Polytechnic University. His research interests include low-power wide-area networks, localization, mobile computing.

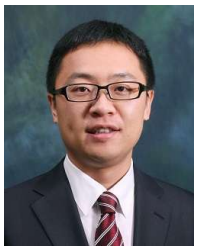

Yuanqing Zheng (Senior Member, IEEE) received the B.S. degree in Electrical Engineering and the M.E. degree in Communication and Information System from Beijing Normal University, Beijing, China, in 2007 and 2010 respectively. He received the PhD degree in School of Computer Engineering from Nanyang Technological University in 2014. He is currently an Associate Professor with the Department of Computing in Hong Kong Polytechnic University. His research interests include Wireless Networking and Mobile Computing, Ubiquitous Computing, Internet of Things, and Embedded AI. He received the Best Paper Award in IEEE INFOCOM 2020, and the Test of Time Award in ACM SenSys 2022. He served on the editorial board of IEEE Transactions on Wireless Communications (TWC), ACM Transactions on Sensor Networks (TOSN), Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT), and Computer Networks.

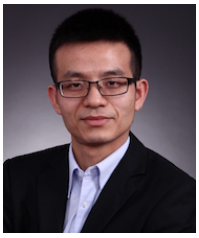

Jiliang Wang (Senior Member, IEEE) received the B.E. degree in computer science and technology from the University of Science and Technology of China and the PhD degree in computer science and engineering from the Hong Kong University of Science and Technology. He is currently an associate professor with the School of Software, Tsinghua University. His research interests include Internet of Things and mobile computing.