A Comparative Study of DFT-Based Detectors for Mitigating Multipath Interference in LoRa via Cyclic Prefix

Ahmet Kerem Yumusak * and Mehmet Bulut

Posted Date: 20 May 2026

doi: 10.20944/preprints202605.1366.v1

Keywords: cyclic prefix; discrete Fourier transform; equalization; LoRa; multipath fading

 doi:10.20944/preprints202605.1366.v1

Article

A Comparative Study of DFT-Based Detectors for Mitigating Multipath Interference in LoRa via Cyclic Prefix

Ahmet Kerem Yumuşak * and Mehmet Bulut

Department of Electrical and Electronics Engineering, Atılım University, Ankara, Turkey

* Correspondence: yumusak.ahmetkerem@student.atilim.edu.tr

Abstract

Long Range (LoRa) is a chirp spread spectrum (CSS) physical-layer technology that has become a leading candidate for low-power wide-area network (LPWAN) connectivity in the Internet of Things (IoT). At the receiver, the standard demodulator multiplies the incoming signal with a conjugate reference chirp and applies a one-dimensional discrete Fourier transform (DFT), reducing symbol detection to peak search in the frequency domain. While this non-coherent baseline is simple and robust under additive white Gaussian noise (AWGN), its symbol error rate (SER) degrades significantly in frequency-selective multipath channels, where parasitic spectral peaks distort the dominant tone. This paper presents a unified comparative study of seven LoRa detectors for spreading factor seven, six of which share a common one-dimensional DFT engine while a matched-filter bank operates directly in the time domain, with the six DFT detectors differing in their per-bin frequency-domain weighting and decision rule. The detector set spans the standard non-coherent DFT, a non-coherent matched filter bank, two coherent equalizers in the frequency domain (zero-forcing and minimum mean-square error), a phase-only equalizer, a maximal-ratio combiner with non-coherent decision, and an exhaustive maximum-likelihood detector that serves as a near-optimal reference under the same preamble-based CSI. To mitigate inter-symbol interference in the multipath case, every transmitted symbol is preceded by a cyclic prefix that converts the linear convolution with the channel into a circular convolution, enabling per-bin frequency-domain processing. Throughout the paper a deployment-realistic receiver model is adopted: the per-bin channel response is estimated by a frequency-domain least-squares estimator from a short preamble, and the noise variance is estimated blindly from the preamble residuals. The quality of the noise-variance estimator is reported separately as a diagnostic. Each detector is evaluated under both AWGN and a two-tap Rayleigh multipath channel through Monte Carlo simulation, and its execution time per call is recorded to provide a complementary view of computational cost. The framework introduced here clarifies how coherent processing, diversity combining, equalization, and exhaustive search trade detection performance against complexity within a single DFT-centric LoRa receiver architecture. The principal quantitative finding is that, under the two-tap Rayleigh multipath channel, the MMSE equalizer reaches SER $\approx 4.4\times10^{-5}$ at SNR = −5 dB and tracks the exhaustive maximum-likelihood detector within 0.1 dB across the full SNR sweep, while costing only 1.26× the per-symbol time of the standard DFT receiver. Conversely, the standard non-coherent baseline hits an irreducible 16% error floor and the unregularized zero-forcing equalizer fails to reach the 10^{-2} SER level at any SNR considered, isolating MMSE as the recommended choice in the multipath regime at every SNR for which a LoRa link is operationally viable.

Keywords: cyclic prefix; discrete Fourier transform; equalization; LoRa; multipath fading

 doi:10.20944/preprints202605.1366.v1

I. Introduction

LOW-POWER wide-area network (LPWAN) technologies have emerged as a central enabler of large-scale Internet of Things (IoT) deployments, providing kilometer-range connectivity to battery-operated sensor nodes at very low data rates and very low energy cost [1]. Among the available standards, LoRa has gained particularly broad commercial and academic traction, largely because its medium-access layer (LoRaWAN) is openly specified and its physical layer can be deployed in unlicensed sub-GHz industrial, scientific, and medical (ISM) bands. The LoRa physical layer is built on chirp spread spectrum (CSS) modulation, in which the information symbol is encoded as a cyclic time shift of a baseband linear chirp [2]. By selecting a spreading factor (SF) between 6 and 12, the system can trade throughput against link sensitivity, an attribute that is particularly attractive for heterogeneous IoT scenarios.

Although the modulation scheme itself is patented and was never publicly described in detail by the manufacturer, several authors have reverse-engineered and analyzed it from a signal-processing perspective. Vangelista showed that LoRa can be interpreted as frequency shift chirp modulation and that an optimal low-complexity receiver consists of dechirping the received signal with the conjugate of the reference up-chirp followed by a one-dimensional discrete Fourier transform (DFT) [2]. Chiani and Elzanaty later confirmed and refined this view, providing a rigorous waveform and spectral characterization that has since become a de facto reference for analytical work on LoRa [3]. The DFT is therefore the mathematical method on which essentially every practical LoRa receiver relies, and is also the core analytical tool employed throughout this study.

Most existing analytical and simulation studies focus on a single receiver structure—typically the standard non-coherent DFT-based detector under AWGN or flat Rayleigh fading [4,5]. In practical deployments, however, the channel is often frequency-selective: the transmitted chirp arrives along multiple delayed paths whose superposition introduces parasitic peaks at the DFT output [6,7]. Under such conditions, the standard detector is no longer optimal. In parallel, the LoRa literature has investigated coherent detection more broadly: practical low-complexity coherent receivers under AWGN [8] and coherent detectors evaluated under AWGN with co-channel interference [9]. For the multipath case specifically, alternative receivers have been proposed in the form of matched-filter or RAKE structures that combine multipath energy constructively [7]. To the best of the present authors' knowledge, however, very few works place all of these architectures on a common testbench, evaluate them under identical channel and synchronization assumptions, and report not only their symbol error rate (SER) but also their measured execution time. The contribution of this paper is to fill that gap.

Specifically, this work investigates seven LoRa detectors for SF = 7 (i.e., N = 128 samples per symbol) and reports their symbol error rate and execution time under both AWGN and a two-tap Rayleigh multipath channel. The seven detectors are: the standard DFT detector, which acts as a non-coherent baseline reflecting the deployed-receiver reference; a non-coherent matched filter bank operating directly on the received vector, included as a complexity control against the DFT-based receivers; a coherent zero-forcing equalizer in the frequency domain, included as a cautionary baseline that exposes the noise-enhancement penalty motivating MMSE regularization; a coherent minimum mean-square error (MMSE) equalizer, the proposed candidate of this study; a phase-only equalizer that compensates only the channel phase, included as a variant that is robust to amplitude-estimation errors; a maximal-ratio combining (MRC) front end followed by a non-coherent magnitude detector, included as a diversity-gain fallback for situations where coherent decision is unreliable; and an exhaustive maximum-likelihood (ML) detector that serves as the near-optimal performance reference under the same preamble-based CSI assumption rather than a true genie-aided upper bound. To mitigate inter-symbol interference (ISI) in the multipath case, each transmitted symbol is preceded by a cyclic prefix (CP) longer than the maximum tap delay, which converts the linear convolution with the channel into a circular convolution and enables single-tap per-bin frequency-domain processing [10]. The use of a CP with LoRa signals was previously explored by Li et al. [11], who combined it with a RAKE receiver to mitigate ISI under ITU and COST-

 doi:10.20944/preprints202605.1366.v1

207 multipath profiles; the present work extends that line by adopting CP within a unified DFT-based detector comparison framework. To remain faithful to deployment conditions, no detector is given oracle channel-state information (CSI) or oracle noise variance: every channel-aware detector estimates the per-bin channel response by a frequency-domain least-squares (LS) estimator from a short preamble, and every detector that requires a noise-variance estimate (only the MMSE equalizer in the present set) recovers it blindly from the preamble residuals. The accuracy of the noise-variance estimator is reported separately as a numerical diagnostic so that the reader can independently judge the validity of the SER curves. Cast in these terms, the central technical contribution of this paper is the introduction of OFDM-style coherent equalization (zero-forcing, MMSE, and phase-only) and per-bin diversity combining to LoRa under a deployment-realistic preamble-based estimation pipeline, made possible by the cyclic-prefix front end.

The role of artificial intelligence in this project is intentionally null. No machine-learning models, neural-network detectors, or learned equalizers are used. The motivation for this choice is twofold. First, the present study is meant to isolate the impact of classical signal-processing principles—coherent versus non-coherent detection, equalization, diversity combining, and exhaustive search—on LoRa SER, free from the additional variance introduced by data-driven methods. Second, the comparative framework developed here is intended to serve as a baseline against which future learning-based LoRa detectors—such as the hybrid CNN / matched-filter architecture recently proposed by Dakic et al. [12]—can be benchmarked. The reader interested in broader perspectives on LoRaWAN networking, deployment, and physical-layer detection is referred to the survey literature [1], and to the recent surveys and tutorials on LoRaWAN scalability and chirp spread spectrum modulation [13,14].

The remainder of this paper is organized as follows. Section II reviews related work on LoRa detector design, equalization strategies, and performance analysis under fading. Section III formally defines the engineering problem, the system boundaries, and the input–output relationships of the receiver. Section IV develops the mathematical foundation of the one-dimensional DFT and derives the per-bin frequency-domain channel model that underlies every detector considered. Section V describes the simulation environment, the seven detector implementations, and the cyclic-prefix-based transmission scheme. Section VI presents the simulation results and analysis, and Section VII concludes the paper.

II. Literature Review

Theoretical analysis of LoRa modulation began with Vangelista's signal-processing formulation, in which the LoRa symbol is described as a cyclically time-shifted up-chirp and the optimal low-complexity demodulator is shown to consist of multiplication by the complex conjugate of the reference chirp followed by a DFT and a peak search [2]. Chiani and Elzanaty independently developed a closely related formulation, additionally analyzing the spectral occupation of the modulated waveform and confirming the DFT-based receiver as essentially optimal in the AWGN regime [3]. Together, these two works established the analytical framework upon which the present study—and most subsequent LoRa detection literature—rests.

Closed-form approximations of LoRa bit-error and symbol-error rates were subsequently derived for AWGN and flat Rayleigh fading channels by Elshabrawy and Robert [4]. Their expressions confirm that LoRa exhibits a clear waterfall behavior in AWGN and an inverse-SNR-like behavior over Rayleigh fading, mirroring classical M-ary orthogonal modulation. The framework of [4] is widely used as a sanity check for Monte Carlo simulation results and is also adopted, in adapted form, in Section V of this work for cross-validation against the AWGN baseline.

Practical receiver architectures and their non-idealities have been investigated by Ghanaatian et al., who described the basic LoRa receiver chain, including the effect of carrier and sampling frequency offsets, and proposed compensation schemes [15]. Their analysis underlines the dominance of the standard non-coherent DFT detector in deployed receivers, due primarily to its robustness to phase impairments. Nguyen et al. [8] complemented this line of work from the

 doi:10.20944/preprints202605.1366.v1

transmitter side, proposing an orthogonal chirp generator (OCG) that supports multiple spreading factors and bandwidths at low hardware cost, providing a closed-form analysis of both coherent and non-coherent CSS demodulation under AWGN, and introducing a phase-shifted CSS (PS-CSS) variant in which extra information is encoded in the starting phase of each chirp symbol. The authors of [8] explicitly defer the design of the corresponding timing, frequency, and phase synchronization receiver to a follow-up paper, leaving the practical realization of coherent CSS detection an open problem at the time of writing.

The performance gap between coherent and non-coherent LoRa detection in interference-limited scenarios was recently characterized by Afisiadis et al. [5,9]. In [5], the authors derived rigorous SER and frame-error-rate expressions for LoRa in the presence of a second co-channel LoRa interferer, and in [9] they showed that, while the coherent gain over AWGN alone is only of the order of a fraction of a dB — approximately 0.6 dB at SF = 6 as reported in [8], consistent with classical M-ary orthogonal modulation theory — the same coherent processing yields up to 10 dB of gain at low signal-to-interference ratios. This non-trivial benefit motivates the inclusion of explicitly coherent detectors—zero-forcing, MMSE, and phase-only—in the comparative study presented here.

LoRa performance over multipath channels has received less attention than the AWGN and flat-fading cases, but two recent contributions are particularly relevant. Demeslay et al. [6] derived a semi-analytical SER approximation for LoRa under frequency-selective fading using the standard DFT-based detector, demonstrating that significant echoes produce parasitic peaks in the DFT output and degrade the detector. The same authors proposed in [7] a matched-filter / RAKE receiver that combines multipath energy constructively at the symbol-index frequency of the DFT, thereby improving SER under non-trivial multipath conditions. A closely related contribution by Li et al. [11] introduced a cyclic prefix to the LoRa frame and showed, through Monte Carlo simulations over the ITU Pedestrian-A, Vehicular-A, COST-207 Typical Urban, and Hilly Terrain profiles, that the combined CP + RAKE receiver eliminates the parasitic peaks of the standard detector and recovers more than 10 dB at low SER targets in heavily dispersive channels. The MRC-based detector studied in this paper is conceptually related: it weights each frequency bin by the conjugate of the per-bin channel response — the frequency-domain dual, under cyclic prefix, of the time-domain path-weighting performed by a classical RAKE — but takes a non-coherent magnitude decision after combining, exploiting frequency diversity without requiring a coherent decision rule. Chen et al. [16] further studied per-bin frequency-domain equalization for frequency shift chirp modulation under a two-tap multipath model essentially identical to the one adopted in this paper, applying both ZF and MMSE weights and combining them with Hamming and LDPC channel coding to extract bit-likelihood information from the post-DFT amplitudes; their results independently confirm both the noise-enhancement penalty of unregularized ZF and the favorable trade-off achieved by the MMSE regularizer of (11), and provide a concrete soft-decision starting point for the extensions discussed in Section VII-C.

Equalization techniques for orthogonal frequency-division multiplexed (OFDM)-like systems with cyclic prefix have been studied for decades. The classical zero-forcing and MMSE equalizers operate per-subcarrier in the frequency domain and trade off noise enhancement against residual interference [10]. Although LoRa is not an OFDM system, the cyclic prefix permits the same per-bin diagonalization of the channel matrix once the received signal is processed by a DFT. Brennan's seminal work on linear diversity combining [17] further provides the theoretical basis for maximal-ratio combining, which is here adapted to the per-bin frequency-domain combining of LoRa chirps over multipath channels.

In summary, the existing literature has established (i) the optimality of DFT-based detection in AWGN [2,3]; (ii) closed-form SER expressions for AWGN and flat Rayleigh fading [4]; (iii) the practical dominance of the non-coherent DFT receiver in deployed systems [15]; (iv) the benefit of coherent detection in interference-limited regimes [9]; and (v) the parasitic-peak problem of the standard detector in multipath, together with first proposals for matched-filter / RAKE-type receivers and CP-based front ends for LoRa [6,7,11]. What is missing—and what motivates the present work—

 doi:10.20944/preprints202605.1366.v1

is a unified comparison that places non-coherent, partially coherent, fully coherent equalized, diversity-combining, and exhaustive ML detectors on the same testbench, with the same channel realization and the same cyclic-prefix front end, and that reports both SER and per-detector execution time. The innovative contributions of this study are therefore the following:

1) a unified one-dimensional DFT framework that supports six of the seven heterogeneous detectors, the seventh being a non-coherent matched-filter bank that operates directly in the time domain and is included as a complexity reference, with identical signal generation, channel realization, and cyclic-prefix processing;

2) the introduction of cyclic-prefix-based transmission for LoRa SF = 7, which is not part of the standard but which transforms the multipath channel into a per-bin diagonal channel and enables OFDM-style equalization;

3) the proposal and evaluation of explicitly coherent LoRa detectors (zero-forcing, MMSE, phase-only) and a non-coherent MRC variant in addition to the standard non-coherent baseline, all operating on a deployment-realistic preamble-estimated channel and a residual-based blind noise-variance estimate; and

4) the systematic measurement of per-detector execution time alongside SER, providing a joint accuracy/complexity view that is rarely reported in the LoRa detection literature.

III. Problem Definition

A. Engineering Problem and Motivation

The engineering problem addressed in this work is the symbol-level demodulation of a LoRa CSS signal received over a possibly frequency-selective wireless channel, under perfect synchronization but with channel-state information (CSI) and noise variance estimated from a short preamble at the receiver, and limited to a one-dimensional DFT as the underlying transformation. The motivation is twofold. First, the DFT is the only frequency-domain tool used by the de facto LoRa receiver [2,3], and any architectural change—coherent equalization, diversity combining, exhaustive search—must in practice still rest on it for computational reasons. Second, real LoRa deployments increasingly operate in indoor, urban, and industrial environments where the assumption of a flat channel does not hold; the multipath channel introduces parasitic peaks at the DFT output that degrade the standard detector [6]. The problem is therefore to identify which detector, among a representative set of DFT-based receivers, achieves the best trade-off between SER and execution time when a cyclic prefix is added to combat ISI and the channel and noise statistics are estimated from a short preamble, as detailed in Section IV-E.

B. Physical and Theoretical System Description

The system under study is a baseband-equivalent LoRa link with spreading factor SF = 7, bandwidth BW = 125 kHz, and chip rate equal to the bandwidth. Each symbol consists of $N = 2^{SF} = 128$ complex baseband samples that constitute a cyclically time-shifted up-chirp. A cyclic prefix of length $L_{CP} = 8$ samples is prepended to every symbol to absorb the multipath spread of the channel. The transmitter generates the chip sequence in the digital baseband, and no carrier or sampling-frequency offsets are present at the receiver. The channel is modeled either as a pure additive white Gaussian noise (AWGN) channel or as a two-tap frequency-selective channel with delays of zero and three samples (≈ 0 and 24 μs at BW = 125 kHz) and Rayleigh-distributed amplitudes, in addition to AWGN. The receiver removes the cyclic prefix and applies one of the seven candidate detectors, all of which produce an estimate of the transmitted symbol index.

C. Input–Output Relationships and System Boundaries

The transmitted vector for symbol m is denoted $s_m[n]$, n = 0, …, N − 1. Cyclic-prefix insertion produces the length-$(N + L_{CP})$ augmented vector $\tilde{s}_m[n] = s_m[(n - L_{CP}) \bmod N]$, n = 0, …, $N + L_{CP} - 1$, which copies the last L_{CP} samples of $s_m[n]$ to the front, yielding a 136-sample transmitted block at SF

 doi:10.20944/preprints202605.1366.v1

= 7. After cyclic-prefix insertion, the channel output is the linear convolution of the prepended waveform with the discrete-time channel impulse response h[n], plus complex Gaussian noise. After cyclic-prefix removal at the receiver, the input–output relation in the frequency domain reduces to a per-bin product Y[k] = H[k] · S[k] + η[k], where uppercase letters denote N-point DFTs and η[k] is the DFT of the noise vector. The boundary of the system considered here is therefore the digital baseband: the analog front end, sampling, and synchronization stages are explicitly excluded and assumed ideal. This boundary choice isolates the effect of the detector itself from upstream impairments, which is precisely the comparative purpose of this work.

IV. Mathematical Method

A. The One-Dimensional Discrete Fourier Transform

The mathematical method underlying six of the seven detectors evaluated in this paper (the exception, the non-coherent matched-filter bank, operates directly in the time domain and is treated separately in Section IV-B) is the one-dimensional discrete Fourier transform of length N. For a complex-valued sequence x[n] of length N, the DFT and its inverse are defined as

$$X[k] = \sum_{n=0}^{N-1} x[n].\, e^{(-j2\pi kn/N)},\ k = 0, \ldots, N-1, \tag{1}$$

$$x[n] = \frac{1}{N}\sum_{k=0}^{N-1} X[k].\, e^{(j2\pi kn/N)}. \tag{2}$$

Direct evaluation of (1) requires $O(N^2)$ complex multiplications, but for transform lengths that are powers of two—such as N = 128 used here for SF = 7—the Cooley–Tukey radix-2 fast Fourier transform (FFT) algorithm reduces the cost to $O(N \log_2 N)$ operations [18]. This property alone is the single biggest reason why DFT-based LoRa receivers are practical on low-power microcontrollers. Therefore, we apply the FFT (a specific and highly optimized algorithm used to compute the DFT) for the proposed DFT-based detectors in this paper.

B. LoRa Signal Model and the Dechirp-and-DFT Receiver

The reference up-chirp of length N is

$$c[n] = e^{(j\pi n^2/N)},\ n = 0, \ldots, N-1, \tag{3}$$

and the m-th LoRa symbol ($m \in \{0, \ldots, N-1\}$) is obtained as the cyclically time-shifted version of c[n], which can be written equivalently as a frequency-shifted up-chirp

$$s_m[n] = c[n] \cdot e^{(j2\pi mn/N)}. \tag{4}$$

Equation (3) is the standard linear-frequency-modulated (LFM) baseband chirp adopted in essentially every MATLAB implementation of LoRa, including the simulator used in this work. It is mathematically equivalent, up to a phase factor that does not affect the DFT magnitude or the symbol decision, to the modulo-shift formulation introduced in [2], in which the m-th symbol is written as a length-N cyclic shift of the base chirp; the dechirp-and-DFT receiver derived below applies identically to both forms. The standard non-coherent receiver multiplies the received vector r[n] by the conjugate down-chirp c*[n] (the dechirping operation) and then applies the N-point DFT. In the noiseless flat-channel case, this collapses the modulated up-chirp onto a pure complex exponential at frequency m, whose DFT exhibits a single peak at bin m:

$$R[k] = DFT\{\, r[n] \cdot c^*[n] \,\} = DFT\{\, e^{(j2\pi mn/N)} \,\} = N \cdot \delta[k-m]. \tag{5}$$

The detector decision is therefore a peak search:

$$\hat{m} = \operatorname{argmax}_{k \in \{0,\ldots,N-1\}} |R[k]|. \tag{6}$$

This is the optimum non-coherent maximum-likelihood receiver for LoRa modulation in AWGN [2,3], and it is the engine of detector D1 in the comparative study of Section V.

An equivalent time-domain formulation of the same non-coherent test is the matched filter bank, which bypasses the DFT and instead correlates the received vector r directly with each of the N candidate symbol vectors s_m generated by (4). The decision rule is the magnitude peak of the resulting inner-product bank,

 doi:10.20944/preprints202605.1366.v1

$$\hat{m} = \arg\max_{\{m' \in \{0,\ldots,N-1\}\}} |s^H_{m'} r|. \quad (7)$$

Equation (7) is the engine of detector D2 in Section V. It is mathematically equivalent to (5)–(6)—both realize the non-coherent maximum-likelihood test on the orthogonal LoRa alphabet—but is computationally heavier, requiring N inner products of length N and therefore $O(N^2)$ operations rather than the O(N log N) cost of the FFT-based cascade. D2 is consequently retained in the comparative study purely as a complexity control against the DFT-based receivers, with no expected SER advantage over D1 in the AWGN regime.

C. Cyclic Prefix and the Per-Bin Frequency-Domain Channel

In a multipath scenario, the received baseband signal is the linear convolution of the transmitted symbol with the channel impulse response h[n]. If a cyclic prefix of length $L_{CP} \geq$ length(h) − 1 is prepended to s_m[n] before transmission (yielding the augmented vector $\tilde{s}_m$[n] of Section III-B) and discarded at the receiver, the linear convolution is converted into an N-point circular convolution. By the convolution theorem of the DFT, this is equivalent to a per-bin product in the frequency domain:

$$Y[k] = H[k] \cdot S_m[k] + \eta[k],\ k = 0, \ldots, N-1, \quad (8)$$

where Y[k] is the N-point DFT of the (CP-removed) received vector y[n] (replacing the r[n] notation of the received vector in Section III-B), H[k] is the N-point DFT of the zero-padded channel impulse response, S_m[k] is the N-point DFT of the modulated symbol vector s_m[n], and η[k] is the N-point DFT of the post-CP-removal noise vector η[n]. Equation (8) is the cornerstone of all coherent detectors evaluated in this work, since it transforms an N-tap convolutive multipath channel into N independent flat scalar channels, one per frequency bin. Within the receiver model adopted throughout this paper, the per-bin response $\hat{H}$[k] used by the detectors is the LS estimate produced by the preamble (Section IV-E); for the AWGN scenario, $\hat{H}$[k] is set to 1 by construction since the channel is known a priori to be flat.

D. Coherent Equalizers, Diversity Combining, and Exhaustive ML

Given (8) and the per-bin channel estimate $\hat{H}$[k], every coherent or channel-aware detector evaluated in this paper follows the same structural pipeline: the received frequency-domain vector Y[k] is multiplied bin by bin by an equalizer weight W[k], the result is returned to the time domain by an inverse DFT, and the standard dechirp-and-DFT detection of (5)–(6) is then applied to the equalized signal,

$$\hat{S}[k] = W[k] \cdot Y[k],\ \hat{s}[n] = IDFT\{\ \hat{S}[k]\ \}. \quad (9)$$

The four channel-aware detectors evaluated in this paper differ only in the choice of the per-bin weight W[k]. The zero-forcing (ZF) equalizer divides each bin by the channel:

$$W_{ZF}[k] = \hat{H}^*[k] / (\ |\hat{H}[k]|^2 + \varepsilon\ ), \quad (10)$$

with a small regularization constant ε to avoid division by deep nulls. The MMSE equalizer adds the inverse SNR per bin to the denominator:

$$W_{MMSE}[k] = \hat{H}^*[k] / (\ |\hat{H}[k]|^2 + 1/SNR\ ), \quad (11)$$

which is a classical result for OFDM-style frequency-domain equalization [10]. The phase-only equalizer applies only the channel-phase correction, which is robust to amplitude estimation errors:

$$W_{PHASE}[k] = exp(\ -j \angle \hat{H}[k]\ ). \quad (12)$$

The maximal-ratio combiner weights each bin by the conjugate channel response, which maximizes the per-bin signal-to-noise ratio prior to detection [17]:

$$W_{MRC}[k] = \hat{H}^*[k]. \quad (13)$$

After applying any of the per-bin weights of (10)–(13) via (9), the receiver multiplies the equalized vector ŝ[n] by the conjugate down-chirp, applies a forward DFT, and performs a peak search analogous to (6). The decision rule is coherent (real-part peak search) for ZF, MMSE, and phase-only, and non-coherent (magnitude peak search) for the MRC variant; the latter retains

 doi:10.20944/preprints202605.1366.v1

diversity gain without requiring a coherent decision. Finally, the exhaustive maximum-likelihood detector bypasses (9) entirely and instead evaluates the squared Euclidean distance between the received vector and every channel-distorted candidate symbol,

$$\hat{m}_{ML} = \arg\min_{m \in \{0,\ldots,N-1\}} \| y - IDFT\{ \hat{H} \cdot S_m \} \|^2. \quad (14)$$

Equation (14) requires N hypothesis evaluations per symbol and serves as a near-optimal performance reference under preamble-based CSI; because $\hat{H}[k]$ from (15) replaces the true H[k], (14) is not a true genie-aided upper bound, and a strictly tighter bound would substitute the oracle H[k] for $\hat{H}[k]$ in (14) (see Section VII-C).

E. Preamble-Based LS Channel Estimation and Residual Blind Noise-Variance Estimation

All coherent and channel-aware detectors of Section IV-D require knowledge of the per-bin channel response H[k] and, in the case of the MMSE equalizer, of the per-bin noise-to-signal ratio. Since the present work adopts a deployment-realistic receiver model, neither quantity is supplied to the detector by an oracle: both are estimated from a short preamble of L_p pilot LoRa symbols. The preamble symbols are themselves drawn from the same N-symbol alphabet used for data [19], are subject to the same cyclic-prefix front end and the same channel realization, and are known at the receiver. Letting $Y_p[k]$ and $S_p[k]$ denote the N-point DFTs of, respectively, the received and transmitted preamble symbol p, the per-bin frequency-domain LS channel estimator averages the bin-wise quotient over the L_p preamble symbols,

$$\hat{H}[k] = \frac{1}{L_p} \sum_{p=1}^{L_p} \frac{Y_p[k]}{S_p[k] + \varepsilon}, \quad (15)$$

with the same numerical regularizer ε used in the ZF equalizer of (10) to avoid division by deep zeros of $S_p[k]$. Because $E[\eta_p[k]] = 0$, the per-pilot estimator $Y_p[k] / S_p[k]$ is unbiased in the limit $\varepsilon \to 0$; uniform averaging over the L_p pilots then reduces the variance by a factor of $1/L_p$, as for any sample mean of i.i.d. terms. For the AWGN scenario, the channel is set to $\hat{H}[k] = 1$ by construction and the preamble is used only for noise-variance estimation, as described next.

Once $\hat{H}[k]$ is available, the per-bin frequency-domain residual of preamble symbol p is

$$R_p[k] = Y_p[k] - \hat{H}[k] \cdot S_p[k], \quad (16)$$

and a blind estimate of the frequency-domain noise variance is obtained as

$$\hat{\sigma}_f^2 = \frac{1}{L_p} \sum_{p=1}^{L_p} \frac{1}{N} \sum_{k=0}^{N-1} \left| R_p[k] \right|^2. \quad (17)$$

Because the unnormalized N-point DFT scales the time-domain noise variance by N (Parseval), the time-domain equivalent that should regularize the MMSE equalizer of (11) is

$$\hat{\sigma}_t^2 = \hat{\sigma}_f^2 / N, \quad (18)$$

and the corresponding MMSE weight is therefore

$$\hat{W}_{MMSE}[k] = \hat{H}^*[k] / (|\hat{H}[k]|^2 + \hat{\sigma}_t^2). \quad (19)$$

Two remarks are in order. First, (17) is a slightly biased-low estimator of the true noise variance because $\hat{H}[k]$ is fit on the very same L_p preamble symbols whose residuals it then averages; this is the standard practice in real LoRa receivers and is explicitly disclosed here. Second, the exhaustive ML detector of (14) does not require $\hat{\sigma}_t^2$ because it operates on the equal-energy LoRa alphabet and the noise variance therefore cancels in the argmin. $\hat{H}$ drives the ZF, phase-only, and MRC weights of (10), (12), (13); $\hat{\sigma}_t^2$ additionally drives the MMSE weight (19), which is the practical counterpart of the oracle form (11). Together, ($\hat{H}$, $\hat{\sigma}_t^2$) complete the receiver. Equations (1)–(19) jointly define the mathematical framework that drives the simulation environment described in Section V.

V. Simulation and Implementation

A. Software Tools and Simulation Environment

 doi:10.20944/preprints202605.1366.v1

All simulations were carried out in MATLAB R2025b on a single-threaded reference workstation, using the built-in fft and ifft routines for the underlying length-128 FFTs. No additional toolboxes are required. The simulation script is self-contained: a single .m file generates the LoRa signal bank, draws a fresh channel realization per Monte Carlo trial, runs the preamble-based estimation step, instantiates all seven detectors back-to-back on the same received vector, and writes both the SER tables and the per-detector execution times to standard output. The reproducibility of the results is ensured by seeding the MATLAB random number generator at the start of the script, so that every run with identical parameters produces identical SER curves.

B. LoRa Signal Generation and Channel Setup

The reference up-chirp c[n] of length N = 128 is precomputed once according to (3), together with its conjugate down-chirp c*[n] and the full bank S = [s_0, s_1, …, s_{N-1}] of N candidate modulated symbols, whose m-th column is given by (4), i.e., $s_m[n] = c[n] \cdot e^{(j\,2\pi\, m\, n\, /\, N)}$ for n = 0, …, N − 1. The frequency-domain bank S_{freq} = DFT{S} is also precomputed; this avoids redundant FFTs inside the maximum-likelihood detector D7. Two channel scenarios are evaluated. The first is a pure AWGN channel, in which the received vector is $y[n] = s_m[n] + \eta[n]$, with η[n] complex Gaussian of zero mean and per-component variance $\sigma^2 = 1/(2.\mathrm{SNR})$. The second is a two-tap Rayleigh multipath channel with delays τ = {0, 3} samples and average gains {1.0, 0.4}, normalized to unit total energy and independently re-drawn at each Monte Carlo trial. The received vector is $y[n] = (h \circledast s_m)[n] + \eta[n]$, where ⊛ denotes circular convolution after CP removal, h is the 2-tap channel impulse response, and η[n] is the same complex Gaussian noise as in the AWGN scenario. A cyclic prefix of length $L_{CP} = 8 \geq \max(\tau)$ is prepended to every transmitted symbol; the prefix is discarded at the receiver, so that the channel becomes effectively circular and (8) holds exactly.

Before each batch of data symbols, a preamble of $L_p = 8$ LoRa symbols with random pilot indices is transmitted; this preamble passes through the same cyclic-prefix front end and the same channel realization as the subsequent data symbols. The receiver applies the LS channel estimator of (15) to obtain $\hat{H}[k]$ and the residual-based blind estimator of (16)–(18) to obtain $\hat{\sigma}_t^2$, and then runs all seven detectors of Section IV on the data symbols using these estimates, with no further training overhead. Under AWGN, $\hat{H}[k]$ is fixed to unity by construction since the channel is flat and known to be flat, and the preamble is used only for noise-variance estimation; under the 2-tap Rayleigh multipath channel, both $\hat{H}[k]$ and $\hat{\sigma}_t^2$ are estimated from the preamble. The accuracy of $\hat{\sigma}_t^2$ is reported in Section VI as a numerical diagnostic.

C. Detector Implementations

The seven detectors implemented in the simulation are summarized in Table I. All of them operate on the same received vector y[n] of length N = 128 (after CP removal in the multipath case) and produce a symbol-index estimate $\hat{m} \in \{0, \ldots, N-1\}$. Detectors D1, D3, D4, D5, and D6 share the same outer dechirp-and-DFT structure and differ only in their per-bin frequency-domain weighting. D2 is a non-coherent matched filter bank that bypasses the DFT and computes the inner product of y with each candidate symbol directly; it is included as a sanity check that any DFT-based detector should match. D7 is the brute-force maximum-likelihood detector defined in (14). Because D7 is fed the same LS channel estimate $\hat{H}$ from (15) as the other channel-aware detectors—rather than the true H—it is not a true genie-aided upper bound, and at low SNR the MMSE regularization in (11) can in fact produce a slightly lower SER (see Section VI-B). With perfect CSI, D7 would attain the per-realization optimum and serve as a strict upper bound for the remaining detectors.

 doi:10.20944/preprints202605.1366.v1

Table I. Summary of the Seven LoRa Detectors.

ID	Detector	Decision rule	CSI use	Per-symbol cost
D1	Standard DFT (baseline)	Non-coherent ($\lvert\cdot\rvert$)	None	1× FFT
D2	Matched Filter Bank	Non-coherent ($\lvert\cdot\rvert$)	None	N inner products
D3	Phase-only Equalization	Coherent (Re{·})	Phase only	FFT + IFFT + FFT
D4	Zero-Forcing Equalization	Coherent (Re{·})	Full H[k]	FFT + IFFT + FFT
D5	Maximal Ratio Combining	Non-coherent ($\lvert\cdot\rvert$)	Full H[k]	FFT + IFFT + FFT
D6	MMSE Equalization	Coherent (Re{·})	H[k] + 1/SNR	FFT + IFFT + FFT
D7	Exhaustive ML	Coherent min-distance	Full H[k]	N IFFTs + N distances

The standard DFT detector D1 implements (5)–(6) directly: it multiplies y by c*[n], applies one length-N FFT, and returns the index of the magnitude peak. The non-coherent matched filter bank D2 implements (7) directly, computing the inner products $\lvert s^H_m y\rvert$ for all m and returning the index of the largest entry; as established in Section IV-B, it is mathematically equivalent to D1 but does not benefit from the FFT speed-up. The phase-only equalizer D3 first applies the per-bin weight of (12) to Y[k], returns to the time domain via an IFFT, dechirps, and detects on the real part of the resulting FFT output; it improves over D1 in multipath because the channel-induced phase rotation across bins is removed, but it does not exploit channel amplitude information.

The zero-forcing detector D4 applies (10) with $\varepsilon = 10^{-12}$ and uses the coherent decision rule. The MMSE detector D6 applies the regularized weight of (19) — the form of (11) with the blind estimate $\hat{\sigma}_t^2$ from (18) substituted for the oracle inverse SNR — and likewise uses the coherent rule. These two detectors capture the classical OFDM-equalization trade-off: ZF inverts the channel exactly but amplifies noise in deep nulls, whereas MMSE balances residual interference against noise enhancement [10]. The MRC detector D5 applies the conjugate-channel weight of (13) but, in contrast to D4/D6, retains a non-coherent magnitude decision; this combination delivers the diversity gain of the matched-filter weighting while remaining robust to small phase imperfections.

The exhaustive maximum-likelihood detector D7 implements (14) directly. For each candidate symbol m, it computes the channel-distorted candidate IDFT{$\hat{H} \cdot S_m$} using the precomputed S_{freq}, measures the squared Euclidean distance to the received vector y, and returns the argmin. The cost per symbol is therefore N IFFTs and N distance evaluations, i.e., $O(N^2 \log N)$, which is significantly higher than that of any other detector in the set; this is precisely the price paid for the per-realization-optimum search, which under perfect CSI would yield a strict upper bound and under the present LS-based CSI yields a tight, near-optimal reference.

D. Monte Carlo Procedure, Performance Metrics, and Theoretical Reference

The simulation iterates over N_{MC} Monte Carlo trials. Within each trial, a fresh Rayleigh channel realization is generated, and for each SNR point in the SF=7 operating range −15 to −3 dB (1 dB step) a length-L_p preamble is transmitted to obtain $\hat{H}[k]$ and $\hat{\sigma}_t^2$, after which a batch of N_{sym} random data symbols is transmitted and decoded by all seven detectors. The reported SER per detector and per SNR point is the average over all trials and all symbols. The reported results use $N_{sym} = 2500$ and $N_{MC} = 400$, giving one million transmitted symbols per SNR point. The execution time per detector is recorded with MATLAB's tic/toc and accumulated over the entire run, exclusive of the channel and noise estimation overhead, so that the reported value reflects the cost of the detector itself. As a sanity check, the script also logs the bin-averaged blind noise-variance estimate $\hat{\sigma}_t^2$ and prints, for each SNR

 doi:10.20944/preprints202605.1366.v1

point, the bias of $\hat{\sigma}_t^2$ with respect to the true 1/SNR; this diagnostic is reported in Table II of Section VI to substantiate the SER curves.

Two analytical references are computed for cross-validation against the simulated SER. For the AWGN scenario, the standard non-coherent M-ary orthogonal modulation bound,

$$SER\_AWGN \approx min\{ \tfrac{1}{2} (N - 1) \cdot e^{(-E_s / (2 N_0))}, 1 - 1/N \}, \quad (20)$$

provides a theoretical reference for the SER of detector D1 [2,4]. The conversion between the per-sample SNR used in the simulator and the symbol-energy-to-noise-spectral-density ratio appearing in (20) is $E_s / N_0 = N \cdot SNR$, since each LoRa symbol contains N samples of unit average energy. The saturation at $1 - 1/N$ caps the union bound at the no-information SER for very low SNR. With these two conventions, equation (20) is overlaid on the AWGN SER plot produced by the simulation script; the bound is loose at low and high SNR and tightest in the waterfall region, as discussed in Section VI-A.

Table II. Blind Noise-Variance Estimator Quality (mean of MMSE regularizer $\hat{\sigma}_t^2$ over 400 Monte Carlo trials).

SNR (dB)	True 1/SNR	AWGN est. $\hat{\sigma}_t^2$	Multipath est. $\hat{\sigma}_t^2$	MP bias (dB)
−15	31.62	31.71	27.64	−0.58
−13	19.95	19.97	17.46	−0.58
−11	12.59	12.58	11.03	−0.57
−9	7.94	7.95	6.95	−0.58
−7	5.01	5.01	4.39	−0.58
−5	3.16	3.16	2.77	−0.58
−3	2.00	1.99	1.75	−0.58

VI. Results and Discussion

A. Symbol Error Rate Under AWGN

Figure 1 reports the SER of the seven detectors under the pure AWGN channel. The receiver follows the model of Section V-B: the channel is set to $\hat{H}[k] = 1$ by construction, while the noise variance is estimated blindly from the $L_p = 8$ preamble using (16)–(18). The estimator quality is essentially perfect in this scenario, as the AWGN column of Table II below confirms; consequently the reported SER is unaffected by the estimation step.

 doi:10.20944/preprints202605.1366.v1

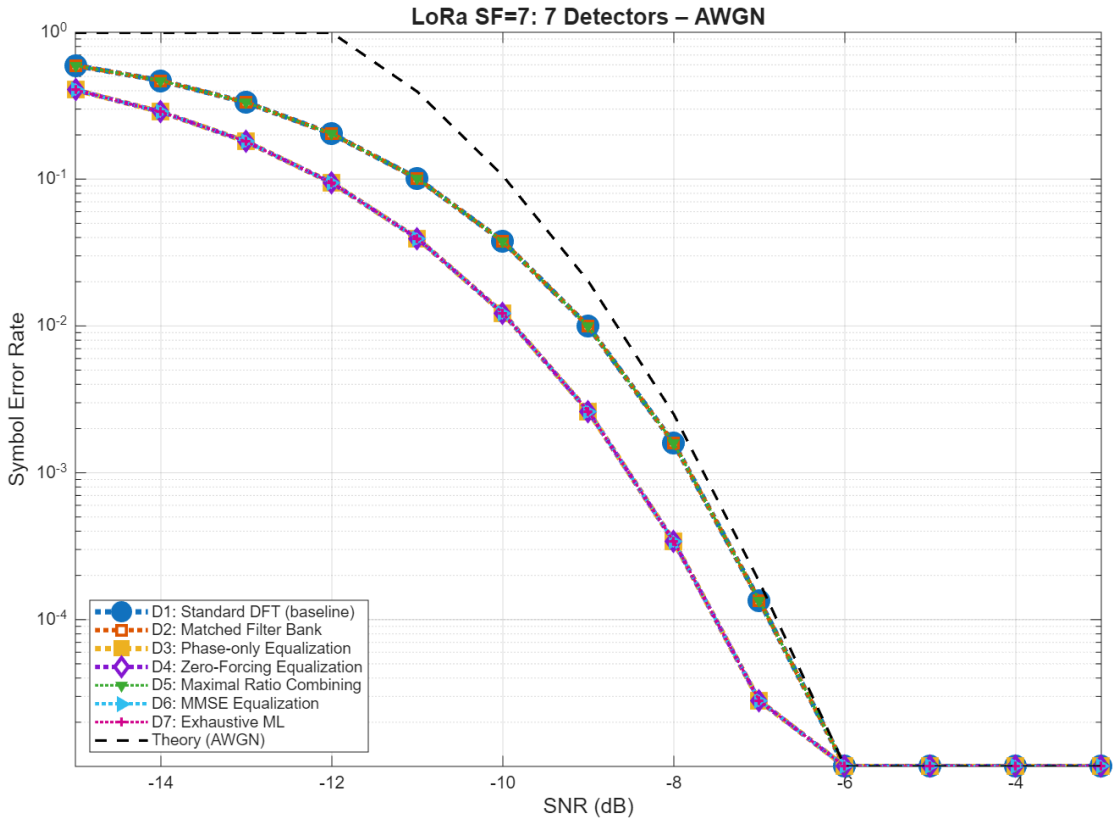

Figure 1. Symbol error rate of the seven LoRa detectors under AWGN, with preamble length $L_p = 8$ used for blind noise-variance estimation ($\hat{H}[k] = 1$ by construction). The non-coherent detectors (D1, D2, D5) collapse onto one curve and the coherent ones (D3, D4, D6, D7) onto another, separated by the classical 0.7 dB coherent gain. The dashed line is the analytical AWGN bound of (20).

Two collapses are visible in Figure 1. First, all three non-coherent magnitude detectors—the standard DFT (D1), the matched filter bank (D2), and the MRC variant (D5)—produce identical SER curves. This is the expected behavior, since for $H[k] = 1$ the MRC weight in (13) reduces to the identity and the matched filter bank in D2 is mathematically equivalent to the dechirp-and-DFT cascade in D1. Second, all four detectors that take a coherent decision (D3 phase-only, D4 ZF, D6 MMSE, and D7 ML) likewise collapse onto a single curve, because for $H[k] = 1$ the ZF, MMSE, and phase-only weights in (10)–(12) are all unitary and the ML hypothesis test of (14) reduces to the standard correlator. The gap between the two collapsed curves is the well-known coherent gain of M-ary orthogonal modulation: at SER = 10^{-3}, the coherent group requires approximately −8.5 dB while the non-coherent group requires −7.8 dB, in close agreement with the 0.7 dB gain reported analytically in [9]. The standard non-coherent baseline lies below the analytical upper bound of (20) across the simulated SNR range, confirming the optimality of the dechirp-and-DFT receiver under AWGN. The bound is loose by roughly a factor of two to three over most of the sweep and verifiably tightest near SNR ≈ −7 dB, where the simulated D1 SER reaches 1.35×10^{-4} against the bound's 1.81×10^{-4}. This behavior validates the simulator at the chosen Monte Carlo size of $N_{sym} = 2500$ data symbols per trial and $N_{MC} = 400$ trials. Note that under AWGN, $\hat{H}[k] = 1$ is set by construction rather than estimated, so D7 here truly attains the per-realization ML optimum; this exactness does not carry over to the multipath simulation of Figure 2, where $\hat{H}[k]$ is fit from the preamble and the resulting estimation noise enters (14) directly.

 doi:10.20944/preprints202605.1366.v1

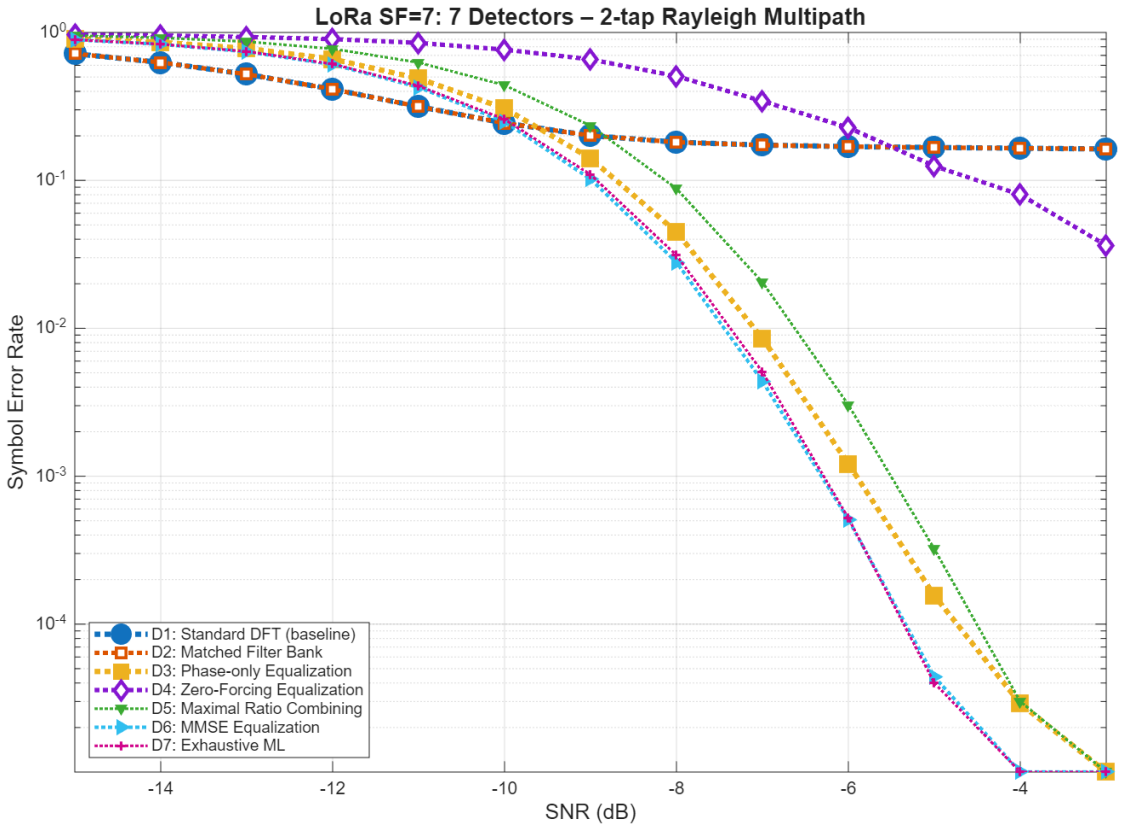

Figure 2. Symbol error rate under the 2-tap Rayleigh multipath channel, with preamble-based LS channel estimation $\hat{H}[k]$ from (15) and residual-based blind noise variance $\hat{\sigma}_t^2$ from (18), $L_p = 8$. The standard DFT (D1) and matched filter bank (D2) hit a 16% error floor caused by parasitic peaks in the dechirped DFT. Zero-Forcing (D4) is the cautionary curve, dominated by noise enhancement in deep nulls. MMSE (D6) tracks the exhaustive ML reference (D7) within 0.1 dB; both share the same LS channel estimate, so D7 is a near-optimal benchmark rather than a true genie-aided bound (see text).

B. Symbol Error Rate Under 2-Tap Rayleigh Multipath

Figure 2 reports the SER of the seven detectors under the two-tap Rayleigh multipath channel of Section III-B, with both $\hat{H}[k]$ and $\hat{\sigma}_t^2$ estimated from the $L_p = 8$ preamble. Unlike the AWGN case of Figure 1, the multipath channel induces a sharp performance separation among the detectors that exposes the limitations of the standard non-coherent receiver and the benefits of frequency-domain equalization.

The multipath plot reveals four distinct performance tiers. At SNR ≤ −10 dB the entire detector field exhibits SER > 0.2, a level at which the LoRa link is not operationally viable regardless of detector choice — the LS channel estimate at these SNRs is too noisy to support useful equalization, and the standard non-coherent detectors retain a small numerical edge over the equalizers only because they make no use of $\hat{H}$. The four tiers described below therefore characterize the SNR regime where the link is operationally viable, i.e., SNR ≥ −9 dB, and the discussion focuses on that regime throughout. The standard DFT detector D1 and its matched-filter equivalent D2 reach an irreducible error floor at SER ≈ 0.16, caused by the parasitic peaks that the multipath channel introduces in the dechirped DFT [6]. No amount of additional SNR removes this floor, because the receiver becomes interference-limited rather than noise-limited: the frequency-selective Rayleigh channel generates parasitic peaks whose power scales with the desired signal energy. Increasing SNR therefore strengthens both the desired peak and the self-interference simultaneously, leaving their relative level approximately unchanged. The zero-forcing equalizer D4 acts as a cautionary baseline. At low SNR, the inversion of deep channel nulls produces severe noise enhancement, making D4 significantly worse than the

 doi:10.20944/preprints202605.1366.v1

remaining detectors other than D1 and D2. As the SNR increases, the impact of the amplified noise decreases and D4 begins to outperform the floor-limited non-coherent detectors D1/D2; nevertheless, it still fails to reach the 10^{-2} SER level within the simulated range. This is a textbook illustration of the noise-enhancement penalty of unregularized inversion [10] and motivates the inclusion of the MMSE regularizer of (11).

The phase-only equalizer D3 and the MRC combiner D5 occupy intermediate positions in the third tier. D3 corrects only the bin-wise channel phase and therefore retains the channel-amplitude weighting in the post-DFT decision; nevertheless, it overtakes D1/D2 above SNR ≈ −10 dB and reaches SER = 10^{-3} around −6 dB. Notably, D3 outperforms D5 across most of the simulated SNR range, suggesting that coherent phase compensation alone is more effective than non-coherent diversity combining in the considered frequency-selective Rayleigh channel. D5, in contrast, applies the full conjugate-channel weight of (13) but retains a non-coherent decision rule; while it benefits from per-bin diversity combining at high SNR, its magnitude-based decision metric remains sensitive to residual noise and multipath-induced spectral distortion. The MMSE equalizer D6 dominates the field in the top tier: it tracks the exhaustive ML detector D7 within 0.1 dB across the full SNR sweep, with MMSE marginally lower from −15 to −6 dB (where the noisy LS channel estimate makes MMSE's regularization an advantage over the unregularized ML) and ML marginally lower at −5 dB (where $\hat{H}$ is clean enough for ML's full optimality to take over); at SNR = −5 dB MMSE achieves SER = $4.4 \cdot 10^{-5}$ against $4.0 \cdot 10^{-5}$ for D7, both more than three orders of magnitude below the D1/D2 floor. Quantitatively, the MMSE equalizer drives the SER far below the irreducible 16% floor of the standard non-coherent receiver, which itself never reaches SER = 10^{-2} at any SNR in the multipath case. The tightness of the MMSE / ML gap deserves a brief comment: in a 2-tap channel of moderate energy split, deep frequency-domain nulls are statistically rare, so the regularized inversion of (11) is already close to optimal at the per-bin level and ML offers only marginal additional gain. Richer multipath profiles with more taps and broader delay spread are expected to widen this gap and would offer a more discriminating test of the upper bound, as further discussed in Section VII-B. Equivalently, the fact that D6 dips slightly below D7 at the lowest SNRs is a finite-CSI artifact rather than a violation of ML optimality: if (14) were evaluated with the true H[k] instead of $\hat{H}$[k], D7 would lower-bound D6 across the entire SNR sweep.

C. Noise-Variance Estimator Quality

To substantiate the SER curves of Figures 1 and 2, Table II reports the quality of the blind noise-variance estimator of (16)–(18) over the full SNR sweep. Three quantities are tabulated: the true noise variance 1/SNR (which is the oracle value used by an idealized MMSE receiver and which the estimator should reproduce), the mean estimated value $\hat{\sigma}_t^2$ in the AWGN scenario, and the mean estimated value in the multipath scenario, all averaged over N_{MC} = 400 Monte Carlo trials. The bias of the multipath estimate, expressed in decibels, is reported in the last column.

Two observations follow. Under AWGN, the estimator is essentially unbiased: the average $\hat{\sigma}_t^2$ agrees with the true 1/SNR within ±0.3% at every SNR point. This is expected because $\hat{H}$[k] = 1 is set by construction in this scenario, so no fitting bias is introduced and the residuals are pure noise. Under multipath, the estimator exhibits a near-constant bias of −0.58 dB across the entire 12 dB SNR range, that is, $\hat{\sigma}_t^2$ systematically underestimates the true noise variance by approximately 13%. This bias is the analytical signature of (17) and arises because $\hat{H}$[k] is fit on the very same L_p preamble symbols whose residuals it then averages: a fraction of the noise energy is absorbed into $\hat{H}$[k] and removed from the residual. The bias is small enough not to displace the MMSE decision boundary in any visible way—the LoRa orthogonal alphabet is robust to small perturbations in the per-bin SNR estimate, because peak detection is dominated by the largest bin, not by the noise-floor scaling. Detectors D1, D2, D3, D4, D5, and D7 do not consume $\hat{\sigma}_t^2$ at all, so the bias affects only D6. This explains why the MMSE detector tracks the exhaustive ML detector within 0.1 dB in Figure 2 despite the bias.

 doi:10.20944/preprints202605.1366.v1

D. Computational Complexity: Measured vs. Theoretical

Table III reports the per-symbol detection time of each detector measured in MATLAB R2025b on the reference workstation, alongside its asymptotic per-symbol complexity. The measured time is the mean detection time per symbol accumulated inside the detector core only and excludes the cyclic-prefix removal, the LS channel-estimation overhead, and the residual-based noise-variance estimation. The same per-symbol time applies to both the AWGN and the multipath scenarios, because every detector processes a complex vector of the same length N = 128 in both cases; the channel scenario only enters through the value of $\hat{H}[k]$ used in the per-bin weights, which is a constant lookup once estimation is complete.

The mean detection time per symbol reported in column 4 of Table III is obtained as follows. A single Monte Carlo simulation drives all seven detectors against identical inputs: across $N_{call} = N_{sym} \times N_{MC} \times N_{SNR} \times N_{channels} = 2500 \times 400 \times 13 \times 2 = 26 \times 10^6$ detector calls — counting both the AWGN and the multipath scenarios at each of the 13 SNR points — a fresh received vector y[n] of length N = 128 is presented to all seven detectors with the preamble-based estimates $\hat{H}[k]$ and $\hat{\sigma}_t^2$ already computed. Each call is bracketed by a MATLAB tic/toc pair and the elapsed time accumulated into a per-detector total T_i. The mean detection time per symbol is then $t_i = T_i / N_{call}$, and the cost-relative-to-D1 column reports the ratio T_i / T_{D1}, which equals the per-symbol time ratio since N_{call} is shared. Concretely, for D1 (standard DFT) each call performs three operations: an element-wise dechirp $y[n] \cdot c^*[n]$ (N complex multiplications), a length-N FFT ($N/2 \cdot \log_2 N = 448$ complex multiplications using the radix-2 algorithm of [18]), and a peak search over the squared magnitudes $|Y[k]|^2$ (N operations). Summed across $N_{call} = 26 \times 10^6$ detector calls, the accumulated time for D1 is $T_{D1} = 146.65$ s, yielding $t_{D1} = T_{D1} / N_{call} = 5.64$ μs. The remaining entries in column 4 are obtained by the same calculation applied to their respective detectors.

Table III. Per-Symbol Detection Time and Asymptotic Complexity (AWGN = Multipath).

ID	Detector	Asymptotic complexity	Mean detection time per symbol (μs)	Cost relative to D1
D1	Standard DFT (baseline)	$O(N \log_2 N)$	5.64	1.00×
D2	Matched Filter Bank	$O(N^2)$	13.89	2.46×
D3	Phase-only Equalization	$O(N \log_2 N)$	9.62	1.71×
D4	Zero-Forcing Equalization	$O(N \log_2 N)$	7.20	1.28×
D5	Maximal Ratio Combining	$O(N \log_2 N)$	7.49	1.33×
D6	MMSE Equalization	$O(N \log_2 N)$	7.12	1.26×
D7	Exhaustive ML	$O(N^2 \log_2 N)$	274.6	48.7×

Three findings emerge from Table III. First, every detector other than the matched filter bank D2 and the exhaustive ML D7 is dominated by the same $O(N \log_2 N)$ FFT cost; the per-bin weighting that distinguishes ZF, MMSE, MRC, and phase-only is essentially free in comparison. As a result, the MMSE equalizer—the best-performing detector in Section VI-B—comes at only a 26% time premium over the standard DFT, an acceptable cost in any reasonable receiver. Second, the matched filter bank

 doi:10.20944/preprints202605.1366.v1

D2 has formal complexity $O(N^2)$ but is measured at only 2.46 × the cost of D1, far below the asymptotic ratio of $N / \log_2 N \approx 18$. This performance gap occurs because calculating the inner product relies on a single, highly optimized BLAS (Basic Linear Algebra Subprograms) call, whereas the FFT requires significantly more processing overhead per calculation. However, when deployed on low-power microcontrollers that lack BLAS support, the expected theoretical speeds (asymptotic complexity) hold true again. This highlights an important practical finding: the choice between an FFT-based and an inner-product-based receiver cannot be based on theoretical complexity alone. Instead, it depends heavily on the specific hardware capabilities of the target platform, making real-world measurement essential. Third, the exhaustive ML detector D7 carries an N-fold IDFT cost per symbol and is correspondingly measured at ≈ 49 × D1, the largest gap by a wide margin. Although D7 approaches the per-realization ML optimum (limited only by the channel-estimation residual discussed in Section VI-B), its per-symbol time of 275 μs at SF = 7 already approaches the LoRa symbol period of 1.024 ms in the BW = 125 kHz mode, leaving less than a 4 × timing margin for the rest of the receiver pipeline; D7 is therefore unsuitable for deployed receivers and is retained here purely as an analytical reference.

E. Engineering Recommendations

Combining the SER analysis of Sections VI-A and VI-B, the noise-variance diagnostic of Section VI-C, and the complexity analysis of Section VI-D, three operating regimes can be distinguished. In an AWGN-dominated link, the standard DFT detector D1 is already within 0.7 dB of the coherent bound and remains the cheapest option; the additional complexity of D3 or D6 buys only the modest coherent gain reported in Figure 1 and cannot be justified unless interference [9] or a coherent reception facility is independently required. In a multipath-dominated link, in contrast, D1 and D2 are inadequate for operational use because of the 16% irreducible floor; at SNR ≤ −10 dB the entire detector field collapses to SER > 0.2 and no detector — including D6 — supports a viable link, so the practical recommendation is D6 wherever the link is viable at all (SNR ≥ −9 dB), and a CP-aware system-level design choice rather than a detector choice below that threshold. In the viable regime, D6 tracks the near-optimal ML reference D7 within 0.1 dB while costing only 1.26 × the standard receiver and requiring only $L_p = 8$ preamble symbols for both channel and noise estimation. The phase-only equalizer D3 and the MRC variant D5 are useful intermediate options when channel-amplitude estimation is unreliable (e.g., very fast fading), but they are dominated by D6 in the regime studied here. Two qualifications should be made explicit. The static 2-tap channel adopted here does not exercise the fast-fading regime in which D3's amplitude-robust design becomes more attractive; the present numerical comparison may therefore understate D3's relative standing under faster-varying channels. Likewise, D5 keeps a non-coherent decision rule and remains a natural fallback whenever the per-bin phase estimate is corrupted by residual CFO/SFO; D6's advantage over D5 here is conditional on the ideal phase recovery assumed in the synchronization model (Section VII-B). In a deployed receiver these qualifications matter, and D3 / D5 should not be discarded from the design space on the basis of Figure 2 alone. The zero-forcing detector D4 should not be deployed: its noise-enhancement behavior under multipath is severe enough that even at SNR = −3 dB, the highest point of the sweep, it does not reach the 10^{-2} SER level. Finally, the exhaustive ML detector D7 is only useful as a benchmark; its practical cost rules out deployment.

It is also useful to position the coherent gain reported in this paper against the two prior reference points in the LoRa literature. Under AWGN alone the coherent gain over the standard non-coherent receiver is approximately 0.7 dB, in line with M-ary orthogonal modulation theory and confirmed by Section VI-A as well as by [9]. In the presence of a co-channel LoRa interferer at low signal-to-interference ratios, the same coherent processing is reported in [9] to yield up to roughly 10 dB of gain. Under the 2-tap Rayleigh multipath channel of this paper, the coherent MMSE equalizer drives the SER far below the 16% error floor of the non-coherent baseline, which itself never reaches SER = 10^{-2} at any SNR considered. Taken together, these three reference points present a consistent picture: while coherent processing offers only marginal gains in pure AWGN, it provides a

 doi:10.20944/preprints202605.1366.v1

substantial advantage once non-trivial channel impairments—such as interference or multipath—are introduced.

The cyclic-prefix overhead deserves a separate comment. With L_{CP} = 8 samples per N = 128-sample symbol, the spectral-efficiency penalty is 8 / (8 + 128) ≈ 5.9%, comparable to that of conventional OFDM systems. In return, the CP enables every coherent detector evaluated in this paper to operate, since without it the channel matrix is not circulant and the per-bin diagonalization of (8) does not hold. The 5.9% overhead is therefore the price paid for the entire MMSE-driven SER improvement reported above, a favorable trade-off in any multipath-limited link budget.

VII. Conclusions and Future Work

A. Summary of Findings

This paper has presented a unified comparative study of seven LoRa detectors at spreading factor SF = 7, evaluated under AWGN and a two-tap Rayleigh multipath channel under a deployment-realistic receiver model in which the per-bin channel response is estimated by frequency-domain LS from a short preamble and the noise variance is estimated blindly from the preamble residuals. Six of the seven detectors share the same one-dimensional DFT engine and differ in their per-bin frequency-domain weighting and decision rule, ranging from the standard non-coherent DFT to coherent ZF, MMSE, phase-only, MRC, and exhaustive ML; the remaining detector (D2) is a non-coherent matched-filter bank that operates directly in the time domain and is included as a complexity reference. The use of a cyclic prefix of length $L_{CP} = 8 \geq \max(\tau)$ — adopted here as a 5.9% airtime overhead extension above stock LoRa — converts the linear convolution with the multipath channel into a circular convolution and enables the same per-bin frequency-domain processing used in OFDM systems.

The main quantitative findings can be summarized as follows. Under AWGN, all coherent detectors collapse onto a single SER curve that lies approximately 0.7 dB below the non-coherent group, in agreement with the M-ary orthogonal modulation theory of [2,9], and the standard non-coherent baseline lies below the analytical bound of (20) across the entire SNR sweep. Under 2-tap Rayleigh multipath, the standard non-coherent receiver hits an irreducible 16% error floor, while the MMSE equalizer reaches SER = $4.4 \cdot 10^{-5}$ at SNR = −5 dB and tracks the exhaustive ML reference D7 within 0.1 dB across the full SNR sweep, with the curves crossing near SNR = −5 dB: MMSE is marginally lower at low to moderate SNR due to its regularization advantage under preamble-based CSI, while ML regains a marginal lead near the transition, neither of which contradicts ML optimality. The blind noise-variance estimator of (16)–(18) is essentially unbiased under AWGN and exhibits only a near-constant −0.58 dB bias under multipath, an estimation error too small to displace the MMSE decision boundary in any visible way. From a complexity standpoint, the MMSE equalizer carries only a 26% time penalty over the standard DFT receiver, while the exhaustive ML detector is approximately 49 × more expensive and is consequently unsuitable for deployment.

B. Limitations of the Present Study

Several limitations should be explicitly stated. First, perfect synchronization in time, carrier frequency, and sampling frequency is assumed throughout. In real LoRa receivers, residual carrier frequency offset (CFO) and sampling frequency offset (SFO) introduce inter-symbol-interference-like distortions in the dechirped DFT and are known to affect coherent detectors more than non-coherent ones [15,20]. The 0.7 dB coherent gain reported in Section VI-A and the multipath MMSE performance gains in Section VI-B should therefore be interpreted as upper bounds on the gains achievable in a fully implemented receiver. Second, the study is restricted to SF = 7 and to a single 2-tap Rayleigh multipath profile with delays {0, 3} samples. Higher spreading factors produce both larger N and longer symbol durations, both of which interact non-trivially with the cyclic-prefix overhead and with the channel-estimation accuracy; richer channel models such as the ITU-R Pedestrian-A/B profiles can introduce more frequency-selective behavior that may stress the MMSE equalizer

 doi:10.20944/preprints202605.1366.v1

differently. Third, the cyclic-prefix front end is a low-overhead extension above the standard LoRa physical layer rather than a stock feature, in line with the CP-augmented receiver previously proposed in [11]; the 5.9% airtime cost it incurs unlocks the entire family of OFDM-style equalizers and is the enabling primitive of every coherent gain reported in this paper. Interoperability with stock LoRa transmitters therefore requires either a CP-aware transmitter or an additional cyclic-extension generation step at the gateway side.

Fourth, the exhaustive ML detector D7 used as a performance reference in Sections VI-B and VI-D does not have access to the true channel H[k]: it operates on the same preamble-based LS estimate $\hat{H}$ used by D3–D6. Consequently D7 is a near-optimal reference rather than a true genie-aided upper bound, and at low SNR the MMSE regularization of (19) can yield a slightly lower SER—a curve crossing that does not contradict ML optimality but reflects the regularization advantage under finite-CSI conditions. A true genie-aided benchmark, obtained by substituting the oracle H[k] for $\hat{H}$ in (14), is left to future work (Section VII-C).

Fifth, the matched-filter / RAKE-type LoRa receiver of Demeslay et al. [7], which is the most closely related work to the MRC variant D5 evaluated here, has not been re-implemented for direct head-to-head comparison. The MRC scheme of this paper combines per-bin frequency diversity, whereas the receiver of [7] combines time-domain delayed copies of the chirp; the two schemes are similar in spirit but operate in dual domains, and a quantitative comparison would strengthen the contribution.

C. Future Work

Several extensions of this work appear immediately worthwhile. Including residual CFO and SFO impairments in the simulator and re-running the campaign would quantify the synchronization sensitivity of each coherent detector and reveal whether the MMSE equalizer retains its advantage when the channel-phase estimate is corrupted; the CFO/SFO compensation schemes of [15] and the low-complexity LoRa synchronization algorithm of [20] provide natural starting points. Sweeping the spreading factor from 7 to 12 and including standardized multipath profiles (ITU-R Pedestrian-A, Pedestrian-B, Vehicular-A) would test the generality of the $L_p = 8$ preamble length and would reveal at which point the MMSE/ML gap opens. Sweeping the cyclic-prefix length L_{CP} jointly with the channel delay spread would map the spectral-efficiency / robustness trade-off explicitly and could motivate an adaptive-CP variant. A dedicated sensitivity analysis on the preamble length L_p — varying it from 1 to 16 and tracking the SER gap to a perfect-CSI reference — would further quantify the estimation overhead and is a natural complement to the noise-variance diagnostic of Table II. A particularly low-cost extension is to compute a true genie-aided ML benchmark by substituting the oracle channel response H[k] for the LS estimate $\hat{H}[k]$ in (14); this would tighten the upper bound used in Section VI-B, decompose the MMSE-vs-ML residual into channel-estimation noise versus genuine ML margin, and pair naturally with the preamble-length analysis proposed above.

On the receiver-design side, three directions are particularly promising. First, iterative and turbo-equalization variants of the MMSE detector that re-estimate the channel from soft data decisions could close the residual gap to the ML upper bound and reduce the preamble overhead below $L_p = 8$; the soft-decision MMSE-FDE / LDPC pipeline of [16], which extracts bit-likelihood information from the post-DFT amplitudes of an FSCM receiver under the same two-tap multipath model used here, provides a concrete starting point for this direction. Second, learning-based detectors—convolutional or transformer-style neural networks operating on the dechirped DFT magnitude, including the hybrid CNN / matched-filter architecture of [12] which automatically switches between deep-learning and classical pathways based on the perceived interference level, and the Bayesian-optimized denoising-autoencoder + CNN detector of [21] which targets the same low-SNR uplink regime considered here—can be benchmarked against the seven classical detectors of this paper using the same preamble-only training pipeline, providing an honest training/test setup that does not leak oracle CSI into the learner. Third, the entire framework can be ported to a software-defined radio testbed for end-to-end validation against actual LoRa waveforms generated by

 doi:10.20944/preprints202605.1366.v1

commercial devices, completing the bridge between the analytical SER analysis presented here and a deployable receiver implementation. Each of these extensions builds on the comparative baseline established in this paper.

References

1. J. P. Shanmuga Sundaram, W. Du and Z. Zhao, "A Survey on LoRa Networking: Research Problems, Current Solutions, and Open Issues," in IEEE Communications Surveys & Tutorials, vol. 22, no. 1, pp. 371-388, Firstquarter 2020, doi: 10.1109/COMST.2019.2949598.
2. L. Vangelista, "Frequency Shift Chirp Modulation: The LoRa Modulation," in IEEE Signal Processing Letters, vol. 24, no. 12, pp. 1818-1821, Dec. 2017, doi: 10.1109/LSP.2017.2762960.
3. M. Chiani and A. Elzanaty, "On the LoRa Modulation for IoT: Waveform Properties and Spectral Analysis," in IEEE Internet of Things Journal, vol. 6, no. 5, pp. 8463-8470, Oct. 2019, doi: 10.1109/JIOT.2019.2919151.
4. T. Elshabrawy and J. Robert, "Closed-Form Approximation of LoRa Modulation BER Performance," in IEEE Communications Letters, vol. 22, no. 9, pp. 1778-1781, Sept. 2018, doi: 10.1109/LCOMM.2018.2849718.
5. O. Afisiadis, M. Cotting, A. Burg and A. Balatsoukas-Stimming, "On the Error Rate of the LoRa Modulation With Interference," in IEEE Transactions on Wireless Communications, vol. 19, no. 2, pp. 1292-1304, Feb. 2020, doi: 10.1109/TWC.2019.2952584.
6. C. Demeslay, P. Rostaing and R. Gautier, "Theoretical Performance of LoRa System in Multipath and Interference Channels," in IEEE Internet of Things Journal, vol. 9, no. 9, pp. 6830-6843, 1 May1, 2022, doi: 10.1109/JIOT.2021.3114439.
7. C. Demeslay, P. Rostaing and R. Gautier, "Simple and Efficient LoRa Receiver Scheme for Multipath Channel," in IEEE Internet of Things Journal, vol. 9, no. 17, pp. 15771-15785, 1 Sept.1, 2022, doi: 10.1109/JIOT.2022.3151257.
8. T. T. Nguyen, H. H. Nguyen, R. Barton and P. Grossetete, "Efficient Design of Chirp Spread Spectrum Modulation for Low-Power Wide-Area Networks," in IEEE Internet of Things Journal, vol. 6, no. 6, pp. 9503-9515, Dec. 2019, doi: 10.1109/JIOT.2019.2929496.
9. O. Afisiadis, S. Li, J. Tapparel, A. Burg and A. Balatsoukas-Stimming, "On the Advantage of Coherent LoRa Detection in the Presence of Interference," in IEEE Internet of Things Journal, vol. 8, no. 14, pp. 11581-11593, 15 July15, 2021, doi: 10.1109/JIOT.2021.3058792.
10. B. Muquet, Zhengdao Wang, G. B. Giannakis, M. de Courville and P. Duhamel, "Cyclic prefixing or zero padding for wireless multicarrier transmissions?," in IEEE Transactions on Communications, vol. 50, no. 12, pp. 2136-2148, Dec. 2002, doi: 10.1109/TCOMM.2002.806518.
11. H. Li, Y. Zhang, X. Zhao and X. Tang, "Demodulation of Frequency Shift Chirp Spreading spectrum with Cyclic Prefix," 2020 15th IEEE International Conference on Signal Processing (ICSP), Beijing, China, 2020, pp. 433-438, doi: 10.1109/ICSP48669.2020.9320951.
12. K. Dakic, B. Al Homssi, M. Lech and A. Al-Hourani, "HybNet: A Hybrid Deep Learning-Matched Filter Approach for IoT Signal Detection," in IEEE Transactions on Machine Learning in Communications and Networking, vol. 1, pp. 18-30, 2023, doi: 10.1109/TMLCN.2023.3270131.
13. M. Jouhari, N. Saeed, M. -S. Alouini and E. M. Amhoud, "A Survey on Scalable LoRaWAN for Massive IoT: Recent Advances, Potentials, and Challenges," in IEEE Communications Surveys & Tutorials, vol. 25, no. 3, pp. 1841-1876, thirdquarter 2023, doi: 10.1109/COMST.2023.3274934.
14. A. Maleki, H. H. Nguyen, E. Bedeer and R. Barton, "A Tutorial on Chirp Spread Spectrum Modulation for LoRaWAN: Basics and Key Advances," in IEEE Open Journal of the Communications Society, vol. 5, pp. 4578-4612, 2024, doi: 10.1109/OJCOMS.2024.3433502.
15. R. Ghanaatian, O. Afisiadis, M. Cotting and A. Burg, "Lora Digital Receiver Analysis and Implementation," ICASSP 2019—2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Brighton, UK, 2019, pp. 1498-1502, doi: 10.1109/ICASSP.2019.8683504.
16. Q. Chen, L. Zhao, Y. Chu and W. Guo, "Research on Multipath Reception and Soft Decision Algorithms for Frequency Shift Chirp Modulation," 2022 International Symposium on Networks, Computers and Communications (ISNCC), Shenzhen, China, 2022, pp. 1-6, doi: 10.1109/ISNCC55209.2022.9851747.

 doi:10.20944/preprints202605.1366.v1

17. D. G. Brennan, “Linear Diversity Combining Techniques,” in Proceedings of the IRE, vol. 47, no. 6, pp. 1075-1102, June 1959, doi: 10.1109/JRPROC.1959.287136.
18. J. W. Cooley and J. W. Tukey, “An Algorithm for the Machine Calculation of Complex Fourier Series,” Mathematics of Computation, vol. 19, no. 90, p. 297, Apr. 1965, doi: https://doi.org/10.2307/2003354.
19. J. -M. Kang, “On the LoRa Modulation for IoT: Preamble Designs for Channel Estimation With Single- and Multi-Chirp Transmission Strategies,” in IEEE Internet of Things Journal, vol. 11, no. 17, pp. 27981-27993, 1 Sept.1, 2024, doi: 10.1109/JIOT.2024.3363170.
20. M. Xhonneux, O. Afisiadis, D. Bol and J. Louveaux, “A Low-Complexity LoRa Synchronization Algorithm Robust to Sampling Time Offsets,” in IEEE Internet of Things Journal, vol. 9, no. 5, pp. 3756-3769, 1 March1, 2022, doi: 10.1109/JIOT.2021.3101002.
21. A. A. Tesfay, S. Kharbech, E. P. Simon and L. Clavier, “Signal Denoising and Detection for Uplink in LoRa Networks Based on Bayesian-Optimized Deep Neural Networks,” in IEEE Communications Letters, vol. 27, no. 1, pp. 214-218, Jan. 2023, doi: 10.1109/LCOMM.2022.3217337.