



HAL
open science

Concilier débit et efficacité énergétique dans les réseaux étendus à basse consommation

Laurent Chasserat

► To cite this version:

Laurent Chasserat. Concilier débit et efficacité énergétique dans les réseaux étendus à basse consommation. Library and information sciences. Université Paul Sabatier - Toulouse III, 2023. English. NNT : 2023TOU30171 . tel-04347640

HAL Id: tel-04347640

<https://theses.hal.science/tel-04347640v1>

Submitted on 15 Dec 2023

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.

Université Fédérale



Toulouse Midi-Pyrénées

THÈSE

En vue de l'obtention du

DOCTORAT DE L'UNIVERSITÉ FÉDÉRALE TOULOUSE MIDI-PYRÉNÉES

Délivré par :

l'Université Toulouse 3 Paul Sabatier (UT3 Paul Sabatier)

Présentée et soutenue le *02/10/2023* par :

LAURENT CHASSERAT

**Trading-off Throughput and Energy-efficiency in Long Range Wide
Area Networks**

JURY

D. DRAGOMIRESCU	Professeur d'Université	Président du Jury
F. THEOLEYRE	Directeur de Recherche	Membre du Jury
N. MITTON	Directeur de Recherche	Membre du Jury
S. CÉSPÉDES	Professeur Assistant	Membre du Jury
P. BERTHOU	Maître de Conférences	Membre du Jury
N. ACCETTURA	Chargé de Recherche	Membre du Jury

École doctorale et spécialité :

MITT : Domaine STIC : Réseaux, Télécoms, Systèmes et Architecture

Unité de Recherche :

Laboratoire d'analyse et d'architecture des systèmes

Directeur(s) de Thèse :

Pascal BERTHOU et Nicola ACCETTURA

Rapporteurs :

Fabrice THEOLEYRE et Nathalie MITTON

Remerciements

Alors que ces quatre longues années de thèse touchent à leur fin, me voilà arrivé au moment d'écrire cette section si particulière des remerciements. Ces deux pages sont évidemment bien trop courtes pour mentionner toutes les personnes qui m'ont aidé durant ce périple, je vais cependant tâcher d'être aussi exhaustif que possible. Je n'en aurais pas vu le bout sans votre soutien, il me paraît donc juste que vous figuriez en préambule de ce manuscrit qui est aussi le vôtre.

Je tiens tout d'abord à remercier mes directeurs de thèse Nicola Accettura et Pascal Berthou pour leur confiance et la qualité de leur encadrement. Votre exigence aura su me pousser à donner le meilleur de moi-même, et je suis très fier du travail que nous avons accompli ensemble. Je remercie également tous les membres du jury pour le temps et la rigueur investis dans l'évaluation de mon travail de thèse. Je tiens aussi à exprimer ma gratitude envers le LAAS-CNRS et tous les membres permanents de l'équipe SARA pour leur accueil chaleureux dans cet environnement scientifique de qualité. Enfin, merci à tous les doctorants et stagiaires qui ont partagé mon quotidien au laboratoire, au gré de discussions philosophico-techniques et de moments de convivialité.

Je remercie également ma famille dont le soutien n'a jamais failli, et ce même si je n'ai pas été très présent ces dernières années. Tout d'abord merci à toi ma maman chérie, pour la force et le courage dont tu fais preuve depuis toujours. Tu m'as élevé presque seule, me protégeant sans cesse des aléas de la vie, et grâce à toi j'ai pu grandir dans les meilleures conditions. Cette thèse est bien dérisoire en comparaison mais j'espère qu'elle te rendra fière, tout comme j'aime imaginer qu'elle aurait rendu papa fier s'il était encore parmi nous. Merci à vous Julien et Manu, mes deux grands frères adorés. Vous avez au fil des années endossé les rôles de pères de substitution, de confidents, de professeurs, de camarades de jeu, et n'avez jamais cessé d'être mes modèles, mes héros. Vous avez fait de moi qui je suis aujourd'hui, et il m'est difficile de trouver des mots assez forts pour exprimer l'amour que je vous porte ou la reconnaissance que j'ai de vous avoir comme mentors. Cela étant dit, un autre grand merci à Marion et Stéphanie pour votre soutien, ainsi que pour votre capacité à supporter mes frères au quotidien ! ;-)

Je tiens aussi à adresser une pensée à tous mes amis de longue date pour lesquels j'ai été moins disponible durant ce doctorat. Merci à Clément, Coco et TBZ de m'avoir supporté pendant les années du collège et du lycée, et merci à Victor et Mathieu d'avoir pris le relais sur les bancs de l'INSA. Un autre grand merci à Paul Pidou, qui voyant mon intérêt naissant pour la recherche m'a pris sous son aile le temps d'un stage, et à l'issue d'innombrables discussions métaphysiques m'a finalement aidé à rassembler le courage de devenir docteur. Enfin, merci à toute la gang de Montréal pour le semestre de folie qui a précédé la thèse! J'espère que vous saurez pardonner mon absence, et sachez que j'ai bien l'intention de rattraper

le temps perdu !

Vient maintenant le moment de remercier les personnes qui m'accompagnent au quotidien, et qui m'ont plus que quiconque aidé à surpasser les difficultés de la thèse, à savoir mes colocataires Théis, Yoyo, Paulo, Tonio et les deux Simon. Plus que des amis, vous êtes ma famille d'adoption, et j'ai une chance infinie de pouvoir jour après jour faire face à la vie au sein d'une telle équipe. Il est clair que je n'aurais jamais terminé ce manuscrit sans votre soutien, alors pour ça, et (surtout) pour tout le reste, merci. Je vous aime, et j'ai hâte de découvrir toutes les belles histoires que l'avenir nous réserve.

Je ne peux cependant guère m'arrêter aux portes de la Suze ! En effet, comment oublier toutes les personnes exceptionnelles qui gravitent autour, qu'ils soient de la Smooth ou de l'Ormeau, de Komet ou de S.oror, de la zik ou de la grimpe, Touloutchazins ou bretons exilés ? Chacun à votre manière, vous enchantez ma vie, et vous compter parmi mes amis est ma plus grande fierté.

Enfin, merci à toi Juliette. Ces quatre années de thèse ont certes été laborieuses, mais à tes côtés la fin de cette aventure fut emplie de bonheur.

Abstract

Low-Power Wide-Area Networks (LPWAN) connect a large number of battery-powered wireless devices over long distances. Replacing the batteries of low-power devices spread over a large geographical area is difficult and costly, their lifetime should therefore be maximized. Among available LPWAN technologies, Long Range Wide Area Networks (LoRaWAN) have prevailed in the past years, offering an excellent tradeoff between coverage range, data rate, power consumption and ease of use. This technology implements a pure ALOHA medium access scheme to save device energy by minimizing the radio usage. However, using such a simple strategy induces frame collisions that restrain the network's ability to handle high traffic loads. In this context, synchronization can be used to exploit the available bandwidth more efficiently by controlling the timing of frame transmissions and reducing the collision probability. In that, this manuscript first introduces a slotted scheme specifically tailored for LoRa networks. This design allows the implementation of a simple slotted ALOHA strategy, which increases the network capacity at the cost of an extra energy demand due to the inherent overhead. As a first research outcome, we show that the synchronized version can also be more energy efficient than pure ALOHA in congested networks if the power consumption dedicated to synchronization beacon receptions is handled carefully. Yet, the legacy access remains preferable in low traffic situations. This finding highlights the need for traffic-aware protocols able to adapt the access strategy to the network load. As a result, this thesis then introduces a switching mechanism capable of dynamically synchronizing or desynchronizing the network depending on the probed conditions. Through simulations, we show that such a system is able to follow the variations of the input load to always select the most energy-efficient strategy. We then bridge the gap between simulations and experiments with the implementation of an energy-efficient synchronization mechanism on real hardware. The robustness of our approach is proven thanks to an experimental study of the device clock drift. Hardware experiments additionally reveal that common ALOHA models are unable to accurately represent the real-world network performances, notably due to occurrences of the capture effect. As a final contribution, we thus fall back to the mathematical approach and include the features observed in real networks into our performance evaluation models. This ultimate step allows to *(i)* extend the protocol adaptation capabilities by showing how the slot size can be adjusted to the network load to maximize energy efficiency and *(ii)* account for capture effect occurrences to faithfully picture the testbed throughput. In a nutshell, this thesis explores trade-offs between throughput and energy efficiency in LoRa networks, and highlights the need for traffic-aware protocols. We additionally show that state-of-the-art performance models do not fully grasp the behavior of real networks yet, which opens future research on more accurate and versatile models to drive such adaptive access strategies.

Keywords: LPWAN, LoRaWAN, MAC Protocols, Scalability, Energy Efficiency, Adaptability, Synchronization, Capture Effect

Résumé

Les réseaux étendus à basse consommation (LPWAN pour Low-Power Wide-Area Networks) connectent un grand nombre de dispositifs sans fil fonctionnant sur batterie et disséminés sur des zones étendues. Le remplacement des batteries d'alimentation de tels objets connectés étant coûteux et difficile, il devient crucial de maximiser leur espérance de vie. Parmi les technologies de type LPWAN, les réseaux à longue portée (LoRaWAN pour Long Range Wide Area Networks) se sont imposés en tête du marché, offrant un excellent compromis entre portée de couverture, débit, consommation d'énergie et facilité d'implémentation. Cette technologie fait appel à une stratégie d'accès au médium de type pure ALOHA, qui permet d'économiser de l'énergie en minimisant l'utilisation des radios des terminaux. Cependant, l'utilisation d'une telle stratégie implique un fort taux de collision de paquets qui limite la capacité du réseau à supporter de fortes charges de trafic. Dans ce contexte, la synchronisation est une piste de recherche intéressante permettant de faire un meilleur usage de la bande passante en réduisant les probabilités de collision. En cela, ce manuscrit aborde tout d'abord la conception d'un mécanisme d'accès synchrone adapté aux réseaux LoRa. Ce mécanisme permet la mise en place d'une stratégie de type slotted ALOHA, qui étend la capacité réseau au prix d'une hausse de la consommation d'énergie dédiée au mécanisme de synchronisation. Le premier aboutissement de cette recherche est la découverte que l'accès synchrone peut aussi être plus énergétiquement efficace que pure ALOHA dans les réseaux congestionnés si la réception des balises de synchronisation est gérée intelligemment. Cependant, le protocole de base reste préférable pour les faibles charges de trafic. Cette découverte met en évidence le besoin d'un protocole conscient du trafic capable d'adapter le mode d'accès à la charge réseau. En conséquence, un mécanisme capable de dynamiquement synchroniser et désynchroniser les terminaux est introduit afin de sélectionner le protocole le plus efficace en fonction des conditions de trafic sondées. Au travers de simulations, nous montrons qu'un tel système est capable de suivre les variations de la charge réseau et ainsi de maximiser l'efficacité énergétique du déploiement. Nous comblons ensuite le fossé entre simulations et expériences réelles en implémentant un mécanisme de synchronisation efficace sur un banc de tests. La robustesse de notre approche repose sur une analyse expérimentale du décalage d'horloge des terminaux. Les résultats issus du banc de tests révèlent additionally que les modèles mathématiques les plus communément utilisés pour décrire le comportement des réseaux ALOHA sont très imprécis dans notre cas de figure, car l'effet capture survient très fréquemment. En guise de dernière contribution, nous proposons donc une approche mathématique incluant les effets observés sur le banc de tests à nos modèles d'évaluation de performances. Cette dernière étape permet (i) d'étendre les capacités d'adaptation du protocole en montrant comment la taille des slots peut être ajustée à la charge réseau pour maximiser l'efficacité énergétique, et (ii) de considérer l'effet capture pour modéliser fidèlement le débit réel du banc de tests. En un mot, cette thèse explore des compromis entre débit et efficacité énergétique dans les réseaux LoRa, et met en évidence le besoin

de protocoles conscients de la charge de trafic. Nous montrons additionnellement que les modèles d'évaluation de performances de l'état de l'art ne représentent pas encore parfaitement le comportement des réseaux réels, ce qui ouvre la voie vers de nouvelles pistes de recherche pour améliorer cette modélisation et ainsi piloter de tels protocoles adaptatifs.

Mots clés : *LPWAN, LoRaWAN, Protocoles de Contrôle d'Accès au Médium, Passage à l'Echelle, Efficacité Energétique, Adaptabilité, Synchronisation, Effet Capture*

Contents

Glossary	xv
Introduction	1
1 Technological background and thesis problematic	7
1.1 The LoRa physical layer and its features	9
1.1.1 The Chirp Spread Spectrum modulation	9
1.1.2 The physical-layer packet structure and Time on Air	10
1.1.3 The low-power and long-range capabilities of LoRa compared against other LPWAN technologies	11
1.1.4 Experimental insights on the LoRa transmission range	13
1.2 The LoRaWAN MAC-layer specification	15
1.2.1 The LoRaWAN classes	15
1.2.2 Regulations	17
1.2.3 The MAC-layer frame format	18
1.2.4 The LoRaWAN data rates and Adaptive Data Rate mechanism	19
1.3 The challenge of trading-off throughput and energy efficiency when designing a MAC layer scheme for LoRaWAN	20
2 MAC strategies for LoRa and Thesis Methodology	23
2.1 Overview of Medium Access Control strategies	24
2.2 State-of-the-Art Research on Access Strategies for LoRa	27
2.2.1 Asynchronous protocols and Carrier Sense Multiple Access schemes	27
2.2.2 Synchronous protocols: Slotted ALOHA and Scheduled Ac- cess Schemes	29
2.3 Methodological approach	36
3 From Theory to Simulations: designing a slotted access that fits LoRaWAN	39
3.1 Theoretical motivation: Pure and Slotted ALOHA throughput in ideal networks	41
3.2 From Class B to Class S: defining a slotted framework tailored for LoRa communications	42
3.2.1 Protocol design	44
3.2.1.1 Advantages of beacon-based synchronization	44
3.2.1.2 Slotframe structure	44
3.2.2 Simulation-based performance evaluation of a Slotted ALOHA scheme over Class S	46
3.2.2.1 Simulation setup	46

3.2.2.2	Simulation results	47
3.3	TREMA: a traffic-aware MAC protocol that maximizes energy efficiency in LoRa networks	49
3.3.1	Design of TREMA	50
3.3.1.1	A time-synchronized scheduled access over Class S	51
3.3.1.2	Prior fingerprinting of the deployment	52
3.3.1.3	Network probing and decision mechanism	53
3.3.1.4	Signaling protocol	54
3.3.2	Performance evaluation	56
3.3.2.1	Preliminary deployment fingerprinting	58
3.3.2.2	Online testing	59
3.3.2.3	Overall performance evaluation	61
3.3.2.4	Switching duration	61
3.3.2.5	Performance gain	61
3.4	Conclusions and future works	62
4	Experimental Approach: Achieving Synchronization under real-world constraints	63
4.1	LoRaSync requirements and design	64
4.1.1	Device drift measurement and modeling	65
4.1.2	Spreading Factor selection	66
4.1.3	Payload size standardization	67
4.1.4	LoRaSync design	67
4.2	Proof of Concept on real hardware	69
4.2.1	Architectural challenges	70
4.2.2	Software challenges	70
4.2.3	Clock correction demonstration and discussion	72
4.2.4	Impact of the Capture Effect on the testbed performances	72
4.3	Conclusion and perspectives	73
5	Back to Theory: Accurate Modeling to optimize the protocol performances	75
5.1	Modeling and optimizing LoRaSync	77
5.1.1	Model validation and analysis	79
5.1.2	Slot size optimization	81
5.2	Modeling the Capture Effect	83
5.2.1	Background on the Capture Effect	83
5.2.2	Throughput modeling with capture	85
5.2.2.1	Transmission scenarios and their occurrence probabilities	86
5.2.2.2	Experimental setup and capture coefficient measurements	87
5.2.3	Experimental results and discussion	88
5.2.3.1	Model validation with experimental data	88

5.2.3.2	Impact on the network fairness	89
5.2.3.3	Discussion	91
5.3	Conclusions and perspectives	91
6	Conclusion and Future Works	93
6.1	Thesis contributions	94
6.2	Lessons learned	95
6.3	Research perspectives	96
6.3.1	Improving the fingerprinting process	96
6.3.2	TREMA improvements and implementation on the testbed	96
6.3.3	Towards more accurate and versatile throughput models	97
A	Performance models of MAC protocols	99
A.1	Throughput models	99
A.1.1	Throughput models for theoretical Random Access Protocols	99
A.1.1.1	Infinite number of devices	99
A.1.1.2	Finite number of devices	100
A.1.2	Throughput models for LoRa Access schemes	100
A.2	Power consumption models	101
A.3	Energy efficiency models	102
	Scientific production	103
	Bibliography	105

List of Figures

1.1	Topology of a LoRaWAN deployment	8
1.2	Chirp Spread Spectrum modulation example with 2 bits per symbol	9
1.3	LoRa physical frame structure	10
1.4	RSSI and range	14
1.5	The LoRaWAN <i>Class A</i>	15
1.6	The LoRaWAN <i>Class B</i>	17
1.7	The LoRaWAN <i>Class C</i>	17
1.8	The LoRaWAN MAC frame format	18
1.9	Challenge of trading-off throughput and energy efficiency	21
2.1	Coarse taxonomy of MAC protocols for wireless networks	24
2.2	Throughput models for several Random Access Protocols ($a=0.01$) .	26
2.3	Research methodology with complementary tools	37
3.1	Vulnerable periods with theoretical Pure and Slotted ALOHA	41
3.2	Theoretical throughput for Pure and Slotted Aloha with $n = 1000$ devices.	43
3.3	From <i>Class B</i> to <i>Class S</i> : adapting the beacon window to slotted uplink communications	45
3.4	Throughput of <i>Class A</i> and <i>Class S</i> LoRa deployments.	48
3.5	Energy efficiency of devices using <i>Class A</i> and <i>Class S</i>	49
3.6	Switching mechanism	51
3.7	Fingerprint of the simulated deployment as a function of the gener- ated traffic	52
3.8	TREMA's switching mechanism flowchart	53
3.9	General-case packet exchange for class switching.	54
3.10	Desynchronization beacon broadcasting.	55
3.11	TREMA's switching mechanism signaling protocol.	55
3.12	TREMA testing with one example scenario ($n_{skip} = 0$).	59
3.13	Switching duration.	60
3.14	Increase in T and E as a function of the generated traffic ($n_{skip} = 3$). .	61
4.1	Device drift measurements	65
4.2	LoRaSync beacon period and slot structure.	68
4.3	Experimental testbed infrastructure.	70
4.4	<code>lock_wait_execute</code> flowchart.	71
4.5	Device drift measurements	72
4.6	Experimental pure and slotted ALOHA throughput	73
5.1	Throughput model and simulation results	80

5.2	Energy efficiency model and simulation results	81
5.3	Energy efficiency as a function of the slot size (2000 devices)	82
5.4	LoRa receiver synchronization and capture scenarios	84
5.5	Pure ALOHA capture coefficient measurement	89
5.6	Throughput model validation with experimental results (10 devices).	90
5.7	Throughput per device	91
5.8	Jain's fairness index	92

List of Tables

1.1	SNR lower bound as a function of the SF	10
1.2	Comparison of the leading LPWAN technologies	12
1.3	Log-distance path loss model parameters in several places	13
1.4	Features for the different LoRaWAN data rates.	19
2.1	Theoretical network capacity for different random access strategies (from [Kleinrock 1975c])	26
2.2	Access scheme strategies in the literature	35
3.1	Duration of each portion of the <i>Class B</i> and <i>S</i> beacon period	44
3.2	Simulation parameters	46
3.3	Simulation parameters	57
4.1	Slotframe intervals	68
5.1	Simulation parameters	79
5.3	Measured success coefficients	87

Glossary

CAD	Channel Activity Detection
CCA	Clear Channel Assessment
CR	Coding Rate
CRC	Cyclic Redundancy Check
CSMA	Carrier Sense Multiple Access
CSS	Chirp Spread Spectrum
DC	Duty Cycle
DCF	Distributed Coordinated Function
DIFS	DCF Inter-Frame Space
DR	Data Rate
ETSI	European Telecommunications Standards Institute
FHDR	Frame Header
ISM	Industrial, Scientific, Medical
LoRa	Long Range
LoRaWAN	Long Range Wide Area Network
LPWAN	Low Power Wide Area Network
MAC	Medium Access Control
MHDR	MAC Header
MIC	Message Integrity Code
RFU	Reserved for Future Usage
RSSI	Received Signal Strength Indication
RTS/CTS	Ready To Send/Clear To Send
RX	Reception
SF	Spreading Factor
SNR	Signal-to-Noise Ratio
ToA	Time on Air
TX	Transmission

Introduction

– *Version française plus bas* –

In the last decade we have witnessed relevant technological progresses in radio communications and electronic miniaturization techniques. Together they unveiled the possibility to monitor physical phenomena at large scale [Akyildiz 2002]. Low power wireless devices have been then leveraged to deploy widely spread radio networks. However, battery replacement can be very expensive, and in most cases unfeasible in networks with potentially thousands of devices, so maximizing their lifetime becomes the only viable solution. Herein, the Low Power Wide Area Network (LPWAN) architecture has been largely adopted for long range, low throughput and energy efficient data collection with relatively simple and cheap devices [Mekki 2019]. Several LPWAN technologies have emerged, including Long Range (LoRa) networks. They are promoted by the LoRa Alliance, which is a consortium gathering more than 500 companies to drive the open development of the LoRa Wide Area Network (LoRaWAN) specification [LoRa Alliance 2020a]. LoRa has prevailed among all LPWAN technologies, as it offers an excellent tradeoff between transmission range, power consumption and payload size [Rizzi 2017b]. LoRa deployments additionally display many interesting features such as bidirectional communications and end-to-end encryption, making them suitable for a large variety of applications. Furthermore, LoRaWANs are built upon unlicensed Industrial, Scientific, Medical (ISM) radio bands enabling easy prototyping and deployment.

The LoRaWAN Medium Access Control (MAC) scheme was designed as pure ALOHA for uplink communications, meaning that upon frame generation, devices immediately start transmitting without checking whether the radio channel is free. Such a random access scheme does not require synchronization, thus saving power. However, this scheme suffers from high collision rates that naturally limit the channel throughput to 18% at most [Abramson 1970]. A low throughput is thus traded-off for a low power consumption. The increasing availability of cheap wireless devices and the consequent growth of communication traffic have raised strong concerns regarding the scaling capabilities of LoRa deployments [Mikhaylov 2016, Bor 2016, Bankov 2016, Georgiou 2017, Adelantado 2017, Van den Abeele 2017, Yousuf 2018]. As a result, a large panel of contributions have been proposed in the recent years to improve the performances of the LoRa MAC layer. However conciliating network scalability with an energy-efficient management of the device batteries is difficult task. We show that no solution in the current literature is able to do so while also remaining robust to topology changes.

In this manuscript, we therefore target the design of MAC protocols for LoRa networks able to trade-off throughput and energy-efficiency. After an extensive literature review, we argue that the most promising solution to extend the capacity of LoRa networks with the current technological limitations is to leverage synchroniza-

tion. Moreover, we claim that a beacon-based synchronization is the only scalable option given the Duty-Cycle regulations operating over the ISM bands used by LoRa, and the half-duplex nature of the LoRa gateways. As a result we introduce *Class S*, a new device class allowing the use of slotted access schemes in LoRa networks. Through simulations, we show that a beacon-skipping mechanism is crucial to ensure energy-efficiency with such a strategy.

A major finding of this first contribution is that the slotted access allows to increase the maximum achievable throughput in LoRa deployments, and may also be more energy efficient than pure ALOHA in congested networks. However the legacy access performs better under a certain traffic load threshold. In that, we see that there is no one-fits-all solution, and that the most energy-efficient access scheme depends on the situation. As a result, we introduce a switching mechanism able to change the MAC protocol depending on the probed traffic conditions. In details, the Traffic-aware Energy-efficient MAC (TREMA) mechanism can dynamically synchronize or desynchronize all devices in the network to select the preferable strategy. Simulation results reveal that such a scheme is able to follow variations of the input load and improve the network performances in terms of throughput and energy efficiency.

After that, we perform a proof-of-concept implementation of *Class S* on a real-hardware testbed in order to prove the feasibility of our solution. This task pushed us to consider hardware-specific constraints in more details, such as the clock drifting of the low-power LoRa devices. Herein, guard times in each slot have been linked to the worst-case skew of the testbed nodes, and to the aforementioned beacon-skipping mechanism. This robust, yet energy-efficient synchronization scheme has been named *LoRaSync*. Experiments results ultimately show that *LoRaSync* is able to successfully correct timing errors and ensure a reliable synchronization. However, the testbed implementation also revealed that the capture effect occurs in LoRa deployments, thus increasing the observed throughput with respect to the expected one computed through classical ALOHA models.

In a final contribution, we therefore include all features observed on the testbed into novel performance evaluation models that faithfully picture the behavior of real networks. First, we mathematically approach the throughput and energy-efficiency of *LoRaSync*-operated networks with varying margin sizes. This modeling shows that the slot size should be adjusted depending on the traffic load to maximize energy efficiency. In detail, the *LoRaSync* transmission margins should be shortened as the offered load increases, and the synchronization periodicity should thus be adapted accordingly. This finding can be used to enhance the TREMA adaptive mechanism even further. Second, this last contribution tackles the elaboration of experimental models accounting for the capture effect in order to accurately picture the testbed throughput. This allows us to check that the observed network behavior fits our expectations, thus validating the implementation of our synchronous and asynchronous schemes. Further than improving the overall throughput, experiments reveal that the capture effect negatively affects the fairness of the network.

All in all, this thesis highlights the need for adaptive mechanisms to support

the development of scalable, yet energy-efficient access schemes for LoRa. Several tools are leveraged to support this claim. In the first steps, theoretical modeling is used to forecast the network behavior with synchronous and asynchronous access strategies. Simulations are then leveraged to adapt these access schemes to the features of the LoRa technology and evaluate the resulting protocols. After that, an experimental testbed is used to assert the viability of the studied approach, and highlights real-world phenomena such as clock drifting or the capture effect. As a means to thoroughly evaluate the impact of these effects on any deployment, these observations are finally fed back into theoretical models that fit the real network behavior. Our modeling is capable of accurately describing our specific setup, however it does not yet generalize to other deployment layouts. This paves the way for future research so that results extracted from the mathematical approach, simulations and experiments finally converge, independently from the studied network topology. Deriving such versatile models would notably benefit in driving the decision-making mechanism in traffic-aware protocols.

The rest of this manuscript is organized as follows. Chapter 1 provides an overview of the LoRa modulation and LoRaWAN MAC protocol in its current state-of-the-art. Based on this description, we justify the reasons why trading-off throughput and energy efficiency is a complex task in such a system. Then, Chapter 2 presents a taxonomy of the MAC strategies most commonly used in wireless networks, and discusses their application to LoRa deployments. Following these premises, we then provide an extensive literature review of alternative MAC layer protocols for LoRa, and show that none is able to improve the network capacity without weighing too much on the device battery consumption or on the system flexibility. Based on this extensive state-of-the-art, we thereby justify the approach developed in this manuscript. In Chapter 3, we present the theoretical behavior of the pure ALOHA protocol used in LoRaWAN, and of its synchronous equivalent, slotted ALOHA, forecasting the benefits of using such a slotted access. We then introduce the novel *Class S* that enables defining slotted access schemes in large-scale LoRa networks, and compare the performances of pure and slotted ALOHA deployments through simulations. To conclude this chapter, we present the TREMA switching mechanism able to dynamically synchronize or desynchronize the network depending on the probed network load, and demonstrate its benefits with the same simulator. In Chapter 4 we bridge the gap between simulations and experiments, and implement *Class S* on a research testbed. The LoRaSync synchronization scheme is thereby introduced to efficiently cope for clock drifting issues. Experimental results assess the viability of our solution, and highlight the impact of the capture effect on real networks. Chapter 5 closes the loop and feeds the experimental observations (*i.e.*, clock drift and capture effect) back into our performance models to accurately picture the network behavior. Finally, concluding remarks and research perspectives are disclosed in Chapter 6.

De significatives avancées en communications radio et en techniques de miniaturisation électronique ont été réalisées au cours de ces dix dernières années. Cela a ouvert la possibilité de surveiller l'évolution de phénomènes physiques à large échelle [Akyildiz 2002]. Des terminaux sans fil à basse puissance sont ainsi utilisés pour déployer des réseaux radio très étendus. Cependant, le remplacement des batteries d'un grand nombre de terminaux disséminés sur une large zone géographique peut s'avérer très coûteux, voire impossible. Ainsi, maximiser leur durée de vie reste la seule solution viable. L'architecture LPWAN (pour Low-Power Wide-Area Network) a été largement adoptée pour la collection de données sur des zones étendues avec de faibles exigences de débit [Mekki 2019]. Parmi les technologies LPWAN disponibles, les réseaux LoRa (Long Range) se sont imposés sur le marché. Cette technologie est soutenue par la LoRa Alliance, un consortium industriel promouvant le développement ouvert de la spécification LoRaWAN (Long Range Wide Area Networks) [LoRa Alliance 2020a] décrivant la couche d'accès au medium utilisée dans de tels réseaux. La modulation LoRa s'est également imposée car elle propose un excellent compromis entre portée de transmission, consommation d'énergie et taille de la charge utile [Rizzi 2017b]. De plus, les déploiements LoRaWAN présentent de nombreuses fonctionnalités intéressantes telles que des communications bidirectionnelles ou un cryptage de bout-en-bout, ce qui les rend adaptés à une large variété d'applications. Enfin, ils utilisent les bandes de fréquence sans licence ISM (Industrial, Scientific, Medical) ce qui facilite le prototypage et le déploiement de réseaux LoRa par de nombreux acteurs.

Le mécanisme de contrôle d'accès au medium des LoRaWAN repose sur une stratégie de type pure ALOHA. Cela signifie que lorsqu'un terminal génère un message, il le transmet instantanément sans se synchroniser ou sonder le medium de communication. Un tel mécanisme économise ainsi de l'énergie de par la simplicité du comportement des dispositifs sans fil. Cependant, ce mode d'accès souffre d'un fort taux de collisions qui limite naturellement la capacité réseau à un maximum de 18% [Abramson 1970]. Un bas débit est ainsi le prix à payer pour une basse consommation d'énergie. La disponibilité grandissante de terminaux sans fil à faible coût et la conséquente hausse de la charge de trafic a soulevé de fortes inquiétudes concernant les capacités de passage à l'échelle des déploiements LoRa [Mikhaylov 2016, Bor 2016, Bankov 2016, Georgiou 2017, Adelantado 2017, Van den Abele 2017, Yousuf 2018]. En conséquence, de nombreux travaux de recherche ont depuis tenté d'améliorer les performances de la couche de contrôle d'accès au medium de ces réseaux. Cependant, concilier une grande capacité réseau avec une gestion efficiente des batteries des terminaux est un tâche difficile. Nous montrons qu'aucun protocole de l'état de l'art actuel n'en est capable tout en tolérant des changements de topologie.

Dans ce manuscrit, nous abordons donc la conception de protocoles d'accès pour les réseaux LoRa capables de concilier des exigences de débit et d'efficacité énergétique. Après une revue de littérature extensive, nous argumentons que la solution la plus prometteuse pour étendre la capacité réseau avec les limitations technologiques actuelles est la mise en place d'un accès synchrone. De plus, une synchronisation

basée sur la diffusion de balises par les stations de base est la seule solution scalable en raison des limitations de rapport cyclique imposées sur les bandes ISM. En conséquence, nous introduisons la Classe S, une nouvelle catégorie de dispositifs LoRa qui permet l'utilisation de modes d'accès avec slots. Par le biais de simulations, nous montrons également qu'un mécanisme pour réduire la périodicité de réception des balises de synchronisation est nécessaire pour garantir l'efficacité énergétique d'une telle approche.

Une découverte majeure de cette première contribution est que l'accès synchrone permet d'étendre le débit maximum tout en étant également plus énergétiquement efficient que le protocole de base dans les réseaux congestionnés. Cependant pure ALOHA reste préférable pour de faibles charges de trafic. Ainsi, nous montrons qu'aucune solution ne convient à toutes les situations. En conséquence, nous introduisons un mécanisme conscient du trafic capable d'adapter la stratégie d'accès à la charge réseau sondée. Spécifiquement, notre protocole est capable de dynamiquement synchroniser et désynchroniser les terminaux pour toujours sélectionner l'approche la plus efficiente. Des simulations sont réalisées pour démontrer que notre solution peut suivre les variations de trafic et offre un meilleur débit maximal et une meilleure efficacité énergétique que le protocole de base.

Ensuite, nous implémentons la Classe S sur un banc de tests réel afin de démontrer la faisabilité de notre approche. Cette tâche nous a amenés à considérer des contraintes matérielles telles que le décalage d'horloge des terminaux LoRa. Nous avons ainsi ajouté des marges de transmission basées sur l'erreur de synchronisation correspondant au pire cas de décalage. La taille de ces marges sert aussi à adapter la périodicité de réception des balises de synchronisation afin de minimiser la consommation d'énergie. Nous avons nommé ce mécanisme de synchronisation robuste et efficient LoRaSync. Les résultats expérimentaux montrent que LoRaSync est capable de corriger les erreurs de synchronisation correctement. Cependant, les mesures de performances révèlent aussi que l'effet capture intervient fréquemment au sein du banc de tests, augmentant significativement le débit par rapport aux estimations obtenues via les modèles ALOHA classiques.

Dans une dernière contribution, nous implémentons toutes les spécificités observées sur le banc de tests dans de nouveaux modèles d'évaluation de performances. D'abord, nous modélisons le débit et l'efficacité énergétique de réseaux utilisant le mécanisme LoRaSync avec différentes tailles de slots. Nous montrons ainsi que la longueur des marges de transmission devrait être ajustée en fonction de la charge réseau pour optimiser la consommation énergétique. Spécifiquement, les slots devraient être raccourcis lorsque la charge de trafic augmente, et la périodicité de synchronisation devrait donc être adaptée en conséquence. Cette découverte pourra être utilisée pour améliorer notre mécanisme adaptatif dans de futures contributions. Ensuite, nous élaborons des modèles expérimentaux pour considérer l'effet capture dans notre estimation du débit. Nous montrons que cette approche converge avec les performances réelles du banc de tests, validant donc notre implémentation. De plus, les traces réseau révèlent que l'effet capture impacte l'équité du réseau.

En un mot, cette thèse met en évidence le besoin de mécanismes adaptatifs pour

supporter le développement de mécanismes d'accès scalables et énergétiquement efficaces pour les réseaux LoRa. De nombreux outils sont utilisés pour vérifier ce postulat. Lors des premières étapes, des modélisations mathématiques sont utilisées pour anticiper le comportement du réseau avec des stratégies d'accès synchrones et asynchrones. Ensuite, des simulations permettent d'adapter ces stratégies aux spécificités de la technologie LoRa et d'évaluer les protocoles résultants. Après cela, un banc de tests est utilisé pour vérifier la faisabilité de notre approche, et mettre en évidence des phénomènes physiques tels que le décalage d'horloge ou l'effet capture. Pour évaluer l'impact de tels phénomènes, ces observations sont finalement intégrées dans de nouveaux modèles théoriques qui décrivent fidèlement les performances réelles. Cette modélisation ne peut cependant pas encore être généralisée à d'autres topologies. Cela ouvre la voie vers de futures contributions pour faire converger les résultats issus des modèles, des simulations et des expériences indépendamment de la forme du déploiement. De tels modèles pourraient notamment guider le processus de décision du protocole adaptatif présenté dans ce manuscrit.

Technological background and thesis problematic

Dans ce chapitre, nous présentons d'abord l'environnement technologique des réseaux LoRa. Spécifiquement, nous abordons la modulation LoRa, et spécifions les particularités qui font que cette couche physique propose un excellent compromis entre portée de transmission, consommation d'énergie et débit. Ensuite nous détaillons la spécification LoRaWAN, qui définit les caractéristiques de la couche d'accès au médium de ces réseaux. Enfin, nous mettons en évidence les raisons pour lesquelles il est dans ce contexte difficile de trouver un compromis entre capacité de passage à l'échelle et efficacité énergétique. D'une part, la multiplication du nombre d'objets communicants répartis sur de larges zones géographiques met en péril le mode accès asynchrone décrit dans la spécification. D'autre part, l'élaboration de mécanismes d'accès plus sophistiqués risque de peser lourdement sur l'efficacité énergétique des terminaux qui ne transmettent quotidiennement que de très faibles quantités de données. Ce constat justifie la problématique de recherche de ce manuscrit, qui vise à développer des protocoles capables de concilier débit et efficacité énergétique dans les réseaux LoRa.

Contents

1.1	The LoRa physical layer and its features	9
1.1.1	The Chirp Spread Spectrum modulation	9
1.1.2	The physical-layer packet structure and Time on Air	10
1.1.3	The low-power and long-range capabilities of LoRa compared against other LPWAN technologies	11
1.1.4	Experimental insights on the LoRa transmission range	13
1.2	The LoRaWAN MAC-layer specification	15
1.2.1	The LoRaWAN classes	15
1.2.2	Regulations	17
1.2.3	The MAC-layer frame format	18
1.2.4	The LoRaWAN data rates and Adaptive Data Rate mechanism	19
1.3	The challenge of trading-off throughput and energy efficiency when designing a MAC layer scheme for LoRaWAN	20




Figure 1.1: Topology of a LoRaWAN deployment

Low-Power Wide-Area Networks (LPWAN) are designed to provide Internet access to widely-spread, battery-powered wireless devices [Centenaro 2016]. Such networks are tailored for large-scale sensing applications, with any device transmitting very small amounts of data per day. Among all the available LPWAN solutions, Long Range Wide Area Networks (LoRaWAN) have gathered a major interest from the research and industrial communities, featuring robust bidirectional low-power communications with very long transmission ranges over unlicensed Industrial, Scientific and Medical (ISM) bands. In the typical LoRaWAN architecture, end-devices communicate with a set of gateways through LoRa communications, and gateways are in charge of forwarding the messages to a server on the Internet using IP. This topology is represented in Figure 1.1.

The term LoRa specifically designates the physical layer modulation, which is proprietary and owned by Semtech Corporation. On the other hand, LoRaWAN refers to the MAC protocol, which benefits from an open documentation. Indeed, the LoRa Alliance [®] is a non-profit association of industrial actors that aims at the global adoption of the LoRaWAN standard. This consortium therefore promotes the transparent development of the MAC layer specification [LoRa Alliance 2020a] and provides a common view on its implementation.

In this context, this chapter presents the scientific context of the thesis and introduces its research question. It is organized as follows. The LoRaWAN technology, central case study of our work, is introduced in details faithfully to its state-of-the-art specification. In details, the LoRa physical-layer modulation is presented in Section 1.1 and the LoRaWAN MAC layer standard is detailed in Section 1.2. On this basis, Section 1.3 presents the scalability issues experienced in current LoRa deployments, as well as the difficulty to solve this problem while guaranteeing an energy-efficient management of the device batteries.

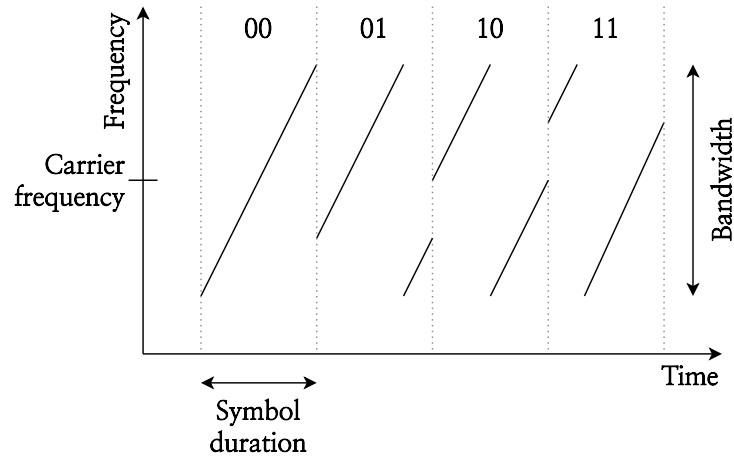


Figure 1.2: Chirp Spread Spectrum modulation example with 2 bits per symbol

1.1 The LoRa physical layer and its features

1.1.1 The Chirp Spread Spectrum modulation

The LoRa physical layer relies on a Chirp Spread Spectrum (CSS) modulation [Reynders 2016]. With this technique, each symbol is represented by a linear frequency sweep (*i.e.*, *Chirp*) over the available bandwidth. The different symbol values are coded with frequency offsets at the start of the Chirp. A very simple example of such a modulation is represented in Figure 1.2, in the particular case where each symbol represents 2 bits. Herein, $2^2 = 4$ offset values are possible, which results in the 4 represented symbols.

Chirps can sweep from the minimum to the maximum frequency (up-chirps), or the opposite (down-chirps). This allows a quadrature inversion, which is used to differentiate the uplink (device to gateway) and downlink (gateways to device) traffic in LoRaWAN.

The following information is gathered from the LoRa Modulation Basics Application Note [Semtech Corporation 2015] provided by Semtech. The number of bits per symbol is called the Spreading Factor (SF) of the transmission. As a result, each symbol may take 2^{SF} values. The symbol duration T_{sym} , bandwidth BW and Spreading Factor SF are linked by the following equation:

$$T_{\text{sym}} = \frac{2^{SF}}{\text{BW}} \quad (1.1)$$

With a fixed bandwidth, a higher SF therefore induces longer symbols. In LoRa, the admitted SF values span from 7 to 12, and the bandwidth is usually equal to 125 kHz. The LoRa modulation includes a Cyclic Redundancy Check (CRC), enabling corruption detection and error correction thanks to redundancy bits. The coding rate CR therefore indicates the proportion of useful bits in the payload, and is expressed as such:



Figure 1.3: LoRa physical frame structure

$$\text{CR} = \frac{4}{4 + k_{\text{CR}}}, \text{ with } k_{\text{CR}} \in \llbracket 1; 4 \rrbracket \quad (1.2)$$

The physical bit rate B of the modulation can easily be derived from these parameters:

$$B = \text{SF} \cdot \frac{\text{BW}}{2^{\text{SF}}} \cdot \text{CR}, \quad (1.3)$$

All in all, **the bit rate decreases as the SF increases**. Indeed, a greater SF induces a smaller number of bits per symbol as well as a longer chirp duration.

A LoRa frame may be demodulated by a receiver if its Received Signal Strength Indication (RSSI) is greater than the receiver sensitivity [Semtech Corporation 2020]. In section 1.1.4 we will see how the RSSI can be estimated as a function of the distance with experimental path loss models. From [Semtech Corporation 2015], we know that the receiver sensitivity S depends on the receiver noise floor NF (a 6 dB value is commonly admitted), and the SNR lower bound allowing a correct demodulation. It is expressed as such:

$$S = -174 + 10 \cdot \log_{10}(\text{BW}) + \text{NF} + \text{SNR} \quad (1.4)$$

Interestingly, SNR depends on the Spreading Factor used for the transmission. Exact values are provided in Table 1.1. In detail, we see that using a greater SF results in a lower SNR limit, which in turn improves the receiver sensitivity as shown by equation 1.4. That is why **the transmission range increases as the SF increases**.

Table 1.1: SNR lower bound as a function of the SF

SF	7	8	9	10	11	12
SNR (dBm)	-7.5	-10	-12.5	-15	-17.5	-20

1.1.2 The physical-layer packet structure and Time on Air

Until now, we have established that the several LoRa SFs represent different trade-offs between bit rate and transmission range. We will now introduce the physical frame structure [Semtech Corporation 2013] (*c.f.* Figure 1.3) and show how to derive the frame Time on Air (ToA).

The first field is the preamble which announces the beginning of the transmission with a series of up-chirps with no offset: 8 for most packets and 10 for beacons (*c.f.* Section 1.2), to help the devices perform a low power beacon search. This number of up-chirps is noted n_{pre} . Then comes the synchronization word, which allows the receiver to precisely align its demodulator with the arriving symbols. This word is composed of 2 up-chirps and 2.25 down chirps. After that, the physical header PHDR indicates the data size. The PHDR_CRC indicates whether a Cyclic Redundancy Check (CRC) will be used or not. PHYPayload contains the physical payload, and CRC carries the optional CRC value.

We note PL the number of payload bytes. LoRa admits an implicit mode, which reduces the packet overhead by omitting the header. H indicates whether the header is enabled ($H = 0$) or disabled ($H = 1$). Finally, DE is used to express whether a parameter called "low data rate optimization" is used (1 if it is enabled, else 0). This parameter allows to increase the robustness of the transmission for SFs 11 and 12. Indeed these data rates are associated with very long ToAs, so such an optimization is helpful to avoid clock drift or oscillator imprecision issues (due to motion or temperature changes for instance). All in all, the packet ToA is expressed as such:

$$\text{ToA} = \left[n_{\text{pre}} + 4.25 + 8 + \max \left(\left\lceil \frac{8\text{PL} - 4\text{SF} + 28 + 16 - 20\text{H}}{4 \cdot (\text{SF} - 2\text{DE})} \right\rceil \cdot (k_{\text{CR}} + 4), 0 \right) \right] \cdot T_{\text{sym}} \quad (1.5)$$

1.1.3 The low-power and long-range capabilities of LoRa compared against other LPWAN technologies

The LoRa modulation is particularly robust to in-band and out-of-band interferences. It is resilient to multipath fading and Doppler effect issues [Semtech Corporation 2015]. All this combined with the excellent receiver sensitivity of LoRa (from -124.5 for SF7 to -137 dBm for SF12, *c.f.* Equation 1.4) results in a long range for a relatively low transmission power. As a result, LoRa has been widely adopted for wireless sensing purposes.

It should also be noted that the different SFs are considered orthogonal, meaning that using several SFs in parallel in the same network will allow to seamlessly increase the deployment capacity. This orthogonality has been questioned by the research community [Croce 2018], but remains a reasonable assumption in most cases.

In order to position LoRa in regards to other LPWAN technologies, we compare its characteristics against those two of its main competitors (*i.e.* Sigfox and NB-IoT) in Table 1.2. Supply current and nominal voltage information have been taken from the datasheets of a typical transceiver for each selected technology, namely the uBlox SARA N210 [uBlox 2019] for NB-IoT, the Atmel ATA 8520E [Atmel 2016] for Sigfox and the SX1272 [Corporation 2019b] for LoRa. These specific transceivers have been selected because they are widely spread on the market, and feature good

performances that have thoroughly been evaluated. Provided data relates to the physical layer perspective, meaning that MAC layer considerations (*e.g.* collisions, retransmissions...) are not accounted for in this comparison. Indicative ranges for rural and urban environments have been provided by the comparative study of Mekki *et al.* [Mekki 2019]. However this parameter varies a lot depending on the exact nature of the environment. More details on about the range estimation of LoRa transmissions are provided in Section 1.1.4.

Table 1.2: Comparison of the leading LPWAN technologies

Technology	NB-IoT	LoRa	Sigfox
Example transceiver	uBlox SARA N210	Semtech SX1272	Atmel ATA 8520E
Supply current			
TX	220 mA	125 mA	10.4 mA
RX	46 mA	11.2 mA	32.7 mA
IDLE	6 mA	1.4 mA	0.05 mA
SLEEP	3 μ A	0.1 μ A	0.15 μ A
Nominal voltage	3.6 V	3.3 V	3 V
Power consumption			
TX	792 mW	412.5 mW	31.2 mW
RX	165.6 mW	36.96 mW	98.1 mW
IDLE	21.6 mW	4.62 mW	0.15 mW
SLEEP	10.8 μ W	0.33 μ W	0.15 μ W
Max. payload size	1600 bytes	250 bytes	12 bytes
Frequency bands	LTE	ISM	ISM
Indicative range [Mekki 2019]			
Urban environment	1 km	5 km	10 km
Rural environment	10 km	20 km	40 km

This comparison shows that LoRa proposes a nice compromise between range, power consumption and payload size. Indeed, Sigfox specializes in very long range capabilities but the 12 bytes limitation on the payload is a limiting factor for many use-cases. Indeed, like LoRa, Sigfox operates over licence-free ISM (Industrial, Scientific, Medical) radio bands which are subject to Duty-Cycle limitations (*c.f.* Section 1.2.2). This means that the number of transmissions a node is able to perform everyday is limited, which makes the payload size limitation an even greater problem. Conversely, NB-IoT allows much larger payloads but its range is much more limited. This technology benefits from the infrastructure and frequency bands of cellular Long-Term Evolution (LTE) networks. It therefore naturally fits higher throughput applications in urban environments where an LTE infrastructure is already in place. All in all, LoRa offers a nice tradeoff between all these considerations, and can be used for many applications in urban and rural environments. Furthermore its open specification and operation over license-free ISM bands make it an easy and affordable solution, which explains why it has been adopted by many

research and industrial actors.

Information provided in Table 1.2 can be leveraged to easily estimate the battery lifetime of an IoT device. Let us consider a LoRa node powered with a 250 mAh battery that transmits 10 messages per day. Each message is transmitted with SF7BW125, and carries the maximum payload length allowed, 250 bytes. If we consider $CR = \frac{4}{8}$ and explicit header enabled, this results in a packet ToA of 626.94 ms (*c.f.* Equation 1.5). The SX1272 LoRa transceiver has a supply current of 125 mA in a transmission state and $0.1 \mu\text{A}$ in a sleeping state. The node therefore has an average current draw over the day of $\frac{(10 \cdot 0.62694) \cdot 125 + (86400 - 10 \cdot 0.62694) \cdot 0.0001}{86400} \approx 0.00917$ mA. Considering a typical IoT device battery with a capacity of 250 mAh, we can expect a lifetime of $\frac{250}{0.00917} \approx 27263$ hours, which represents over 3 years.

1.1.4 Experimental insights on the LoRa transmission range

Table 1.3: Log-distance path loss model parameters in several places

Ref.	Location	SF	n	B	d_0 (m)
[Petajajarvi 2015]	Oulu, Finland	12	2.32	128.95	1000
[Jörke 2017]	Dortmund, Germany	12	2.65	132.25	1000
[Seye 2018]	Dakar, Senegal	12	2.8	133.6	1000
[Callebaut 2020]	Leuven, Belgium				
	Urban area	7, 9, 12	2.75	74.85	1
	Coastal area	7, 9, 12	3.62	43.96	1
	Forest area	7, 9, 12	2.03	95.52	1
[Rademacher 2021]	Bonn, Germany	12	1.58	132.41	1000
[Bianco 2021]	Bletterbach canyon, Italy	7	5.51	0.9	1

LoRa devices being almost always placed near the ground, the Friis free-space propagation equation can not be used to accurately estimate the transmission range. Instead, the approach most commonly used in the literature to evaluate the range of LoRa transmissions is the log-distance path loss model. With B the path loss intercept, n the path loss exponent and d_0 a reference distance (typically 1 km), the path loss PL (in dBm) for a distance d is expressed as such:

$$PL(d) = B + 10 \cdot n \cdot \log_{10} \left(\frac{d}{d_0} \right) \quad (1.6)$$

Many contributions [Petajajarvi 2015, Jörke 2017, Seye 2018, Callebaut 2020, Rademacher 2021, Bianco 2021] use this model by collecting a large number of RSSI samples with moving devices, and then approximate the model parameters b and n . The model can then be used to estimate the path loss in this specific environment. Table 1.3 lists the model coefficients found in the literature, and measured in different environments by different research teams. When using these models we must remember that they were derived from experimental measurements. Therefore, we can only expect that they generalize well for the distance range in which measurements have been performed. Moreover, we must be aware that they

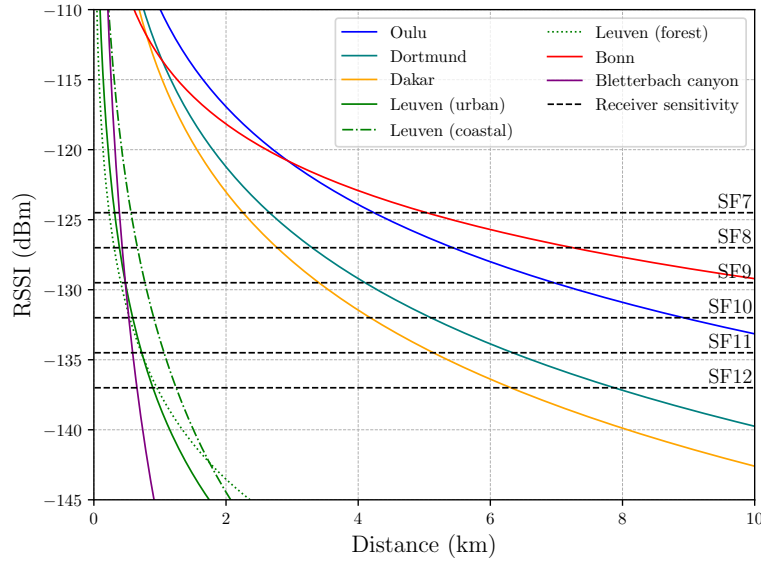



Figure 1.4: RSSI and range

may be very imprecise if the number of samples is low. Most of the time SF12 was used by researchers to derive the studied experimental models, but the impact of the Spreading Factor on the propagation models should be investigated as well.

By simply subtracting these path loss models to the adequate transmission power and antenna gain values, the RSSI of a signal at any given distance of the transmitter can be estimated. By noting P_{TX} the transmission power and G the antenna gain, the RSSI (in dBm) at a distance d is thus expressed as:

$$\text{RSSI}(d) = P_{TX} + G - \text{PL}(d) \quad (1.7)$$

With this model, the estimated maximum transmission range is the distance d for which the RSSI is equal to the receiver sensitivity for the considered SF. RSSI for all models referenced in Table 1.3 have been plotted in Figure 1.4 as a function of the distance, assuming a 14 dBm TX power and a 5 dBm antenna gain. The receiver sensitivities for all SFs are represented as well. The crossing points between the different models and sensitivity thresholds therefore provide the associated transmission ranges. Naturally, a higher SF results in a lower sensitivity, which lengthens the range. Herein, Figure 1.3 allows to compare the transmission ranges estimated through the different experimental path loss models derived in the literature. This thus provides perspective on the indicative transmission ranges provided in Table 1.2. This comparison indeed reveals that the transmission range greatly depends on the environment, and that specific studies are required to gather precise estimations of the transceiver capabilities in different geographical areas.

Figure 1.5: The LoRaWAN *Class A*

1.2 The LoRaWAN MAC-layer specification

In this section we present the LoRaWAN specification, which provides a common view on the MAC layer implementation of LoRaWAN deployments. This specification is in fact composed of three documents. The main specification [LoRa Alliance 2020a] describes the access schemes associated with 3 device classes. It also details the MAC frame formats, MAC commands and device activation procedures. The backend interfaces document [LoRa Alliance 2020b] describes the traffic flows on the server side in the case where the different server components are decentralized. Finally, the regional parameters document [LoRa Alliance 2021] provides the channel settings and access regulations for ISM license-free bands in the different regions of the world.

1.2.1 The LoRaWAN classes

The LoRaWAN specification [LoRa Alliance 2020a] defines 3 device classes: *Class A*, *Class B* and *Class C*. These modes are characterized by different MAC layer access schemes, and are a central cornerstone on which the contributions of this manuscript have been built. We will therefore now present these classes in great details.

Class A is the default operational mode that all LoRaWAN devices should implement. The radio medium is accessed by the device in a pure ALOHA fashion [Abramson 1970]. This means that when a device generates a frame it transmits it immediately, without getting synchronized beforehand or checking whether the radio channel is free. Then, two receiving (RX) slots, RX1 and RX2, are opened at fixed delays after the uplink transmission (1 and 2 seconds by default for Europe). RX2 is only opened if nothing was received during RX1. In the case of a confirmed message, if no downlink acknowledgment was received during RX1 or RX2, the message will be re-transmitted after a random delay. Receiving windows are sufficiently long to let the device detect a LoRa preamble (30 ms). If a preamble is detected, then the window will be kept open to demodulate the whole message. A downlink generated by the server may only be transmitted synchronously with the RX1 (or RX2) opened by the recipient end-device after this latter has sent an uplink.

Class B devices still implement *Class A* for uplink communications but must periodically open additional RX slots, namely *ping slots*, to enable server initiated messages. This strategy thus allows to bound the downlink delay. Ping slots are 30 ms long, as *RX1* and *RX2*, to correctly capture a LoRa preamble. In order to open such slots, devices need to get synchronized with the network. They do so by listening to periodic beacons that are broadcast by the gateways. Figure 1.6 represents the *Class B* slot frame. Each beacon is sent at the beginning of a 128 seconds long *BEACON_PERIOD* composed of:

1. The 2.120 s *BEACON_RESERVED* interval. The gateway broadcasts a beacon at the beginning of this interval.
2. The 122.880 s *BEACON_WINDOW* interval. It contains exactly 4096 ping slots, each 30 ms-long.
3. The 3 s *BEACON_GUARD* interval. It allows the transmission of a frame started during the last ping slot of the *BEACON_WINDOW*, and is sufficiently long to let the gateway to prepare the next beacon transmission.

A *Class B* device has to turn its radio on during *pingNb* ping slots in every beacon window, with:

$$\text{pingNb} = 2^k, k \in \llbracket 1; 7 \rrbracket \quad (1.8)$$

This value can be parameterized to fit several delay requirements. The ping slot period *pingPeriod*, expressed in number of slots, is therefore:

$$\text{pingPeriod} = \frac{4096}{\text{pingNb}} \quad (1.9)$$

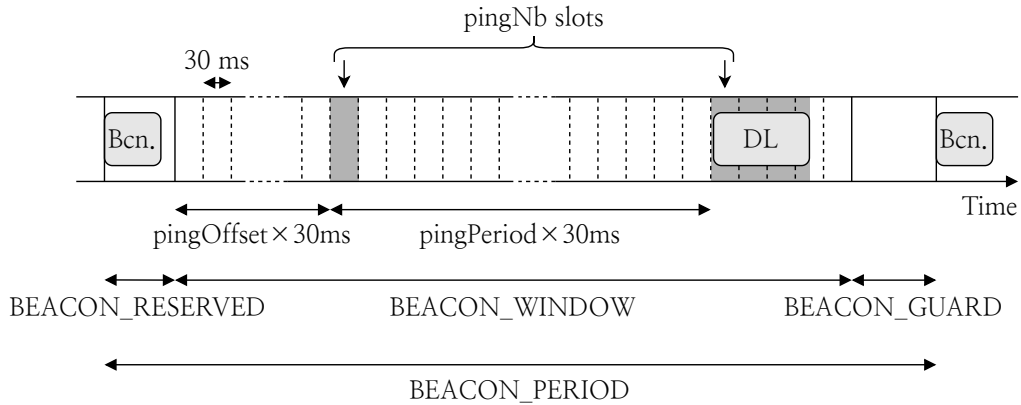
The first ping slot occurs with an offset *pingOffset*, also expressed in number of slots, that follows:

$$0 \leq \text{pingOffset} < \text{pingPeriod} \quad (1.10)$$

Herein, a given device opens its first ping slot after *pingOffset* slots (randomly chosen), and then periodically after *pingPeriod* slots until the end of the beacon window. Assuming t_{bcn} the beacon reception timestamp in seconds, the ping slot opening timestamp for the k th ping slot $t_{\text{ping slot}}(k)$ ($k \in \llbracket 0, \text{pingNb} \rrbracket$) can be computed as such:

$$t_{\text{ping slot}}(k) = t_{\text{bcn}} + \text{BEACON_RESERVED} + \text{pingOffset} \cdot 30 \cdot 10^{-3} + k \cdot \text{pingPeriod} \cdot 30 \cdot 10^{-3} \quad (1.11)$$

In Figure 1.6, a device opens 2 ping slots, and a downlink message is received in the second one. The device radio is therefore kept active until the end of the reception, which overlaps over several ping slots. Clearly, the opening of ping slots and beacon receptions imply an additional energy consumption for end-devices. It should therefore be reserved to nodes that require a bounded downlink delay. Besides, it must be noted that using a higher *pingNb* value reduces the worst-case downlink delay, but also further increases the energy consumption because the

Figure 1.6: The LoRaWAN *Class B*Figure 1.7: The LoRaWAN *Class C*

device must open more ping slots. All devices join the network using *Class A* and can switch to *Class B* after negotiation with the network server.

Finally, *Class C* devices are always listening when not transmitting. This allows to reduce processing delays to the minimum, but at the cost of a considerable energy consumption. In fact, *Class C* targets actuators or repeaters.

1.2.2 Regulations

In Europe, the European Telecommunications Standards Institute (ETSI) defines specific regulations for [ETSI 2012] different sub-bands composing the radio spectrum. A notable LoRaWAN feature is the use of the Industrial, Scientific and Medical (ISM) license-free radio spectrum, which facilitates the prototyping and deployment of distributed applications. However, the sub-bands assigned to LoRa are subject to Duty Cycle (DC) limitations. This means that any device transmitting a message for a given duration must then remain silent for a certain time. More specifically, if a channel is subject to a DC of x %, a device transmitting during t seconds is then forbidden to transmit again during $(100 - x) \cdot t$ seconds. The exact DC value for the different ISM sub-bands differs depending on the region of the world. The LoRaWAN specification thus provides channel access rules in their regional parameters document [LoRa Alliance 2021], that ensure legal compliance with regional regulations.

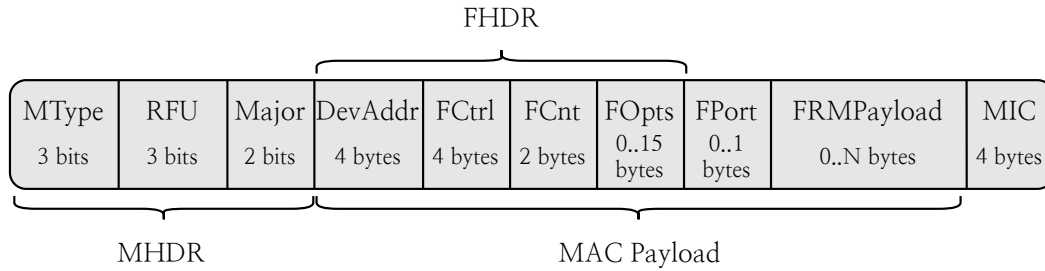


Figure 1.8: The LoRaWAN MAC frame format

1.2.3 The MAC-layer frame format

In this part we present the LoRaWAN MAC frame, and briefly describe its different fields. This information will be useful to know what can be retrieved in each frame, and design access schemes that comply with the legacy protocol. Devices in LoRaWAN need to be activated before exchanging data frames. The activation procedures are disclosed in the specification, but we will not describe them in this manuscript. We present here the frame structure for activated devices, so please note that the following does not apply for Join-Request or Join-Accept types of messages. The LoRaWAN MAC frame format is provided in Figure 1.8. This whole MAC frame is in fact encapsulated in the physical layer frame, and is therefore equivalent to the PHYPayload field in Figure 1.3.

The first three fields form the MAC header (MDHR). The **Mtype** field holds the message type, coded on 3 bits. It indicates whether the frame is associated with data traffic or with the joining procedure. For data frames, it also specifies if the frame is an uplink or a downlink, and if it requires an acknowledgment. **RFU** means "Reserved for Future Usage". Therefore, these 3 bits can be used to support the development of new features. The **Major** field indicates the utilized LoRaWAN version.

After the MDHR comes the MAC payload, which begins with 4 fields forming the Frame Header (FHDR). **DevAddr** holds the end-device address, coded on 4 bytes. Then comes the **FCtrl** frame control byte, that is notably related to the adaptive data rate mechanism (presented in Section 1.2.4) and the retransmission procedure. It also informs about the length of the subsequent **FOpts** field. **FCnt** is used to disclose the frame counter on 2 bytes. **FOpts** provides up to 15 bytes to transport MAC commands. The detail of all MAC commands is provided in the standard. **FPort** indicates the frame port, and **FRMPayload** the MAC payload.

Finally, a Message Integrity Code (**MIC**) is added in the final 4 bytes. It relies on the Advanced Encryption Standard (AES) encryption to produce a hash of the frame content. The MIC value is first computed by the transmitter and added to the frame, then the receiver computes it again over the received frame and checks that the results are identical. This allows to make sure that the frame content is not corrupted during the transmission.

1.2.4 The LoRaWAN data rates and Adaptive Data Rate mechanism

We have seen in Section 1.1 that the different SFs offered different trade-offs between bit rate and transmission range. As a result, the LoRaWAN Regional Parameters document [LoRa Alliance 2021] defines several data rates (DR), that associate several {SF, bandwidth} tuples to specific payload size limitations in order to bound the maximum frame ToA. Table 1.4 discloses the features of these different DRs for the EU863-870 region. Note that an additional DR (DR7) is defined as well, but we have not included it in the table because it uses a Frequency Shift Keying physical-layer modulation instead of LoRa.

DR ¹	SF/BW ² (BW in kHz.)	S ³ (dBm.)	Range ⁴ (km.)	PL_{max} ⁵ (bytes)	ToA _{max} ⁶ (ms.)	B ⁷ (bps.)
0	12/125	-137	7.8	59	4071.42	293
1	11/125	-134.5	6.3	59	2297.86	537
2	10/125	-132	5.1	59	1017.86	977
3	9/125	-129.5	4.1	123	1033.22	1758
4	8/125	-127	3.3	250	1106.43	3125
5	7/125	-124.5	2.3	250	626.94	5470
6	7/250	-124.5	2.3	250	313.47	11000

Table 1.4: Features for the different LoRaWAN data rates.

Notice that the higher data rates are also featured with the shortest ToA, which implies less time spent transmitting for the devices and a lower energy consumption. In fact the only drawback of using the highest DR is the reduced range. As a result, the standard defines an Adaptive Data Rate (ADR) mechanism, which aims at optimizing the DR and transmission power of each ADR-enabled device to minimize its power consumption. It operates in 4 steps:

1. The server collects 20 uplink transmissions from a single node.
2. The server uses the collected transmissions to estimate a worst-case SNR value. It then computes the margin between this worst-case SNR and the

¹LoRaWAN Data Rate number.

²Spreading Factor and Bandwidth (in kHz) combination.

³Receiver sensitivity in dBm (*c.f.* Equation 1.4 and Table 1.1).

⁴Range approximation computed with the log-distance path loss model (*c.f.* Equation 1.6), using the coefficients measured in Dortmund [Jörke 2017], and assuming a transmission power of 14 dBm, a 5 dBm antenna gain, a 6 dBm noise figure, and the known sensitivity for the considered SF and bandwidth (expressed in kilometers).

⁵Maximum MAC payload allowed by the Regional Parameters [LoRa Alliance 2021] for the European EU863-870 MHz band (and in the absence of the FOpts optional field, *c.f.* Figure 1.8).

⁶Associated maximum frame ToA in milliseconds (*c.f.* Equation 1.5), computed with $CR = \frac{4}{8}$ and explicit header enabled. For the payload size, we take length allowed in the absence of the FOpts field provided in the Regional Parameter document [LoRa Alliance 2020a], to which we add the 5 bytes of MHDR and MIC (*c.f.* Figure 1.8).

⁷Physical bit rate B in bits per second (*c.f.* Equation 1.3), computed with $CR = \frac{4}{5}$.

gateway sensitivity for the current DR used by the device.

3. If the margin is sufficiently large, the server commands the node to use a higher DR, thus increasing its energy efficiency.
4. Finally, the server may command the device to reduce its transmission power to reduce the device energy consumption even further.

The LoRa physical layer and LoRaWAN MAC protocol have now been described in details. In the next Section, we describe the research question studied in this manuscript, and justify its relevance in this current state of the art.

1.3 The challenge of trading-off throughput and energy efficiency when designing a MAC layer scheme for LoRaWAN

LoRaWANs face new challenges in terms of MAC protocol design, which can be generalized to all LPWAN technologies. On the one hand, the **low-power** nature of these networks indicates that we are dealing with energy-constrained devices running on batteries. We have seen in Section 1.1.3 that devices mainly consume power when transmitting or receiving frames (*c.f.* Table 1.2). For this reason the LoRaWAN standard defines a pure-ALOHA access scheme for Class A devices. This means that upon frame generation, devices immediately start transmitting without sensing the radio channel or network state beforehand. This strategy saves power because devices only need to be listening during the two RX windows (*c.f.* Section 1.2.1).

On the other hand, LoRaWANs are also characterized by their **wide-area** nature. Indeed these networks are featured with long transmission range capabilities, which leads to very large numbers of devices sharing the same collision area. It is well known that with a pure-ALOHA access, the maximum theoretical throughput that can be achieved when the number of devices tends to infinity is only 18% of the available bandwidth [Abramson 1970] because of frame collisions. In that, the legacy LoRaWAN Class A access dwarfs the network's ability to handle high traffic loads. Considering the ever increasing number of LoRa devices since the creation of the technology, the research community has raised strong concerns regarding the scalability of LoRa deployments with such an access scheme [Bankov 2016, Bor 2016, Georgiou 2017, Mikhaylov 2016, Adelantado 2017, Van den Abeele 2017, Yousuf 2018].

The vast majority of related contributions addressing this scalability issue rely on some kind of synchronization to use the available bandwidth more efficiently. Indeed, sharing a common time reference among all devices allows to define transmission timeslots which can be leveraged to design more sophisticated access schemes. Remarkably, even a very simple slotted-ALOHA access theoretically provides a maximum achievable throughput of 36%, which is twice as much as its unslotted

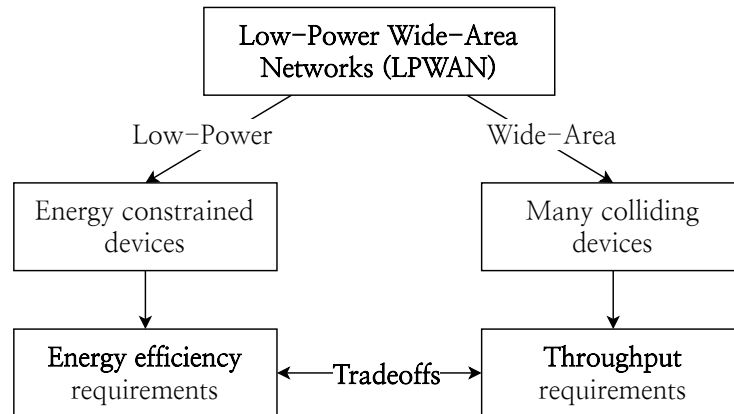


Figure 1.9: Challenge of trading-off throughput and energy efficiency

equivalent. However we must recall that LoRaWANs are typically characterized by huge numbers of devices that each transmit very small amounts of data because of DC limitations. As a result, the power consumption dedicated to synchronization purposes may very easily outweigh the power consumption dedicated to data traffic. Given that such devices are drastically energy-constrained, energy efficiency must fundamentally be considered when designing the slotted access.

All in all, it is difficult to support high traffic loads generated by a great number of devices without weighing on their energy consumption. Herein, this manuscript explores different MAC improvements to trade-off throughput and energy efficiency in LoRaWANs. We will see that there is no "one fits all" solution, and that the most efficient strategy varies depending on the traffic load.

Medium Access Control strategies for LoRa and Thesis Methodology

Dans ce chapitre, nous dressons d'abord un état de l'art des stratégies d'accès au médium utilisées dans le domaine des communications sans fil. Ensuite, nous abordons l'application aux réseaux LoRa de ces différentes stratégies d'accès par l'étude de la littérature scientifique actuelle sur le sujet. Enfin, nous décrivons l'approche abordée dans ce manuscrit de thèse en la positionnant vis-à-vis de cet état de l'art. Spécifiquement, nous prônons l'exploration de stratégies d'accès synchrones pour étendre la capacité réseau. De plus, nous défendons l'utilisation d'un mécanisme de synchronisation reposant sur la diffusion de balises par la station de base, qui fonctionneront indépendamment du nombre de terminaux à synchroniser. Nous décrivons aussi notre approche méthodologique qui repose sur différents outils, à savoir la modélisation mathématique, les simulations et les expériences réelles. Ces options sont toutes associées à des avantages et défauts, et nous montrons que leur utilisation complémentaire assure la solidité de notre travail de recherche. Enfin, nous détaillons les contributions de ce manuscrit.

Contents

2.1	Overview of Medium Access Control strategies	24
2.2	State-of-the-Art Research on Access Strategies for LoRa	27
2.2.1	Asynchronous protocols and Carrier Sense Multiple Access schemes	27
2.2.2	Synchronous protocols: Slotted ALOHA and Scheduled Access Schemes	29
2.3	Methodological approach	36

This chapter presents an overview of well-known access strategies used in wireless communications, and explores related works relating the application of these strategies to LoRa networks. Based on these premises, we justify the approach taken in this thesis. It is organised as follows. First, Section 2.1 builds a taxonomy of MAC protocols inherited from the scientific literature. Section 2.2 then references significant related works about Medium Access Control (MAC) improvements for

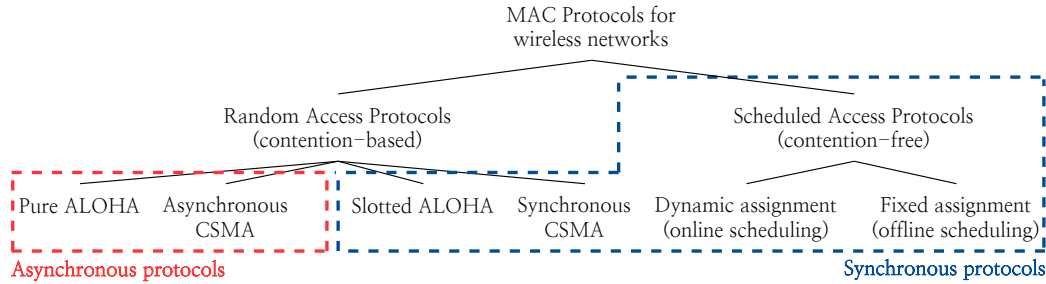


Figure 2.1: Coarse taxonomy of MAC protocols for wireless networks

such networks, and shows that the research problematic presented in chapter 1 is timely. Finally, Section 2.3 positions the thesis approach in this general context, and presents its main contributions.

2.1 Overview of Medium Access Control strategies

This Section presents a classification of channel access strategies used in radio networks. This taxonomy provides theoretical background on MAC protocols in order to later approach their application to LoRa networks with sufficient perspective.

Most MAC protocols found in modern networks are derived from a set of theoretical access strategies which have been extensively studied in the literature. In this part the well-known access schemes are inventoried, which will later allow to clearly position related contributions into the landscape of MAC protocols. Herein, a coarse taxonomy of access strategies is provided in Figure 2.1. These strategies can be divided into two great categories: *Random Access Protocols* and *Scheduled Access Protocols*. Random Access Protocols are associated with scenarios where no coordination is assumed between terminals. Transmissions occur at random times, and no priority is given to an end-device more than another. On the other hand, Scheduled Access Protocols rely on reservation systems to coordinate the transmissions of all terminals. These approaches generally rely on a centralized scheduler to organize the network traffic.

Random Access Protocols have been extensively studied and modeled by the research community. The simplest instance of such a protocol is pure ALOHA [Abramson 1970]. When a pure ALOHA device generates a frame, it transmits it immediately without getting synchronized or sensing the transmission medium beforehand. In that, the device logic is kept as lightweight as possible. However, the overall network capacity is bounded to a theoretical maximum of 18% due to frame collisions.

This asynchronous approach was later improved by Carrier Sense Multiple Access (CSMA) strategies, which consist in optimizing the channel throughput by letting devices sense the transmission medium availability before triggering transmissions. In [Kleinrock 1975c] Kleinrock and Tobagi describe 3 main CSMA vari-

ants:

1. *1-persistent CSMA*, in which a device that senses a busy channel will persist in transmitting its frame as soon as the ongoing transmission has ended. This strategy ensures that the channel will never go idle if some terminal has a message to send.
2. *Non-persistent CSMA*, in which a device that senses a busy channel will reschedule its transmission after a random time. It reduces the interference probability compared to 1-persistent CSMA, but may induce idle gaps between consecutive non-overlapped transmissions.
3. *p-persistent CSMA*, in which a device that senses a busy channel will transmit its message with a probability p once the ongoing transmission has ended.

These Carrier Sense strategies have been later improved and applied to many network types. Notable examples include CSMA/CD (Collision Detection, used in the first versions of Ethernet) and CSMA/CA (Collision Avoidance, used in WiFi).

The aforementioned access schemes assume no synchronization between end-devices. However, sharing a common time reference among all terminals allows to align all transmissions on shared timeslots. Using such timeslots enables a more efficient use of the available bandwidth. The most simple slotted access scheme is slotted ALOHA [Roberts 1975]. When a slotted ALOHA device generates a frame, it triggers its transmission at the start of the first timeslot that follows the generation. A slotted ALOHA access remarkably doubles the maximum achievable throughput compared to its asynchronous equivalent. Interestingly, asynchronous CSMA strategies may be adapted to a slotted context, also improving their performances.

With Scheduled Access Protocols, multiple access is performed over the transmission medium by assigning parts of the available bandwidth to the competing terminals. Such strategies generally rely on a central scheduler that knows the state of the whole network, and performs the bandwidth assignment. This way, a fine-grained transmission schedule can be established, capable of completely preventing collisions in the network. Scheduled access protocols are most of the time synchronized, allowing the definition of shared transmission timeslots which can be easily assigned to terminals. The multiplexing of transmissions can also be performed with several frequency channels in parallel, or with the simultaneous use of several orthogonal codes (*e.g.* the Spreading Factors in LoRa). The transmission schedule can be established either *offline* or *online*. An offline scheduling implies that the slot assignment is static, and fixed before starting the network. A static slot assignment can be performed when the whole deployment layout is known *a priori*, and where the traffic is predictable. It is typically associated with industrial *ad-hoc* networks with time-criticality constraints. On the other hand, an online scheduling designs a dynamic bandwidth allocation, performed in parallel with the data traffic. It allows much more flexibility, notably allowing devices to seamlessly

join and leave the network. However, it requires the transmission of control messages to share the transmission schedule with all involved terminals, which may weigh on the data traffic and on the energy consumption.

Throughput models for several access schemes presented in this Section are provided in Appendix A. For the sake of comparison, they have been plotted in Figure 2.2 ($a=0.01$ for CSMA schemes, *cf.* Equations A.3 and A.4). These theoretical models assume an infinite number of devices, but provide a reasonable estimation of the throughput observed in large-scale networks. Besides, they picture a best-case scenario in which perfect synchronization and perfect scheduling are assumed. The network capacity designates the maximum throughput that can be achieved with a given access scheme. Theoretical network capacities for the considered strategies are provided in Table 2.1.

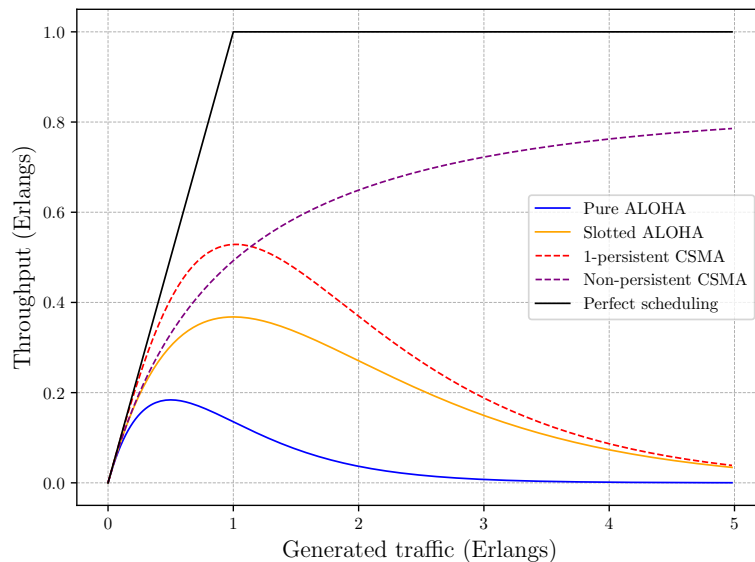


Figure 2.2: Throughput models for several Random Access Protocols ($a=0.01$)

Table 2.1: Theoretical network capacity for different random access strategies (from [Kleinrock 1975c])

Protocol	Network capacity
Pure ALOHA	0.184
Slotted ALOHA	0.368
1-Persistent CSMA	0.529
Non-Persistent CSMA	0.815
Perfect scheduling	1

This concludes our overview of theoretical MAC strategies. In the following literature review, we explore different ways these strategies were applied to LoRa networks in related contributions.

2.2 State-of-the-Art Research on Access Strategies for LoRa networks

2.2.1 Asynchronous protocols and Carrier Sense Multiple Access schemes

Let us first explore how Carrier Sense Multiple Access (CSMA) schemes [Kleinrock 1975c] have been applied to LoRa communications. First of all, it must be underlined that the first applications of these listen-before-talk approaches to LoRaWAN are still quite recent. Indeed, LoRa transceivers prior to the SX126x generation (released in 2017) were only capable of detecting preamble symbols transmitted at the beginning of a transmission, and not the data symbols that compose the rest of the frame. After that, Semtech implemented a Channel Activity Detection (CAD) feature [Corporation 2019a] allowing them to detect the presence of data symbols as well, thus enabling viable carrier sensing capabilities. Yet, it is well-known that simple CSMA protocols are sensitive to the hidden terminal problem [Tobagi 1975]. This issue occurs when two terminals are in range of the same gateway, but they are too far away from one another to be able to sense their respective transmissions. It is a threat to CSMA protocols because it leads to collisions when these terminals simultaneously have frames to transmit. This problem is emphasized in large scale LoRa networks because devices are numerous and featured with very long transmission ranges, and many different terminals may be hidden from any node.

In [Pham 2018], Pham investigates how this CAD feature can be used to adapt the 802.11 CSMA protocol and its Distributed Coordinated Function (DCF) to LoRa communications. In a nutshell, it defines a DCF Inter-Frame Space (DIFS) to normalize the sensing operation. If a node senses that the medium is free for the duration of the DIFS, then it will be able to send its packet. Otherwise, the transmission is deferred for a duration determined by a random backoff mechanism. In the author's implementation, the DIFS duration corresponds to 9 successive CAD calls, which form the Clear Channel Assessment (CCA) procedure. This preliminary evaluation reveals that the CCA lacks reliability when the distance increases, which greatly affects the good operation of the listen-before-talk protocol. The author proposes to alleviate this issue by lengthening the DIFS duration to fit the maximum frame ToA for the considered LoRa settings, but does not show that this strategy will solve the CAD issue. Moreover, additional results are required to evaluate the impact of such a strategy on the throughput, energy consumption and delay in a real-world deployment.

In [Pham 2021], the same author studies this CAD unreliability in more details. Experiments confirm that the procedure becomes unstable from 400 meters and above, and unfortunately no improvement has been seen with the newer SX1262 LoRa chips in which the CAD procedure was supposed to be ameliorated. Besides, the experiments focus on SF12 which results in the longest frame ToA and range, so we can expect that the CAD performances would be even worse when using

higher data rates. The fact that this procedure is unreliable makes the network particularly vulnerable to the aforementioned hidden terminal problem. In WiFi networks, the hidden terminal problem has been addressed by an extension of the CSMA/CA with Ready to Send/Clear to Send (RTS/CTS) messages. However this idea can not be applied to LoRa because of two reasons. First, gateways are subject to Duty Cycle (DC) constraints on the downlink channel, which puts a boundary on the amount of CTS messages that could be sent. Second, most current LoRa gateways are half-duplex [Di Vincenzo 2019], and need to stop listening on all channels when transmitting a downlink message. This means that the transmission of CTS messages would impact the reception of uplink frames, ultimately jeopardizing the network capacity in large-scale deployments. All this considered, Pham proposes to extend the CSMA protocol with short RTS messages that carry the length of the data packet to come, used in conjunction with guard times before transmissions. However this approach does not solve the hidden terminal problem, weighs on the device power consumption, and results are insufficient to prove the protocol robustness in large-scale scenarios.

In [Ahsan 2019], Ahsan *et al.* compare 3 random backoff strategies (namely binary exponential delay, binary exponential backoff and binary exponential hybrid) to decide the retransmission delay value. However all results are obtained with MATLAB simulations, so the system performance is not assessed with the CAD unreliability experienced in real deployments.

In [O’Kennedy 2020], O’Kennedy *et al.* proposes another practical evaluation of the CAD performance in order to adapt CSMA to wildlife monitoring networks. They notably show that using several successive CAD measurements allow to increase the reliability of the CCA procedure. In their experimental setup, they achieve a satisfying CCA sensitivity with 8 CAD measurements. However they also conclude that a more elaborated mechanism is required to adapt this strategy to environments with higher radio interference, and that more CADs could be needed. In that, the power consumption dedicated to the channel sensing is drastically increased.

In [Gamage 2020], Gamage *et al.* introduce another CSMA-based access scheme, in which several channel/SF pairs can be handled in parallel. Outdoor experiments are conducted with 16 nodes, but terminals are not sufficiently spread to witness the CAD unreliability issue. Therefore the protocol ability to handle bigger traffic loads in large-scale scenarios remains unsure.

Finally, additional measurements performed in [Subbaraman 2022] confirm the CAD unreliability, and show that CSMA-based approaches can therefore not be easily applied in urban deployments. Here, the proposed solution is to design new gateways with full-duplex capabilities, and have them transmit a busy signal everytime an uplink message is being demodulated. This protocol solves the hidden terminal problem and offers a better throughput than Pure ALOHA, however it requires considerable changes to the gateway hardware, and the use of a downlink channel that is not subject to DC limitations. For these reasons, the application of this solution to LoRaWAN in the near future is compromised.

In this Section, we saw that there have been several attempts aiming at applying CSMA strategies to LoRa in order to improve the MAC layer performances. However, the first real-world experiments have unveiled that the CAD mechanism available in modern LoRa chips becomes unreliable when the distance increases. Besides since LoRa gateways are half-duplex, and forced to stop listening on all channels when transmitting a downlink message, RTS/CTS-based mechanisms cannot be used to solve the hidden terminal problem. Another option to get a reliable CCA is to increase the number of successive CAD calls to assess the channel state. However this option heavily weighs on the transmission delay and on the node power consumption, and these aspects are rarely studied in depth in these preliminary contributions. For all these reasons, CSMA-based approaches are not very promising to improve the MAC-layer performances in LoRaWAN. We believe that access schemes relying on synchronization offer more interesting perspectives, since the network topology already benefits from gateways capable of broadcasting synchronization beacons and a centralized server that can handle scheduling operations. In the rest of this Section, we therefore focus on contributions that leverage a synchronization strategy to exploit the available bandwidth more efficiently.

2.2.2 Synchronous protocols: Slotted ALOHA and Scheduled Access Schemes

Let us now focus on the access schemes that rely on synchronization as a means to use the available bandwidth more efficiently. Sharing a common time reference with all devices allows the definition of transmission timeslots, which are then used to organize the transmissions. Such timeslots can be accessed in a random manner (corresponding to a Slotted ALOHA scheme) or they can be assigned to certain devices (which refers to scheduling strategies).

Wireless devices are generally featured with low-precision embedded clocks that need to be periodically realigned with the shared time reference to setup a synchronized access. Different strategies exist to perform this operation, each associated with its advantages and inconveniences. We identify three main synchronization methods:

- The *beacon-based synchronization* relies on the periodical broadcast of beacon messages containing the reference timestamp. This strategy is the one used for the LoRaWAN *Class B*, and has the advantage of allowing the synchronization of several devices simultaneously. However with half-duplex base-stations such as LoRa gateways it prevents the reception of data frames overlapping with the beacon transmission. Besides a device has first to join the base station in a pure ALOHA fashion when joining the synchronized access for the first time in order to be informed of the next beacon transmission timestamp.
- With the *ack-based synchronization* every device has to perform a synchronization request with an uplink frame and the base station embeds the reference timestamp within the acknowledgment. With this strategy, any device

may get synchronized on-demand whenever it needs it, and does not need to keep track of the beacon receptions. However in LPWAN where a huge number of end-devices depend on a few half-duplex gateways, the transmission of individual synchronization frames greatly impacts the uplink throughput. Moreover in LoRa, the DC limitations that also apply on the downlink traffic limit the number of devices that may be handled by a single base station.

- Finally, the *out-of-band synchronization* refers to strategies where reference timestamps are shared on a separate communication medium than data frames. This strategy allows to seamlessly handle the synchronization task without impacting the other communications. However it generally requires the addition of specific hardware to the end-devices in order to receive these messages, and protocols developed in such a fashion can therefore not be used with legacy devices.

In the rest of this section, we explore synchronized access schemes tailored for LoRa networks found in the research literature. In regards to the considerations above and the taxonomy presented in Section 2.1, we state the strategy used in each of these propositions, including (i) the synchronization strategy, (ii) whether the access is random or scheduled and (iii) the nature of the scheduling (online or offline).

Let us first focus on contributions that target a slotted ALOHA access. In [Polonelli 2019], the synchronization is performed with individual acknowledgments. The estimated drift of the device real-time embedded clocks is used to estimate the worst-case synchronization error over time. A hardware implementation is used to validate the protocol design on a few devices. However, an ack-based synchronization such as this one is not scalable to large-scale deployments because of the gateway half-duplex nature and DC requirements.

Another synchronization strategy has been explored in [Beltramelli 2021]. In this case, out-of-band FM-RDS signals have been used to share the common time reference and enable a slotted ALOHA access. Timing errors have been measured on real hardware, and simulations have been used to assess the performances of the proposed protocol. The main drawback of this solution is that it requires the addition of specific hardware to the nodes in order to run the protocol.

Many contributions have proposed offline scheduling algorithms. These approaches generally target industrial networks where the topology is known *a-priori*, but are subject to delay constraints. In [Rizzi 2017a], Rizzi *et al.* propose a first attempt at using Time Slotted Channel Hopping (TSCH) within LoRaWAN, relying on very large slots encapsulating *Class A* exchanges. It sets a scheduled access where the schedule is built offline with a set of devices known *a-priori*, and where several frequency channels are leveraged. However the idea of reserving very large slots for full *Class A* transactions that include very spaced out reception windows result in a waste of the available bandwidth. Herein, the protocol targets soft real time industrial applications by bounding the delay between messages, at the cost

of dwarfing the achievable throughput. Experiments have been performed on real devices, but the problem of distributed synchronization is not addressed.

In [Islam 2018], several offline scheduling algorithms are proposed in order to reach real-time constraints with duty-cycle requirements. Hardware experiments have been conducted, however since the algorithms are offline the protocol does not allow devices to join or leave the network. It is also possible that the system does not find a feasible schedule for a given layout. Besides, the question of distributed synchronization is avoided by equipping devices with external real-time clocks.

Leonardi *et al.* introduce an hybrid protocol designed to support both real-time and non real-time traffic flows [Leonardi 2018]. It is based on a beacon synchronization, and the beacon period is split onto two sub-periods: one contention-based and one contention-free. Simulation results validate that different quality of service requirement can be handled with this strategy. During the contention-free period, the schedule is established offline with a simple TDMA strategy. While this strategy allows to mix-up real-time and non real-time traffic, it reduces the overall network throughput compared to the legacy access. Besides, it induces a higher energy consumption on all devices because of beacon receptions.

Another strategy has been explored in [Piyare 2018], that introduce an offline-scheduled data collection relying on out-of-band synchronization through low-power wake-up receivers. The system can be activated on-demand, and is capable of waking-up the devices to setup a TDMA-based data collection. A proof-of-concept hardware implementation has been realized to validate the protocol design. The main drawbacks in this solution is that the network layout has to be fixed *a-priori*, and that devices need to be equipped with specific hardware to receive the wake-up messages out-of-band.

Another offline scheduling solution is explored in [Tessaro 2018], that targets industrial networks featured with periodic, predictable traffic. This time, synchronization is performed with beacon broadcasts. Clock drift measurements and corrections are notably emphasized in this last contribution. Hardware experiments validate that the clock error can be maintained within the requirements. However this protocol cannot handle general-purpose LoRa deployments with unpredictable traffic and layout changes.

Finally, let us consider contributions that introduce an online scheduling of the transmissions. The authors of [Reynders 2018] argue that the capture effect [Leentvaar 1976] results in a loss of fairness that disadvantage the nodes placed the furthest away from base stations. Indeed, the capture effect refers to a phenomenon that allows the correct demodulation of a frame when several frames are transmitted simultaneously. We have also studied this effect when dealing with our hardware testbed, therefore more information about it is available in Chapter 5. The occurrence of capture events implies that the devices placed near gateways benefit from a greater share of the bandwidth. In order to alleviate this issue, they propose a scheduling scheme in which coarse timeslots are used to group devices using similar transmission powers and SF. Synchronization is handled with beacons that additionally contain the transmission power and SF schedule. Devices are then

able to adapt their transmission parameters accordingly. The protocol is evaluated with the NS-3 simulator, and results show that the fairness is improved thanks to this strategy that reduces the occurrence probability of capture events. However the overall throughput is not significantly increased, and a significant power consumption is induced by the beacon receptions.

In [Abdelfadeel 2018], another strategy is explored in which devices bufferize their messages for long periods of time. They are then woken-up when a gateway is available, and data is collected in a bulk with a fine-grained online scheduling scheme. This solution targets remote areas that could be reached periodically by a moving gateway. The contribution contains a joining procedure called when the devices are woken-up, and a scheduling algorithm to organize the data collection phase. Interestingly, the joining procedure first uses an ack-based synchronization to scan the devices, and then a beacon-like synchronization just before the data collection, when sharing the fine-grained schedule. In this case, the schedule is built online, so the protocol is able to handle any discovered device layout. Simulations are used to evaluate the protocol, and show that the bulk data collection performs well in terms of delay, especially in small networks. This work was then extended by the same authors to handle continuous traffic flows [Zorbas 2020]. This time, periodic beacon-like frames are used to maintain synchronization and perform group-acknowledgments. These SACK (Synchronization and ACKnowledgment) frames also contain the encrypted timeslot allocation schedule. Simulations and an hardware implementation have been realized to evaluate the protocol, and results show that it greatly improves the packet delivery ratio compared to the legacy protocol. However, the SACK packet receptions greatly increase the device power consumption. Authors claim that they offer a better energy efficiency than legacy LoRaWAN, but they only consider the acknowledged mode where devices have to spend energy for ACK receptions.

In [Haxhibeqiri 2019], Haxhibeqiri *et al.* introduce an online scheduling protocol that utilizes individual acknowledgments for synchronization. The central server is in charge of periodically computing the schedule, which is then transmitted to devices with an efficient Bloom filter-based encoding. This contribution additionally presents a full scheduling mechanism built on top of *Class A*. Simulations are used to validate the protocol design, and show that an optimal scheduling could be achieved in between synchronization phases. However, an overlooked aspect in the paper is the scalability issue experienced with the ack-based synchronization. Indeed, the DC constraints on the gateway bound the number of devices that can be handled by this strategy. This contribution has later been extended with a hardware implementation based on real clock drift measurements [Garrido-Hidalgo 2021].

This overview of the most relevant synchronized access schemes for LoRa found in the research literature confirms that improving the MAC layer strategy is a timely research topic. A summary of all these references is provided in Table 2.2. We highlight that no contribution proposes a scalable synchronization strategy capable of handling high traffic loads while focusing on energy efficiency as a prime requirement. In the next Section, we position the manuscript in this context and

justify our approach in regard to the state-of-the-art.

Ref.	Protocol type	Synchronization strategy	Tools
[Pham 2018]	CSMA	None	Hardware
[Pham 2021]	CSMA	None	Hardware
[Ahsan 2019]	CSMA	None	Simulation
[O’Kennedy 2020]	CSMA	None	Hardware
[Gamage 2020]	CSMA	Coarse beacon-based synchronization	Hardware
[Subbaraman 2022]	CSMA	None	Sim, HW
[Polonelli 2019]	Slotted ALOHA	Acknowledgments	HW, modeling
[Beltramelli 2021]	Slotted ALOHA	Out-of-band (FM-RDS signals)	Sim, HW
[Rizzi 2017a]	Offline scheduling with large slots	Required but not addressed.	Hardware
[Islam 2018]	Offline scheduling	External real-time clocks	Sim, HW

[Leonardi 2018]	Partial offline scheduling	Beaconing	Simulation
[Piyare 2018]	On-demand offline scheduling (TDMA)	Out-of-band (wake-up receivers)	Hardware
[Tessaro 2018]	Offline scheduling	Beaconing	Hardware
[Reynders 2018]	Lightweight online scheduling	Beaconing	Simulation
[Abdelfadeel 2018]	Online scheduling for bulk collection.	Beaconing	Simulation
[Zorbas 2020]	Online scheduling with acks	Beaconing	Sim, HW
[Haxhibeqiri 2019]	Online scheduling	Acknowledgments	Simulation
[Garrido-Hidalgo 2021]	Online scheduling	Acknowledgments	Hardware

Table 2.2: Access scheme strategies in the literature

2.3 Methodological approach

In Chapter 1 we introduced the LoRaWAN technology in details, and we highlighted the scalability issues experienced with the current Pure ALOHA protocol in large scale deployments. On this regard, other access strategies have been proposed to improve the network performances. In this Chapter we therefore presented the current state-of-the-art regarding Medium Access Control in LoRa networks. In Section 2.1 we built a classification of well-known MAC strategies, and in Section 2.2 we inventoried relevant alternative protocols that were introduced by the research community. In this landscape, we show that no proposition is able to alleviate the scalability issues while maintaining a good device energy efficiency and overall network flexibility. This finding is problematic since LoRa targets battery-powered devices that are supposed to operate over the span of several years. In that, this manuscript targets the design of scalable, yet energy efficient MAC protocols for LoRa deployments and LPWAN in general.

Let us now explicit the chosen approach for this thesis. We have identified three main alternatives to the legacy pure ALOHA access: *(i)* slotted ALOHA, *(ii)* CSMA techniques and *(iii)* scheduling protocols. We have seen that CSMA approaches are not yet applicable to LoRa networks due to an unreliable channel activity detection feature, that prevents devices from performing an adequate channel sensing when interfering nodes are too far away. Besides, we have seen that the use of RTS/CTS messages can not be used to work around the hidden terminal problem because:

1. The downlink channel in LoRa is subject to a 10% Duty Cycle [LoRa Alliance 2021], which bounds the number of CTS messages that could be transmitted.
2. LoRa gateways are half duplex [Di Vincenzo 2019], and need to stop listening on all channels before sending a downlink frame. In consequence, the transmission of downlink messages greatly affects the uplink throughput [Pop 2017].

As a result, we believe that the synchronized approaches are more promising than carrier sense protocols. We have seen that synchronization can be performed with three main strategies: *(i)* with individual acknowledgments, *(ii)* with beacons and *(iii)* out-of-band. For the same reasons that prevent the use of RTS/CTS messages, individual acknowledgments are not a scalable synchronization approach for LoRa networks. Out-of-band synchronization requires the addition of specific hardware to the LoRa nodes, which prevents legacy devices from utilizing these protocols. However, a beaconing strategy can be applied to any network size because the time reference is shared with periodical broadcasts. Remarkably, the LoRaWAN *Class B* already relies on a beacon-based synchronization, which would facilitate the integration of protocols using a similar approach. As a result, this manuscript focuses on synchronized access schemes that rely on beacon messages to share the common time reference with all nodes.

Beacon synchronization is scalable, however it increases the device power consumption. Indeed the device radios need to be periodically switched to a listen-

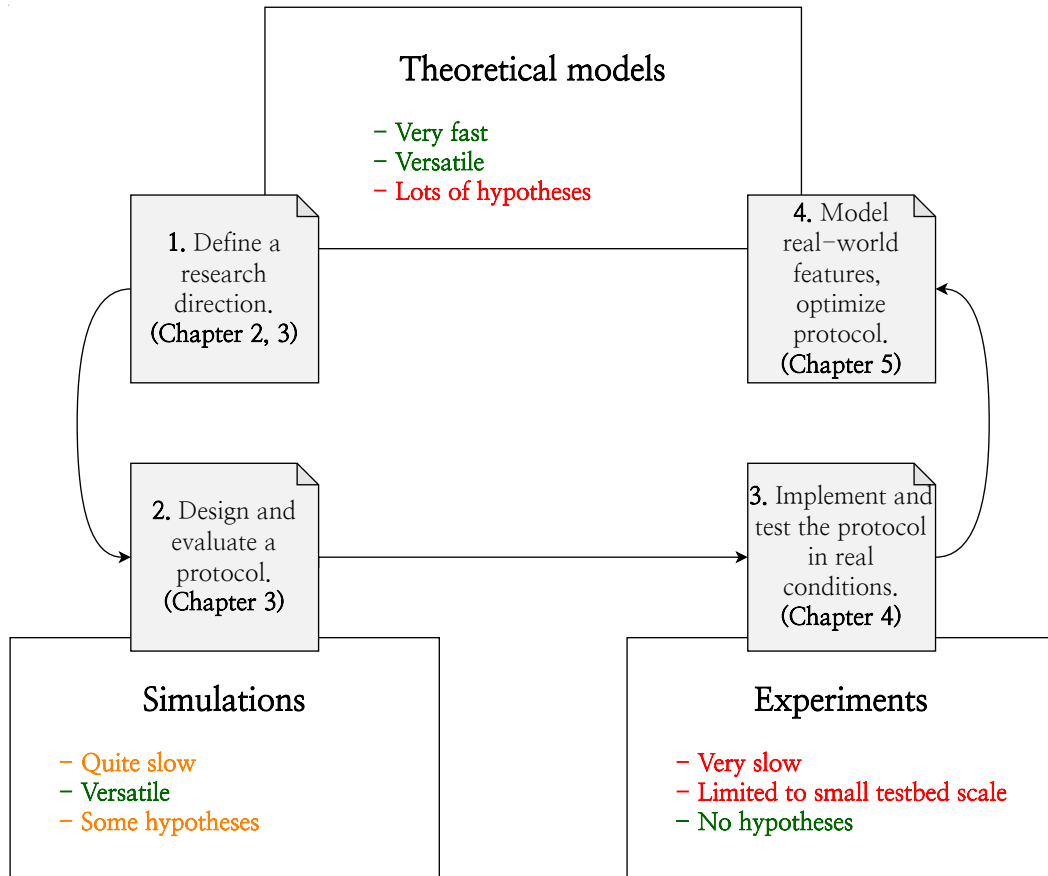


Figure 2.3: Research methodology with complementary tools

ing state in order to receive the beacon broadcasts, which consumes a significant amount of power (*c.f.* Table 1.2). The reception of synchronization beacons must therefore be handled wisely to ensure a long battery lifetime. This aspect is most often overlooked in the literature. In this manuscript, we leverage beacon-skipping mechanisms to enable slotted accesses for LoRa communications while ensuring a satisfying device energy-efficiency. We leverage this strategy to setup a slotted ALOHA access, however it may also be used for scheduling schemes.

The tools leveraged in this manuscript include modeling, simulations and hardware implementation, which unveil a clear understanding of real LoRa networks and allow the extrapolation of this knowledge to large-scale scenarios. Let us now go through the strengths and weaknesses of each approach, and see how they can be used to obtain solid results. The research pace when leveraging models is very fast, because a large spectrum of scenarios may be evaluated simply by changing the model parameters. However a lot of hypotheses are generally made with this approach, and results thus represent an approximation of reality. This is the approach we first used to define a research direction. Simulations allow to evaluate the behavior of large scale networks much more precisely than when using the base

models, but the process of programming and running the simulation campaigns is slower. We leveraged simulations to design our protocol and evaluate its performances. Finally, experiments force the consideration of all constraints experienced in real networks. This strategy allows the collection of real network traces, but limits the study to a small scale testbed topology. In the latest steps of our research, we leverage observations made on the testbed to update our models. The resulting gain in accuracy with these tools ultimately allows to extrapolate results to other scenarios and fine-tune the protocol parameters. This methodological approach is schematized in Figure 2.3.

All in all, this manuscript presents a set of contributions to support the development of a scalable, yet energy efficient MAC layer for LoRa networks. It shows that traffic-aware protocols are required to wisely select the optimal strategy depending on the network conditions. In Chapter 3 we introduce *Class S*, a new device class that leverages a beacon-based synchronization to allow the use of timeslots for uplink communications. We enlighten the need for such a beacon-skipping mechanism to guarantee the energy efficiency of the protocol. We also provide models to evaluate the performances of the protocol, and validate them with a simulations. Results show that *Class S* allows to improve the network capacity, and also ameliorates its energy efficiency for high traffic loads. However, models also reveal that *Class A* remains more efficient for low traffic situations. As a result, we then introduce a switching mechanism able to alternate between asynchronous and synchronous access schemes depending on the situation. It relies on probing and fingerprinting techniques to evaluate the current network state, and provides the packet signaling to trigger protocol changes when needed. Simulations are realized to show that this solution can improve the capacity and energy efficiency of deployments with varying traffic loads. In Chapter 4, we additionally consider real hardware constraints to support the implementation of the protocol on a LoRa testbed. In details, we measure the clock drift on real devices, and improve *Class S* to be robust to the worst-case skew. To do so, guard times are added to every slot, and the size of these margins is directly linked to the beacon-skipping mechanism. Performance evaluation on the hardware testbed notably unveils that the capture effect has a huge impact on the overall throughput, and that the models used until then actually underestimate the real network performances. Herein, we dedicate Chapter 5 to the addition of real-world observations to our theoretical modeling of the network performances. First, the LoRaSync slot margins used to cope for clock drifting issues are integrated to throughput and energy efficiency models. Such a modeling then allows to derive the most energy efficient slot size for every traffic load. Second, we account for the Capture Effect in experimental models that rely on measurements performed on our testbed to accurately picture the real-world throughput.

From Theory to Simulations: designing a slotted access that fits LoRaWAN

Dans ce chapitre, nous utilisons la modélisation des modes d'accès pure et slotted ALOHA comme point de départ pour mettre en évidence les avantages de la stratégie synchrone.

Ensuite, nous proposons une nouvelle classe de dispositifs LoRa qui permet la conception de stratégies d'accès synchronisées malgré les contraintes de cette technologie. Des campagnes de simulations sont réalisées pour évaluer les performances d'une telle approche et la comparer à la version de base du protocole. Ces premiers résultats révèlent que slotted ALOHA peut proposer un meilleur débit et une meilleure efficacité énergétique que son équivalent asynchrone dans des réseaux gérés. Cependant, pure ALOHA reste préférable pour les faibles charges de trafic.

En conséquence, nous introduisons ensuite un mécanisme adaptatif capable d'alterner entre la version synchrone et asynchrone du protocole en fonction des conditions réseau sondées. Cette approche consciente du trafic est testée en simulation, confirmant que le mode d'accès peut être adapté de manière transparente en sélectionnant toujours la stratégie la plus efficace.

Contents

3.1	Theoretical motivation: Pure and Slotted ALOHA throughput in ideal networks	41
3.2	From Class B to Class S: defining a slotted framework tailored for LoRa communications	42
3.2.1	Protocol design	44
3.2.1.1	Advantages of beacon-based synchronization	44
3.2.1.2	Slotframe structure	44
3.2.2	Simulation-based performance evaluation of a Slotted ALOHA scheme over Class S	46
3.2.2.1	Simulation setup	46
3.2.2.2	Simulation results	47
3.3	TREMA: a traffic-aware MAC protocol that maximizes energy efficiency in LoRa networks	49

3.3.1	Design of TREMA	50
3.3.1.1	A time-synchronized scheduled access over Class S	51
3.3.1.2	Prior fingerprinting of the deployment	52
3.3.1.3	Network probing and decision mechanism	53
3.3.1.4	Signaling protocol	54
3.3.2	Performance evaluation	56
3.3.2.1	Preliminary deployment fingerprinting	58
3.3.2.2	Online testing	59
3.3.2.3	Overall performance evaluation	61
3.3.2.4	Switching duration	61
3.3.2.5	Performance gain	61
3.4	Conclusions and future works	62

In chapter 1 we have presented LoRaWANs in details, and showed that these networks will be subject to scalability issues in the near future. Indeed, the asynchronous pure ALOHA access employed in these deployments limits the maximum achievable throughput to 18% because of frame collisions [Abramson 1970]. We have demonstrated why among all MAC strategies found in the literature, synchronized schemes seem to be the most promising solution to this issue. In this chapter, we therefore explore the application of a slotted ALOHA scheme to LoRa deployments.

First, we provide theoretical background on ideal pure and slotted ALOHA protocols in section 3.1 to highlight the benefits of the synchronized approach.

Second, we introduce *Class S* in section 3.2 to enable such a slotted access in the context of LoRa networks. *Class S* relies on the LoRaWAN *Class B* beaconing mechanism for synchronization and defines a slotframe able to support uplink communications. The analysis is backed with a performance evaluation of the throughput and energy efficiency of *Class S*-enabled devices in large-scale scenarios with the LoRaWAN-Sim network simulator. Herein, we verify that even a very simple slotted ALOHA scheme nearly doubles the maximum achievable throughput in this context. Besides, results show that slotted ALOHA also outperforms pure ALOHA in terms of energy efficiency at high traffic loads if devices efficiently handle the reception of synchronization beacons. In a nutshell, pure ALOHA performs better at low traffic loads and slotted ALOHA becomes preferable in highly congested situations.

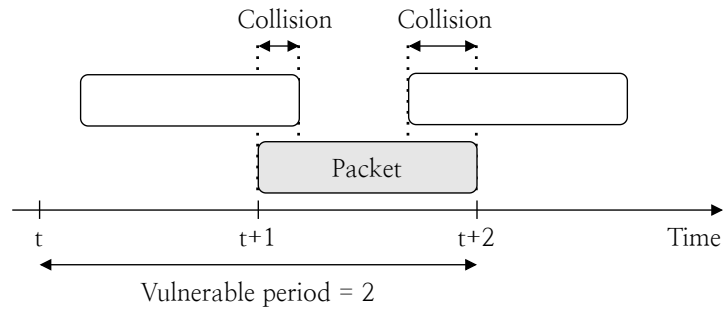
Finally, in section 3.3 we present the TREMA switching mechanism to dynamically synchronize and desynchronize devices depending on the variations of the offered load, and thus always select the most energy efficient strategy. Simulation results show that such an adaptive scheme outperforms the legacy protocol when the traffic load varies.

3.1 Theoretical motivation: Pure and Slotted ALOHA throughput in ideal networks


In order to motivate our research on synchronized access schemes for LoRa networks, this section provides theoretical background on ideal pure and slotted ALOHA protocols.

As presented in section 2.1, the Pure ALOHA protocol is the simplest random access based on an asynchronous scheme. Indeed, a pure ALOHA node that needs to transmit a frame does it instantly without getting synchronized or sensing the transmission medium beforehand. No coordination is made between devices competing for the same channel. Such a scenario is pictured in Figure 3.1a. We assume a device that transmits a frame featured with a Time-on-Air of 1 at $t = 1$. Here, the frame may collide with any interferer between $t = 0$ and $t = 2$.

However with slotted ALOHA we assume all nodes are synchronized and forced to wait for the next slot start before triggering a transmission. This scenario is pictured in Figure 3.1b. Similarly to before, a device transmits a frame featured with a Time-on-Air of 1 at $t = 1$. Yet this time the transmission may only collide with interferers that generated a frame between $t = 0$ and $t = 1$. In that, the vulnerable period length is remarkably divided by two.



(a) Pure ALOHA



(b) Slotted ALOHA

Figure 3.1: Vulnerable periods with theoretical Pure and Slotted ALOHA

In order to approach these concepts more formally, let us now introduce the most usual ALOHA throughput modeling with a finite number of devices n , that was

first presented in [Kleinrock 1975a]. The number of frames generated by any device during a duration t is modeled by a Poisson distribution of parameter λ . Besides, we make the hypothesis that all frames in the network have the same length, defining the reference time unit of the modeling. Additionally, Pure ALOHA devices do not use any kind of buffering, and slotted ones use a buffer of size 1 that simply allows them to wait for the start of the next slot. A buffer of size 1 is assumed, therefore terminals ignore new packet arrivals while they are busy trying to transmit a packet. Therefore, the probability p that a pure ALOHA device generates a frame during a time unit follows the cumulative distribution function of the exponential law [Brand 2002]:

$$p = p[X \leq 1] = 1 - e^{-\lambda} \quad (3.1)$$

The probability that exactly k among n devices generate a frame during a time unit can then be derived with the binomial coefficient:

$$P[k \text{ in } n] = \frac{n!}{k!(n-k)!} p^k (1-p)^{n-k} \quad (3.2)$$

For Pure ALOHA, the average throughput $T_p(n, p)$ (in erlangs) is equal to the probability that exactly one device generates a frame during a time unit, and that none of the other $(n-1)$ devices do during the previous one:

$$T_p(n, p) = P[1 \text{ in } n] \cdot P[0 \text{ in } (n-1)] = np(1-p)^{2(n-1)} \quad (3.3)$$

For Slotted ALOHA, the average throughput $T_s(n, p)$ (in erlangs) is simply equal to the probability that exactly one device generates a frame during a time unit:

$$T_s(n, p) = P[1 \text{ in } n] = np(1-p)^{(n-1)} \quad (3.4)$$

These models have been plot in Figure 3.2 in the particular case of $n = 1000$ devices. They confirm that the network capacity (*i.e.*, the maximum achievable throughput) experienced with pure and slotted ALOHA protocols is respectively 18% and 36% [Roberts 1975]. Intuitively, this results is in accordance with the fact that the vulnerable period is divided by 2 in the slotted version.

We see that synchronization allows to double the network capacity, which motivates the idea of setting slotted LoRa communications to alleviate scalability issues. In the next section, we will see how such an approach can be adapted to the specificities of the LoRa technology.

3.2 From Class B to Class S: defining a slotted framework tailored for LoRa communications

We have demonstrated why among all MAC strategies found in the literature, synchronized schemes seem to be the most promising solution to scalability issues in LPWAN. We have also argued that among all available synchronization strategies,

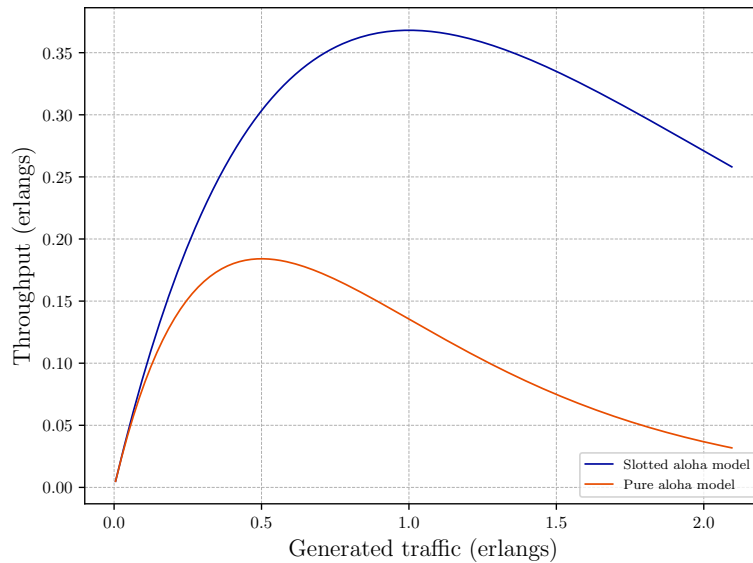


Figure 3.2: Theoretical throughput for Pure and Slotted Aloha with $n = 1000$ devices.

beacon broadcasting was the best choice for these networks. Interestingly, a beacon-based synchronization is already present in the LoRaWAN specification, as it is part of the LoRaWAN *Class B*. However the purpose of *Class B* is to bound the delay of downlink communications, which does not help solving the scalability limitation. In this chapter we introduce *Class S*, which takes advantage of the *Class B* beacon synchronization to setup a slotted access designed for uplink communications. In a nutshell, the beacon period is divided into timeslots whose length is equal to the maximum frame ToA allowed with the selected transmission parameters of supporting different kinds of slotted accesses, from slotted ALOHA schemes to scheduled access protocols, which in turn will allow to improve the overall network capacity. We also provide a performance evaluation of networks operating the legacy pure ALOHA *Class A* and a slotted ALOHA scheme setup with *Class S*. It shows that the slotted ALOHA scheme can outperform pure ALOHA in terms of both throughput and energy-efficiency in congested networks. This is however only possible if the reception of synchronization is handled wisely. Indeed, the reception of synchronization beacons forces devices to regularly switch their radios to a listening state, which may consume a considerable amount of power. In the particular case of LPWAN in which every node transmits only a few messages per day, staying synchronized therefore represents a considerable fraction of the overall energy consumption. In that, we show that a beacon-skipping mechanism is highly desired in order to guarantee a satisfying energy efficiency, and provide a long lifetime expectancy to the batteries powering the LoRa devices. To the best of the author's knowledge, this finding has never been mentioned in the related literature before. This beacon-skipping mechanism will therefore be an essential cornerstone of the

contributions introduced in the following sections.

3.2.1 Protocol design

3.2.1.1 Advantages of beacon-based synchronization

Let us first recall the reasons why designing a beacon-based synchronized access for LoRa communications is highly desired in the current state-of-the-art. In section 2.2, we have seen that CSMA schemes were not applicable to existing LoRa networks. Indeed, current LoRa chips display an unreliable Channel Activity Detection procedure, which makes the deployment dramatically vulnerable to the hidden terminal problem. The research targeting MAC layer enhancements for LoRaWAN is therefore naturally pushed towards synchronized access schemes.

Given that LoRa networks operate over Duty-Cycled ISM radio bands, gateways are limited in terms of downlink transmission capabilities. As a result, the strategy that consists in synchronizing devices by the means of individual acknowledgments does not scale well when the number of end-devices grows. Moreover LoRa gateways are half-duplex, so using such a strategy would additionally weigh on the uplink throughput. Interestingly, these are the same reasons why a RTS/CTS mechanism cannot be used to address the hidden terminal problem experienced with CSMA schemes.

Another option would be to synchronize devices using out-of-band communications. However this would require the addition of specific hardware to the LoRa terminals, thus increasing their complexity and preventing backward compatibility. As a result, sharing the common time reference using beacon broadcasts seems to be the best solution to synchronize LoRa networks.

In the rest of this Section, we therefore show how the *Class B* beaconing mechanism can be adapted to support a slotted access for uplink communications.

3.2.1.2 Slotframe structure

BEACON_PERIOD portion	Duration
BEACON_RESERVED	2.12 s.
BEACON_WINDOW	122.88 s.
BEACON_GUARD	3 s.

Table 3.1: Duration of each portion of the *Class B* and *S* beacon period

This proposal relies on the LoRaWAN *Class B* synchronization scheme to slice the beacon window into timeslots large enough to match the Time-on-Air (ToA) of the longest LoRa frame allowed with the selected transmission parameters. As the *Class B* access scheme enhances *Class A* to bound the downlink delay, the introduced *Class S* (“Slotted”) is meant to enhance *Class B* to increase the uplink throughput. In that, the length of a *Class S* timeslot is set to be a multiple of the 30 ms duration of a *Class B* ping slot, thus facilitating backward compatibility with




Figure 3.3: From *Class B* to *Class S*: adapting the beacon window to slotted uplink communications

the current LoRaWAN standard. It is defined as such (in seconds):

$$L_{slot} = \left(\left\lceil \frac{ToA_{max}}{30 \cdot 10^{-3}} \right\rceil + 1 \right) \cdot 30 \cdot 10^{-3} \quad (3.5)$$

We have chosen to define the slot length as the minimal multiple of 30 ms required to fit the longest frame ToA allowed, to which we add 30 additional ms. The reason behind these additional 30 ms is (i) to anticipate clock drift issues that will occur when implementing *Class S* on real hardware, and (ii) to allow the implementation of possible sensing windows in future improvements of the protocol.

Herein, the slot size computation has been done for SF7 with BW125 (DR5, *c.f.* section 1.2.4) as it offers the smallest power consumption and the highest data rate. Yet, similar considerations can be drawn to handle several orthogonal SFs in parallel. All other parameters are set to configure the longest frame ToA allowed with this SF. The maximum MAC payload length allowed for DR5 has been considered, *i.e.*, 250 bytes [LoRa Alliance 2021], to which the 5 bytes of MAC header and Message Integrity Code are added [LoRa Alliance 2020a]. CR is set to 4, and explicit header is enabled.

All in all, the maximum frame ToA is 626.94 ms (*cf.* equation 1.5). Hence, $\lceil 626.94/30 \rceil + 1 = 22$ ping slots are required to fit the corresponding transmission. This corresponds to a timeslot duration of $L_{slot} = 22 \cdot 30 = 660$ ms. Following such design guidelines, $\lceil 122880/660 \rceil = 187$ timeslots can fit into the BEACON_WINDOW, with the last one overlapping onto the BEACON_GUARD interval. The *Class S* slot-frame depicted above has been represented in Figure 3.3, along with the LoRaWAN

Class B for the sake of comparison.

3.2.2 Simulation-based performance evaluation of a Slotted ALOHA scheme over Class S

The slotted scheme introduced with *Class S* is general enough to be used by any slotted access, random or scheduled. This section describes a preliminary performance evaluation of such a strategy in the particular case of unconfirmed uplink communications randomly initiated at the beginning of any *Class S* timeslot. In this way, the resulting access scheme is slotted ALOHA.

3.2.2.1 Simulation setup

Table 3.2: Simulation parameters

Parameter	Value
Simulation duration	24 hours
Number of seeds	10
Frame time-on-air	626.94 ms
Beacon time-on-air	173.06 ms
Timeslot size	660 ms
Data rate	DR5 (SF7 with bandwidth 125 kHz)
Frame generation	Poisson with avg. 1 pkt./h.
Number of channels	1
Number of devices	From 100 to 9000
Device buffer size	1
n_{skip}	{0, 4, 9}
Uplink Duty Cycle	1%
Device buffer size	1
Downlink data messages	Disabled
Acks and retransmissions	Disabled
Sensor voltage	3.3 V
Sensor current intensity	20 mA TX, 10.8 mA RX [Semtech Corporation 2019]

The proposed *Class S* framework has been evaluated with the LoRaWAN-Sim simulation environment, which is an ad-hoc LPWAN simulator developed at LAAS-CNRS. This method allows us to precisely implement LoRaWAN features, such as Duty Cycle, beacon receptions, reception slots or device buffers. In that, simulation is more accurate than a formal modeling, and contrary to a testbed implementation it grants the possibility to estimate the behavior of large-scale deployments.

We need to check that the simulator operates as expected in order to ensure that the results can be trusted. For this purpose, usual ALOHA schemes have been implemented in LoRaWAN-Sim, and they were validated with well-known associated models. After this step, all features of the LoRaWAN *Class A* were implemented, as well as a slotted ALOHA scheme built on top of the *Class S*

presented above. The reception slots RX1 and RX2 have been implemented in both cases, according to the LoRaWAN specification.

Networks operating in *Class A* are compared with ones using Slotted ALOHA over *Class S*. For these simulations, only one radio channel has been used and each deployment is evaluated for a 1 day duration. In compliance with European regulations, end devices are subject to a 1% duty-cycling limitation. Each device pseudo-randomly generates Poisson traffic with an average of 1 frame per hour, matching typical LoRaWAN-based application requirements. The number of end devices has been varied from 100 to 9000. For each device density, 10 seeds have been used to feed the random number generator, thus obtaining 10 different realizations of the same network configuration. The network throughput and the energy efficiency have been evaluated for each scenario, by averaging the results over those 10 realizations. All simulation parameters are summed-up in table 3.2.

From Section 1.1.3, we know that a LoRa chip consumes a significant amount of power in reception state. Therefore, receiving synchronization beacons will naturally impact the energy efficiency of LoRa devices, especially if performed every 128 seconds as indicated in the *Class B* specification [LoRa Alliance 2020a]. Indeed, re-synchronizing at this rate would seem exaggerated compared to the fact that LoRa devices generally transmit only a few messages per day because of Duty Cycle regulations. Interestingly, the standard also indicates that end devices should be able to maintain a beacon-less *Class-B* operation for 2 hours. If we consider low-quality, cheap embedded clocks featured with a substantial drift of ± 30 ppm RTC XTAL, we know that in the worst case a skew of $30 \cdot 10^{-6} \cdot 128000 = 3.84$ ms may be experienced per beacon window. In Section 3.2.1 we have defined *Class S* with 660 ms slots for a maximum ToA of 626.94 ms, which leaves a margin of 33.06 ms between the transmissions of two successive slots. As a result, it is possible for devices to skip the reception of a certain amount of beacons without generating overlaps between transmissions assigned to different subsequent slots. In the simulations, we have therefore studied three different cases, in which devices receive either (i) all beacons, (ii) one beacon out of 5 or (iii) one beacon out of 10. We note n_{skip} the number of beacons that are skipped by devices between each actual beacon reception. We will see that such a beacon-skipping mechanism allows to significantly improve the overall network energy efficiency.

3.2.2.2 Simulation results

Let us now study the results obtained following the simulations presented above. In all curves, 95% errorbars have been computed with Student's t law to evaluate the variability of the results obtained with the different seeds. Given that the simulated time is quite long (an entire day), this variability qualitatively seems relatively small.

The simulated throughput is evaluated in Erlangs, which means that it is expressed as the ratio between the channel time dedicated to useful transmissions over the overall time. It has been plotted in Figure 3.4 for *Class A* Pure ALOHA, and

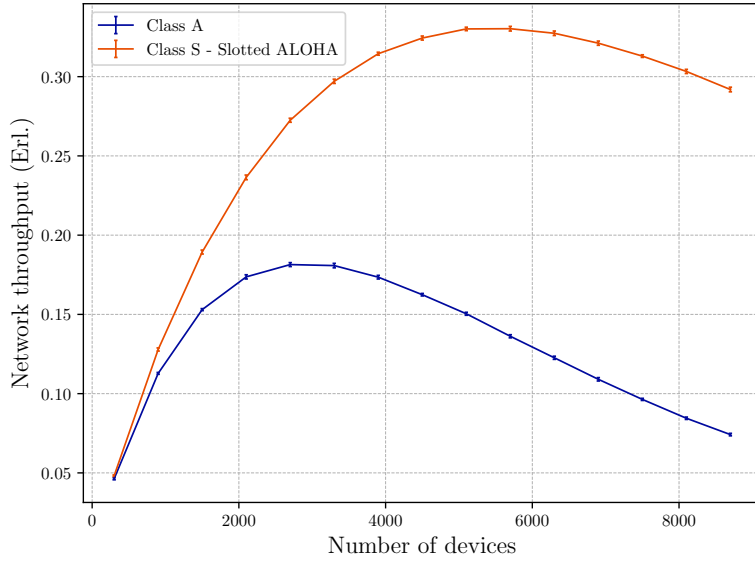


Figure 3.4: Throughput of *Class A* and *Class S* LoRa deployments.

the Slotted ALOHA scheme built over *Class S*.

The maximum achievable throughput is 0.18 at ~ 2750 devices for the *Class A*. Instead, the slotted ALOHA access implemented on the top of *Class S* permits a maximum achievable throughput of ~ 0.33 at about 5500 devices. Such a value is lower than the theoretical $e^{-1} \approx 0.36$ for slotted ALOHA networks, because no transmissions can be triggered during the BEACON_RESERVED and most of the BEACON_GUARD intervals. With these preliminary results, we verify that synchronization allows to increase network capacity by better exploiting the available bandwidth.

Energy efficiency is evaluated in Bytes per Joules [Björnson 2018], and obtained through the division of the total useful transmitted bytes by the overall energy spent for transmission and synchronization purposes. It is important to specify that the energy consumed for a transmission attempt is evaluated as the product of the frame ToA, the current drawn by a typical SX1276 LoRa transceiver in TX mode [Semtech Corporation 2019], 20mA, and a typical battery voltage, 3.3V. Instead, the energy spent for synchronization is evaluated as the product of the total time a device keeps its radio on for listening to and receiving beacons, the typical current drawn by a SX1276 transceiver in RX mode, 10.8mA, and the same typical battery voltage. The overall network energy efficiency has been plotted for all access schemes and n_{skip} values in Figure 3.5.

Results highlight that a beacon-skipping mechanism is crucial for energy efficiency. For low densities, *Class A* is more efficient as collisions are rare and spending energy for synchronization is energy-wasting. However, when 80% of the beacons are skipped (i.e., when 1 beacon out of 5 is still listened to by end devices) and the density of devices increases beyond 4000, the slotted access results more energy-

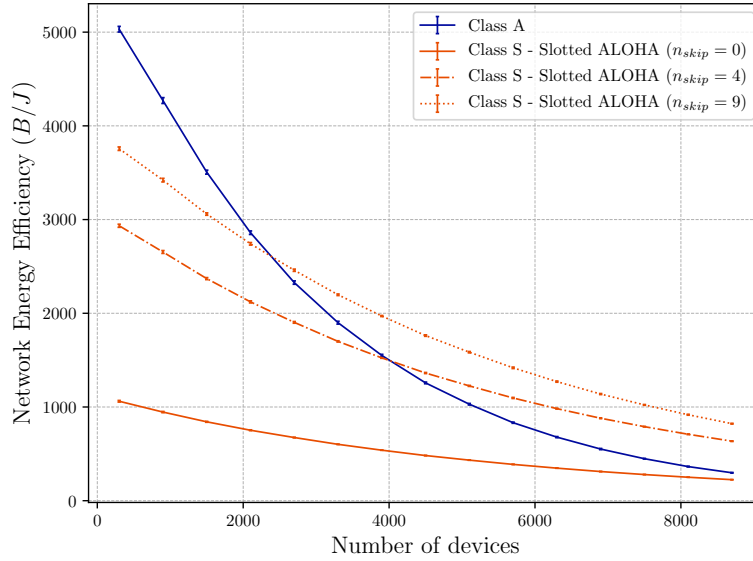


Figure 3.5: Energy efficiency of devices using *Class A* and *Class S*.

efficient than the default *Class A* communication pattern. More interestingly, when end devices keep listening to beacons once every 10, and the density of end devices is the one that maximizes the throughput when they are operating in *Class A* (i.e., around 2750 end devices per unit area, see Fig. 3.4), the *Class S* communication pattern outperforms *Class A* both in terms of throughput and energy-efficiency.

These results unveil the need for a mechanism allowing to switch between communication patterns. Indeed, being able to alternate between asynchronous and synchronous schemes depending on the offered load would increase the overall network capacity while always employing the most energy efficient strategy. Such a mechanism has therefore been developed and is presented in the next section.

3.3 TREMA: a traffic-aware MAC protocol that maximizes energy efficiency in LoRa networks

In the previous section, we have introduced *Class S* as a means to enable a slotted access for uplink communications in LoRa networks. Preliminary simulation results revealed that further than improving the maximum achievable throughput, a Slotted ALOHA scheme could also be more energy efficient than the legacy *Class A* under certain traffic loads if a beacon-skipping mechanism was implemented. We concluded that a mechanism to dynamically adapt the access scheme to the traffic load was desirable in order to optimize the energy efficiency of LoRa deployments in all situations. Indeed traffic flows in such networks can be variable, or even unpredictable, depending on the needs of the monitoring applications. As an example, one of the many LoRaWAN applications is pollution tracking. Interestingly, the

air quality in a city varies much more during the daytime than during the night [Jalava 2015]. Efficient monitoring of this kind of phenomena would thus require a time-varying number of measures. In the considered example scenario, switching between the asynchronous pure ALOHA access scheme and a time-synchronized one would permit a higher data reporting rate during daytime, and conversely energy savings during night hours. In such a context, this chapter introduces the Traffic-aware energy efficient Medium Access Control (TREMA) protocol for LoRa networks, capable of seamlessly switching between asynchronous and synchronous schemes according to the probed traffic variations. TREMA ultimately increases the maximum capacity of LoRa deployments while always selecting the most energy efficient access scheme.

Remarkably, configuring TREMA requires the deployment fingerprint, defined as the characterization of several network metrics as a function of the traffic load. A prior performance evaluation has therefore been carried-out for both the asynchronous and synchronous access schemes, and then leveraged by TREMA to maximize energy efficiency in all situations. In more details, the design of TREMA includes (i) a *time-synchronized scheduled access* to be used when the traffic load is high, (ii) a detailed description of the required deployment *fingerprint*, (iii) a *probing strategy* that aims at estimating the generated traffic, (iv) a *decision mechanism* that determines whether or not a switching should be triggered, and (v) the *signaling protocol* required to adapt the access scheme of all devices from asynchronous to synchronous and *vice versa*.

The performance of such a protocol is assessed through simulations. Then, TREMA is tested in an example scenario to show how the network capacity is adapted in reaction to traffic load variations. A final performance assessment has been pursued to feature its behavior under any traffic condition. To the best of the authors' knowledge, TREMA is the first protocol aiming at dynamically synchronizing and desynchronizing LoRa deployments according to the traffic load variations, thus adapting the network capacity and always maximizing energy efficiency.

3.3.1 Design of TREMA

The goal of TREMA is to adapt the LoRaWAN capacity to the traffic load while always maximizing energy efficiency. Herein, a time-synchronized scheduled access is presented to increase the achievable throughput. It provides a good energy efficiency when the generated traffic is high. On the other hand, the legacy pure ALOHA asynchronous scheme is more efficient when the generated traffic is low. TREMA thus implements a mechanism capable of seamlessly switching between the two aforementioned schemes depending on the situation. To achieve this goal, the radio medium is probed to estimate the current traffic load. The network server then uses this estimation to select the most appropriate access scheme. To do so, this decision mechanism refers to a pre-established deployment fingerprint that relates the network behavior under different traffic loads. When the access scheme needs to be changed, the signaling protocol embedded in TREMA is leveraged to




Figure 3.6: Switching mechanism

transmit switching commands to the end devices. This behavior is schematized in Figure 3.6.

3.3.1.1 A time-synchronized scheduled access over Class S

The proposed time-synchronized resource scheduling protocol is built upon *Class S*, the slotting scheme introduced earlier in Chapter 3. *Class S* takes advantage of the *Class B* synchronization beacons broadcast by the gateways to define transmission timeslots. These slots fit the longest frame time-on-air of the considered deployment to ensure that collisions may only occur between frames transmitted within the same slot. The highest priority in LoRa networks is to maximize the lifetime of the device batteries. The transmission time must therefore be minimized in order to reduce the energy required to send a frame and the collision probability. This is achieved by using the highest LoRaWAN data rate available, DR5, which corresponds to the smaller SF. Results can however safely be extended to the other orthogonal SFs by following a similar reasoning to further increase the capacity. The *Class S* slot size is identical to the one computed in section 3.2.1.2.

Given the focus on delay tolerant applications that require only unconfirmed uplink traffic in a single-gateway deployment, a simple resource scheduling scheme has been designed to maximize the achievable network throughput. Specifically, all devices are assigned to a specific slot within the slot frame based on their network joining index. This guarantees that all devices are evenly spread out between all available slots. Each device is thus competing with less neighbors, but for a shorter duration. The frame collision probability is therefore lower than when using slotted ALOHA if all devices are active and generating frames at similar rates. A round-robin channel hopping is used to prevent a possible channel fading from making a given device unreachable. This guarantees that the frames generated by a device may collide with frames from a maximum number of $\lceil \text{NumDevices} / (\text{NumSlots} \cdot \text{NumChannels}) \rceil$ other devices.

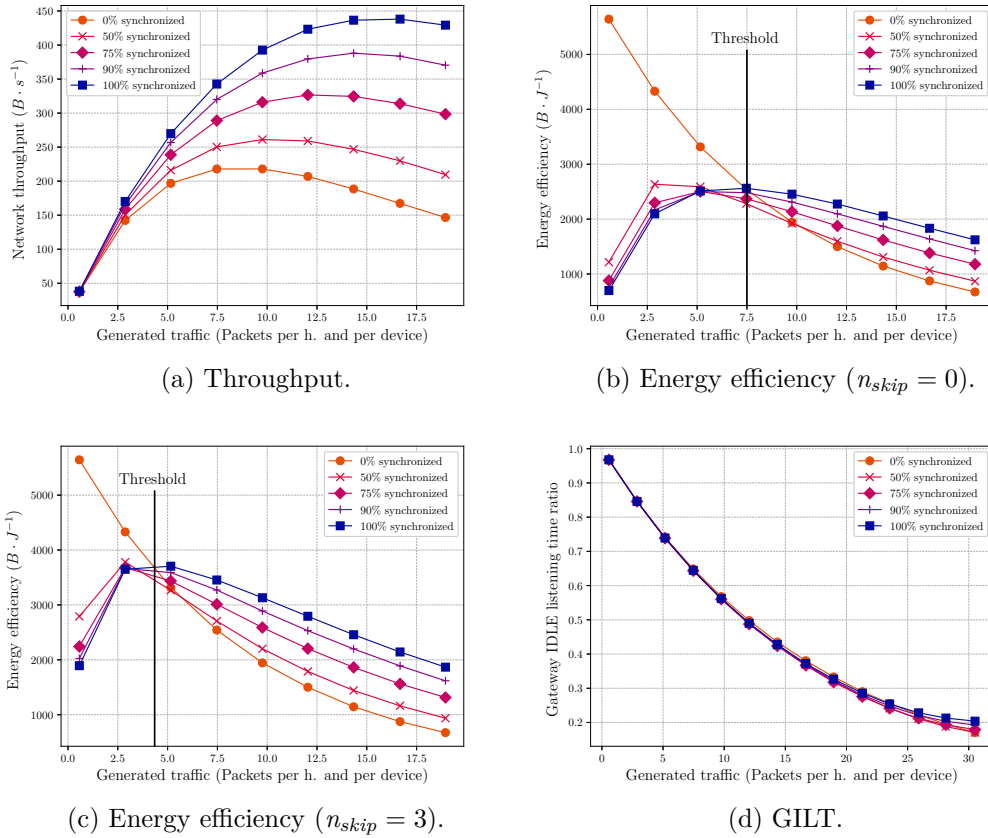


Figure 3.7: Fingerprint of the simulated deployment as a function of the generated traffic

In the given scenario, an uplink message is sent by a device assigned to the second slot. As specified in chapter 3, the last slot (which corresponds to the 187th when using 660 ms-wide slots) overlaps onto the BEACON_GUARD period. The uplink transmission is centered within the slot to prevent a possible clock drift from letting it overlap onto another slot.

3.3.1.2 Prior fingerprinting of the deployment

In order to wisely select the most energy efficient access scheme for a given network load, it is beforehand necessary to study the deployment behavior for each scheme depending on the frame generation rate g . This analysis ultimately allows to derive g from the probed traffic information, and provides a threshold value for g which allows to decide which access scheme should be chosen. The use of this deployment fingerprint in TREMA's switching mechanism is depicted in the flowchart of Figure 3.8. The **throughput** T , expressed in bytes per second, represents the amount of data successfully transmitted during a given time interval. It must be assessed in order to understand the maximum capacity of the network. The **Gateway Idle**

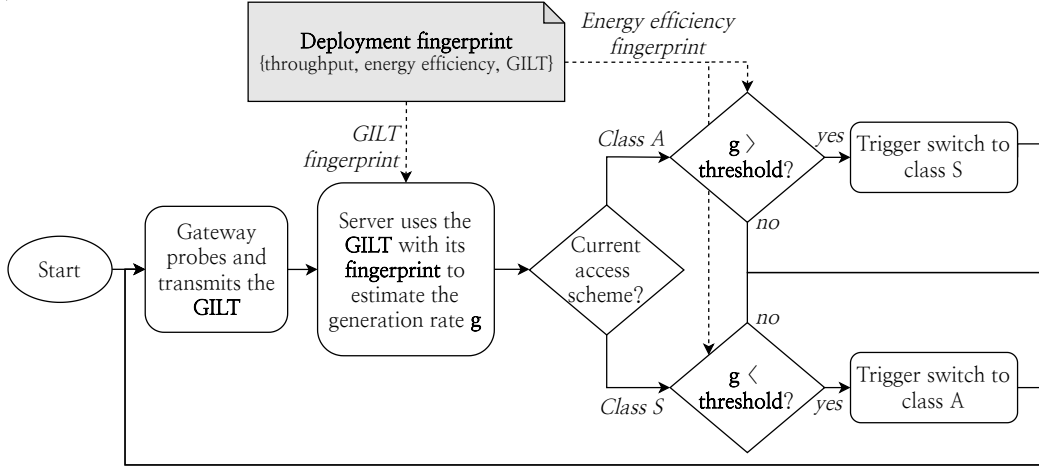


Figure 3.8: TREMA’s switching mechanism flowchart

Listening Time (GILT) metric is defined as the average time when the gateway is not receiving any packet, computed over all uplink channels. It plays a role in the traffic load estimation, which is detailed hereafter in Section 3.3.1.3. The **energy efficiency** E [Björnson 2018] is obtained by dividing the throughput T by the average power consumption. In order to increase E , devices may skip some of the synchronization beacons. Indeed, the LoRaWAN specification states that devices should be able to maintain beacon-less operation during at least two hours if the RX slots are widened according to the worst-case drift time since the last synchronization. Hereafter, this upper bound to the allowed drift is indicated with $drift_{max}$, while the parameter n_{skip} is introduced to represent the number of beacons skipped by devices. For instance, if $n_{skip} = 1$, the device will listen to one beacon out of two. The relationship between the maximum time-on-air within the network and the chosen *Class S* slot size determines $drift_{max}$, which in turn bounds the maximum value for n_{skip} . As we will show later on in Section 3.3.2.1, the fingerprint of E allows to determine a traffic load tipping point, under which pure ALOHA is the most energy efficient access scheme, and above which the synchronized access becomes preferable. This tipping point defines the frame generation rate **threshold** guiding TREMA’s decision mechanism.

3.3.1.3 Network probing and decision mechanism

TREMA’s probing strategy aims at estimating the generated traffic thanks to a measurable value. The probed metric must respect two conditions: (i) the generated traffic should be deductible from its value, and (ii) all gateways must be able to compute it easily. The GILT is an excellent choice for this task, as it checks both requirements. Condition (i) is respected by any metric that proves to be a bijective function of the considered generation rate range. Indeed, the Gateway Idle Listening Time (GILT) strictly decreases when the traffic load increases. The GILT also

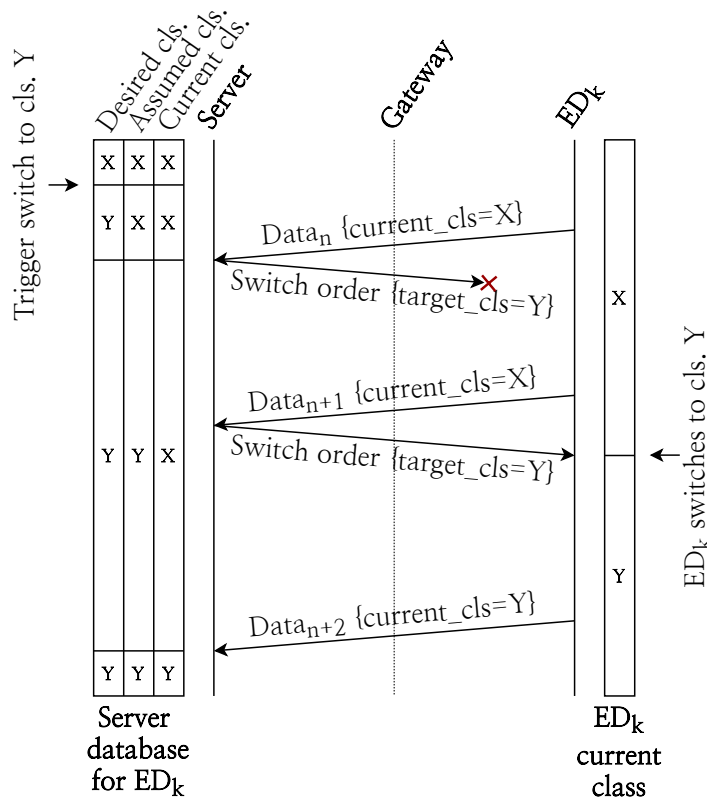


Figure 3.9: General-case packet exchange for class switching.

matches (ii), as all gateways can keep track of their idle time. This is why this metric was chosen over the collision time for instance, that checks (i) as well, but would require discriminating the successes from the failures among the total reception duration. It is interesting to note that the introduction of the GILT within the fingerprint was necessary, as the T matches only (i), while E complies only with (ii). The gateway keeps track of the GILT and transmits it to the server periodically. Upon reception of the probed GILT, the server uses it to deduce an estimation of the offered traffic with the GILT fingerprint. Such an estimate is checked against the threshold value to decide whether the access scheme should be adapted following the reasoning depicted in Figure 3.8.

3.3.1.4 Signaling protocol

It will be shown in Section 3.3.2.1 that the transient states in which access schemes coexist are sub-optimal in terms of energy efficiency. This contribution will therefore focus on switching between 100% *Class A* and 100% *Class S* scenarios. The switching command requires the definition of a `SwitchCommand` MAC command that will be transmitted through the `FOpts` field within the MAC payload [LoRa Alliance 2020a]. As there are only two possible access modes for now, one

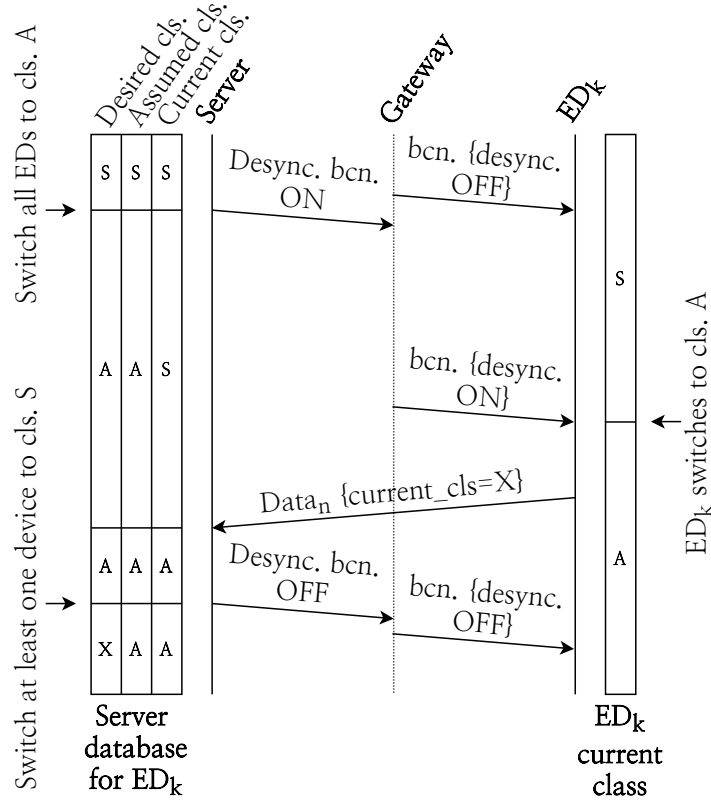


Figure 3.10: Desynchronization beacon broadcasting.

Figure 3.11: TREMA’s switching mechanism signaling protocol.

bit will suffice to transmit the command value. The server keeps track of three class-related variables for each registered device. **DESIRED_CLASS** is updated when a class switching event is triggered and contains the target access scheme for this device. **CURRENT_CLASS** keeps track of the class the considered device was using when it sent the last uplink frame. **ASSUMED_CLASS** is updated when a switching command has been sent, but with no guarantee that it has been correctly received. This last one is used to provide an optimistic estimation of the current class division within the deployment, which is necessary to estimate the metrics more accurately when using the fingerprint as will be shown in Section 3.3.2.1. Indeed, when a device receives a switching order, it does not send an acknowledgement immediately but waits the next uplink frame to inform the server that it has switched. This restricts the traffic induced by the mechanism, at the cost of adding some delay to the server database updates. Fig 3.9 holds a switching command transmission chronogram. First, a switch from any class X to any class Y is triggered, so the

server updates the DESIRED_CLASS entry. With no prior assumption on the node's current class, a RX window has to be used to transmit the order. The server thus waits for an uplink transmission $Data_n$ to send the order, which is done every time the values in CURRENT_CLASS and DESIRED_CLASS differ. When doing so, it updates ASSUMED_CLASS, but then the downlink frame is lost. At the next uplink message $Data_{n+1}$, the current class of the device is still X so another switching order is transmitted, and this one arrives. The device updates its access mode, and when finally it sends $Data_{n+2}$ CURRENT_CLASS is updated. When switching from *Class A* to *Class S*, the *Class A* RX windows are used to transmit the switching commands. The server must therefore wait for an uplink transmission from a given device before triggering its class switching. This signaling induces a lot of downlink traffic and may take some time to converge depending on the device transmission rate, but it is expected to be more energy efficient than forcing devices to track beacons to determine what access scheme they should be using. This aspect will be proved in future contributions. The switching process can however be sped up in the particular case where devices are switched from *Class S* to *Class A*. Indeed, all synchronized devices receive beacons periodically, so using them to broadcast a desynchronization command allows the process to converge much faster. This process is schematized in Figure 3.10. When the switching event is triggered, the server informs the gateway that the desynchronization beacon flag should be set. At the next beacon reception, the device updates its class, and at its next data frame it informs the server that it has switched. Then, a switching to another class distribution is triggered and the desynchronization beacon flag is disabled.

3.3.2 Performance evaluation

In order to assess the goodness of the TREMA design, an ad-hoc event-based simulator has been used. To kick off TREMA, the deployment fingerprint was established beforehand through large simulation campaigns using the considered network parameters. This allowed to quantify the achieved performance gain. After that, an example scenario in which the frame generation rate varies over time has been used to feature the online TREMA behavior.

The simulation environment and the considered network that served as a basis for these results will now be presented. First of all, it is important to consider that IoT traffic cannot be simulated with the same tools as standard Internet traffic. From this observation, the Third Generation Partnership Project (3GPP) proposed the first model to reproduce it accurately [3GPP 2012], and subsequent works have proposed improvements [Laner 2013] and comparisons [Smiljkovic 2014],[Zarrini 2018]. It may follow three typical shapes: (i) Periodic Update, mostly used for monitoring purposes, (ii) Event Driven, where traffic may be randomly triggered by external events (*e.g.* emergency detection systems), and (iii) Payload Exchange, when transmitting large files such as firmware updates [Nikaein 2013]. The aggregated outcome may be modeled by a combination of these patterns, and adapted to the applications using the network [Farooq 2018].

All in all, this shows that such traffic is subject to important variations that may be spatially and/or temporally correlated. In the scope of this paper that mainly focuses on the design and testing of the switching mechanism, the frame-arrival process is modeled as a Poisson process with a variable rate to represent the fluctuations of the aggregated load. The network features have been set to resemble a large scale single-gateway deployment of 1000 devices. All frames are considered to be featured with the maximum time-on-air allowed for DR5, 626.94 ms (cf. Section 3.3.1.1). All devices use homogeneous rates in each scenario, which is a reasonable simplification as the fingerprint is established as a function of the overall traffic. Herein, the simulator frame-arrival process is modeled as a Poisson process with a variable rate to represent the fluctuations of the aggregated load. These rates vary from ~ 0.5 to ~ 20 packets per hour. The traffic is restrained to uplink data communications only, thus the downlink data frames have been disabled along with acknowledgments and re-transmissions. In this way, the downlink capabilities are left free to be used without interference by the switching mechanism. Impact of downlink traffic on switching delays and overall performances will be tackled in future works. Only the three mandatory LoRaWAN channels are implemented, which are subject to a 1% Duty Cycle.

Table 3.3: Simulation parameters

Parameter	Value
Simulation duration	24 hours
Frame time-on-air	626.94 ms
Beacon time-on-air	173.06 ms
Timeslot size	660 ms
Data rate	DR5 (SF7 with bandwidth 125 kHz)
Frame generation rate	Varies from ~ 0.5 to ~ 19 pkts. / h.
Channels	868.1, 868.3 and 868.5 MHz
Number of devices	1000
Device buffer size	1
n_{skip}	0
Uplink Duty Cycle	1%
Downlink Duty Cycle	10%
Downlink messages	Switching commands only
Acks and retransmissions	Disabled
Sensor voltage	3.3 V
Sensor current intensity	20 mA TX, 10.8 mA RX [Semtech Corporation 2019]


3.3.2.1 Preliminary deployment fingerprinting

In order to assess which access scheme performs better for different traffic loads, the deployment fingerprint was established when using *Class A* and the scheduled access built over *Class S*. Different degrees of coexistence of the two modes were also assessed to evaluate the network behavior during the transient states of the switching. For the sake of statistical significance, for each configured scenario, 10 different realizations have been considered by feeding the pseudo-random number generator with 10 different seeds. T has been plotted in Figure 3.7a, for different percentages of synchronized devices. The figure shows that for this deployment, *Class A* is capable of offering a maximum of $220 B \cdot s^{-1}$, while the scheduled access peaks at $440 B \cdot s^{-1}$. It therefore provides a 100% capacity increase, which is better than the 88% increase offered by slotted ALOHA over *Class S* in Chapter 3.


Another notable observation is that the advantage of the scheduled access is evident when a significant portion of the devices are synchronized. Indeed, a capacity increase of solely 18% is observed when 50% of devices operate over *Class S*, and with 90% of synchronized devices the curve peaks at $390 B \cdot s^{-1}$ which corresponds to a relative improvement of simply 77%. When using *Class S*, the chosen slot size sets the value of $drift_{max}$, and therefore an upper bound for n_{skip} , because any transmission overlapping onto another slot should be strictly prohibited. A typical low-cost LoRa device crystal has a 30 ppm quality, which means that in the worst-case scenario, it may drift of 1 ms every 33.3 s. With our parameters, the transmission is framed by a margin of $drift_{max} = (660 - 626.94)/2 = 16.53$ ms within its slot. It would take a minimum of 550.5 seconds to get a similar drift, so the maximum number of synchronization-less beacon periods that a device can maintain is $\lfloor 550.5/BEACON_PERIOD \rfloor = 4$, which bounds n_{skip} to 3. A higher energy efficiency may be achieved by widening the slots to enlarge the margin, but this would have the effect to reduce the maximum achievable throughput. Once again, designing the deployment requires trading-off network capacity for battery life. From these considerations, E has been plotted for the minimum and maximum values of n_{skip} , respectively 0 and 3, in Figure 3.7b and 3.7c. When using the maximum value, the RX slots opened for beacon receptions have been widened by w , defined as:

$$w = 2 \cdot n_{skip} \cdot BEACON_PERIOD \cdot 3 \cdot 10^{-5} \quad (3.6)$$


With $3 \cdot 10^{-5}$ representing the 30 ppm crystal quality, $n_{skip} \cdot BEACON_PERIOD$ the synchronization-less time induced by the skipping, and the multiplication by 2 accounts for a possible positive or negative drift. TREMA aims at maximizing the energy efficiency at all times. A first remark that could be made on Figure 3.7 is that the maximum value for E is always attained with either 0% or 100% of synchronized devices. Therefore, the coexistence of the two access schemes should remain transitory. The point at which the 0% and 100% curves cross defines the frame generation rate threshold over which the scheduled access becomes more energy efficient than pure ALOHA. This is the tipping point used by TREMA to



(a) Generation rate variation and evolution of the number of devices in each class, with TREMA enabled.



(b) Evolution of the throughput over time, with and without TREMA.



(c) Evolution of the energy efficiency over time, with and without TREMA.

Figure 3.12: TREMA testing with one example scenario ($n_{skip} = 0$).

decide whether the network should be synchronized or not. An interesting impact of setting n_{skip} to 3 is that it shifts the threshold to the left. Indeed, without beacon skipping, the scheduled scheme becomes more efficient than pure ALOHA when generating 7.5 pkt/h per device, but when receiving only 1 beacon out of 4 the threshold is reduced to 4.5 pkt/h per device. This allows to profit from the gain in network capacity induced by synchronization at a lower rate, and therefore to ameliorate the frame delivery ratio between 4.5 and 7.5 pkt/h per device. Finally, the GILT fingerprint has been plotted in Figure 3.7d. With the specific parameters used in this deployment, curves for all class divisions overlap. However in the general case, the network class division is used to interpolate the appropriate traffic load from the closest class division curves. Further studies about the impact the network parameters on the GILT shall be led in future works.

3.3.2.2 Online testing

In order to test TREMA, a 24 hour scenario has been simulated to employ day and night frame generation rates. In the interval between midnight and 8 AM, a low rate of 1.5 pkt/h per device is used, lower than the switching threshold. Then,

the devices employ a higher rate of 10 pkt/h per device until 8 PM, higher than the tipping point, and finally go back to the night rate. No beacons are skipped, therefore the threshold is 7.5 pkt/h per device. The network behavior is evaluated through the plotting of T and E over time using a sliding window averaging over the last 15 minutes, with and without the switching mechanism. When enabled, the number of devices using each access mode is plotted over time as well. All these test results are pictured in Figure 3.12.

Without TREMA, the network always operates with *Class A*. When the rate increases, the throughput stabilizes at $220B \cdot s^{-1}$ and the energy efficiency at $2000 B \cdot J^{-1}$, which were indeed the values observed for the 0% synchronized fingerprint curves in figs. 3.7a and 3.7 for 10 pkt/h per device. But, when switching is enabled, the rate increase is automatically detected by the server that starts synchronizing the devices. The process is not instantaneous, and takes about 3 hours, as seen in Figure 3.12a. This transition time explains the tipping point observed at around 9 AM in the energy efficiency curve. Once again these values are perfectly consistent with the fingerprint data. This time, the throughput T remains stable at $400 B \cdot s^{-1}$, while the energy efficiency E sets up at $2500 B \cdot J^{-1}$. When the rate decreases at nightfall, the desynchronization beacon allows to switch the whole deployment back to *Class A* instantly, and the metrics return to their initial values. Two additional remarks can be made about the E plot. First, the energy efficiency drops faster when the rate increases if TREMA is enabled. This is due to the additional energy consumption induced by the reception of the switching commands. On the other hand, when the rate decreases, E increases a bit later when TREMA is enabled than when it is disabled. This is because the system takes some time to probe the traffic load reduction and trigger the switching to *Class A*.

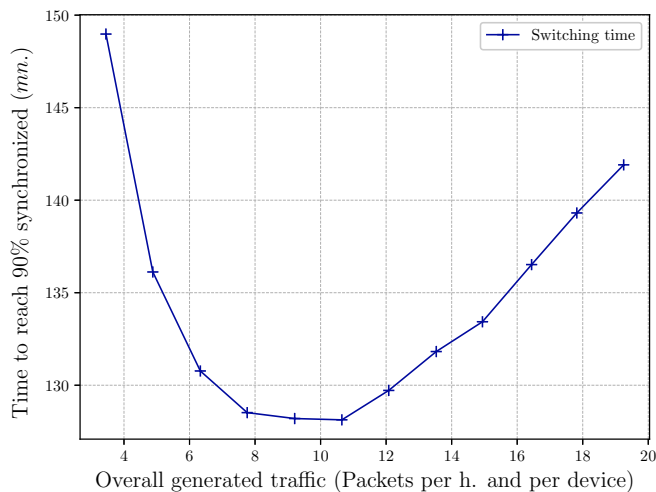


Figure 3.13: Switching duration.

3.3.2.3 Overall performance evaluation

This part aims at estimating the performance of TREMA for the deployment presented in Section 3.3.2.

3.3.2.4 Switching duration

First, the switching duration from *Class A* to *Class S* has been plotted in Figure 3.13. It is defined as the time between the triggering and the moment when 90% of the devices have changed their access mode. A first remark that can be made is that its value is always within the two and three hours range for this specific deployment. Another interesting fact is that it starts by decreasing between 0 and 9 pkt/h per device, because the increase of the *Class A* throughput reduces the average delay the server has to wait before sending a switch command to a device. The value of 9 pkt/h per device actually corresponds to the rate at which the maximum pure ALOHA throughput is attained in Figure 3.7a. It then increases again because the number of collisions causes the uplink throughput to decrease. The switching duration from *Class S* to *Class A* has not been plotted because it may easily be estimated thanks to the desynchronization beacon mechanism. In the worst case scenario, it is $(n_{skip} + 1) \cdot \text{BEACON_PERIOD}$. With $n_{skip} = 3$ it represents a duration of 512 seconds only, which is much shorter than the switching from *Class A* to *Class S*.




Figure 3.14: Increase in T and E as a function of the generated traffic ($n_{skip} = 3$).

3.3.2.5 Performance gain

In order to quantify the performance gain achieved when using TREMA compared to the legacy LoRaWAN access, the increase in terms of throughput and energy efficiency has been plotted in Figure 3.14 as a function of the generated traffic, when

using $n_{skip} = 3$. When the rate is below the 4.5 pkt/h per device threshold, the gains are 0% because TREMA selects the *Class A*. However, when the rate is above the threshold, the relative gains increase because the synchronized access performs better than pure ALOHA with these high traffic loads. The energy efficiency curve is continuous because the 100% and 0% synchronized curves cross in the E fingerprint, which is not the case for the throughput.

3.4 Conclusions and future works

In this chapter we have started from the theoretical pure and slotted ALOHA protocols to justify the need for synchronization in LoRaWAN. We have then introduced *Class S* to support the development of slotted schemes in LoRa networks. A slotted ALOHA scheme has been compared to the legacy protocol with large simulations campaigns, able to capture all the specificities of the LoRa technology. Notable results include (i) that a beacon-switching mechanism allows to achieve a satisfying energy-efficiency for the slotted access, and (ii) that pure ALOHA performs best at low traffic loads while slotted ALOHA is more energy-efficient in congested scenarios. The TREMA protocol was developed to dynamically adapt the LoRaWAN MAC layer to traffic load variations. More specifically, the legacy pure ALOHA is used when the generated traffic is low, and the access is switched to a time-synchronized scheduling of the transmissions when the network gets more congested. A threshold-based decision mechanism is used to seamlessly switch between the schemes according to the probed traffic load. Results show that this mechanism increases the maximum achievable throughput while always maximizing the device energy efficiency, by synchronizing and desynchronizing the devices.

One limitation of this preliminary approach is that it requires large simulation campaigns to establish the deployment fingerprint before implementing the mechanism. Future contributions will therefore evaluate the tools capable of establishing such fingerprints on the fly for any network parameters. Research directions include probability models and machine-learning based approaches. Further studies should also investigate the mitigation of network churn under unstable traffic conditions through hysteresis, assuring that a switching event is triggered only when observing a significant traffic load change for a sufficient amount of time, thus justifying the access scheme adaptation. In the next chapter, a real-hardware implementation of *Class S* will be performed in order to shed some light on the behavior of time-synchronized access schemes in real-world deployments, with the handling of device clock drifting.

Experimental Approach: Achieving Synchronization under real-world constraints

Dans le chapitre précédent, nous avons proposé une nouvelle classe de dispositifs LoRa pour mettre en place des stratégies d'accès synchrones. Nous avons jusque là supposé une synchronisation parfaite obtenue par le biais de la diffusion de balises par les stations de base. Dans ce chapitre, nous levons cette hypothèse en implémentant notre protocole synchrone sur un banc de tests réel. Dans ce but, nous réalisons une étude du décalage d'horloge éprouvé par de vraies puces LoRa, et définissons un mécanisme de synchronisation robuste aux imprécisions du matériel grâce à l'ajout de marges entre les transmissions. De plus, nous établissons un lien entre la longueur de ces marges et la périodicité de synchronisation de chacun des dispositifs LoRa afin de minimiser leur consommation d'énergie dédiée à cette tâche. Notre preuve de concept prouve que même du matériel à bas coût est capable de maintenir une synchronisation relativement précise. Les résultats issus du banc de tests révèlent additionally que l'effet capture se produit de manière consistante dans les réseaux réels, ce qui remet en question nos modèles d'évaluation de performances.

Contents

4.1	LoRaSync requirements and design	64
4.1.1	Device drift measurement and modeling	65
4.1.2	Spreading Factor selection	66
4.1.3	Payload size standardization	67
4.1.4	LoRaSync design	67
4.2	Proof of Concept on real hardware	69
4.2.1	Architectural challenges	70
4.2.2	Software challenges	70
4.2.3	Clock correction demonstration and discussion	72
4.2.4	Impact of the Capture Effect on the testbed performances	72
4.3	Conclusion and perspectives	73

In chapter 3 we have introduced *Class S* to enable slotted communications in LoRa networks. We have shown that such a slotted access could outperform the legacy pure ALOHA protocol in terms of throughput and energy efficiency in congested networks, and proposed a traffic-aware mechanism to seamlessly switch between the synchronous and asynchronous modes depending on the network load. Yet, we have assumed that a perfect synchronization was achieved through beacon broadcasts. In this chapter, we raise this hypothesis and introduce *LoRaSync*, an energy-efficient synchronization scheme that allows the implementation of *Class S* on real LoRa hardware. We measured the clock skew of typical LoRa hardware to evaluate how fast devices drift away from a given time reference. Based on this preliminary experimental evaluation, we designed *LoRaSync* with slots large enough to cope with clock errors. In that, the slot size and the maximum clock drift are used to compute the maximum amount of beacons that may be safely skipped while keeping devices synchronized. Such a beacon skipping mechanism is crucial to minimize the power consumption due to frame receptions, and eventually maximize the lifetime of feeding batteries.

In order to prove the soundness of the conceived *LoRaSync* mechanism, a real-world testbed has been setup to implement our synchronization scheme. This proof-of-concept demonstrates that the timing error can successfully be controlled despite the low-quality clocks found on such devices. In addition, this implementation allows us to highlight that the capture effect impacts the real-world throughput. We have therefore studied this phenomenon in more detail in Chapter 5.

From these premises, this chapter is organized as follows. Section 4.1 introduces the *LoRaSync* design, based on real clock skew measurements. Then, the preliminary evaluation of *LoRaSync* on our research testbed is presented in Section 4.2 through the description of a proof-of-concept implementation. Finally, Section 4.3 discloses concluding remarks and research perspectives.

4.1 LoRaSync requirements and design

With *LoRaSync*, we propose an energy-efficient and scalable synchronization approach tailored for LoRa deployments. It is designed to support the development of time-slotted access schemes for such networks, with an effort to minimize the device power consumption. This Section first provides an experimental evaluation of the clock drift experienced with low-cost LoRa devices. Second, we detail the reasons why we recommend to parameterize all transmissions with SF7, and why we concentrate on using the largest payload length allowed in the scope of this chapter. Based on this knowledge, we then present the *LoRaSync* design which includes (i) the synchronization strategy, (ii) a slotframe structure resilient to the device clock skew and (iii) a beacon-skipping mechanism designed to save battery power.




Figure 4.1: Device drift measurements

4.1.1 Device drift measurement and modeling

The *LoRaSync* synchronization mechanism aims at efficiently correcting the natural clock drift of cheap wireless devices to allow slotted uplink communications. Therefore, it is necessary to understand how this drift can be modeled before tackling the *LoRaSync* design. To do so, we use the testbed presented in Section 4.2, and program LoPy devices¹ to periodically transmit their local time. By comparing the intervals between the subsequent transmissions as seen by the gateway (which benefits from a very accurate GPS time), and the estimation of the same intervals done by the devices, we are able to evaluate the variations of the timing error. This clock skew has been plot in Figure 4.1 for two devices.

In most related works found in the literature [Beltramelli 2021, Polonelli 2019, Tessaro 2018], the clock drift is modeled as a linear function with the assumption that the sensor temperature is constant. This is referred to as the Simple Skew Model (SKM) [Veitch 2004], and it is verified in our experiment as we can see on Figure 4.1 that the device drift follows a linear trend on the long term. For cheap devices equipped with low-quality clock crystals, a 20 parts per million (ppm) worst-case drift coefficient is typically assumed.

However, it should also be noted that the clock skew displays a substantial noise around this overall trend, which must be estimated and is accounted for when designing the *LoRaSync* margins. This jitter has been mentioned in [Mallat 2014], but the authors chose not to include it in their modeling for simplicity reasons. Such a noise is indeed rarely considered in practice, even though for such cheap hardware it can result in significant errors for short time spans. As a result, we define a clock

¹Pycom website: <https://pycom.io/product/lopy4/>

drift interval model in which d is the drift coefficient in ppm, and ν the maximum deviation from the linear model due to noise. Therefore, we are assured that a device featured with such parameters which hasn't been synchronized for a duration of t will experience a drift comprised within the interval $(d \cdot t) \pm \nu$. The value of ν has been experimentally obtained by observing the maximum deviation from the linear trend found in the samples collected by the two devices and rounding it to the higher value. Following this methodology, we find $\nu = 11$ milliseconds. Such a model has been plot next to the real drift on Figure 4.1 with d adapted to the drift coefficient of each of the two devices, respectively -1.5 and 3.6 ppm. This preliminary drift evaluation will later be used to fine tune the *LoRaSync* parameters. Besides, we note that the worst-case 20 ppm drift hypothesis is checked for the two considered devices.

4.1.2 Spreading Factor selection

We have seen in Chapter 1 that the SF is a key parameter of the LoRa modulation, which has an impact on the range and ToA of frame transmissions. In this part, we justify our choice to recommend using only the smallest SF, SF7.

First of all, using the other SFs is detrimental in terms of energy efficiency. In detail, a higher SF results in a lower data rate and longer ToA, the only benefit being the improved transmission range. The consequence of having a longer ToA is that devices will have to use their radio longer to transmit the same amount of data, and will thus consume a greater amount of energy. Indeed we know that given SF the transmission Spreading Factor and BW its bandwidth, the symbol duration T_{Sym} can be computed with equation 1.1. As a result, the ToA magnitude grows proportionally to 2^{SF} , meaning that high SFs drastically affect the device's energy efficiency. We have designed LoRaSync on top of SF7, because it showcases the best performances in terms of both data rate and power consumption. Yet, the same designing logic proposed throughout this paper for SF7 can be applied to work with other SFs: for each additional SF, a new timeslot duration must be defined to accommodate the maximum frame size related to that SF. The resulting slotted structures (each on a separate SF) can herein coexist on a given channel.

SFs are presented as orthogonal by Semtech [Semtech Corporation 2015], meaning that transmissions occurring on the same channel but different SFs should be able to overlap without interference. However this claim has been put into question by the research community [Croce 2018]. Specifically, Croce *et al.* measure a 16 dB co-channel inter-SF rejection threshold. Such a RSSI difference is very likely to occur in near-far conditions, which are commonly experienced in large scale LoRaWANs, especially if high SFs featured with long transmission ranges and frame lengths are used. Multi-SF LoRa deployments can therefore not be studied as the simple superposition of independent networks. This research finding lets us prefer deploying multi-gateway single-SF LoRaWANs to a single-gateway multiple-SF LoRaWAN, if the goal is to cover a wider geographical area. Yet, this hypothesis is out of scope for this contribution and its validity and application scope will be

investigated in future works. When the use of several non-orthogonal SFs is still the preferred option, the availability of a slotted structure for each SF (of course, featured by a SF-specific timeslot size) would be beneficial to the definition of access schemes (ranging from slotted Aloha to collision-free scheduling) able to bound possible inter-SF collisions. This is also a research topic to be investigated in future works.

4.1.3 Payload size standardization

In this contribution, the timeslot length is derived from the maximum frame ToA allowed in the network, ToA_{\max} (*c.f.* Equation 4.3). This parameter depends on the SF and payload size. We consider the maximum payload size authorized by the LoRaWAN regional parameters [LoRa Alliance 2021] (*i.e.*, 250 bytes in Europe). LoRaSync can still be used with heterogeneous payload sizes, but optimal performances are obtained when every frame is featured with the maximum ToA.

The advantage of such a design in our situation is twofold. First, it allows us to use the available bandwidth more efficiently when a slotted access is utilized. Indeed, having frames with a length smaller than the maximum would create gaps between slotted transmissions, thus wasting a portion of the bandwidth. In the worst case, a slotted access could even showcase a lower throughput than Pure ALOHA. Conversely, an homogeneous frame length ensures the gaps in between frames are kept as small as possible, and guarantees that the bandwidth is used more efficiently.

Additionally, this hypothesis is also beneficial to the energy efficiency of devices. In fact when the maximum payload length is selected, the proportion of frame overhead compared to the amount of payload bytes is minimized. Therefore, a larger portion of the device energy is spent transmitting useful bytes, thus ameliorating the network energy efficiency from the application layer perspective. Such a requirement can be easily fulfilled by forcing devices to buffer their messages, and concatenate them into a single payload when the sufficient amount of data has to be accumulated. The detrimental aspect of such an approach is that the delay between frames increases because of such a buffering mechanism. However LoRa networks are typically not used for time-critical applications, so we consider this hypothesis to be reasonable.

In a nutshell, enforcing the transmission of frames with the maximum allowed size is beneficial to both extend the network capacity and foster device energy efficiency.

4.1.4 LoRaSync design

Building a slotted access requires signaling messages to be exchanged in the network to share a common time reference among all devices. To do so, *LoRaSync* exploits the LoRaWAN *Class B* beaconing mechanism. Therefore, this scheme can run on typical LoRa hardware with a few adjustments to the firmware (*c.f.* Section




Figure 4.2: LoRaSync beacon period and slot structure.

Table 4.1: Slotframe intervals

Interval	Length
<i>BEACON_PERIOD</i>	128 s.
<i>BEACON_RESERVED</i>	2.12 s.
<i>BEACON_WINDOW</i>	122.88 s.
<i>BEACON_GUARD</i>	3 s.

4.2.2). A synchronization event is defined as a tuple $(ts_{\text{bcn}}, cnt_{\text{bcn}})$, where ts_{bcn} is the beacon’s UTC timestamp and cnt_{bcn} is the device’s internal counter value at the beginning of the beacon reception. Such event is saved by the device at each beacon reception. UTC timestamps are defined as the total number of microseconds elapsed since the Unix epoch. Device counters are incremented every microsecond as well, therefore it is possible to add and subtract counter values to timestamps. Hence, the timestamp ts associated with any counter value cnt may be computed with the following equation:

$$ts(cnt) = ts_{\text{bcn}} + (cnt - cnt_{\text{bcn}}) \quad (4.1)$$

The precision of this estimator depends on both the internal clock quality and the time elapsed since the last synchronization event. Indeed, indicating the device’s worst-case clock drift with d (expressed in parts per million, ppm) and ν the noise margin, the worst-case timing error err may be computed for any counter value cnt with the following equation:

$$err(cnt, d, \nu) = (cnt - cnt_{\text{bcn}}) \cdot d + \nu \quad (4.2)$$

This worst-case error is notably used to enlarge the width of the reception slot each time a device needs to receive a beacon. This feature allows the radio to be in a

listening state when the beacon is sent, regardless of its current clock misalignment. *LoRaSync* implements *Class S* (c.f. Chapter 3), meaning that the beacon window is divided into timeslots dedicated to uplink transmissions. As a means to preserve bidirectional communications, *Class S* transmissions are followed by the RX1 and RX2 reception slots just like *Class A* ones. In order to adapt *Class S* to real hardware constraints, a maximum clock offset threshold δ_{\max} is defined. The slot size computation thus accounts for δ_{\max} and the maximum frame Time on Air (ToA) allowed:

$$L_{\text{slot}} = \text{ToA}_{\max} + 2 \cdot \delta_{\max} \quad (4.3)$$

With such a slot size, a device may not trigger a transmission that overlaps over two different slots as long as the current clock skew of this device is smaller than δ_{\max} . The associated slotframe layout within the *BEACON_PERIOD* is pictured in figure 4.2. It is divided in three intervals: *BEACON_RESERVED* which is dedicated to the beacon reception, *BEACON_WINDOW* in which slots are laid-out, with the last slot overlapping onto *BEACON_GUARD*. These intervals are identical to the ones used in the LoRaWAN *Class B* (c.f. Table 4.1) in order to facilitate the implementation of *Class S*. According to the scheme pictured so far, the number of slots in the slotframe is:

$$n_{\text{slots}} = \left\lceil \frac{\text{BEACON_WINDOW}}{L_{\text{slot}}} \right\rceil \quad (4.4)$$

The beacon-skipping mechanism has been implemented in order to ameliorate the energy efficiency of the protocol. In particular, n_{skip} is defined as the maximum number of beacons that can be skipped by a device before it needs to get synchronized again. Given the worst-case drift coefficient d , the noise margin ν and the maximum offset allowed δ_{\max} , n_{skip} is equal to the biggest $k \in \mathbb{N}$ that fits the following inequality:

$$(k + 1) \cdot \text{BEACON_PERIOD} \cdot d + \nu \leq \delta_{\max} \quad (4.5)$$

This means that increasing the maximum offset allowed (and thus the slot size) reduces the synchronization beacon periodicity, eventually decreasing the device power consumption induced by beacon receptions.

4.2 Proof of Concept on real hardware

LoRaSync was implemented on a small-scale research testbed. This proof-of-concept demonstrates that our mechanism is capable of bounding the synchronization error despite real hardware constraints. We additionally provide hands-on insights on how to handle the hardware and software challenges faced when synchronizing low-cost LoRa devices. low-cost LoRa devices. Finally, we provide a performance evaluation of the Pure and Slotted ALOHA access schemes implemented on the testbed with 10 devices, which reveal a discrepancy between the models and reality. In that, the

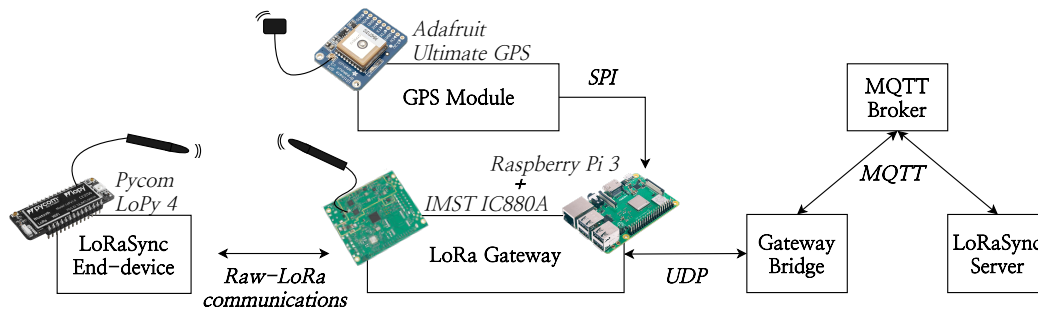


Figure 4.3: Experimental testbed infrastructure.

hardware implementation allows us to demonstrate the impact of capture effect on the channel throughput.

4.2.1 Architectural challenges

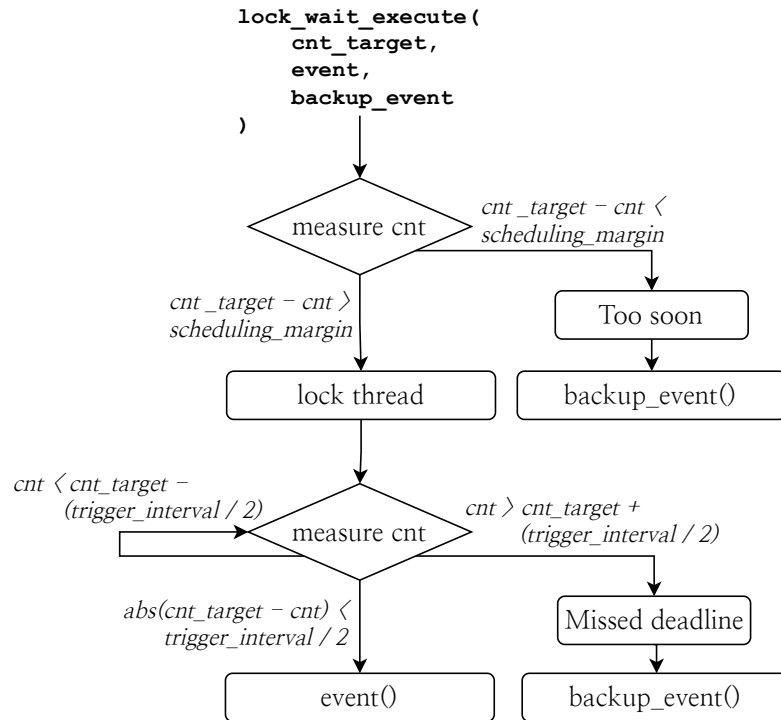
A LoRa testbed was set up according to the typical LPWAN architecture, *i.e.* a server, a gateway and devices, as depicted in Figure 4.3.

The server side mimicks the ChirpStack architecture, in which a MQTT broker centralizes the communications of all gateways. The ChirpStack gateway bridge allows to translate the MQTT streams into UDP datagrams understandable by the gateways. Finally, a custom *LoRaSync* server was developed to control and monitor the network with MQTT messages. The gateway is composed of a Raspberry Pi 3 running the Semtech Packet Forwarder software, used in conjunction with an IMST IC880A LoRa concentrator. The testbed devices are LoPy 4 chips developed by Pycom, offering quick prototyping capabilities for a relatively low cost.

By default, LoRa gateways are only capable of triggering downlink transmissions *(i)* immediately, or *(ii)* after a short delay following an uplink transmission (*Class A* RX windows). In order to allow timestamp-based downlink transmissions, an Adafruit Ultimate GPS chip was connected to the Raspberry Pi through Serial Peripheral Interface (SPI). It provides time and location information to the gateway program, and also sends a direct Pulse Per Second (PPS) signal to the concentrator enabling the triggering of transmissions with a 10 nanoseconds precision [Industries 2021]. This GPS chip requires a direct line-of-sight with the sky, as a result our gateway has been setup outdoors. This module ultimately allows the periodic broadcasting of the *LoRaSync* synchronization beacons.

4.2.2 Software challenges

This Section details the firmware adjustments that had to be made to implement *LoRaSync* on LoPy 4 chips. It must be pointed out that this hardware does not support *Class B*, and that a big portion of its behavior had first to be reproduced. The main challenge faced when programming *LoRaSync* on LoPy devices was to

Figure 4.4: `lock_wait_execute` flowchart.

schedule event executions with absolute timestamps, even though the embedded time library only allows callbacks after relative delays. To achieve this goal, the device logic was first split onto four threads, each assigned to a specific task: (i) message generation, (ii) frame reception, (iii) handling of received messages, and (iv) frame transmission. In this way, events requiring a precise timing such as frame transmissions and receptions may lock the CPU when needed, while the other tasks may proceed their execution freely during non-critical time.

Additionally, a software overlay was developed in order to handle critical method executions. Once synchronized, a node is able to estimate its current timestamp with its internal counter according to equation (4.1). When a timestamp event is scheduled, relative callbacks are used to trigger a call to the `lock_wait_execute` method, capable of handling time-critical events. This method is in charge of locking the hardware resources for the current thread, and waiting until the internal counter falls within a precise confidence interval around the target timestamp (*c.f.* Figure 4.4). It then triggers the event execution with a ~ 100 microseconds precision. In the case where resources could not be locked in time for the critical method call, a backup event passed as parameter is executed in order to correct the failure.

Once the device is synchronized, this triggering process allows to track beacons by turning the radio on just before the next broadcast. It also enables uplink timestamp-based transmissions, which ultimately allows to implement the *Class S* timeslots.

4.2.3 Clock correction demonstration and discussion




Figure 4.5: Device drift measurements

In order to demonstrate the good operation of *LoRaSync*, the clock drift of a periodically synchronized device is plot in Figure 4.5. In order to facilitate the experiments, the node has been set indoor, and is not in direct line of sight with our outdoors gateway (*c.f.* Section 4.2.1). It notably ensures that the temperature is relatively constant for the duration of the experiment. This single device is setup with $\delta_{\max} = 39.16$ ms, a worst-case drift coefficient $d = 20$ ppm and a noise margin $\nu = 11$ ms. This setting results in $n_{\text{skip}} = 10$, meaning that a beacon is received every $11 \times 128 = 1408$ seconds. Beacon receptions are represented on the graph by black vertical dashed lines.

Here it is clear that the device never crosses the maximum drift allowed line $\delta_{\max} = 39.16$. The margin here appears particularly wide because the actual drift coefficient of the chosen device is $d = -1.5$ ppm, which is very small compared to the worst-case estimation of 20 ppm. The drift model interval has been plot over the actual drift in order to highlight the synchronization events. It shows that in this particular case, the noise plays a bigger role in clock errors than the linear component of the drift.

All in all, this proof of concept shows that the synchronization error is kept below the worst-case drift threshold δ_{\max} . *LoRaSync* therefore operates as expected in a realistic context.

4.2.4 Impact of the Capture Effect on the testbed performances

Additionally to this proof of concept with one device, we have evaluated the performances of the testbed under Pure and Slotted ALOHA access schemes for various

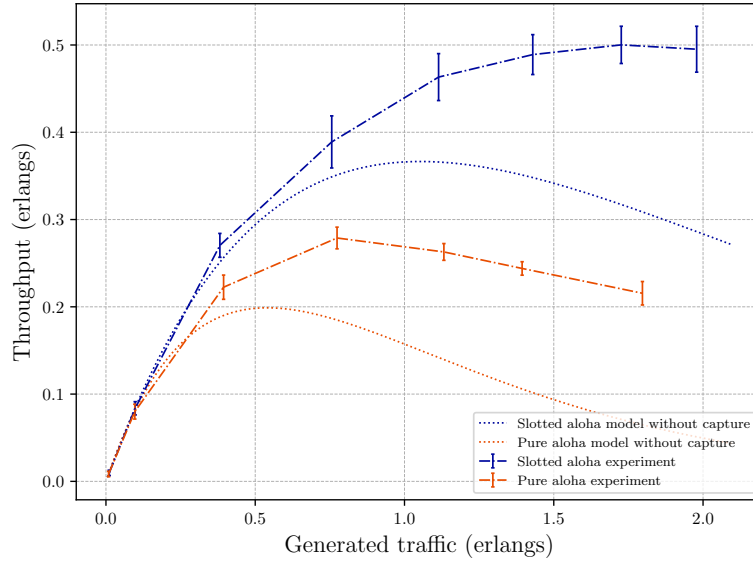


Figure 4.6: Experimental pure and slotted ALOHA throughput

generation rates with 10 devices. All devices are setup indoors at the same location. For the Slotted ALOHA scheme, we use *LoRaSync* parameterized with $\delta_{\max} = 39.16$, which results in the same parameters as in Section 4.2.3. For each generation rate the network runs for 30 minutes, and the throughput is sampled every minute. Errorbars represent 95% confidence intervals computed with Student's t law. Results are plotted in Figure 4.6, along with ALOHA models presented in chapter 3 for the sake of comparison. Here, it is obvious that the models underestimate the real-world throughput. As a result, we have thoroughly investigated this finding in the next chapter, which presents experimental models able to accurately picture the testbed throughput. One of topics explored in this complementary study is that the impact of capture effect on the network heavily depends on the deployment layout.

4.3 Conclusion and perspectives

In this chapter we introduced *LoRaSync*, a synchronization scheme designed for LoRa networks that aims at maximizing the device energy efficiency. Synchronization was used to setup a Slotted ALOHA access providing capacity improvements compared to the legacy Pure ALOHA scheme. *LoRaSync* slots may however be leveraged to implement any slotted MAC protocol. Synchronization is achieved through beacon broadcasts, and is therefore scalable to any network size. The *LoRaSync* slots contain margins to cope for the worst-case device clock drift, ensuring a reliable operation even on low-cost hardware. A beacon skipping mechanism is implemented so that devices only realign their internal clock when their current drift is likely to exceed the margin length. This way, the energy devices spend

listening to incoming beacons is minimized.

LoRaSync was implemented on an experimental testbed, showing its capability to successfully correct the timing errors and keep the clock drift below the maximum threshold allowed even on a real-hardware setup. This proof-of-concept also reveals the impact of capture effect in a realistic setup.

All in all, *LoRaSync* provides a robust and energy efficient basis to allow development of slotted schemes over LoRaWAN. It must be pointed out that the simple Slotted ALOHA scheme used in this paper to demonstrate the viability of the mechanism is not flawless. For instance, a severe competition occurs during the first timeslot due to traffic accumulating during the beacon reserved and guard intervals. In that, more advanced schemes such as scheduling techniques should be developed, and *LoRaSync* has been designed as robust cornerstone able to support these future improvements. Besides, the TREMA switching mechanism allowing to select synchronous or asynchronous access schemes depending on the traffic load should be implemented, and improved to also adapt the slot size. Finally, all these contributions should be extended to operate in multi-gateway scenarios as well.

In the next chapter, we will feed the main observations made on the testbed back into performance models as a means to provide accurate performance evaluation tools for our protocols. In details, we integrate all the specificities of LoRaSync-operated networks to the usual ALOHA theory, which allows us to derive the most energy efficient slot size for any offered load. As a second contribution, we account for the Capture Effect in our throughput modeling thanks to an experimental evaluation of capture occurrences.

Back to Theory: Accurate Modeling to optimize the protocol performances

Dans le chapitre précédent, nous avons montré par l'expérience que les modèles utilisés jusque là pour évaluer les performances réseau ne correspondent pas à la réalité. En conséquence, ce chapitre propose de mettre à jour notre modélisation mathématique avec tous les éléments observés sur le banc de tests. Tous d'abord, nous introduisons des modèles de débit et d'efficacité énergétique adaptés aux réseaux utilisant notre mécanisme de synchronisation. Cela permet ensuite de trouver la taille de slot optimale pour n'importe quelle charge de trafic. Ce résultat pourra être ainsi utilisé pour améliorer le mécanisme adaptatif présenté au chapitre 3 dans de futures contributions.

Dans la deuxième partie de ce chapitre, nous réalisons une étude expérimentale de l'effet capture afin de décrire le comportement observé sur le banc de tests. Plusieurs scénarios de transmission sont décrits et reproduits un grand nombre de fois afin d'estimer la probabilité de capture dans chaque situation. Les coefficients sont ensuite agrégés pour obtenir un modèle de débit décrivant fidèlement nos traces réseau. Nous montrons enfin que de futures contributions sont encore nécessaires pour généraliser cette modélisation à d'autres topologies de déploiements. En effet, des modèles d'une grande précision seront également nécessaires pour piloter le mécanisme de décision du protocole adaptatif présenté au chapitre 3.

Contents

5.1	Modeling and optimizing LoRaSync	77
5.1.1	Model validation and analysis	79
5.1.2	Slot size optimization	81
5.2	Modeling the Capture Effect	83
5.2.1	Background on the Capture Effect	83
5.2.2	Throughput modeling with capture	85
5.2.2.1	Transmission scenarios and their occurrence probabilities	86
5.2.2.2	Experimental setup and capture coefficient measurements	87
5.2.3	Experimental results and discussion	88

5.2.3.1	Model validation with experimental data	88
5.2.3.2	Impact on the network fairness	89
5.2.3.3	Discussion	91
5.3	Conclusions and perspectives	91

In the previous chapter, we have implemented Pure and Slotted ALOHA on real LoRa devices, leveraging the LoRaSync synchronization mechanism for Slotted ALOHA. In order to assess the performances of large-scale *LoRaSync*-enabled networks, this chapter first tackles the establishment of associated throughput and energy efficiency models. These models notably account for slot size variations, reception slot enlargements and the proposed beacon skipping mechanism. They additionally have been validated through the LoRaWAN-sim simulation environment. With these models, we evaluate the impact of the *LoRaSync* slot size on the overall energy efficiency of the network. Remarkably, the slot size has a twofold impact on the network behavior. On the one hand, increasing the slot length reduces the number of slots that fit into the slotframe, thus reducing the maximum achievable throughput. On the other hand, using large slots also allows devices to stay synchronized for longer time periods without receiving beacons, thus reducing their overall power consumption. As a consequence, we discovered that setting up the proper slot size value is functional to fit the best trade-off between throughput and power consumption. We therefore provide a methodology to determine the most energy efficient slot size for any traffic load.

A notable finding of the previous chapter is that our usual ALOHA models failed to accurately picture the real-world throughput. Indeed, these models [Rom 1991] are based on the hypothesis that two colliding frames are inevitably lost. However, recent findings [Bor 2016, Noreen 2018] prove that the capture effect occurs in such deployments. This physical layer phenomenon describes the possibility that when several frames are transmitted simultaneously on a radio medium, the one featured with the highest Received Signal Strength Indication (RSSI) at the receiver may successfully be demodulated. The second part of this chapter therefore provides new throughput models that picture the network performances more faithfully by taking the capture effect into account. Due to the unpredictable nature of the radio medium, the occurrence of capture events is subject to randomness. To quantify this random component and take it into account in our modeling, we combine the theoretical approach with testbed measurements. For this purpose, we first of all establish the theoretical occurrence probabilities of several transmission scenarios. Then, each scenario is reproduced for a large number of iterations on an experimental testbed in order to estimate the capture probability in such situation. All scenarios are finally aggregated, and the resulting experimental models prove to be consistent with real throughput measures gathered on the same testbed. In an effort assess present and future MAC protocols for LoRa, the Pure and Slotted ALOHA access schemes are both evaluated. Besides an overall increase of the maximum achievable throughput, the occurrence of capture events imply that some devices

benefit from a greater share of the total throughput than others. Indeed, transmitters featured with the highest average signal power at the receiver will prevail in most capture scenarios. Their frames therefore have more chances to be successfully demodulated than the average. In that, the network fairness is degraded. In an effort to quantify this phenomenon, Jain's fairness index [Jain 1991] has been computed with our experimental data. Results show that the fairness degrades as the traffic load increases, for both the Pure and Slotted ALOHA schemes.

The rest of this chapter is organized as follows. Section 5.1 presents the modeling of LoRaSync and shows how it can be leveraged to optimize the slot size. Then, section 5.2 presents our methodology to establish experimental models that accurately describe the occurrence of capture events on our testbed setup. Section 5.3 draws concluding remarks and discloses research perspectives.

5.1 Modeling and optimizing LoRaSync

This section presents the throughput, power consumption and energy efficiency models used to evaluate the performances of *LoRaSync*-operated networks. For the sake of comparison, both *Class A* (*i.e.* Pure ALOHA) and a Slotted ALOHA scheme running on top of *LoRaSync* will be assessed. It has been shown that bidirectional traffic drastically weighs on the throughput of LoRa networks [Pop 2017]. For this reason, transmissions do not require acknowledgments in the scope of this contribution. In the scope of this section we do not account for the Capture Effect, we therefore assume that overlapping frames are necessarily lost.

The most usual ALOHA throughput modeling with a finite number of devices, has been presented in section 3.1. In order to realize a similar modeling in *LoRaSync*-operated networks, we introduce q as the probability that exactly one device generates a frame for the duration of a slot. We once again leverage the exponential law:

$$q = p \left[X \leq \frac{L_{\text{slot}}}{ToA_{\text{max}}} \right] = 1 - e^{-\lambda \frac{L_{\text{slot}}}{ToA_{\text{max}}}} \quad (5.1)$$

Besides, the k_s coefficient is introduced as a means to represent the actual transmission time available per slotframe:

$$k_s = \frac{n_{\text{slots}} \cdot ToA_{\text{max}}}{BEACON_PERIOD} \quad (5.2)$$

Using once again the binomial coefficient to compute the probability that exactly one device generates a frame during the duration of a slot, the average slotted ALOHA throughput T_s (in erlangs) is equal to:

$$T_s = k_s n q (1 - q)^{n-1} \quad (5.3)$$

In order to model the energy efficiency of the devices, their power consumption (P) must first be computed. This is done by considering its value in the trans-

mission, reception and sleep states separately. We note P_{TX} , P_{RX} and P_{SLEEP} the power consumption of a typical SX1276 LoRa transceiver in transmission, reception and sleeping states respectively, when equipped with a 3.3V battery voltage [Semtech Corporation 2019]. It is computed with the transceiver's supply current in the considered state, respectively 20, 10.8 and $0.2 \cdot 10^{-3}$ mA. Our modeling accounts for the LoRaWAN 30 ms. reception slots following the uplink transmissions (*i.e.*, RX1 and RX2). For this purpose, we introduce the slot listening rate ρ_s :

$$\rho_s = \lambda \cdot \frac{2 \cdot 0.03}{ToA_{\text{max}}} \quad (5.4)$$

The overall network power consumption therefore is:

$$P_p = [\lambda P_{\text{TX}} + \rho_s P_{\text{RX}} + (1 - \lambda - \rho_s) P_{\text{SLEEP}}] \cdot n \quad (5.5)$$

For the *Class S* model we need to additionally consider the time spent in a receiving state for beacon reception purposes. We define the device beacon reception period T_{bcn} as:

$$T_{\text{bcn}} = \text{BEACON_PERIOD} \cdot (n_{\text{skip}} + 1) \quad (5.6)$$

The beacon time on air is noted ToA_{bcn} . With *LoRaSync*, the beacon reception window is enlarged to cope for the worst-case drift that can occur during T_{bcn} with a drift coefficient d and a noise margin ν . Indeed, the device transmission may be shifted by an offset ranging from $-d \cdot T_{\text{bcn}} - \nu$ to $d \cdot T_{\text{bcn}} + \nu$. At the end of the beacon reception, the device switches its radio to sleep mode, so that the elapsed listening time may range from ToA_{bcn} to $ToA_{\text{bcn}} + 2 \cdot (d \cdot T_{\text{bcn}} + \nu)$ depending on the transmission offset. We assume that the drift coefficient d of all devices is uniformly distributed in the interval $\pm d$, therefore the average time spent in listening mode is $d \cdot T_{\text{bcn}} + \nu$. As a result, the beacon listening rate ρ_b is:

$$\rho_b = \frac{ToA_{\text{bcn}} + d \cdot T_{\text{bcn}} + \nu}{T_{\text{bcn}}} \quad (5.7)$$

The power consumption model can then be expressed as such:

$$P_s = [(\rho_s + \rho_b) P_{\text{RX}} + \lambda P_{\text{TX}} + (1 - \rho_s - \rho_b - \lambda) P_{\text{SLEEP}}] \cdot n \quad (5.8)$$

The energy efficiency (E) is expressed in Bytes per Joule and can be obtained by dividing the throughput by the power consumption. The throughput must however be expressed in bytes per second first, which is achieved by the term $\frac{\text{bytes}_{\text{pkt}}}{ToA_{\text{pkt}}}$ in the equations below. The energy efficiency models for Pure and Slotted ALOHA, respectively E_p and E_s , are therefore:

$$E_p = \frac{T_p}{P_p} \cdot \frac{\text{bytes}_{\text{pkt}}}{ToA_{\text{pkt}}} \quad (5.9)$$

$$E_s = \frac{T_s}{P_s} \cdot \frac{\text{bytes}_{\text{pkt}}}{\text{ToA}_{\text{pkt}}} \quad (5.10)$$

5.1.1 Model validation and analysis

Table 5.1: Simulation parameters

Parameter	Value
Simulation duration	24 hours
Number of seeds	10
Frame time-on-air	389.376 ms
Beacon time-on-air	173.06 ms
δ_{\max}	2.56 to 53.76 ms
Data rate	DR5 (SF7 with bandwidth 125 kHz)
Frame generation rate	Varies from ~ 0.5 to ~ 2.5 pkts. / h.
Number of channels	1
Number of devices	2000
Device buffer size	1
n_{skip}	Fixed according to δ_{\max}
Uplink Duty Cycle	Disabled
Downlink data messages	Disabled
Acks and retransmissions	Disabled
Sensor voltage	3.3 V
Sensor current intensity	20 mA TX, 10.8 mA RX [Semtech Corporation 2019]

LoRaSync targets large-scale deployments that experience scalability issues with the legacy *Class A*. Given the difficulty to approach such a large scale scenario with a research testbed, the models presented above have been validated using the LoRaWAN-Sim simulation environment to instantiate a network of 2000 devices. This analysis is completed by a proof-of-concept implementation on a few devices (*c.f.* Section 4.2) to demonstrate the feasibility of our approach. The LoRaWAN *Class A* and a Slotted ALOHA access over *LoRaSync* have been replicated in the simulator. Additionally, LoRaWAN-Sim integrates a linear clock drift feature that allows us to replicate the *LoRaSync* clock correction. The SF has been set to 7 for the reasons stated in Section 4.1.2. All transmissions occur on a single channel, and are parameterized with a coding rate of 4/5 with explicit header and cyclic redundancy check enabled. As specified in Section 4.1.3, all frames carry the maximum payload size allowed by the LoRaWAN specification (*i.e.* 250 bytes) to minimize the amount of overhead transmitted by the devices, which results in a ToA of 389.376 ms. *LoRaSync* was set up with a worst-case drift coefficient $d = 20$ ppm, which is a typical value for low-quality crystals found in such cheap hardware. Throughout this analysis, error bars represent 99% confidence intervals along the x and y




Figure 5.1: Throughput model and simulation results

axes computed with Student’s t law. In order to facilitate the reproduction of our experiments, detailed simulation parameters are provided in Table 5.1.

The simulated throughput for Pure and Slotted ALOHA are compared to the models in Figure 5.1, and the energy efficiency equivalent is displayed in Figure 5.2. In all cases, the model curves match the simulation results, showing the consistency between the modeling and the protocol implementation on LoRaWAN-Sim. In each plot *Class A* is compared to *LoRaSync* for several slot sizes. This comparison notably shows that the models faithfully capture the impact of the transmission margin size on the overall performances, which will be relevant for the analysis of Section 5.1.2. Specifically, the chosen slot sizes correspond to the δ_{\max} values of 2.56, 12.8, 28.16 and 53.76 milliseconds. Since the worst-case drift coefficient d has been fixed to 20 ppm and that the simulator drift follows a linear trend without noise, the n_{skip} values can be very easily deduced. Indeed we know that any device is subject to drift of a maximum of $20 \cdot 10^{-6} \cdot 128000 = 2.56$ milliseconds in the duration of a single beacon window. Therefore, we have in this case $n_{\text{skip}} = (\lfloor \delta_{\max}/2.56 \rfloor - 1)$. Interestingly, δ_{\max} values lower than 2.56 milliseconds can not be allowed because it is the drift experienced between the smallest possible interval between two beacon receptions. The selected δ_{\max} then result in n_{skip} values of 0, 4, 10 and 20.

When focusing on the observed performances, Figure 5.1 also shows the good operation of the *LoRaSync* synchronization, since the slotted access allows to nearly double the maximum achievable throughput as expected. Regarding energy efficiency we find the result already presented in Chapter 3, which is that *Class A* performs better for low generation rates and the slotted access becomes preferable when traffic rate is higher than the threshold $\tau_{p/s1}$. With this network configuration, and if we consider 53.76 ms to be the upper bound to δ_{\max} , $\tau_{p/s1}$ is equal to




Figure 5.2: Energy efficiency model and simulation results

0.34 erlangs. This value strongly depends on the hypothesis that devices are featured with a 20 ppm. drift, which was checked with our devices as we saw in Figure 4.1. Using devices of higher quality will naturally increase the Slotted ALOHA efficiency and move the crossing point towards a lower rate. More interestingly, Figure 5.2 shows that the value of δ_{\max} , and therefore the slot size, has a strong impact on the network energy efficiency. Indeed, for a generated traffic range of 0.34 to 0.6 erlangs the maximum energy efficiency is attained when $\delta_{\max} = 53.76$ ms ($n_{\text{skip}} = 20$). Then the $\delta_{\max} = 28.16$ ms ($n_{\text{skip}} = 10$) curve is the best until 1.2 erlangs and finally $\delta_{\max} = 12.8$ ms ($n_{\text{skip}} = 4$) prevails above that point. These intervals are represented by the $\tau_{s1/s2}$ and $\tau_{s2/s3}$ thresholds in Figure 3.7. We remark that the curve associated with $\delta_{\max} = 2.56$ ms ($n_{\text{skip}} = 0$) is never optimal in terms of energy efficiency. This is consistent with the preliminary results presented in chapter 3 that already showed that the beacon-skipping mechanism was unavoidable when trying to optimize energy efficiency. We also notice that the optimal amount of skipped beacons decreases as the traffic load increases. In order to further explore this finding, the next part sheds some light on the relationship between slot size and energy efficiency.

5.1.2 Slot size optimization

This part aims at evaluating the impact of the *LoRaSync* slot size on the network performances. Indeed the slot margin length δ_{\max} has an effect on the maximum achievable throughput, but also on the synchronization periodicity which in turn alters the device power consumption. This makes it a relevant parameter to consider when optimizing energy efficiency. In order to represent a realistic scenario, the

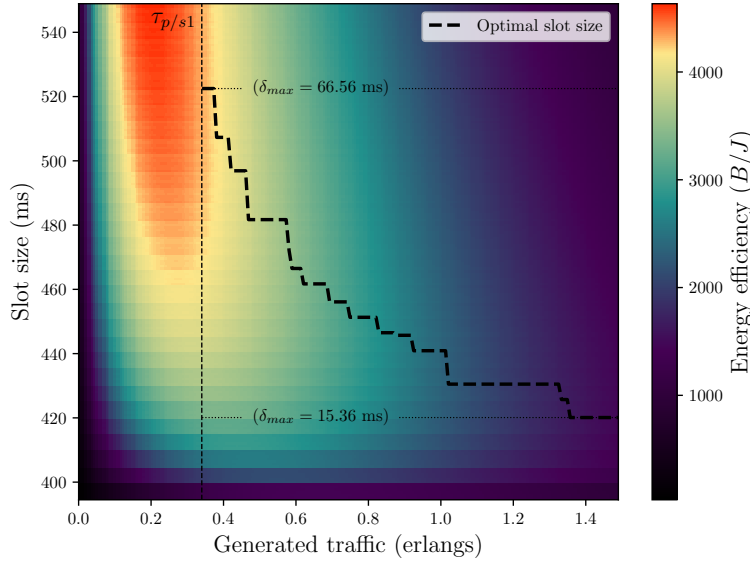


Figure 5.3: Energy efficiency as a function of the slot size (2000 devices)

models presented in the previous Section has been used to analyze the same large-scale deployment of 2000 devices.

From equation 4.3 we know that the slot size depends on δ_{\max} , which represents the maximum device drift allowed. Increasing δ_{\max} will result in larger slots, ultimately decreasing the number of slots per slotframe and the overall throughput. On the other hand the bigger δ_{\max} , the more devices are able to drift. This means that they are able to skip more beacons. Each additional beacon skipped reduces the overall device power consumption. Energy efficiency, defined as the ratio between the throughput and power consumption, thus strongly depends on the slot size. It has been plot as a function of the slot size and generated traffic in Figure 5.3.

At each rate, the slot size associated with the maximum energy efficiency is represented by the black dashed line. The line is only traced for generation rates bigger than the $\tau_{p/s1}$ threshold introduced in Figure 5.2. Indeed the asynchronous Pure ALOHA access should be preferred in terms of energy efficiency for rates lower than this threshold, it is therefore useless to consider the slot size there. This line shows that the optimal slot size decreases as the network load increases. Remarkably, with the same maximum ToA, smaller slots mean that devices are allowed to drift of a smaller amount. As a result a shorter synchronization period is required, which implies consuming more power in beacon receptions. Such short margins are therefore only suitable for high transmission rates. Conversely, wide slots allow to space out synchronization events, but reduce the number of slots available per slotframe. This allows to save power at the cost of a reduced maximum achievable throughput, which is not a problem when the traffic load is low. In a nutshell the more traffic there is, the more the margins should be reduced.

All in all, this result shows that the slot size should ideally be adapted to the

traffic load in order to maximize energy efficiency. We presented in Chapter 3 the TREMA mechanism to dynamically switch between Pure and Slotted ALOHA depending on the probed traffic conditions. Interestingly, this work could easily be extended to dynamically adapt the slot size to the generated load in order to further optimize the energy efficiency of the network.

5.2 Modeling the Capture Effect

This section presents our experimental models to faithfully picture the occurrences of the Capture Effect observed on our research testbed. It is organized as follows. First, section 5.2.1 provides background on the capture effect and its impact on LoRa deployments. Our experimental throughput models are introduced in Section 5.2.2. This modeling is then validated with testbed measures in Section 5.2.3. Insights on the impact of such effect on the network fairness and a discussion of our results are given as well.

5.2.1 Background on the Capture Effect

When several frames collide in a wireless network, not all of them are necessarily lost. Indeed, the RSSI difference and transmission offset between the involved frames may allow one transmission to be successfully demodulated by the receiver. This phenomenon is referred to as *capture effect*, and has been observed in several types of wireless networks. Semtech typically assumes a 6 dB co-channel rejection figure [Semtech Corporation 2019]. As a result, some related works [Georgiou 2017, Goursaud 2015] define a 6 dB RSSI difference between colliding frames as a requirement to benefit from the capture effect. However it was shown in [Croce 2018] and [Fernandes 2019] that capture events could be observed even with threshold values between a 0 and 5 dB. In [Elshabrawy 2018] this threshold is evaluated in greater details, and expressed as a function of the signal-to-noise ratio and spreading factor of the considered transmission. These findings therefore reveal that using a fixed 6 dB value leads to underestimating the actual throughput and coverage probability of LoRa deployments.

During a capture event, the demodulator behavior differs depending on whether the strongest packet arrives first or last at the receiver. We will now precisely describe these scenarios in order to approach related contributions with insightful details. The LoRa physical frame structure (as presented in the LoRaWAN specification [LoRa Alliance 2020a]) is recalled in Figure 1.3. In such a frame, the preamble allows the receiver to detect the presence of a packet, and the synchronization word is used to precisely align the demodulator with the arriving symbols [Daniel 2019]. After that, the physical header PHDR indicates the data size. The PHDR_CRC indicates whether a Cyclic Redundancy Check (CRC) will be used or not. PHYPayload contains the physical payload, and CRC carries the optional CRC value. In the event of a stronger-first capture scenario (*c.f.* Figure 5.4a), the demodulator synchronizes to the packet featured with the highest RSSI. Then,

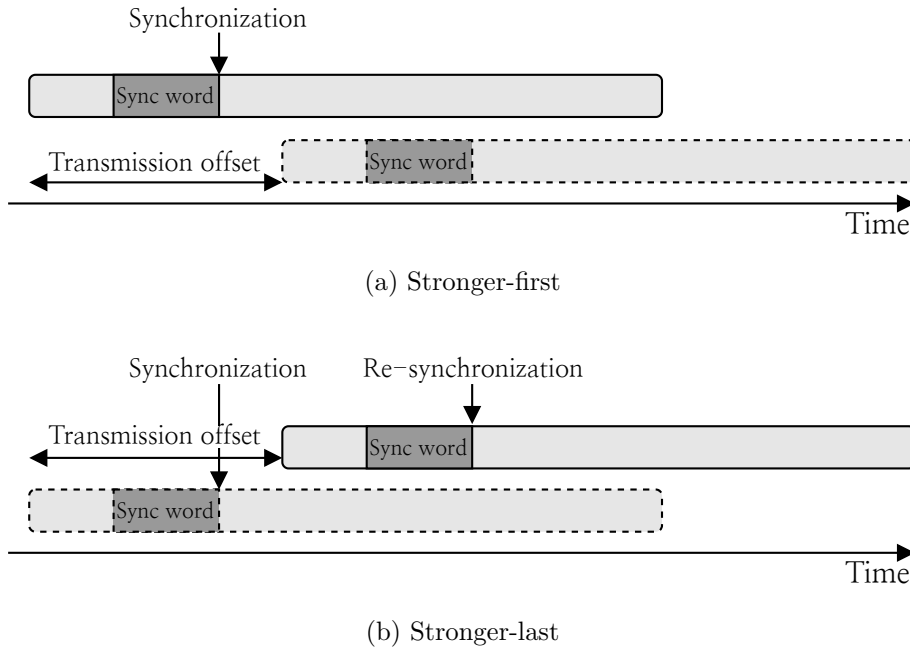


Figure 5.4: LoRa receiver synchronization and capture scenarios

it may retrieve its content without being affected by other frames if their relative power is sufficiently low to avoid interference. However when the frame featured with the lowest RSSI value arrives first (*c.f.* Figure 5.4b), the receiver starts its processing and then detects the preamble of the other one. At this point, it may be able to re-synchronize to the second frame and process it normally if its relative RSSI is sufficiently high.

These two scenarios have notably been described in [Whitehouse 2005], where the authors exploit the capture effect with a collision detection and recovery technique. They show that this method allows to retrieve information in the failing frame header for the stronger-last case. Such data can then be used to identify the terminal that lost a frame and setup a recovery procedure (*e.g.* trigger a re-transmission or adapt the device’s transmission parameters). In [Bor 2016], Bor *et al.* experimentally measured the probability of occurrence of capture events in a LoRa network. They notably demonstrated that the capture probability strongly depended on the transmission offset between the two concurring frames in all capture scenarios. This fact must therefore be considered when modeling the asynchronous pure ALOHA access scheme, that leads to frame offset variations. In [Noreen 2018], the CSS technique that characterizes the LoRa modulation has been mathematically modeled while accounting for the capture effect. Results show that a successive interference cancellation technique could also take advantage of the capture effect to decode even the weaker LoRa signal, thus further improving the network performances. This interference mitigation strategy has also been explored

in [Petroni 2022], which additionally proposes to combine information from all receiving gateways to achieve a more reliable decoding.

The impact of capture effect on the throughput of wireless networks is a well-known topic, and has been modeled in the context of Pure [Roberts 1975] and Slotted [Arnbak 1987] ALOHA access schemes. Interestingly, several recent contributions have studied the interactions between the capture effect and some characteristics of LoRa networks. In [Bankov 2017], Bankov *et al.* evaluate the Packet Error Rate with numerical results, and validate them with simulations. Contrary to us, they focus on LoRaWANs using acknowledgments and retransmissions. Besides, we validate our models with experimental data instead of simulations. The authors of [Sorensen 2019] propose another model that additionally considers the multiple demodulating paths available on typical LoRa gateways. Indeed most gateway concentrators are capable of demodulating up to 8 frames in parallel, which has an impact on the network capacity especially when several orthogonal Spreading Factors (SF) are used. The authors additionally tackle the problem of Spreading Factor allocation, and use simulations to validate their numerical results. Contrary to them, we focus on experiments with the smallest SF because it offers the highest data rate and lowest energy consumption. An experimental evaluation of the capture effect has also been led in [Fernandes 2019], in which the capture probability is assessed for a range of RSSI differences between several LoRa frames. Experimental data is then fed into a simulator to evaluate the overall network performances, but no model is derived. To the best of the author’s knowledge, the mathematical and experimental approaches have never been jointly applied to LoRa networks in the existing literature. This paper therefore provides new models that accurately describe the LoRa performances thanks to an experimental evaluation of the capture effect. We additionally use Jain’s index to quantify the fairness drop induced by capture events.

The modeling methodology used in this paper is in fact inspired of [Kosunalp 2015]. Indeed, Kosunalp *et al.* evaluated the capture effect in IEEE 802.15.4 networks (for Pure ALOHA only) by experimentally assessing several transmission scenarios and modeling the occurrence probability of each. We apply this idea to Pure and Slotted ALOHA access schemes, and extend it to account for single packet failures as well.

5.2.2 Throughput modeling with capture

The most common ALOHA throughput models [Rom 1991] are derived considering that two or more (even partially) overlapped frame transmissions do not allow the correct decoding of any frame. In a real-world deployment, these hypotheses are not necessarily met. For the case of a single ongoing transmission, unfavorable channel conditions may lead to a failure. This is emphasized by the fact that LoRa operates over unlicensed ISM bands, in which many concurring networks are potentially interfering with each other. On the other hand, when several transmissions overlap, the strongest one in terms of Received Signal Strength Indication (RSSI) may still

be demodulated by the gateway according to the capture effect.

In order to determine more accurate Pure and Slotted ALOHA throughput models, we conduct an experimental study inspired from the methodology presented in [Kosunalp 2015]. We consider the overlapping transmission of 1 up to 3 frames for Pure ALOHA, and up to 5 for Slotted ALOHA. By noting i the number of overlapping transmissions, we define the *occurrence probabilities* P_i (for Pure ALOHA) and P_i^* (for Slotted ALOHA) of such events (*c.f.* Section 5.2.2.1). These scenarios are then reproduced for a large number of iterations on the real hardware testbed in order to evaluate their respective *success coefficients* C_i and C_i^* , representing the average probability that a packet can successfully be demodulated in each scenario (*c.f.* Section 5.2.2.2).

The Pure ALOHA throughput T (in Erlangs) is naturally obtained by aggregating the success probabilities of all the three scenarios, weighted by their respective success coefficients:

$$T = \sum_{i=1}^3 (P_i C_i) \tag{5.11}$$

Setting up a Slotted ALOHA scheme requires a means to share a common time reference with all terminals. For this purpose, we leverage the LoRaSync mechanism introduced in Chapter 4 because it ensures an energy efficient synchronization while being robust to clock drift issues. With LoRaSync, margins are added to each slot based on the worst-case drift that may occur between two synchronization events. A portion of the time is also reserved for the reception of synchronization beacons. When computing the Slotted ALOHA throughput (in Erlangs), we must therefore consider the fraction of time available for transmissions α :

$$T^* = \alpha \cdot \sum_{i=1}^5 (P_i^* C_i^*) \tag{5.12}$$

5.2.2.1 Transmission scenarios and their occurrence probabilities

For this modeling, we make the hypothesis of a finite number of devices n generating a Poisson traffic with parameter λ . All frames are featured with the same Time on Air (ToA), which is used as the reference time unit. For Pure ALOHA, we know the probability p that any device generates a frame during such time unit from equation 3.1. We consider the scenarios where the transmissions of 1, 2 and 3 frames are overlapped. This access scheme operates smoothly for low traffic loads, where higher-order overlaps are very rare and can be neglected. For the single frame transmission scenario, we introduce the probability P_1 that one and exactly one device transmits a frame during a time unit. It is in fact equal to the classical Pure ALOHA throughput model presented in section 3.1, in which capture effect is not considered.

$$P_1 = np(1 - p)^{2(n-1)} \tag{5.13}$$

Table 5.3: Measured success coefficients

C_1	C_2	C_3	C_1^*	C_2^*	C_3^*	C_4^*	C_5^*
0.88	0.42	0.23	0.88	0.49	0.44	0.25	0.19

The probabilities that 2 and 3 packets overlap during a time unit, respectively P_2 and P_3 , have been computed by Kosunalp *et al.* [Kosunalp 2015], and their expressions are:

$$P_2 = n(n-1)p^2 \left(\frac{(1-p)^{2(n-2)}}{2} + (1-p)^{2n-3} \right) \quad (5.14)$$

$$P_3 = \frac{n(n-1)p^3(1-p)^{2(n-2)}}{2}(2n-3) \quad (5.15)$$

As mentioned before, LoRaSync uses margins to handle the device clock drift. For this reason, the LoRaSync slot size L_{slot} is bigger than the frame ToA. We therefore define p^* the probability that any device generates a frame during such a slot:

$$p^* = 1 - e^{-\lambda L_{\text{slot}}} \quad (5.16)$$

For Slotted ALOHA the transmission scenarios are easier to model and experimentally reproduce because frames always arrive simultaneously on the transmission medium. We therefore consider the cases of 1 to 5 simultaneous transmissions. The probability P_i^* that exactly i frames are generated for the duration of a LoRaSync slot can be expressed with the binomial coefficient:

$$P_i^* = P(i \text{ in } n) = \frac{n!}{i!(n-i)!} p^{*i} (1-p^*)^{n-i} \quad (5.17)$$

5.2.2.2 Experimental setup and capture coefficient measurements

In this part we detail all experiments used to measure the success coefficients C_i and C_i^* , which represent the probability to correctly demodulate a frame in several transmission scenarios. All resulting values are provided in Table 5.3. An experimental LoRa testbed has been setup with LoPy 4 devices, a Raspberry Pi 3 gateway equipped with an IMST IC880A LoRa concentrator and a custom network server. In this contribution we are not studying the impact of very heterogeneous RSSI values on the capture coefficients. Indeed, as mentioned in Section 5.2.1, it was shown that the capture effect could be observed even with RSSI gaps below 6 dB. We therefore minimize the RSSI variation by placing all devices next to each other, and setting the transmission power to the maximum 14 dBm value on all of them. The following experiments show that capture events can be witnessed even when all devices are setup identically. We need to synchronize all devices on the same time reference to reproduce the following experiments consistently. For this purpose, we use the LoRaSync synchronization mechanism (c.f. Chapter 4). This

way we are able to tightly control transmission offsets while keeping the timing error below 5 ms. Herein, we use an outdoor gateway equipped with a GPS chip in order to trigger synchronization beacons with great precision. To facilitate experiments, devices are placed indoor and therefore not in a direct line of sight with the receiver. It should be noted that even with this setup where all devices are setup identically, we still witness an uneven RSSI distribution as we will see in Section 5.2.3.2. These received power differences may be imputed to small differences in the positioning of antennas, multipath fading and uncertainties in each hardware component. In any case, our approach consists in estimating the capture coefficients for our setup through the repetition of capture scenarios, and for that we do not need to know the exact RSSI differences between frames nor what causes them. SF7 is used in conjunction with the typical 125 kHz bandwidth for all experiments throughout this paper, and only the 868.1 MHz channel is used. In an effort to minimize the amount of transmitted overhead, all frames are featured with the maximum ToA allowed for SF7, *i.e.* 389.376 ms.


For the asynchronous transmissions occurring when using a Pure ALOHA access scheme, the capture probability depends on the degree of overlap between the interfering packets [Bor 2016, Kosunalp 2015]. In order to evaluate the 2 packets overlap case, the experiment represented in Figure 5.5a has been reproduced for 15 evenly spread offset values. At each offset, over 200 repetitions have been realized for statistical relevance. We then compute 15 intermediary capture coefficients by dividing the number of capture events (*i.e.*, a frame coming from one of the two devices has successfully been demodulated) by the number of iterations. C_2 is finally obtained by averaging all intermediary coefficients. For the 3-packets overlap case, the experiment represented in Figure 5.5b has been realized. This time we vary the offset between the second and third packet, and a fixed 50 ms offset is used between the first two. A similar approximation has been done in [Kosunalp 2015], because varying both offsets would result in an extremely large number of cases to be evaluated. Such assumption proved to be reasonable, as we will see in Section 5.2.3 that the resulting model fits experimental results. Like before, C_3 is obtained by averaging the coefficients of all offsets.

In order to measure the C_i^* coefficients, we trigger slotted transmissions with i simultaneous devices. Since this time no offset variation is required, running the experiments is much faster. This allowed us to measure the coefficients up to $i = 5$.


5.2.3 Experimental results and discussion

5.2.3.1 Model validation with experimental data

In this Section we compare the models established above to performance measurements performed on the real-hardware testbed. 10 devices are used to generate a wide range of traffic loads. Both Pure ALOHA and a Slotted ALOHA access built over LoRaSync are evaluated. For each generation rate the network runs for 30 minutes, and the throughput is sampled every minute. In each plot errorbars rep-



(a) 2-packets capture experiment



(b) 3-packets capture experiment

Figure 5.5: Pure ALOHA capture coefficient measurement

resent 95% confidence intervals computed with Student's t law. The experimental throughput (in Erlangs) is computed by dividing the successful transmission time by the considered duration:

$$T_{\text{exp.}} = \frac{\sum_{\text{pkt} \in \text{received}} ToA_{\text{pkt}}}{\text{duration}} \quad (5.18)$$

The experimental throughput for Pure and Slotted ALOHA are compared to our models in Figure 5.6. The usual models that do not take capture effect into account [Rom 1991] are also provided for the sake of comparison. It is clear that our experimental models faithfully represent the experimental data, while the ones that do not account for capture effect widely underestimate the real-world throughput. For low rates the curves perfectly match, and when the rate increases the model becomes more pessimistic. This is due to the fact that at high rates the probability of a capture event involving many concurring devices increases, and that we have only modeled up to 3 simultaneous devices for Pure ALOHA and 5 for Slotted ALOHA. For generation rates between 0 and 1.5 Erlangs the model falls within experimental errorbars, and is therefore considered to be validated.

5.2.3.2 Impact on the network fairness

It has been shown that the capture effect has a negative impact on the fairness of LoRa networks [Reynders 2018]. To quantify this side effect in our testbed, the contribution of each device to the overall Pure and Slotted ALOHA throughput has been displayed in Figures 5.7a and 5.7b respectively. The average RSSI of the successful transmissions of each device is provided and has been used to sort them. It clearly appears that the devices featured with the strongest average signal power

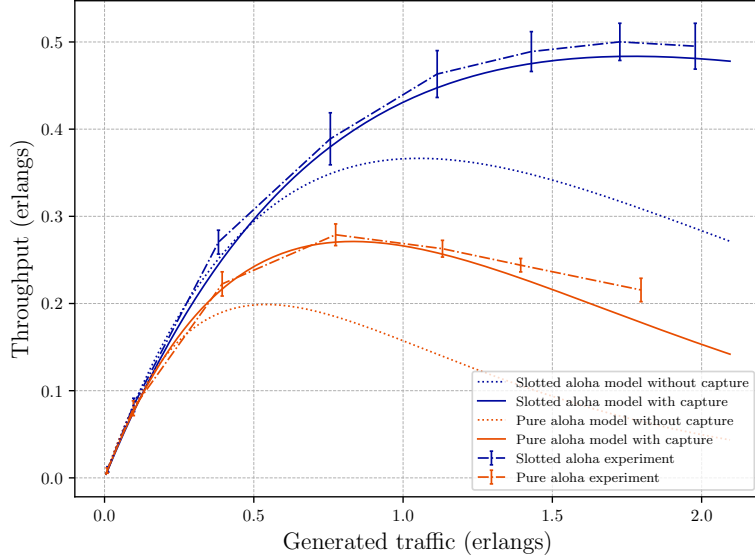


Figure 5.6: Throughput model validation with experimental results (10 devices).

at the receiver are favored in terms of throughput. For instance, for a generated load of 1 Erlang, the weakest of the 10 devices participates up to 5.5% of the useful slotted traffic while the strongest produces 16.6%. Therefore, even in our simplistic testbed where all devices are placed at the same location and use the same transmission power, one device may benefit from three times as many successfully demodulated frames as another. This supports the claim that the capture effect may jeopardize the overall fairness of a LoRa network, especially because in a more realistic deployment devices would be much more spread apart which would result in larger RSSI differences.

In order to quantify this loss of fairness, we additionally compute Jain’s index [Jain 1991] with our experimental data. Such index is independent of the throughput scale, continuous, bounded between 0 and 1 and applies for any population size. As a result, it is widely used in the networking community to analyze how fairly a given bandwidth is shared between traffic flows. By noting n the number of devices and T_i the success rate for device i , Jain’s Fairness Index J is computed as such:

$$J = \frac{(\sum_{i=1}^n T_i)^2}{n \cdot \sum_{i=1}^n T_i^2} \tag{5.19}$$

J has been plotted in Figure 5.8 for the Pure and Slotted ALOHA access schemes. First of all, we notice that no major difference can be witnessed between the Pure and Slotted ALOHA accesses. Besides, we observe that the Jain’s index general trend is to decrease when the traffic load increases. This is due to the fact that when more frames are generated, the occurrence probability of capture events increases. Therefore, the devices featured with the highest average RSSI values are able to successfully transmit more frames than the others, which in turns decreases

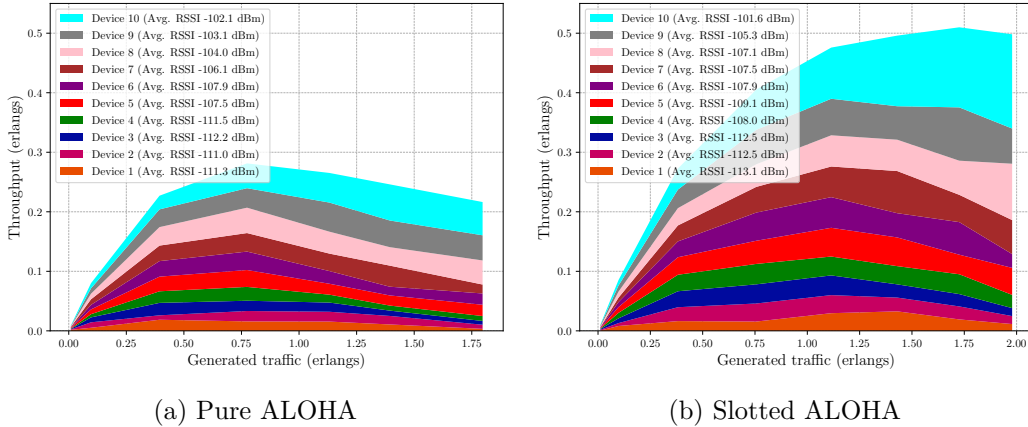


Figure 5.7: Throughput per device

the overall fairness.

5.2.3.3 Discussion

The models in this contribution are established with the strong assumption that the capture coefficients are constant. In reality, each coefficient should depend on the RSSI difference between the concurrent packets. This simplification is acceptable for our testbed because all devices are all similar and close one to another, therefore their RSSI difference does not vary much. Besides we focus our analysis on SF7, and it was shown in [Elshabrawy 2018] that this parameter affects the capture coefficients as well. In fact in this preliminary work we are not investigating the causes of the capture effect, but simply quantifying its impact in a simple setup. Further modeling is therefore desired in future contributions to adapt these coefficients to other deployments. It remains interesting to witness that the small RSSI differences occurring in the testbed are sufficient to trigger capture events. We can safely expect that the capture effect would even be more impacting in a realistic large-scale deployment in which devices are spread far apart.

5.3 Conclusions and perspectives

In the first section of this chapter, throughput, power consumption and energy efficiency models were established to assess the performances of the Slotted ALOHA scheme built on top of *LoRaSync*. This access scheme was compared to the asynchronous Pure ALOHA access, and the LoRaWAN-Sim simulator was used to verify the consistency of the results. Such modeling ultimately allows to find the most efficient slot size for any traffic load. Results show that the more traffic is generated, the smaller margins should be in order to maximize energy efficiency at all times. Such a finding is important to define energy efficient MAC protocols that adaptively adjust their behavior to the network load. Indeed, including a slot size




Figure 5.8: Jain's fairness index

adaption to the TREMA mechanism presented in chapter 3 would be an interesting research perspective to optimize the protocol even further.

In the second section of this chapter, we provide experimental models to accurately picture the performances of a LoRa network subject to the capture effect. To achieve this goal, several transmission scenarios are described and their occurrence probabilities are established. Each scenario is repeated for a large number of iterations on an experimental testbed in order to measure the probability of a success. The resulting models prove to be consistent with the performances observed on the testbed. All in all, the capture effect increases the maximum achievable throughput compared to the expectations provided by more classical Pure and Slotted ALOHA models. However the existence of such effect seriously jeopardizes the network fairness in LoRa deployments. As a matter of fact, even in our small scale setup with minimal RSSI variations, Jain's fairness index significantly drops when the traffic load increases. A more thorough analysis of the capture coefficients is highly needed in future contributions. Indeed we made the hypothesis that these coefficients were constant, when in reality they should be expressed as a function of the RSSI difference between the concurring frames. Tackling this challenge will be mandatory in order to accurately model the throughput of real-life, large-scale LoRa deployments.

Conclusion and Future Works

Ce dernier chapitre conclut ce manuscrit de thèse et relate les pistes de recherches futures identifiées au cours de son élaboration. En un mot, nos travaux montrent que la conception de modes d'accès scalables et énergétiquement efficaces pour les réseaux LoRa nécessite de mettre en place des mécanismes adaptatifs capables de réagir aux variations de la charge réseau. Spécifiquement, le protocole doit être capable de synchroniser et désynchroniser les terminaux dynamiquement en fonction de la charge réseau sondée, sélectionnant ainsi toujours la stratégie la plus efficace.

Dans le processus d'élaboration d'un tel protocole, nous avons tiré parti de modélisations mathématiques, de simulations et d'expériences réelles. Chaque approche a ses avantages, ainsi les utiliser de manière complémentaire nous a permis de construire une représentation fidèle des performances réseau et d'optimiser les paramètres de notre stratégie d'accès. Cependant, d'additionnelles recherches sont toujours requises pour élaborer des modèles versatiles capables de représenter une large variété de topologies réseau. De tels outils seraient capables de piloter le mécanisme de décision de notre protocole adaptatif et ainsi de permettre son implémentation dans différents environnements.

Contents

6.1	Thesis contributions	94
6.2	Lessons learned	95
6.3	Research perspectives	96
6.3.1	Improving the fingerprinting process	96
6.3.2	TREMA improvements and implementation on the testbed	96
6.3.3	Towards more accurate and versatile throughput models	97

By providing cheap devices with low power and long range communication capabilities [Mekki 2019], Low Power Wide Area Networks (LPWAN) keep gaining relevance in industrial and research environments. In particular, Long Range Wide Area Networks (LoRaWANs) have emerged as a very promising technology to support a wide range of distributed sensing applications [Rizzi 2017b]. In the current state-of-the-art, a very simple Pure ALOHA scheme is used in such networks. Such a scheme dwarfs the network capacity to 18% [Abramson 1970]. Given the growth of IoT traffic that has been witnessed in the past years, the research

community has raised strong concerns regarding the scalability of such deployments [Mikhaylov 2016, Bor 2016, Bankov 2016, Georgiou 2017, Adelantado 2017, Van den Abeele 2017, Yousuf 2018]. On the other hand, low-power LoRa devices are typically powered by batteries whose lifetime must be maximized. In that, the protocol design must consider energy efficiency as a prime requirement. In this thesis, we therefore tackle the design of a scalable, yet energy efficient MAC protocol for LoRa networks.

6.1 Thesis contributions

In chapter 1, we present the LoRa modulation and LoRaWAN protocol, and highlight the reasons behind the aforementioned scalability limitations. Chapter 2 then presents the main existing access strategies for wireless networks, and takes stock of relevant contributions introducing MAC layer improvements for LoRa. We show that no solution is able to conciliate network capacity, energy efficiency and adaptability in this current landscape. We argue that among the possible access strategies, the best solution to address this challenge with the current technical limitations is to use synchronization to setup a slotted scheme. Additionally, we state that the only scalable synchronization strategy that does not require additional hardware is to leverage beacon broadcasts. Indeed, using individual acknowledgments would not be applicable to large networks because of downlink Duty Cycle limitations, and because LoRa gateways are half-duplex.

In Chapter 3, we first provide theoretical background on pure and slotted ALOHA protocols. Then, we introduce *Class S* as an extension of the LoRaWAN **Class B** that allows the development of slotted access schemes in accordance with the specificities of LoRa. A notable conclusion of this first contribution is that a beacon-skipping mechanism is highly desired in order to guarantee a satisfying energy efficiency for end-devices. If such a mechanism is implemented, we show that the slotted access is able to significantly improve the overall network capacity while being more efficient than the legacy *Class A* for a wide range of traffic rates. The Pure ALOHA strategy only remains preferable for very low network loads. In that, we state that a switching mechanism is desired to adapt the access scheme to the traffic conditions and select the most energy efficient strategy at all times. As a result, we tackle the design of such a switching mechanism, named TREMA. This system relies on a prior fingerprinting of a given deployment to assess its performances under a wide range of traffic loads. TREMA leverages this fingerprint to probe the current network conditions, and decide whether a synchronous or asynchronous scheme should be used. A signaling protocol is also introduced to trigger the protocol switch when needed. Simulation results show that our switching mechanism is able to follow variations of the offered load, and maximize the overall energy efficiency at all times.

In Chapter 4, we tackled the implementation of *Class S* on real-hardware, which required the consideration of additional constraints. A LoRa testbed has been setup

for this extent. In details, the clock drift of LoPy 4 chips has been measured and modeled. Guard times are added to *Class S* slots based on the worst-case drift derived from this modeling, and the length of these margins is directly linked to the beacon-skipping mechanism introduced in Chapter 3. The resulting *LoRaSync* system is therefore able to provide a robust synchronization to LoRa networks of any size while guaranteeing an energy-efficient management of the end-devices. To complete the contribution, LoRaSync has been implemented on our testbed. The output of this implementation is twofold. First, we show that we are able to correctly control the timing errors and validate that the synchronization operates as expected on real conditions. On the other hand, the evaluation of the testbed performances operating Pure and Slotted ALOHA schemes reveal that the models used until then underestimate the actual network throughput.

Chapter 5 then tackles the establishment of realistic mathematical models in accordance with the observations made on the testbed. First, the performances of *LoRaSync*-operated deployments have been modeled and validated with simulations. Resulting models have been leveraged to find optimal guard times, and this analysis concludes that the synchronization periodicity should be adapted to the traffic load. Such a finding can be used to further improve the TREMA adaptive mechanism. Second, the main finding we made when studying the testbed was that the capture effect frequently occurs in real-world LoRa networks. As a result, we propose experimental models to quantify the impact of such an effect on the overall throughput. These models are consistent with the observed network performances, which confirms that we master the behavior of our network.

6.2 Lessons learned

All in all, this manuscript provides several building blocks to support the implementation of novel energy-efficient, yet scalable MAC layer schemes for LoRa networks. We have seen that there is no one-fits-all solution, and that future protocols must be able to adapt to traffic variations to always maximize the end-devices' energy-efficiency. In details, the network should be able to dynamically synchronize or desynchronize, and adjust the slot size and synchronization periodicity when the slotted access is utilized.

To support this claim, different tools have been utilized during the thesis, all associated with their advantages and inconvenients. We have leveraged simulations to represent the behavior of large-scale networks. This approach can be as accurate as desired granted that all desired features are implemented in the simulator. However the more complex the behavior gets, the longer it takes to generate the results. Modeling has been used to validate our simulations, and then leveraged to fine-tune the parameters of our protocol. Once validated, it is very fast to generate performance estimations with models. However it is much more difficult to represent complex behavior, so a lot of simplifications are generally made. Finally, we have implemented our synchronization scheme on real-hardware, proving the

feasibility of the mechanism in real conditions. Of all methods, hardware implementation takes the most development and running time, but no hypotheses have to be made. This is what unveiled some unexpected results such as occurrences of the capture effect. One weakness of the testbed approach is that it remains difficult to extrapolate results to large scale deployments. Through the combination of all three methods, we can reasonably affirm that the soundness of our contribution has been proven.

6.3 Research perspectives

6.3.1 Improving the fingerprinting process

Many research perspectives have been opened during the thesis. A notable point of interest regarding the switching mechanism is that its decision mechanism relies on a fingerprint established *a-priori*. In details, the throughput, energy efficiency and gateway idle listening time (GILT) must be characterized as a function of the generated traffic. To this end, we have performed extensive simulation campaigns on a specific network topology. In order to run TREMA on any device layout, it would be beneficial to be able to establish such a fingerprint much more quickly. Several options are envisaged to do so:

- Mathematical modeling could be used to derive the network fingerprint very quickly. However we have seen that finding accurate models can be tricky in real deployments, since many factors (such as the capture effect) intervene. A first step in this direction has been performed in Chapter 5, with the accurate modeling of the testbed throughput. Yet, more work is still required to adapt the model to any network topology (more details about this point are provided in Section 6.3.3). The energy efficiency can easily be derived from the throughput and generation rate, however a model for the GILT remains to be found as well.
- Machine Learning methods could also be leveraged to refine the fingerprint on-the-fly while the network is running. However, this would require the collection of samples making the correspondence between the generated traffic and the different metrics. To do so, we need to precisely estimate the generated traffic, which can be done by either through the GILT modeling, or by having devices transmit additional information about all their transmission trials.

6.3.2 TREMA improvements and implementation on the testbed

As a major achievement, TREMA should be implemented on the hardware testbed to demonstrate that the network probing and dynamic protocol adaptation operates as expected with real devices.

Its threshold-based switching mechanism should also be improved with hysteresis, to prevent the network from switching back and forth when the traffic load is close to the threshold. Indeed, the energy cost of the switching procedure should be studied in more details, so that it is triggered only if the benefits in terms of throughput is worth the transmission and reception of all switching commands. Besides, we have showed in Chapter 4 that further than dynamically synchronizing and desynchronizing the network, TREMA should also be able to adapt the slot size and beacon reception periodicity to the traffic load.

6.3.3 Towards more accurate and versatile throughput models

Another significant finding of the thesis is that we lack models able to accurately picture the throughput of ALOHA-operated LoRa networks subject to capture effect. Indeed, most models in the literature that consider this effect rely on RSSI difference thresholds between colliding frames to determine whether a capture event occurs or not. Experiment results show that in reality, capture events are triggered with a probability that depends on the RSSI difference between the frames. We were able to accurately model the throughput in our testbed by measuring these probabilities (the capture coefficients) in a simplified layout where all devices are placed at the same location and use the same transmission power. However, these coefficients should be studied in more details to adapt this modeling to more realistic network topologies in which devices are widely spread, and where the RSSI values are thus far more heterogeneous. Such a contribution would remarkably allow to estimate the network fingerprint to be fed into TREMA's switching mechanism in realistic conditions. Ultimately, the resulting protocol should be adapted to operate on multi-gateway deployments as well.

Performance models of MAC protocols

Several models have been used throughout this manuscript to picture the performances of several MAC protocols. This Appendix regroups all of these models and links them to their reference publication.

A.1 Throughput models

A.1.1 Throughput models for theoretical Random Access Protocols

For this modeling, we assume that frame generations follow a Poisson process with rate λ frames per second, all featured with the same length.

A.1.1.1 Infinite number of devices

Given G the normalized offered load in Erlangs accumulating frame generations and retransmissions, and S the throughput in Erlangs,

Pure ALOHA [Abramson 1970]

$$S(G) = Ge^{-2G} \quad (\text{A.1})$$

Slotted ALOHA [Roberts 1975]

$$S(G) = Ge^{-G} \quad (\text{A.2})$$

For CSMA schemes, we define a the ratio between the packet propagation time and its time-on-air.

1-persistent CSMA [Kleinrock 1975b]

$$S(G, a) = \frac{G[1 + G + aG(1 + G + (aG)/2)]e^{-G(1+2a)}}{G(1 + 2a) - (1 - e^{-aG}) + (1 + aG)e^{-G(1+a)}} \quad (\text{A.3})$$

Non-persistent CSMA [Kleinrock 1975b]

$$S(G, a) = \frac{Ge^{-aG}}{G(1 + 2a) + e^{-aG}} \quad (\text{A.4})$$

A.1.1.2 Finite number of devices

We now consider a finite number of devices n , each generating traffic according to a Poisson process with parameter λ . Given $p = 1 - e^{-\lambda}$ the probability that a given device generates a frame during such time unit and T the overall throughput in Erlangs,

Pure ALOHA [Kleinrock 1975a]

$$T(n, p) = np(1 - p)^{2(n-1)} \quad (\text{A.5})$$

Slotted ALOHA [Kleinrock 1975a]

$$T(n, p) = np(1 - p)^{(n-1)} \quad (\text{A.6})$$

Pure ALOHA with Duty Cycle [Accettura 2020]

$$T(n, g, q) = ngq^{n-1} \quad (\text{A.7})$$

Given ε the Duty Cycle parameter, such that a node transmitting for a duration d shall then remain silent for $d \cdot (\varepsilon - 1)$. The average transmission rate g is expressed as:

$$g = \frac{\lambda}{1 + \lambda\varepsilon} \quad (\text{A.8})$$

Given c the number of channels, the probability q that an end device does not interfere on the same channel with an ongoing transmission is:

$$q = 1 - \frac{1}{c} \frac{\lambda \min(\varepsilon, 2) + 1 - e^{\lambda \min(\varepsilon - 2, 0)}}{1 + \lambda\varepsilon} \quad (\text{A.9})$$

The proof for this model is provided in [Accettura 2020].

Slotted ALOHA with Duty Cycle

With a similar approach as the one presented in [Accettura 2020], we can derive the slotted ALOHA throughput model with Duty Cycle:

$$T(n, g, q) = ngq^{n-1} \quad (\text{A.10})$$

$$g = \frac{1 - e^{-\lambda}}{1 + (\varepsilon - 1)(1 - e^{-\lambda})} \quad (\text{A.11})$$

$$q = 1 - \frac{1}{c} \frac{1 - e^{-\lambda}}{1 + (\varepsilon - 1)(1 - e^{-\lambda})} \quad (\text{A.12})$$

A.1.2 Throughput models for LoRa Access schemes

Class A, Class B, Class C

For LoRaWAN Classes A, B and C, the throughput is exactly the same as pure ALOHA with a finite number of devices. It is provided in equations A.5 (without DC) and A.7 (with DC). We note it T_p in the rest of this chapter.

LoRaSync - Slotted ALOHA without Duty Cycle [Chapter 4]

$$T_s = k_s n q (1 - q)^{n-1} \quad (\text{A.13})$$

With:

$$k_s = \frac{n_{\text{slots}} \cdot ToA_{\text{max}}}{BEACON_PERIOD} \quad (\text{A.14})$$

$$q = p \left[X \leq \frac{L_{\text{slot}}}{ToA_{\text{max}}} \right] = 1 - e^{-\lambda \frac{L_{\text{slot}}}{ToA_{\text{max}}}} \quad (\text{A.15})$$

Pure ALOHA with capture effect [Chapter 5]

$$T = \sum_{i=1}^3 (P_i C_i) \quad (\text{A.16})$$

With:

$$P_1 = n p (1 - p)^{2(n-1)} \quad (\text{A.17})$$

$$P_2 = n(n-1)p^2 \left(\frac{(1-p)^{2(n-2)}}{2} + (1-p)^{2n-3} \right) \quad (\text{A.18})$$

$$P_3 = \frac{n(n-1)p^3(1-p)^{2(n-2)}}{2} (2n-3) \quad (\text{A.19})$$

LoRaSync - Slotted ALOHA with capture effect [Chapter 5]

$$T^* = k_s \cdot \sum_{i=1}^5 (P_i^* C_i^*) \quad (\text{A.20})$$

With:

$$P_i^* = P(i \text{ in } n) = \frac{n!}{i!(n-i)!} p^{*i} (1-p^*)^{n-i} \quad (\text{A.21})$$

$$p^* = 1 - e^{-\lambda L_{\text{slot}}} \quad (\text{A.22})$$

A.2 Power consumption models

We note P_{TX} , P_{RX} and P_{SLEEP} the power consumption of a typical SX1276 LoRa transceiver in transmission, reception and sleeping states respectively, when equipped with a 3.3V battery voltage [Semtech Corporation 2019]. It is computed with the transceiver's supply current in the considered state, respectively 20, 10.8 and $0.2 \cdot 10^{-3}$ mA.

Class A [Chapter 4]

$$P_p = [\lambda P_{\text{TX}} + \rho_s P_{\text{RX}} + (1 - \lambda - \rho_s) P_{\text{SLEEP}}] \cdot n \quad (\text{A.23})$$

With:

$$\rho_s = \lambda \cdot \frac{2 \cdot 0.03}{ToA_{\text{max}}} \quad (\text{A.24})$$

LoRaSync - Slotted ALOHA [Chapter 4]

$$P_s = [(\rho_s + \rho_b)P_{RX} + \lambda P_{TX} + (1 - \rho_s - \rho_b - \lambda)P_{SLEEP}] \cdot n \quad (\text{A.25})$$

With:

$$\rho_b = \frac{ToA_{bcn} + d \cdot T_{bcn} + \nu}{T_{bcn}} \quad (\text{A.26})$$

A.3 Energy efficiency models

Class A [Chapter 4]

$$E_p = \frac{T_p}{P_p} \cdot \frac{\text{bytes}_{\text{pkt}}}{ToA_{\text{pkt}}} \quad (\text{A.27})$$

LoRaSync - Slotted ALOHA [Chapter 4]

$$E_s = \frac{T_s}{P_s} \cdot \frac{\text{bytes}_{\text{pkt}}}{ToA_{\text{pkt}}} \quad (\text{A.28})$$

Scientific production

International conference publications

- Laurent Chasserat, Nicola Accettura and Pascal Berthou. *Short: Achieving energy efficiency in dense LoRaWANs through TDMA*. In IEEE WoWMoM 2020, Cork, Ireland, August 2020. [Chasserat 2020]
- Laurent Chasserat, Nicola Accettura, Balakrishna Prabhu and Pascal Berthou. *TREMA: A traffic-aware energy efficient MAC protocol to adapt the LoRaWAN capacity*. In 2021 International Conference on Computer Communications and Networks (ICCCN), July 2021. ISSN: 2637-9430 [Chasserat 2021]
- Laurent Chasserat, Nicola Accettura and Pascal Berthou. *Experimental throughput models for LoRa networks with capture effect*. In IEEE STWiMob 2022, Thessaloniki, Greece, October 2022. [Chasserat 2022a]

Journal publication

- Laurent Chasserat, Nicola Accettura and Pascal Berthou. *LoRaSync: energy efficient synchronization for scalable LoRaWAN*, 2022. [Chasserat 2022b] **(under reviewing process with minor revision)**

Bibliography

- [3GPP 2012] 3GPP. *3GPP Study on RAN Improvements for Machine-Type Communications*. Technical Report TR 37.868, 3GPP, 2012. (Cited in page 56.)
- [Abdelfadeel 2018] Khaled Q. Abdelfadeel, Dimitrios Zorbas, Victor Cionca, Brendan O’Flynn and Dirk Pesch. *FREE - Fine-grained Scheduling for Reliable and Energy Efficient Data Collection in LoRaWAN*. arXiv:1812.05744 [cs], December 2018. (Cited in pages 32 and 35.)
- [Abramson 1970] Norman Abramson. *THE ALOHA SYSTEM: another alternative for computer communications*. In Proc. of AFIPS ’70 (Fall), AFIPS ’70 (Fall), pages 281–285, Houston, Texas, November 1970. ACM. (Cited in pages 1, 4, 15, 20, 24, 40, 93, and 99.)
- [Accettura 2020] Nicola Accettura and Balakrishna Prabhu. *Ubiquitous multi-gateway LoRa networks: models and performance evaluation*, August 2020. (Cited in page 100.)
- [Adelantado 2017] Ferran Adelantado, Xavier Vilajosana, Pere Tuset-Peiro, Borja Martinez, Joan Melia-Segui and Thomas Watteyne. *Understanding the Limits of LoRaWAN*. IEEE Communications Magazine, vol. 55, no. 9, pages 34–40, September 2017. (Cited in pages 1, 4, 20, and 94.)
- [Ahsan 2019] Shahzeb Ahsan, Syed Ali Hassan, Ahsan Adeel and Hassaan Khaliq Qureshi. *Improving Channel Utilization of LoRaWAN by using Novel Channel Access Mechanism*. In 2019 15th International Wireless Communications & Mobile Computing Conference (IWCMC), pages 1656–1661, June 2019. (Cited in pages 28 and 34.)
- [Akyildiz 2002] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam and E. Cayirci. *Wireless sensor networks: a survey*. Computer Networks, vol. 38, no. 4, pages 393–422, March 2002. (Cited in pages 1 and 4.)
- [Arnbak 1987] J. Arnbak and W. van Blitterswijk. *Capacity of Slotted ALOHA in Rayleigh-Fading Channels*. IEEE Journal on Selected Areas in Communications, vol. 5, no. 2, pages 261–269, February 1987. Conference Name: IEEE Journal on Selected Areas in Communications. (Cited in page 85.)
- [Atmel 2016] Atmel. *ATA8520E datasheet*, 2016. (Cited in page 11.)
- [Bankov 2016] Dmitry Bankov, Evgeny Khorov and Andrey Lyakhov. *On the Limits of LoRaWAN Channel Access*. In 2016 International Conference on Engineering and Telecommunication (EnT), pages 10–14, November 2016. (Cited in pages 1, 4, 20, and 94.)

- [Bankov 2017] Dmitry Bankov, Evgeny Khorov and Andrey Lyakhov. *Mathematical model of LoRaWAN channel access with capture effect*. In 2017 IEEE 28th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC), pages 1–5, October 2017. (Cited in page 85.)
- [Beltramelli 2021] Luca Beltramelli, Aamir Mahmood, Patrik Osterberg, Mikael Gidlund, Paolo Ferrari and Emiliano Sisinni. *Energy efficiency of slotted LoRaWAN communication with out-of-band synchronization*. IEEE Transactions on Instrumentation and Measurement, page 1–1, 2021. (Cited in pages 30, 34, and 65.)
- [Bianco 2021] Giulio Maria Bianco, Romeo Giuliano, Gaetano Marrocco, Franco Mazzenga and Abraham Mejia-Aguilar. *LoRa System for Search and Rescue: Path-Loss Models and Procedures in Mountain Scenarios*. IEEE Internet of Things Journal, vol. 8, no. 3, pages 1985–1999, February 2021. (Cited in page 13.)
- [Björnson 2018] E. Björnson and E. G. Larsson. *How Energy-Efficient Can a Wireless Communication System Become?* In 52nd Asilomar Conference on Signals, Systems, and Computers, pages 1252–1256, 2018. (Cited in pages 48 and 53.)
- [Bor 2016] Martin C. Bor, Utz Roedig, Thiemo Voigt and Juan M. 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, New York, NY, USA, 2016. ACM. event-place: Malta, Malta. (Cited in pages 1, 4, 20, 76, 84, 88, and 94.)
- [Brand 2002] Alex Brand and Hamid Aghvami. *Multiple Access Protocols for Mobile Communications*. John Wiley & Sons Ltd, April 2002. (Cited in page 42.)
- [Callebaut 2020] Gilles Callebaut and Liesbet Van der Perre. *Characterization of LoRa Point-to-Point Path Loss: Measurement Campaigns and Modeling Considering Censored Data*. IEEE Internet of Things Journal, vol. 7, no. 3, pages 1910–1918, March 2020. (Cited in page 13.)
- [Centenaro 2016] Marco Centenaro, Lorenzo Vangelista, Andrea Zanella and Michele Zorzi. *Long-range communications in unlicensed bands: the rising stars in the IoT and smart city scenarios*. IEEE Wireless Communications, vol. 23, no. 5, pages 60–67, October 2016. (Cited in page 8.)
- [Chasserat 2020] Laurent Chasserat, Nicola Accettura and Pascal Berthou. *Short: Achieving energy efficiency in dense LoRaWANs through TDMA*. In IEEE WoWMoM 2020, Cork, Ireland, August 2020. (Cited in page 103.)

- [Chasserat 2021] Laurent Chasserat, Nicola Accettura, Balakrishna Prabhu and Pascal Berthou. *TREMA: A traffic-aware energy efficient MAC protocol to adapt the LoRaWAN capacity*. In 2021 International Conference on Computer Communications and Networks (ICCCN), pages 1–6, July 2021. ISSN: 2637-9430. (Cited in page 103.)
- [Chasserat 2022a] Laurent Chasserat, Nicola Accettura and Pascal Berthou. *Experimental throughput models for LoRa networks with capture effect*. In IEEE STWiMob 2022, Thessaloniki, Greece, October 2022. (Cited in page 103.)
- [Chasserat 2022b] Laurent Chasserat, Nicola Accettura and Pascal Berthou. *LoRaSync: energy efficient synchronization for scalable LoRaWAN*, 2022. (Cited in page 103.)
- [Corporation 2019a] Semtech Corporation. *AN1200.48 SX126x CAD Performance Evaluation*, 2019. (Cited in page 27.)
- [Corporation 2019b] Semtech Corporation. *SX1272 datasheet*, January 2019. (Cited in page 11.)
- [Croce 2018] Daniele Croce, Michele Gucciardo, Stefano Mangione, Giuseppe Santaromita and Ilenia Tinnirello. *Impact of LoRa Imperfect Orthogonality: Analysis of Link-Level Performance*. IEEE Communications Letters, vol. 22, no. 4, pages 796–799, April 2018. (Cited in pages 11, 66, and 83.)
- [Daniel 2019] Belo Daniel, Ricardo Correia, Yuan Ding, Spyridon Nektarios Daskalakis, George Goussetis, Apostolos Georgiadis and Nuno Carvalho. *IQ Impedance Modulator Front-End for Low-Power LoRa Backscattering Devices*. IEEE Transactions on Microwave Theory and Techniques, vol. PP, pages 1–8, October 2019. (Cited in page 83.)
- [Di Vincenzo 2019] Valentina Di Vincenzo, Martin Heusse and Bernard Tourancheau. *Improving Downlink Scalability in LoRaWAN*. In ICC 2019 - 2019 IEEE International Conference on Communications (ICC), pages 1–7, May 2019. ISSN: 1938-1883. (Cited in pages 28 and 36.)
- [Elshabrawy 2018] Tallal Elshabrawy and Joerg Robert. *Analysis of BER and Coverage Performance of LoRa Modulation under Same Spreading Factor Interference*. In 2018 IEEE 29th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), pages 1–6, September 2018. (Cited in pages 83 and 91.)
- [ETSI 2012] ETSI. *ETSI EN 300 220-1 V2.4.1 Electromagnetic compatibility and Radio spectrum Matters (ERM); Short Range Devices (SRD); 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*, 2012. (Cited in page 17.)

- [Farooq 2018] Muhammad Omer Farooq and Dirk Pesch. *Evaluation of Multi-Gateway LoRaWAN with Different Data Traffic Models*. In 2018 IEEE 43rd Conference on Local Computer Networks (LCN), pages 279–282, October 2018. ISSN: 0742-1303. (Cited in page 56.)
- [Fernandes 2019] Rui Fernandes, Rui Oliveira, Miguel Luís and Susana Sargento. *On the Real Capacity of LoRa Networks: The Impact of Non-Destructive Communications*. IEEE Communications Letters, vol. 23, no. 12, pages 2437–2441, December 2019. (Cited in pages 83 and 85.)
- [Gamage 2020] Amalinda Gamage, Jansen Christian Liando, Chaojie Gu, Rui Tan and Mo Li. *LMAC: efficient carrier-sense multiple access for LoRa*. In Proceedings of the 26th Annual International Conference on Mobile Computing and Networking, MobiCom '20, pages 1–13. Association for Computing Machinery, September 2020. (Cited in pages 28 and 34.)
- [Garrido-Hidalgo 2021] C. Garrido-Hidalgo, J. Haxhibeqiri, B. Moons, J. Hoebeke, T. Olivares, F. J. Ramirez and A. Fernández-Caballero. *LoRaWAN Scheduling: From Concept to Implementation*. IEEE Internet of Things Journal, 2021. (Cited in pages 32 and 35.)
- [Georgiou 2017] Orestis Georgiou and Usman Raza. *Low Power Wide Area Network Analysis: Can LoRa Scale?* IEEE Wireless Communications Letters, vol. 6, no. 2, pages 162–165, April 2017. (Cited in pages 1, 4, 20, 83, and 94.)
- [Goursaud 2015] Claire Goursaud and Jean-Marie Gorce. *Dedicated networks for IoT: PHY / MAC state of the art and challenges*. EAI Endorsed Transactions on Internet of Things, vol. 1, page 150597, October 2015. (Cited in page 83.)
- [Haxhibeqiri 2019] J. Haxhibeqiri, I. Moerman and J. Hoebeke. *Low Overhead Scheduling of LoRa Transmissions for Improved Scalability*. IEEE Internet of Things Journal, vol. 6, no. 2, pages 3097–3109, April 2019. (Cited in pages 32 and 35.)
- [Industries 2021] Adafruit Industries. *Adafruit Ultimate GPS Breakout v3 datasheet*, 2021. (Cited in page 70.)
- [Islam 2018] M. T. Islam, B. Islam and S. Nirjon. *Duty-Cycle-Aware Real-Time Scheduling of Wireless Links in Low Power WANs*. In 2018 14th DCOSS, pages 53–60, June 2018. (Cited in pages 31 and 34.)
- [Jain 1991] Raj Jain. *Art of Computer Systems Performance Analysis Techniques For Experimental Design Measurements Simulation And Modeling*. John Wiley & Sons, May 1991. (Cited in pages 77 and 90.)

- [Jalava 2015] P. I. Jalava, Q. Wang, K. Kuuspallo, J. Ruusunen, L. Hao, D. Fang, O. Väisänen, A. Ruuskanen, O. Sippula, M. S. Hoppo, O. Uski, S. Kasurinen, T. Torvela, H. Koponen, K. E. J. Lehtinen, M. Komppula, C. Gu, J. Jokiniemi and M. R. Hirvonen. *Day and night variation in chemical composition and toxicological responses of size segregated urban air PM samples in a high air pollution situation*. Atmospheric Environment, vol. 120, pages 427–437, November 2015. (Cited in page 50.)
- [Jörke 2017] Pascal Jörke, Stefan Böcker, Florian Liedmann and Christian Wietfeld. *Urban channel models for smart city IoT-networks based on empirical measurements of LoRa-links at 433 and 868 MHz*. In 2017 IEEE 28th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC), pages 1–6, October 2017. (Cited in pages 13 and 19.)
- [Kleinrock 1975a] L. Kleinrock and S. Lam. *Packet Switching in a Multiaccess Broadcast Channel: Performance Evaluation*. IEEE Transactions on Communications, vol. 23, no. 4, pages 410–423, April 1975. (Cited in pages 42 and 100.)
- [Kleinrock 1975b] L. Kleinrock and F. Tobagi. *Packet Switching in Radio Channels: Part I - Carrier Sense Multiple-Access Modes and Their Throughput-Delay Characteristics*. IEEE Transactions on Communications, vol. 23, no. 12, pages 1400–1416, December 1975. (Cited in page 99.)
- [Kleinrock 1975c] Leonard Kleinrock and Fouad Tobagi. *Random access techniques for data transmission over packet-switched radio channels*. In Proceedings of the May 19-22, 1975, national computer conference and exposition, AFIPS '75, pages 187–201. Association for Computing Machinery, May 1975. (Cited in pages xiii, 24, 26, and 27.)
- [Kosunalp 2015] Selahattin Kosunalp, Paul D. Mitchell, David Grace and Tim Clarke. *Experimental Study of Capture Effect for Medium Access Control with ALOHA*. ETRI Journal, vol. 37, no. 2, pages 359–368, 2015. (Cited in pages 85, 86, 87, and 88.)
- [Laner 2013] Markus Laner, Philipp Svoboda, Navid Nikaein and Markus Rupp. *Traffic Models for Machine Type Communications*. In ISWCS 2013; The Tenth International Symposium on Wireless Communication Systems, pages 1–5, August 2013. (Cited in page 56.)
- [Leentvaar 1976] K. Leentvaar and J. Flint. *The Capture Effect in FM Receivers*. IEEE Transactions on Communications, vol. 24, no. 5, pages 531–539, May 1976. (Cited in page 31.)
- [Leonardi 2018] Luca Leonardi, Filippo Battaglia, Gaetano Patti and Lucia Lo Bello. *Industrial LoRa: A Novel Medium Access Strategy for LoRa in Industry 4.0 Applications*. In IECON 2018 - 44th Annual Conference of the

- IEEE Industrial Electronics Society, pages 4141–4146, October 2018. (Cited in pages 31 and 35.)
- [LoRa Alliance 2020a] LoRa Alliance. *LoRaWAN® L2 1.0.4 Specification (TS001-1.0.4)*, October 2020. (Cited in pages 1, 4, 8, 15, 19, 45, 47, 54, and 83.)
- [LoRa Alliance 2020b] LoRa Alliance. *TS2-1.1.0 LoRaWAN Backend Interfaces Specification*, 2020. (Cited in page 15.)
- [LoRa Alliance 2021] LoRa Alliance. *RP002-1.0.3 LoRaWAN® Regional Parameters*, 2021. (Cited in pages 15, 17, 19, 36, 45, and 67.)
- [Mallat 2014] A. Mallat and L. Vandendorpe. *Joint estimation of the time delay and the clock drift and offset using UWB signals*. In 2014 IEEE International Conference on Communications (ICC), pages 5474–5480, June 2014. ISSN: 1938-1883. (Cited in page 65.)
- [Mekki 2019] Kais Mekki, Eddy Bajic, Frederic Chaxel and Fernand Meyer. *A comparative study of LPWAN technologies for large-scale IoT deployment*. ICT Express, vol. 5, no. 1, pages 1–7, March 2019. (Cited in pages 1, 4, 12, and 93.)
- [Mikhaylov 2016] K Mikhaylov, J Petäjajarvi and T Haenninen. *Analysis of Capacity and Scalability of the LoRa Low Power Wide Area Network Technology*. In European Wireless 2016; 22th European Wireless Conference, Proceedings of EW '16, pages 1–6, Oulu, Finland, May 2016. (Cited in pages 1, 4, 20, and 94.)
- [Nikaein 2013] Navid Nikaein, Markus Laner, Kaijie Zhou, Philipp Svoboda, Dejan Drajić, Milica Popovic and Srdjan Krco. *Simple Traffic Modeling Framework for Machine Type Communication*. In ISWCS 2013; The Tenth International Symposium on Wireless Communication Systems, pages 1–5, August 2013. (Cited in page 56.)
- [Noreen 2018] Umber Noreen, Laurent Clavier and Ahcene Bounceur. *LoRa-like CSS-based PHY layer, Capture Effect and Serial Interference Cancellation*. In European Wireless 2018; 24th European Wireless Conference, pages 1–6, May 2018. (Cited in pages 76 and 84.)
- [O’Kennedy 2020] Morgan O’Kennedy, Thomas Niesler, Riaan Wolhuter and Nathalie Mitton. *Practical evaluation of carrier sensing for a LoRa wildlife monitoring network*. In 2020 IFIP Networking Conference (Networking), pages 614–618, June 2020. (Cited in pages 28 and 34.)
- [Petajarvi 2015] Juha Petajarvi, Konstantin Mikhaylov, Antti Roivainen, Tuomo Hanninen and Marko Pettissalo. *On the coverage of LPWANs: range evaluation and channel attenuation model for LoRa technology*. In 2015 14th

- International Conference on ITS Telecommunications (ITST), pages 55–59, December 2015. (Cited in page 13.)
- [Petroni 2022] Andrea Petroni and Mauro Biagi. *Interference Mitigation and Decoding Through Gateway Diversity in LoRaWAN*. IEEE Transactions on Wireless Communications, pages 1–1, 2022. (Cited in page 85.)
- [Pham 2018] Congduc Pham. *Investigating and experimenting CSMA channel access mechanisms for LoRa IoT networks*. In 2018 IEEE Wireless Communications and Networking Conference (WCNC), pages 1–6, April 2018. (Cited in pages 27 and 34.)
- [Pham 2021] Congduc Pham and Muhammad Ehsan. *Dense Deployment of LoRa Networks: Expectations and Limits of Channel Activity Detection and Capture Effect for Radio Channel Access*. Sensors, vol. 21, no. 3, page 825, January 2021. (Cited in pages 27 and 34.)
- [Piyare 2018] Rajeev Piyare, Amy L. Murphy, Michele Magno and Luca Benini. *On-Demand TDMA for Energy Efficient Data Collection with LoRa and Wake-up Receiver*. In 2018 14th WiMob, pages 1–4, October 2018. (Cited in pages 31 and 35.)
- [Polonelli 2019] Tommaso Polonelli, Davide Brunelli, Achille Marzocchi and Luca Benini. *Slotted ALOHA on LoRaWAN-Design, Analysis, and Deployment*. Sensors, vol. 19, no. 4, page 838, January 2019. (Cited in pages 30, 34, and 65.)
- [Pop 2017] Alexandru-Ioan Pop, Usman Raza, Parag Kulkarni and Mahesh Sooriyabandara. *Does Bidirectional Traffic Do More Harm Than Good in LoRaWAN Based LPWA Networks?* arXiv:1704.04174 [cs], December 2017. (Cited in pages 36 and 77.)
- [Rademacher 2021] Michael Rademacher, Hendrik Linka, Thorsten Horstmann and Martin Henze. *Path Loss in Urban LoRa Networks: A Large-Scale Measurement Study*, September 2021. (Cited in page 13.)
- [Reynders 2016] Brecht Reynders and Sofie Pollin. *Chirp spread spectrum as a modulation technique for long range communication*. In 2016 Symposium on Communications and Vehicular Technologies, pages 1–5, November 2016. (Cited in page 9.)
- [Reynders 2018] B. Reynders, Q. Wang, P. Tuset-Peiro, X. Vilajosana and S. Pollin. *Improving Reliability and Scalability of LoRaWANs Through Lightweight Scheduling*. IEEE Internet of Things Journal, vol. 5, no. 3, pages 1830–1842, June 2018. (Cited in pages 31, 35, and 89.)

- [Rizzi 2017a] M. Rizzi, P. Ferrari, A. Flammini, E. Sisinni and M. Gidlund. *Using LoRa for industrial wireless networks*. In 2017 IEEE 13th WFCS, pages 1–4, May 2017. (Cited in pages 30 and 34.)
- [Rizzi 2017b] Mattia Rizzi, Paolo Ferrari, Alessandra Flammini and Emiliano Sisinni. *Evaluation of the IoT LoRaWAN Solution for Distributed Measurement Applications*. IEEE Transactions on Instrumentation and Measurement, vol. 66, no. 12, pages 3340–3349, December 2017. Conference Name: IEEE Transactions on Instrumentation and Measurement. (Cited in pages 1, 4, and 93.)
- [Roberts 1975] Lawrence Roberts. *ALOHA Packet System With and Without Slots and Capture*. ACM SIGCOMM Computer Communication Review, vol. 5, pages 28–42, April 1975. (Cited in pages 25, 42, 85, and 99.)
- [Rom 1991] R. Rom and M. Sidi. *Multiple Access Protocols: Performance and Analysis*. SIGMETRICS Perform. Evaluation Rev., vol. 18, page 11, 1991. (Cited in pages 76, 85, and 89.)
- [Semtech Corporation 2013] Semtech Corporation. *AN1200.13 SX1272/3/6/7/8: LoRa Modem Designer’s Guide*, 2013. (Cited in page 10.)
- [Semtech Corporation 2015] Semtech Corporation. *AN1200.22 LoRa Modulation Basics Application Note*, 2015. (Cited in pages 9, 10, 11, and 66.)
- [Semtech Corporation 2019] Semtech Corporation. *SX1276/77/78/79 datasheet*, January 2019. (Cited in pages 46, 48, 57, 78, 79, 83, and 101.)
- [Semtech Corporation 2020] Semtech Corporation. *Predicting LoRaWAN Capacity*, May 2020. Technical paper. (Cited in page 10.)
- [Seye 2018] Madoune R. Seye, Bassirou Ngom, Bamba Gueye and Moussa Diallo. *A Study of LoRa Coverage: Range Evaluation and Channel Attenuation Model*. In 2018 1st International Conference on Smart Cities and Communities (SC-CIC), pages 1–4, July 2018. (Cited in page 13.)
- [Smiljkovic 2014] Katerina Smiljkovic, Vladimir Atanasovski and Liljana Gavrilovska. *Machine-to-Machine traffic characterization: Models and case study on integration in LTE*. In 2014 4th International Conference on Wireless Communications, Vehicular Technology, Information Theory and Aerospace Electronic Systems (VITAE), pages 1–5, May 2014. (Cited in page 56.)
- [Sorensen 2019] Rene Brandborg Sorensen, Nasrin Razmi, Jimmy Jessen Nielsen and Petar Popovski. *Analysis of LoRaWAN Uplink with Multiple Demodulating Paths and Capture Effect*. In ICC 2019 - 2019 IEEE International Conference on Communications (ICC), pages 1–6, May 2019. (Cited in page 85.)

- [Subbaraman 2022] Raghav Subbaraman, Yeswanth Guntupalli, Shruti Jain, Rohit Kumar, Krishna Chintalapudi and Dinesh Bharadia. *BSMA: scalable LoRa networks using full duplex gateways*. In Proceedings of the 28th Annual International Conference on Mobile Computing And Networking, MobiCom '22, pages 676–689, New York, NY, USA, October 2022. Association for Computing Machinery. (Cited in pages 28 and 34.)
- [Tessaro 2018] Luca Tessaro, Cristiano Raffaldi, Maurizio Rossi and Davide Brunelli. *Lightweight Synchronization Algorithm with Self-Calibration for Industrial LORA Sensor Networks*. In 2018 Workshop on Metrology for Industry 4.0 and IoT, pages 259–263, April 2018. (Cited in pages 31, 35, and 65.)
- [Tobagi 1975] F. Tobagi and L. Kleinrock. *Packet Switching in Radio Channels: Part II - The Hidden Terminal Problem in Carrier Sense Multiple-Access and the Busy-Tone Solution*. IEEE Transactions on Communications, vol. 23, no. 12, pages 1417–1433, December 1975. (Cited in page 27.)
- [uBlox 2019] uBlox. *SARA N210 datasheet*, November 2019. (Cited in page 11.)
- [Van den Abeele 2017] Floris Van den Abeele, Jetmir Haxhibeqiri, Ingrid Moerman and Jeroen Hoebeke. *Scalability Analysis of Large-Scale LoRaWAN Networks in ns-3*. IEEE Internet of Things Journal, vol. 4, no. 6, pages 2186–2198, December 2017. Conference Name: IEEE Internet of Things Journal. (Cited in pages 1, 4, 20, and 94.)
- [Veitch 2004] Darryl Veitch, Satish Babu and Attila Pásztor. *Robust synchronization of software clocks across the internet*. In Proceedings of the 4th ACM SIGCOMM conference on Internet measurement, IMC '04, pages 219–232, New York, NY, USA, October 2004. Association for Computing Machinery. (Cited in page 65.)
- [Whitehouse 2005] K. Whitehouse, A. Woo, F. Jiang, J. Polastre and D. Culler. *Exploiting the capture effect for collision detection and recovery*. In The Second IEEE Workshop on Embedded Networked Sensors, 2005. EmNetS-II., pages 45–52, May 2005. (Cited in page 84.)
- [Yousuf 2018] Asif M. Yousuf, Edward M. Rochester, Behnam Ousat and Majid Ghaderi. *Throughput, Coverage and Scalability of LoRa LPWAN for Internet of Things*. In 2018 IEEE/ACM 26th International Symposium on Quality of Service (IWQoS), pages 1–10, June 2018. ISSN: 1548-615X. (Cited in pages 1, 4, 20, and 94.)
- [Zarrini 2018] Morteza Zarrini and Abdorasoul Ghasemi. *Loss and delay analysis of non-Poisson M2M traffic over LTE networks*. Transactions on Emerging Telecommunications Technologies, vol. 29, no. 2, page e3273, 2018. (Cited in page 56.)

- [Zorbas 2020] Dimitrios Zorbas, Khaled Abdelfadeel, Panayiotis Kotzanikolaou and Dirk Pesch. *TS-LoRa: Time-slotted LoRaWAN for the Industrial Internet of Things*. *Computer Communications*, vol. 153, pages 1–10, March 2020. (Cited in pages 32 and 35.)