



HAL
open science

The Impact of Downlink Scheduling Policy on the Capacity of LoRaWAN

Christelle Caillouet, Alexandre Guitton, Oana Iova, Fabrice Valois

► **To cite this version:**

Christelle Caillouet, Alexandre Guitton, Oana Iova, Fabrice Valois. The Impact of Downlink Scheduling Policy on the Capacity of LoRaWAN. GLOBECOM 2024 - IEEE Global Communications Conference, IEEE, Dec 2024, Cape Town / Le Cap, South Africa. hal-04694132

HAL Id: hal-04694132

<https://hal.science/hal-04694132v1>

Submitted on 11 Sep 2024

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.

The Impact of Downlink Scheduling Policy on the Capacity of LoRaWAN

Christelle Caillouet*, Alexandre Guitton†, Oana Iova‡, Fabrice Valois‡

*Université Côte d’Azur, I3S, Inria, 06903 Sophia Antipolis, France

†UCA, CNRS, Mines de Saint-Étienne, Clermont-Auvergne-INP, LIMOS, 63000 Clermont-Ferrand, France

‡INSA Lyon, Inria, CITI, UR3720, 69621 Villeurbanne, France

Abstract—LoRaWAN is a widely used wireless communication standard for the Internet of Things that enables the collection of measurement data from numerous monitoring applications (such as smart metering, pollution, and asset tracking). Most research work has been focused on the performance of LoRaWAN as a function of the uplink traffic, as this is the most common scenario. Nevertheless, downlink traffic is a fundamental building block of the LoRaWAN standard, and a crucial part of applications such as smart healthcare, where reliability is extremely important. In this paper, we study the impact of the downlink traffic and different scheduling policies on the performance of LoRaWAN. We investigate: (i) the impact of an optimal schedule for the downlink traffic, (ii) the choice of the physical layer parameters used to send downlinks to end-devices during the second reception window, and (iii) the choice of the gateway that sends the downlinks. Our results show that even when using an optimal schedule, the presence of downlink traffic reduces the capacity of a LoRaWAN network up to 20%. The most limiting factor is the gateway, due to its duty cycle and half-duplex characteristics.

I. INTRODUCTION

LoRa technology [1] and Long-Range Wide Area Networks (LoRaWAN) [2] have become a cornerstone in the Internet of Things (IoT), enabling the collection of measurement data from a plethora of sensors, such as air quality, parking, and smart metering. Its success is due to its energy efficient communication over long distances [3]. While most of the traffic is uplink (UL) (from end-devices to gateways), downlink (DL) traffic (from a gateway to end-devices) should not be neglected as it represents a fundamental building block in LoRaWAN networks, enabling: (i) network joining procedure for end-devices, (ii) adapting the data rate of end-devices, (iii) confirming correctly received data through acknowledgments (Ack), (iv) sending commands to actuator devices, and (v) sending firmware updates to end-devices.

Recent works showed that DL communications can severely limit the capacity of LoRaWAN, and even impair UL communications, as confirmed in real-life experiments [4]. For example, the capacity of a LoRaWAN network drops to just a few hundred end-devices when the traffic model is changed from unconfirmed traffic to confirmed (traffic that needs to be acknowledged) [5], its performance being severely impaired with the growth of the number of confirmed frames [6], [7].

The root of the problem stands both in the hardware and software components of LoRaWAN: (i) LoRaWAN is a half-duplex communication system, as end-devices and gateways cannot transmit and receive at the same time, (ii) a

LoRaWAN gateway is physically capable of receiving several communications in parallel from up to 8 different end-devices [8], but it is only able to respond to a single end-device at a given time, (iii) traffic sent by end-devices is not scheduled, so UL and DL communications can overlap [9], (iv) the duty cycle of the gateway is limited, so the number of DL frames that can be sent to end-devices per unit of time is also limited, (v) LoRaWAN defines two distinct receive windows for end-devices during which the gateway can send a DL frame. However, these windows cannot be fully exploited when the ULs were received at the same time.

The aim of our work is to quantify the impact of an optimal schedule for DL traffic on the capacity of a LoRaWAN network. To that end, we develop an optimization model that takes into account a fine characterization of the LoRa chipset, the radio channel in an urban setting, and the spatial diversity of LoRaWAN deployments, which is not the case for other theoretical works [5], [10]–[12]. We make the following contributions: 1) we propose an optimal schedule for DL traffic, 2) we study the impact on LoRaWAN performance when increasing the number of gateways, 3) we study the causes for dropped DL frames, 4) we compare different scheduling policies by taking into account the choice of the physical layer parameters (such as the spreading factor) used to send the DLs, and the choice of the gateway sending the DL traffic.

The rest of this paper is organized as follows. Section II provides an overview of the LoRaWAN protocol and related state of the art. Section III describes our system, with the optimization model being detailed in Section IV. Section V discusses our simulation results. Finally, Section VI concludes this work and discusses future work.

II. LORA AND LORAWAN BACKGROUND

LoRa is a proprietary wireless technology based on a chirp spread spectrum (CSS) modulation, where the time on air (ToA) of a frame and the data rate depend on several parameters: 1) *Spreading Factor (SF)*: defines how many chirps make up a symbol. The higher the SF, the better the signal-to-noise ratio, as well as the communication range. 2) *Bandwidth (BW)*: determines how spread out the signal is, providing higher rates for higher bandwidth. 3) *Coding Rate (CR)*: sets the forward error coding (FEC) rate in use. A typical configuration for high data rate and energy efficiency is SF7–BW125–CR4/5.

LoRaWAN standard specifies the network architecture and the upper layers protocols to be used on top of LoRa. Aloha is used by end-devices to send their data to all the gateways in range, which then relay it to the network server. Depending on the configuration, a DL can be sent to end-devices by choosing one of the gateways. While gateways are more simplistic, the network server plays the most important role in the network, being responsible for: data rate adaptation, network joining procedure, UL frame deduplication and DL scheduling. The adaptive data rate (ADR) mechanism makes end-devices optimize their data rate by gradually decreasing their SF, while checking the received signal quality. To test if the channel conditions did not change, end-devices have to periodically send a confirmed frame. If no Ack is received, the end-device must try to regain connectivity by gradually increasing its SF.

We focus here on the Class A specification, as it is the most used class. After each UL transmission, end-devices open two reception windows – RX1 and RX2 – before going back in sleep mode to preserve energy. If a DL needs to be sent to the end-device, it should be scheduled to arrive in one of these reception windows, otherwise it is missed. LoRaWAN specifies that if a DL is sent on RX1, then the gateway needs to use the same frequency channel and SF as the UL. If the DL is sent on RX2, a different frequency (that allows a higher transmission power) and SF can be used, to increase the probability of frame reception: in practice, the LoRa Alliance recommends the use of S12 for reliability, whereas The Things Networks (TTN) uses SF9.

Just like others low-power wireless technologies, LoRaWAN uses the ISM unlicensed bands. Such band is subjected to regional regulations [13]. For example, in Europe (see Table I), LoRa needs to uphold a maximum transmission power and a duty cycle to ensure that the medium is not being overused and saturated, allowing fair access to all end-devices.

To evaluate the capacity of LoRaWAN under different conditions, several mathematical models have been proposed in the literature, with the majority of them focusing solely on UL communications [14]. Even when DL traffic was considered, most of the works only modeled the case of a single gateway scenario [5], [11], which is uncommon in LoRaWAN deployments. While spatial diversity is one of the strengths of a LoRaWAN deployment, accounting for multiple gateways in a mathematical model is much more complex, as each gateway serves a specific area, with some overlapping zones. Each gateway has a different capacity, making the schedule of DLs more complicated. Zorbas [12] models the DL capacity problem in LoRaWAN in the context of multiple gateways, however it does not consider the intrinsic LoRa characteristics such as the SF imperfect orthogonality or the capture effect. To the best of our knowledge, we are the first ones to present an optimal schedule for DL traffic in a LoRaWAN network by accurately modelling the LoRa physical layer in a spatially diverse deployment with up to 4 gateways.

	Channel Frequency [MHz]	Max TX power	Duty Cycle	SF	BW [kHz]
UL	868.{10, 30, 50}	14 dBm (25 mW)	1%	7-12	125
DL RX1					
DL RX2	869.525	27 dBm (500 mW)	10%	12	

TABLE I: LoRa Parameters in Europe for EU863-870.

III. SYSTEM DESCRIPTION

Our system has two components. First, a simulator that generates the network topologies and the traffic model while accounting for the realistic characteristics of the physical layer in an urban scenario. Second, we use an optimization model that receives as input the confirmed traffic from the simulator, and creates an optimal schedule of the DL frames. We note that the unconfirmed traffic is not an input of our optimization problem, as it does not need to be acknowledged, but is considered in the evaluation of the DL impact.

We start by presenting the simulator that we developed to accurately model the LoRa chipset [8] and the signal propagation in an urban scenario [15]. We simulate a square area of $1\text{km} \times 1\text{km}$. End-devices are deployed randomly within the whole area w.r.t. a connectivity constraint. Gateways are regularly distributed on a circle centered at the center of the area, and whose radius is 250m, to obtain a full coverage of the topology. In the specific case where only one gateway is deployed, it is located at the center of the area.

End-devices can send confirmed and unconfirmed frames. In both cases, traffic is generated so that the average number of frames per end-device is set to $n_{frames} = 10$. The starting time of the first frame for an end-device is equal to a random value chosen uniformly in $[0; t_{sim}/n_{frames}]$. Then, each next frame is generated following a Poisson distribution of intensity n_{frames}/t_{sim} . Successive frame starting times that are too close with respect to the duty-cycle of 1% are delayed. We do not consider frame retransmission.

The channel follows the Okumura-Hata model for urban areas [15]. Path loss is given by: $PL(d) = PL(d_0) + 10.a.\log_{10}(d/d_0)$, using the constants: $d_0 = 1$ m, $PL(d_0) = 74.85$ and $a = 2.75$. The shadow fading is a Gaussian variable of zero mean and variance σ , with $\sigma = 11.25$. The fast fading is a random variable of unit mean and with exponential distribution.

An UL frame f is received by a gateway g if the four following conditions occur: (i) the received signal strength of f for g exceeds the sensitivity of g , which depends on the SF of f , (ii) the received signal strength of f exceeds the total received power of all other frames overlapping in time with f and using the same SF as f , by a given threshold which depends on the SF (capture effect), (iii) the received signal strength of f exceeds the total received power of all frames overlapping in time with f but using others SF, by a given threshold which depends on the involved SFs (imperfect orthogonality of the SFs), and (iv) each gateway

can demodulate up to eight frames simultaneously [8]. Note that when the overlap between two frames is partial, only the corresponding fraction of the power of the interfering frame is considered. The value of each threshold is set according to the specifications of LoRaWAN [2].

The above conditions ensure that we closely follow the specification of the LoRa chipset and the signal propagation conditions in an urban setting, being as close as possible to a real LoRaWAN deployment. All the correctly received confirmed frames are then sent as input to the optimization model for creating the optimal DL scheduling.

IV. OPTIMIZATION MODEL

In this section we present our optimization model for DL scheduling by taking into account the complex characteristics of a spatially diverse deployment with multiple gateways. We consider a set N of end-devices that send UL frames in the form of confirmed and unconfirmed traffic. We seek to determine an optimal schedule of the DL frames that send the Acks to the received confirmed frames to evaluate the impact of DL scheduling on the network capacity.

A. Input Values

The model considers only the set of confirmed frames received by the network server (as given from the simulator) $\cup_{n_i \in N} \text{ul}(n_i)$ for all end-devices $n_i \in N$. Given t_j^i the j th frame $\in \text{ul}(n_i)$ of end-device n_i , we can determine when n_i will open both its receive windows RX1 and RX2 to listen to the Ack of t_j^i , as well as the DL duration (derived from the SF used by n_i). In the following, we denote $t1_j^i$ (respectively $t2_j^i$) the possible DL frame on RX1 (resp. RX2) to acknowledge t_j^i .

We then use the following sets of binary variables:

- $y1(g, t_j^i) = 1$ if t_j^i is acknowledged by gateway g during RX1 window, 0 otherwise.
- $y2(g, t_j^i) = 1$ if t_j^i is acknowledged by gateway g during RX2 window, 0 otherwise.
- $y(t_j^i) = 1$ if t_j^i collides during its complete transmission with a DL sent by the gateways, 0 otherwise.

B. Linear Program (LP)

The objective of our optimization problem aims at maximizing the number of DLs sent by the gateways in response to correctly received confirmed frames from all end devices. To give priority on RX1 over RX2, we introduce a parameter $\alpha = 10^{-5}$ and optimize

$$\max \sum_{n_i \in N} \sum_{t_j^i \in \text{ul}(n_i)} \sum_{g \in G} ((1 + \alpha)y1(g, t_j^i) + y2(g, t_j^i)) \quad (1)$$

where G is the set of gateways in the network. α is chosen in such a way that $\sum_{n_i \in N} \sum_{t_j^i \in \text{ul}(n_i)} \sum_{g \in G} (1 + \alpha)y1(g, t_j^i)$ stays strictly lower than 1.

We derive next the set of constraints of the LP modeling the DL scheduling.

1) *Half-duplex property of the gateways:* Additionally to the reception conditions presented in Section III, we consider that a confirmed frame t_j^i of $\cup_{n_i \in N} \text{ul}(n_i)$ is lost if all the gateways that can decode the frame (specified by our simulator) have a DL to send during the reception of t_j^i . In that case, since the gateways are half-duplex, they cannot listen to the channel while emitting. We also consider that if a gateway has a DL transmission scheduled, then the overlapping UL frame (confirmed or not) is lost even if the gateway has already locked on it [8].

Let $Gr(t_j^i) \subseteq G$ be the subset of gateways correctly receiving frame t_j^i , determined by our simulator following the radio model presented in Section III. If all the gateways of $Gr(t_j^i)$ have a DL scheduled during the reception of t_j^i , we set the binary variable $y(t_j^i)$ to 1 to indicate that t_j^i collides and has not to be acknowledged anymore. Therefore, we must ensure that $y(t_j^i) = 1$ if $\sum_{g \in Gr(t_j^i)} (y1(g, t_{j'}^i) + y2(g, t_{j'}^i)) = |Gr(t_j^i)|$, with $t1_{j'}^i \cap t_j^i \neq \emptyset$ and $t2_{j'}^i \cap t_j^i \neq \emptyset$. This means that each gateway of $Gr(t_j^i)$ has to send $t1_{j'}^i$ (any DL on RX1 for any end-device $n_{j'}$ on any SF) or $t2_{j'}^i$ (any DL on RX2 for any end-device $n_{j'}$) while receiving t_j^i . Otherwise there is at least one free gateway to receive and decode t_j^i , so in our model there is no constraint on the value of $y(t_j^i)$. In that case $\sum_{g \in Gr(t_j^i)} (y1(g, t_{j'}^i) + y2(g, t_{j'}^i)) < |Gr(t_j^i)|$, and so $\sum_{g \in Gr(t_j^i)} (y1(g, t_{j'}^i) + y2(g, t_{j'}^i)) - |Gr(t_j^i)| + 1 \leq 0$.

We thus obtain the following set of constraints:

$$\sum_{g \in Gr(t_j^i)} (y1(g, t_{j'}^i) + y2(g, t_{j'}^i)) - |Gr(t_j^i)| + 1 \leq y(t_j^i) \quad (2)$$

$\forall t_j^i, t1_{j'}^i \cap t_j^i \neq \emptyset, t2_{j'}^i \cap t_j^i \neq \emptyset$.

2) *Duty cycle of the gateways:* The duty cycle regulations for end-device are included in our simulator. We need here to derive the duty cycle constraints for the gateways, as specified in Table I. For each DL on RX1 using $\text{sf}(n_i)$, we must ensure that no other DL starts before $100 * \text{ToADL}(\text{sf}(n_i))$ (where ToA is the time on air using $\text{sf}(n_i)$). And for each DL on RX2, the gateway should wait $10 * \text{ToADL}(\text{SF}_{\text{RX2}})$, leading to:

$$y1(g, t_j^i) + y1(g, t_{j'}^i) \leq 1 \quad (3)$$

$\forall g \in G, t1_j^i = \{s1_j^i, e1_j^i\}$ and $t1_{j'}^i = \{s1_{j'}^i, e1_{j'}^i\}$ such that $s1_j^i < s1_{j'}^i < s1_j^i + 100 * \text{ToADL}(\text{sf}(n_i))$. Similarly for RX2 :

$$y2(g, t_j^i) + y2(g, t_{j'}^i) \leq 1 \quad (4)$$

$\forall g \in G, t2_j^i = \{s2_j^i, e2_j^i\}$ and $t2_{j'}^i = \{s2_{j'}^i, e2_{j'}^i\}$ such that $s2_j^i < s2_{j'}^i < s2_j^i + 10 * \text{ToADL}(\text{SF}_{\text{RX2}})$.

3) *One DL at a time:* This set of constraints looks globally at both reception windows to ensure that 2 overlapping DLs cannot be scheduled on the same gateway, even if one is during the RX1 window of one end-device and the other is during the RX2 window of another end-device.

$$y1(g, t_j^i) + \sum_{t1_{j'}^i \cap t1_j^i \neq \emptyset} y1(g, t_{j'}^i) + \sum_{t2_{j'}^i \cap t1_j^i \neq \emptyset} y2(g, t_{j'}^i) \leq 1 \quad (5)$$

$$y2(g, t_j^i) + \sum_{t_1^{i'} \cap t_2^{i'} \neq \emptyset} y1(g, t_{j'}^{i'}) + \sum_{t_2^{i'} \cap t_2^{i'} \neq \emptyset} y2(g, t_{j'}^{i'}) \leq 1 \quad (6)$$

$\forall g \in G, t_j^i \in \cup_{n_i \in \text{Nul}(n_i)}$.

4) *DL only if UL correctly received*: If a confirmed frame is correctly received by the network server, then at most one gateway can schedule the Ack either during RX1 or RX2.

$$\sum_{g \in G} (y1(g, t_j^i) + y2(g, t_j^i)) \leq 1 - y(t_j^i), \quad \forall t_j^i \quad (7)$$

For every t_j^i considered, only one DL can be sent either during RX1 or RX2 ($y1(g, t_j^i) + y2(g, t_j^i) \leq 1$), except if the UL has collided with a DL and is lost. In that case $y(t_j^i) = 1$, so all DL variables are forced to be equal to 0.

By changing the set G in the sum of constraints (7), we can evaluate different scheduling policies for sending the DL traffic in the network (see Section V for more details).

V. PERFORMANCE EVALUATION

In this section we evaluate our system model, by determining how gateways can acknowledge the largest number of UL frames that have been correctly received by the network server. We also evaluate the impact of different scheduling policies on the performance of LoRaWAN considering both confirmed and unconfirmed traffic. We present first our simulation setup, followed by a discussion on the obtained results.

A. Simulation Setup

The simulation duration is set to 3600 s. Each end-device transmits frames with a 20 bytes payload, at 14 dBm. The SF for each end-device is chosen based on the average power of 20 receptions plus a margin, such that its data rate is optimized. The bandwidth is set to 125 kHz and the coding rate to 4/5. The traffic and channel models are the ones presented in Section III. We consider two types of end-devices:

- (i) End-devices belonging to a high-reliability application whose frames need to be acknowledged. Their traffic is made only of confirmed frames.
- (ii) End-devices with ADR enabled belonging to a low-reliability application that do not need to receive any Ack. Still, since the ADR imposes a periodic confirmed frame transmission to recheck the connectivity with the network, we claim that 95% of the frames sent by these end-devices are unconfirmed frames, while the remaining 5% are confirmed frames.

We define a simulation configuration by the triple $\{N, G, x\}$, where $N \in \{100, 200, 300, 400\}$ is the number of end-devices, $G \in \{1, 2, 3, 4\}$ is the number of gateways, and $x \in \{0, 25, 50, 75, 100\}$ is the percentage of end-devices with only confirmed traffic in the network. We run 100 random scenarios for each configuration.

We also define a baseline scenario where we consider that 50% of end-devices have only confirmed traffic to send. The SF chosen to send the Ack during the RX2 window is fixed at




Fig. 1: Mean FDR with and without DL traffic (50% of confirmed traffic, SF12 on RX2).

SF12 (default value in the standard). The network server can choose any gateway to send the DL, even if it did not receive the corresponding UL. We present next our results.

B. Overall Performance

Our optimization model is solved using IBM Cplex solver 20.1 until the optimal solution is reached. For topologies with 1 or 2 gateways, the program is solved in a few seconds. For more gateways, and when the number of confirmed frames increases, the resolution time reaches around one minute for large cases with an input set of approximately 2200 confirmed frames. This shows the effectiveness of our model to obtain the optimal number of DL frames supported by the network.

Figure 1 studies the impact of sending DL traffic on the Frame Delivery Ratio (FDR), which is defined as the number of UL frames received at the network server divided by the number of frames sent. More specifically, it shows the mean FDR as a function of the number of end-devices, for several gateways. The solid curves represent the mean FDR when there is no confirmed traffic (except for the ADR frames), whereas the dotted curves show the mean FDR with the DL frame scheduling. Sending DL traffic reduces the FDR significantly in configurations with only 1 gateway because the gateway becomes unavailable for unconfirmed and confirmed frame reception every time it sends a DL. Moreover confirmed frames that cannot be acknowledged are considered lost because the end-devices need to resend them. Deploying more gateways mitigates this phenomenon since the possibility to receive UL traffic and send DLs at the same time is leveraged.

Figure 2 presents the percentage of DLs successfully sent by the gateways, as a function of the number of end-devices (in each box), of the percentage of confirmed traffic (x-axis of each box), and of the number of gateways (colored boxplot inside each box). Each boxplot shows the mean value (black triangle), the median value (orange line), the first and third quartile (box), and the whiskers extend from the box by $1.5 \times$ the inter-quartile range (length of the box) of the 100 scenarios solved for each triple $\{N, G, x\}$. This figure shows the effectiveness of deploying more than one gateway in the network to drastically increase the percentage of DLs sent. Increasing G is especially important when N and x increase.

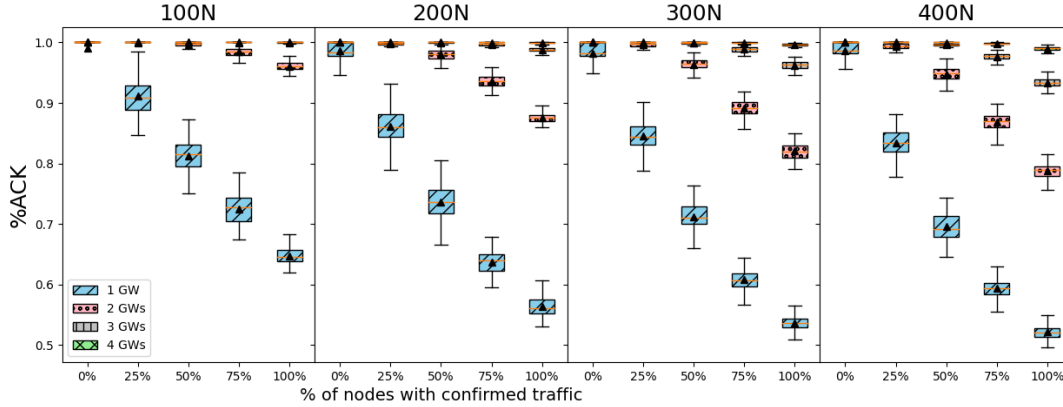


Fig. 2: Percentage of downlinks sent by the gateways.

C. Impact of RX2 Scheduling Policy


We compare here different scheduling policies in the choice of the SF used by the gateways to send DLs during the RX2 window. We consider the three following scheduling policies: (i) SF12: the gateways use SF 12 to send DLs during RX2. The default value proposed in the standard. (ii) SF9: the gateways use SF9 to send DLs during RX2, as proposed by TTN. We consider here that DLs sent with this SF can be received by the end-devices, even if they are allocated to a larger SF for an UL with this gateway. (iii) SF(ul)-2 : in this policy, we dynamically allocate the SF of the DL to $\max(SF(ul) - 2, 7)$ for the RX2 window. The last two scheduling policies are made possible in practice because the gateways use a different band for RX2 (as specified in Table I), which allows a larger transmission power, which compensates for the difference in spreading factors.

Figure 3 shows the percentage of DLs that gateways can send for the three RX2 scheduling policies, as a function of the number of end-devices, and for a single gateway (Figure 3a) or two gateways (Figure 3b). Recall that our model gives priority to DLs sent during the RX1 window over the RX2 window, meaning that we look for an optimal solution (in term of the total number of DLs sent) with the maximum number of DLs sent during the RX1 window. For a single gateway, it can be seen that the baseline policy using SF12 for RX2 can be significantly improved by using either SF9 for RX2, or SF(ul)-2. The improvement varies between 10% and 20%. For two gateways, all the DLs can be sent and the mean FDR is almost equal to 1 for all scheduling policies. The dynamic allocation of SF(ul)-2 for the DLs does not show any improvement compared to the SF9 strategy. This is due to the two following reasons that compensate each other. First, DLs corresponding to ULs made with SF12 use SF10, which improves the probability of good reception for the end-devices compared to the SF9 policy. Second, DLs corresponding to ULs made with SF7 to SF9 use SF7, which increases the probability of collisions of SF7 frames, compared to the SF9 policy.


D. Impact of the Choice of the Gateway

We study here how different scheduling policies in the choice of the gateway impacts the performance of the LoRaWAN network. We compare our baseline scenario in which any gateway in the network can send a DL, with a more restricted scheduling policy in which only the subset of gateways that correctly received the UL are allowed to schedule the corresponding DL. For each frame t_j^i , we denote by $Gr(t_j^i)$ the subset of gateways that have received t_j^i . In the constraints (7) of our model, instead of summing over all gateways $g \in G$, we restrict the sum to all $g \in Gr(t_j^i)$. In order to finely analyze the impact of the choice of the gateway, we focus on the DL policy, which is how confirmed frames are treated. Indeed, confirmed frames can be either acknowledged on RX1, acknowledged on RX2, not acknowledged due to the half-duplex property, or not acknowledged due to the duty-cycle limitation. Recall that the half-duplex property indicates that if a DL is scheduled while receiving an UL frame, then the UL is lost because the gateway cannot listen to the channel if it is sending. So an UL frame that is said to be correctly received by one or several gateways with our simulator (before scheduling the DL frames), can now be lost if all the gateways that are able to receive the UL frame are busy sending DLs. Also recall that the duty-cycle limitation indicates that gateways have to fulfill the duty cycle regulations, meaning that they have to remain silent for a certain amount of time after having sent a DL. So if a DL has to be sent but all the gateways are waiting (99% of the previous DL duration on RX1, and 90% on RX2), then the DL cannot be scheduled and is thus lost.

Figure 4 presents the distribution of the DL policy between the baseline scenario (Figure 4a) and a scenario using the restricted scheduling policy (Figure 4b), as a function of the number of end-devices. Confirmed frames that are acknowledged are depicted in light green or in light blue, depending on the fact that they are acknowledged on RX1 or RX2 respectively. Confirmed frames that are not acknowledged are depicted in red if it is due to the duty-cycle limitation (DC), or in yellow if it is due to the half-duplex property (HD). Even if




(a) One gateway

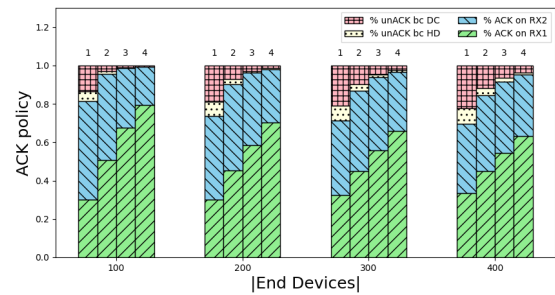


(b) Two gateways

Fig. 3: Percentage of downlinks sent by the gateways with different RX2 scheduling policies for 50% of confirmed traffic.



(a) All gateways (1, 2, 3, or 4 GW) can send the DL.



(b) Only gateways receiving ULs can send the DL.

Fig. 4: Downlink policy for the restricted scheduling policy of gateway selection (50% of confirmed traffic, SF12 on RX2).

any gateway can send the DL (see Figure 4a), the percentage of DLs that a gateway can send during the RX1 window is very limited: it reaches around 30% of the confirmed frames for one gateway, and between 50% and 60% for two gateways. Note that the duty cycle limitation does not mean that the gateway actually reaches the 1% occupation of the channel used for transmissions (especially during RX1). However, the number of confirmed frames to acknowledge can be too high during short intervals so that none of the gateways are available to send the acknowledgment at the correct time. Finally, allowing all gateways to send DL improves the overall traffic capacity.

VI. CONCLUSION

In LoRaWAN, downlink transmissions reduce network capacity up to 20%, even when they are sparse. In this paper, we proposed an optimal scheduling of DL transmissions given a realistic model of the LoRa physical layer. We analyzed the causes for dropped DL frames, and observed that the first reception window can be a priority to send DL in an optimal strategy, even though it has a duty-cycle ten times smaller than the duty-cycle of the second reception window. We studied the impact of the spreading factor of the second reception window of LoRaWAN, and the choice of the gateway sending the DLs. Overall, we believe that correctly setting up the strategy for downlink traffic can enhance the network performance, which can be evaluated through our optimization framework.¹

¹This research has received support from ANR-21-CE25-0002-01 project.

REFERENCES

- [1] Semtech, “LoRa technology.” <https://www.semtech.com/loro>, 2023.
- [2] LoRa Alliance Technical Committee, “LoRaWAN 1.1 specification,” standard, LoRa Alliance, October 2017. Final release.
- [3] J. Petäjäjärvi, K. Mikhaylov, A. Roivainen, T. Hanninen, and M. Pettissalo, “On the coverage of LPWANs: range evaluation and channel attenuation model for LoRa technology,” in *ITST*, 2015.
- [4] K. Mikhaylov, J. Petäjäjärvi, and A. Pouttu, “Effect of downlink traffic on performance of LoRaWAN LPWA networks: Empirical study,” in *IEEE PIMRC*, 2018.
- [5] M. Capuzzo, D. Magrin, and A. Zanella, “Confirmed traffic in LoRaWAN: Pitfalls and countermeasures,” in *Med-Hoc-Net*, 2018.
- [6] A.-I. Pop, U. Raza, P. Kulkarni, and M. Sooriyabandara, “Does bidirectional traffic do more harm than good in LoRaWAN based LPWA networks?,” in *IEEE GLOBECOM*, Dec. 2017.
- [7] M. Centenaro, L. Vangelista, and R. Kohno, “On the impact of downlink feedback on LoRa performance,” in *IEEE PIMRC*, 2017.
- [8] Semtech, “LoRa SX1301 datasheet v2.4,” 2017.
- [9] R. Saroui, A. Guitton, O. Iova, and F. Valois, “Uplink and downlink are not orthogonal in LoRaWAN!,” in *IEEE VTC-Fall*, Sept. 2022.
- [10] F. V. den Abeele, J. Haxhibeqiri, I. Moerman, and J. Hoebeke, “Scalability analysis of large-scale LoRaWAN networks in ns-3,” *IEEE Internet of Things Journal*, vol. 4, pp. 2186–2198, 2017.
- [11] D. Magrin, M. Capuzzo, A. Zanella, and M. Zorzi, “A configurable mathematical model for single-gateway lorawan performance analysis,” *IEEE Trans. on Wireless Comm.*, vol. 21, no. 7, pp. 5049–5063, 2022.
- [12] D. Zorbas, “Downlink spreading factor selection in lorawan,” *Elsevier Computer Communications*, vol. 215, pp. 112–119, 2024.
- [13] LoRa Alliance Technical Committee, “LoRaWAN 1.1 regional parameters,” January 2018. Revision B.
- [14] Z. Sun, H. Yang, K. Liu, Z. Yin, Z. Li, and W. Xu, “Recent advances in LoRa: A comprehensive survey,” *ACM TOSN*, vol. 18, nov 2022.
- [15] G. Callebaut and L. Van der Perre, “Characterization of LoRa point-to-point path loss: Measurement campaigns and modeling considering censored data,” *IEEE IoT Journal*, vol. 7, no. 3, pp. 1910–1918, 2020.