Open-Source LoRa PHY for Medium-Range IoT: Reverse-Engineering SF 5 and SF 6

Sascha Rösler, Anatolij Zubow and Falko Dressler
School of Electrical Engineering and Computer Science, TU Berlin, Germany
{roesler,zubow,dressler}@tkn.tu-berlin.de

***Abstract*—LoRa's extension into the 2.4 GHz ISM band introduced the new spreading factor (SF) 5 and SF 6, together with wider bandwidths of up to 1.625 MHz enabling higher data rates. However, these lower spreading factors remain unsupported in existing open-source LoRa implementations, leaving a critical gap for researchers and developers. In this paper, we reverse-engineer the framing and header encoding specific to SF 5 and SF 6 as implemented in the Semtech SX1280 transceiver. We identify three key differences from higher SFs: (i) modified synchronization field values, (ii) an added fine synchronization field, and (iii) disabled low data rate optimization (LDRO) for headers with SF $\leq$ 6. Based on these findings, we extend the open-source LoRa PHY implementation of Xu et al. [1] to fully support the additional two SFs. We validate our implementation through simulations in additive white Gaussian noise (AWGN) channel, wired attenuation measurements, and over-the-air experiments using legacy LoRa hardware. Results reveal signal to noise ratio (SNR) sensitivity losses consistent with higher SFs.**

***Index Terms*—IoT, LoRa, long-range narrow-band**

I. Introduction

The rapid proliferation of internet of things (IoT) paradigm has created a demand for wireless communication technologies that balance range, data rate, and energy efficiency. Low-power wide area networks (LPWANs) have emerged as a key enabler for large-scale IoT deployments, with LoRa being one of the most widely adopted radio technology. LoRa's chirp spread spectrum (CSS) modulation allows for flexible trade-offs between communication range and data rate through the spreading factor (SF) parameter, making it suitable for a broad range of applications.

Although high SFs (SF 10–12) are well suited for long-range, low-data-rate applications such as remote environmental monitoring, lower SFs allow for shorter air times by reducing the communication range (see Figure 1). Indeed, applications like dense urban sensing and indoor asset tracking represent scenarios, where devices are separated by up to a few hundred meters. These applications fall into the technology gap between SF $\geq$ 7 LoRa and Bluetooth (see Figure 2) as the communication range offered by SF 7 is not needed and inefficient. The lower spreading factors SF 5 and SF 6, introduced alongside the Semtech SX1280 transceiver for the 2.4 GHz ISM band, directly address this need. Thus, with SF 5 airtime and energy can be saved while its communication range is reduced. Together with the support of bandwidths of up to 1.625 MHz LoRa achieves shorter symbol air-times, enables higher throughput and reduces the collision probability in medium-range deployments, while retaining LoRa's inherent resilience to interference [2].

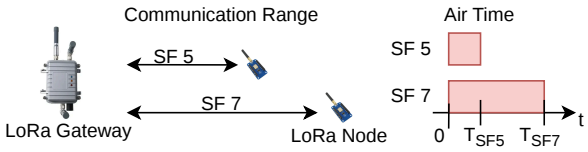

Figure 1. Trade-off between communication range and packet air time for LoRa: lower SFs yield shorter air times and higher data rates at the cost of reduced link range, positioning SF 5 and SF 6 for medium-range IoT.

Despite their practical relevance, SF 5 and SF 6 remain unsupported in existing open-source LoRa implementations. This is a critical limitation, as open-source PHY implementations are essential to enable reproducible experimentation, interference mitigation strategies, network architectures, signal emulation [4] and prototyping new MAC protocols. These tasks are infeasible with closed, proprietary firmware. Prior reverse-engineering efforts have made significant progress in this direction: gr-lora [5] and Marquet et al. [6] demonstrated the feasibility of software-defined LoRa receivers, Tapparel et al. [7] introduced robust center frequency offset compensation, and Xu et al. [1] closed the remaining SNR performance gap. However, none of these implementations support SF 5 and SF 6, leaving a gap for the medium-range IoT operating regime.

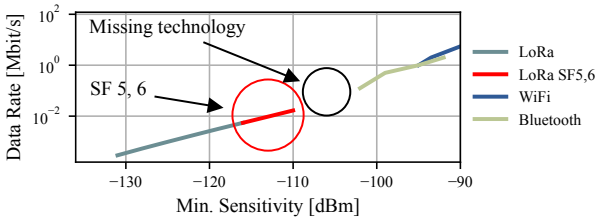

Figure 2. Receiver sensitivity versus achievable data rate for different wireless technologies in 2.4 GHz band. LoRa SF 5+6 substantially narrow the sensitivity gap between LoRa and Bluetooth, however, still this gap is not yet closed. Figure adjusted form [3]

In this paper, we address this gap by reverse-engineering the framing and header encoding of SF 5 and SF 6 as implemented in the SX1280 transceiver. We identify three specific differences from higher SFs, namely changes in synchronization field values, the introduction of a fine synchronization field, and the disabling of low data rate optimization (LDRO) for SF ≤ 6 headers. Based on these findings, we extend the open-source implementation of Xu et al. [1] to fully support SF 5 and SF 6, and validate it through simulations and real-world experiments. Our implementation is made publicly available to foster reproducibility and further research on LoRa.

II. Related Work

LoRa PHY reverse engineering: The proprietary PHY layer of LoRa has motivated a number of reverse-engineering efforts. Knight and Seeber [5] presented gr-lora, the first open-source GNU Radio implementation of a LoRa receiver, revealing the core structure of CSS modulation, Hamming coding, interleaving, and whitening. Marquet et al. [6], [8] independently validated these findings and extended the evaluation to Rayleigh fading channels, yet these show robustness gaps compared to commercial Semtech chipsets under real world scenarios [1]. Tapparel et al. [7] were the first to explicitly address center frequency offset (CFO) estimation and compensation, substantially improving demodulation robustness. Subsequently, Xu et al. [1] closed the remaining signal to noise ratio (SNR) performance gap in the low-SNR regime and demonstrated reliable decoding of SF 11/12, spreading factors where all prior implementations failed in experimental evaluations. Despite this progress, none of these implementations support SF 5/6, which we address in this work.

LoRa at 2.4 GHz: With the introduction of the Semtech SX1280 transceiver, LoRa operation was extended to the globally available 2.4 GHz ISM band, removing regional duty-cycle restrictions and enabling wider bandwidths and new spreading factors. Janssen et al. [9] provided an early characterization of the LoRa 2.4 GHz link, deriving achievable communication ranges for different SFs and bandwidth combinations, and showing that the lower SFs offer substantially higher data rates of up to 254 kbit/s but at the cost of reduced range. Falanji et al. [10] further investigated range and capacity in simulated urban deployments, confirming that shorter air-times at low SFs translate directly into reduced collision probability and higher network throughput. Hernández et al. [11] evaluated 2.4 GHz LoRa in indoor environments and found that low SF and high bandwidth configurations such as SF 5 and SF 6 achieve competitive packet delivery performance at short to medium distances, directly motivating their use in dense indoor IoT scenarios. These studies collectively underscore the practical relevance of SF 5 and SF 6.

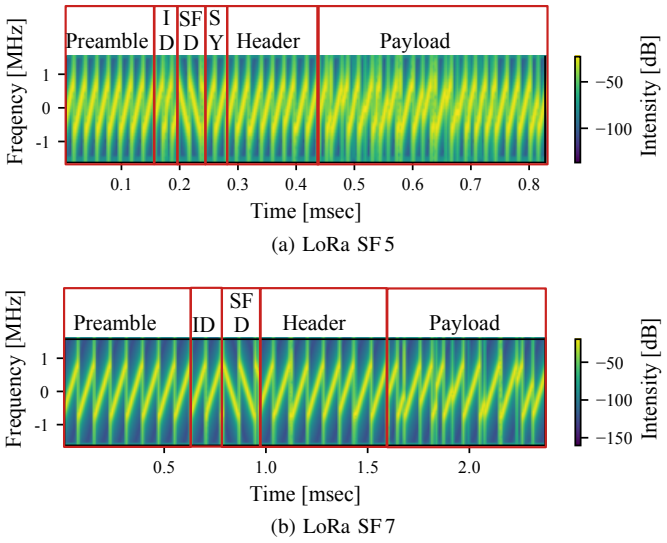

(a) LoRa SF 5

(b) LoRa SF 7

Figure 3. Spectrograms of recorded LoRa signals at SF 5 (a) and SF 7 (b) at 1.625 MHz bandwidth with annotated frame fields. The SF 5 chirp has a $4\times$ steeper slope than SF 7, and the fine synchronization field (SY) is clearly visible after the start frame delimiter (SFD) which is absent in SF 7.

LoRa scalability and network-layer research: A key motivation for open-source LoRa PHY implementations is the enabling of network-level research beyond closed firmware. Indeed, Liando et al. [2] conducted a large-scale measurement study with commercial hardware revealing fundamental trade-offs between SF, range, energy, and collision probability. However, for a deeper view they started looking into the PHY. Building on open PHY access, recent works have exploited the structure of CSS modulation to recover colliding packets [12], [13], improve demodulation at ultra-low SNR via deep learning [14] and demodulating on multiple channels [15]. Also non-linear chirps have been studied [16], [17].

III. LoRa Primer

A. Modulation

LoRa modulates data using CSS, where each symbol is represented by a chirp, i.e. a sinusoidal signal whose instantaneous frequency sweeps linearly across the channel bandwidth BW. Each chirp consists of 2^{SF} chips of duration $1/\mathrm{BW}$, so the total symbol duration is $T_s = 2^{\mathrm{SF}}/\mathrm{BW}$. The symbol value is encoded in the starting frequency f_0: with 2^{SF} distinct starting frequencies available, each symbol carries exactly SF bits of information. For an upchirp, the instantaneous frequency increases monotonically from f_0 until it reaches $+\mathrm{BW}/2$, wraps around to $-\mathrm{BW}/2$, and continues until the chirp period is complete. Demodulation is performed by multiplying the received signal by the complex conjugate of a baseline upchirp (i.e., a downchirp), then applying a fast Fourier transformation (FFT) [1], [2]. The frequency bin of the dominant peak in the

resulting folded spectrum directly yields the transmitted symbol value. A key consequence of CSS is that the SF controls the trade-off between spectral efficiency and link robustness: lower SFs yield shorter air-times and higher data rates at the cost of reduced SNR sensitivity, while higher SFs improve noise resilience at the expense of longer transmissions and higher collision probability. [1], [2]

B. Frame structure and encoding

A LoRa packet consists of a preamble, a header, and a payload, as illustrated in Figure 3b. The preamble comprises a configurable number of upchirps followed by two upchirps for the network identifier (ID) and two downchirps forming the SFD, which enables symbol-level synchronization at the receiver. The encoding pipeline, which is shown in Figure 4, processes header and payload separately but through the same sequence of operations. Payload bits are first whitened to decorrelate the data, then arranged into blocks of SF nibbles (4-bit half-bytes). For $\mathrm{SF} > 10$, low data rate optimization (LDRO) is automatically enabled (in green in Figure 4): the two least significant bits per nibble, which are most vulnerable to frequency drift over long symbol durations, are suppressed, reducing the effective block size to $\mathrm{SF} - 2$ nibbles and using one bit as a parity check. Hamming coding then expands each nibble to $\mathrm{CR} + 4$ bits, where $\mathrm{CR} \in \{1, 2, 3, 4\}$ is the configurable code rate. A diagonal interleaver rearranges the coded bits to spread burst errors across multiple symbols, producing $\mathrm{CR} + 4$ LoRa symbols of SF bits each. Finally, Gray mapping is applied before modulation to minimize bit errors from adjacent symbol confusions. The header follows the same pipeline but is encoded in LDRO mode and occupies a fixed number of symbols, regardless of the payload SF. It is always coded with the highest code rate of $4/8$. [1], [7]

C. Relevance to SF 5 and SF 6

With $2^5 = 32$ and $2^6 = 64$ chips per symbol respectively, SF 5 and SF 6 produce the steepest chirp slopes and shortest air-times among all LoRa configurations. At a bandwidth of 1.625 MHz, a single SF 5 symbol lasts only $\approx 20\,\mu$s, enabling data rates of 254 kbit/s [18], i.e. nearly two orders of magnitude beyond what SF 12 can achieve. However, the reduced symbol duration has a direct impact on the header encoding: as we show in Section IV, the standard LDRO-based header format cannot accommodate the required 40 bits within 8 SF 5 or SF 6 symbols, necessitating a deviation from the encoding used for higher SFs.

IV. LoRa SF 5 and SF 6 Encoding

We reverse-engineered the encoding of LoRa SF 5 and SF 6 using the SX1280 chip on the iM282A[1] board as the reference hardware. Therefore, we captured raw IQ samples using a universal software radio peripheral (USRP) device and analyzed the frame structure by comparing spectrograms of SF 5 and SF 7 transmissions, as shown in Figure 3. The steeper chirp slope visible for SF 5 is consistent with the $2^5 = 32$ chips per symbol, compared to $2^7 = 128$ chips for SF 7, resulting in a fourfold reduction in symbol duration. At a bandwidth of 1.625 MHz, this translates to raw data rates of 254 kbit/s and 152 kbit/s for SF 5 and SF 6, respectively [18]. We observe no changes to the underlying chirp modulation or the spreading mechanism. However, a systematic comparison of the frame structure against $\mathrm{SF} \geq 7$ recordings reveals three distinct encoding differences, which we describe in detail below.

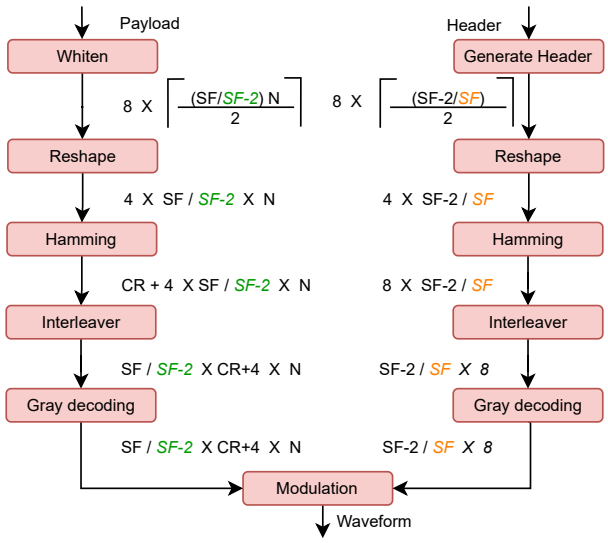

Figure 4. LoRa encoding pipeline with intermediate data structure sizes. The green path shows low data rate optimization (LDRO) active for $\mathrm{SF} > 10$. The orange path highlights the header encoding for $\mathrm{SF} \leq 6$, where LDRO is disabled to fit the 40-bit header into 8 symbols.

A. Change 1: Modified synchronization field values

In the preamble of LoRa $\mathrm{SF} \geq 7$, two network ID symbols (ID Figure 3b) are transmitted to allow the receiver to identify the target network and reject packets from foreign networks. According to the patent of Seller and Sornin [19], these fields are redefined in the extended LoRa standard as synchronization symbols with nominal values of 4 and 8. In our SF 5 recordings, we observe synchronization symbols at the corresponding positions, but with values of 8 and 16 rather than 4 and 8. This doubling of the symbol values appears to be a deliberate design choice specific to the lower spreading factors, though no public documentation explains the rationale. We verified this finding across multiple captures and confirmed the values are consistent and independent of payload content.

[1]https://wireless-solutions.de/data/Wireless\%20Solutions/Downloads/iM282A-L/iM282A_Datasheet.pdf

B. Change 2: New fine synchronization field

For SF 5 and SF 6, a new fine synchronization field (SY) is appended immediately after the SFD, as visible in Figure 3a but absent in the SF 7 capture. This field, also anticipated by Seller and Sornin [19], consists of two additional base chirp symbols and is intended to improve sub-symbol timing alignment at the receiver. This can be required by the extremely short symbol duration of SF 5 ($\approx 20\,\mu$s at 1.625 MHz) reducing the tolerance for timing offset compared to higher SFs. In our recordings, both fine synchronization symbols consistently carry a value of 1. This field must be correctly detected and stripped by the decoder before header processing begins; failure to account for it causes a systematic misalignment of all subsequent symbol boundaries.

C. Change 3: LDRO disabled for the header

For SF $\geq$ 7, the LoRa header is always encoded in LDRO mode, which encodes $\mathrm{SF}-2$ bits per nibble block to improve robustness against frequency drift over long symbol durations. Under this scheme, the 40-bit header requires $\lceil 40/(\mathrm{SF}-2)\rceil$ symbols: 8 symbols for SF $\geq$ 7. However, for SF 5, LDRO would allow only 3 bits per symbol, yielding a capacity of just 24 bits across 8 header symbols which is insufficient for the 40-bit header. For SF 6, the situation is analogous: LDRO yields only 32 bits across 8 symbols. We confirmed experimentally that correct header decoding is achieved by disabling LDRO for SF ≤ 6 (in orange in Figure 4), reverting to the full SF bits per symbol. This allows 8 SF 5 symbols to carry exactly $8 \times 5 = 40$ bits, perfectly matching the header length. Importantly, the payload encoding pipeline is unaffected: LDRO remains disabled for SF 5 and SF 6 payload symbols as well, since the short symbol duration eliminates the frequency drift concerns that originally motivated LDRO.

D. Summary

Table I summarizes the three encoding differences between SF 5/6 and the higher SFs. The decoding pipeline is the reverse of the encoding process described above. We extend the open-source implementation of Xu et al. [1] with all three changes and make the code publicly available to the community[2].

V. Performance Evaluation

We evaluate the correctness and performance of our SF 5/6 implementation through three complementary experiments: (i) an AWGN simulation, (ii) a wired attenuation measurement, and (iii) an over-the-air test. For all experiments, we configure LoRa to use a bandwidth of 1.625 MHz and a code rate of $4/5$, matching the maximum bandwidth supported by the SX1280 chipset. We transmit packets of 10 Byte, assuming data within a sensor network. The metric is packet error ratio (PER) and packet delivery ratio (PDR) measured as a function of SNR or signal attenuation across SF 5 through SF 10. The key question is whether the extended implementation achieves SNR sensitivity consistent with the theoretical expectation for CSS modulation, i.e., an approximately constant sensitivity loss per SF step [20].

[2]https://github.com/tkn-tub/LoRaPHY

Table I
Encoding differences between LoRa SF 5/6 and SF $\geq$ 7

Property	SF $\geq$ 7	SF 5 and SF 6
Sync field values	network ID	8, 16
Fine sync field values	absent	1, 1
Header LDRO	enabled	disabled
Payload LDRO	SF $>$ 10 only	disabled

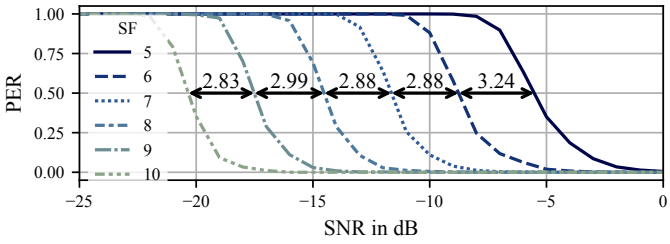

Figure 5. PER versus SNR for SF 5-10 under an AWGN channel (BW=1.625 MHz, code rate $4/5$). The annotated values denote the SNR gap at PER = 0.5 between adjacent SFs. The SF 5 and SF 6 curves are consistent with the envelope established by SF $\geq$ 7.

A. Simulation under AWGN

We first validate our implementation in a controlled simulation in MATLAB using both the extended transmitter and receiver under an additive white Gaussian noise (AWGN) channel. Figure 5 shows the PER as a function of SNR for all evaluated SFs. The curves exhibit the characteristic steep transition of around 4-6 dB from near-zero to near-unity PER that is typical of CSS modulation, confirming correct demodulation behavior. The SNR sensitivity loss which is defined as the SNR gap between adjacent SFs at a PER of 0.5 is approximately 2.9 dB for SF 6 and 3.2 dB for SF 5, compared to a reference loss of 2.88 dB at SF 7. The marginally larger loss at SF 5 (0.3 dB above the SF $\geq$ 6) is within tolerance. Importantly, all SF 5 and SF 6 curves fall within the envelope established by the higher SFs, confirming that our changes to encoding do not introduce systematic performance degradation.

B. Wired measurement

To validate the implementation against real hardware under controlled channel conditions, we connect a USRP transmitter to a commercial iM282A board (Semtech SX1280) via a coaxial cable and a programmable attenuator, operating at

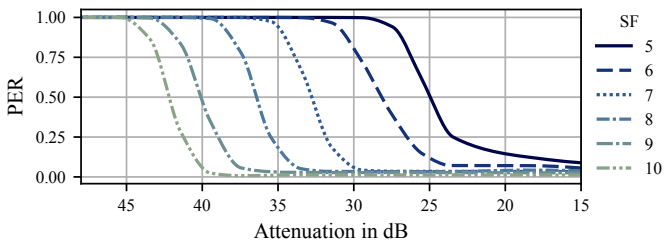

Figure 6. PER versus cable attenuation for SF 5-10 (B=1.625 MHz, code rate $4/5$), measured in a setup with USRP transmitter and a commodity SX1280 receiver. Transition midpoints are spaced approximately 3–5 dB per SF step.

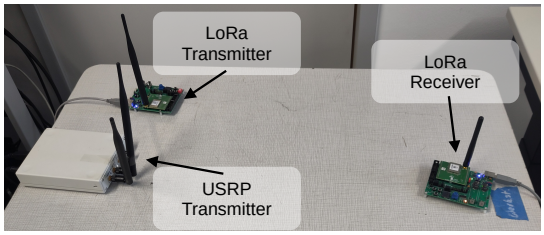

Figure 7. Over-the-air measurement setup with USRP node (left) and two commodity iM282A boards with Semtech SX1280 transceivers (top and right) serving as legacy LoRa reference devices, all operating in the 2.4 GHz ISM band.

2.4 GHz. Figure 6 presents the measured PER across different attenuation values for SF 5 through SF 10. The overall behavior mirrors the simulation results, with each SF showing a well-defined attenuation threshold beyond which PER drops rapidly. The measured curves are slightly irregular compared to the simulation, which we attribute to hardware imperfections including local oscillator frequency offsets, coupled interferences and measurement resolution. Notably, the attenuation gap between SF 5 and SF 6 is compressed relative to the simulated prediction, while the gap between SF 6 and SF 7 is larger than expected, averaging each other out. Nevertheless, the USRP-generated SF 5 and SF 6 signals are correctly decoded by the commercial SX1280 receiver across the full attenuation range, confirming bitwise compatibility with the proprietary Semtech implementation.

C. Over-the-air measurement

Finally, we evaluate end-to-end performance in a realistic indoor propagation environment. The measurement setup, shown in Figure 7, consists of a USRP and two iM282A commodity LoRa boards, placed several meters apart in a laboratory room with sufficient multipath reflections. We measure the PDR for three transmission configurations: (i) USRP transmitter to iM282A receiver (USRP $\rightarrow$ LoRa), (ii) iM282A transmitter to iM282A receiver (LoRa $\rightarrow$ LoRa), and (iii) iM282A transmitter to USRP receiver (LoRa $\rightarrow$ USRP). In our experiment we send 100 times 40 packets. The results in Figure 8 show a PDR close to 100 % for all configurations and SFs, confirming that both the extended transmitter and receiver operate correctly in a real wireless propagation channel. Minor exceptions are observed at the two lowest spreading factors: SF 5 exhibits a slightly reduced PDR in the LoRa $\rightarrow$ LoRa configuration, and SF 6 shows a similar mild drop in the USRP $\rightarrow$ LoRa configuration. The fact that both anomalies occur exclusively at the lowest SFs and across different hardware configurations points likely to external interference in the 2.4 GHz ISM band as the most likely cause, rather than any implementation-specific artifact. As SF 5 and SF 6 offer the lowest gain due to less chips among all evaluated configurations, they are more susceptible to co-channel interference from other 2.4 GHz technologies such as WiFi and Bluetooth. This result underlines an important practical consideration for deploying SF 5 and SF 6 in dense indoor environments: while our implementation achieves correct decoding under normal conditions, link reliability at the lowest spreading factors may require careful channel planning and interference mitigation to ensure robust operation in the crowded 2.4 GHz band.

VI. Conclusion

In this paper, we addressed a gap in both open-source implementations and the published literature by presenting a reverse-engineered characterization of LoRa SF 5 and SF 6. Here we identified three key encoding differences specific to these lower spreading factors: modified synchronization field values, the introduction of a fine synchronization field, and the disabling of low data rate optimization (LDRO) for SF $\leq$ 6 headers. Building on these findings, we extended the open-source LoRa PHY library of Xu et al. [1] to fully support the two lower SFs. Simulative evaluations in AWGN channel, wired attenuation experiments, and over-the-air measurements with real hardware confirm the correctness of our implementation.

Future work will focus on two directions. First, we plan to improve the robustness of the LoRa receiver implementation

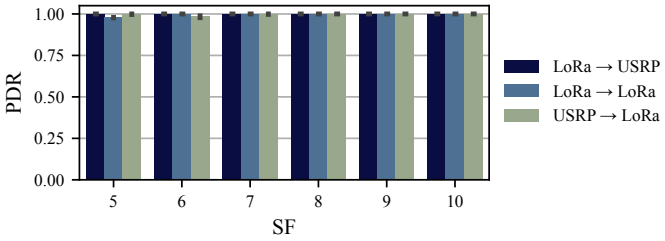

Figure 8. PDR over the air for three configurations across SF 5-10. All configurations achieve PDR close to 100 %, validating the extended implementation. Minor drops at SF 5/6 are likely due to co-channel interference from other 2.4 GHz technologies.

under challenging channel conditions, including frequency and timing offsets, to close the remaining performance gap with commercial LoRa hardware. Second, we intend to investigate the suitability of SF 5 and SF 6 in dense IoT deployments, where their shorter air-times may significantly reduce collision probability and energy consumption. Scenarios, such as successive interference cancellation, which may benefit particularly from the higher throughput offered by the lower spreading factors.

Acknowledgements

This work was supported by the German Research Foundation (DFG) within the project ResCTC under grants DR 639/30-1 and ZU 235/4-1.

References

[1] Z. Xu, S. Tong, P. Xie, and J. Wang, "From Demodulation to Decoding: Toward Complete LoRa PHY Understanding and Implementation," *ACM Transactions on Sensor Networks*, vol. 18, no. 4, pp. 1–27, Jan. 2023.

[2] J. C. Liando, A. Gamage, A. W. Tengourtius, and M. Li, "Known and Unknown Facts of LoRa: Experiences from a Large-scale Measurement Study," *ACM Transactions on Sensor Networks*, vol. 15, no. 2, pp. 1–35, May 2019.

[3] A. Zubow, I. von Stebut, S. Rösler, and F. Dressler, "ResCTC: Resilience in Wireless Networks through Cross-Technology Communication," in *IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC 2024)*, Valencia, Spain: IEEE, Sep. 2024, pp. 1–6.

[4] P. Gawłowicz, A. Zubow, and F. Dressler, "Wi-Lo: Emulation of LoRa using Commodity 802.11b WiFi Devices," in *IEEE International Conference on Communications (ICC 2022)*, Seoul, South Korea: IEEE, May 2022, pp. 4414–4419.

[5] M. Knight and B. Seeber, "Decoding LoRa: Realizing a Modern LPWAN with SDR," in *6th GNU Radio Conference*, Boulder, CO: GNU Radio Foundation, Sep. 2016.

[6] A. Marquet, N. Montavont, and G. Z. Papadopoulos, "Towards an SDR implementation of LoRa: Reverse-engineering, demodulation strategies and assessment over Rayleigh channel," *Elsevier Computer Communications*, vol. 153, pp. 595–605, Mar. 2020.

[7] J. Tapparel, O. Afisiadis, P. Mayoraz, A. Balatsoukas-Stimming, and A. Burg, "An Open-Source LoRa Physical Layer Prototype on GNU Radio," in *21st International Workshop on Signal Processing Advances in Wireless Communications (SPAWC 2020)*, Virtual Conference: IEEE, May 2020.

[8] A. Marquet, N. Montavont, and G. Z. Papadopoulos, "Investigating Theoretical Performance and Demodulation Techniques for LoRa," in *20th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM 2019)*, Washington, D.C.: IEEE, Jun. 2019, pp. 1–6.

[9] T. Janssen, N. BniLam, M. Aernouts, R. Berkvens, and M. Weyn, "LoRa 2.4 GHz Communication Link and Range," *Sensors*, vol. 20, no. 16, Aug. 2020.

[10] R. Falanji, M. Heusse, and A. Duda, "Range and Capacity of LoRa 2.4 GHz," in *EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services (MOBIQUITOUS 2022)*, Pittsburgh, PA: Springer, Nov. 2022, pp. 403–421.

[11] C. F. Hernández, G. Hochet Derévianckine, A. Guitton, O. Iova, and F. Valois, "Indoor Performance Evaluation of LoRa® 2.4 GHz," in *IEEE Wireless Communications and Networking Conference (WCNC 2023)*, Glasgow, United Kingdom: IEEE, Mar. 2023, pp. 1–6.

[12] M. O. Shahid, M. Philipose, K. Chintalapudi, S. Banerjee, and B. Krishnaswamy, "Concurrent interference cancellation: decoding multi-packet collisions in LoRa," in *ACM SIGCOMM 2021*, Virtual Conference: ACM, Aug. 2021, pp. 503–515.

[13] J. Álamos, T. C. Schmidt, and M. Wählisch, "CoRa: A Collision-Resistant LoRa Symbol Detector of Low Complexity," in *44th IEEE International Conference on Computer Communications (INFOCOM 2025)*, London, United Kingdom: IEEE, May 2025, pp. 1–10.

[14] C. Li, H. Guo, S. Tong, X. Zeng, Z. Cao, M. Zhang, Q. Yan, L. Xiao, J. Wang, and Y. Liu, "NELoRa: Towards Ultra-Low SNR LoRa Communication with Neural-Enhanced Demodulation," in *19th ACM Conference on Embedded Networked Sensor Systems (SenSys 2021)*, Coimbra, Portugal: ACM, Nov. 2021, pp. 56–68.

[15] P. Robyns, P. Quax, W. Lamotte, and W. Thenaers, "A Multi-Channel Software Decoder for the LoRa Modulation Scheme," in *3rd International Conference on Internet of Things, Big Data and Security (IoTBDS 2018)*, vol. 1, Vancouver, Canada: SciTePress, Feb. 2018, pp. 41–51.

[16] C. Li, Z. Cao, and L. Xiao, "CurveALOHA: Non-linear Chirps Enabled High Throughput Random Channel Access for LoRa," in *41st IEEE International Conference on Computer Communications (INFOCOM 2022)*, Virtual Conference: IEEE, May 2022, pp. 520–529.

[17] H. Shi, Y. Yang, X. Guo, W. Meng, C. Gu, and S. He, "Understanding and Optimizing Nonlinear Chirp Spread Spectrum Modulation in LoRa Networks," in *26th IEEE International Conference on High Performance Computing and Communications (HPCC 2024)*, Wuhan, China: IEEE, Dec. 2024, pp. 876–884.

[18] Semtech, "SX1280 Datasheet," Semtech, Data Sheet DS.SX1280-1.W.APP, Mar. 2020, pp. 1–158.

[19] O. B. A. Seller and N. Sornin, "Low complexity, low power and long range radio receiver," USPTO, US Patent US10305535B2, May 2019.

[20] M. J. Faber, K. M. van der Zwaag, W. G. V. dos Santos, H. R. d. O. Rocha, M. E. V. Segatto, and J. A. L. Silva, "A Theoretical and Experimental Evaluation on the Performance of LoRa Technology," *IEEE Sensors Journal*, vol. 20, no. 16, pp. 9480–9489, Aug. 2020.