

LORA DIGITAL RECEIVER ANALYSIS AND IMPLEMENTATION

Reza Ghanaatian, Orion Afisiadis, Matthieu Cotting, and Andreas Burg

Telecommunication Circuits Laboratory, École polytechnique fédérale de Lausanne, Switzerland

ABSTRACT

Low power wide area network technologies (LPWANs) are attracting attention because they fulfill the need for long range low power communication for the Internet of Things. LoRa is one of the proprietary LPWAN physical layer (PHY) technologies, which provides variable data-rate and long range by using chirp spread spectrum modulation. This paper describes the basic LoRa PHY receiver algorithms and studies their performance. The LoRa PHY is first introduced and different demodulation schemes are proposed. The effect of carrier frequency offset and sampling frequency offset are then modeled and corresponding compensation methods are proposed. Finally, a software-defined radio implementation for the LoRa transceiver is briefly presented.

Index Terms— LPWAN, IoT, LoRa, carrier frequency offset, sampling frequency offset

1. INTRODUCTION

Internet of Things (IoT) has sparked a lot of interest. It is foreseen that by 2025 Internet nodes will reside in every day objects, e.g., furniture, packages, etc [1] and [2]. IoT nodes often target energy autonomy with a battery lifetime of 10+ years, which creates business opportunities for new services in various fields such as home automation, traffic control, environmental monitoring, personal health care, etc.

Low Power Wide Area Network (LPWAN) technologies have recently emerged to complement existing communication standards. While short-range wireless networks such as Bluetooth, Zigbee, and Wireless Local Area Networks (WLANs) are designed to cover short distances with different rates and cellular networks are deployed to provide high rate with global coverage, LPWANs are meant to provide low data rate, wide area coverage and high energy efficiency [2] for the IoT. Several LPWAN technologies, such as different generations of 3GPP standards, e.g., NB-IoT, as well as proprietary ultra-low-rate standards, e.g., LoRa and SigFox, are predicted to co-exist in the future to provide connectivity for billions of nodes.

In this paper we focus specifically on the LoRa technology. LoRa is a proprietary physical layer (PHY) standard, which was developed by Cycleo and acquired by Semtech in 2012 [3]. LoRaWANTM is an open standard proposed by the LoRaTM Alliance [4] that defines the network architecture and layers above the LoRa PHY. The work of [5] and [6] provides a high-level system architecture overview of LoRa. Furthermore, several attempts were recently made to systematically

explain the LoRa PHY properties. In [7] a description of the reverse engineered proprietary LoRa PHY is presented and a software decoder using the GNU radio framework is provided. The work of [8] details the modulation and encoding elements that comprise the LoRa PHY. Although many aspects of the LoRa PHY are known, an in-depth analysis and detailed algorithmic description of a LoRa receiver is so far missing in the literature.

In this paper, a detailed analysis of the LoRa PHY is provided by studying multiple aspects of a LoRa receiver. In particular, we present different structures for LoRa demodulation, we study the synchronization process in LoRa receivers and we analyze the effect of carrier frequency offset (CFO) and sampling frequency offset (SFO) on the receiver performance. Finally, a software defined radio (SDR) implementation for LoRa is briefly shown.

2. LORA PHY OVERVIEW

LoRa employs chirp spread spectrum (CSS) modulation, which provides variable data rates by changing the *spreading factor*. Therefore, this modulation allows to trade throughput for coverage and/or energy consumption [5].

In this section, we first present the frame structure of LoRa, we then provide a high-level overview of the LoRa PHY and finally focus on the modulation and demodulation as the most important blocks of a LoRa transmitter and receiver.

2.1. LoRa Frame Structure

The LoRa PHY frame structure is defined in [4] and illustrated in Fig. 1. A frame is composed of a preamble with a number of preamble N_{pre} upchirps and 4.25 LoRa symbols as frame delimiters for synchronization, a PHY-header containing the frame information, a variable-length PHY-payload, and a cyclic redundancy check (CRC). The PHY-header and the CRC are optional.

2.2. LoRa PHY Block Diagram

Fig. 1 shows the block diagram of a LoRa transceiver. On the transmitter side, the input bits are first encoded using a Hamming code. Then whitening, interleaving, and Gray indexing are applied before modulation. LoRa uses CSS modulation for the preamble and the data. The LoRa CSS modulation is explained in more detail in the next subsection. The receiver performs synchronization and frequency-offset estimation and compensation prior to demodulation. Gray indexing, de-interleaving, de-whitening, and Hamming decoding are carried out to recover the information.

2.3. LoRa Modulation and Demodulation

2.3.1. LoRa Modulation

A LoRa CSS modulated symbol with spreading factor $SF \in \{6, 7, \dots, 12\}$ is defined as

$$x_S(t) = \begin{cases} e^{j2\pi\left(\frac{BW}{2T_s}t^2 + (f(S) - \frac{BW}{2})t\right)}, & 0 \leq t < t_{\text{fold}}, \\ e^{j2\pi\left(\frac{BW}{2T_s}t^2 + (f(S) - \frac{3BW}{2})t\right)}, & t_{\text{fold}} \leq t < T_s, \end{cases} \quad (1)$$

where $BW \in \{125, 250, 500\}$ kHz is the bandwidth, $T_s = \frac{2^{SF}}{BW}$ is the symbol duration, $f(S)$ is the initial frequency of a chirp, which depends on the data symbol $S \in \{0, 1, \dots, 2^{SF} - 1\}$ and is defined as $f(S) = S \cdot \frac{BW}{2^{SF}}$, and $t_{\text{fold}} = \frac{2^{SF} - S}{BW}$. LoRa is a spread spectrum technology, which indicates that 2^{SF} samples are transmitted per LoRa symbol to convey SF bits. To encode the data, the two-sided baseband bandwidth is split into 2^{SF} frequency steps. The symbol frequency starts at $f(S) - \frac{BW}{2}$ and increases linearly with time until it reaches the Nyquist frequency $\frac{BW}{2}$ at $t = t_{\text{fold}}$, where a *frequency fold* to $-\frac{BW}{2}$ occurs. Setting $S = 0$, results in an *upchirp*, whose frequency continuously increases during the symbol duration.

The discrete-time equation for a LoRa symbol is derived by replacing $t = \frac{n}{f_s}$ in (1), where n is the sample index and f_s is the sampling frequency, as

$$x_S[n] = \begin{cases} e^{j2\pi\left(\frac{1}{2 \cdot 2^{SF}}\left(\frac{BW}{f_s}\right)^2 n^2 + \left(\frac{S}{2^{SF}} - \frac{1}{2}\right)\left(\frac{BW}{f_s}\right)n\right)}, & n \in N_1, \\ e^{j2\pi\left(\frac{1}{2 \cdot 2^{SF}}\left(\frac{BW}{f_s}\right)^2 n^2 + \left(\frac{S}{2^{SF}} - \frac{3}{2}\right)\left(\frac{BW}{f_s}\right)n\right)}, & n \in N_2, \end{cases} \quad (2)$$

where $N_1 = \{0, \dots, n_{\text{fold}} - 1\}$, $N_2 = \{n_{\text{fold}}, \dots, 2^{SF} - 1\}$, and $n_{\text{fold}} = t_{\text{fold}} f_s$. By setting $f_s = BW$, (2) is simplified to

$$x_S[n] = e^{j2\pi\left(\frac{n^2}{2 \cdot 2^{SF}} + \left(\frac{S}{2^{SF}} - \frac{1}{2}\right)n\right)}, n \in \{0, 1, \dots, 2^{SF} - 1\} \quad (3)$$

2.3.2. LoRa Demodulation

The received signal is given by

$$y_S[n] = h x_S[n] + z[n], \quad (4)$$

where h denotes the block-fading channel, and $z[n]$ is the zero-mean white Gaussian complex-valued noise with variance σ^2 . A non-coherent demodulator applies 2^{SF} matched filters with the candidate reference symbols

$$X_k = \sum_{n=0}^{2^{SF}-1} y_S[n] x_k^*[n] = \sum_{n=0}^{2^{SF}-1} e^{j2\pi\left(n\left(\frac{S-k}{2^{SF}}\right)\right)}, \quad (5)$$

and retrieves the maximum-likelihood symbol estimate \hat{S} with

$$\hat{S} = \underset{k}{\operatorname{argmax}}(|X_k|). \quad (6)$$

Unfortunately, the complexity of the above scheme is high due to the required 2^{SF} convolutions in (5). Another, less complex formulation of the LoRa demodulator first multiplies the

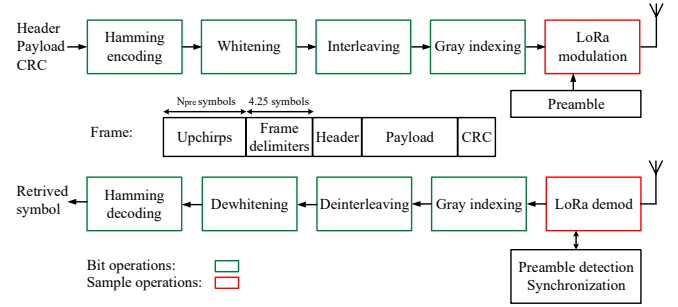


Fig. 1: LoRa frame structure [4] and PHY block diagram.

received symbol with the complex-conjugate of an upchirp, i.e., *dechirping*, as $y_{\text{dc}}[n] = y_S[n]x_0^*[n]$ followed by a discrete Fourier transform (DFT) of the dechirped signal. In this way, (5) can be written as

$$[X_1, X_2, \dots, X_{2^{SF}-1}] = \text{DFT}(y_{\text{dc}}[n]), \quad (7)$$

which is fully equivalent, but less complex to compute. We consider (7) for the analysis provided in Section 3.

3. LORA RECEIVER ANALYSIS

In the previous section, we assumed perfect synchronization and no frequency offset during the LoRa demodulation. In this section, we discuss the necessary initial synchronization and the receiver in the presence of CFO and SFO. Specifically, we propose a synchronization algorithm, we study the effect of CFO on the synchronization and demodulation, and we provide an analysis on the effect of SFO on error-rate performance.

3.1. Preamble Detection and Synchronization

The first step in the receiver is to detect the preamble and to synchronize to the frame boundary. To this end, we exploit the repeating upchirps in the preamble. More specifically, the synchronization is done in two steps: The receiver collects a block of samples¹ and computes (7). If at least one of the frequency bins has a magnitude that exceeds a given threshold, the index of the bin with the largest magnitude is noted as \hat{S}_{pre} and the process is continued. If a LoRa preamble is present, this value remains the same during the next DFT blocks and $N_{\text{pre}} - 1$ equal indices \hat{S}_{pre} are detected. Since the preamble consists of consecutive upchirps, the index \hat{S}_{pre} indicates the time offset in samples between the start of the block collected for synchronization and the start of the received preamble symbol. Therefore, the receiver synchronizes to the start of the header by skipping $2^{SF} - \hat{S}_{\text{pre}}$ samples and the initial frame delimiters.

3.2. CFO Formulation and Robustness Analysis

Low-cost crystal oscillators have an inherent mismatch with their nominal frequency value and therefore the down-conversion is performed with a different frequency than the

¹We note that by scanning the received signal power level, the receiver has a mean to detect a transmission, while the start of the preamble still needs to be detected.

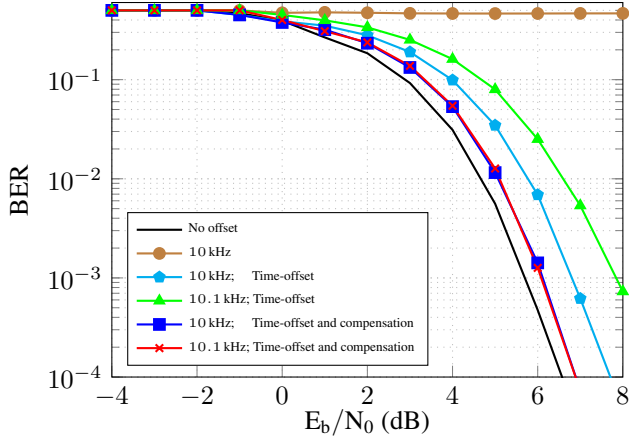


Fig. 2: BER for CFO of 10 kHz and 10.1 kHz.

up-conversion. In this subsection, we analyze the effect of the resulting CFO on the receiver performance.

We denote with f_{c1} the carrier frequency that is used for up-conversion, and with f_{c2} the frequency that is used for down-conversion. The difference $\Delta f_c = f_{c1} - f_{c2}$ is the carrier frequency offset. For the simplicity of notation we consider only the signal without the noise. A LoRa symbol after up- and down-conversion and dechirping can be written as

$$\tilde{y}_{dcS}[n] = \tilde{y}_S[n]x_0^*[n] = e^{j2\pi n \left(\frac{S}{2SF} - \frac{\Delta f_c}{f_s} \right)}, \quad (8)$$

where $\tilde{y}_S[n] = y_S[n] \cdot e^{j2\pi n \frac{\Delta f_c}{f_s}}$ is the LoRa symbol with CFO. The CFO results in a frequency shift, which can introduce an error in the decision of (6). More specifically, if the offset is sufficiently large to displace the peak in the Fourier domain by more than half a bin, i.e., $\frac{|\Delta f_c|}{f_s} > \frac{1}{2 \cdot 2^{SF}}$, a demodulation error will occur even without noise. This offset stays constant during the frame if the CFO stays constant over time.

The LoRa demodulator described in the previous section is, however, partially robust against a CFO since the CFO leads to a time offset in the synchronization equal to $\lfloor \frac{\Delta f_c}{f_s} \cdot 2^{SF} \rfloor$ samples, where $\lfloor \cdot \rfloor$ denotes the rounding operation. This time offset partially mitigates the CFO introduced in (8). However, a residual CFO remains, which can still result in a significant error-rate performance degradation. This offset needs to be compensated to prevent a symbol lying in between two frequency bins, which increases the sensitivity to noise.

The CFO results in a phase offset $\Delta\phi = 2\pi \frac{\Delta f_c}{f_s} \cdot 2^{SF}$ between two samples with the same index in consecutive upchirps. The residual part of this offset can be estimated by taking the average across the entire symbol as

$$\hat{\Delta\phi} = \arg \left(\sum_{n=0}^{2^{SF}-1} \tilde{y}_S[n] \cdot \tilde{y}_S^*[n + 2^{SF}] \right). \quad (9)$$

We note that the estimation becomes more accurate by taking the average among all the upchirps in the preamble. The

residual CFO can be compensated as

$$\hat{y}_S[n] = \tilde{y}_S[n] \cdot e^{jn \frac{\hat{\Delta\phi}}{2^{SF}}}. \quad (10)$$

To study the effect of a CFO and the proposed compensation method, we use Monte-Carlo simulations, where we generate a LoRa signal and add the CFO according to (8). Fig. 2 shows the bit error rate (BER) of the system illustrated in Fig. 1 with $SF = 8$ and Hamming (4, 8). The offset values chosen in the simulation are 10 kHz and 10.1 kHz, which both correspond approximately to 10 ppm. As can be seen, the synchronization with the time-offset can improve the performance for both CFO values, while the residual CFO is different for them. This residual CFO is then compensated by using the algorithm as in (10), where the estimation is performed using the entire preamble. We note that the small remaining difference in comparison to the ideal system without CFO is due to the small remaining inter-symbol interference (ISI) because of the synchronization with a time-offset.

3.3. SFO Formulation and Performance Analysis

The mismatch between the transmitter and receiver oscillators also results in different sampling frequencies, and therefore, the sampled LoRa signal at the receiver experiences a sampling frequency offset. In this subsection, we study the effect of this offset on the receiver performance.

The received LoRa symbols undergo low-pass filtering with bandwidth BW and are sampled with the receiver sampling frequency $f'_s \neq BW$. We consider again the signal without noise. To describe the impact of the SFO throughout the entire frame, we introduce the LoRa symbol index $d \in \{0, 1, \dots\}$ and write the ISI free samples of the d -th symbol as

$$y_S^d[n] = \begin{cases} e^{j2\pi \left(\frac{BW}{2T_s} \left(\frac{n}{f'_s} + \Delta T_d \right)^2 + \left(S \cdot \frac{BW}{2SF} - \frac{BW}{2} \right) \left(\frac{n}{f'_s} + \Delta T_d \right) \right)} \\ e^{j2\pi \left(\frac{BW}{2T_s} \left(\frac{n}{f'_s} + \Delta T_d \right)^2 + \left(S \cdot \frac{BW}{2SF} - \frac{3BW}{2} \right) \left(\frac{n}{f'_s} + \Delta T_d \right) \right)}, \end{cases} \quad (11)$$

where $\Delta T_d = d \left(\frac{2^{SF}}{f'_s} - T_s \right)$ is the time offset of the d -th symbol, $\lfloor -\Delta T_d f'_s \rfloor < n < \lceil (T_s - \Delta T_d) f'_s \rceil$, and the conditions of each equation are pruned for notational simplicity.

Even if we ignore the ISI, ideally, the signal needs to be resampled to be able to be demodulated as in (5) and (6) due to the sampling rate mismatch between f'_s and BW . However, resampling is a complex operation. Instead, we only assume that the receiver is able to generate a reference upchirp that matches the bandwidth of the transmitter upchirp as $y'_{ref}[n] = e^{j2\pi \left(\frac{BW}{2T_s} \left(\frac{n}{f'_s} \right)^2 - \frac{BW}{2} \left(\frac{n}{f'_s} \right) \right)}$. The dechirped signal for d -th LoRa symbol is then derived by multiplying the received signal $y_S^d[n]$ with the complex-conjugate of $y'_{ref}[n]$. After some simplifications and removing the constant phase offsets we obtain

$$y_{dcS}^d[n] = \begin{cases} e^{j2\pi n \left(\frac{S}{2SF} \left(\frac{BW}{f'_s} \right) + d \left(\frac{BW^2}{f_s'^2} - \frac{BW}{f'_s} \right) \right)} \\ e^{j2\pi n \left(\left(\frac{S}{2SF} - 1 \right) \left(\frac{BW}{f'_s} \right) + d \left(\frac{BW^2}{f_s'^2} - \frac{BW}{f'_s} \right) \right)}. \end{cases} \quad (12)$$

Subsequently, the DFT of the dechirped signal is computed as

$$\begin{aligned}
X_k^d &= \text{DFT}(y_{\text{dc}_S}^d[n]) = \sum_{n=0}^{2^{\text{SF}}-1} y_{\text{dc}_S}^d[n] \cdot e^{-j2\pi n \frac{k}{2^{\text{SF}}}} \\
&= \sum_{n \in N_1^d} e^{j \frac{2\pi n}{2^{\text{SF}}} \left(S \frac{BW}{f_s'} + 2^{\text{SF}} d \left(\frac{BW^2}{f_s'^2} - \frac{BW}{f_s'} \right) - k \right)} \\
&\quad + \sum_{n \in N_2^d} e^{j \frac{2\pi n}{2^{\text{SF}}} \left((S-2^{\text{SF}}) \frac{BW}{f_s'} + 2^{\text{SF}} d \left(\frac{BW^2}{f_s'^2} - \frac{BW}{f_s'} \right) - k \right)},
\end{aligned} \quad (13)$$

where N_1^d and N_2^d are the sets of ISI free indices before and after folding for d -th symbol, respectively. For a small frequency offset such that $e^{j2\pi \frac{BW}{f_s'}} \approx 1$, the above equation is simplified to

$$X_k^d \approx \sum_{n=0}^{2^{\text{SF}}-1} e^{j \frac{2\pi n}{2^{\text{SF}}} \left(\left[S \frac{BW}{f_s'} + 2^{\text{SF}} d \left(\frac{BW^2}{f_s'^2} - \frac{BW}{f_s'} \right) \right] - k \right)}. \quad (14)$$

We can observe the effect of the SFO by considering the term $\frac{BW}{f_s'}$ in (14). We first consider the case of no offset, i.e., $\frac{BW}{f_s'} = 1$. In such a case, X_k^d will be maximized and equal to 2^{SF} for $k = S$ and zero elsewhere. However, for the case of an offset, we observe a re-scaling of the symbol location to $2\pi S \frac{BW}{f_s'}$ with respect to the receiver frequency axis (i.e., a symbol-dependent offset of $S(\frac{BW}{f_s'} - 1)$) and a frequency-offset $2\pi d(\frac{BW^2}{f_s'^2} - \frac{BW}{f_s'})$ that depends only on the index of the received symbol d . The effect of the re-scaling of the frequency axis is negligible and does not change throughout the received frame, but the frequency-offset term leads to *side-lobes* in the frequency bins other than the desired bin. These sidelobes render the decision in (6) suboptimal and lead to an increase in the noise sensitivity. The mismatch between f_s' and BW causes a drift of the sidelobes as d increases and will eventually lead to a drift of the peak in the Fourier domain into the adjacent symbol to S which results in a constant demodulation error, i.e., an error floor. Specifically, assuming that $f_s' > BW$ with no loss of generality, the sample drift occurs at the first sample n in d -th symbol that satisfies $\frac{n+1+d2^{\text{SF}}}{f_s'} < \frac{n+d2^{\text{SF}}}{BW}$. Therefore, the frames with number of symbols larger than d experience such an error floor.

Motivated by the above, we propose to discard a sample whenever half of a sample has drifted into the adjacent symbol. More specifically, the receiver can find the indices n and d that satisfy $\frac{n+1/2+d2^{\text{SF}}}{f_s'} < \frac{n+d2^{\text{SF}}}{BW}$ and discard sample n from d -th symbol. This measure *re-aligns* the symbol boundaries and prevents the accumulation of error due to a SFO. We note that the re-alignment resolution is half of a sample duration, i.e., $\frac{1}{2f_s'}$. Oversampling can improve this resolution and thus lower the error rate due to the sample drift from the adjacent symbol.

To study the effect of a SFO and the proposed compensation method, we use Monte-Carlo simulation, where a LoRa signal is generated with bandwidth BW at the transmitter and is sampled with frequency f_s' at the receiver. Fig. 3 demon-

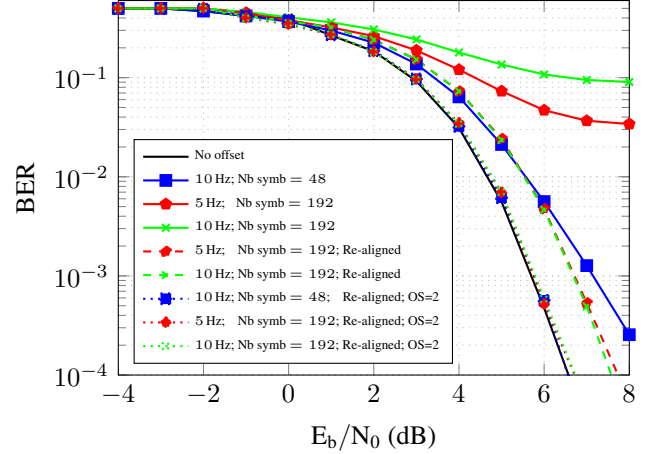


Fig. 3: BER for SFO of 5 Hz and 10 Hz.



Fig. 4: Implementation set-up and the received LoRa signal.

strates the BER of the system illustrated in Fig. 1 with $\text{SF} = 8$ and Hamming (4, 8). The offset values chosen in the simulation are 5 Hz and 10 Hz, which correspond to 20 and 40 ppm, respectively. As can be seen, the performance degradation becomes more severe by increasing the offset value as well as the number of symbols in the frame. However, re-aligning the symbol boundaries according to the proposed method prevents the error-floor in the frame with large number of symbols. Furthermore, oversampling with only a factor of 2 results in a sufficient re-alignment resolution that almost entirely compensates the performance degradation.

4. SDR IMPLEMENTATION

The LoRa transceiver was implemented using National Instrument (NI) universal software radio peripherals (USRPs) and was tested with the commercial HOPERF RFM95 LoRa radio chip. We have used NI-USRP 2920 devices to reverse-engineer and understand the LoRa PHY according to the block diagram in Fig. 1. To this end, we receive and analyze different transmitted messages from a LoRa radio that have special structures and can help to extract the parameters of the different blocks. The reversed-engineered LoRa transmitter and receiver are finally verified by decoding the LoRa frames transmitted by the commercial radio, as shown in Fig. 4.

5. CONCLUSION

In this paper, we have provided an in-depth analysis of the LoRa PHY by analytically studying the algorithmic aspects of a LoRa transceiver. We showed that the LoRa receiver is robust against the CFO by synchronizing with a time-offset to the preamble, while a residual offset needs to be compensated. Further, the SFO effect on the LoRa demodulation was modeled as an increasing phase mismatch, which can be prevented by re-aligning the symbol boundaries. Finally, we showed that LoRa can be implemented and tested on a USRP platform.

6. REFERENCES

- [1] S. Li, L. Da Xu, and S. Zhao, “5G internet of things: A survey,” *Journal of Industrial Information Integration*, 2018.
- [2] 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.
- [3] Semtech, “AN1200.22, LoRa modulation basics,” *Application Note*, 2015.
- [4] N. Sornin, M. Luis, T. Eirich, T. Kramp, and O. Hersent, “LoRaWAN specification,” *LoRa alliance*, 2015.
- [5] M. Centenaro, L. Vangelista, A. Zanella, and M. Zorzi, “Long-range communications in unlicensed bands: The rising stars in the IoT and smart city scenarios,” *IEEE Wireless Communications*, vol. 23, no. 5, pp. 60–67, 2016.
- [6] R. S. Sinha, Y. Wei, and S.-H. Hwang, “A survey on LPWA technology: LoRa and NB-IoT,” *Ict Express*, vol. 3, no. 1, pp. 14–21, 2017.
- [7] P. Robyns, P. Quax, W. Lamotte, and W. Thenaers, “A multi-channel software decoder for the LoRa modulation scheme,” 2018.
- [8] M. Knight and B. Seeber, “Decoding lora: Realizing a modern LPWAN with SDR,” in *Proceedings of the GNU Radio Conference*, vol. 1, no. 1, 2016.