

A Maximum-Likelihood-based Multi-User LoRa Receiver Implemented in GNU Radio

Mathieu Xhonneux,[†] Joachim Tapparel,^{*} Orion Afisiadis,^{*} Alexios Balatsoukas-Stimming,[‡]
and Andreas Burg^{*}

^{*}Telecommunication Circuits Laboratory, EPFL, Switzerland

[†]ICTEAM, UCLouvain, Belgium

[‡]Department of Electrical Engineering, Eindhoven University of Technology, The Netherlands

Abstract—LoRa is a popular low-power wide-area network (LPWAN) technology that uses spread-spectrum to achieve long-range connectivity and resilience to noise and interference. For energy efficiency reasons, LoRa adopts a pure ALOHA access scheme, which leads to reduced network throughput due to packet collisions at the gateways. To alleviate this issue, in this paper we analyze and implement a LoRa receiver that is able to decode LoRa packets from two interfering users. Our main contribution is a two-user detector derived in a maximum-likelihood fashion using a detailed interference model. As the complexity of the maximum-likelihood sequence estimation is prohibitive, a complexity-reduction technique is introduced to enable a practical implementation of the proposed two-user detector. This detector has been implemented along with an interference-robust synchronization algorithm on the GNU Radio software-defined radio (SDR) platform. The SDR implementation shows the effectiveness of the proposed method and also allows its experimental evaluation. Measurements indicate that our detector inherently leverages the time offset between the two colliding users to separate and demodulate their signals.

I. INTRODUCTION

With the rise of the Internet of Things (IoT), new low-power wide-area network (LPWAN) standards have emerged [1]. Unlike 3GPP cellular standards such as NB-IoT, LPWAN standards rely on low-complexity PHY and MAC layers. In the last years, LoRa has become one of the most popular LPWAN protocols [2]. At the PHY layer, LoRa uses a chirp spread-spectrum modulation that enables receivers to attain sensitivities as low as -130 dBm [3]. The spreading gain of the LoRa modulation is determined by the spreading factor (SF), which allows to trade off transmission time and data rate for coverage. For energy efficiency reasons, LoRaWAN relies on a non-slotted ALOHA multiple access scheme. As the user end-nodes are not synchronized, interference between users is very common in large-scale LoRa networks, posing a threat to the scalability of future massive LoRa systems.

Interference between LoRa users can be one of two very different types: *inter-SF interference* and *same-SF interference*. Due to spreading, the impact of inter-SF interference is different than the impact of same-SF interference. In particular, inter-SF interference can be treated as additional AWGN [4] that slightly degrades the signal-to-noise ratio (SNR). On the other hand, interfering transmissions which use the same SF lead to strong interference at the receiver as they are far from orthogonal. Nevertheless, colliding LoRa packets with the same SF exhibit a capture effect, such that the strongest

signal can often be demodulated at the expense of the weakest signal [5] if the signal-to-interference ratio (SIR) and the SNR are sufficient. Furthermore, experimental measurements have shown that for a 3 dB difference in power between two users, there is a 97% chance of demodulating the packet of the strongest user [6], but the packet of the weaker user is generally lost.

To improve the throughput of LoRa networks, several LoRa multi-user receivers have been suggested recently [7]–[10]. In [7], a successive interference cancellation (SIC) receiver that relies on the conventional single-user detector is proposed. A multi-user receiver based on non-stationary signal scaling to separate superimposed transmissions is presented in [8]. In [9] and [10], the authors suggest to first demodulate all symbols and then to assign them to their respective user based on transmitter-specific features (e.g., time offset and power). The aforementioned works have established the challenge of designing robust multi-user LoRa receivers that will allow LoRaWAN to scale by exploiting the concept of non-orthogonal multiple access (NOMA). However, a multi-user receiver that is derived based on the maximum likelihood (ML) principle, as well as its practical implementation and experimental evaluation, are still lacking in the literature.

Contributions: In this work, we present a two-user LoRa detector derived from the maximum likelihood (ML) principle. Due to the high complexity of the ML detector, we also propose complexity reduction techniques that enable a practical implementation of the receiver. The proposed detector requires estimates of the power, the carrier frequency offset (CFO), and the sampling time offset (STO) of each user. We therefore also design a synchronization algorithm capable of detecting the arrival of a new user and of estimating its parameters, even in the presence of an on-going transmission with the same SF. We implement the proposed multi-user detector and the synchronization algorithm with a software-defined radio (SDR). Our GNU Radio implementation of the multi-user LoRa receiver is publicly available at [11]. We also present a performance evaluation of the proposed detector, showing the capability of the receiver to demodulate two overlapping LoRa users with low error rates. Finally, we discuss the impact of integer and fractional sampling time offsets between the colliding users, which have an important impact on the ability to separate the two signals.

II. PRINCIPLES OF THE LoRa PHY

In this section, we provide a brief summary of the LoRa modulation, the single-user demodulation stage, and the structure of the preamble. We then explain the baseband-equivalent model of two superimposed LoRa signals.

A. Modulation and Demodulation

LoRa is a chirp spread-spectrum modulation operating in the ISM bands, with typical passband bandwidth values $B \in \{125, 250, 500\}$ kHz. LoRa symbols are chirps, i.e., signals whose instantaneous frequency increases linearly and spans the entire bandwidth. Every chirp consists of $N = 2^{\text{SF}}$ chips and carries SF bits of information, where $\text{SF} \in \{7, \dots, 12\}$ is called the *spreading factor*. When sampled at the Nyquist frequency $f_s = B$, the discrete-time baseband-equivalent of a symbol is split into N frequency steps [12], [13]. For a symbol $s \in \{0, \dots, N-1\}$, the first chip starts at an initial frequency of $(\frac{sB}{N} - \frac{N}{2})$. The instantaneous frequency of the chirp increases by $\frac{B}{N}$ in every chip. When the Nyquist frequency $\frac{B}{2}$ is attained, i.e., at chip index $n_f = N - s$, a folding to the frequency $-\frac{B}{2}$ occurs. The corresponding discrete-time baseband-equivalent equation of a LoRa symbol s can be expressed as [12], [13]

$$x_s[n] = \begin{cases} e^{j2\pi\left(\frac{n^2}{2N}\left(\frac{B}{f_s}\right)^2 + \left(\frac{s}{N} - \frac{1}{2}\right)\left(\frac{B}{f_s}\right)n\right)}, & 0 \leq n < n_f, \\ e^{j2\pi\left(\frac{n^2}{2N}\left(\frac{B}{f_s}\right)^2 + \left(\frac{s}{N} - \frac{3}{2}\right)\left(\frac{B}{f_s}\right)n\right)}, & n_f \leq n < N. \end{cases} \quad (1)$$

When the transmission takes place over an AWGN channel with a complex-valued channel gain $h \in \mathbb{C}$, the received LoRa symbol is represented by $y[n] = hx_s[n] + z[n]$, where $z[n] \sim \mathcal{CN}(0, \sigma^2)$ is complex AWGN with variance $\sigma^2 = \frac{N_0}{2N}$ and N_0 is the single-sided noise power spectral density. A receiver first performs a point-by-point multiplication of the sampled signal $y[n]$ with $x_0^*[n]$, the complex conjugate of an unmodulated symbol $s = 0$, yielding the *dechirped* signal

$$\tilde{y}[n] = y[n]x_0^*[n] = \sqrt{P}e^{j2\pi n\frac{s}{N} + \theta} + \tilde{z}[n], \quad (2)$$

where $P = |h|^2$ is the power of the signal at the receiver, $\theta = \angle h$ represents the phase introduced by the channel, and $\tilde{z}[n] = z[n]x_0^*[n]$. For a perfectly synchronized receiver, $\tilde{y}[n]$ contains a single complex tone of frequency $\frac{s}{N}$ and AWGN.

Let $Y[i] = \sum_{n=0}^{N-1} \tilde{y}[n]e^{-j2\pi\frac{ni}{N}}$ be the i -th bin from the N -point discrete Fourier transform (DFT) of the dechirped signal. In the frequency domain, the single complex tone of frequency $\frac{s}{N}$ translates into a Kronecker delta at position $i = s$. Typical receivers perform non-coherent detection by selecting the DFT bin with the largest magnitude

$$\hat{s} = \arg \max_s |Y[s]|. \quad (3)$$

B. LoRa Preamble Structure

Every LoRa packet starts with a specific preamble for synchronization purposes. The preamble consists of N_{pr} repetitions of an unmodulated upchirp (i.e., $x_0[n]$), succeeded by two symbols acting as network identifiers, and 2.25 downchirps $x_0^*[n]$. The structure of the preamble is illustrated in Fig. 1. It is well known that the combination of upchirps




Fig. 1: Structure of a LoRa preamble.




Fig. 2: Two-user scenario where user B has a time offset τ with respect to user A.

and downchirps can be exploited by a receiver to estimate the CFO and STO of a user [14], [15]. As explained in Section III, our two-user receiver also leverages this preamble in its synchronization stage.

C. Signal Model for Two Interfering Users

In this section, we describe the model of superimposed signals from two users with the same SF, namely user A and user B. Since LoRa uses a non-slotted ALOHA multiple access scheme, the users are neither synchronized among themselves nor to the gateway. Let us define $s_A^{(k)}$ and $s_B^{(k)}$ as the k -th symbols sent by users A and B, respectively. To simplify the explanation and the mathematical formulation of the multi-user detector, we assume that the gateway is perfectly synchronized in frequency and time to user A, whose packet arrives first. This can be achieved with a standard synchronization procedure [14]–[16]. We define $\tau \in [0, N)$ as the relative chip-level time offset between the first chip of a symbol transmitted by user A and the first chip of the next symbol of user B, as illustrated in Fig. 2. This offset can be split into an integer part $L_{\text{STO}} = \lfloor \tau \rfloor$ and a non-integer part $\lambda_{\text{STO}} = \tau - \lfloor \tau \rfloor$ [5].

Since the transmission of user B experiences an STO τ with respect to user A, the first $\lfloor \tau \rfloor$ samples of symbol $s_A^{(k)}$ overlap with symbol $s_B^{(k-1)}$ and the last $N - \lfloor \tau \rfloor$ samples of symbol $s_A^{(k)}$ overlap with symbol $s_B^{(k)}$. The contribution of user B to the k -th window of N samples $y^{(k)}[n]$ can therefore be split into two parts, namely $y_{B,1}^{(k)}[n]$ for $n \in \mathcal{N}_1 = \{0, \dots, \lfloor \tau \rfloor - 1\}$




Fig. 3: Finite state machine representation of the receiver. The receiver always synchronizes to the strongest user.

and $y_{B,2}^{(k)}[n]$ for $n \in \mathcal{N}_2 = \{\lceil \tau \rceil, \dots, N-1\}$ [17], with

$$y_{B,1}^{(k)}[n] = e^{j2\pi \left(\frac{(n+N-\tau)^2}{2N} + (n+N-\tau) \left(\frac{s_B^{(k-1)}}{N} - \frac{1}{2} - u[n-n_{f,1}^{(k)}] \right) \right)}, \quad (4)$$

$$y_{B,2}^{(k)}[n] = e^{j2\pi \left(\frac{(n+N-\tau)^2}{2N} + (n+N-\tau) \left(\frac{s_B^{(k)}}{N} - \frac{1}{2} - u[n-n_{f,2}^{(k)}] \right) \right)}. \quad (5)$$

Prior to synchronization, both users are affected by distinct carrier frequency offsets relative to the receiver, namely $\Delta f_{c,A}$ and $\Delta f_{c,B}$. However, since we assume the receiver to be synchronized to user A, there is a single effective CFO $\Delta f_c = \Delta f_{c,B} - \Delta f_{c,A}$ that only affects the signal from user B.

Finally, the users have different transmit powers and experience independent channels h_A and h_B . Thus, the baseband-equivalent model of the sampled signal contained in the k -th window of N samples is

$$y^{(k)}[n] = h_A x_{s_A^{(k)}}[n] + z[n] + \begin{cases} h_B c[n] y_{B,1}^{(k)}[n], & n \in \mathcal{N}_1, \\ h_B c[n] y_{B,2}^{(k)}[n], & n \in \mathcal{N}_2, \end{cases} \quad (6)$$

where $c[n] = e^{j2\pi n \frac{\Delta f_c}{f_s}}$ is the effective CFO affecting user B.

III. A PRACTICAL ML-DERIVED MULTI-USER RECEIVER

As previously explained, a LoRa user can start or stop transmitting at any time during the transmission of another user. The gateway must hence be capable of tracking the arrival and departure of successive users. To this end, the proposed receiver implements a three-state finite state machine (FSM) which is illustrated in Fig. 3. Each state corresponds to the current number of colliding users, up to two.

In the no-user or single-user states, a synchronization algorithm is constantly running to detect the preamble of a potential arriving user. When a new user is detected, the algorithm estimates its parameters, i.e., its CFO, STO, and power. For the first user, the receiver synchronizes in time and frequency to this user, and the non-coherent detector of (3) is used to demodulate its symbols. Upon the arrival of a second user, the receiver switches to a two-user detection rule capable of jointly demodulating both users. Our two-user detector requires the receiver to re-synchronize to the strongest

user. It also requires the parameter estimation of the second user in the presence of interference. In the following, we first explain the parameter estimation and synchronization stages and then discuss the two-user detection.

A. Robust Synchronization in the Presence of Interference

Contrary to the algorithms of [15] and [16], the synchronization stage of our two-user receiver has to estimate the parameters of a new user even in the presence of a colliding user. As explained in [14], the estimation of the integer offsets L_{CFO} and L_{STO} requires the demodulation of an upchirp and a downchirp in the preamble. However, the conventional single-user detector is typically not able to correctly demodulate these symbols in the presence of strong interference.

Hence, to perform a robust estimation of the CFO and STO, we propose instead to leverage the repetition of the upchirps and the almost orthogonal relationship between the 2.25 downchirps in the preamble of the second user and the modulated symbols of the first user [18]. The proposed algorithm implements the following steps:

- 1) Instead of demodulating each upchirp of the preamble separately, the magnitudes of the DFTs of $N_{pr} - 1$ consecutive windows of N samples are averaged with a geometric mean.
- 2) The largest bin of the averaged DFTs is selected as demodulated upchirp. The algorithm uses this demodulated value to estimate λ_{CFO} over the $N_{pr} - 1$ upchirps with the estimator from [14]. The fractional CFO is corrected for the subsequent steps.
- 3) The demodulation of the two downchirps is performed with a cross-correlation of length $2NR$ on an over-sampled version of the received signal, where R is the oversampling factor. The polyphase that yields the largest output is used as an estimate of λ_{STO} . The magnitude of the largest output is used to estimate the received power of the user (including the channel).
- 4) The demodulated aggregated upchirp and the downchirps are used to estimate L_{CFO} and L_{STO} , as described in [14].

If the second detected user is the strongest one, the receiver re-synchronizes to this user by selecting the closest polyphase of the oversampled signal and correcting the estimated CFO.

B. Maximum-Likelihood-based Two-User Detector

We now present the detector used to jointly demodulate the symbols from two superimposed LoRa users. The proposed detection rule is derived from the maximum likelihood sequence detector for the signal model given by (6).

Without loss of generality, we assume that both users send a frame of M symbols, and that the receiver is synchronized to user A. This two-user M -symbol asynchronous signal can be decomposed into M successive windows of N samples, and a final window of $\lceil \tau \rceil$ samples. Since every symbol from user B is contained in two consecutive windows, the joint maximum likelihood decision of all symbols $\{s_A^{(k)}, s_B^{(k)}\}$ for $k \in \{0, \dots, M-1\}$ implies to process all $M+1$ windows at once. The maximum likelihood sequence can be computed

with reduced complexity using the Viterbi algorithm [19]. However, since there are N^2 possible states and N^3 possible state transitions, the complexity of this maximum likelihood sequence detector is still prohibitive, especially for large SFs (i.e., large N).

To avoid such a complex receiver, we suggest to move from a joint demodulation of all symbols to individual decisions bounded to a single window of N samples. For each window k and the corresponding received signal $y^{(k)}$, an individual decision amounts to detecting the overlapping symbols $\mathbf{s}^{(k)} = \{s_A^{(k)}, s_B^{(k-1)}, s_B^{(k)}\}$. Let P_A and P_B represent the powers of user A and user B at the receiver, respectively. Similarly, we define $\theta_A^{(k)}$ and $\theta_B^{(k)}$ as the initial phase of the symbols $s_A^{(k)}$ and $s_B^{(k)}$ at the receiver. It can be shown that¹ the individual ML detection for a single window is achieved by selecting the symbols $\hat{\mathbf{s}}^{(k)}$ that maximize the following metric

$$\Lambda_{\text{ML}}(\bar{\mathbf{s}}^{(k)}) = \exp \left[\sqrt{P_A} \Re \left(e^{-j\theta_A^{(k)}} Y^{(k)} \left[\bar{s}_A^{(k)} \right] \right) + \sqrt{P_B} \Re \left(e^{-j\theta_B^{(k-1)}} M_1^{(k)} \left[\bar{s}_A^{(k)}, \bar{s}_B^{(k-1)} \right] \right) + \sqrt{P_B} \Re \left(e^{-j\theta_B^{(k)}} M_2^{(k)} \left[\bar{s}_A^{(k)}, \bar{s}_B^{(k)} \right] \right) \right], \quad (7)$$

where $\bar{\mathbf{s}}^{(k)} = \{\bar{s}_A^{(k)}, \bar{s}_B^{(k-1)}, \bar{s}_B^{(k)}\}$ are the candidate symbols evaluated by the ML criterion, and $Y^{(k)}$, $M_1^{(k)}$ and $M_2^{(k)}$ correspond to three matched filters that are connected to each of the three symbols. It is worth noting that this detector requires knowledge of the offsets τ and Δf_c , the phases $\{\theta_A^{(k)}, \theta_B^{(k)}\}$ and the received powers of each user.

Since the receiver is synchronized to user A, the symbol $s_A^{(k)}$ translates to a Kronecker delta in the frequency domain. The matched filter $Y^{(k)}$ hence corresponds to the DFT of the dechirped signal $\tilde{y}^{(k)}[n]$. The two remaining terms in (7) evaluate the contribution from user B, which is not synchronized. As the symbols of this user lie partially in a different signal space than $s_A^{(k)}$, the individual ML detector uses two specific matched filters $M_1^{(k)}$ and $M_2^{(k)}$ which depend on the offsets Δf_c and τ . For both matched filters, we first eliminate the presumed contribution $\sqrt{P_A} e^{j\theta_A^{(k)}} e^{j2\pi \frac{\tau}{N} \bar{s}_A^{(k)}}$ of user A, and then compute a partial DFT over the first $\lceil \tau \rceil$ and the remaining $N - \lceil \tau \rceil$ samples in the window, respectively

$$M_1^{(k)} \left[\bar{s}_A^{(k)}, \bar{s}_B^{(k-1)} \right] = \sum_{n=0}^{\lceil \tau \rceil - 1} \left(\tilde{y}^{(k)}[n] - \sqrt{P_A} e^{j\theta_A^{(k)}} e^{j2\pi \frac{\tau}{N} \bar{s}_A^{(k)}} \right) \cdot e^{-j2\pi \frac{\tau}{N} \left(\bar{s}_B^{(k-1)} - \tau + N \frac{\Delta f_c}{f_s} \right)} e^{j2\pi \tau u \left[n - n_{f,1}^{(k)} \right]}, \quad (8)$$

$$M_2^{(k)} \left[\bar{s}_A^{(k)}, \bar{s}_B^{(k)} \right] = \sum_{n=\lceil \tau \rceil}^{N-1} \left(\tilde{y}^{(k)}[n] - \sqrt{P_A} e^{j\theta_A^{(k)}} e^{j2\pi \frac{\tau}{N} \bar{s}_A^{(k)}} \right) \cdot e^{-j2\pi \frac{\tau}{N} \left(\bar{s}_B^{(k)} - \tau + N \frac{\Delta f_c}{f_s} \right)} e^{j2\pi \tau u \left[n - n_{f,2}^{(k)} \right]}. \quad (9)$$

Yet, the matched filters $M_1^{(k)}$ and $M_2^{(k)}$ provide only partial information on the symbols $s_B^{(k-1)}$ and $s_B^{(k)}$, as illustrated in

Fig. 2. Detecting these symbols using only the information from the k -th window is clearly suboptimal. To efficiently demodulate the symbols $s_B^{(k-1)}$ and $s_B^{(k)}$, the detector must also use the partial information provided by the matched filters $M_2^{(k-1)}$ and $M_1^{(k+1)}$ of the preceding and following windows, respectively. We hence suggest to estimate $s_A^{(k)}$ by marginalizing over all candidates for $s_B^{(k-1)}$ and $s_B^{(k)}$ independent of the previous and next window. We then decide on $s_B^{(k-1)}$ based on the prior decisions on $s_A^{(k-1)}$ and on the latest decision on $s_A^{(k)}$ as follows:

$$\hat{s}_A^{(k)} = \max_{\bar{s}_A^{(k)}} \max_{\bar{s}_B^{(k-1)}, \bar{s}_B^{(k)}} \Lambda_{\text{ML}}(\bar{s}_A^{(k)}, \bar{s}_B^{(k-1)}, \bar{s}_B^{(k)}), \quad (10)$$

$$\hat{s}_B^{(k-1)} = \max_{\bar{s}_B^{(k-1)}} \left[M_2^{(k-1)} \left[\hat{s}_A^{(k-1)}, \bar{s}_B^{(k-1)} \right] + M_1^{(k)} \left[\hat{s}_A^{(k)}, \bar{s}_B^{(k-1)} \right] \right]. \quad (11)$$

The decision on $s_B^{(k)}$ is deferred to the next time step. This rule requires that the receiver keeps in memory the vector containing the N matched filter outputs $M_2^{(k-1)} \left[\hat{s}_A^{(k-1)}, \bar{s}_B^{(k-1)} \right]$ of the previous window.

The individual ML detection rule given in (7) requires the knowledge of the phases $\theta_A^{(k)}$, $\theta_B^{(k-1)}$ and $\theta_B^{(k)}$. In practical systems, a residual CFO or a drifting STO often arise and modify the initial phases of the symbol during the transmission of the packet. To avoid a continuous tracking of the phase of each user, we further propose to marginalize over all three phase terms in (7). The marginalization of $\theta_A^{(k)}$ is however challenging, as it is used in all three parts of $\Lambda_{\text{ML}}(\bar{\mathbf{s}}^{(k)})$. In a single-user scenario, the phase θ of a demodulated symbol $\hat{s}$ can be estimated by using the phase of the DFT bin $\hat{s}$, i.e., $\hat{\theta} = \arctan(Y[\hat{s}])$ [5]. We propose to use the same estimator with the DFT $Y^{(k)}$ to obtain an estimate $\hat{\theta}_A^{(k)}$ of the initial phase of the candidate symbol $\bar{s}_A^{(k)}$. Let $\tilde{\Lambda}(\bar{\mathbf{s}}^{(k)})$ be a modified version of the function $\Lambda_{\text{ML}}(\bar{\mathbf{s}}^{(k)})$ where the variable $\theta_A^{(k)}$ is replaced by the estimate $\hat{\theta}_A^{(k)}$ in the expression of the matched filters $M_1^{(k)}$ and $M_2^{(k)}$. The remaining occurrence of $\theta_A^{(k)}$ in $\tilde{\Lambda}(\bar{\mathbf{s}}^{(k)})$ is marginalized along with $\theta_B^{(k-1)}$ and $\theta_B^{(k)}$, yielding

$$\begin{aligned} \Lambda(\bar{\mathbf{s}}^{(k)}) &= \iiint_{-\pi}^{\pi} \tilde{\Lambda}(\bar{\mathbf{s}}^{(k)}) d\theta_A^{(k)} d\theta_B^{(k-1)} d\theta_B^{(k)} \\ &= I_0 \left(\sqrt{P_A} \left| Y^{(k)} \left(\bar{s}_A^{(k)} \right) \right| \right) \cdot \\ &I_0 \left(\sqrt{P_B} \left| M_1^{(k)} \left(\bar{s}_A^{(k)}, \bar{s}_B^{(k-1)} \right) \right| \right) \cdot \\ &I_0 \left(\sqrt{P_B} \left| M_2^{(k)} \left(\bar{s}_A^{(k)}, \bar{s}_B^{(k)} \right) \right| \right), \end{aligned} \quad (12)$$

where $I_0(x)$ is the first order modified Bessel function of the first kind. This function is akin to e^x , and has a practical closed-form expression. The accuracy of the estimate $\hat{\theta}_A^{(k)}$ is increased when $P_A > P_B$, i.e., when the synchronized user is the strongest user. This behavior is the principal motivation for the receiver to synchronize to the strongest user in the two-user state.

¹We skip the detailed derivation due to space constraints.




Fig. 4: Multi-user testbed with two USRPs as Tx and one USRP as Rx.

IV. SOFTWARE-DEFINED RADIO IMPLEMENTATION

The two-user detector and the synchronization algorithm previously described have been implemented on the GNU Radio software-defined radio platform. We use this implementation in a testbed to experimentally assess the error rate performance of the proposed receiver. The implementation of our multi-user receiver is open-source and available at [11].

A. Architecture of the Implementation

The implementation of the two-user receiver is split into three stages. The first stage performs the preamble detection and parameter estimation of new users on an oversampled signal, following the algorithm described in Section III. Upon detection of a new user, the preamble detection stage provides the estimated CFO, STO, and received power of the user to the synchronization stage. The synchronization stage implements the FSM of Fig. 3. This stage stores in memory the parameters of the current users and synchronizes the receiver to the strongest user. The synchronized signal is split into windows of N samples, and each window is fed to the demodulation stage. In the presence of two users, the receiver jointly demodulates both users using the metric $\Lambda(\bar{s}^{(k)})$ of (12), and otherwise uses the non-coherent single-user detector of (3).

B. Testbed Description

The testbed uses three National Instruments (NI) 2920 USRP transceivers, with two acting as transmitters and the third one implementing the two-user receiver. The transmission of the LoRa frames is implemented using the open-source GNU Radio LoRa prototype described in [16], which has been tested to be compatible with commercial LoRa devices. The nominal carrier frequency used by both transmitters is 868MHz. To avoid interference from other sources in the 868MHz band, the transmitters are connected to the multi-user receiver with an RF combiner. Two -30 dB attenuators are inserted between the combiner and the receiver to attain the low SNR regions of interest. Both transmit USRPs share the same reference clock, which enables GNU Radio to define the STO τ and the effective CFO Δf_c between the transmitters.




Fig. 5: Simulation and experimental SERs of two users with $\tau = 64.0$, $\Delta f_c = 0$, $P_A - P_B = 3$ dB and SF = 7.

The CFO and STO between the transmitters and the receiver are not controlled in the testbed, but are corrected by the receiver when it synchronizes to the strongest transmitter.

V. PERFORMANCE EVALUATION

In the following, we provide measurements to evaluate the performance of the GNU Radio multi-user receiver using the testbed described in Section IV. The experiments yield per-user symbol error rates (SER) when the receiver jointly demodulates two superimposed users. While many parameters influence the SERs, in this paper we only focus on the impact of the STO τ between the users due to space constraints.

The results are obtained by creating collisions between the two transmit USRPs. In each experiment, both transmitters send each one LoRa packet of $N_P = 32$ random symbols to the receiver, with SF = 7. The starting time of the transmission of the second user is delayed by 15 symbols and τ samples with respect to the first user. The transmit power of the second user is always 3 dB stronger than the power of the first user. The effective CFO Δf_c between the transmitters is set to zero. The receiver samples the received signal with an oversampling factor $R = 8$. The SER vs SNR curves are obtained by sweeping through the transmit gains of the transmit USRPs. A total of 20 000 experiments are performed for each SNR level. An experiment is considered to be valid if both packets are detected by the receiver and the estimated power of the second user exceeds the estimated power of the first user. Invalid experiments are not taken into account in the SER evaluation. Finally, only the 15 overlapping payload symbols are used for the computation of the SERs.

Fig. 5 shows the experimental SERs of both users along with Monte-Carlo simulation results for an STO between the users of $\tau = 64.0$ samples. The simulation assumes perfect synchronization, i.e., the received power, CFO, and STO of each user are known. The SER differences between the simulation and experimental results mainly illustrate the




Fig. 6: Experimental SERs of two users for different values of τ , with $\Delta f_c = 0$, $P_A - P_B = 3\text{dB}$ and $\text{SF} = 7$.

non-idealities of the synchronization stage. We observe that for both users and in all SNR regimes, there is a loss of approximately 1 dB between the simulation and the SDR implementation. This result indicates that the synchronization algorithm is capable of estimating the parameters of both users.

We subsequently compare three scenarios with $\tau = 16.0$, $\tau = 16.5$ and $\tau = 64.0$ to study the influence of the STO on the demodulation. The experimental SERs in all three scenarios are shown in Fig 6. The case $\tau = 16.0$ is used as a baseline for the following discussion. Regarding first the impact of the integer offset L_{STO} , we clearly observe that a larger integer STO (up to $\frac{N}{2}$) reduces the SERs of both users. For the strongest user, the required SNR to attain a 10^{-3} SER, is 4 dB lower for $\tau = 64.0$ compared to $\tau = 16.0$. A similar behavior can be observed for the fractional part of the STO. With $\lambda_{\text{STO}} = 0.5$, the weakest user reaches a SER of 10^{-3} at -5 dB SNR, whereas for $\lambda_{\text{STO}} = 0$ the SER is much worse and even levels off an error floor around a SER of 10^{-3} .

The observed behavior illustrates that the more the interfering users are desynchronized in time, the easier it is for the two-user receiver to separate and demodulate them. This effect can be explained by the contribution of each user to the DFT of the dechirped signal. While the contribution of the strongest and synchronized user is always a Kronecker delta, the signal space of the contribution of the weakest user depends on the STO τ . For $\tau = 0$, both users share the same signal space and the matched filter $M_1^{(k)}$ is identical to the DFT $Y^{(k)}$. In this case, it is difficult to distinguish the symbols of the two users. In the presence of an STO $\tau \neq 0$, the contribution of the weakest user is no longer a single peak but a bell-shaped function scattered across several DFT bins [5]. The maximum likelihood detector inherently leverages both parts of the STO to separate the contribution of each user. Therefore, integer or fractional STOs close to $\frac{N}{2}$ or 0.5, respectively, improve the

performance of the proposed two-user detector.

VI. CONCLUSION

Multi-user receivers are required to overcome the scalability limitations of LoRa networks. In this paper, we presented a receiver able to demodulate two colliding LoRa users. This receiver is derived from the maximum likelihood two-user detector and relies on a novel synchronization algorithm robust to interference. Instead of resorting to a costly maximum likelihood sequence estimation, we propose a complexity reduction technique that enables the receiver to process one symbol at a time. The proposed synchronization algorithm and two-user detector have been implemented on the GNU Radio SDR platform. Experimental measurements show that the detector inherently leverages the time offset between the two interfering users to separate and demodulate the contribution of each user.

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.
- [2] J. Haxhibeqiri *et al.*, "A survey of LoRaWAN for IoT: From technology to application," *Sensors*, vol. 18, no. 11, p. 3995, 2018.
- [3] A. Augustin *et al.*, "A study of LoRa: Long range & low power networks for the Internet of Things," *Sensors*, vol. 16, no. 9, p. 1466, 2016.
- [4] D. Croce, M. Gucciardo, S. Mangione, G. Santaromita, and I. Tinnirello, "LoRa technology demystified: From link behavior to cell-level performance," *IEEE Transactions on Wireless Communications*, vol. 19, no. 2, pp. 822–834, 2019.
- [5] O. Afisiadis, M. Cotting, A. Burg, and A. Balatsoukas-Stimming, "On the error rate of the LoRa modulation with interference," *IEEE Transactions on Wireless Communications*, vol. 19, no. 2, pp. 1292–1304, 2019.
- [6] R. Fernandes, R. Oliveira, M. Luís, and S. Sargento, "On the real capacity of LoRa networks: the impact of non-destructive communications," *IEEE Communications Letters*, vol. 23, no. 12, pp. 2437–2441, 2019.
- [7] M. A. B. Temim *et al.*, "An enhanced receiver to decode superposed LoRa-like signals," *IEEE Internet of Things Journal*, 2020.
- [8] S. Tong, J. Wang, and Y. Liu, "Combating packet collisions using non-stationary signal scaling in LPWANs," in *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*, 2020, pp. 234–246.
- [9] B. Hu, Z. Yin, S. Wang, Z. Xu, and T. He, "SCLoRa: Leveraging multi-dimensionality in decoding collided LoRa transmissions," in *IEEE 28th International Conference on Network Protocols*. IEEE, 2020, pp. 1–11.
- [10] X. Xia, Y. Zheng, T. Gu *et al.*, "FTrack: Parallel decoding for LoRa transmissions," *IEEE/ACM Transactions on Networking*, 2020.
- [11] "Multi-User LoRa repository," <https://www.epfl.ch/labs/tcl/resources-and-sw/lor-multi-user-receiver/>.
- [12] M. Chiani and A. Elzanaty, "On the LoRa modulation for IoT: Waveform properties and spectral analysis," *IEEE Internet of Things Journal*, vol. 6, no. 5, pp. 8463–8470, 2019.
- [13] R. Ghanaatian, O. Afisiadis, M. Cotting, and A. Burg, "LoRa digital receiver analysis and implementation," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 1498–1502.
- [14] C. Bernier, F. Dehmas, and N. Deparis, "Low complexity LoRa frame synchronization for ultra-low power software-defined radios," *IEEE Transactions on Communications*, vol. 68, no. 5, pp. 3140–3152, 2020.
- [15] M. Xhonneux, D. Bol, and J. Louveaux, "A low-complexity synchronization scheme for LoRa end nodes," *arXiv preprint 1912.11344*, 2019.
- [16] J. Tapparel, O. Afisiadis, P. Mayoraz, A. Balatsoukas-Stimming, and A. Burg, "An open-source LoRa physical layer prototype on GNU Radio," *IEEE International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, 2020.
- [17] O. Afisiadis *et al.*, "On the advantage of coherent LoRa detection in the presence of interference," *arXiv preprint arXiv:2010.00507*, 2020.
- [18] Z. Xu *et al.*, "FlipLoRa: Resolving collisions with up-down quasi-orthogonality," in *17th Annual IEEE International Conference on Sensing, Communication, and Networking*. IEEE, 2020, pp. 1–9.
- [19] S. Verdú *et al.*, *Multiuser detection*. Cambridge university press, 1998.