UniChirp: Unwrapping In-Chirp Phase Misalignment for Weak LoRa Signal Demodulation

***Abstract*—LoRa is widely adopted in low-power wide-area networks (LPWANs) for long-range and energy-efficient IoT communications, yet its demodulation performance degrades severely in challenging deployment scenarios with extremely low signal-to-noise ratios (SNRs), such as obstructed or industrial environments. A key limiting factor is in-chirp phase misalignment, which prevents coherent energy aggregation in conventional FFT-based LoRa demodulation and leads to weak spectral peaks. In this paper, we present *UniChirp*, a practical LoRa demodulation framework that improves robustness under weak-signal conditions by explicitly addressing the in-chirp phase misalignment. *UniChirp* derives a lightweight analytical model that captures chirp phase deviations induced by hardware offsets and modulation artifacts, enabling deterministic phase compensation without exhaustive search. Building on this model, the framework incorporates a preamble-guided phase correction mechanism and a structure-aware spectral filtering strategy to enhance symbol-level energy aggregation. The proposed method is implemented entirely in software and requires no modifications to existing LoRa end devices or protocol stacks. We deployed the system in real industrial production environments, addressing weak signal issues caused by environmental occlusions. Extensive evaluations using commercial LoRa hardware across controlled experiments demonstrate that *UniChirp* enhances symbol-level decoding robustness, achieving a 1.3 dB improvement in SNR sensitivity over current state-of-the-art approaches.**

***Index Terms*—LoRa, low-SNR demodulation, frequency-domain compensation, phase modeling, IoT**

I. Introduction

LoRa, as a representative Low Power Wide Area Networks (LPWANs) technology, has been widely adopted in various Internet of Things (IoT) applications [1]. It supports long-range communication with battery-powered devices and requires low deployment and maintenance cost, making it well suited for large-scale industrial monitoring systems [1], where sensors are distributed over wide areas and installed at hard-to-access locations. In practical industrial environments, however, reliable LoRa communication remains challenging. Sensors are often deployed inside pipelines or reaction vessels, introducing significant attenuation, multipath fading, and interference, substantially resulting in extremely low signal-to-noise ratios (SNRs) [2]–[5]. When the SNR falls below the threshold for the lowest data rate setting (SF12, CR=4/8), LoRa fails to maintain connectivity, revealing a gap between its design aspirations and its actual performance in complex environments.

To investigate the root cause of LoRa link failures in low-SNR scenarios, we revisit its physical-layer modulation mechanism. LoRa employs frequency-swept chirp signals, where each symbol is encoded by varying the initial frequency of an up-chirp within a fixed bandwidth. Due to the cyclic nature of the spectrum, frequencies exceeding the upper bound of the band are wrapped around to the lower edge, resulting in a discontinuity that splits each symbol into two chirp segments, as illustrated in Fig. 1. Ideally, LoRa achieves robust demodulation by coherently aggregating signal energy across the entire chirp duration. However, we observe a critical impairment: the two chirp segments within a physical-layer symbol have a phase discontinuity. This discontinuity disrupts the phase coherence required for energy accumulation, significantly limiting demodulation performance in low SNR conditions.

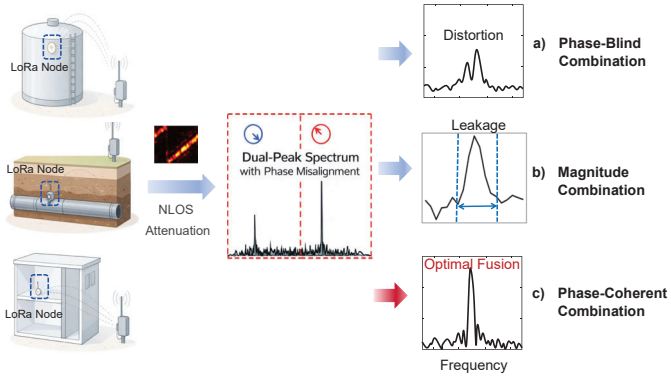

Fig. 1: Illustration of dual-peak combination mechanisms in LoRa demodulation with and without phase coherence

Existing efforts to recover symbol energy under phase discontinuity primarily follow two strategies. The first approach leverages spectrum aliasing by setting the physical sampling frequency equal to the chirp bandwidth [6]. This causes the two chirp segments within a symbol to be aliased onto the same frequency bin after dechirping and FFT. However, in the presence of phase discontinuity, such aliasing results in destructive interference due to misaligned phases, as illustrated in Fig. 1, exacerbating energy leakage and degrading demodulation performance. The second strategy attempts to compensate for the phase offset between the two chirp segments prior to demodulation [6], [7]. They estimate the phase discontinuity and adjust the signal phase to enable coherent energy fusion. These methods, however, treat the phase offset as an unknown variable and perform exhaustive search across a wide range of candidate values. This brute-force search introduces substantial computational overhead and limits their applicability in real-time systems.

To address these limitations, we introduce a novel phase discontinuity compensation mechanism that enables reliable

energy aggregation for robust LoRa demodulation under extremely low SNR. Our key observation is that the phase jump within a LoRa symbol is not random, but predictable based on the chirp modulation structure. Leveraging this insight, we construct a phase prediction model that infers the compensation phase directly, eliminating the need for brute-force search. Built on this foundation, we propose *UniChirp*, a lightweight demodulation framework that mitigates frequency-domain phase discontinuity and enables coherent energy fusion, significantly improving the identifiability of spectral energy peaks and enhancing LoRa decoding performance in low-SNR scenarios.

Putting the idea of UniChirp into practice, however, remains challenging due to the following key issues: *(1) How to accurately model intra-symbol phase jumps?* Phase discontinuities arise from symbol misalignment and frequency offsets, and require an analytical model that can explicitly capture their deterministic relationship for precise compensation. *(2) How to reliably aggregate energy under low-SNR conditions?* Severe noise leads to unpredictable phase deviations across symbols, making it difficult to extrapolate and align dual peaks for coherent energy fusion. *(3) How to suppress out-of-band noise before demodulation?* LoRa's reception accumulates noise beyond the chirp's encoding band, which distorts the spectral profile and weakens energy concentration after FFT.

To tackle the core challenges of *UniChirp*, we design a unified demodulation framework that enables precise phase modeling, robust energy aggregation, and effective noise suppression for weak LoRa signals. First, to address the challenge of intra-symbol phase jumps, we develop a closed-form analytical model that accurately captures phase discontinuities, caused by symbol misalignment and carrier frequency offset. This model explicitly links phase evolution to timing and frequency parameters, allowing efficient compensation of phase jumps without brute-force search, and facilitating consistent dual-peak alignment across symbols. Second, to enable complete energy aggregation within each chirp, we propose a preamble-guided extrapolation algorithm that constructs a phase compensation trajectory by fitting the preamble's spectral peaks. To ensure robustness, the algorithm detects noise-induced outliers using slope continuity constraints and local phase deviation metrics, thereby enabling reliable compensation across payload symbols. Third, to suppress out-of-band noise before FFT-based demodulation, we design a structure-aware frequency-domain filter that adapts to the encoding pattern of LoRa chirps. By selectively retaining energy within the expected spectral trajectory and attenuating non-conforming components, this filter mitigates noise folding and improves spectral concentration in the encoded bins. Together, these components form the foundation of *UniChirp*, enabling accurate demodulation of weak LoRa signals that would otherwise fall below the decoding threshold.

We implement *UniChirp* on software-defined radios and evaluate its performance through extensive experiments with weak LoRa signals. The results demonstrate that *UniChirp* significantly enhances the demodulation success rate under extremely low-SNR conditions, pushing the decoding threshold beyond conventional capabilities.

In summary, our contributions are as follows:

- We present *UniChirp*, the first demodulation framework that explicitly models intra-symbol phase jumps to enable reliable LoRa decoding at ultra-low SNR. To the best of our knowledge, *UniChirp* is the first to analytically characterize intra-symbol phase discontinuities as a function of timing and frequency offsets, providing a closed-form compensation method without exhaustive search.
- We present a demodulation framework for weak-signal LoRa reception that addresses three key challenges: First, we develop a closed-form model to characterize intra-symbol phase jumps caused by timing and frequency offsets, enabling accurate phase compensation. Second, we propose a preamble-guided phase extrapolation scheme that achieves robust phase-coherent energy aggregation by filtering noise-induced deviations across symbols. Third, we design a structure-aware spectral filter to suppress out-of-band noise before demodulation, thereby improving weak-signal decodability.
- We implement *UniChirp* on commodity LoRa hardware, demonstrating reliable packet demodulation at SNRs as low as -17.3 dB, 1.3 dB lower than state-of-the-art methods, significantly exceeding the decoding limits of existing LoRa systems. Building on this capability, we further deploy *UniChirp* in real industrial systems, embedding sensor nodes inside pipelines and reaction vessels, where severe signal attenuation previously prevented connectivity, and show that it effectively enables communication in these challenging environments.

The remainder of this paper is organized as follows. Section II presents the background and fundamentals of the LoRa physical layer. Section III introduces our analytical model for intra-symbol phase jumps. Section IV describes the overall system architecture. Section V details the key components and algorithms of the system design. Section VI presents the implementation details. Section VII covers the evaluation methodology and results. Section VIII reviews related works. Finally, Section IX concludes the paper.

II. Background

A. LoRa Physic Primer

LoRa employs Chirp Spread Spectrum (CSS) modulation at the physical layer, where each symbol is represented by a chirp whose frequency increases linearly over time across the entire bandwidth BW. The base up-chirp is defined as:

$$C(t) = e^{j2\pi\left(-\frac{BW}{2}+\frac{BW}{2T}t\right)t}, \tag{1}$$

where T is the chirp duration. To encode a symbol $x \in \{0, 1, ..., 2^{SF}-1\}$, LoRa applies a cyclic frequency shift to the base chirp by:

$$f_s = \frac{x \cdot BW}{2^{SF}}, \tag{2}$$

where SF denotes the spreading factor. The resulting modulated chirp is:

$$y_e(t) = C(t) \cdot e^{j2\pi f_s t}. \tag{3}$$

As frequency shifting pushes part of the chirp beyond the upper edge of the bandwidth, LoRa wraps exceeding frequencies back to the lower bound $-BW/2$, forming a two-segment waveform. This circular frequency shift introduces a phase discontinuity between the two segments within a symbol. At the receiver, demodulation is performed in two steps. First, the received signal is dechirped by multiplying it with a base down-chirp, i.e., the complex conjugate of the base up-chirp. This operation transforms the chirp into a single-tone signal whose frequency corresponds to the symbol's frequency shift. Then, a Fast Fourier Transform (FFT) is applied to the dechirped signal, yielding a peak at the frequency bin associated with the transmitted symbol.

B. Motivation

Conventional LoRa demodulation adopts a dechirp-then-FFT approach, which treats each symbol as a purely frequency-domain entity and relies solely on energy detection. However, this method fundamentally ignores the intra-symbol phase discontinuities caused by LoRa's step-wise frequency modulation. As a result, signal energy fails to align coherently in the FFT domain, leading to spectral leakage and energy dispersion, particularly detrimental under low SNR conditions. This inherent limitation motivates our work: we aim to explicitly model the phase transitions within each chirp symbol, enabling accurate intra-symbol phase alignment and coherent energy aggregation. By recovering phase continuity across the entire chirp duration, we significantly enhance demodulation robustness in weak-signal scenarios.

III. Phase Model

This section models the phase characteristics of LoRa chirps, revealing how hardware and synchronization errors distort ideal phase trajectories. The phase model enables precise compensation for improved decoding in noisy environments.

In an ideal LoRa system, the modulated signal generated by a transmitter should exhibit a continuous phase trajectory over time. In practice, however, this phase continuity is disrupted by imperfections in hardware and synchronization between the transmitter and receiver. These factors introduce systematic phase offsets in the received signal, manifesting as intra-symbol phase discontinuities. Following, we identify three dominant sources of chirp phase distortions:

(1) Carrier Frequency Offset (CFO). CFO originates from oscillator mismatches between the transmitter and receiver, introducing a constant frequency offset Δf. This leads to a linear phase shift over time:

$$\varphi_{\text{CFO}}(t) = 2\pi\Delta f \cdot t. \tag{4}$$

The resulting phase offset accumulates gradually across the chirp duration but remains strictly continuous. That is, CFO causes a uniform rotation of the signal's phase without breaking the intra-symbol phase trajectory. Therefore, it does not disrupt the coherent structure of the chirp, nor degrade the energy concentration in the FFT domain. At the receiver, compensating for CFO simply involves correcting the frequency shift in the dechirped signal.

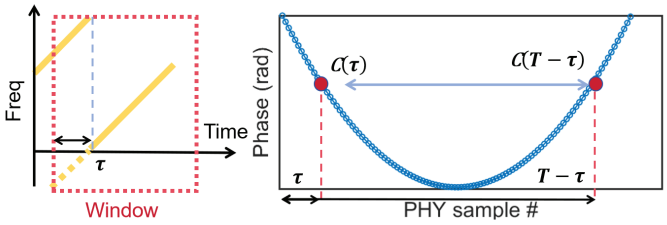

Fig. 2: Intra-symbol phase discontinuity caused by misalignment

(2) Time Offset (TO). TO refers to the misalignment between the receiver's FFT window and the actual start of the chirp symbol. It introduces a frequency-dependent phase rotation:

$$\varphi_{\text{TO}}(f) = -2\pi f\tau, \tag{5}$$

where τ is the timing offset between the received signal and the FFT window. Due to the cyclic structure of LoRa chirps, each symbol is composed of two concatenated segments at different frequencies. These segments experience distinct phase rotations under the same τ, resulting in a phase discontinuity at the boundary between them.

Following, we analyze how the time offset τ quantitatively leads to intra-symbol phase discontinuity. As illustrated in Fig. 2, each modulated LoRa symbol within consists of two concatenated chirp segments. The first segment starts from the encoded frequency, and the second wraps around to $-BW/2$ after reaching the band edge. Let $C(t)$ denote the instantaneous phase of a base chirp that starts at time $t = 0$. Then, at $t = 0$, the observed phase for the first segment corresponds to a base chirp that started at $t = -T + \tau$, while the second segment corresponds to a chirp that started at $t = -\tau$. The resulting intra-symbol phase discontinuity is given by:

$$\varphi_{\text{diff}} = C(T - \tau) - C(\tau) \tag{6}$$

Since $C(t)$ is quadratic with respect to time (due to linear frequency modulation), we can simplify the expression to:

$$\varphi_{\text{diff}} = -2\pi BW \cdot \tau + M, \tag{7}$$

where M is a constant determined by system parameters including bandwidth BW, symbol duration T, and the carrier frequency. This result provides an intuitive interpretation: the phase discontinuity between the two chirp segments increases linearly with the intra-symbol time offset τ. This predictable relationship enables accurate compensation in signal processing, thereby restoring intra-symbol phase continuity crucial for downstream demodulation.

(3) Sampling Frequency Offset (SFO). SFO occurs when the receiver's sampling rate differs from the transmitter's nominal rate F_s, introducing a gradual drift in sampling time. This causes the chirp-window timing offset τ_n growing across chirps:

$$\tau_n = n \cdot \Delta t + \tau_0, \tag{8}$$

where n is the chirp index, Δt is the timing drift per chirp caused by SFO, and τ_0 is the initial time offset. Substituting this time-varying chirp-window offset (i.e., τ_n) into Eq. 7, we can obtain the time-varying phase jump within each symbol:

$$\varphi_n = -2\pi \cdot \mathrm{BW} \cdot \tau_n + M = -2\pi \cdot \mathrm{BW} \cdot (n \cdot \Delta t + \tau_0) + M \quad (9)$$

Since τ_0 maps to a fixed chirp-window time offset, its phase contribution can be absorbed into a constant, giving a linear phase evolution:

$$\varphi_n = -2\pi \cdot \mathrm{BW} \cdot n \cdot \Delta t + \varphi_0 \quad (10)$$

This shows that SFO introduces a linear phase trend across chirps, where the larger the SFO is, the faster the phase jump increases over time.

We now integrate the impacts of CFO, TO, and SFO into a unified phase misalignment model:

$$\varphi_{mis}[n] = -2\pi \cdot \mathrm{BW} \cdot n \cdot \frac{SFO}{F_s} - 2\pi \cdot \mathrm{BW} \cdot TO + \varphi_0 \quad (11)$$

This model serves as the analytical foundation for subsequent phase misalignment compensation.

In summary, CFO does not affect in-chirp phase alignment; TO introduces a constant in-chirp phase offset; and SFO leads to a linearly increasing phase misalignment within each chirp. This phase model provides a quantitative foundation for compensating in-chirp phase distortion in low-SNR LoRa demodulation.

IV. System Overview

We propose *UniChirp*, a LoRa modulation framework that achieves optimal chirp energy concentration by modeling and compensating in-chirp phase misalignment. As shown in Fig. 3, the decoding pipeline consists of three main stages: **(1) Phase model construction**: Our theoretical analysis shows that the in-chirp phase misalignment is determined by the timing offset (TO) and sampling frequency offset (SFO) of the transmitter. Based on this insight, *UniChirp* builds a precise phase model for each received packet. It first estimates TO and SFO from the preamble symbols, and then refines the model through a fitting process that aligns the predicted phase evolution with the actual preamble phases. **(2) Peak alignment and fusion**: While the phase model reveals the phase offset between two chirp segments within the same symbol, these segments naturally overlap in frequency and cannot be separated in conventional demodulation. To solve this, *UniChirp* introduces oversampling to create two distinct energy peaks in the frequency domain. By applying model-guided phase compensation to one peak and then summing the two, *UniChirp* achieves coherent energy fusion and maximizes peak intensity. **(3) Interference suppression**: Oversampling inevitably introduces wideband noise beyond the LoRa signal bandwidth. To mitigate this, *UniChirp* applies a spectrum-aligned bandpass filter that removes out-of-band noise while preserving in-band signal phase, ensuring clean and reliable demodulation.

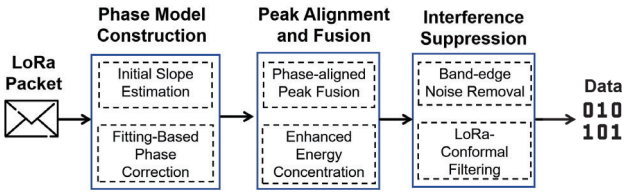

Fig. 3: Workflow of *UniChirp* for weak signal demodulation.

Together, these modules form a systematic demodulation pipeline that reconstructs chirp phase misalignment, separates and compensates spectral components for ideal energy concentration, and suppresses external interference. We detail the implementation of each module in Section V.

V. Design

A. Phase Model Construction

According to the analysis in Section III, we establish a direct relationship between the phase discontinuity and the sampling frequency offset (SFO). Based on this, a phase drift model is constructed from the preamble symbols, enabling accurate prediction and compensation of the phase evolution across received symbols. This approach effectively addresses distortions caused by frequency offset and sampling mismatch, improving demodulation performance even under challenging conditions such as phase wrapping or abrupt jumps.

The preamble provides a sequence of identical symbols with stable structure, making it suitable for observing phase variation over time. By processing the preamble symbols, we are able to extract per-symbol phase measurements and construct a phase trend model that can be applied to subsequent payload symbols.

Specifically, the received preamble symbols are first dechirped and transformed into the frequency domain. For each preamble symbol, the dominant frequency component is identified, and the relative phase between the corresponding spectral components is computed. This yields a sequence of phase observations, one per preamble symbol.

These phase observations are then fitted with a linear model, capturing the overall phase drift across the preamble. The resulting slope and intercept characterize the phase evolution and are used to predict the phase of subsequent payload symbols.

$$\hat{\varphi}(i) = \hat{s} \cdot i + \hat{b} \quad (12)$$

where $\hat{s}$ and $\hat{b}$ denote the estimated slope and intercept, respectively.

This preamble-based phase model allows the receiver to compensate phase variation in a simple and consistent manner, enabling coherent processing in later demodulation stages.

As illustrated in Fig. 4, although the raw phase measurements may exhibit wrapping and fluctuations under low SNR, the estimated phase trend remains smooth and stable, which is sufficient for the purposes of phase compensation in our system.

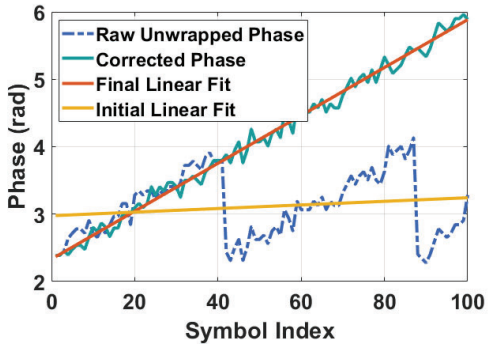

Fig. 4: Example of phase trend estimation using preamble symbols.

B. Peak Alignment and Fusion

In oversampled LoRa receivers where the sampling frequency F_s exceeds the signal bandwidth BW, a single chirp symbol manifests as a characteristic dual-peak structure in the frequency domain after dechirping. The dechirping process is performed by multiplying the received signal $y_r(t)$ with a time-aligned base down-chirp $C^{-1}(t)$:

$$y_s(t) = y_r(t) \cdot C^{-1}(t) = e^{j2\pi f_s t}, \tag{13}$$

which effectively transforms the frequency-modulated chirp into an approximately single-tone signal. The FFT of $y_s(t)$ reveals a spectral peak corresponding to the encoded frequency.

Under oversampling conditions ($F_s > BW$), spectral folding due to FFT periodicity introduces two dominant peaks in the frequency domain: a primary peak at f_x, and a mirrored secondary peak at

$$f_x' = f_x + (F_s - BW). \tag{14}$$

These two peaks originate from different segments of the cyclically shifted chirp in the time domain and are separated in frequency by $F_s - BW$.

Although negligible under high SNR, the energy split between the two peaks becomes problematic at low SNR, where a failure to effectively aggregate the spectral energy can result in diminished peak prominence and symbol misclassification.

Naively summing these two peaks without addressing their phase difference leads to destructive interference or even energy cancellation. To address this, we propose a phase-aligned dual-peak energy aggregation mechanism. The phase offset $\Delta\phi$ between the two peaks, which arises from their time-segment differences, has been analytically modeled in Section III using preamble-based training.

We perform complex-domain phase compensation on the secondary peak before aggregation:

$$X_{\text{agg}} = X[f_x] + X[f_x'] \cdot e^{-j\Delta\phi}, \tag{15}$$

where $X[f_x]$ and $X[f_x']$ denote the FFT amplitudes at the primary and secondary peak positions, respectively. The compensated secondary peak is coherently added to the primary, yielding a reinforced peak energy:

$$E_{\text{agg}} = |X_{\text{agg}}|^2 = |A_1 + A_2|^2, \tag{16}$$

where A_1 and A_2 are the complex amplitudes of the primary and compensated secondary peaks.

This method not only enhances the peak energy and improves the robustness of FFT-based symbol detection, but also mitigates the impact of phase distortion under adverse noise conditions, significantly improving decoding reliability in ultra-low SNR environments.

C. Interference Suppression

In practical LoRa reception, the received signal is often contaminated by wideband noise whose bandwidth far exceeds the actual spectral range of the signal. Since LoRa symbols have spectra strictly confined within a narrow band centered at zero frequency, i.e., $[-BW/2, BW/2]$, this spectral localization provides theoretical basis and implementation space for band-limited filter design.

To enhance demodulation under low SNR, a digital low-pass filter is applied to the received baseband signal prior to FFT demodulation. The filtering operation is given by:

$$y[n] = \sum_{m=0}^{M} h[m]\, x[n-m], \tag{17}$$

where $h[m]$ denotes the filter impulse response. This process preserves the in-band signal components while effectively attenuating out-of-band noise.

From a theoretical perspective, the received signal can be modeled as

$$x(t) = s(t) + n(t),$$

where the signal $s(t)$ has spectral support strictly within $[-BW/2, BW/2]$, and the noise $n(t)$ is zero-mean Gaussian white noise with power spectral density $N_0/2$, uniformly distributed over the sampling bandwidth $[-F_s/2, F_s/2]$. Therefore, the total noise power is

$$P_n = N_0 F_s.$$

After band-limited filtering, the noise power is confined to

$$P_{n,\text{filtered}} = N_0 BW,$$

while the signal power remains approximately unchanged. The corresponding SNRs before and after filtering are

$$\text{SNR}_{\text{raw}} = \frac{P_s}{N_0 F_s}, \quad \text{SNR}_{\text{filtered}} = \frac{P_s}{N_0 BW},$$

yielding an SNR gain of

$$G_{\text{SNR}} = 10 \log_{10} \frac{F_s}{BW}.$$

In the frequency domain, the filtered noise remains a complex Gaussian process with amplitude following a Rayleigh distribution [8]:

$$f_{|N|}(x) = \frac{x}{\sigma_k^2} \exp\left(-\frac{x^2}{2\sigma_k^2}\right), \quad x \geq 0,$$

where σ_k^2 is proportional to the noise power. For an FFT length N, the maximum noise peak approximately follows

$$M \approx \sigma_k \sqrt{2 \ln N}.$$

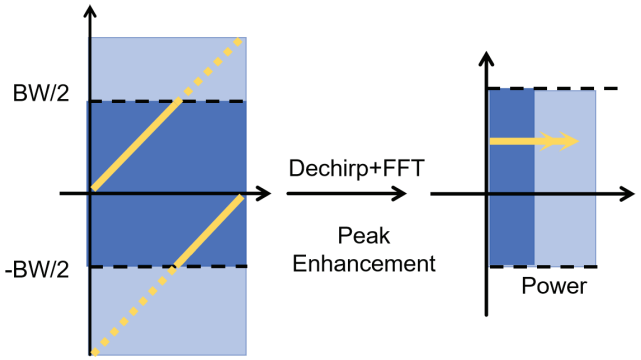

(a) FFT spectrum without filtering.

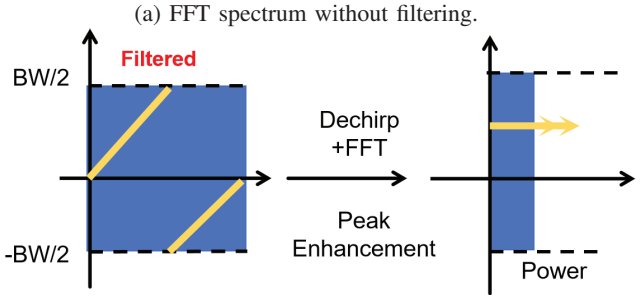

(b) FFT spectrum after low-pass filtering.

Fig. 5: Comparison of demodulation performance before and after applying low-pass filtering. Filtering suppresses out-of-band noise, enhancing peak detectability.

The band-limited filtering reduces noise power, thereby decreasing σ_k and lowering the maximum noise peak M, which significantly reduces false detection probability and enhances decoding robustness under low SNR conditions.

VI. Implementation

A. Experimental Setup

We evaluate the proposed method using both software-defined radios and commercial LoRa devices to reflect practical deployment conditions. Our testbed includes USRP N210 and B210 platforms with SBX daughterboards, connected to host PCs for baseband processing implemented in GNU Radio and MATLAB.

In addition to controlled experiments, we conduct field measurements using off-the-shelf SX1276-based LoRa nodes in representative real-world environments. As illustrated in Fig. 6, a campus-scale outdoor testbed is used to evaluate performance under typical urban propagation conditions, including line-of-sight and non-line-of-sight links around buildings and vegetation. To further assess robustness under harsh industrial impairments, we deploy LoRa nodes in an operational industrial environment, as shown in Fig. 7, where strong electromagnetic interference, metallic structures, and dense machinery introduce severe attenuation and multipath effects.

Experiments cover spreading factors from 7 to 12 and bandwidths of 125, 250, and 500,kHz, with a sampling rate of 1,MHz to enable dual-peak spectral fusion.All transmissions operate at a center frequency of 470,MHz, and controlled low-SNR conditions are generated using programmable attenuators and AWGN injection.

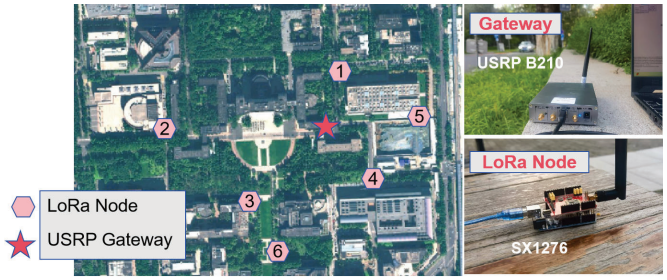

Fig. 6: Bird view of outdoor testbed and hardware setup.

Fig. 7: Representative industrial deployment environments in an operational factory.

We compare *UniChirp* with several representative LoRa weak signal demodulation mechanisms:

- **Ostinato** [9] uses a brute-force search to align dual peaks by scanning phase compensation values.
- **Choir** [10] merges dual peaks via downsampling without phase alignment.
- **LoRaTrimmer** [7] trims noisy FFT windows and applies probabilistic modeling to mitigate frequency-induced phase discontinuities.
- **NELoRa** [11] uses neural network for weak LoRa signal demodulation.

Evaluation metrics include demodulation success rate, residual phase error, and spectral energy gain, with each experiment repeated on 200 randomized packets for statistical reliability.

VII. Evaluation

A. Overall Performance

In our evaluation, the reported SNR refers to the pre-despreading signal-to-noise ratio measured at the raw complex baseband (IQ) signal. Additive white Gaussian noise (AWGN) is injected directly into the received IQ samples prior to dechirping and symbol-level processing. Unless otherwise stated, we report uncoded symbol error rate (SER), measured before channel coding and interleaving, enabling a consistent comparison across different demodulation methods under iden-

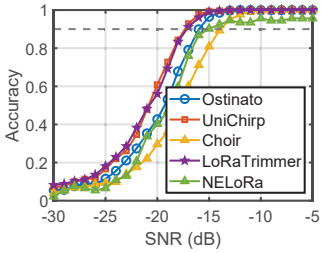

Fig. 8: Symbol accuracy versus SNRs.

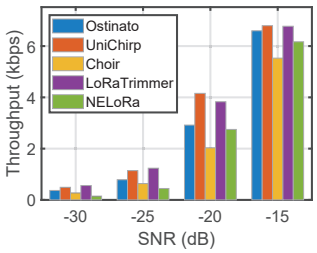

Fig. 9: PHY-layer throughput at SNR points.

tical raw channel conditions. For clarity, symbol accuracy (1 - SER) is reported in the figures.

To evaluate robustness under extremely low SNR conditions, we compare the proposed method with representative spectrum-based and learning-based baselines, including *Ostinato*, *Choir*, *LoRaTrimmer*, and *NELoRa*, under the configuration of SF=7 and BW=125 kHz. As shown in Fig. 8, the comparison is based on the minimum SNR required to achieve a decoding accuracy of 90%, which serves as a practical indicator of noise tolerance. The corresponding impact on effective throughput at representative SNR points is shown in Fig. 9.

The results show that *UniChirp* consistently reaches the target decoding accuracy at lower SNR levels than all considered baselines, indicating improved robustness in weak-signal regimes. This advantage primarily stems from the explicit modeling and compensation of in-chirp phase distortion, which enables coherent energy concentration even when noise dominates. In contrast, Ostinato and similar methods rely on phase-agnostic or search-based aggregation, whose effectiveness rapidly degrades when noise perturbs phase estimates.

Compared with learning-based approaches such as *NELoRa*, *UniChirp* avoids dependence on training data distribution and generalization behavior. By exploiting explicit signal structure and preamble-guided phase modeling, *UniChirp* maintains stable and predictable decoding performance across the ultra-low SNR regime. Similarly, while *LoRaTrimmer* achieves comparable decoding performance, it incurs higher computational complexity due to matrix-based operations that scale with the spreading factor, whereas *UniChirp* achieves robustness using a lightweight linear phase compensation strategy.

Overall, these results demonstrate that the proposed design achieves superior robustness under extreme noise conditions by directly addressing the fundamental phase misalignment issue in chirp-based demodulation, rather than relying on search heuristics or data-driven fitting.

B. Parameter Robustness

To evaluate the robustness and generalization capability of the proposed method under different LoRa configurations, we compare it with the baseline method *Ostinato*. In the spreading factor (SF) comparison, the bandwidth (BW) is fixed at 125 kHz, while the bandwidth comparison fixes the SF at 10 for all experiments.

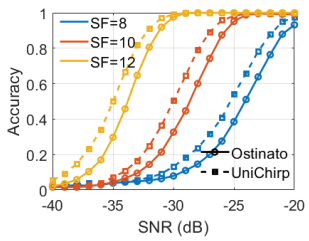

(a) Impact of SF

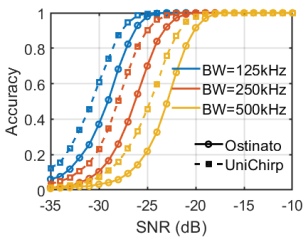

(b) Impact of BW

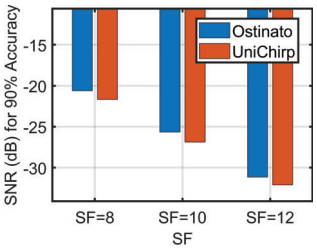

(c) SNR limits across SFs

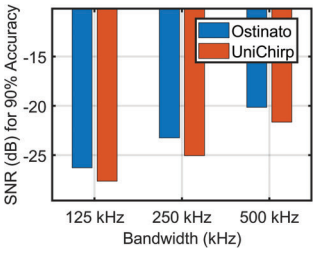

(d) SNR limits across BWs

Fig. 10: Performance comparison under different spreading factors (SFs) and bandwidths (BWs).

Figure 10a shows the decoding accuracy under various SNRs, and Fig. 10c reports the minimum SNR required to achieve 90% accuracy. Across all tested SF values, the proposed method consistently outperforms the baseline. This improvement is mainly attributed to the proposed analytical modeling of in-chirp phase misalignment, which remains valid across different symbol durations induced by varying SFs. By explicitly compensating phase distortion at the symbol level, *UniChirp* preserves coherent energy aggregation even under long-symbol configurations. In contrast, *Ostinato* relies on a search-based spectral aggregation strategy, which becomes more susceptible to noise under small-SF configurations where the inherent processing gain is limited, where the reduced processing gain leaves limited margin for reliable phase search under noise.

Similarly, the comparison under different bandwidth settings (Fig. 10b and Fig. 10d) further demonstrates the robustness of *UniChirp* . Since the proposed phase modeling and compensation mechanism does not depend on specific bandwidth values, *UniChirp* maintains stable decoding performance and reaches the target accuracy at lower SNRs across a wide range of bandwidth configurations. Overall, these results indicate that the proposed modeling-based design generalizes well across diverse LoRa parameter settings, offering a more scalable and robust alternative to search-based aggregation approaches for practical deployments in heterogeneous and dynamic network environments.

C. Computational Cost and Latency

Beyond decoding robustness, we further evaluate the computational cost and symbol-level latency of *UniChirp* and compare it with representative baseline methods, including Ostinato, Choir, and LoRaTrimmer. All experiments in this subsection are conducted under the configuration of SF=12

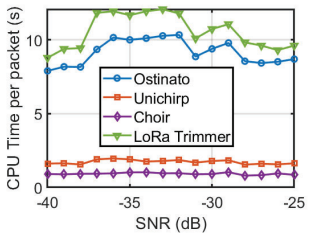

Fig. 11: CPU processing time.

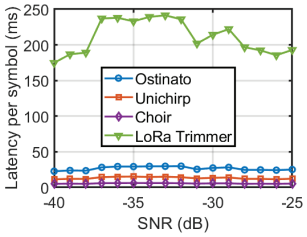

Fig. 12: Per-symbol processing latency.

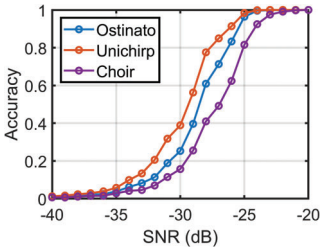

Fig. 13: Symbol-level accuracy under industrial conditions.

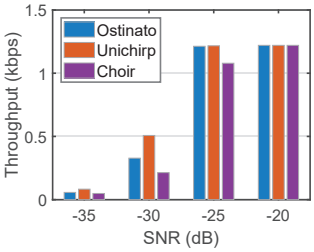

Fig. 14: Throughput at representative SNRs in the industrial field trial.

and BW=125 kHz, which represents a typical long-range LoRa setting under weak-signal conditions. These metrics are critical for practical LoRa gateways that must process concurrent transmissions in real time.

Fig. 11 shows the average CPU processing time per packet as a function of SNR. *UniChirp* consistently exhibits lower processing time across the evaluated SNR range. While Ostinato and LoRaTrimmer incur high computational overhead due to repeated spectral aggregation and matrix-based operations, *UniChirp* maintains a stable and low processing cost. This efficiency mainly stems from its closed-form phase estimation and linear compensation, which avoid exhaustive search or high-dimensional computations. Fig. 12 further reports the average per-symbol processing latency over the payload, showing that *UniChirp* also achieves lower and more stable symbol-level latency across different SNR levels, indicating consistent runtime behavior even in noise-dominated operating regimes.

Although Choir exhibits relatively low processing latency due to its simplified spectral aggregation pipeline, its decoding robustness degrades rapidly under severe noise conditions at this operating point. Taken together, these results indicate that *UniChirp* achieves a more favorable balance between decoding robustness and computational efficiency across different operating regimes. By combining lightweight phase compensation with efficient spectral processing, *UniChirp* enables reliable low-SNR demodulation without incurring excessive computational or latency overhead.

D. Industrial Field Trial

To further examine the behavior of *UniChirp* under industrial propagation conditions, we conduct a field trial in an operational industrial environment. Compared with the campus-scale outdoor testbed, this evaluation emphasizes deployment scenarios with dense metallic structures and complex layouts, which introduce severe non-line-of-sight propagation and strong multipath effects.

During the field trial, LoRa sensor nodes periodically transmit packets using the same physical-layer configuration as in the testbed experiments. For the received signals captured at the gateway, we analyze the decoding performance using the uncoded SER metric defined in Section VII. This evaluation focuses on assessing the robustness of the demodulation process under realistic industrial propagation conditions, rather than end-to-end system-level delivery performance.

Fig. 13 shows that *UniChirp* consistently achieves lower uncoded SER than the baseline method across the evaluated industrial deployment scenarios. This performance improvement is attributed to the proposed model-guided phase compensation, which mitigates in-chirp phase mis-alignments, as well as the structure-aware spectral filtering that enhances the prominence of valid FFT peaks under strong industrial impairments. As a result, the improved decoding reliability directly translates into higher effective throughput at representative SNR points, as illustrated in Fig. 14.

Overall, this industrial field trial complements the controlled testbed evaluation by demonstrating that the proposed demodulation framework remains effective in representative industrial environments. It further highlights the potential of *UniChirp* for enhancing LoRa demodulation robustness in complex real-world deployment scenarios.

VIII. Related Work

A. Enhancing LoRa Communication Range and Efficiency

To address LoRa's limited communication range and severe signal attenuation in urban and industrial environments, recent research has explored signal enhancement techniques at the receiver side. For example, CHOIR [10] leverages up to 36 co-located nodes to collaboratively transmit the same data, thereby enhancing signal strength at the gateway. CHARM [5] synchronizes 2–8 spatially distributed gateways to enable energy joint decoding of weak signals. CHIME [12] uses periodic heartbeat packets to estimate channel conditions and selects frequency channels with lower interference. XCopy [13] applies coherent merging techniques to optimize retransmitted packets, while NEPHALAI [14] compresses received signals and uploads them to the cloud, using sparse reconstruction to enable demodulation. Tong et al. [15] address link instability in urban LoRa deployments caused by building obstructions and multipath fading. They propose link dynamics modeling and adaptive connection management to enhance communication robustness and reliability. These methods significantly improve reliability and decoding coverage in complex environments, but often rely on infrastructure-level changes or incur high computational complexity and system delays due to cloud offloading. From a network and system-level perspective, several studies have explored multi-gateway cooperation and cloud-assisted architectures to further extend LoRa coverage. ALOBA [16] aggregates uplink signals from multiple gateways to enable joint decoding and improved reception diversity, while cloud-based LoRa RAN designs [17] offload signal

processing to centralized servers for enhanced scalability and reliability. These approaches highlight the benefits of spatial and infrastructure diversity, but often introduce additional backhaul overhead and processing latency.

B. Trends and Challenges in Enhanced LoRa Demodulation

In addition, several recent works have aimed to improve demodulation performance under low SNR conditions. NELORA [11] introduces a neural-enhanced demodulation framework that learns denoised representations to recognize ultra-low SNR signals effectively. LORATRIMMER [7] proposes a chirp clipping technique to compress energy and improve the detectability of short or attenuated chirps. Meanwhile, OSTINATO [9] presents an oversampling-based dechirping technique with phase alignment that mitigates SNR loss caused by symbol frequency shifts. XU ET AL. [6] systematically restructured the LoRa PHY layer and demonstrated the profound impact of hardware non-idealities on spectral features. These studies highlight the importance of modeling and compensating for abnormal physical-layer behavior in real-world environments and underscore the need for further exploration of signal phase evolution and frequency-domain structure recovery. In parallel, concurrent transmission decoding has emerged as an important direction for improving LoRa demodulation under high network load. CoLoRa [18] and PLoRa [19] exploit frequency and timing diversity to recover colliding packets, while subsequent studies [20], [21] further investigate the limits of collision resolution in practical deployments. These works emphasize the critical role of precise frequency and phase modeling, yet often require stringent assumptions on signal separability or collision patterns.

IX. CONCLUSION

We present *UniChirp*, a practical LoRa demodulation framework that significantly improves robustness under weak-signal conditions. By explicitly modeling phase deviations caused by hardware offsets and modulation artifacts, *UniChirp* derives a lightweight analytical formulation enabling deterministic phase compensation. Building on this, we develop a preamble-guided phase extrapolation mechanism and a structure-aware spectral filtering strategy to enhance symbol-level demodulation reliability. We evaluate the framework using commercial LoRa hardware in controlled low-SNR experiments, outdoor testbeds, and real industrial production environments. Results demonstrate that *UniChirp* provides a lightweight, deployable solution that enables robust LoRa communication in challenging industrial propagation environments, achieving reliable demodulation at SNRs as low as -17.3 dB without increasing system complexity.

REFERENCES

[1] L. Alliance, "A technical overview of lora and lorawan.(2015)," *https://lora-alliance. org/wp-content/uploads/2020/11/what-is-lorawan. pdf¿. Citado*, vol. 2, p. 24, 2015.

[2] L. Liu, Y. Yao, Z. Cao, and M. Zhang, "Deeplora: Learning accurate path loss model for long distance links in lpwan," in *IEEE INFOCOM 2021-IEEE Conference on Computer Communications*. IEEE, 2021, pp. 1–10.

[3] S. Demetri, M. Zúñiga, G. P. Picco, F. Kuipers, L. Bruzzone, and T. Telkamp, "Automated estimation of link quality for lora: A remote sensing approach," in *Proceedings of the 18th International Conference on Information Processing in Sensor Networks*, 2019, pp. 145–156.

[4] Y. Yao, Z. Ma, and Z. Cao, "Losee: Long-range shared bike communication system based on lorawan protocol." in *EWSN*, 2019, pp. 407–412.

[5] A. Dongare, R. Narayanan, A. Gadre, A. Luong, A. Balanuta, S. Kumar, B. Iannucci, and A. Rowe, "Charm: Exploiting geographical diversity through coherent combining in low-power wide-area networks," in *2018 17th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, 2018, pp. 60–71.

[6] Z. Xu, S. Tong, P. Xie, and J. Wang, "From demodulation to decoding: Toward complete lora phy understanding and implementation," *ACM Transactions on Sensor Networks*, vol. 18, no. 4, pp. 1–27, 2023.

[7] J. Du, Y. Liu, Y. Ren, L. Liu, and Z. Cao, "Loratrimmer: Optimal energy condensation with chirp trimming for lora weak signal decoding," in *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*, 2024, pp. 1104–1118.

[8] D. Roy, "Discrete rayleigh distribution," *IEEE transactions on reliability*, vol. 53, no. 2, pp. 255–260, 2004.

[9] Z. Xu, P. Xie, J. Wang, and Y. Liu, "Ostinato: Combating lora weak links in real deployments," in *2022 IEEE 30th International Conference on Network Protocols (ICNP)*. IEEE, 2022, pp. 1–11.

[10] 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.

[11] C. Li, H. Guo, S. Tong, X. Zeng, Z. Cao, M. Zhang, Q. Yan, L. Xiao, J. Wang, and Y. Liu, "Nelora: Towards ultra-low snr lora communication with neural-enhanced demodulation," in *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems*, 2021, pp. 56–68.

[12] A. Gadre, R. Narayanan, A. Luong, A. Rowe, B. Iannucci, and S. Kumar, "Frequency configuration for {Low-Power}{Wide-Area} networks in a heartbeat," in *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*, 2020, pp. 339–352.

[13] 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.

[14] J. Liu, W. Xu, S. Jha, and W. Hu, "Nephalai: towards lpwan c-ran with physical layer compression," in *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, 2020, pp. 1–12.

[15] S. Tong, Z. Shen, Y. Liu, and J. Wang, "Combating link dynamics for reliable lora connection in urban settings," in *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*, 2021, pp. 642–655.

[16] X. Guo, L. Shangguan, Y. He, J. Zhang, H. Jiang, A. A. Siddiqi, and Y. Liu, "Aloba: Rethinking on-off keying modulation for ambient lora backscatter," in *Proceedings of the 18th conference on embedded networked sensor systems*, 2020, pp. 192–204.

[17] A. Balanuta, N. Pereira, S. Kumar, and A. Rowe, "A cloud-optimized link layer for low-power wide-area networks," in *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*, 2020, pp. 247–259.

[18] 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.

[19] Y. Peng, L. Shangguan, Y. Hu, Y. Qian, X. Lin, X. Chen, D. Fang, and K. Jamieson, "Plora: A passive long-range data network from ambient lora transmissions," in *Proceedings of the 2018 conference of the ACM special interest group on data communication*, 2018, pp. 147–160.

[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. Yang, Z. Liu, K. Tang, X. Yin, C. Zhuo, Q. Wei, and F. Qiao, "Breaking the energy-efficiency barriers for smart sensing applications with "sensing with computing" architectures," *Science China Information Sciences*, vol. 66, no. 10, p. 200409, 2023.