

# Evaluating the Scalability of LoRaWAN Gateways for Class B Communication in ns-3

Joseph Finnegan

Department of Computer Science  
Maynooth University  
Maynooth, Ireland  
Email: joseph.finnegan@mu.ie

Stephen Brown

Department of Computer Science  
Maynooth University  
Maynooth, Ireland  
Email: stephen.brown@mu.ie

Ronan Farrell

Department of Electronic Engineering  
Maynooth University  
Maynooth, Ireland  
Email: ronan.farrell@mu.ie

**Abstract**—New wireless technologies have been developed in recent years which enable applications that require the transmission of small amounts of data over long distances in an energy efficient manner. One of these technologies, LoRaWAN, includes a server-initiated communication mode named Class B which provides a deterministic latency for downlink communications. In this paper, we model Class B of LoRaWAN in ns-3 to explore the limits of scale at which this form of bi-directional communication remains feasible in large networks. The simulation results show that the principle restriction on scalability is caused by the duty cycle limits that the gateway must adhere to. In addition, we identify a limitation in the protocol which in certain configurations allows a gateway node to block the future transmission of its own beacon frames. Our contributions are the development of the first implementation and simulation of LoRaWAN Class B in ns-3, and an evaluation of the scalability limits of Class B.

## I. INTRODUCTION

In recent years, a number of new wireless technologies have been developed that allow the energy efficient transmission of small amounts of data over a long range, enabling the development of new pervasive Internet of Things applications in areas such as smart metering, smart agriculture, and smart buildings. These technologies have been designated LPWAN (Low Power Wide Area Networks).

The LPWAN technology which has attracted significant attention in research is LoRaWAN [1], which is based on LoRa [2], a physical layer technology developed by Semtech. The LoRaWAN protocol separates devices into three distinct classes. Class A supports device-initiated uplink and downlink traffic. Class B extends this by providing support for applications which require server-initiated communication and a deterministic latency on downlink communications, as regular receive windows are scheduled for potential downlink traffic. Class C is for devices without strict energy efficiency constraints, as devices remain in continuous receive mode.

Recent work has highlighted a number of potential issues with LoRaWAN, particularly in terms of scalability and QoS guarantees [3], [4]. In a typical LoRaWAN network, a gateway is expected to be able to serve thousands of end devices. To evaluate the scalability in more detail, mature and scalable simulation software is required. A widely used simulator in wireless research is ns-3. ns-3 is a discrete-event network simulator which enables the simulation of heterogeneous networks

consisting of thousands of nodes. Support is currently available for a number of wireless and wired protocols, including LTE, Wi-Fi, and IEEE 802.15.4.

A module was introduced in [5] to enable the simulation of Class A devices in ns-3. In our work we extend that module to support the simulation of Class B devices. The LoRa Alliance mandates that for a device to be LoRaWAN-compliant, only the implementation of Class A is mandatory. However, as Class B provides further downlink feedback possibilities for devices, it will provide support for evolving and reactive IoT applications, and thus requires further analysis. Using the extended module, we investigate the scalability of Class B of LoRaWAN, and our results show that the principle restriction on scalability is due to the duty cycle limits that the gateway must adhere to. We simulate an example application which requires the sending of a single 8 byte packet to every node every two and a half hours. A maximum of 64 nodes using the slowest data rate for uplink and downlink can be supported while maintaining a Packet Delivery Ratio (PDR) of >95%. Using the fastest data rate, this rises to over 1000 nodes. In addition, we identify a limitation in the protocol which in certain configurations allows a gateway node to block the future transmission of its own beacon frames.

This paper is outlined as follows: Section II describes LoRaWAN. Section III provides an overview of related work. Section IV outlines the internals of Class B of LoRaWAN, and the simulator implementation. Section V describes the experiments run and results. Section VI concludes the paper.

## II. LORAWAN

LoRa is a proprietary Chirp Spread Spectrum modulation scheme which enables long range, low data rate communication over the license-free sub-1GHz ISM bands. LoRa range depends on the link budget, which can be modified through changes in bandwidth, coding scheme, transmission power, carrier frequency, and spreading factor (SF), with six different quasi-orthogonal SFs available. Increasing the SF makes the signal more robust to noise, but decreases the data throughput.

LoRaWAN defines an upper layer protocol for LoRa. LoRaWAN is an ALOHA-based protocol which organises networks in a star-of-stars topology, where End Devices communicate only with their nearest Gateway, and Gateways

act purely as relays to a central Network Server. LoRaWAN organises sets of SFs and bandwidth into defined data rates (DR); these are outlined in Table I. LoRaWAN also defines an Adaptive Data Rate (ADR) scheme which optimises the choice of data rate for each node based on its recent packet delivery rate. The LoRaWAN overhead per packet is 13 bytes.

TABLE I: LoRaWAN Data Rates for EU863-870

DR	SF	Bandwidth	Bit rate	Receiver Sensitivity
0	12	125kHz	250	-136dBm
1	11	125kHz	440	-133dBm
2	10	125kHz	980	-132dBm
3	9	125kHz	1760	-129dBm
4	8	125kHz	3125	-126dBm
5	7	125kHz	5470	-123dBm
6	7	250kHz	11000	-120dBm

As transmission is in the unlicensed ISM bands, LoRaWAN devices must follow the regulations defined by the region the network is deployed in. In Europe devices are required to only transmit a maximum of a certain percentage of the time (with the limit depending on the particular subband of the channel); these regulations are described in Table II. In the ALOHA-based Class A of LoRaWAN, End Devices receive downlink frames through the use of two receive windows, which are by default opened in the following two seconds after the completion of the transmission. The Network Server may choose to respond in either of the two receive windows, if there is downlink data to send.

TABLE II: ETSI Spectrum Access per subband [6]

Subband	Spectrum Access	Edge Frequencies	Max EIRP
g	1 % or LBT AFA	865-868MHz	10mW
g1	1 % or LBT AFA	868-868.6MHz	25mW
g2	0.1% or LBT AFA	868.7-869.2MHz	25mW
g3	10 % or LBT AFA	869.4-869.65MHz	500mW
g4	No Requirement	869.7-870MHz	5mW
g4	1 % or LBT AFA	869.7-870MHz	25mW

### III. RELATED WORK

#### A. Scalability of LoRaWAN

LoRaWAN spreading factors have been shown to not be fully orthogonal, and so co-spreading factor interference limits the scalability of networks [7], but LoRaWAN does scale better than pure ALOHA due to the potential receive of both colliding packets under certain conditions, and the existence of the capture effect [8]. The ADR feature has been shown to be essential for scalable Class A uplink communication [9]. Some identified potential scalability issues in the protocol are the size of the LoRaWAN overhead (especially for applications with small payloads), and ambiguity in following: the retransmission policy, the selection of the DR for retransmissions, and the choice of DR in the Class A downlink [10].

#### B. Downlink in LoRaWAN

In Class A networks, the most common cause of lost downlink packets is due to the duty cycle regulations applied to the Gateway, and because of this, the lack of a requested

acknowledgement to an uplink frame does not usually indicate poor link quality. This lack of reliable knowledge about the link quality by End Devices can thus cause a large drop in network goodput through the use of redundant retransmissions if the retransmission limit of devices per packet is not set appropriately [11]. In addition, LoRaWAN network performance severely degrades when all downlink frames are required to be confirmed with an acknowledgement frame [12].

#### C. Class B of LoRaWAN

Research focusing on Class B is to our knowledge currently limited to a single paper [13], where the delivery delay of acknowledged Class B LoRaWAN packets is modelled analytically based on a Markov chain. The duty cycle regulations and uncoordinated scheduling of receive windows are highlighted as factors which may limit the scalability of Class B networks.

#### D. LoRaWAN Simulation Tools

Simulation tools for LoRaWAN are currently under active development. Modules implementing the physical layer of LoRa and Class A of LoRaWAN have been developed in ns-3 [5], [14], [15], and OMNET++ [16]. Simulations of large scale networks have been performed, with [5] in particular focusing on downlink feedback, [14] on urban environments, and [16] on the modelling of the ADR algorithm.

### IV. CLASS B SIMULATION

Class B of LoRaWAN is designed to enable server-controlled bi-directional communication while still maintaining a device lifetime suitable for IoT devices. End Devices schedule receive windows, referred to as “ping slots”, in which to potentially receive downlink frames. This requires precise time synchronisation across the network, and so every Gateway periodically broadcasts a timing beacon, and Class B devices synchronise their clocks to the included timestamp. Devices then schedule ping slots in the time between beacons. At each occurrence of a ping slot the device reads for an incoming frame, and the Network Server can potentially send a downlink frame. For uplink the Class B device transmits using Class A-style channel access, with just one bit changed in the frame header to indicate the device class.

Class A simulation support [5] has been extended to support Class B for this research<sup>1</sup>. In particular, the simulator’s MAC layer (which acts as the driver of the PHY layer and is shared by both LoRaWAN End Devices and Gateways) and the application layers (which implement the Network Server, Gateway, and End Device-specific functionality) have been modified. As Class B is a purely MAC-level mode, minimal changes to the PHY layer were required.

#### A. Modifications to Network Server & Gateway

In the implementation described in [5], the Network Server is a singleton object which encapsulates the functionality of a LoRaWAN Network Server, Application Server and Join

<sup>1</sup><https://github.com/ConstantJoe/ns3-lorawan-class-B>

Server, as well as performing the generation of downlink traffic.

To implement Class B, the information the Network Server stores about each End Device has been extended to include each device's ping periodicity, and Class B-specific Channel, Data Rate, and Code Rate. As the join procedure is not yet modelled in the simulator, an initial Gateway is also assigned to every End Device at the start of each simulation, providing the Network Server a prospective Gateway to use to send Class B downlink frames to that device.

The exact periodicity of beacons is defined as being every 128s since the beginning of GPS time (i.e. 00:00:00 6th Jan 1980) plus a delay of  $1.5ms \pm 1\mu s$ . Gateway beacons are thus synchronised across all LoRaWAN networks. Every 128s, a beacon is received by the End Device from its nearest Gateway. The next 2.120s are reserved for the frame Time-On-Arrival, and for potential network management frames. The following 122.880s is the beacon window period, which is split into  $2^{12} = 4096$  30ms ping slots. The final 3s constitutes a guard period. The format of the beacon period is shown in Fig. 1.




Fig. 1: Beacon Period Format

To implement this, an additional timer is added to the Network Server, which, if the device is in Class B mode, fires every 128 seconds and causes the Network Server to generate the beacon frame.

Preamble (10 symbols)	NetID (3 bytes)	Time (4 bytes)	CRC (1 byte)	GwSpecific (7 bytes)	CRC (2 bytes)
--------------------------	--------------------	-------------------	-----------------	-------------------------	------------------

Fig. 2: Class B Beacon Frame Format

The beacon format in Europe is displayed in Fig. 2, and consists of a preamble of 10 unmodulated symbols, followed by a 17 byte payload without a LoRa physical header or CRC. The *Time* field is the least significant 32 bits of GMT time. *GwSpecific* is Gateway-dependent, and may contain location data for the closest Gateway. For stationary devices, the timestamp is enough to continue to operate in Class B mode. In the simulator, the Network Server generates the leading eight bytes of the beacon frame, which is passed to the individual Gateways which add the Gateway-specific information and make a call to the underlying MAC layer to send the packet using the parameters defined in the LoRaWAN standard for beacon sends. In Europe, beacons are transmitted using *DR3* and a coding rate of  $4/5$ . The default broadcast

frequency is  $869.525MHz$ , which is inside the subband with the highest spectrum access allocation (10%) and highest maximum effective radiated power (500mW).

Next, the Network Server calculates the ping slots to be used by each device in the beacon period. The number of slots for each device is prearranged by the Network Server, and must be a power of two, up to a maximum of 128 (i.e. a ping slot every second):

$$period = 2^{32}/slots \quad (1)$$

The initial offset value  $O$  for a device in a beacon period is calculated using the first two bytes of the result of applying the AES-128 encryption algorithm to a padded block consisting of the beacon timestamp (*Time*) and the device address (*DevAddr*), using a key of all zeroes:

$$R = AES128\_enc(16*(0x00), Time|DevAddr|pad16) \quad (2)$$

$$O = (R[0] + R[1] * 256)\%period \quad (3)$$

As an End Device has no knowledge of the scheduled ping slots of other devices, the inclusion of the timestamp in the Equation 2 prevents continuous collisions between devices across beacon periods. The timings of ping slots are then equally spaced in the broadcast period, starting from the offset:

$$timings = \{O + x * period \mid x < slots, x \in \mathbb{N}\} \quad (4)$$

In the simulator, each Gateway has an array of 4096 vectors and the Network Server appends the End Device ID to the  $i$ th vector in order to signify that that device is to potentially have a frame sent to it in the  $i$ th ping slot of the beacon period. A new event is then generated to fire at the time of that ping slot.

During each event, if the Network Server has a data packet queued for the associated device, and if there is no device ahead of that device in the ping slot vector with a data packet queued, and if the Gateway may immediately send on the associated channel (i.e. the Gateway is not currently transmitting or receiving, and can send without breaking the duty cycle limits), then the queued packet is transmitted. By default, in Europe downlink Class B communications are sent at a frequency of  $869.525MHz$ , the same channel as the beacon frames.

### B. Modifications to End Device

The End Device application class has been modified to include the ping periodicity of each End Device, as well as the Class B-specific channel, Data Rate, and Code Rate. If a device is in Class B, at the start of a simulation an event is fired using a timer every 128 seconds to put the device into receive mode, in order to receive the incoming beacon frame.

If a beacon is received, the timestamp is used to generate the exact timings of the pingslots to be used in this beacon period, in an identical fashion to the Network Server. The

device then schedules events for each of those times, which cause the device to wake up to potentially receive a downlink frame.

### C. Modifications to MAC Layer

The MAC layer of the simulator has also been modified to implement Class B behaviour. Two additional receive states have been added to the End Device, for the beacon and Class B downlink receives respectively. The MAC layer state changes can be seen in Fig. 3, with new states shown in red.




Fig. 3: MAC layer FSM changes

The new Finite State Machine (FSM) reflects the uncoordinated nature of Class A and Class B communication in a LoRaWAN device. An End Device may only attempt to receive a beacon or downlink ping when it is in the idle state. Similarly, a device may not attempt to send an uplink packet when currently receiving a Class B downlink or beacon frame.

The transition to *RX\_BEACON* or *RX\_CLASS\_B* causes the device to attempt to receive the expected length of LoRa symbols as the preamble of the incoming packet, on the pre-set channel and data rate. If the preamble is detected, the packet is received and handled based on the MAC header and the device transitions to idle mode. If no preamble is detected the device transitions directly into idle mode.

### D. Modifications to PHY Layer

No changes have been made to the PHY layer of the simulator, beyond the use of a longer preamble for beacon frames.

## V. SCALABILITY ANALYSIS

The developed simulator was used to assess the performance of a LoRaWAN network with various downlink requirements.

Each experiment consists of a LoRaWAN network with one Gateway and a virtual Network Server. A number of LoRaWAN nodes operating in Class B mode are created and randomly placed within a Gateway-centred disk with a radius of 6100m (enabling transmissions from nodes using any possible DR to reach the Gateway). The analysed throughput requirements were an uplink data generation period of 900s, and a downlink period of 9000s. This is consistent with the expected throughput requirements of LoRaWAN applications, where uplink forms the vast majority of network traffic. 8

byte payloads are transmitted, both in the uplink and downlink. Packets are generated once randomly inside each data generation period. Simulations were ran for one day.

In the simulations, all downlink communication takes place using the Class B ping slots (i.e. no use of Class A-style receive windows). In each simulation, the chosen DR is equal across all nodes and remains static for the duration of the experiment. All other LoRaWAN parameters were set to the default in the standard.

As shown in Table II, the spectrum access and max transmission power of a LoRaWAN channel depends on the subband. In Europe, the LoRa Alliance mandates that 3 channels inside the g1 subband must be implemented. In addition, by default beacon frames and Class B downlink frames are transmitted in a channel inside the g3 subband. The core of our analysis focuses on how the protocol performs when using these subbands. The particular aspects of LoRaWAN Class B to be analysed were as follows:

- The scalability of a network of LoRaWAN Class B devices, where downlink pings are sent in a channel in a subband with a duty cycle limit of 1% (e.g. g1).
- The scalability of a network of LoRaWAN Class B devices, where downlink pings are sent in a channel in the subband with a duty cycle limit of 10% (i.e. g3), which is shared with the transmission of beacon frames.

### A. Transmission in channel with 1% limit

Fig. 4 shows the percentage of Class B beacons that were received by devices. It can be seen that, due to the transmission of beacon frames in a separate channel, these frames are reliably sent by the Gateway and received by End Devices. Any failed reception of beacon frames in this case are due to the simultaneous transmission by an End Device during the time in which a beacon frame is due to be received. For each DR, the *delivered throughput* is consistently  $>95\%$ , as can be seen in Fig. 5; End devices wake up at the correct time to receive packets, and packets sent by the Gateway are successfully delivered. Lost packets are due to a “collision” between the ALOHA-based MAC layer of the uplink and the scheduled traffic of the Class B downlink on the End Device (i.e. an End Device attempts to transmit an uplink frame during the time in which a Class B downlink packet is due to be received), with a nominal amount due to on-air collisions.

However, as can be seen in Fig. 6, in terms of the actual *generated traffic* the PDR drops quickly with scale. As the number of nodes increases, the number of downlink packets queued for transmission increases proportionally, but the number of packets actually transmitted remains limited by the duty cycle regulations. By comparing Figs. 5 and 6, it can be seen that the vast majority of lost packets are never actually sent. Because of the infrequency of IoT traffic, the network is not yet dense enough that on-air collisions causes a significant drop in the PDR, but the network is large enough that the gateway has reached the duty cycle limits of the band and can no longer transmit. For this particular traffic pattern, only a maximum of about 64 Class B nodes can be supported while

maintaining a PDR of  $>90\%$  when the slowest DR is used. By comparison, for the same traffic pattern the fastest DR can support over 1000 nodes. As can be seen in Fig. 7, the PDR of uplink frames remains high, even with ten times the amount of uplink data being sent compared to downlink.




Fig. 4: Beacon Delivery Ratio - 1% band




Fig. 5: PDR of sent downlink packets - 1% band

### B. Transmission in channel with 10% limit

In this case the single 10% channel is shared for the transmission of beacon frames and Class B downlink frames, enabling the Gateway to spend a longer amount of time transmitting. When the channel is shared with the beacon frames, in larger networks (where there is regularly data queued at the Gateway for transmission) the Gateway may transmit a frame that causes a time-off-band period that extends to when the next beacon frame is due to be transmitted. The beacon guard between beacon periods prevents this for the faster data rates, but is not sufficiently long enough to prevent this blocking of beacon transmissions for the slower data rates. This is shown by the notch in the curve of Fig. 8. The Class B definition




Fig. 6: PDR of generated downlink packets - 1% band

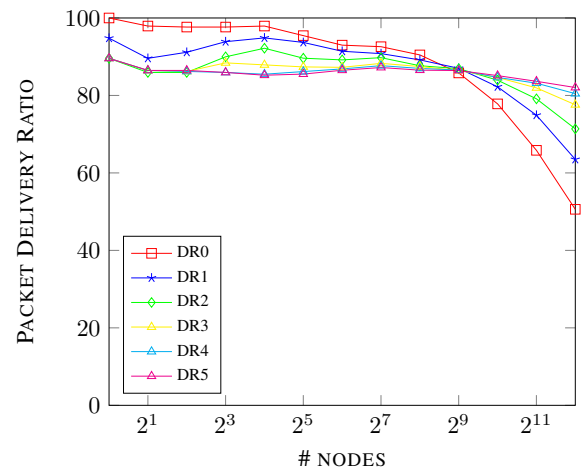


Fig. 7: PDR of sent uplink packets - 1% band

currently does not prevent a Gateway from sending a downlink ping that will block itself from transmitting the next beacon. When no beacon is sent, no ping slots on devices are scheduled so many packets are lost, as indicated in Fig. 9. However, as can be seen in Fig. 10 despite this the overall PDR of generated packets is higher because of the greater amount of time that can be spent transmitting by the Gateway.

## VI. CONCLUSIONS

LoRaWAN is a recently developed prominent LPWAN technology. In this paper, we concentrate on the downlink-focused Class B of the protocol. We implement Class B of LoRaWAN in ns-3 to enable the simulation of large scale networks that require server-initiated communication and deterministic downlink latency. Our simulation results indicate that the primary bottleneck in the transmission of downlink frames in Class B mode is the duty cycle regulations. In addition, while the use of the default g3 subband with a higher spectrum access limit does provide a greater maximum throughput, we identify that currently in the protocol there is no mechanism




Fig. 8: Beacon Delivery Ratio - 10% band

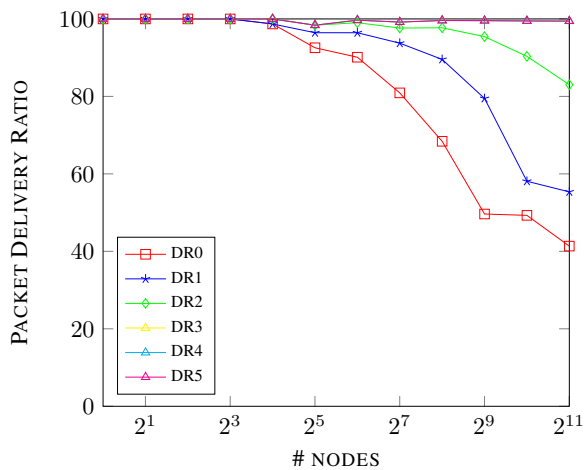


Fig. 9: PDR of sent packets - 10% band

to prevent gateways from transmitting packets in this subband that will block the future transmission of beacon frames. Our uplink frame results are consistent with previous Class A simulation results. As future work, we plan to extend the simulator to model the join procedure of LoRaWAN and Class B transition mechanism, in addition to the MAC layer commands. We also plan to integrate an energy modelling framework into our module, to enable the simulation of the energy consumption of LoRaWAN devices in Class A and Class B.

#### ACKNOWLEDGMENT

This publication has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) and is co-funded under the European Regional Development Fund under Grant Number 13/RC/2077.

#### REFERENCES

- [1] LoRa Alliance. LoRaWAN Specification v1.1, 1 2017.
- [2] A. Augustin, J. Yi, T. Clausen, and W. Townsley. A Study of LoRa: Long Range & Low Power Networks for the Internet of Things. *Sensors*, 16(9), 2016.

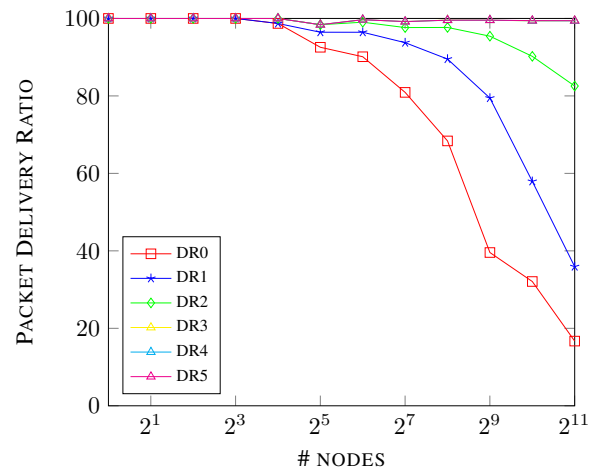


Fig. 10: PDR of generated packets - 10% band

- [3] F. Adelantado, X. Vilajosana, P. Tuset, B. Martinez, J. Melia-Segui, and T. Watteyne. Understanding the Limits of LoRaWAN. *IEEE Communications Magazine, Publication: Article in proceeding*, June 2017.
- [4] E. De Poorter, J. Hoebeke, M. Strobbe, I. Moerman, S. Latre, M. Weyn, B. Lannoo, and J. Famaey. Sub-GHz LPWAN Network Coexistence, Management and Virtualization: An Overview and Open Research Challenges. *Wireless Personal Communications*, 95(1):187–213, 7 2017.
- [5] F. Van den Abeele, J. Haxhibeqiri, I. Moerman, and J. Hoebeke. Scalability Analysis of Large-Scale LoRaWAN Networks in ns-3. *IEEE Internet of Things Journal*, 4(6):2186–2198, Dec 2017.
- [6] ETSI. Electromagnetic compatibility and Radio spectrum Matters; Short Range Devices; Radio equipment to be used in the 25 MHz to 1 000 MHz frequency range with power levels ranging up to 500 mW; Part 1: Technical characteristics and test methods, Jan 2012.
- [7] O. Georgiou and U. Raza. Low Power Wide Area Network Analysis: Can LoRa Scale? *IEEE Wireless Communications Letters*, PP(99):1–1, 2017.
- [8] Jettmir Haxhibeqiri, Floris Van den Abeele, Ingrid Moerman, and Jeroen Hoebeke. LoRa Scalability: A Simulation Model Based on Interference Measurements. *Sensors*, 17(6), 2017.
- [9] M. Bor, U. Roedig, T. Voigt, and J. Alonso. Do LoRa Low-Power Wide-Area Networks Scale? In *Proceedings of the 19th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems, MSWiM '16*, pages 59–67, 2016.
- [10] D. Bankov, E. Khorov, and A. Lyakhov. On the Limits of LoRaWAN Channel Access. In *2016 International Conference on Engineering and Telecommunication (EnT)*, pages 10–14, Nov 2016.
- [11] A. Pop, U. Raza, P. Kulkarni, and M. Sooriyabandara. Does Bidirectional Traffic Do More Harm Than Good in LoRaWAN Based LPWA Networks? In *GLOBECOM 2017 - 2017 IEEE Global Communications Conference*, pages 1–6, Dec 2017.
- [12] M. Centenaro, L. Vangelista, and R. Kohno. On the impact of downlink feedback on LoRa performance. In *2017 IEEE 28th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*, pages 1–6, Oct 2017.
- [13] F. Delobel, N. El Rachkidy, and A. Guitton. Analysis of the delay of confirmed downlink frames in Class B of LoRaWAN. In *85th Vehicular Technology Conference (VTC)*, Sydney, Australia, June 2017. IEEE.
- [14] D. Magrin, M. Centenaro, and L. Vangelista. Performance evaluation of LoRa networks in a smart city scenario. In *2017 IEEE International Conference on Communications (ICC)*, pages 1–7, May 2017.
- [15] Brecht Reynders, Qing Wang, and Sofie Pollin. A LoRaWAN Module for Ns-3: Implementation and Evaluation. In *Proceedings of the 10th Workshop on Ns-3, WNS3 '18*, pages 61–68. ACM, 2018.
- [16] Mariusz Slabicki, Gopika Preamsankar, and Mario Di Francesco. Adaptive Configuration of LoRa Networks for Dense IoT Deployments. In *2018 IEEE/IFIP Network Operations and Management Symposium (NOMS)*, 2018.