Available online at UofKEJ Website Here
UofKEJ Vol. 14 Issue 1 pp. 40-48 (2026)

UNIVERSITY OF KHARTOUM ENGINEERING JOURNAL (UofKEJ)

Spreading Factor-Dependent Carrier Frequency Offset Tolerance in LoRa: Empirical Characterisation and Oscillator Selection for Adaptive Data Rate Networks

Osemwegie Sam Omorogiuwa[1*], Imonje Godswill Oseiwe[1], and Oduware Okosun[1]

[1]Department of Electrical Engineering, Faculty of Engineering, University of Benin, Nigeria.
Corresponding author (E-mail: sam@uniben.edu)

ARTICLE INFO

Keywords
LoRa, carrier frequency offset, spreading factor, Adaptive Data Rate, crystal oscillator, software-defined radio, implementation headroom

Article History:
Received on: December 25, 2025
Accepted on : February 23, 2025.

Article Type: Research Article

DOI: 10.53332/uofkej.v14i1.330

ABSTRACT

Carrier frequency offset (CFO) tolerance in Long-Range (LoRa) modulation decreases with increasing spreading factor because longer symbol durations accumulate greater phase rotation per unit frequency error. This paper presents a systematic empirical characterization of CFO tolerance across the LoRa spreading factor range under controlled high-margin indoor conditions using a software-defined radio receiver. Five spreading factors (SF7–SF10 and SF12) are directly measured; SF11 is estimated by log-linear interpolation owing to a software decoder limitation. Measured tolerance ranges from more than 20 kHz at SF7 to 0.5 kHz at SF12—a 40-fold variation—and exhibits two distinct behavioral regimes: a tracking-limited plateau at SF7–SF8 and a coherence-failure regime at SF9–SF12 described by an exponential decay model. Measured tolerances exceed Semtech theoretical predictions by factors of 66–189, attributed to receiver-side implementation techniques including fractional frequency estimation and soft-decision decoding. These measured margins are then compared with the frequency drift of standard ±10 ppm crystal oscillators (8.68 kHz at 868 MHz) to identify a graduated vulnerability threshold: SF7–SF9 operates with adequate margin, SF10 lies in a marginal zone where forward error correction compensated for physical-layer errors under the controlled conditions of this study, and SF11–SF12 operates below the commodity oscillator drift limit. Because Adaptive Data Rate (ADR) algorithms may increase the spreading factor dynamically, oscillator adequacy at low spreading factors does not guarantee adequacy across the full ADR range. An evidence-based oscillator selection framework is derived from these results. SF11 CFO tolerance is reported as an estimate obtained by log-linear interpolation, clearly labelled throughout, owing to a decoder limitation of the software demodulator used.

1. INTRODUCTION

Long-Range (LoRa) modulation has become a dominant physical-layer technology for Low-Power Wide-Area Networks (LPWANs), enabling communication over ranges exceeding 10 km with battery lifetimes measured in years [1]. These properties make LoRa attractive for large-scale sensor deployments in smart agriculture, industrial monitoring, and urban sensing applications.

A practical constraint in any radio system is the tolerance to carrier frequency offset (CFO)—the difference $\Delta f = f_{\mathrm{tx}} - f_{\mathrm{rx}}$. between transmitter and receiver centre frequencies. In LoRa, Δf arises primarily from crystal oscillator inaccuracy. At 868 MHz, a commodity ±10 ppm crystal introduces a worst-case drift of ±8.68 kHz, while a ±5 ppm mid-grade crystal contributes ±4.34 kHz, and a temperature-compensated oscillator (TCXO) at ±2 ppm contributes ±1.74 kHz. For deployments of 10^5 sensors, the unit cost difference between commodity and TCXO components is significant; yet choosing an inadequate oscillator risks systematic link failure.

The LoRa spreading factor (SF) parameter controls the trade-off between data rate and receiver sensitivity. The symbol duration T_s scales as $T_s = 2^{SF}/BW$, where BW denotes the signal bandwidth (125 kHz in this work). A CFO of Δf induces a phase rotation of $\Delta\phi = 2\pi\Delta f T_s$ radians per symbol. Because T_s increases by a factor of $2^5 = 32$ from SF7 to SF12, the same physical frequency error produces 32 times more phase rotation at SF12 than at SF7. Theoretical analysis by Semtech [2] predicts that CFO tolerance falls as $CFO_{max} \approx 0.25 \times BW/2^{SF}$, yielding approximately 244 Hz at SF7 and 7.6 Hz at SF12. Empirical studies of CFO tolerance, however, consistently report tolerances substantially higher than Semtech's predictions [3], indicating headroom introduced by receiver implementation techniques not captured by the theoretical model.

Adaptive Data Rate (ADR) is a network-layer mechanism that adjusts the spreading factor in response to observed link quality [4]. A node operating satisfactorily at SF7 with a commodity crystal may be instructed by an ADR server to transmit at SF10 or higher during periods of poor reception. If the crystal's frequency drift exceeds the CFO tolerance at the assigned spreading factor, systematic packet loss would be expected—a failure mode not explicitly addressed in ADR protocol design [5]. This interaction is an inferred engineering implication; an end-to-end ADR control loop was not experimentally evaluated in this study.

Despite extensive LoRa research, a systematic characterisation of CFO tolerance spanning all six spreading factors under a uniform methodology has not been reported. The sole prior empirical study examined SF7 alone [6], leaving intermediate and high spreading factors entirely uncharacterised and the ADR-oscillator interaction unquantified. This paper addresses that gap.

1.1 Research Problem

ADR mechanisms may increase the spreading factor to improve coverage under degraded link conditions. CFO tolerance, however, decreases with increasing spreading factor. No prior study has quantified whether standard commodity oscillators remain adequate when ADR increases the spreading factor, nor has a systematic transition zone been identified where performance shifts from adequate to unreliable.

1.2 Research Objective

This study characterises CFO tolerance empirically at SF7–SF10 and SF12 under controlled high-margin conditions, estimates SF11 tolerance by interpolation, determines the spreading factor at which commodity oscillator drift exceeds measured tolerance, and develops a practical oscillator selection framework for ADR-enabled deployments.

1.3 Contributions

This work makes the following contributions to the LoRa literature:

C1. Direct measurement of CFO tolerance at SF7–SF10 and SF12 under a uniform experimental methodology, with SF11 estimated by validated log-linear interpolation, filling the gap left by studies limited to endpoint spreading factors.

C2. Identification of a two-regime behavioural model: a tracking-limited plateau at SF7–SF8 and a coherence-failure regime at SF9–SF12 following an exponential decay.

C3. Quantification of a graduated vulnerability threshold where commodity ±10 ppm oscillators transition from adequate (SF7–SF9) through marginal (SF10) to insufficient (SF11–SF12).

C4. Quantification of implementation headroom—measured tolerances exceed Semtech predictions by 66–189×—and attribution to receiver-side implementation techniques.

C5. An oscillator selection framework for ADR-enabled networks with deployment-scenario guidance.

2. BACKGROUND AND RELATED WORK

2.1 LoRa Chirp Spread Spectrum and CFO Sensitivity

LoRa modulation encodes data as linear frequency sweeps (chirps) across a bandwidth BW. Each symbol spans 2^{SF} chips, so the symbol duration is $T_s = 2^{SF}/BW$. Table 1 defines the notation used throughout this paper.

A carrier frequency offset Δf shifts the received chirp in the frequency domain. The demodulator correlates the received signal against a reference chirp; a frequency-shifted received chirp decorrelates from the reference as the phase error $\Delta\phi = 2\pi\Delta f T_s$ increases. When $\Delta\phi$ approaches π, inter-symbol interference arises and the demodulator begins to fail. This mechanism predicts that CFO tolerance scales inversely with T_s, and therefore inversely with 2^{SF}—the fundamental relationship motivating this investigation.

Table 1: Notation and abbreviations used in this paper.

Symbol / Term	Definition
Δf (CFO)	Carrier frequency offset: $f$tx−$f$rx (Hz)
BW	Signal bandwidth; 125 kHz in this study
ϵ	Normalized CFO: Δf/BW
Ts	Symbol duration: 2SF/BW (s)
$\Delta\phi$	Phase error per symbol: $2\pi\Delta fTs$ (rad)
SF	Spreading factor; {7, 8, 9, 10, 11, 12}
CFO_{max}	Maximum CFO satisfying stated criterion
BER	Bit error rate
PDRsweep	Packet delivery ratio across full 0–20 kHz sweep
PDR≤CFO-max	Packet delivery ratio at offsets ≤ CFOmax only
M	Link margin (dB): received power minus receiver sensitivity
FSPL	Free-space path loss (dB)
ADR	Adaptive Data Rate
FEC	Forward error correction
LPWAN	Low-Power Wide-Area Network
TCXO	Temperature-compensated crystal oscillator
a, b	Exponential model parameters: CFOmax=a e−b SF
$R2$	Coefficient of determination
ZMQ	ZeroMQ inter-process messaging library
SDR	Software-defined radio

2.2 Signal Model and CFO Normalisation

The normalised CFO is defined as

$$\epsilon = \frac{\Delta f}{\mathrm{BW}}, \tag{1}$$

so that $\epsilon = 1$ corresponds to an offset equal to the full signal bandwidth. Over one symbol period, the accumulated phase error is

$$\Delta\phi = 2\pi\,\Delta f\,T_s = 2\pi\,\epsilon\,2^{\mathrm{SF}} \tag{2}$$

which grows linearly with Δf and exponentially with SF. The Semtech theoretical model [2] treats demodulation failure as occurring when $\Delta\phi$ exceeds a fixed threshold, yielding

$$\mathrm{CFO}_{\mathrm{max}}^{\mathrm{theory}}(\mathrm{SF}) = \frac{0.25\,\mathrm{BW}}{2^{\mathrm{SF}}} \tag{3}$$

This model assumes ideal demodulation and does not account for fractional frequency estimation, soft-decision decoding, or oversampling, all of which extend practical tolerance beyond Eq. 3.

For the coherence-failure regime observed empirically (Section 4), we fit a two-parameter exponential decay model

$$\mathrm{CFO}_{\mathrm{max}}(\mathrm{SF}) = a\,e^{-b\,\mathrm{SF}} \tag{4}$$

where a and b are determined by least-squares regression over the measured data points, and R^2 quantifies goodness of fit.

2.3 Prior Work

Table 2 places this study in the context of prior CFO characterisations. Semtech [2] and Chiani and Elzanaty [7] established the theoretical framework but assumed ideal receiver operation. Croce et al. [3] and To and Duda [8] used simulation to predict higher tolerance than theory, but simulations cannot capture proprietary chipset enhancements. Van den Abeele et al. [6] provide the only prior empirical CFO measurement, limited to SF7; no empirical data for SF8 through SF12 has previously been published. Polonelli et al. [9] examined LoRaWAN MAC-layer behaviour and provide deployment context but did not report CFO tolerance measurements across spreading factors.

ADR capacity optimisation has been studied by Rahmadhani and Kuipers [5] and Reynders et al. [4], but neither work examined oscillator requirements as a function of the assigned spreading factor. The interaction between ADR-driven SF assignment and crystal oscillator adequacy has not previously been reported.

Table 2: Comparison of prior CFO tolerance studies. †SF7–SF10 and SF12 measured; SF11 estimated by interpolation.

Reference	**Method**	**SFs**	**Gap**
Van den Abeele et al. [6]	Empirical	SF7	SF8–SF12 absent
Polonelli et al. [9]	MAC context	—	No CFO measurement
Croce et al. [3]	Simulation	SF7, SF12	No hardware validation
This study	Empirical	SF7–SF12†	Complete coverage

3. METHODOLOGY

3.1 Hardware and Configuration

The transmitter was a LILYGO LoRa32 T3 V1.6.1 development board (ESP32 microcontroller with SX1276 transceiver), configured to transmit at +5 dBm on an 868 MHz centre frequency. Arduino firmware swept the transmit frequency from 868.000 MHz to 868.020 MHz in 500 Hz increments, transmitting a 64-byte ground-truth payload at each offset with a 3-second inter-packet delay, yielding 41 packets per spreading factor. The receiver was an RTL-SDR V3 dongle running GNU Radio 3.10.12.0 with the gr-lora 0.5.7 demodulator [10], configured for 125 kHz bandwidth, coding rate 4/5, explicit header mode, CRC enabled, and soft-decision decoding. Both ends used 5 dBi dipole antennas at a height of 1 m. Link geometry was 5.69 m line-of-sight, indoors.

3.2 Link Margin

The quantity reported as "94.3 dB" is a *link margin relative to receiver sensitivity*, not a directly measured in-band SNR. It is computed as

$$M = P_{\text{tx}} + G_{\text{tx}} + G_{\text{rx}} - \text{FSPL} - L_{\text{conn}} - S_{\text{min}} \quad (5)$$

where $P_{\text{tx}} = +5$ dBm, $G_{\text{tx}} = G_{\text{rx}} = 5$ dBi, free-space path loss at 5.69 m and 868 MHz is

$$\text{FSPL} = 20\log_{10}(d) + 20\log_{10}(f) + 32.45 = 40.2\,\text{dB} \quad (6)$$

connector and implementation losses $L_{\text{conn}} \approx$ 0.8 dB, and SF7 sensitivity $S_{\text{min}} = -123$ dBm [2], yielding $M = 94.3$ dB. This large margin—substantially above the minimum required for demodulation—was intentional: by operating well above the sensitivity floor, thermal noise effects on BER are minimised, allowing CFO to be isolated as the primary degradation variable. Actual received power, measured at −28.7 dBm, is consistent with the free-space path loss prediction.

3.3 Measurement Procedure and Performance Metrics

A Python script received demodulated payloads from GNU Radio via ZeroMQ (ZMQ)—a high-performance asynchronous messaging library used as a software transport between the GNU Radio flowgraph and the analysis process—and performed bit-wise XOR comparison against the known 64-byte ground-truth payload. Two complementary performance criteria are used throughout this paper:

Physical-layer reliability: $\text{CFO}_{\text{max}}^{(\text{BER})}$ is the highest tested offset at which the cumulative BER remains below 10^{-3}. This criterion captures the physical-layer demodulation limit and is used as the primary metric for vulnerability classification and oscillator margin calculations.

Application-layer viability: PDR is the fraction of the 41 transmitted packets received without error. PDR reflects the combined effect of physical-layer errors and forward error correction (FEC) at coding rate 4/5, and is reported as a supplementary indicator. Tables and figures explicitly state which criterion applies.

.

3.4 SF11 Estimation Methodology

At SF11, the gr-lora v0.5.7 decoder exhibited approximately 45% BER at zero offset—a known decoder defect unrelated to CFO [10]. This precluded direct CFO tolerance measurement at SF11. An estimate was obtained by fitting the exponential model of Eq. 4 to the three measured Regime 2 data points (SF9, SF10, SF12) and evaluating the fitted curve at SF11. All SF11 values are labelled "(est.)" in tables and figures throughout this paper.

A sensitivity analysis was conducted by varying the SF11 estimate by ±20%. The central estimate is 1.55 kHz (0.18× the ±10 ppm drift limit of 8.68 kHz). At the lower bound of 1.24 kHz (0.14×) and the upper bound of 1.86 kHz (0.21×), SF11 remains well below the commodity oscillator drift limit in all cases. The qualitative vulnerability classification—SF11 in the failure zone—is therefore unchanged by estimation uncertainty.

4. RESULTS

4.1 Overview

Table 3 summarises CFO tolerance, PDR, and oscillator margin for all spreading factors. Figure 1 and Figure 2 show the BER versus CFO curves and the CFO_{max} trend, respectively.

4.2 SF7: Tracking-Limited Regime

All 41 packets transmitted at SF7 were received without error across the complete 0–20 kHz offset sweep, yielding $\text{CFO}_{\text{max}} > 20$ kHz (> 23.0 ppm). Zero bit errors were observed at any tested offset, giving $\text{PDR}_{\text{sweep}}$ = 100% (95% CI: 91.6%–100%, Clopper–Pearson exact). SF7 operates well within the tracking bandwidth of the gr-lora demodulator;

the measured limit reflects the tracking range rather than symbol coherence, consistent with Regime 1 behaviour. The oscillator margin relative to the ±10 ppm drift limit is >2.30×.

4.3 SF8: Transition into Degradation

At SF8, no bit errors were observed for offsets up to 15.5 kHz (31 of 41 packets), with the first error appearing at 16.0 kHz. $\mathrm{CFO}_{\max}^{(\mathrm{BER})} = 19.0$ kHz (21.9 ppm), and the per-packet BER at the maximum tested offset of 20 kHz was 8/512 = 0.016. Implementation headroom over the Semtech prediction is 156×. The oscillator margin is 2.19×, indicating safe operation for commodity oscillators. SF8 marks the lower boundary of the tracking-limited plateau.

4.4 SF9: Entry into the Coherence-Failure Rgime

At SF9, error onset occurred at 9.0 kHz (packet 19 of 41), with $\mathrm{CFO}_{\max}^{(\mathrm{BER})} = 11.5$ kHz (13.2 ppm). The per-packet BER reached 27/512 = 0.053 at 20 kHz. The oscillator margin is 1.32×—still above unity, but approaching the boundary of adequate operation. SF9 is the first spreading factor at which the gap between measured tolerance and commodity oscillator drift is narrow enough to warrant attention in production deployments exposed to temperature cycling or component aging.

4.5 SF10: Marginal Zone

At SF10, the first error onset was at 4.5 kHz (packet 10), and $\mathrm{CFO}_{\max}^{(\mathrm{BER})} = 5.5$ kHz (6.3 ppm). The per-packet BER at 20 kHz was 61/512 = 0.119, which approaches the theoretical correction limit of convolutional FEC at rate 4/5 (≈10% BER). Despite this, all 41 packets were delivered ($\mathrm{PDR}_{\leq \mathrm{CFOmax}}$ = 100%, 95% CI: 91.6%–100%), because FEC successfully corrected the accumulated bit errors under the controlled conditions of this experiment. The calculated oscillator margin is 0.63×—below unity, indicating that the ±10 ppm drift budget exceeds the BER-based tolerance threshold.

This apparent contradiction—sub-unity BER margin yet 100% PDR—is explained by three factors specific to the laboratory conditions of this study: (i) FEC at rate 4/5 can recover from per-packet BERs up to approximately 10%, and the observed SF10 BER at 20 kHz approached but did not exceed this limit; (ii) the controlled indoor temperature (≈25°C) eliminated thermal drift, effectively reducing the actual oscillator contribution below the ±10 ppm worst-case figure; and (iii) the 94.3 dB link margin ensured that thermal noise did not compound the CFO degradation. Production deployments typically operate at 10–20 dB lower link margins with greater thermal variation. SF10 is therefore characterised as a marginal zone: functional under favourable conditions but lacking the robustness required for high-availability deployments.

4.6 SF11: Estimated Tolerance

As described in Section 3.4, SF11 CFO tolerance could not be measured directly owing to a gr-lora decoder defect. The log-linear interpolation from the Regime 2 fitted model yields an estimated $\mathrm{CFO}_{\max}^{(\mathrm{est.})}(\mathrm{SF11}) = 1.55$ kHz (1.78 ppm), representing 0.18× the commodity oscillator drift limit. Even at the upper sensitivity bound of 1.86 kHz, the margin remains 0.21×—substantially below the threshold for adequate operation. SF11 is classified as a failure zone for commodity oscillators under this study's conditions.

4.7 SF12: Observed Packet Loss

At SF12, bit errors appeared at the first non-zero offset tested (0.5 kHz, packet 2). $\mathrm{CFO}_{\max}^{(\mathrm{BER})} = 0.5$ kHz (0.58 ppm), and the per-packet BER at 20 kHz reached 94/512 = 0.184. Of 39 transmitted packets, only 1 was received without error ($\mathrm{PDR}_{\mathrm{sweep}}$ = 2.6%, 95% CI: 0.1%–13.5%), providing direct experimental evidence of link failure. Implementation headroom over Semtech theory is 66×. The calculated oscillator margin is 0.06×, indicating that the commodity crystal frequency drift exceeds the tolerance by approximately 17-fold.

4.8 Two-Regime Behavioural Model

The combined dataset reveals two distinct behavioural regimes (Figure 2):

Regime 1 (SF7–SF8): CFO tolerance remains high (>19 kHz), limited by the tracking bandwidth of the demodulator rather than symbol coherence. Tolerance does not decrease sharply with spreading factor in this regime.

Regime 2 (SF9–SF12): Tolerance decreases approximately exponentially with spreading factor. Fitting Eq. 4 to the three directly measured Regime 2 points (SF9: 11.5 kHz; SF10: 5.5 kHz; SF12: 0.5 kHz) yields

$$\mathrm{CFO}_{\max}(\mathrm{SF}) = 1.945 \times 10^{5} \cdot e^{-1.067\,\mathrm{SF}}\ \mathrm{kHz} \quad (7)$$

with $R^2 = 0.998$. The pre-exponential factor (1.945×10^5 kHz) is a regression constant with no independent physical interpretation; it represents the extrapolated intercept of the fitted curve at SF = 0,

far outside the measurement range. The decay rate $b = 1.067$ reflects the per-unit-SF reduction in tolerance. The measured values exceed the Semtech theoretical predictions by factors of 66–189.

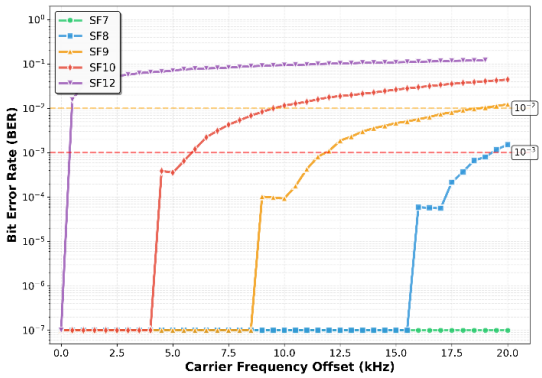

Figure 1: Bit error rate (log scale) versus carrier frequency offset for SF7–SF12. Error onset shifts from >20 kHz (SF7) to 0.5 kHz (SF12), demonstrating the 40-fold tolerance variation. The 10^{-3} threshold (dashed) defines $\mathrm{CFO}_{\max}^{(\mathrm{BER})}$.

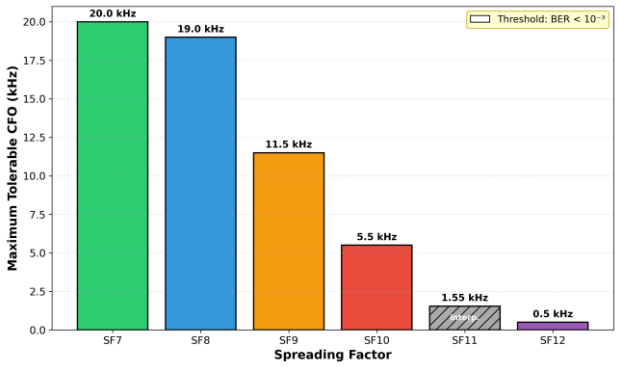

Figure 2: Measured $\mathrm{CFO}_{\max}^{(\mathrm{BER})}$ for the five directly characterized spreading factors (SF7, SF8, SF9, SF10, SF12). Tolerance decreases from >20 kHz at SF7 to 0.5 kHz at SF12; SF11 is absent because the gr-lora decoder defect precluded direct measurement (see Section 3.4). The BER < 10^{-3} threshold criterion applies throughout.

Table 3: Summary of experimental results. $\mathrm{CFO}_{\max}$ uses BER < 10^{-3} criterion. $\mathrm{PDR}_{\mathrm{sweep}}$: delivery across full 0–20 kHz sweep (n = 41 per SF, n = 39 for SF12). $\mathrm{PDR}_{\leq \mathrm{CFOmax}}$: delivery at offsets ≤ $\mathrm{CFO}_{\max}$ only. 95% CIs are Clopper–Pearson exact. SF11 is an interpolation estimate (est.). $\mathrm{PDR}_{\mathrm{sweep}}$ falls below 100% for SF8–SF12 because packets at high offsets (> $\mathrm{CFO}_{\max}$) fail. $\mathrm{PDR}_{\leq \mathrm{CFOmax}}$ isolates delivery within the usable operating range; 100% in all measured SFs confirms FEC maintains link integrity below the BER threshold under controlled conditions. SF12 $\mathrm{PDR}_{\leq \mathrm{CFOmax}}$ CI is wide owing to n = 1 qualifying packet.

SF	Ts (ms)	CFOmax (kHz / ppm)	PDRsweep [95% CI]	PDR≤CFOmax [95% CI]	H/room	Osc. Margin
SF7	2.048	>20.0 / >23.0	100% [91.6, 100]	100% [91.6, 100]	82×	>2.30×
SF8	4.096	19.0 / 21.9	82.9% [67.9, 92.8]	100% [79.4, 100]	156×	2.19×
SF9	8.192	11.5 / 13.2	48.8% [32.9, 64.9]	100% [75.3, 100]	189×	1.32×
SF10	16.384	5.5 / 6.3	24.4% [12.4, 40.3]	100% [69.2, 100]	183×	0.63×
SF11 (est.)	32.768	1.55 / 1.78	—	—	~111×	0.18×
SF12	65.536	0.5 / 0.58	2.6% [0.1, 13.5]	100% [2.5, 100]	66×	0.06×

5. DISCUSSION

Scope and generalisability. The reported CFO tolerance curves are empirical characterisations for the specific radio configuration and experimental conditions tested: a single transmitter–receiver device pair, an indoor high-link-margin setting, and the gr-lora software decoding chain. While the qualittive SF-dependent trend is expected to hold broadly—being grounded in the phase accumulation physics of chirp spread spectrum—the numerical thresholds should be treated as device- and implementation-dependent until replicated across additional radio front-ends, oscillator lots, and propagation conditions.

5.1 Implementation Headroom

The measured CFO tolerances exceed Semtech theoretical predictions by factors of 66–189 (Table 3). This headroom arises from receiver-side implementation techniques not modelled by the analytical framework of Eq. 3: fractional frequency estimation enables sub-symbol tracking beyond the integer-chirp limit; soft-decision decoding provides an additional 2–3 dB equivalent margin; oversampling and interpolation improve frequency resolution; and FEC at rate 4/5 maintains packet integrity through elevated BER up to approximately 10%. The nonlinear compression of headroom from 189× at SF9 to 66× at SF12 suggests that these techniques approach their operating limits as symbol duration exceeds approximately 30 ms, consistent with the phase accumulation analysis of Eq. 2.

Comparison with prior work is limited to SF7, where Van den Abeele et al. [6] provide the only published empirical reference point. Our SF7 tolerance exceeds that value, attributable to the higher link margin in this study, which enables the demodulator to distinguish CFO-induced phase errors from thermal noise more effectively. No prior empirical SF12 CFO measurement is available for direct comparison.

5.2 SF-Dependent Vulnerability Threshold

The principal finding of this study is a graduated vulnerability threshold rather than a binary bounry:

Safe zone (SF7–SF9): Measured tolerances are 1.32× or more above the ±10 ppm drift limit. Under realistic temperature variation and component aging, SF7–SF9 deployments with commodity oscillators are expected to operate reliably. The 94.3 dB link margin used in this study is higher than typical deployments; the threshold should be validated under lower-margin conditions before generalisation.

Marginal zone (SF10): The BER-based oscillator margin is 0.63×. This study observed 100% application-layer $PDR_{\leq CFOmax}$ at SF10 under controlled conditions, but this outcome should not be interpreted as indicating adequate commodity oscillator performance for production deployment. The conditions under which this result was obtained—stable temperature, high link margin, new components—are unlikely to be representative of deployed networks. SF10 with commodity oscillators requires case-by-case evaluation of the deployment environment.

Failure zone (SF11–SF12): Margins of 0.18× and 0.06×, respectively, indicate that commodity oscillator drift substantially exceeds the measured physical-layer tolerance. The 95% CI on SF12 PDR_{sweep} of [0.1%, 13.5%] confirms that packet delivery at SF12 with commodity oscillators is unreliable under the conditions of this study. ADR algorithms that assign SF11 or SF12 to nodes equipped with commodity oscillators are therefore expected to produce systematic link failures, even at nominally adequate SNR.

5.3 Oscillator Selection Framework

The following guidance is derived from the experimental results. It applies to the hardware and frequency band characterised here; users are encouraged to validate against their specific chipset and environment.

For deployments operating at SF7–SF9 (typical of smart agriculture and urban sensing applications with short-to-medium range links), commodity ±10 ppm oscillators are indicated by the measured margins. Mid-grade ±5 ppm components provide additional environmental margin at modest incremental cost.

For SF10 deployments (indoor penetration and extended range applications), mid-grade ±5 ppm oscillators are advisable given the sub-unity commodity margin. Stable indoor environments with controlled temperature may tolerate commodity oscillators, but this determination requires site-specific validation.

For SF11–SF12 deployments (maximum-range applications), TCXOs at ±2 ppm or better are necessary. The measured margins indicate that commodity and mid-grade oscillators cannot reliably support demodulation at these spreading factors under the conditions tested.

For ADR-enabled networks, two approaches are available: specify TCXOs universally to accommodate the full ADR spreading factor range, or implement spreading-factor-aware ADR that restricts commodity-oscillator nodes to SF ≤ 9. The latter approach is preferable where traffic distribution favours lower spreading factors, as it avoids the cost of universal TCXO adoption.

5.3 Cost Sensitivity and Assumptions

The oscillator selection guidance above has cost implications for large deployments. Indicative component prices—±10 ppm at approximately $0.15, ±5 ppm at approximately $0.50, and ±2 ppm TCXO at approximately $2.50 per unit—imply that specifying TCXOs for a 100,000-sensor deployment incurs an incremental cost of approximately $235,000 compared with commodity oscillators. These figures are sensitive to pricing, procurement scale, and regional

availability; a ±25% price variation does not alter the qualitative conclusion that cost optimisation is achievable at lower spreading factors. Networks operating exclusively at SF7–SF9 may adopt commodity oscillators with confidence, while networks using SF11–SF12 incur TCXO cost regardless of deployment scale.

These cost estimates do not include maintenance costs associated with oscillator-induced link failures, which will depend on network topology, service-level requirements, and fault-diagnosis capability. Such costs are outside the scope of this experimental study.

6. LIMITATIONS

Several limitations of this study should be noted. First, measurements were conducted with a single transmitter–receiver pair; results may vary with alternative LoRa chipsets, antenna configurations, or outdoor environments subject to multipath and interference. Second, the controlled laboratory temperature eliminated thermal drift; production environments experiencing cycling from −20°C to +70°C may expand effective oscillator drift beyond the nominal ±10 ppm specification. Third, the SF11 CFO tolerance is an interpolation estimate rather than a direct measurement; the associated sensitivity analysis bounds the uncertainty in the vulnerability classification but does not substitute for hardware validation. Fourth, ADR vulnerability conclusions are derived from the static measurements of this study; no dynamic ADR loop experiment was conducted, and the stated failure modes are analytical inferences rather than observed ADR events.

ADR implication (Modeled). The interaction with ADR is an inferred engineering implication based on the measured CFO tolerance bounds. The scenario in which an ADR server increases the spreading factor beyond the node's oscillator capability is plausible and consistent with the data, but it was not directly observed or experimentally triggered in this study. Validation requires an end-to-end ADR experiment with instrumented oscillator drift.

7. CONCLUSIONS

This paper has presented a systematic empirical characterisation In LoRa, **CFO arises** primarily from crystal oscillator inaccuracy. of CFO tolerance at SF7–SF10 and SF12 under controlled high-margin indoor conditions, with SF11 estimated by log-linear interpolation. The principal results are as follows.

CFO tolerance decreases by a factor of 40 from SF7 (measured >20 kHz) to SF12 (measured 0.5 kHz), following a two-regime pattern: a tracking-limited plateau at SF7–SF8 and an exponential coherence-failure decay at SF9–SF12 described by Eq. 7. Measured tolerances exceed Semtech theoretical predictions by 66–189×, attributed to receiver-side implementation techniques.

Comparison with the 8.68 kHz drift of ±10 ppm oscillators at 868 MHz reveals a graduated vulnerability threshold: SF7–SF9 operates with adequate margin (≥1.32×), SF10 is marginal (0.63×, functional under controlled conditions but unreliable in production), and SF11–SF12 lies below the drift limit (≤0.18×), with SF12 packet loss confirmed experimentally. Because ADR algorithms may dynamically increase the spreading factor, a node adequate at SF7 is not necessarily adequate at higher spreading factors assigned by the ADR server.

An oscillator selection framework derived from these results recommends commodity ±10 ppm components for SF7–SF9 deployments, mid-grade ±5 ppm components for SF10, and TCXOs at ±2 ppm or better for SF11–SF12. For ADR-enabled networks, either universal TCXO specification or spreading-factor-aware ADR ceiling is advised.

Future work should validate these findings across multiple hardware platforms and chipset generations, characterise the low-SNR regime where thermal noise and CFO interact, conduct temperature-sweep measurements to quantify combined drift effects, and validate the SF11 estimate through hardware correction of the decoder defect. The numerical thresholds reported here should be treated as device- and implementation-dependent until replicated across additional radio front-ends and propagation conditions.

DATA AND CODE AVAILABILITY

A versioned reproducibility artefact package (including raw CFO sweep logs, firmware, GNU Radio receiver flowgraph, and analysis scripts used to reproduce the reported figures) has been archived in a restricted Zenodo draft for this study and can be provided to the editor or reviewers upon request during peer review. A public release will be made upon acceptance.

ACKNOWLEDGEMENTS

The author(s) acknowledge the University of Benin for providing laboratory facilities and equipment used in this experimental research.

REFERENCES

[1]. U. Raza, P. Kulkarni, and M. Sooriyabandara, "Low power wide area networks: An overview," IEEE Communications Surveys & Tutorials, vol. 19, no. 2, pp. 855–873, 2017, doi: 10.1109/COMST.2017.2652320.

[2]. Semtech Corporation, "AN1200.13: SX1272/73/76/77/78/79 modem design LoRa guide," Semtech Corporation, Application Note Rev. 3.1, 2021.

[3]. D. Croce, M. Gucciardo, S. Mangione, G. Santaromita, and I. Tinnirello, "Impact of LoRa imperfect orthogonality: Analysis of link-level performance," IEEE Communications Letters, vol. 22, no. 4, pp. 796–799, 2020, doi: 10.1109/LCOMM.2018.2797057.

[4]. B. Reynders, Q. Wang, P. Tuset-Peiro, X. Vilajosana, and S. Pollin, "Improving reliability and scalability of LoRaWAN through lightweight scheduling," IEEE Internet of Things Journal, vol. 5, no. 3, pp. 1830–1842, 2018, doi: 10.1109/JIOT.2018.2815150.

[5]. A. Rahmadhani and F. Kuipers, "When LoRaWAN frames collide," in Proceedings of the international workshop on wireless network testbeds, experimental evaluation & characterization, 2018, pp. 89–97. doi: 10.1145/3212480.3212483.

[6]. F. Van den Abeele, J. Haxhibeqiri, I. Moerman, and J. Hoebeke, "Scalability analysis of large-scale LoRaWAN networks in NS-3," in IEEE internet of things journal, 2017, pp. 2186–2198. doi: 10.1109/JIOT.2017.2768498.

[7]. M. Chiani and A. Elzanaty, "On the LoRa chirp spread spectrum modulation: Signal properties and digital waveform," IEEE Transactions on Wireless Communications, vol. 18, no. 7, pp. 3625–3637, 2019, doi: 10.1109/TWC.2019.2913594.

[8]. T.-H. To and A. Duda, "Simulation of LoRa in NS-3: Improving LoRa performance with CSMA," in Proceedings of the IEEE international conference on communications, 2018, pp. 1–7. doi: 10.1109/ICC.2018.8422221.

[9]. T. Polonelli, D. Brunelli, A. Marzocchi, and L. Benini, "Slotted ALOHA overlay on LoRaWAN: A distributed synchronization approach," in Proceedings of the IEEE international conference on embedded wireless systems and networks, 2018, pp. 85–90.

[10]. P. Robbe, "gr-lora: GNU Radio LoRa receiver." GitHub repository, 2017. Available: https://github.com/rpp0/gr-lora