

Received August 2, 2021, accepted August 18, 2021, date of publication August 26, 2021, date of current version September 22, 2021.

Digital Object Identifier 10.1109/ACCESS.2021.3108021

Interference Cancellation for LoRa Gateways and Impact on Network Capacity

DOMENICO GARLISI¹, STEFANO MANGIONE², (Member, IEEE), FABRIZIO GIULIANO¹, DANIELE CROCE¹, GIOVANNI GARBO¹, AND ILENIA TINNIRELLO¹

Department of Engineering, University of Palermo, 90133 Palermo, Italy
Consorzio Nazionale Inter-universitario per le Telecomunicazioni (CNIT), 43124 Parma, Italy

Corresponding author: Domenico Garlisi (domenico.garlisi@unipa.it)

This work was supported in part by the Research Project Support, Alerting, Early Warning (SAWE) through Piano Organizzativo (PO) Fondo Europeo di Sviluppo Regionale (FESR) Sicilia 2014–2020 under Grant 08PA9511000101 (Codice Identificativo Progetto (CUP): G78I18000960007).

ABSTRACT In this paper we propose LoRaSyNc (LoRa receiver with SyNchronization and Cancellation), a second generation LoRa receiver that implements Successive Interference Cancellation (SIC) and time synchronization to improve the performance of LoRa gateways. Indeed, the chirp spread spectrum modulation employed in LoRa experiences very high capture probability, and cancelling the strongest signal in case of collisions can significantly improve the cell capacity. An important feature of LoRaSyNc is the ability to track the frequency and clock drifts between the transmitter and receiver, during the whole demodulation of the interfered frame. Due to the use of low-cost oscillators on end-devices, a signal cancellation scheme cannot result accurate without such a tracking, especially at the lower data rates. We validate the performance of LoRaSyNc in presence of collisions by implementing a receiver prototype on software-defined-radios, and perform several experiments in different realistic scenarios, by also comparing our receiver with commercial gateways. Finally, we simulate a cell deployment with one or more gateways, showing that the proposed scheme improves performance by almost 50% compared to a traditional receiver.

INDEX TERMS Capture effect, interference, interference cancellation, LoRa, LoRaWAN, LPWAN, scalability, spreading factor, synchronized signals.

I. INTRODUCTION

In the last years, we have assisted to a steep increase of Internet of Things (IoT) applications and devices. According to the IoT Analytics forecast,¹ the IoT market has seen an acceleration since 2018: the number of IoT devices is expected to be higher than 22 billion by 2025. In many application scenarios, the smart objects require low-rate and low-power wireless technologies to be connected to the Internet. Therefore, the design of efficient solutions for Low-Power Wide Area Networks (LPWANs) is an important challenge. Currently, common LPWAN technologies include Long Range Wide Area Network (LoRaWAN), Sigfox, and NB-IoT [1], which are widely used in smart agriculture, smart industry

and smart city environments, just to name a few application areas.

In this work, we consider the LoRaWAN technology and associated LoRa modulation, a typical example of wireless IoT technology designed for connecting simple devices in unlicensed bands. LoRa modulation has been demonstrated to be very robust to interference, while the LoRaWAN MAC protocol is simple and energy efficient. Indeed, LoRaWAN is based on a simple star topology and on the Aloha MAC protocol, where a central gateway (GW) collects frames randomly transmitted by smart devices in its coverage area, and devices can sleep for long intervals.

Because of the simple MAC protocol employed, when the number of devices active in a particular region grows significantly, the capacity of the network will quickly deteriorate. This phenomenon is critical despite the availability of multiple non-interfering channels, which can be implemented by allocating different carrier frequencies or modulation parameters (called spreading factors) to devices [2]. Indeed, optimal

¹The associate editor coordinating the review of this manuscript and approving it for publication was Yan Huo.

¹<https://iot-analytics.com/state-of-the-iot-update-q1-q2-2018-number-of-iot-devices-now-7b/>

network planning is difficult in unlicensed bands and node density can be so high to saturate all the available channels.

In this paper, we consider the problem of optimizing the LoRaWAN network capacity by improving the receiver architecture of the GWs, in order to enable the parallel reception of multiple colliding signals. Indeed, experimental studies of receiver performance have demonstrated that different Spreading Factors (SFs) are not perfectly orthogonal [3]. On the other side, collisions between signals modulated at the same SF often result in the correct demodulation of the strongest signal: the demodulation of one colliding signal can be successful even when the signal strength differences are as little as 1 dB [4], provided that the strongest signal is transmitted first or within a limited interval from the start of the weakest signal. This phenomenon, called channel capture, can significantly improve the network capacity [5].

Starting from these considerations, we present a new LoRa receiver architecture, named LoRaSyNc (LoRa receiver with Synchroization and Cancellation), that extends the demodulation procedure in case of collisions. By adding a synchronization and clock tracking scheme during the whole reception of the frame, LoRaSyNc is able to perform signal cancellation in case the strongest frame is correctly received, and can also recover the weakest frame. A few other studies have considered the possibility of implementing signal cancellation mechanisms in LoRa (e.g., [6], [7]). Different from these studies, we consider the impact of deploying multiple SIC-enabled (Successive Interference Cancellation) GWs on the performance of a LoRaWAN network. Moreover, the proposed LoRaSyNc receiver is accompanied by a clock drift tracking scheme which, as we will show, is fundamental to keep the residual cancellation noise at a minimum and thus improve performance when working with low-cost oscillators (typical of LoRaWAN devices). Additionally, the proposed SIC strategy can be used in combination with other Non Orthogonal Multiple Access (NOMA) strategies, such as [8], [9]. The design, implementation and validation of LoRaSyNc is performed both in simulation and in real experiments, using USRP Software Defined Radios (SDRs) [10] to generate repeatable and controlled frame collisions and to run the LoRaSyNc receiver. After a thorough analysis of LoRaSyNc's characteristics, we simulate a LoRa network in presence of one or more LoRaSyNc GWs and show that, compared to traditional receivers, the proposed scheme can increase the system capacity by almost 50%. The contributions presented in this paper are summarized as follows:

- we present LoRaSyNc, an accurate SIC scheme with precise time synchronization and clock tracking mechanisms;
- we perform a thorough evaluation of the proposed system exploiting both simulations and experiments using a USRP testbed;
- we analyze the capacity gains achievable in a LoRa cell, considering also the case of multiple GWs covering the same area.

The rest of this paper is organized as follows. Section II briefly reviews some related work and Section III introduces the LoRaWAN architecture and the relevant features of the Chirped Spread Spectrum (CSS) modulation used in LoRa. In Section IV, we present the LoRaSyNc demodulator, comprised of the clock tracking module and the SIC algorithm. In Section V we analyze the performance of our receiver in controlled colliding events, by also comparing results with a commercial GW (able to demodulate at most one frame in case of collisions). Simulation-based results on system capacity are presented in Section VI. Finally, conclusions are drawn in Section VII.

II. RELATED WORK

A. PERFORMANCE OF LoRa

In the last years several studies have been published on the evaluation of LoRa link-level or cell-level performance [5]. A general description of LoRaWAN applications and implications for IoT scenarios can be found in [11]. In [12] the performance of LoRa links is compared with ultra-narrowband (Sigfox-like) networks, concluding that ultra-narrowband has a larger coverage but LoRa modulation are less sensitive to interference. The LoRa modulation is presented in [13], where the authors also describe an implementation of a LoRa transceiver, called *gr-lora*, based on GNU Radio. Also, the description of an improved LoRa receiver is provided in [14].

Link-level studies of the LoRa PHY are mostly based on the experimental characterization of coverage and interference rejection capabilities [15]. The power reception thresholds for different SFs and the Signal to Interference Ratio (SIR) required for rejecting interfering LoRa signals are quantified in [3]. In [4] it is also shown that, because of the imperfect orthogonality between different SFs, a LoRa network cannot be studied as a simple super-position of independent networks working on independent SFs or channels. Indeed, [5] shows that non-orthogonality of the SFs can deteriorate significantly the performance, especially of higher SFs, and show that capture effects can significantly improve the performance of LoRa networks. In this paper, we focus on capture effects of frames having the same SF, confirming channel rejection thresholds presented in previous works, and exploiting captures to develop an accurate cancellation procedure in order to recover overlapped frames.

B. INTERFERENCE CANCELLATION

One of the pioneering works on interference cancellation is [16], which proposes the concept of cancellation to support concurrent transmissions and improve MAC capabilities. A proof-of-concept experiment using USRP showed that interference cancellation is feasible in practice, but signal processing algorithms need to be adjusted to this new model. Regarding the LoRa technology, a mathematical model that takes into account the capture effect is presented in [17]. The preliminary work in [18] provides two algorithms for

decoding overlapped LoRa signals that are slightly desynchronized and completely synchronized respectively. Through numerical results, the authors prove that throughput improves significantly for the first case, while the second algorithm is able to decode only the strongest signal. More specifically, concerning the slightly desynchronized case, the proposed algorithm solves the decoding issue of two overlapped frames in the following way: first, for preamble detection, when an overlap occurs the receiver recognizes that two signals are being received and is able to deduce the symbol frontiers for both signals; then, for the data decoding, at each frontier the receiver updates the previous frequencies and compares these new frequencies with the previous. Based on this comparison, the receiver knows if a new symbol has started and which of the transmitters it belongs to. Our approach is different because in LoRaSyNc we rebuild the frame with the strongest power and subtract it from the whole received trace. Moreover, [18] shows only numerical results, while we validate the proposed technique also through experimental results using USRPs. In [19] and [20] the scalability of a LoRa network is studied considering the impact of the capture effects. These results could significantly be improved considering also the interference cancellation scheme proposed in this paper. A preliminary study in this direction is presented in [21] where some simulations on the effectiveness of SIC are reported. In this paper, we extend this study with a thorough analysis of capture effects and the impact of SIC on the capacity of a LoRa cell deployment with one or more GWs. A few other SIC algorithms have also been presented in the literature, such as [6], [7]. In [6], the mLoRa protocol is presented to decode multiple collided frames, and experiments on a USRP testbed have demonstrated that up to three concurrent transmissions can be simultaneously received. Similarly, the work in [7] proposes a receiver capable of processing LoRa collisions using commercialized LoRa chips. Although these studies present good results in realistic settings, none of them draw conclusions on the performance that such schemes could achieve in a scenario with thousands of active end-devices and many GWs covering the same area. In this paper, instead, we study the impact of deploying multiple SIC GWs on the performance of a LoRa cell. Moreover, the proposed LoRaSyNc receiver is accompanied with a precise clock drift tracking scheme which, as we will show, is fundamental to keep the residual cancellation noise at a minimum and thus improve performance. Indeed, clock drifts are significant in low-cost commercial devices, especially with high SF values, and this problem is not analyzed in [6] and [7]. Finally, [9] and [8] leverage power, frequency and/or time differences to develop NOMA schemes without the need of a SIC receiver. In particular, FTrack [9] is a scheme to resolve LoRa collisions in both time and frequency domains. FTrack is implemented in SDRs and the performance improves of LoRaWAN is shown in real testbeds. Similarly, in [8] time and power differences are exploited to distinguish between LoRa transmissions belonging to different EDs. Although




FIGURE 1. LoRaWAN architecture.

promising, these techniques are orthogonal to the LoRaSyNc approach presented in this paper and could be included in future works to further improve the performances of an enhanced LoRa GW.

III. BACKGROUND

A. LoRaWAN NETWORK ARCHITECTURE

LoRaWAN is a standard for IoT applications based on the LoRa modulation and promoted by the LoRa Alliance [22]. Fig. 1 shows the main system components:

- 1) the End Devices (ED), which are the low-power smart objects, such as sensors and actuators, employing the LoRa modulation for communicating towards a GW;
- 2) the Gateway (GW), which forwards frames received on the wireless channel from EDs to a Network Server through an IP backhaul network (and vice-versa). Multiple GWs can be placed in the same geographical area to improve coverage and cell capacity.
- 3) the Network Server (NS), which is responsible for decoding and de-duplicating the frames received by the GWs and delivering them to the relevant Application Servers. Application Servers can belong to external service operators and can offer heterogeneous services on the same network infrastructure. The NS can also send frames back to the EDs in case of downlink traffic or for enforcing some specific device configurations.

Different from traditional cellular networks, EDs are not associated with a particular GW. Instead, GWs serve simply as link layer relayers and forward all the packets received from EDs to the NS, adding information regarding the reception quality. Thus, the same packet can be received and forwarded by multiple GWs and the NS is responsible for detecting duplicate packets and choosing the appropriate GW for transmitting downlink packets (if any). GWs are usually equipped with multiple transceivers for receiving simultaneously on multiple (configurable) frequency channels, while EDs can dynamically select a transmission channel, from among those available, at each transmission attempt. Some channels are reserved for data transmission, one channel is reserved for GW's responses, while some other channels are usually reserved for control information and in particular for transmitting network join requests from EDs.

For accessing the wireless channel, LoRaWAN defines a MAC layer based on a simple Aloha protocol. This choice is devised to minimize the protocol complexity and the energy




FIGURE 2. Example chirps with SF = 9. Instantaneous frequency for one basic upchirp and three symbols: 128, 256 and 384.

consumption of EDs. Moreover, EDs listen to the medium for receiving frames only in special time windows after uplink transmissions (class A devices), or at regular time intervals (class B devices). Only class C devices have continuously active receivers. LoRaWAN operates in the unlicensed ISM (Industrial, Scientific and Medical) radio bands that are available worldwide and a limited duty cycle is allowed to EDs: depending on the channel, EDs are granted a duty cycle of 0.1%, 1.0% and 10% per day.

B. LoRa MODULATION AND FRAMING

LoRa implements a Chirp Spread Spectrum (CSS) modulation. CSS modulation has been demonstrated to be very robust against in-band or out-band interference, which can be very critical when operating in ISM bands. In particular, LoRa employs an M-ary modulation scheme based on chirps [23]. Basic chirps are constant envelope signals whose frequency is linearly modulated sweeping from f_{min} to f_{max} (up-chirp) or from f_{max} to f_{min} (down-chirp). Chirps are cyclically-shifted to produce different symbols, and this cyclical shift carries the information. A symbol, whose length is divided in K equal time intervals called chips, can be cyclically shifted from 0 to $K - 1$ positions. The reference position is given by the un-shifted (base) symbol at the beginning of the LoRa frame, which is also used for building the frame preamble.

For a given bandwidth $B = f_{max} - f_{min}$, the symbol time is given by a SF parameter. The SF defines two modulation features: i) the time duration of each chirp (or, equivalently, the slope of the linear frequency sweep), which is given by 2^{SF} chip intervals; and ii) the number of raw bits encoded by that symbol, equal to SF. The Data Rate (DR) thus depends on the bandwidth B in Hz, the SF and the Coding Rate (CR) as:

$$DR = SF \cdot \frac{B}{2^{SF}} \cdot CR \tag{1}$$

where $1/B$ is the chip interval, the factor $B/2^{SF}$ provides the symbol rate and the coding rate $CR = 4/(4 + RDD)$ depends on the number of redundancy bits (RDD, from 1 to 4) used for Hamming code forward error correction. The bandwidth

can be configured as 125kHz, 250kHz and 500kHz (typically 125kHz is used in the 868MHz ISM band). Fig. 2 shows the modulating signal used for a basic upchirp and three examples of circular shifts obtained for SF = 9: the symbol time is $T = 512 T_c$, while the three exemplary shifts encode the symbols 128, 256 and 384 respectively. Formally, the instantaneous frequency of an unmodulated (base) LoRa chirp can be written as:

$$f_i^{(0)}(t) = -\mu \frac{B}{2} + \mu \frac{B}{T} t \tag{2}$$

where $\mu = +1$ gives an upchirp and $\mu = -1$ a downchirp, $T = 2^{SF} T_c$ is the symbol time and $T_c = 1/B$ the chip duration, $0 \leq t < T$. There are $K = 2^{SF}$ possible symbols, each representing a cyclicly shifted version of the base upchirp. The instantaneous frequency of symbol k is thus given by:

$$f_i^{(k)}(t) = \begin{cases} +\mu \frac{B}{2} + \mu \frac{B}{T}(t - kT_c) & 0 \leq t < kT_c \\ -\mu \frac{B}{2} + \mu \frac{B}{T}(t - kT_c) & kT_c \leq t < T. \end{cases} \tag{3}$$

The LoRa preamble starts with several repetitions of a base upchirp:

$$f_{pr}(t) = -\frac{B}{2} + B \left(\frac{t}{T} - \left\lfloor \frac{t}{T} \right\rfloor \right). \tag{4}$$

After several consecutive base upchirps, the preamble features two modulated symbols, called *sync* words, for network identification, and 2.25 downchirps which are useful for accurate synchronization. Following the preamble, the payload header, the payload and an optional frame check sequence are transmitted by using the cyclicly-shifted M-ary modulation.

C. LoRa DEMODULATION AND CO-CHANNEL REJECTION

LoRa demodulation can be implemented with a very simple receiver architecture: the received symbol is multiplied with the synchronized base down-chirp for obtaining a signal comprising only two frequencies: $-k/T$ and $-B - k/T$. Both frequencies can be aliased to the same frequency by down-sampling at the rate B . Finally, the symbol index k can be estimated by considering the position of the peak at the output of an FFT, as described in [14].

An interesting feature of LoRa is the quasi-orthogonality of signals modulated under different SFs. Indeed, two concurrent transmissions, modulated using different SFs, can be easily separated thanks to the fact that the cross-energy between the signals, say $s_1(t)$ and $s_2(t)$, is almost zero regardless of the overlapping time offset, i.e.:

$$E_{s_1, s_2}(\tau) = \int_0^T s_1(t) \cdot s_2(t - \tau)^* dt \simeq 0$$

where T is the symbol period of the signal with the highest SF. This phenomenon is depicted in Fig. 3, where we show the receiver operation when a reference signal modulated with SF 8 is subject to interference from another signal, modulated

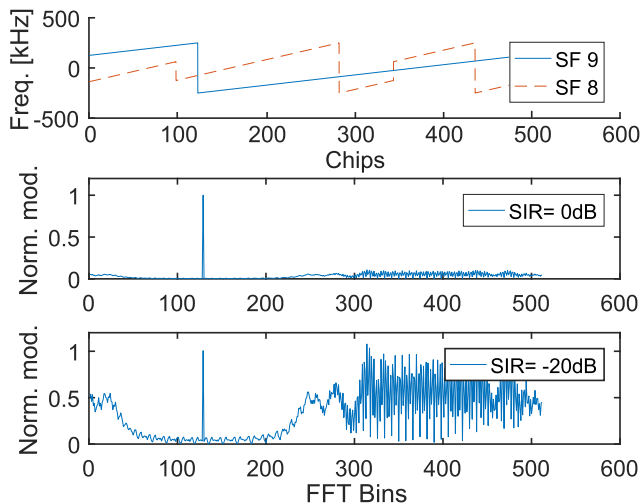


FIGURE 3. Interference of signals modulated with different SF (SF = 9 and SF = 8), and FFT output after multiplication with the base downchirp and downsampling at two SIR values (0dB and -20dB).

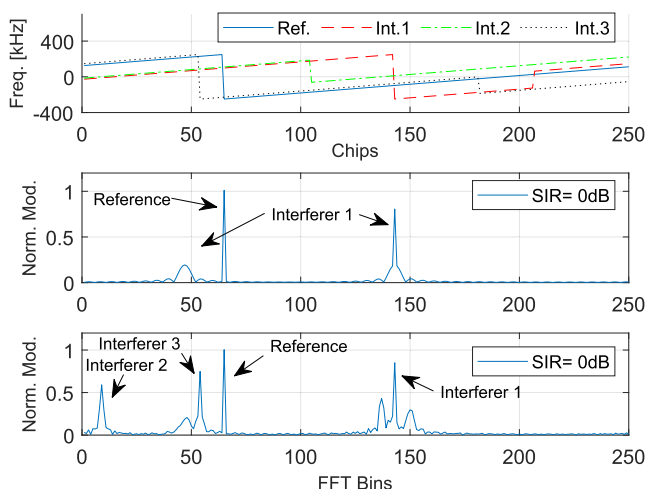


FIGURE 4. Interference of signals modulated with same SF. A LoRa reference symbol (solid line) and 3 partially de-synchronized interfering signals received at same power (SIR = 0dB), and FFT output after multiplication with the base downchirp and downsampling (only one interferer before and all interferers contributes before).

with SF 9. If the reference symbol is correctly synchronized, after multiplying with the down-chirp and performing the FFT, we can easily recognize the peak corresponding to the reference symbol. Only when the interfering signal is 20dB stronger than the reference signal, a simple peak detector can fail. The interference rejection thresholds have been numerically and experimentally derived in [3], [5].

In case the SF of the interfering signal is the same as the one the receiver is listening for, the above receiver will observe multiple peaks at the output of the FFT: a maximum peak corresponding to the reference symbol, and two smaller peaks corresponding to two partially overlapping interference symbols, as shown in the middle diagram of Fig. 4 for two signals received at the same power. In the example, the two interfering peaks are smaller than the reference symbol, thanks to the fact that the receiver is synchronized to the reference

symbol. In other words, even a SIR = 0dB can be sufficient for avoiding ambiguities in the identification of the maximum peak of the reference signal. Obviously, the capability to correctly demodulate the reference signal is much better when SIR > 0. This capability of correctly demodulating the reference signal when interfered by another signal modulated with the same SF is called *channel capture*. Note that, in case of collisions with multiple interfering signals, the receiver operation leads to the identification of multiple peaks: one peak for the reference (synchronized) symbol and two peaks for each other interfering symbol. Since it is very unlikely that the peaks due to the interfering signals are received in the same position, and therefore increase their amplitude, the LoRa receiver appears to work as each interfering signal is contributing separately (i.e., the interfering power of overlapping signals does not sum).

An example of this phenomena is shown in bottom Fig. 4, where four LoRa signals collide, yet the reference signal remains clearly visible compared to the interfering peaks. Obviously, the probability of correctly demodulating the strongest signal is critically affected by the synchronization mechanism implemented in the receiver: for example, if the LoRa receiver is locked to the colliding signal which is received first, the demodulation works only when the first signal is the strongest from among these that are colliding.

IV. THE LoRaSyNc RECEIVER

In this section, we present the design and implementation of the synchronization and cancellation mechanism that we called LoRaSyNc. An important feature of our receiver is the ability to track the frequency and clock drifts between the transmitter and receiver, during demodulation of the entire interfered frame. Indeed, keeping a dynamic synchronization to the strongest signal is essential for improving the probability to correctly demodulate at least one frame in case of collisions. Such a demodulation can be exploited for interference cancellation and iterative demodulation of other colliding frames, as detailed in the next section.

A. CARRIER AND TIME SYNCHRONIZATION

A first, and important, component of our synchronization mechanism is identifying the exact beginning of a preamble in time (the instant it starts) and in frequency (from which value it starts). Indeed, the receiver has to identify the unknown timing of a new frame transmission, as well as the shifts from the nominal carrier values experienced by each transmitter/receiver pair. Fig.5 shows the carrier and timing offsets of the initial part of a LoRa preamble in terms of the instantaneous frequency (continuous blue curve), as observed by an unsynchronized receiver. Assuming a carrier frequency offset equal to Cfo , and a time offset equal to τ (with $|\tau| < T/2$), the instantaneous frequency of the received signal can be written as:

$$f_{rx}(t) = Cfo - \frac{B}{2} + B \left(\frac{t - \tau}{T} - \left\lfloor \frac{t - \tau}{T} \right\rfloor \right). \quad (5)$$




FIGURE 5. Instantaneous frequency of the preamble for an unsynchronized receiver.




FIGURE 6. Instantaneous frequency (red) of the output of a mixer driven by the unsynchronized preamble (blue continuous) and the conjugate of the reference signal (blue dashed).

Our synchronization mechanism is built by exploiting the preamble structure, which includes both upchirp and downchirp transmissions. The idea is mixing (i.e., multiplying) the received signal $f_{rx}(t)$ with the complex conjugate of a reference preamble upchirp $f_{pr}(t)$, as shown in blue dashed line in Fig. 6, obtaining:

$$f_{mix,1}(t) = Cfo - B\frac{\tau}{T} - B\left(\left[\frac{t-\tau}{T}\right] - \left[\frac{t}{T}\right]\right) \quad (6)$$

Since the downchirp has the same absolute slope as the unsynchronized upchirp of the preamble, the frequency of the mixed signal $f_{mix,1}(t) = f_{rx}(t) - f_{pr}(t)$ changes over time as $[(t-\tau)/T] - [t/T]$, which is a square wave with values 0 and ± 1 and duty cycle $|\tau|/T$ (red line in Fig. 6). It follows that the output of the mixer signal features tones at only two frequencies $\nu_1 = Cfo - B\tau/T$ (when $t \geq \tau$) and $\nu_1 \pm B$ (when $t < \tau$). The same mixing mechanism can be applied to the last part of the preamble (constituted by downchirps), by multiplying the signal with base upchirps. The resulting




FIGURE 7. Probability of failing to detect a preamble, using the following parameters: SF = 7, B = 125kHz, $P_{threshold} = -121$ dBm. Dot lines limit the confidence interval.

signal is:

$$f_{mix,2}(t) = Cfo + B\frac{\tau}{T} + B\left(\left[\frac{t-\tau}{T}\right] - \left[\frac{t}{T}\right]\right) \quad (7)$$

which has the same structure as (6), with frequencies $\nu_2 = Cfo + B\tau/T$ and $\nu_2 \pm B$. Assuming that ν_1 and ν_2 are available, the estimated carrier and timing offsets Cfo_{est} and τ_{est} can be computed as:

$$Cfo_{est} = \frac{\nu_1 + \nu_2}{2} \quad (8)$$

$$\tau_{est} = \frac{T}{B} \frac{\nu_2 - \nu_1}{2}. \quad (9)$$

We designed the LoRaSyNc synchronization mechanism on the basis of this estimation approach. For identifying the start of a new preamble, LoRaSyNc acts as follows:

- 1) it samples the received signal $r(t)$ with a sampling frequency $f_s = B \cdot OSF$, i.e., Over Sampling Factor (OSF) times the nominal bandwidth of the signal, obtaining $r_n = r(n/f_s)$
- 2) it multiplies (mixes) a window of $N = K \cdot OSF$ samples with a base downchirp, obtaining:

$$z_n = r_n \exp(j\pi(n/OSF - n^2/(N \cdot OSF))) \quad (10)$$

- 3) it computes the absolute value of the FFT of the mixed signal z_n :

$$\Gamma_k = \left| \sum_{n=0}^{N-1} z_n \exp(-j2\pi nk/N) \right| \quad (11)$$

- 4) it searches for the maximum in Γ_k and saves its position \hat{k}
- 5) if this estimated position \hat{k} is detected continuously for a number of times (e.g., 3 consecutive windows), the receiver understands that an incoming preamble is received.




FIGURE 8. Standard deviation of the Cfo error estimate, using the following parameters: SF = 7, B = 125kHz, P_{threshold} = -121 dBm.

This procedure is executed continuously, for each possible SF, even when the demodulation of a frame is already in progress. This algorithm succeeds in detecting a preamble even several dBs below the sensitivity threshold: indeed, the probability of failing the detection of a preamble is shown in Fig. 7, when the power margin is computed as the difference between the received signal power and the receiver sensitivity declared in the Semtech 1272/73 datasheet [24], in Fig 7, the dot lines limit the confidence interval. From the figure, it is evident that the receiver is able to detect a preamble even when the received power is below that sensitivity threshold. Only when the margin is smaller than 7dB, is the failure probability higher than 1%. Once the preamble is detected, fine estimates of ν_1 and ν_2 can be obtained as follows:

$$\hat{\nu} = \frac{B}{K} \left(\hat{k} + \frac{1}{2} \frac{\Gamma_{\hat{k}-1} - \Gamma_{\hat{k}+1}}{\Gamma_{\hat{k}-1} - 2\Gamma_{\hat{k}} + \Gamma_{\hat{k}+1}} \right) \quad (12)$$

where the values of \hat{k} and Γ_k are obtained from the multiplication with the base downchirp in the first part of the preamble, and with its conjugate for the last 2.25 preamble symbols (for the last portion of the preamble, made of downchirps). Equation (12) is based on a parabolic interpolation around the maximum, and yields excellent results even for very low SNR values. We quantified the effectiveness of the proposed estimation scheme by simulating the receiver operation in MATLAB. Fig. 8 shows the standard deviation of the Cfo estimation error as a function of the power margin defined above, for the exemplary case SF = 7 and B = 125kHz.

B. DRIFTS TRACKER FOR SIGNAL CANCELLATION

A second important component of LoRaSyNc is the continuous tracking of the time and frequency drifts between a given transmitter and the receiver, during the whole reception of an incoming frame. Since LoRa devices employ low-cost crystal oscillators, the duration of symbols and the chirp frequency

sweeps have an inherent mismatch with their nominal values, which need to be tracked for enabling signal cancellation. In other words, even once a preamble has been detected, and initial estimates of Cfo and τ have been performed, the LoRaSyNc receiver needs to update these estimates during the reception of the frame, in order to accurately regenerate the received signal corresponding to the detected frame. This tracking is important especially for frames with long transmission times, for which the initial Cfo and τ estimates can become inaccurate along frame reception. We consider signal cancellation only for frames that are received successfully, i.e., which do not fail the header checksum and payload CRC.² For these frames, LoRaSyNc regenerates the entire modulated signal by using its local oscillator. However, before performing signal cancellation, it tries to compute an average temporal and frequency drift of the received signal in consecutive temporal windows. In particular, we chose a time window of four symbols as a compromise between accuracy and complexity. For each time window, the signal regenerated at the receiver is correlated with three different versions of the originally received signal: the one obtained considering a time offset equal to the initial estimate τ_{est} , and two other versions in which the offsets are equal to $\tau_{est} \pm 1/f_s$, i.e., shifted by plus or minus one sample. The three correlation operations will result in three different maximum values, which will be interpolated using a quadratic function. Finally, the maximum of this parabolic interpolation will provide the time offset used in the current window for signal cancellation. Before performing cancellation, LoRaSyNc also tries to correct possible errors on the frequency estimation. The Cfo tracking mechanism is similar to the clock drift tracking described above: it works by considering both the amplitude and the phase of the signal for interpolation. Once the tracking is completed, the current signal window can be rebuilt and the cancellation can start with a reduced residual error.

Note that, from a complexity point of view, the implementation of LoRaSyNc is mostly based on the computation of FFTs or time correlations (also implemented using the FFT algorithm) and few other vector multiplications (e.g., for linear interpolation), which have negligible complexity. The computational cost of the proposed receiver is therefore in the order of $O(N \log(N))$, being N the number of samples in a window.

For assessing the performance of our tracking scheme, we analyzed the received signal after cancellation, by considering a simple case of two colliding frames. We measured the residual signal (i.e., the difference between the received signal and the regenerated reference frame) in three different type of experiments: i) in simulation, where there is obviously no clock drifting; ii) in real experiments using two USRPs, one for sending frames in a controlled collision scenario and the other running the cancellation scheme without clock drift tracking; iii) in the same USRP testbed, applying the drift

²We omit a detailed description of the receiver procedure and refer the interested reader to [14]

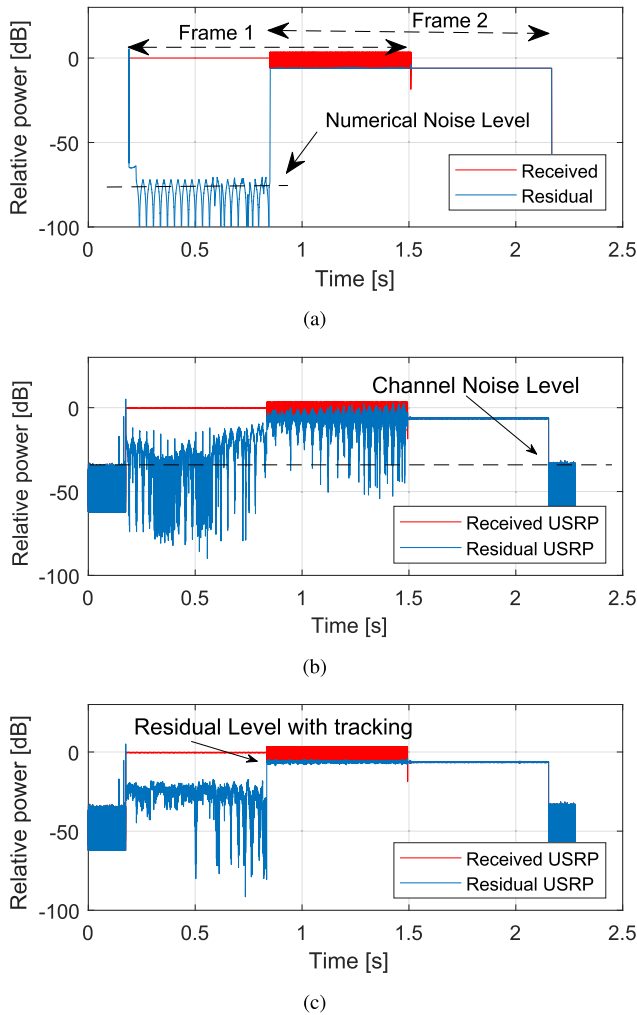


FIGURE 9. Residue in simulation 9(a) and in experimental testbed without clock tracking 9(b) and with clock tracking algorithm 9(c).

tracker algorithms at the receiver in order to solve the drift issues. Fig. 9 compares the results obtained on a sample collision in the above three types of experiments. In particular, Fig. 9(a) shows the result of a collision of two frames modulated at SF 12 and demodulated with the LoRaSyNc scheme simulated in MATLAB, before (received signal, in red) and after cancellation of the first received frame (residual signal, in blue). In ideal conditions, the cancellation residual noise is due solely to numerical approximations, constituting a noise floor of approximately -70dB . Instead, Fig. 9(b) and Fig. 9(c) show the residual signal in real experiments, where both channel noise and clock drift affect the reception. In particular, Fig. 9(b) shows the results of the SIC without applying the tracking algorithm (as done in previous solutions [8], [9]), while Fig. 9(c) shows the same when applying also the clock tracking scheme. From the two figures, it is evident that the residual noise is more pronounced in the former case than in the latter and this phenomenon is more and more relevant as the frame duration grows. More quantitatively, it turns out that the variance of the squared absolute value of the residual signal without clock tracking is about 0.0109, while using

the tracking algorithm the variance drops to $8.5267 \cdot 10^{-5}$. Note also that USRP clocks are more precise than commercial Semtech chips (datasheets report $\pm 2.0\text{ppm}$ for USRP versus $\pm 10\text{ppm}$ of 1301 chipsets, almost an order of magnitude higher), so the impact of clock drifting is more relevant when implementing a SIC algorithm with commercial end devices, reinforcing the need of the proposed tracking mechanisms. Thus, it is possible to assert that SIC performance is definitely improved using the proposed tracking algorithm.

V. DATA EXTRACTION RATE IN CASE OF COLLISIONS

As discussed in the background section, LoRa modulation is very robust against interfering signals, and it therefore is very likely that the frames colliding at a given GW result in the correct reception of the strongest. However, it has been demonstrated that this event also depends on the temporal sequence in which the colliding signals are originated [25]: when the strongest signal does not start first, it may happen that the LoRa receiver is locked for the demodulation of a previous frame and no colliding signal can be received correctly. To study the performance improvements of LoRaSyNc with respect to a commercial GW, we quantified the capture probability in case of collisions, as well as the signal cancellation performance, through both simulation and experimental campaigns. To this end, we artificially generated collision events by using the LoRa cell traffic emulator described in [26], and by considering different collision scenarios (varying the SFs, the SNR/SIR levels, and the time offset between the colliding frames). Although our results can be easily generalized, we specifically studied the case in which collisions are limited to two frames only. Indeed, we assume that in most practical cases, a target station can be interfered by a single colliding signal at time: as discussed in sec. III-C and demonstrated in [5], this assumption is reasonable when the cell works in stable conditions, i.e., when collisions involving multiple overlapping frames have a very low occurrence probability.

A. PERFORMANCE OF COMMERCIAL GATEWAYS

In this section we analyze the performance of a commercial LoRa GW, based on the Semtech 1301 chipset. Since commercial chipsets do not currently implement signal cancellation, we expect that collisions can result at most in the correct reception of one frame. The LoRa traffic generator [26] is used for creating the signal samples corresponding to the overlapping of two frames, with different power levels and temporal shifts. The samples are then transmitted using a USRP B210 SDR platform. The traffic generator has been configured to support a coding rate of 4/5. We created different collision scenarios by varying: i) the overlapping time of the two colliding frames and ii) the SIR value experienced by the first transmitted frame. Unless otherwise specified, the bandwidth is set to 125kHz, the center frequency to 868MHz and the packet length to 20 bytes. Two different experiment campaigns were performed, one involving different SFs and one focusing on collisions between frames with

TABLE 1. Experiment results using Semtech 1301 receiver, for different time offsets and SIR values.

Configuration			Reference Frame			Interferer Frame		
Exp. no.	Offset [%]	SIR[dB]	CRC_OK [%]	CRC_BAD [%]	LOST [%]	CRC_OK [%]	CRC_BAD [%]	LOST [%]
1	no	0	100.0	0.0	0.0	99.0	0.0	1.0
2	no	3	100.0	0.0	0.0	100.0	0.0	0.0
3	85	3	100.0	0.0	0.0	10.0	0.0	90.0
4	15	3	100.0	0.0	0.0	0.0	0.0	100.0
5	85	1	72.0	28.0	0.0	36.0	0.0	64.0
6	15	1	66.0	29.0	5.0	0.0	0.0	100.0
7	85	0.5	53.0	47.0	0.0	59.0	0.0	41.0
8	15	0.5	63.1	23.8	13.1	0.0	0.0	100.0
9	85	-3	0.0	100.0	0.0	100.0	0.0	0.0
10	15	-3	0.0	3.0	97.0	0.0	0.0	100.0
11	85	-1	9.0	91.0	0.0	81.0	1.0	18.0
12	15	-1	4.0	36.0	60.0	0.0	0.0	100.0
13	85	-0.5	18.3	81.7	0.0	75.6	0.0	24.4
14	15	-0.5	7.0	51.0	42.0	0.0	0.0	100.0

TABLE 2. Experiment results using Semtech 1301 receiver, where the second frame (the interferer) is 3 dB stronger.

Configuration			Reference Frame			Interferer Frame		
Exp. no.	Offset [%]	SIR[dB]	CRC_OK [%]	CRC_BAD [%]	LOST [%]	CRC_OK [%]	CRC_BAD [%]	LOST [%]
1	10	-3	0	0	100	1	0	99
2	20	-3	0	11	89	46	0	54
3	30	-3	0	3	97	11	0	89
4	40	-3	0	4	96	0	0	100
5	50	-3	0	34	66	0	0	100
6	60	-3	0	32	68	0	0	100
7	70	-3	1	18	81	0	0	100
8	80	-3	0	20	80	1	0	99
9	90	-3	8	92	0	99	1	0
10	100	-3	100	0	0	99	0	1

the same SF. In the first set of experiments, we generalize the SIR-based capture model provided in [3], for a different CR, namely 4/5 rather than 4/7. We fixed the SF and power level of the reference frame and varied both SF and power level of the interfering frame, in order to identify the interfering power limit corresponding to a Data Extraction Rate (DER, also known as packet delivery ratio – PDR) higher than 90%. The overlapping time was generated uniformly between 0 and the Time on Air (ToA) of the reference frame. Results, not shown here for the sake of space, differ from the case CR = 4/7 of at most 1.5dB, substantially confirming the previously published results in [5]. In the second experiment campaign, we analyze collisions originated by frames transmitted with the same SF. We consider the most critical case SF = 7, which corresponds to the highest modulation rate. We collected results for different offsets between the starting of the two transmissions (85% and 15% of the ToA) and set the SIR to the following values: 0dB, ± 0.5 dB, ± 1 dB, and ± 3 dB. The distance between the nodes is 30 meters and each experiment is repeated over the air 1000 times using the same configuration. Finally, we measured the fraction of packets correctly demodulated at the GW. In particular, we identify three types of results: i) packets received correctly (CRC_OK), ii) packets received but with a failed checksum (CRC_BAD), and

iii) packets lost (LOST). Table 1 reports the results of 14 different experiment settings. In the table, the first column shows an experiment identifier, and summarizes the collision configuration. Moreover, in the first two experiments (number 1 and 2), no collision has been considered to ensure that the channel is clear from external interference and that the setup works correctly. We first analyze the case of positive SIR values (experiments from 3 to 8). In this scenario, corresponding to the case in which the first reference frame is the strongest one, when the SIR is ≥ 3 dB, all the reference frames are correctly received. For lower SIR values, it may happen that the receiver correctly demodulates the interfering frame rather than the reference one, when the overlapping frame starts with an offset of 85% (experiments 5 and 7). Conversely, when negative SIR values are considered (experiments from 9 to 14), the interfering frame captures the channel when the offset is 85% (CRC_OK > 75%), but the receiver is not able to demodulate the strongest interfering frame when the offset between the colliding signals is 15% (in these conditions, both reference and interfering frames are lost). This is likely due to the synchronization mechanism implemented by the GW, as also observed in [25] for frames modulated with different SFs. Indeed, if the receiver is locked after

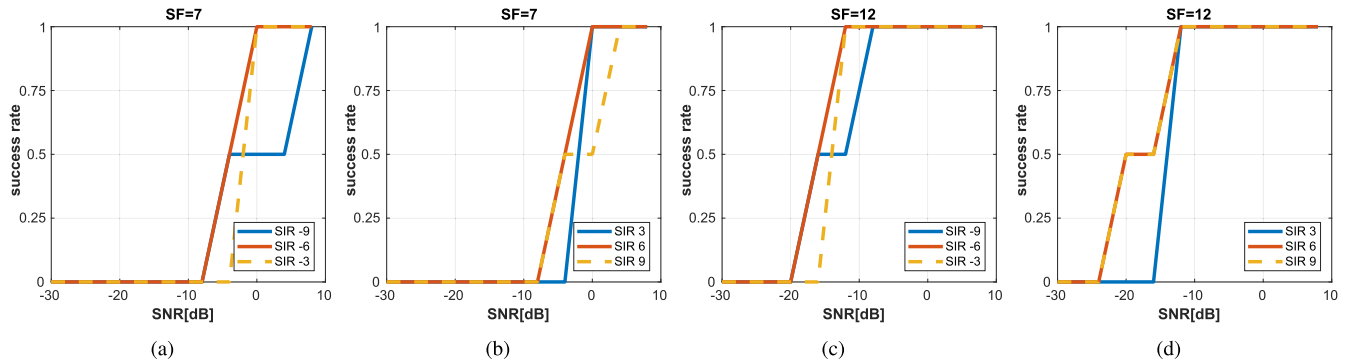


FIGURE 10. Number of demodulated packets varying SIR and SNR for (a) SF = 7 and SIR < 0, (b) SF = 7 and SIR > 0, (c) SF = 12 and SIR < 0 and (d) SF = 12 and SIR > 0.

the first detected preamble, the beginning of the following strongest interfering signal might be missed. Therefore, in Table 2 we focus on the case where the second frame (the interferer) is 3dB stronger. In the table, the offset varies uniformly between 10% and 100% of the ToA, i.e., with collisions involving the preamble (experiment 1), the header (experiment 2 and 3), the rest of the reference frame (experiments 4-9), or do not overlap at all (experiment 10). From the table, it is clear that in this scenario the strongest frame is usually unable to capture the channel and both frames are lost. The only exception are experiments 2-3, where the interferer collides on the header of the reference frame and in experiments 9-10, where the overlap between frames is very little or absent. From the presented results, it is clear that in commercial GWs capture effects are highly dependent on the SIR value of the first received frame, as well as on the time offset between the frames.

B. LoRaSyNc PERFORMANCE

In order to optimize the performance of LoRaSyNc, we implemented a preamble detection function which is executed continuously, even when another frame demodulation is in progress. In other words, with LoRaSyNc, receiver is never locked: in case a stronger preamble is detected by the synchronization module during the demodulation of a frame, the receiver is able to switch to the demodulation of the new detected signal. Indeed, the correct demodulation of the strongest signal is fundamental for increasing the probability to also recover the weakest colliding signal, by means of signal cancellation. We first analyzed the performance of LoRaSyNc in MATLAB, by working on signal traces generated under ideal conditions (i.e., without frequency and time drifts). For evaluating the impact of the data rate on the receiver performance, we considered the limit cases of signals modulated using SF = 7 (high-rate modulation scenario) and SF = 12 (robust modulation scenario). Taking the station transmitting first as reference, we varied the power of the interfering station by considering SIR values in the range within ±9dB (with steps of 3dB). The time offset of the interfering signal has been set to 50% of the frame

duration (recall that LoRaSyNc always tries to synchronize with new frames, so time offset is not a critical parameter). Fig. 10 shows the success rate, obtained by comparing the number of demodulated packets with the number of transmitted ones, as a function of the SNR value of the reference station. We consider different scenarios, in which frames are modulated at SF 7 (sub-figures (a) and (b)) and SF 12 (sub-figures (c) and (d)), and the reference packet is weaker (sub-figures (a) and (c)) or stronger (sub-figures (b) and (d)) than the interfering packet (i.e., the SIR is lower or greater than zero). Since all the received traces include collisions between two frames, it is clear that in most cases LoRaSyNc either recovers both packets (success rate equal to 1) or recovers none of them. In other words, if one of the signals is correctly demodulated (thanks to the capture effect), then the receiver is able to apply the SIC algorithm and successfully recover also the second packet. Note also that, as expected, the minimum SNR for which the receiver is able to properly receive a packet for SF = 7 is different from SF = 12, due to the different robustness of each modulation format. We repeated the same evaluation experimentally in our campus testbed. In each experiment, the combined signal resulting from the overlapping of two LoRa transmissions is synthesized by using the SDR platform USRP B210 with central frequency $f = 868.0\text{MHz}$. A second USRP B210 is configured to receive and collect I/Q samples of 400 consecutive collision events. The main purpose of the over-the-air transmission is to introduce time de-synchronization due to the internal clock drift of transmitter and receiver nodes, which employ non-stabilized oscillators. The acquired traces have been then processed by our MATLAB implementation of LoRaSyNc. Fig. 11 shows the success rate when varying SNR and SIR levels, in both the collision scenarios in which the colliding frames are modulated with SF = 7 and SF = 12. As already observed in simulation, the minimum SNR level for correct reception of a frame using SF = 7 is different from SF = 12. Moreover, results show that the success rate is similar to the one obtained in simulation, especially for SIR = ±9dB (blue line). Note that similar success rates have been obtained by previous solutions only for short frames,

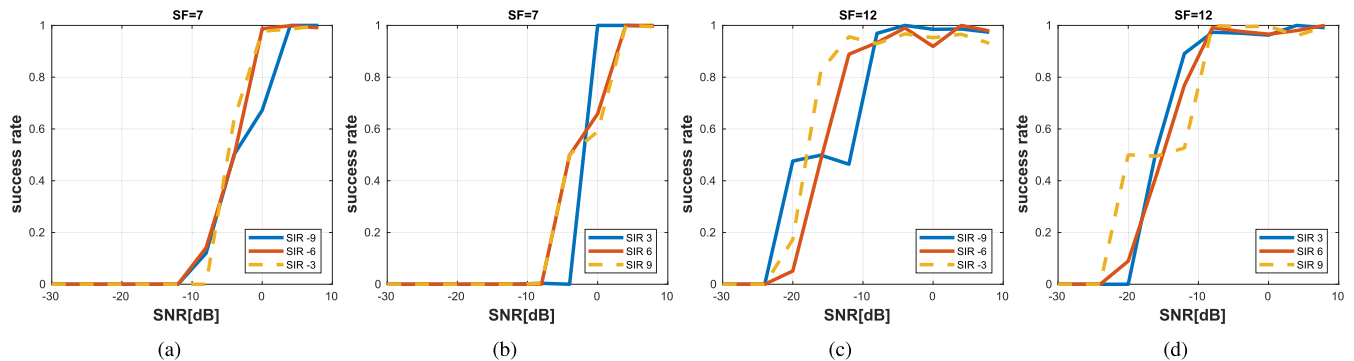


FIGURE 11. Success rate varying SIR and SNR for (a) SF = 7 and SIR < 0, (b) SF = 7 and SIR > 0, (c) SF = 12 and SIR < 0 and (d) SF = 12 and SIR > 0 in real experiments.

TABLE 3. Simulation parameters.

Parameter	Value
Carrier Frequency	868.0 MHz
Transmission Power	14 dBm
Bandwidth	125 kHz
Coding Rate	4/5
Message size	20 Bytes
Message Period	336000 ms
Number of EDs	[100-2000]
Path loss attenuation exponent	4
Simulation time	604800 s

while clock tracking allows our scheme to outperform the other ones for long-lasting frames (e.g., when using SF=12 or transmitting long payloads).

VI. IMPACT OF SIGNAL CANCELLATION ON LoRa CELL DEPLOYMENTS

We now evaluate the performance achieved by LoRaSyNc in terms of capacity improvements, which can be achieved in a cell, where all devices are configured on the same SF. This is the worst case scenario, since collisions involving frames modulated with different SF are much easier to be recovered thanks to the pseudo-orthogonality of the SFs. We measure the improvement margin after several SIC iterations, and analyze the impact of pathloss attenuation. We also compare the results with a commercial GW, able to correctly demodulate at most one single frame in case of collisions (captures only), and with an ideal Aloha system. Finally, we analyze the case of multiple LoRaSyNc GWs covering the same area, varying the SF and the number of nodes in the cell. For this last scenario, we extended the LoRaSim [27] simulator, by implementing a realistic capture model derived by our experiments presented in section V-A. Unless specified otherwise, the parameters used for the simulations are summarized in Table 3. We assume that all the nodes generate packets of the same size and transmit at the same SF, i.e., all frames have a constant transmission time T_{frame} . Fig. 12 shows the performance of a LoRaSyNc gateway in terms of normalized throughput (i.e., average number of packets successfully demodulated in a frame transmission time T_{frame}) versus

normalized offered load (i.e., average number of packets generated during T_{frame}). We used our LoRa traffic generator for synthesizing the traffic traces corresponding to different load scenarios, and then the traces are processed by our MATLAB LoRaSyNc implementation. The figure shows the performance of LoRaSyNc at different stages of the SIC process, after several rounds of cancellation (from 1 to 4 only, since improvements become negligible afterwards). For comparison, the figure also reports the performance that would achieve a commercial GW (exploiting channel captures only) and the theoretical Aloha system. While in the latter case the obtained throughput tends to zero as expected, both the former (traditional GWs) and LoRaSyNc are capable of receiving packets, even under heavily loaded conditions. Indeed, in such scenarios, there is always a chance for nodes closer to the GW to capture the channel and, for LoRaSyNc, to recover more packets after cancellation. Particularly, in extremely congested conditions (offered load higher than 3), LoRaSyNc outperforms traditional GWs almost doubling the obtained throughput. Fig. 13 shows the performance of LoRaSyNc in a cell with only one GW when the attenuation exponent η is equal to 4, 2 and 0, corresponding respectively to strong attenuation (urban scenarios), low attenuation (typical of rural scenarios), and no attenuation (i.e., all signals are received with the same power). From the figure, it is possible to note that LoRaSyNc obtains a normalized throughput of almost 0.7, even when packets are all received with the same power ($\eta = 0$), while in more realistic settings ($\eta = \{2, 4\}$) the maximum throughput is about 1.2, which is possible thanks to the SIC algorithm. Moreover, in scenarios with attenuation $\eta = 4$, the throughput stays above 1 even when faced with offered loads over 10 times the network capacity. In other words, LoRaSyNc preserves the ability of the GW to receive packets in highly saturated conditions and reduces the risk of cell collapse in case of simultaneous awake of devices (which is a very important property e.g., for emergency applications).

In the next set of experiments, we measure the performance of LoRaSyNc by varying the number of nodes and the SF, and analyzing the impact of having one or multiple SIC gateways in the same cell. To this end, we extended the LoRaSim [27] simulator including the capture probabilities

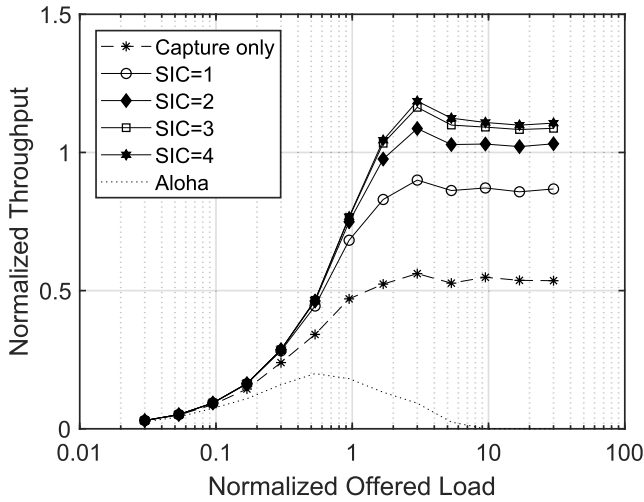


FIGURE 12. Performance of LoRaSyNc with increasing the number of SIC rounds (from 1 to 4), compared to a traditional GW (capture only) and the theoretical Aloha model. The improvement becomes negligible increasing the number of SIC iterations further.




FIGURE 13. Performance of a LoRa cell with single LoRaSyNc GW and different values of path loss attenuation exponent η .

TABLE 4. Capture probability used in simulations.

SIR[dB]	Capture probability
$ SIR > 3$	100%
$ SIR \in [1, 3]$	55%
$ SIR < 1$	0%

observed at different SIR values and for different temporal offsets (presented in section V-A). In particular, we derived a simplified capture probability model, summarized in Table 4, where collisions between two frames lead to the correct reception of both frames provided that the absolute power difference is at least $|SIR| > 3$ dB, while is probabilistic in the range $[1, 3]$ dB (regardless of the order in which the strongest frame is received). We will thus use this receiver model to analyze the performance of a LoRa cell in presence of one or multiple GWs. For the case of a single GW, Fig. 14




FIGURE 14. Performance of LoRaSyNc compared to a traditional receiver (dashed lines) in case of channel captures only.

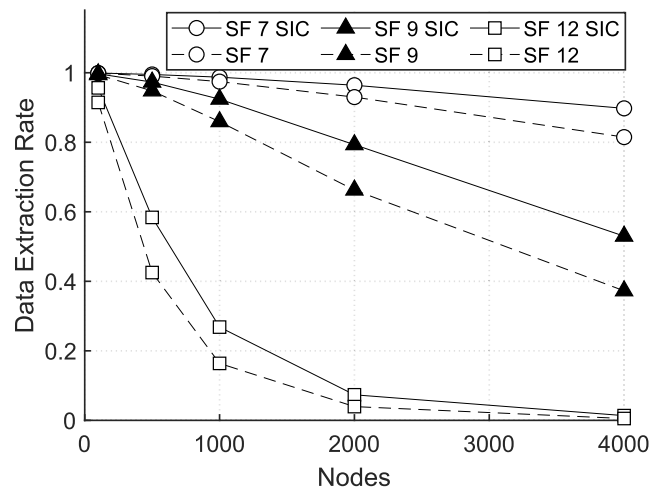


FIGURE 15. Performance of LoRaSyNc in a cell with 4 GWs compared to 4 traditional receivers (dashed lines) in case of channel captures only.

shows the DER results as a function of the number n of EDs configured on a specific SF (we show only a subset of SFs, for the sake of clarity). Each curve refers to an independent cell sub-system with the same number of EDs but different load conditions. The dashed curve depicts the DER as achieved by a standard LoRa GW (captures only, no SIC), which represents a lower performance bound. The results demonstrate that the DER can increase significantly thanks to LoRaSyNc, especially in intermediate scenarios. For example, with $n = 2000$ EDs and SF = 9, the DER increases from about 0.3 to almost 0.45 in presence of LoRaSyNc, thus almost a 50% increase. Obviously, when close to saturation (e.g., 2000 EDs and SF = 12), the probability to receive interference by a node with a similar received power is higher, resulting in a proportionally lower DER increase. Finally, we repeated the experiments using 4 GWs in the same cell and up to $n = 4000$ EDs. Fig. 15 shows the DER obtained when nodes are all using the same SF (again we show only SF = 7, 9 or 12). From the figure, it is clear that LoRaSyNc


outperforms traditional LoRa GWs (dashed lines), e.g., with almost 50% capacity increase with 4000 EDs and $SF = 9$.

VII. CONCLUSION

LoRa modulation has been demonstrated to be very robust to different interference sources, including co-channel interference generated by collisions between multiple overlapping frames. This is relevant for systems working in ISM unlicensed bands, based on a simple random access protocol. In this paper, we discuss the possibility of demodulating the strongest signal in case of collisions, for enabling signal cancellation and iterative decoding. Indeed, this capability can lead to significant improvements of system-level capacity, thus making LoRaWAN technology more scalable. Due to the long intervals required for frame transmissions (especially at high SF) and to the low quality of device oscillators, we discuss the importance of tracking clock and frequency drifts between transmitters and receivers in order to enable signal cancellation. We quantify the capability of demodulating two colliding signals in controlled experiments and compare these results with the performance of commercial GWs. We observe that commercial GWs perform worse than expected, because their capability of demodulating the strongest colliding signal depends on the order and time offsets between the colliding signals. Conversely, our proposed LoRaSyNc receiver is able to demodulate both colliding signals in most of the analyzed scenarios. Finally, we analyze the improvements on the system-level capacity, which can be achieved when deploying one or more LoRaSyNc gateway. We show that the proposed receiver improves performance by almost 50% compared to a traditional receiver.

REFERENCES

- [1] S. Popli, R. K. Jha, and S. Jain, "A survey on energy efficient narrowband Internet of Things (NB-IoT): Architecture, application and challenges," *IEEE Access*, vol. 7, pp. 16739–16776, 2019.
- [2] B. Reynders and S. Pollin, "Chirp spread spectrum as a modulation technique for long range communication," in *Proc. Symp. Commun. Veh. Technol. (SCVT)*, Nov. 2016, pp. 1–5.
- [3] D. Croce, M. Gucciardo, S. Mangione, G. Santaromita, and I. Tinnirello, "Impact of LoRa imperfect orthogonality: Analysis of link-level performance," *IEEE Commun. Lett.*, vol. 22, no. 4, pp. 796–799, Apr. 2018.
- [4] D. Croce, M. Gucciardo, I. Tinnirello, D. Garlisi, and S. Mangione, "Impact of spreading factor imperfect orthogonality in LoRa communications," in *Proc. Int. Tyrrhenian Workshop Digit. Commun. (TIWDC)*, Palermo, Italy, Sep. 2017, pp. 6510–6523.
- [5] D. Croce, M. Gucciardo, S. Mangione, G. Santaromita, and I. Tinnirello, "LoRa technology demystified: From link behavior to cell-level performance," *IEEE Trans. Wireless Commun.*, vol. 19, no. 2, pp. 822–834, Feb. 2020.
- [6] X. Wang, L. Kong, L. He, and G. Chen, "MLoRa: A multi-packet reception protocol in LoRa networks," in *Proc. IEEE 27th Int. Conf. Netw. Protocols (ICNP)*, Chicago, IL, USA, Oct. 2019, pp. 1–11.
- [7] M. A. B. Temim, G. Ferre, B. Laporte-Fauret, D. Dallet, B. Minger, and L. Fuche, "An enhanced receiver to decode superposed LoRa-like signals," *IEEE Internet Things J.*, vol. 7, no. 8, pp. 7419–7431, Aug. 2020.
- [8] Z. Wang, L. Kong, K. Xu, L. He, K. Wu, and G. Chen, "Online concurrent transmissions at LoRa gateway," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, Jul. 2020, pp. 2331–2340.
- [9] X. Xia, Y. Zheng, and T. Gu, "FTrack: Parallel decoding for LoRa transmissions," in *Proc. 17th Conf. Embedded Netw. Sensor Syst.*, New York, NY, USA, Nov. 2019, pp. 192–204.
- [10] *USRP B200/B210 Product Overview*. Accessed: Jan. 8, 2021. [Online]. Available: https://www.ettus.com/content/files/b200-b210_spec_sheet.pdf
- [11] L. Vangelista, A. Zanella, and M. Zorzi, "Long-range IoT technologies: The dawn of LoRa," in *Proc. 1st Int. Conf., Future Access Enablers Ubiquitous Intell. Infrastruct.*, Ohrid, Republic of Macedonia, Sep. 2015, pp. 51–58.
- [12] B. Reynders, W. Meert, and S. Pollin, "Range and coexistence analysis of long range unlicensed communication," in *Proc. 23rd Int. Conf. Telecommun. (ICT)*, Thessaloniki, Greece, May 2016, pp. 1–6.
- [13] M. Knight and B. Seeber, "Decoding LoRa: Realizing a Modern LPWAN with SDR," in *Proc. GNU Radio Conf.*, vol. 1, Sep. 2016, pp. 1–5.
- [14] C. Bernier, F. Dehmas, and N. Deparis, "Low complexity LoRa frame synchronization for ultra-low power software-defined radios," *IEEE Trans. Commun.*, vol. 68, no. 5, pp. 3140–3152, May 2020.
- [15] A. Augustin, J. Yi, T. Clausen, and W. M. Townsley, "A study of LoRa: Long range & low power networks for the Internet of Things," *Sensors*, vol. 16, no. 9, p. 1466, 2016.
- [16] D. Halperin, J. Ammer, T. Anderson, and D. Wetherall, "Interference cancellation: Better receivers for a new wireless MAC," in *Proc. 6th ACM Workshop Hot Topics Netw. (Hotnets)*, 2007, pp. 1–6.
- [17] D. Bankov, E. Khorov, and A. Lyakhov, "Mathematical model of LoRaWAN channel access with capture effect," in *Proc. IEEE 28th Annu. Int. Symp. Pers., Indoor, Mobile Radio Commun. (PIMRC)*, Montreal, QC, Canada, Oct. 2017, pp. 1–5.
- [18] N. El Rachkidy, A. Guitton, and M. Kaneko, "Decoding superposed LoRa signals," in *Proc. IEEE 43rd Conf. Local Comput. Netw. (LCN)*, Oct. 2018, pp. 184–190, doi: [10.1109/LCN.2018.8638253](https://doi.org/10.1109/LCN.2018.8638253).
- [19] A. Mahmood, E. Sisinni, L. Guntupalli, R. Rondon, S. A. Hassan, and M. Gidlund, "Scalability analysis of a LoRa network under imperfect orthogonality," *IEEE Trans. Ind. Informat.*, vol. 15, no. 3, pp. 1425–1436, Mar. 2019.
- [20] O. Georgiou and U. Raza, "Low power wide area network analysis: Can LoRa scale?" *IEEE Wireless Commun. Lett.*, vol. 6, no. 2, pp. 162–165, Apr. 2017, doi: [10.1109/LWC.2016.2647247](https://doi.org/10.1109/LWC.2016.2647247).
- [21] U. Noreen, L. Clavier, and A. Bounceur, "LoRa-like CSS-based PHY layer capture effect and serial interference cancellation," in *Proc. 24th Eur. Wireless (EW)*, Catania, Italy, May 2018, pp. 1–6.
- [22] N. Sornin and A. Yegin. (Oct. 2017). *LoRaWAN 1.1 Specification*. [Online]. Available: <https://loro-alliance.org/resource-hub/lorawan-specification-v11>
- [23] Semtech. (May 2015). *LoRa Modulation Basics. AN1200.22, Revision 2*. [Online]. Available: <https://www.semtech.com/images/datasheet/an1200.22.pdf>
- [24] *Semtech SX1272 Long Range, Low Power RF Transceiver 860-1000 MHz With LoRa Technology*. Accessed: Jan. 8, 2021. [Online]. Available: <https://www.semtech.com/products/wireless-rf/loro-core/sx1272>
- [25] A. Rahmadhani and F. Kuipers, "When LoRaWAN frames collide," in *Proc. 12th Int. Workshop Wireless Netw. Testbeds, Exp. Eval. Characterization*, New York, NY, USA, Oct. 2018, pp. 89–97.
- [26] M. Gucciardo, I. Tinnirello, and D. Garlisi, "Demo: A cell-level traffic generator for LoRa networks," in *Proc. 23rd Annu. Int. Conf. Mobile Comput. Netw.*, New York, NY, USA, Oct. 2017, pp. 474–476, doi: [10.1145/3117811.3119862](https://doi.org/10.1145/3117811.3119862).
- [27] *LoRaSim is a Discrete-Event Simulator Based on SimPy for Simulating Collisions in LoRa Networks and to Analyse Scalability*. Accessed: Jan. 8, 2021. [Online]. Available: <https://www.lancaster.ac.uk/scc/sites/loro/lorasim.html>



DOMENICO GARLISI received the Ph.D. degree in electronic and telecommunication engineering from the University of Palermo, in 2014. He was a Visiting Researcher with the Department of Computer Science, University of California at Los Angeles, in 2013, working on wireless protocol optimizations for vehicular networks. He is currently working a CNIT Researcher with the Palermo Research Unit. His main research interests include wireless networks, software defined radio, sensor networks, the IoT, and wireless systems.



STEFANO MANGIONE (Member, IEEE) received the degree (*cum laude*) in electronics engineering, telecommunications curriculum, in 2000. He is currently an Assistant Professor of telecommunications engineering with the University of Palermo. His research activities have been focused on physical layer aspects of communication systems, from forward error correction coding to equalization strategies for spread spectrum systems. Recently his activity includes automatic methods for registration of nuclear magnetic resonance images, multiuser receiver strategies for LoRa, and the study of underwater communication systems.



GIOVANNI GARBO received the degree (*cum laude*) in electronics engineering, telecommunications curriculum, and the Ph.D. degree in electronics and informatic engineering from the University of Palermo, in 1985 and 1990, respectively. He is currently a Full Professor of telecommunications engineering with the University of Palermo. His research activities have been focused on physical layer aspects of communication systems, convolutional codes over groups, and OFDM/OQAM systems. Recently his activity included automatic methods for registration of nuclear magnetic resonance images.



FABRIZIO GIULIANO was born in Italy, in 1982. He received the master's and Ph.D. degrees in telecommunication engineering, in 2009 and 2014, respectively. He is currently a Researcher with the Consorzio Nazionale Interuniversitario per le Telecomunicazioni (CNIT). He has been involved in several European research projects and in Italian National Project "PON i-NEXT."



DANIELE CROCE received the dual M.Sc. degree in networking engineering from the Politecnico di Torino and EURECOM Institute, Sophia Antipolis, France, in 2006, the Research Master Diploma (ex DEA) degree in networking and distributed systems from UNSA, in 2006, and the Ph.D. degree jointly from the Politecnico di Torino, Turin, Italy, and the Université de Nice-Sophia Antipolis (UNSA), Nice, France, in 2010. He is currently an Assistant Professor with the University of Palermo, Italy. He has long experience of research collaborations, in several European and national research projects, on wireless networks, the Internet of Things, high-quality TV streaming, smart grid communications, and smart cities. He also worked on assistive technologies for visually impaired people and with the Arianna Project. He co-founded the start-up company In.sight s.r.l. (<http://in.sight.srl>), spin-off of University of Palermo.



ILENIA TINNIRELLO received the Ph.D. degree in telecommunications engineering from the University of Palermo, in 2004. She has also been Visiting Researcher with Seoul National University, South Korea, in 2004, and Nanyang Technological University, Singapore, in 2006. She is currently an Associate Professor with the University of Palermo. Her research activities have been focused on wireless networks, and in particular on the design and prototyping of protocols and architectures for emerging reconfigurable wireless networks. Recently, she is also working on the definition of novel services (smart grid, smart metering, and indoor localization) enabled by the pervasive availability of ICT technologies. She has been involved in several European research projects, among which the FP7 FLAVIA project, with the role of technical coordinator, and the H2020 WiSHFUL and Flex5Gware projects.

...