TriLoRa: Improving LoRa PHY Data Rate for High-Density LoRaWAN

Yanbo Zhang*, Qirui Liu†, Songfan Li*, Mo Li*
* *The Hong Kong University of Science and Technology*, Hong Kong SAR, China
† *Peking University*, Beijing, China
yanbozhang@ust.hk, qirui_liu@pku.edu.cn, lisongfan@ust.hk, lim@cse.ust.hk

***Abstract*—As LoRaWAN evolves toward high-density deployments, the rigid, low-data-rate nature of its physical layer becomes a critical bottleneck. This paper introduces TriLoRa, a novel PHY design that significantly enhances throughput via time-domain symbol superposition. The core innovation is Triangle-Chirp Modulation (TCM), which exploits the symmetric spectral features of triangle waveforms to resolve the inherent time-frequency ambiguity of standard linear chirps. By decoupling time delays from frequency encoding, TriLoRa allows multiple symbols to be densely superimposed and reliably distinguished within a single symbol window, effectively converting excess Signal-to-Noise Ratio (SNR) into parallel transmission capacity. Centered on this, we design a hierarchical demodulation architecture and a fine-grained rate adaptation mechanism that maintain robustness even in negative SNR regimes. Experimental results demonstrate that TriLoRa improves point-to-point data rates by up to 30×. Furthermore, trace-driven evaluations based on a commercial-scale campus network reveal a 3.96× gain in aggregate network capacity.**

***Index Terms*—IoT, LPWAN, LoRaWAN, Communication System, Signal Processing, Rate Adaptation**

I. Introduction

Low-Power Wide-Area Network (LPWAN) technologies, particularly LoRaWAN, have become the de facto standard for massive IoT connectivity due to their robust long-range capabilities. While originally architected for wide-area coverage, empirical evidence indicates a paradigm shift in real-world deployments toward high-density [1–5]. This trend is most pronounced in critical vertical domains. In smart utility sectors, electricity metering densities in urban centers are estimated to reach 12,000 nodes/km^2 [1] to enable granular energy monitoring. An experimental deployment at ETH Zurich integrated 390 battery-powered sensors within a single multi-story building, using dense gateway placement (one every 30 meters) to maintain connectivity [4].

To characterize the network dynamics inherent to high-density environments, we conducted a measurement study based on operational traces from a representative campus-scale LoRaWAN deployment. This network covers a 0.6 km^2 area with over 200 gateways serving more than 3,000 nodes for diverse applications ranging from utility metering to environmental sensing. We analyzed a one-week trace containing channel quality metadata (e.g., SNR, RSSI) for over 9 million packets transmitted by 3,227 distinct end-devices and processed by 44 active gateways. As illustrated in Fig. 1, the Signal-to-Noise Ratio (SNR) reveals a long-tail distribution where a substantial fraction of nodes maintain robust link quality (e.g., SNR $>$ 0 dB). This observation is consistent with prior measurements in high-density deployments [5] and stems from the gateway over-provisioning strategies often employed to guarantee deep indoor coverage (e.g., for meters in basements) [6–8]. Consequently, while this redundancy secures connectivity for hard-to-reach nodes, it ensures that the vast majority of devices operate under favorable channel conditions well above the sensitivity floor. However, the current LoRa physical layer is ill-equipped to exploit this channel diversity. The standard defines a rigid set of data rates governed by six spreading factors (SF7 to SF12), among which the most aggressive standard rate (SF7) still operates far below the theoretical channel capacity, leaving significant spectral resources underutilized. In high-density networks, this inefficiency translates directly into prolonged airtime, which exacerbates packet collisions, degrades network capacity, and

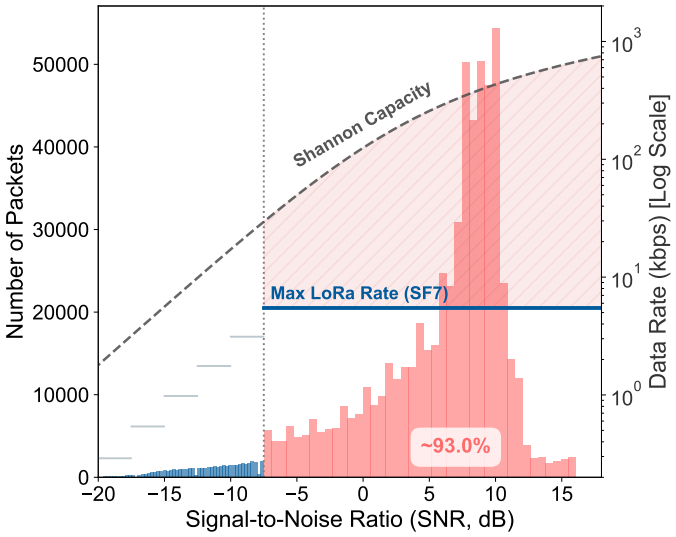


Fig. 1: Analysis of a one-week campus LoRaWAN trace reveals a significant spectrum efficiency gap. With 93% of packets exceeding -7.5 dB SNR, standard LoRa's 5.47 kbps cap (SF7) falls far short of the theoretical Shannon capacity.

This work was supported by the Hong Kong GRF under Grants 16210425 and C6086-25G, and by the National Natural Science Foundation of China under Grant 62502012. Mo Li is the corresponding author.

shortens battery life.

In this paper, we propose *TriLoRa*, a communication system tailored for high-density LoRaWAN environments. To transcend the data rate limitations of the standard physical layer, TriLoRa introduces a *time-domain degree of freedom* to the modulation process. By imposing controlled time shifts on superimposed symbols, the system exploits the intrinsic time-frequency coupling of chirp signals to translate temporal offsets into separable frequency signatures. This mechanism enables the concurrent transmission and reliable demodulation of multiple symbols within a single chirp window, thereby maximizing the spectral efficiency of the constrained channel. Note that this throughput gain is not free: superimposing N symbols divides the per-symbol energy, requiring proportionally higher SNR. TriLoRa therefore functions as an SNR-to-throughput converter: it harvests the excess link margin that is abundant in dense deployments and reinvests it into parallel capacity, while gracefully degrading to standard LoRa ($N = 1$) when the margin is insufficient.

Integrating this idea into a practical LoRa system is non-trivial. Specifically, we identify the following three challenges: (1) Time-Frequency Ambiguity: Standard LoRa employs linear up-chirps where time and frequency are linearly coupled. Consequently, a transmission time delay (Δt) appears identically to a frequency shift (Δf) at the receiver. This ambiguity prevents the receiver from distinguishing whether a frequency shift is caused by a modulated symbol value or a propagation delay, rendering parallel demodulation impossible. (2) Self-Interference: Superimposing multiple symbols within a single window creates significant Inter-Symbol Interference (ISI). Furthermore, since the total transmission power is split among multiple overlapped symbols, the energy per symbol decreases. This imposes a strict requirement on the receiver to cancel interference and recover signals from an elevated noise floor, particularly under low-SINR conditions. (3) Sensitivity to Hardware Impairments: The precise frequency features utilized by TriLoRa are highly sensitive to real-world hardware imperfections. Carrier Frequency Offset (CFO) and Symbol Timing Offset (STO) can distort these spectral signatures. Without precise alignment and compensation, these errors disrupt the decoding logic, leading to demodulation failure.

Our design addresses the challenges through three corresponding techniques: (1) Triangle-Chirp Modulation: To resolve the time-frequency ambiguity, we introduce a symmetric waveform comprising up- and down-segments. Leveraging this symmetry, TriLoRa creates a unique *Frequency-Splitting Effect*, breaking the linear coupling and effectively decoupling time delays from symbol values, laying the foundation for distinguishable parallel transmission. (2) Two-Tier Demodulation: To tackle self-interference and power dispersion, we propose a hierarchical receiver integrating Frequency Matching (FM) and Cyclic Demodulation (CD). The system adaptively switches between FM for computational efficiency in benign conditions and CD for iterative interference cancellation in low-SNR regimes, maximizing decoding robustness. (3) Joint Error Compensation: To neutralize hardware sensitivity, we devise a pre-processing algorithm that exploits the standard preamble to jointly solve the frequency offset equations of the up- and down-chirps. This allows the receiver to simultaneously estimate and calibrate both CFO and STO prior to demodulation, ensuring the integrity of the spectral features.

Our main contributions are summarized as follows:

- We introduce TriLoRa, which leverages Triangle-Chirps to resolve the time-frequency ambiguity of standard LoRa. This innovation allows multiple symbols to be superimposed in the time domain to multiply throughput.
- We present a holistic system design that tightly couples a two-tier demodulation engine with real-time rate adaptation and hardware error cancellation. This ensures robust decoding capabilities even in harsh SNR conditions and under significant hardware-induced spectral distortion.
- Through comprehensive experiments based on the commercial scale network operation traces, we demonstrate an improvement of $30\times$ data rate and $3.96\times$ network capacity, validating the significant performance leap.

II. Related Work

Efforts to transcend the intrinsic data rate limitations of LoRa PHY have spurred significant innovation across both commercial and academic domains. These endeavors can be broadly categorized into bandwidth expansion, coherent superposition, collision resolution, and modulation variants.

Bandwidth Expansion (Commercial Solutions). In the commercial sector, Semtech released the SX1280 chipset [9] to address high-throughput demands. Operating in the 2.4 GHz ISM band, it achieves data rates up to 254 kbps by utilizing a wide bandwidth of 1.6 MHz. However, this performance comes at the cost of communication range. The higher propagation loss at 2.4 GHz, combined with the wider noise bandwidth, significantly shrinks the coverage area compared to Sub-GHz LoRa, making it less suitable for long-range, obstacle-dense IoT scenarios.

Coherent Superposition. To boost data rates without sacrificing spectrum, academic works like HyLink [10] propose superimposing multiple packet payloads into a single transmission. By exploiting phase fingerprints, HyLink increases data rates nearly tenfold. However, such coherent approaches rely on the extraction of precise and stable channel phases over the packet duration. In practice, the low-cost oscillators in LoRa devices, clock drift and channel variations during long packets can destroy phase coherence, leading to decoding failures – a well-documented challenge given that low-cost LoRa oscillators exhibit carrier frequency offsets of up to several ppm, which drift significantly over long packet durations [11].

Collision Resolving and Concurrent Transmission. A substantial body of work focuses on disentangling collided signals to improve network-level throughput [12–22]. Systems like FTrack [12] and CoLoRa [19] utilize edge detection or time-offset features to separate concurrent transmissions. While these methods effectively recover data from random collisions and can operate across a wide SNR range, they are inherently

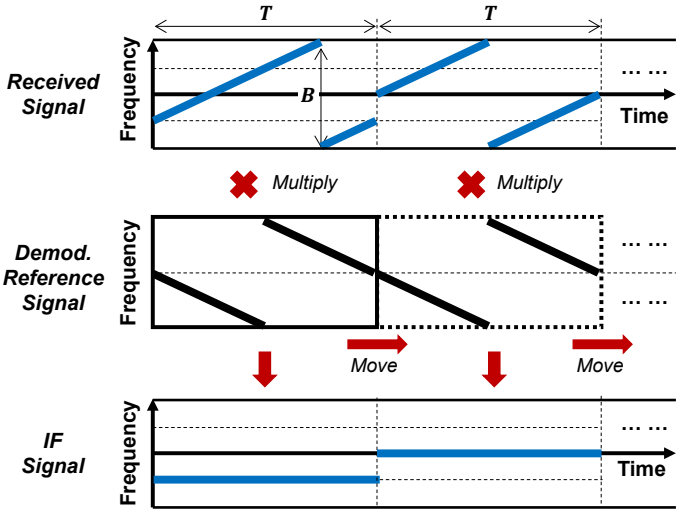


Fig. 2: Demodulation process of standard LoRa.

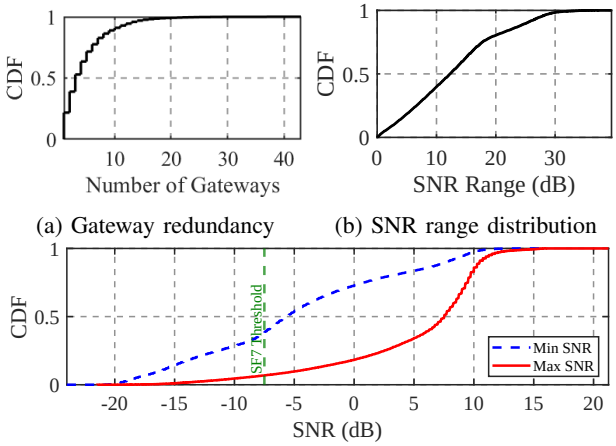


Fig. 3: Analysis of link quality diversity. (a) CDF of gateway redundancy. (b) Per-packet SNR gap distribution. (c) Min/Max SNR comparison.

reactive – improving aggregate network throughput by resolving unintended collisions, rather than *proactively* boosting the data rate of an individual link.

Modulation Variants. Other approaches explore altering the chirp waveform itself, encoding information into chirp slopes or interleaved spreading patterns [23–25]. While theoretically promising, these schemes typically introduce non-orthogonality issues and increased transceiver complexity, limiting the practical scalability of the data rate.

III. Background and Motivation

A. Primer on Standard LoRa PHY

LoRa employs Chirp Spread Spectrum (CSS) modulation, utilizing linear frequency chirps to encode data. A base up-chirp with bandwidth BW and duration $T = 2^{SF}/BW$ is defined by $f(t) = -\frac{BW}{2} + \frac{BW}{T}t$, where SF is the spreading factor. To modulate symbol value S, LoRa applies a cyclic frequency shift:

$$f(t,S) = \left(-\frac{BW}{2} + \frac{BW}{T}t + \frac{S}{2^{SF}}BW\right) \pmod{BW} \quad (1)$$

Demodulation is performed via a "de-chirp" operation as illustrated in Fig. 2, where the received signal is multiplied by a synchronized base down-chirp (the complex conjugate of the base up-chirp). Since the negative frequency slope of the down-chirp cancels the positive slope of the received signal, the multiplication yields a signal with a constant frequency trajectory, referred to as the Intermediate Frequency (IF) signal. A subsequent Fast Fourier Transform (FFT) concentrates the energy of this mono-tone signal into a distinct spectral peak, the index of which uniquely identifies the symbol value.

Crucially, this mechanism relies on strict time-orthogonality. As shown in Fig. 2, the de-chirp window must contain energy from strictly one symbol to ensure the IF signal remains a single pure tone. If symbols were to overlap, the de-chirp process would fail to focus the energy into a single bin, causing Inter-Symbol Interference (ISI). Consequently, standard LoRa enforces a sequential transmission, causing packet airtime to grow linearly with payload size ($N \times T$). This physical constraint caps throughput, preventing LoRa from utilizing good channel conditions.

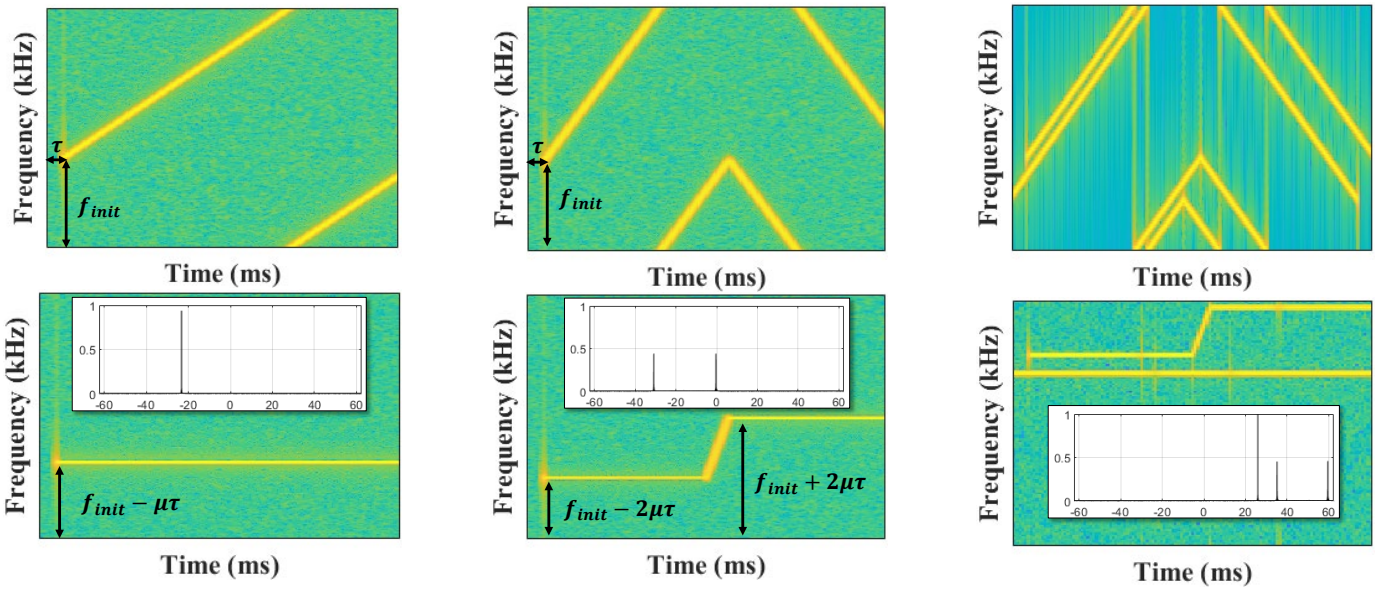
B. Motivation

While the robust design of standard LoRa is essential for edge-of-coverage connectivity, our measurements indicate that it is overly conservative given the massive spatial diversity inherent in high-density deployments.

We provides a detailed analysis of the link quality diversity of the high-density campus LoRaWAN that is introduced in Sec. I. As shown in Fig. 3a, the network exhibits significant macro-diversity: over 90% of packets are decoded by at least two gateways, with a substantial portion captured by 5 to 10 gateways simultaneously. This redundancy manifests as extreme variations in link quality. Fig. 3b reveals that for a single packet, the SNR disparity between the best and worst receiving gateways frequently exceeds 20 - 30 dB. Crucially, Fig. 3c exposes the structural inefficiency: while the minimum SNR (dashed line) often hovers near the sensitivity floor, the maximum SNR (solid line) consistently remains excellent ($\geq$ 10 dB). This implies that standard LoRa applies a worse-case protection strategy even to best-case links, unnecessarily sacrificing valuable airtime that could be reclaimed for higher throughput.

This under-utilization of channel resources is particularly critical as modern IoT applications increasingly demand bandwidths that legacy LoRa cannot provide. Emerging scenarios, such as transmitting voice commands for emergency response or uploading surveillance images for visual verification, generate bursty and data-intensive payloads that require throughput orders of magnitude higher than simple telemetry [10].

Therefore, to unlock the full potential of LoRa's underutilized resource dimensions, reduce channel inefficiency, and empower emerging high-throughput IoT applications, we propose a new transmission framework that exploits the abundant



(a) Standard LoRa: Ambiguity (b) Decoupling via Frequency Splitting (c) Overlapping and Energy Focusing

Fig. 4: The core insight: breaking time-frequency dependency. (a) In Standard LoRa, time offsets translate to frequency shifts, causing demodulation ambiguity. (b) In TriLoRa, time offsets trigger a distinct *Frequency Splitting* effect. (c) Under signal superposition, TriLoRa leverages *Energy Focusing*: the target symbol accumulates into a sharp peak, while interfering symbols are split and suppressed into the noise floor, maximizing the detection SNR.

SNR headroom to support parallel symbol transmissions via controlled time-domain symbol overlapping. However, implementing this concept with standard LoRa faces a fundamental challenge: as detailed in Section III-C, traditional linear chirps exhibit inherent ambiguity when overlapped, where the time offset of a symbol becomes indistinguishable from its encoded value. Consequently, resolving this time-offset ambiguity, namely making time offsets explicitly observable and decodable, emerges as the central challenge for enabling parallel transmissions in LoRa.

C. Our Insight

To effectively increase the data rate, our core strategy is to compress the inter-symbol airtime by enabling *parallel transmission*. Specifically, we explore the feasibility of overlapping multiple symbols within a single symbol time window while maintaining the functionality and robustness of the LoRa physical layer.

However, simply superposing standard linear LoRa chirp symbols introduces a fundamental challenge. Directly aligning the start times of multiple symbols results in the loss of sequential information, making it impossible to reconstruct the data order. Introducing a time delay τ between consecutive symbols to preserve order is also infeasible for standard LoRa. As illustrated in Fig. 4a, the linear nature of standard chirps translates a time delay directly into a frequency shift $\Delta f = -\mu\tau$ on the IF signal. Consequently, the demodulated peak frequency f_{peak} becomes:

$$f_{peak} = \frac{BW}{2^{SF}} S - \frac{BW}{T}\tau \tag{2}$$

This equation exposes a critical *linear coupling problem*: a single observed variable f_{peak} depends on two unknowns – the encoded symbol S and the time offset τ. This ambiguity is exacerbated in a multi-symbol superposition regime, rendering the system mathematically under-determined and unsolvable.

To resolve this ambiguity, we propose a novel use of the *Triangle-Chirp waveform* – composed of symmetric up- and down-chirp segments – for modulation and demodulation. Our key insight is that the intrinsic symmetry of the proposed Triangle-Chirp waveform breaks the linear time-frequency coupling by introducing a unique *Frequency-Splitting Effect*.

As illustrated in Fig. 4b, when a triangle-chirp is demodulated with a time offset τ, the delay induces opposite frequency shifts in the two segments of the IF signal: the Up-segment experiences a negative shift ($-2\mu\tau$), while the Down-segment experiences a positive shift ($+2\mu\tau$). These opposing shifts prevent the signal energy from converging into a single spectral peak. Instead, as shown in Fig. 4b (down), the energy splits into two distinct peaks separated by $4\mu\Delta t$. This physical phenomenon transforms the under-determined problem into a solvable one. By exploiting the splitting behavior, the system effectively decouples the time delay from the symbol frequency, providing theoretically sufficient information to jointly recover both S and τ from the overlapped symbols.

We further validate this insight in a superposition scenario where two triangle-chirp symbols overlap within a single window. Fig. 4c presents the STFT spectrogram of such an overlapped signal, showing two symbols separated by τ. The corresponding FFT spectrum of the demodulated IF signal (Fig. 4c) reveals a distinctive pattern: three peaks are observed – one high-amplitude peak and two lower-amplitude peaks (ap-

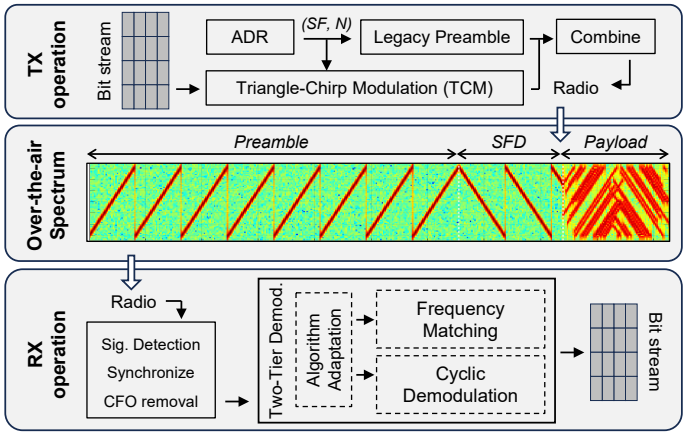


Fig. 5: System overview of TriLoRa.

proximately half the magnitude). This pattern emerges because the demodulation reference is time-aligned with the first symbol ($\tau \approx 0$). Consequently, the first symbol's energy is focused into a single peak, while the second symbol, being misaligned, manifests the *Frequency-Splitting Effect*, distributing its energy into two separate peaks. This observation confirms that the splitting effect persists under superposition. Crucially, the significant amplitude gap between the focused target symbol and the split interfering symbol provides a vital opportunity to suppress self-interference during the demodulation process (detailed in Section IV-B).

IV. Design

TriLoRa redesigns the LoRa physical layer by introducing Triangle-Chirp Modulation (TCM), a novel scheme designed to enable high-density symbol superposition in the time domain while preserving demodulation reliability. Unlike standard LoRa, which encodes information solely on up-chirps, TriLoRa utilizes symmetric triangle waveforms to create distinct frequency-time patterns that facilitate precise peak detection even under heavy overlap. The system architecture (as illustrated in Fig. 5) is composed of three integral modules: (1) a Transmitter that generates and superimposes multiple triangle-chirp symbols (N) within a single symbol window based on available link budget; (2) a Hierarchical Receiver featuring two collaborative demodulation strategies – Frequency Matching (FM) for low-complexity baseline decoding and Cyclic Demodulation (CD) for iterative interference cancellation in low-SNR scenarios; and (3) an Adaptive Data Rate (ADR) mechanism that dynamically optimizes the combination between spreading factor SF and overlapping factor N in real-time.

A. Triangle-Chirp Modulation

The modulation architecture of TriLoRa is designed to transform a serial bit stream into a compact, superposed waveform, thereby breaking the strict sequential constraint of legacy LoRa. As illustrated in the workflow of Figure 6, the process comprises three key stages: Frequency Mapping, Triangle-Chirp Generation, and Delayed Superposition.

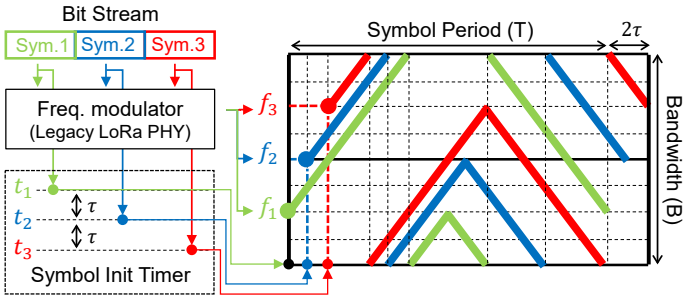


Fig. 6: The TriLoRa Modulation Process.

Stage 1: Frequency Mapping. As shown in the left part of Figure 6, the incoming bit Stream is first segmented into N individual groups, corresponding to symbols $Sym_1, Sym_2, \ldots, Sym_N$. To ensure compatibility with existing hardware logic, we utilize the standard *Frequency Modulator (Legacy LoRa PHY)* to map these bit groups into discrete initial frequency values. For instance, in the figure, three consecutive symbols are mapped to initial frequencies f_1 (green), f_2 (blue), and f_3 (red). These frequencies determine the starting point of the chirp trajectory in the frequency domain.

Stage 2: Triangle-Chirp Generation. Unlike standard LoRa which uses linear up-chirps, TriLoRa encodes these frequencies into a Triangle-Chirp waveform. As depicted in the main symbol area of Figure 6:

- Structure: Each symbol consists of two symmetric segments: an Up-Chirp followed by a Down-Chirp.
- Duration: The total duration of a single symbol remains T, consistent with the standard LoRa symbol period.
- Slope: To traverse the full bandwidth B and return within time T, the chirp rate (slope) is doubled compared to a standard linear chirp.

This symmetric structure creates the frequency-splitting effect as basis for resolving collisions at the receiver.

Stage 3: Delayed Superposition. The core innovation lies in the time-domain arrangement of these symbols. Instead of transmitting them sequentially, TriLoRa employs a *Symbol Init Timer* to trigger transmissions with a precise, staggered delay. Referring to the timing diagram in Figure 6:

1) At time $t_1 = 0$, the first symbol (Green, f_1) is generated.
2) At time $t_2 = \tau$, the timer triggers the second symbol (Blue, f_2). This symbol is superposed onto the first one, sharing the channel resources.
3) At time $t_3 = 2\tau$, the third symbol (Red, f_3) is added to the composite waveform.

As a result, multiple triangle chirps coexist within the same time window. By packing N symbols into a total duration of only $T + (N-1)\tau$ (where $\tau \ll T$), TriLoRa achieves parallel transmission, significantly increasing the effective data rate.

Time Offset (τ) Selection. The time offset τ is a critical system parameter that governs the trade-off between throughput and demodulation resolvability. To determine the lower bound of τ, we consider the fundamental property of Linear Frequency Modulation (LFM) [11, 26]. For a linear chirp

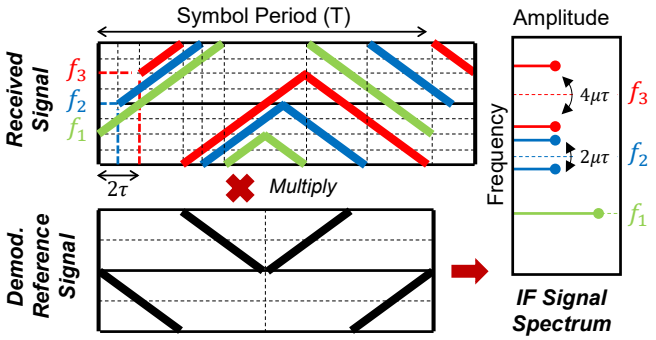


Fig. 7: Spectral representation in Frequency-Offset Matching. The aligned symbol (f_1, green) forms a single peak, while delayed symbols (f_2, blue; f_3, red) manifest as split pairs with predictable separations proportional to their time delays.

with slope μ, a time delay τ mathematically translates into a constant frequency shift $\Delta f = \mu\tau$ during the de-chirping process. In the specific case of our Triangle-Chirp, the waveform contains two segments with opposite slopes ($+\mu$ and $-\mu$). Consequently, a time offset τ causes the Up-segment to shift by $-\mu\tau$ and the Down-segment by $+\mu\tau$, resulting in a total frequency splitting separation of $\Delta f = 2\mu\tau$ in the demodulated spectrum. To ensure this split is resolvable by the receiver's FFT (which has a resolution of $1/T$), the condition $2\mu\tau > 1/T$ must be met. Substituting the triangle chirp slope $\mu = 2B/T$, we derive the theoretical lower bound:

$$\tau > \frac{1}{4B} \tag{3}$$

In practice, τ is set slightly above this bound to provide a safety margin against channel distortions.

B. Two-Tier Demodulation

The demodulation process at the receiver aims to disentangle the superposed waveform and recover the original symbol sequence $Sym_1, \ldots, Sym_N$. Unlike standard LoRa demodulation which searches for a single peak, TriLoRa faces a composite signal where symbol energies are mutually interfering. To address this, we propose two distinct demodulation strategies tailored for different computational and SNR constraints: *Frequency Matching* and *Cyclic Demodulation*.

1) Tier 1: Frequency Matching: This method treats demodulation as a pattern-matching problem within a single spectral view. It relies on the observation that was described in Section III-C: a Triangle-Chirp with a time offset Δt manifests as a pair of split peaks separated by $2\mu\Delta t$.

Figure 7 illustrates the spectral characteristics of the composite signal upon de-chirping with a standard, non-delayed reference aligned at $t = 0$. As observed in the green trace, the first symbol (Sym_1), being perfectly time-aligned with the reference ($\Delta t = 0$), focuses its energy into a single, high-amplitude peak at frequency f_1. In contrast, subsequent symbols transmitted with delays exhibit the frequency-splitting effect due to misalignment. Specifically, the second symbol (Sym_2, shown in blue) with delay τ bifurcates into two distinct peaks separated by $2\mu\tau$ (labeled as $2k\tau$), while the third symbol (Sym_3, shown in red) with delay 2τ splits further into peaks separated by $4\mu\tau$. This progression demonstrates that the frequency separation is directly proportional to the accumulated time delay, providing a distinguishable signature for each superimposed symbol.

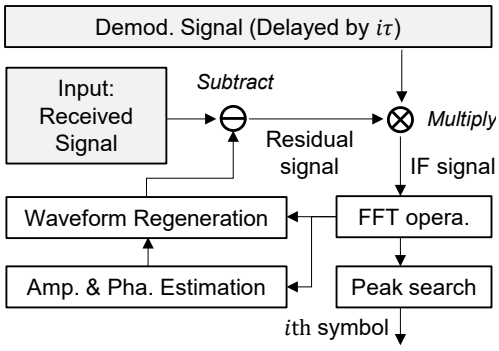


Fig. 8: Workflow of Cyclic Demodulation.

By executing a single FFT, the receiver can identify Sym_1 via the main peak and recover $Sym_2 \ldots Sym_N$ by searching for peak pairs with specific frequency separations (e.g., matching the red and blue dashed lines in Figure 7). This approach is computationally efficient ($O(L \log L)$) but requires high SNR to reliably detect the split, lower-energy peaks.

Mitigating Cross-Symbol Ambiguity. A critical challenge in this spectral analysis is *Cross-Symbol Ambiguity*. In a dense spectrum populated by $2N$ peaks, there is a non-negligible probability that two unrelated peaks – originating from different symbols or noise artifacts – coincidentally exhibit a frequency separation matching a valid target offset ($2\mu k\tau$). Such "phantom pairs" can lead to false decodings. To rigorously mitigate this, our algorithm enforces a *Top-Down Search Strategy*. We iterate from the largest expected time offset ($(N-1)\tau$) down to the smallest. This approach exploits the statistical property that wider spectral separations are unique and least likely to be mimicked by random peak combinations. Furthermore, the immediate *pruning* of identified peak pairs from the candidate set prevents them from contributing to subsequent false associations during the search for smaller offsets, ensuring a globally unique decoding solution.

2) Tier 2: Cyclic Demodulation: To enhance robustness in low-SNR regimes, we introduce a Cyclic Demodulation scheme that incorporates *Delayed Demodulation* and *Successive Interference Cancellation (SIC)*. Instead of analyzing the composite spectrum at once, this method isolates and removes symbols iteratively.

The workflow is detailed in Figure 8. The process operates on the *Residual Signal*, which is the raw received waveform. For the i-th symbol, the loop performs the following steps:

1) Delayed Alignment: The residual signal is multiplied by a reference chirp delayed by $(i-1)\tau$. This operation "focuses" on the energy of the i-th symbol (which was transmitted with delay $(i-1)\tau$) into a single, coherent peak, while leaving other symbols as spread interference.
2) Peak Search & Estimation: An FFT operation identifies the dominant peak to determine the symbol value

(Sym_i). Concurrently, the module estimates the amplitude and phase of the signal.
3) Waveform Regeneration: Using the decoded frequency and estimated channel parameters, the receiver reconstructs the ideal waveform of the i-th symbol.
4) Cancellation: The feedback loop (Fig. 8) subtracts the regenerated waveform from the input signal.

This cycle repeats for all N symbols. By stripping away strong interference layer by layer, this method significantly improves the Signal-to-Interference-plus-Noise Ratio (SINR) for subsequent symbols, enabling reliable decoding even under aggressive overlapping conditions.

Error Propagation Mitigation. A known concern with successive interference cancellation is that early decoding errors can cascade to subsequent iterations. TriLoRa mitigates this through two mechanisms. First, the energy focusing property of Triangle-Chirps provides a natural decoding order: the time-aligned symbol always produces the strongest, most concentrated peak, ensuring that the highest-confidence symbol is decoded and cancelled first. Second, the amplitude and phase estimation step (Step 2) provides a built-in quality metric – if the ratio between the dominant peak and the noise floor falls below a reliability threshold, the system flags the current configuration as unreliable and triggers a rate fallback (Section IV-D), preventing cascading errors from corrupting the entire symbol block.

Computational Complexity Analysis. We analyze the computational overhead of the proposed demodulation strategies, which represents a fundamental trade-off between computational efficiency and noise resilience. The complexity is primarily governed by the number of overlapping symbols (N), the FFT length (L), and the number of iterative cycles (C) for the cyclic approach. Frequency Matching (FM) offers the most lightweight solution, with a complexity of $\mathcal{O}(L \log L)$. Its cost is dominated by a single FFT operation performed on the composite signal, while the subsequent peak pairing involves negligible linear scans. This near-constant complexity, independent of N, makes it ideal for resource-constrained edge devices operating in high-SNR regimes. In contrast, *Cyclic Demodulation (CD)* prioritizes link reliability through an successive interference cancellation process, resulting in a significantly higher complexity of $\mathcal{O}(C \cdot (N^2 L + N \cdot L \log L))$. This heavy computational load arises from the need to perform N independent FFTs for specific time alignments in each cycle, coupled with signal reconstruction and subtraction steps that scale quadratically with N. Consequently, while the *CD* method delivers superior robustness in low-SNR conditions, its computational demand makes it best suited for powerful gateway devices where processing capability is abundant.

C. Hardware Error Compensation

In practical deployments, hardware oscillators inevitably introduce Carrier Frequency Offset (CFO), and propagation delays cause Symbol Timing Offset (STO). For TriLoRa, which relies on precise frequency splitting and temporal alignment, these errors can severely degrade performance if left uncorrected.

Impact Analysis. Let Δf denote the CFO and Δt denote the STO. CFO introduces a global spectral shift, displacing both peaks of a split pair by Δf. If $|\Delta f|$ exceeds half the frequency bin width, it causes symbol mapping errors. STO, manifesting as a time shift, alters the effective delay perceived by the receiver. For the up-segment, STO induces a frequency shift of $-\mu\Delta t$, while for the down-segment, it induces $+\mu\Delta t$. This asymmetry distorts the predictable frequency separation ($\Delta f_{split} = 2\mu\tau$), potentially breaking the pattern-matching logic. These methods rely on precise time alignment to focus symbol energy into a single tone. An uncompensated STO means the local reference chirp is misaligned by Δt, preventing energy concentration and causing spectral leakage, which severely hampers the successive interference cancellation performance in cyclic demodulation.

Joint Estimation and Compensation Design. To mitigate these effects, we implement a lightweight pre-processing stage that leverages the standard LoRa frame structure, which consists of a preamble (consecutive Up-Chirps) and a Start Frame Delimiter (SFD, typically containing Down-Chirps). By leveraging this, we first detect the peak frequency of the de-chirped Preamble ($\hat{f}_{up}$) and the SFD ($\hat{f}_{down}$). Due to the properties of linear chirps, these observed frequencies are linear combinations of the true symbol frequency ($f_{sym} \approx 0$ for preamble), CFO, and STO:

$$\begin{cases} \hat{f}_{up} = \Delta f - \mu \Delta t \\ \hat{f}_{down} = \Delta f + \mu \Delta t \end{cases} \tag{4}$$

By solving this system of linear equations, we isolate the errors:

$$\Delta f = \frac{\hat{f}_{up} + \hat{f}_{down}}{2}, \quad \Delta t = \frac{\hat{f}_{down} - \hat{f}_{up}}{2\mu} \tag{5}$$

Before passing the signal to the TriLoRa demodulator, we apply a digital correction. The CFO is removed by multiplying the received signal $r(t)$ by $e^{-j2\pi\Delta f t}$, and the STO is corrected by applying a fractional delay filter. This calibration ensures that the input to our demodulation algorithms is free of synchronization errors, preserving the integrity of the frequency-splitting features and enabling precise energy focusing. We note that this joint estimation is performed per packet using the standard preamble structure, and thus inherently adapts to device-specific oscillator characteristics. This makes the compensation robust across diverse low-cost and multi-vendor LoRa hardware, as no a priori device-specific calibration profile is required. After compensation, the residual frequency error is bounded within a fraction of the FFT bin width ($1/T$), which preserves the resolvability of split-peak patterns even under aggressive symbol superposition.

D. Adaptive Rate Selection

TriLoRa unlocks a new dimension for rate adaptation: the number of overlapping symbols N. Combined with the traditional spreading factor SF, the system data rate is determined

by the tuple (SF, N). While increasing N linearly boosts throughput, it simultaneously elevates the required SNR due to power division among symbols and inter-symbol interference. This represents an explicit sensitivity-for-throughput tradeoff: each doubling of N costs approximately 3 dB in required SNR. The role of rate adaptation is therefore to precisely manage this tradeoff, ensuring that N is increased only when the link budget provides sufficient margin. To navigate this trade-off, we propose a hybrid rate adaptation algorithm that combines model-based initialization with measurement-based dynamic adjustment.

SNR Threshold Modeling. The foundation of our adaptation logic is an analytical model that predicts the minimum SNR required for a specific configuration (SF, N). We define the threshold $SNR_{TH}(SF, N)$ as:

$$SNR_{TH}(SF, N) = SNR_{base}(SF) + 10\log_{10}(N) + \Delta_{demod} \tag{6}$$

where $SNR_{base}(SF)$ is the sensitivity threshold of standard LoRa (e.g., -20 dB for SF12). The term $10\log_{10}(N)$ represents the power division penalty: as the total transmission power is shared among N superimposed symbols, the energy per symbol decreases by a factor of N (assume uniformed power allocation). Δ_{demod} accounts for the implementation loss and residual interference specific to the chosen demodulation method. This model allows the system to instantly identify a set of feasible rates based on the current estimated channel SNR (SNR_{est}) upon initialization.

Exploit-and-Explore. Real-world channels are dynamic, and model-based prediction alone may suffer from inaccuracies due to fading or external interference. Therefore, we employ a continuous Exploit-and-Explore strategy to fine-tune the rate selection in real-time.

- Exploitation (Stability): The system maintains a Current Best Rate (R_{best}), defined as the (SF, N) pair that maximizes effective throughput. To ensure link stability, 90% of data packets are transmitted using R_{best}.
- Exploration (Optimization): The remaining 10% of packets are allocated to a *Probe Set* (S_{probe}), consisting of configurations adjacent to R_{best} (e.g., $N_{best} \pm k$ or $SF_{best} \pm 1$). These probe packets actively test the channel's capacity for higher throughput.

For every configuration, the system tracks the Packet Delivery Rate (PDR) via exponential weighted moving average. The decision metric is the *Effective Throughput*:

$$Eff.Throughput(R) = \frac{\text{Payload}}{\text{AirTime}(R)} \times PDR(R) \tag{7}$$

The system performs an opportunistic upgrades which updates the current best data rate to a probe rate R_{new} if the estimated effective throughput is $1.1\times$ higher, ensuring the gain outweighs the switching cost.

To prevent connection loss during sudden channel deterioration, a safety mechanism triggers if $PDR_{ewma}(R_{best})$ drops below 50%. The system immediately reverts to a conservative configuration (i.e., standard LoRa mode with $N = 1$) to restore connectivity before resuming exploration.



Fig. 9: Topographical visualization of the campus-scale LoRaWAN deployment, comprising over 3,000 end nodes and 200 gateways. The visualization highlights the high-density deployment characteristic of the testbed environment.

This hybrid approach ensures that TriLoRa aggressively leverages good channel conditions while retaining the ability to gracefully degrade to standard LoRa performance in harsh environments.

V. Evaluation

A. Implementation

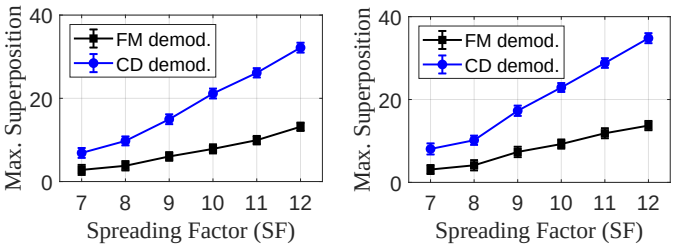
The TriLoRa system was prototyped using USRP X310 devices with UBX-160 daughterboards. In this setup, Triangle-Chirp waveforms are generated in MATLAB and transmitted over the air through the USRP front-ends. On the receiving side, captured raw I/Q samples are processed to execute the full algorithmic pipeline, comprising Joint Synchronization for CFO/STO compensation, the two-tier demodulation strategies, and the rate adaptation mechanism with closed-loop feedback.

Deployment Scope and Practical Feasibility. TriLoRa targets *single-link* rate enhancement: a single transmitter superimposes symbols from its own bit stream within each symbol window. It does not require synchronization between independent nodes and is thus complementary to, rather than a replacement for, multi-node collision resolution schemes [12, 19]. At the transmitter, generating a Triangle-Chirp requires only inverting the frequency slope at the symbol midpoint $(T/2)$ – a minimal logic change implementable on next-generation programmable baseband chips or emerging flexible IoT SoCs. At the gateway, the computationally intensive CD demodulation is designed for modern SDR-based concentrators, which possess sufficient DSP resources for parallel demodulation via firmware updates. The lightweight FM method, with $\mathcal{O}(L \log L)$ complexity comparable to standard LoRa demodulation, remains suitable for resource-constrained scenarios.

B. Experiment Settings

Our evaluation is twofold, comprising controlled link-level and large-scale network-level trace-driven simulations.

Link Performance Benchmarks. We assessed the physical layer performance through controlled link-level experiments.



(a) SER $< 5\%$ (b) SER $< 10\%$

Fig. 10: Performance comparison of symbol integration capability between demodulation method FM (Tier-1) and CD (Tier-2) across different Spreading Factors (SF) under SER thresholds of (a) 5% and (b) 10%, respectively.

The USRP TX and RX ports were linked via coaxial cables and programmable attenuators to create a deterministic channel. This setup facilitated precise control over signal attenuation, thereby mimicking a wide dynamic range of SNRs and equivalent propagation distances.

Network Performance Evaluation. To validate our system in a realistic, high-density environment, we utilized traces from the large-scale LoRaWAN testbed deployed across a university campus. As illustrated in Fig. 9, the network infrastructure consists of over 200 gateways and 3,000 LoRa end devices, covering a diverse range of indoor and outdoor scenarios.

The evaluation is based on a one-week operational dataset comprising over 9 million LoRaWAN packets sent from 3,227 distinct end-devices and processed by 44 active gateways. Each trace entry captures granular, timestamped physical layer metadata, including SNR and RSSI. By replaying this historical traffic profile, we benchmarked the aggregate network performance of TriLoRa against the standard LoRa baseline.

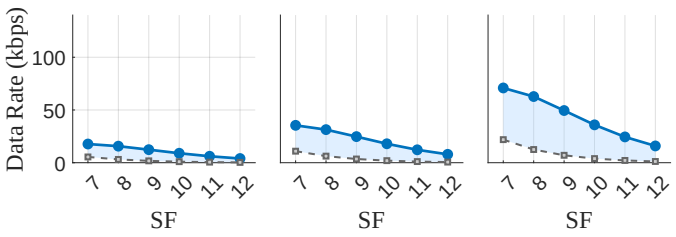
C. Link Performance

1) Symbol Superposition Capabilities: We first evaluate the *maximum number of successfully decoded overlapped symbols* capable of being superimposed within a single symbol window under varying SF settings. Fig. 10 presents the maximum number of overlapped symbols (N) supported while maintaining Symbol Error Rates (SER) below 5% and 10%. The experiment compares two demodulation strategies: *Frequency Matching (FM)* and *Cyclic Demodulation (CD)*.

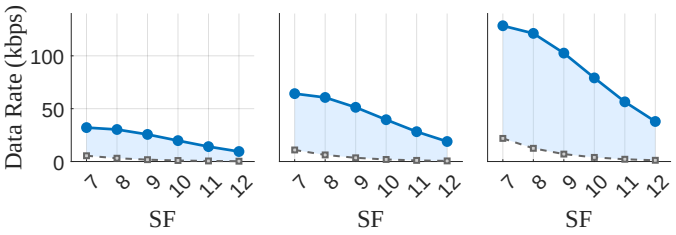
The results indicate that the integration capability increases approximately linearly with SF for both methods. This trend aligns with spread-spectrum theory, where higher spreading gains enhance robustness against interference and improve symbol densities without compromising decoding reliability.

However, a detailed comparison reveals a pronounced performance gap between the two strategies. *CD* consistently outperforms *FM* across all SF configurations. While both methods exhibit linear scaling, *CD* demonstrates a significantly steeper growth rate – offering approximately $2\times$ to $2.5\times$ higher symbol integration capacity than *FM*.

For example, under the stricter reliability constraint (SER $<$ 5%) shown in Fig. 10(a), *FM* supports approximately four symbols at SF7, whereas *CD* supports about seven. This



(a) TriLoRa (with FM demodulation) vs. Standard LoRa



(b) TriLoRa (with CD demodulation) vs. Standard LoRa

Fig. 11: Data rate comparison under varying bandwidths.

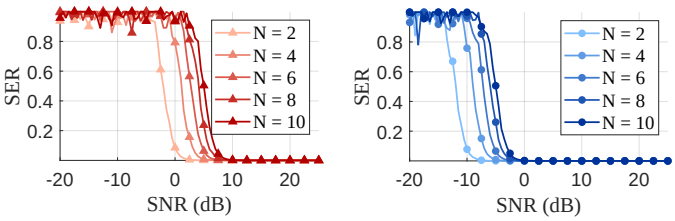
disparity widens at higher spreading factors: at SF12, *FM* reaches a capacity of $\sim$14 symbols, while *CD* dramatically increases to $\sim$33 symbols. These results underscore *CD*'s superior ability to leverage the SNR gains from high SFs, effectively translating them into tangible throughput improvements and higher processing efficiency.

2) PHY Data Rate: We analyzed the data rates of the proposed demodulation methods – *Frequency Matching (FM)* and *Cyclic Demodulation (CD)* – against the Standard LoRa baseline across bandwidths of 125 kHz, 250 kHz, and 500 kHz. The experimental results reveal distinct roles for bandwidth and demodulation mechanisms: while increasing channel bandwidth provides a linear rate gain, the choice of demodulation strategy proves to be the decisive factor in achieving a qualitative leap in communication efficiency.

Specifically, Standard LoRa exhibits a severe rate-coverage trade-off: as the Spreading Factor (SF) increases to extend the communication range, data rates decay exponentially. In contrast, the proposed *CD* method breaks this bottleneck through interference cancellation mechanism within a single symbol window. As illustrated in Fig. 11(b), under identical bandwidth configurations, TriLoRa with *CD* demodulation not only maintains throughput several times higher than Standard LoRa at low SFs but, crucially, achieves a transmission rate up to $32\times$ that of Standard LoRa under extreme high-SF conditions (SF12). Remarkably, the data rate of *CD* at SF12 exceeds even the peak rate of Standard LoRa at SF7.

While TriLoRa with *FM* demodulation (Fig. 11(a)) exhibits more modest gains compared to *CD*, it consistently outperforms the Standard LoRa baseline across the full frequency band, further validating the robustness of the underlying Triangle-Chirp modulation scheme.

3) Noise Tolerance: To further investigate the performance boundaries of TriLoRa under varying SNR conditions, we conducted experimental tests on the Symbol Error Rate (SER) while merging a varying number of symbols ($N \in$



(a) TriLoRa with FM demod. (b) TriLoRa with CD demod.

Fig. 12: Measured SER vs. SNR performance under varying numbers of merged symbols with $SF = 10$.

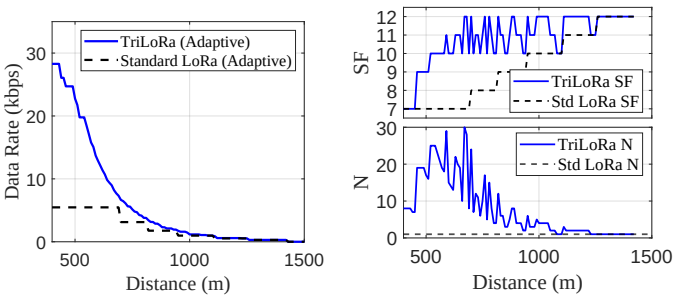
$\{2, 4, 6, 8, 10\}$) at a fixed Spreading Factor (SF10). Fig. 12 presents the measured performance curves for the two proposed demodulation methods.

The experimental data clearly illustrates a general trade-off between symbol density (N) and noise tolerance. This trend stems from two primary factors: 1) energy dispersion, where the effective energy allocated to each symbol decreases as N increases; and 2) the elevation of the noise floor due to signal superposition and inter-symbol interference.

The two methods exhibit distinct capabilities in managing these impairments. The *FM* method (Fig. 12(a)) is more responsive to the cumulative effects of signal stacking. As N increases to 10, its effective operating range naturally shifts, with the demodulation threshold moving into the positive SNR regime (approximately +5 dB). In contrast, the *CD* method (Fig. 12(b)) demonstrates superior robustness. By effectively mitigating interference, it maintains a demodulation threshold well within the negative SNR regime (approximately -5 dB) even when merging 10 symbols, demonstrating its capability for deep symbol integration under severe channel conditions.

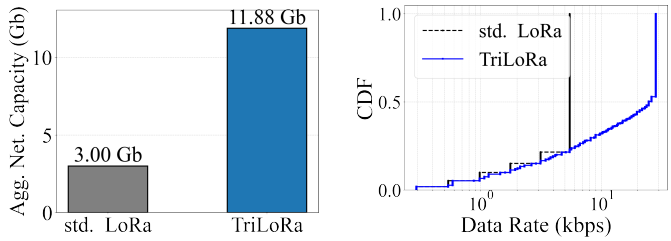
4) Communication range: To demonstrate the system's adaptability in dynamic channel conditions, we evaluated the data rate and parameter configuration (SF and N) over communication distances ranging from 400 m to 1500 m. Fig. 13 illustrates the results of this fine-grained rate adaptation mechanism. Unlike standard LoRa's ADR mechanism, which is constrained to coarse-grained, discrete switching between SF, TriLoRa monitors the channel state (SNR) in real-time and dynamically converts the redundancy in the link budget into a proper combination between SF and N. This approach fills the performance gaps between SF levels, thereby achieving continuous and smooth rate optimization.

Fig. 13a shows that TriLoRa (solid blue line) achieves a smooth varying data rate curve. In the short-range region ($\leq$ 500 m) where channel conditions are favorable, the system selects SF7 and the maximized N ($N = 8$), yielding a throughput of approximately 28kbps, which far exceeds the theoretical limit of standard LoRa (dashed line). As the distance increases, standard LoRa exhibits a step-like rate degradation due to discrete SF switching, dropping to $\sim$0.3 kbps at 1500 m. In contrast, TriLoRa exploits every decibel of SNR margin by precisely adjusting N, maintaining a data rate of approximately 9.5kbps even at the 1500 m edge – a 30$\times$ improvement over standard LoRa. This smooth rate curve demonstrates TriLoRa's ability to maximize throughput while



(a) Data Rate vs. Distance (b) SF & N Adaptation

Fig. 13: Experimental evaluation of rate adaptation over distance.



(a) Aggregate Network Capacity (b) CDF of Packet Data Rates

Fig. 14: Network-wide performance evaluation in aggregated network capacity.

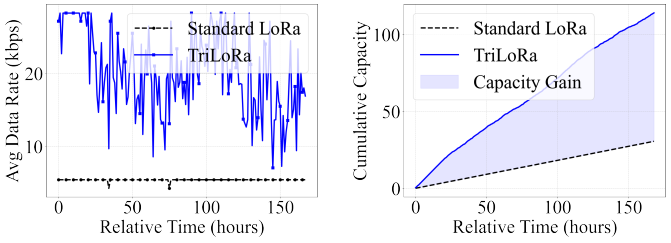
maintaining connection stability in dynamic mobile channels.

Fig. 13b further reveals the internal scheduling mechanism driving this performance. The evolution of N presents a distinct constrained sawtooth pattern. Within the coverage range of a specific SF (e.g., $SF = 9$), N gradually decreases from its capacity peak as the distance increases and SNR fades. Crucially, at the handover points between SF (e.g., from SF9 to SF10), the system leverages the sensitivity gain ($\sim$3dB) to restore N to a higher value. This rate adaptation mechanism effectively fills the performance gaps between SFs, ensuring that the system always operates near its physical capacity limit.

D. Network-Scale Performance Evaluation

1) Aggregate Network Capacity: We define the *Aggregate Network Capacity* (C_{net}) as the summation of the theoretical physical layer data rates for all packets in the trace: $C_{net} = \sum_{i=1}^{M} R_i(SF, N, SNR)$. Since the trace duration is fixed, this metric effectively normalizes the throughput against channel occupancy time (Airtime), serving as a holistic indicator of spectral efficiency.

As illustrated in Fig. 14a, TriLoRa achieves a remarkable aggregate capacity of 11.88 Gb over the one-week period, representing a 3.96$\times$ net gain compared to the 3 Gb baseline of standard LoRa. This significant boost stems from TriLoRa's ability to convert high-SNR margins into parallel symbol transmissions, thereby drastically reducing the airtime required per bit. Fig. 14b further details the CDF of per-packet data rates. Standard LoRa exhibits a discrete, step-like distribution limited by the fixed data rates from SF7 to SF12, with nearly 40% of packets constrained below 3 kbps. In contrast, TriLoRa shifts the entire distribution significantly to the right. The



(a) 7-Day Rate Trend (b) Cumulative Capacity Growth

Fig. 15: Data rate micro-benchmark of the most active node.

smooth curve demonstrates its fine-grained rate adaptation capability, which fills the performance gaps between SFs. Notably, TriLoRa enables over 60% of packets to operate at $\geq$10 kbps, a range almost unreachable by standard LoRa under the same channel conditions.

2) Long-term Temporal Stability: To evaluate robustness under dynamic channel fluctuations, we performed a micro-benchmark on the most active node in the network. Fig. 15a plots the hourly average data rate over the one-week duration. TriLoRa (solid blue line) consistently forms a performance envelope above standard LoRa (dashed black line). Even during periods of channel degradation (e.g., nighttime fading), TriLoRa adapts N to maintain connectivity while maximizing residual throughput. Conversely, standard LoRa suffers from the quantization inefficiencies of SF switching, failing to capitalize on transient SNR peaks.

This instantaneous advantage translates into a compounding long-term benefit, as shown in the cumulative capacity growth in Fig. 14b. The widening gap between the two curves illustrates a snowball effect, where TriLoRa's continuous rate optimization accumulates into a massive total data transfer advantage over time. This confirms that TriLoRa not only provides peak rate bursts but also delivers sustainable, long-term capacity gains suitable for data-intensive IoT applications such as firmware updates and high-frequency sensing.

VI. Conclusion

In this paper, we presented TriLoRa, a PHY-layer innovation enabling high-throughput LoRaWAN via Triangle-Chirp Modulation (TCM). By reliably superposing multiple symbols and employing a hierarchical demodulation architecture – integrating Frequency Matching and Cyclic Demodulation – TriLoRa maintains robust connectivity under severe noise. Trace-driven evaluations on a large-scale commercial deployment demonstrate that TriLoRa's fine-grained rate adaptation effectively exploits underutilized channel margins. Consequently, it achieves a 3.96$\times$ increase in aggregate network capacity and a 30$\times$ data rate improvement for edge nodes. Ultimately, TriLoRa bridges the gap between future IoT connectivity demands and current LPWAN bandwidth limits, paving the way for data-intensive applications. Future research will focus on hardware-level optimization and upper-layer protocol integration to fully unleash TriLoRa's potential in large-scale networks.

References

[1] G. Premsankar, B. Ghaddar, M. Slabicki, and M. D. Francesco, "Optimal configuration of lora networks in smart cities," *IEEE Transactions on Industrial Informatics*, vol. 16, no. 12, pp. 7243–7254, 2020.

[2] N. Varsier and J. Schwoerer, "Capacity limits of lorawan technology for smart metering applications," in *2017 IEEE International Conference on Communications (ICC)*, 2017, pp. 1–6.

[3] Z. Li, T.-B. Nguyen, Q. Lampin, I. Sivignon, and S. Zozor, "Ensuring k-coverage in low-power wide area networks for internet of things," in *2017 International Conference on Computing, Networking and Communications (ICNC)*, 2017, pp. 26–30.

[4] J. Grübel, T. Thrash, L. Aguilar, M. Gath-Morad, D. Hélal, R. W. Sumner, C. Hölscher, and V. R. Schinazi, "Dense indoor sensor networks: Towards passively sensing human presence with lorawan," *Pervasive and Mobile Computing*, vol. 84, p. 101640, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1574119222000700

[5] S. Tong, J. Wang, J. Yang, Y. Liu, and J. Zhang, "Citywide lora network deployment and operation: Measurements, analysis, and implications," in *Proceedings of the 21st ACM Conference on Embedded Networked Sensor Systems*, ser. SenSys '23. New York, NY, USA: Association for Computing Machinery, 2024, p. 362–375. [Online]. Available: https://doi.org/10.1145/3625687.3625796

[6] J. Lorincz, K. Levarda, M. Čagalj, and A. Kukuruzović, "A comprehensive analysis of lora network wireless signal quality in indoor propagation environments," *Journal of Sensor and Actuator Networks*, vol. 14, no. 6, 2025. [Online]. Available: https://www.mdpi.com/2224-2708/14/6/111

[7] R. Fujdiak, K. Mikhaylov, J. Pospisil, A. Povalac, and J. Misurec, "Insights into the issue of deploying a private lorawan," *Sensors*, vol. 22, no. 5, 2022. [Online]. Available: https://www.mdpi.com/1424-8220/22/5/2042

[8] D. Orlovs, A. Rusins, V. Skrastiņš, and J. Judvaitis, "Lpwan technologies for iot: Real-world deployment performance and practical comparison," *Preprints*, November 2025. [Online]. Available: https://doi.org/10.20944/preprints202511.0316.v1

[9] Semtech, "Sx1280," https://www.semtech.com/products/wireless-rf/lora-connect/sx1280, 2024.

[10] X. Xia, Q. Chen, N. Hou, and Y. Zheng, "Hylink: Towards high throughput lpwans with lora compatible communication," in *Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems*, 2022, pp. 578–591.

[11] A. Maleki, H. H. Nguyen, E. Bedeer, and R. Barton, "A tutorial on chirp spread spectrum modulation for lorawan: Basics and key advances," *IEEE Open Journal of the Communications Society*, 2024.

[12] X. Xia, Y. Zheng, and T. Gu, "Ftrack: Parallel decoding for lora transmissions," in *Proceedings of the 17th Conference on Embedded Networked Sensor Systems*, 2019, pp. 192–204.

[13] X. Xia, N. Hou, Y. Zheng, and T. Gu, "Pcube: Scaling lora concurrent transmissions with reception diversities," *ACM Transactions on Sensor Networks*, vol. 18, no. 4, pp. 1–25, 2023.

[14] F. Yu, X. Zheng, D. Xia, L. Liu, and H. Ma, "Enabling reliable lora decoding under cross-channel interference," in *IEEE INFOCOM 2025 - IEEE Conference on Computer Communications*, 2025, pp. 1–10.

[15] X. Xia, Q. Chen, N. Hou, Y. Zheng, and M. Li, "Xcopy: Boosting weak links for reliable lora communication," in *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*, 2023, pp. 1–15.

[16] C. Li, Z. Cao, and L. Xiao, "Curvealoha: Non-linear chirps enabled high throughput random channel access for lora," in *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*. IEEE, 2022, pp. 520–529.

[17] R. Eletreby, D. Zhang, S. Kumar, and O. Yağan, "Empowering low-power wide area networks in urban settings," in *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, 2017, pp. 309–321.

[18] S. Tong, J. Wang, and Y. Liu, "Combating packet collisions using non-stationary signal scaling in lpwans," in *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*, 2020, pp. 234–246.

[19] S. Tong, Z. Xu, and J. Wang, "Colora: Enabling multi-packet reception in lora," in *IEEE INFOCOM 2020-IEEE Conference on Computer Communications*. IEEE, 2020, pp. 2303–2311.

[20] M. O. Shahid, M. Philipose, K. Chintalapudi, S. Banerjee, and B. Krishnaswamy, "Concurrent interference cancellation: Decoding multi-packet collisions in lora," in *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, 2021, pp. 503–515.

[21] X. Wang, L. Kong, L. He, and G. Chen, "mlora: A multi-packet reception protocol in lora networks," in *2019 IEEE 27th International Conference on Network Protocols (ICNP)*. IEEE, 2019, pp. 1–11.

[22] Y. Wang, F. Zhang, X. Zheng, L. Liu, and H. Ma, "Decoding lora collisions via parallel alignment," *ACM Transactions on Sensor Networks*, vol. 19, no. 3, pp. 1–25, 2023.

[23] M. Hanif and H. H. Nguyen, "Slope-shift keying lora-based modulation," *IEEE internet of things journal*, vol. 8, no. 1, pp. 211–221, 2020.

[24] A. Mondal, M. Hanif, and H. H. Nguyen, "Ssk-ics lora: A lora-based modulation scheme with constant envelope and enhanced data rate," *IEEE Communications Letters*, vol. 26, no. 5, pp. 1185–1189, 2022.

[25] T. Elshabrawy and J. Robert, "Interleaved chirp spreading lora-based modulation," *IEEE Internet of Things Journal*, vol. 6, no. 2, pp. 3855–3863, 2019.

[26] B. Reynders and S. Pollin, "Chirp spread spectrum as a modulation technique for long range communication," in *2016 Symposium on Communications and Vehicular Technologies (SCVT)*. IEEE, 2016, pp. 1–5.