



HAL
open science

La rumeur disait faux : ils ne sont pas orthogonaux !

Rachida Saroui, Alexandre Guitton, Oana Iova, Fabrice Valois

► To cite this version:

Rachida Saroui, Alexandre Guitton, Oana Iova, Fabrice Valois. La rumeur disait faux : ils ne sont pas orthogonaux!. CORES 2022 – 7ème Rencontres Francophones sur la Conception de Protocoles, l'Évaluation de Performance et l'Expérimentation des Réseaux de Communication, May 2022, Saint-Rémy-Lès-Chevreuse, France. hal-03657877

HAL Id: hal-03657877

<https://hal.science/hal-03657877>

Submitted on 3 May 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

La rumeur disait faux : ils ne sont pas orthogonaux !

Rachida Saroui¹, Alexandre Guitton^{1,2}, Oana Iova¹, Fabrice Valois¹

¹Univ Lyon, INSA Lyon, Inria, CITI, F-69621 Villeurbanne, France

²Université Clermont-Auvergne, CNRS, Mines de Saint-Étienne, Clermont-Auvergne-INP, LIMOS, 63000 Clermont-Ferrand, France

Les réseaux LoRaWAN font partie des technologies clés pour la collecte de données dans l'Internet des Objets sur des zones étendues. Le protocole LoRaWAN propose principalement une voie montante pour la remontée des informations issues des capteurs, mais utilise également une voie descendante pour transmettre des commandes de paramétrage ou les acquittements des trames de données quand nécessaire. Dans tous les travaux de la communauté, cette voie descendante est supposée orthogonale à la voie montante, en raison de l'utilisation d'une modulation inversée entre ces deux voies. A l'aide d'expérimentations utilisant des radio-logicielles, nous montrons pour la première fois que cette rumeur disait faux : le taux de livraison de trames peut chuter de 20% lorsque des transmissions concurrentes sur les voies montantes et descendantes se produisent.

Mots-clés : LoRaWAN, LoRa, chirp orthogonality, GNU radio, experiments.

1 Introduction

After a decade of research and innovation, the Internet of Things (IoT) became a reality in our everyday life, covering applications like water metering, pollution monitoring, logistics, healthcare, and home automation. To support large coverage for devices with limited energy consumption, low-power wide-area networks (LPWAN) provide a noticeable gain in terms of capacity and deployment cost when compared to classical short-range solutions. For these reasons, the Long Range (LoRa) [Sel17] modulation and the Long Range Wide Area Network (LoRaWAN) [LoR17] protocol became crucial players in the IoT market.

In LoRaWAN, a large part of the traffic uses the uplink to send data collected from end-devices to a remote application server. The downlink is used to send (i) acknowledgments for the confirmed traffic when it is needed, and (ii) control information from the network server to the end-devices, for instance, in order to update the transmission parameters of the end-devices. Many research works have investigated the capacity and the performance of LoRaWAN [KP18, MCZ20], but they always assumed that the uplink is fully independent of the downlink. This assumption relies on the fact that the LoRa modulation scheme of the uplink uses mainly upchirps, while the downlink uses mainly downchirps.

In this work, we investigate the orthogonality of uplinks and downlinks in LoRaWAN. Using software defined radios (*i.e.*, an USRP software and the GNU radio toolkit) and we show that the packet delivery ratio (PDR) on both uplink and downlink communications decreases by up to 20% when the transmissions are concurrent. We also show that the impact depends on the amount of overlap between the uplink and the downlink frames. To the best of our knowledge, this is the first time that a non-orthogonality result of uplinks and downlinks is shown in LoRa.

The remainder of the paper is organized as follows. In Section 2, we briefly describe the LoRa and LoRaWAN technologies. In Section 3, we present our testbed using GNU radio and USRPs. In Section 4, we show and discuss our results. Finally, we conclude our paper in Section 5.

2 LoRa and LoRaWAN in a nutshell

LoRa (Long Range) [Sel17] is a physical layer using a chirp spread spectrum modulation. Each chirp is a linear frequency sweep over a given bandwidth (BW), and for a given duration. A chirp is called upchirp

if the sweep is increasing (see the black chirps in Fig. 2), or downchirp if the sweep is decreasing (see the green chirps in Fig. 2). The duration of the chirp mainly depends on a parameter called spreading factor (SF). The starting frequency of the chirp gives its value, encoded on SF bits. LoRa is able to trade-off the communication range with the throughput through SF: a large SF yields a large communication range, but at the cost of a longer chirp duration, and thus of a smaller throughput. LoRa also uses a forward error correction to detect errors and correct them, denoted coding rate (CR).

A LoRa uplink frame is composed of a preamble and a payload. The preamble contains 8 upchirps of value 0, 2 upchirps encoding the network identification (NI), and 2.25 downchirps of value 0 to indicate the end of the preamble (see Fig. 2). The payload is composed of a varying number of upchirps. It usually contains a header, the actual data, and a cyclic redundancy check. A LoRa downlink frame is also composed of a preamble and a payload. Both the preamble and the payload use an inverse modulation compared to an uplink frame: upchirps are used in place of downchirps, and downchirps in place of upchirps (see Fig. 2).

To standardize the use of LoRa frames, the access to the medium, and the network architecture, LoRa Alliance defined the LoRaWAN protocol [LoR17]. In the LoRaWAN topology, end-devices send uplink frames with LoRa. The frames are captured by gateways, which forward them to a network server, which in turn forwards them to the corresponding application server. When an end-device requests a confirmation, called confirmed traffic, the network server chooses one of the gateways to reply to the end-device: the acknowledgment from the gateway is sent using a downlink LoRa frame. The network server can also send control frames to the end-devices via the gateways using downlink LoRa frames.

End-devices in LoRaWAN typically use an ALOHA mechanism. Notably, end-devices belonging to Class A send their frames without sensing the medium. After each transmission, the end-device opens two short reception windows to listen for any incoming downlink frames. Depending on regional parameters, end-devices might have to implement a duty-cycle (*i.e.*, waiting a specified off time between two consecutive transmissions).

3 Implementation and validation of the testbed

In order to find out if uplink and downlink communications in LoRaWAN are orthogonal or not, we need to be able to study the effects of concurrent transmissions, *i.e.*, we need to send uplink and downlink frames that overlap in time, and see if they are correctly received or not. If the uplink and downlink transmissions are orthogonal, both frames should be correctly received. The tricky part when setting up such an experiment is to be able to control the transmission of a frame with high time precision, so that the effects of overlapping uplink and downlink frames can be studied at the level of a symbol, where the symbol duration is 1.024 ms for the following configuration: $f=868$ MHz, $SF=7$, $BW=125$ kHz.

We hence decided to use Software Defined Radios (SDR). More precisely, we used the Universal Software Radio Peripherals (USRP) from National Instruments[†], and the GNU Radio open-source software development toolkit. In order to send and receive LoRa frames, we need a LoRa implementation module on GNU Radio. However, since LoRa is a proprietary technology, its code is not available in open-source. We evaluated several reversed-engineering implementations and chose the implementation from Tapparel *et al.* [TPA⁺20][‡] that provides a full implementation of LoRa uplink frames, including sampling time offset (STO) and carrier frequency offset (CFO) estimations and corrections. However, the module does not provide the implementation of LoRa downlinks.

We implemented the transmission (resp. reception) of downlink frames by a signal processing technique called I/Q inversion on the output (resp. input) signal modulated for the uplink. This is done mathematically by swapping the I (in-phase) and Q (quadrature) components of the signal, which results into an inversion of the sign of the frequency of the signal, and thus effectively transforms upchirps into downchirps and vice versa. We validated our implementation through the following experiments:

1. *Interactions with real hardware.* The goal of this experiment was to ensure that the uplink and downlink modulations were correctly implemented in the USRP. We made a setup in which we sent LoRa

[†]. National Instruments USRP-2901 (equivalent to the model USRP-B210 from Ettus Research).

[‡]. The module can be downloaded at https://github.com/tapparelj/gr-lora_sdr.




FIGURE 1: (a) Testbed setup using four USRPs. (b) Waterfall diagram for two concurrent uplink and downlink frames, with time flowing downwards. (c) PDR for uplink and downlink traffic. The boxes show the median (in red), the first and third quartile (the box), and the 95% confidence interval (the whiskers).

frames from a commercial LoRa device[§] to our USRP, and vice versa. We were able to successfully decode the frames, both on the USRP and on the commercial device, proving that our USRP implementation is compatible with the LoRa modulation.

2. *Capture effect.* The goal of this experiment was to verify that the capture effect of LoRa works correctly in our USRP implementation. We setup two USRPs to send uplink messages with different transmission power to a third USRP, acting as the receiver. We varied the transmission power, but we fixed all the other parameters (frequency band 868 MHz, SF=7, BW=125 kHz and CR=4/5). We observed the packet error rate as a function of the difference of signal power, and observed that when a signal is at least 6 dB stronger than another, the strong signal is captured. These results are consistent with the state of the art, and show that the capture effect works correctly in our implementation.
3. *Interference-free environment.* The goal of this experiment was to setup a baseline for our tests and verify that our results are not impacted by external interference. We setup one USRP to send 100 uplink frames to another USRP, and observed the packet delivery ratio (PDR). Then, we repeated the same experiment for downlink frames. As we can see in Fig. 1(c), the results showed a PDR of 1, meaning that our experiment was not significantly influenced by external interference.

Considering these three experiments, we assume that our implementation and environment are validated.

4 Results: pseudo-orthogonality of uplink and downlink in LoRa

In this section, we describe our main experiment in order to determine whether uplink communications and downlink communications are orthogonal or not (*hint: they are not!*). We setup four USRPs as follows: 1st is transmitting only uplinks, 2nd is receiving only uplinks, 3rd is transmitting only downlinks, and 4th is receiving only downlinks (see Fig. 1(a)). Both transmitters use the same transmission power, and are located at the same distance from the receivers. We use the LoRa parameters for the highest data rate (SF=7, BW=125 kHz, CR=4/5), which gives us the smallest symbol duration possible (1.024 ms). Note that our findings also apply for the other LoRa parameters.

To investigate uplink and downlink orthogonality, we setup concurrent uplink and downlink communications, and we vary the delay between the two transmitters to control the overlap between the two frames. An example of such concurrent transmission is shown on Fig. 1(b). Note that in this example, the two frames are slightly desynchronized. In order to experience all possible overlaps, we configure a periodic transmission for the 40 downlink frames, and we delay the transmission of the uplink frame with a time offset varying from 0 symbol to 40 symbols with a step of one symbol. For each value of the delay, two repetitions are performed. The average PDR for uplink and for downlink are shown in Fig. 1(c). Surprisingly, the PDR decreased by approximately 10% for both transmissions, whereas it was 100% before. This clearly shows that the two communication streams are not orthogonal as initially thought.

[§]. Dragino LoRa Shield <https://www.dragino.com/products/lora/item/108-lora-gps-shield.html>




FIGURE 2: Concurrent uplink and downlink transmissions as a function of the time offset between the two frames. The red boxes represent collisions between upchirps (in black) and respectively downchirps (in green) from different frames. The PDR is computed as the mean between uplink and downlink PDRs in the specified setup.

We further investigate this result by comparing the signal representations and the PDR of the two signals at different time offsets, as can be seen in Fig. 2. We can notice that as the offset increases from 1 to 4 symbols, the number of chirps that overlap also increases, leading to a decrease in PDR by up to 15%. This is clearly because the uplink downchirps are getting superposed with the downlink downchirps, and the same for the upchirps. For lack of space, we cannot add here the remaining of the figures, but we see a PDR varying from 79% to 99%, depending on the number of symbols that overlap. Furthermore, we also noticed that even when the two transmissions happen exactly at the same time, *i.e.*, the time offset is 0 symbol, there is a slight decrease in PDR (0.97), showing that the uplink and downlink signals are not completely orthogonal. Indeed, for a signal and its spectral inversion to be orthogonal, the components I and Q have to be perfectly orthogonal, which is not always the case because of the IQ imbalance phenomenon [ASdVL11], meaning there is an error on the phase shift and one on the amplitude.

5 Conclusion

LoRa uplinks and downlinks use an inverse modulation, and are therefore believed to be orthogonal. In this paper, we implemented an experimental setup using SDRs in order to verify this assumption. We showed for the first time that uplinks and downlinks are not orthogonal. The PDR on the uplink decreases in the presence of downlink traffic by up to 20%. This new result may significantly impact all the performance studies on the uplink capacity of LoRaWAN in the presence of downlink traffic. ¶

References

- [ASdVL11] M. S. Alavi, R. B. Staszewski, L. C. N. de Vreede, and J. R. Long. Orthogonal summing and power combining network in a 65-nm all-digital RF I/Q modulator. In *IEEE International Symposium on Radio-Frequency Integration Technology*, pages 21–24, 2011.
- [KP18] F. H. Khan and M. Portmann. Experimental evaluation of LoRaWAN in NS-3. In *International Telecommunication Networks and Applications Conference (ITNAC)*, pages 1–8, 2018.
- [LoR17] LoRa Alliance Technical Committee. LoRaWAN 1.1 specification, 2017.
- [MCZ20] D. Magrin, M. Capuzzo, and A. Zanella. A thorough study of LoRaWAN performance under different parameter settings. *IEEE Internet of Things Journal*, 7:116–127, 2020.
- [Sel17] O. Seller. Wireless communication method, 2017. US patent 9647718-B2, Semtech Corp.
- [TPA⁺20] J. Tapparel, O. Afisiadis P., Mayoraz A., Balatsoukas-Stimming, and A. and Burg. An open-source LoRa physical layer prototype on GNU radio. *IEEE 21st International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, pages 1–5, 2020.

¶. This research has received financial support from the project ANR-21-CE25-0002-01.