Simulating scalable Long Range Wide Area Networks for very low power monitoring applications

<u>N. Accettura</u>^a, B. Prabhu^a, T. Monteil^a

^aLAAS-CNRS, Université de Toulouse, CNRS, INSA, Toulouse, France Email: <u>nicola.accettura@laas.fr</u>

The recent emergence of Machine-to-Machine networks and the Internet of Things has Abstract: generated interest in technologies that can underpin networks of low power devices that sporadically exchange data. Both the research and the standardization communities have been actively participating in seeking and defining tech-nologies of the underlying structure of future networks. These Low Power Wide Area Networks (LPWAN) will support large number of nodes with long-range wireless communications on the one hand, while guaranteeing a long life for very cheap sensing devices on the other. Among the many applications that can be supported on these network, there is a wide range of monitoring applications which include low power tracking and radiolocation. The radiolocation application enables a device to determine its position using signals from base stations. It is similar in spirit to GPS except that the signals of radiolocation are provided by terrestrial base stations instead of satellites. This contribution focuses on one such LPWAN technology, called Long Range Wide Area Networks (LoRaWAN) which incorporates the necessary features for radiolocation purposes. A LoRaWAN consists of a central controller which controls several gateways, that in turn listen to incoming data sent by low powered end-devices over sub-gigahertz unlicensed spectra. The geographical placement is such that a signal sent by one end-device is received by at least three gateways. Multilateration for the location of an end-device can be performed by using a combination of both the strength of the signals and the time difference of arrivals measured at the multiple gateways which are in the transmission range of this end-device.

In order to feature LoRaWAN in city scaled scenarios, it is very important to design a proper positioning scheme for gateways and end-devices to assess the related performances in several network conditions. Given the lack of simulation instruments targeting scalable LoRaWAN deployments, this contribution introduces an event based simulation environment that has been developed to characterize the performances of very large LoRaWAN scenarios. Through huge simulation campaigns, the capacity of such network has been thoroughly evaluated in order to predict the impact of low power radiolocation for future tracking applications. Simulation results have also validated the expected probability models for the LoRaWAN throughput and have provided a realistic description of more complex application scenarios.

Keywords: Low power wide area networks, performance evaluation, scalable deployments, simulations

1 INTRODUCTION

The Internet of Things (IoT) will enable new smart services in various contexts such as smart homes, smart cities and smart manufacturing. Billions of IoT objects are expected to generate a huge amount of information (see Warneke et al., 2001), thus requiring Machine-To-Machine (M2M) communications to be energy-efficient over scalable deployments. The IEEE 802.15.4 Standard (2016) provides a technological framework for these IoT objects to communicate with each other. It is intended for low-power and short-range end-devices, and works in lossy wireless communications as well. The end-to-end communication range is extended through relaying capabilities implemented on all devices, at the cost of an additional energy wasted for overhearing the radio channel. This feature coupled with synchronization has been proved by Palattella et al. (2012) to work very well for industrial applications, i.e., in contexts where reliability and time-determinism represent additional constrains on the design of communication protocols.

When time determinism is not a strict requirement for very low traffic applications, low-power long-range communications represent a single-hop low-cost solution for city scaled deployments. This recently led to the definition of Low Power Wide Area Networks (LPWAN) as cellular-type network architecture exploiting the radio resources of unlicensed spectra. In addition, standardization efforts within the IETF lpwan (2016) working group have been triggered for handling interoperable solutions. Among the existing LPWAN technologies, Long Range Wide Area Networks (LoRaWAN) have been introduced by Sornin et al. (2016) to target both network scalability and communication flexibility. LoRaWAN builds up on the top of LoRa hardware technology defined by Semtech specification (2013). In a LoRaWAN, nodes are organized according to a hierarchical architecture. There is one central server controlling several gateways, that in turn control the end-devices. LoRaWAN achieves network scalability and improvements in network lifetime by moving complexity from the end-devices to the gateways and on the network server. It is up to the gateways to decode frames on several channels, and relay them to the server; the server looks for replicas and discards them. Thus, the end-devices just have to transmit their data without dealing with redundancies. Petäjäjärvi et al. (2015) proved that LoRa networks can achieve very long communication ranges by using Chirp Spread Spectrum techniques. In order to achieve a fair use of the unlicensed spectrum among concurrent operators and/or LPWAN technologies, the European Telecommunications Standards Institute (ETSI, 2015) has imposed a duty cycle limitation that has to be hard-coded on end-devices. This duty cycle limitation lets a transmission by an end-device to be followed by a period of silence of certain duration in order to allow other end-devices to use the radio resources. It also means that the data rate of an end-device is limited by this duty cycle policy. Another factor that impacts the data rate of an end-device in a LoraWAN network is the medium access scheme which is based on the ALOHA access scheme (see Bertsekas and Gallager, 1992). With ALOHA, end-devices transmit independently, with two or more simultaneous transmissions resulting in a collision. As a consequence, an end-device cannot grab all the bandwidth because other transmissions would interference with its one ones.

The throughput of LoRaWAN has not been studied for scalable multi-channel multi-gateway deployments. Through a proper placement of the gateways in a city scaled environment, it is possible to find a throughput expression that depends upon the average density of end-devices per area more than upon their absolute number. Such technological configuration has never been investigated from a detailed probabilistic point of view. The aim of this paper is to provide results on the throughput that can be obtained by end-devices in the LoRaWAN for gateways placed according to the honeycomb topology. Such a topology is popular in cellular networks for the placement of base-stations. The end-devices will be assumed to be spread randomly in the coverage area of these gateways. A formula for the throughput in the setting without duty cycle limitation will be provided. This formula is based upon a probabilistic model for LoRaWAN, introduced by Accettura et al. (2017). In that paper, the authors used their probabilistic model to compute the radiolocation throughput for future geofencing and low power tracking applications based on wide LoRaWAN deploymnets: to get the location of the transmitting end-device, a frame has to be decoded by at least three gateways to perform signal multilateration. Instead, this paper focuses also on the investigation of the pure throughput related to the correct reception of frames by at least one gateway, i.e., targeting sensing applications. The rest of the paper is organized as follows. The description of the model and the formula for the throughput is given (see Sec. 2), the ad-hoc simulation environment mimicking LoRaWAN city scaled deployments is described in Sec. 3, and the simulation results are presented in Sec. 4. Sec. 5 eventually concludes the paper and envisages future works.

2 SYSTEM MODEL

As described above, the LoRaWAN medium access scheme is ALOHA-based. Performances of random access schemes for wireless networks have been deeply investigated since fifty years ago, when Abramson (1970)

firstly described Pure ALOHA networks in which end-device could transmit in an asynchronous fashion. Then, Roberts (1975) introduced Slotted ALOHA, that has definitely triggered more research, given its applicability in time-synchronized networks. For example, devices try to join a cellular network by overhearing the radio, until they listen to a beacon frame, that lets them synchronize on a time-slotted structure. The difference in throughput between Slotted and Pure ALOHA networks can be substantial. For example, in a simple network with just one gateway, Slotted ALOHA can carry twice the traffic as Pure ALOHA (see Bertsekas and Gallager, 1992). In a Slotted ALOHA network, due to it synchronous nature, a transmission can collide only with other transmission that began in this slot. On the other hand, in the Pure ALOHA network, a frame can collide with either ongoing transmission or future transmission that may begin during the transmission of this frame. Thus, a frame is successfully transmitted if the channel used is free during a time-interval of two frame lengths, whereas in a Slotted ALOHA the channel needs to be free only for one frame length.

The LoRaWAN architecture permits a frame transmitted by an end-device to several gateways to be successfully received by the network server, if correctly decoded by at least one gateway. The resulting throughput is intuitively higher than in a singlegateway deployment. In this paper, it has been assumed that the gateways are placed according to a hexagonal lattice, as pictured in Fig. 1. A black node represents a gateway. This hexagonal lattice or honeycomb structure is popular in cellular networks. The distance between two gateways is R which is also the transmission range of the end-devices. This fact means that any end-device on the plane will be connected to at least three gateways. For example, the transmissions of an end-device Pwill be received by gateways A, B, and C. This property of the placement of gateways is also useful for radiolocation of end-devices. radiolocation is one of the applications envisaged of LoRaWAN networks as backup to GPS systems: the reception of a frame by at least 3 gateways can be used to perform multilateration, and this is achieved by ensuring that the maximum transmission range R is equal to the distance between two



Figure 1. Topology of LoRaWAN gateways and radio coverage.

nearby gateways. It can also be seen that any end-device in the green area (and other similar areas) will be received by four gateways. For example, for an end-device in the green area, a transmitted frame will be received by gateways A, B, C, and D.

It will be assumed that end-devices are spread on the surface according to a Poisson point process, characterized by an intensity of μ end-devices per unit area. The spatial Poisson process is used as a model for completely random location of end-devices, and it has some nice properties that make the analysis simpler, according to Baccelli and Błaszczyszyn (2010). One of the properties is that the number of end-devices into any area of size A is a Poisson random variable with parameter μA . Furthermore, the process is spatially homogeneous. For two non-overlapping areas, A_1 and A_2 , the number of nodes in each area are mutually independent, and both these quantities have the same Poisson distribution. It has to be noted that the locations of a given number of end-devices into an area are uniformly distributed.

A similar analysis was considered by Bajović et al. (2014), who investigated the throughput of Slotted ALOHA networks with coordinated gateways randomly spread on a given area. They modelled the throughput assuming that end-devices can fall outside the coverage range of gateways. However, they found that the throughput formulation can be only evaluated through approximations and numerical methods. Instead, this contribution targets Pure ALOHA LoRaWAN networks over honeycomb-like deployments of coordinated gateways. In this context, a closed form for the throughput formulation is available and presented in what follows.

Some assumptions have been made on the traffic generation pattern of each end-device: (i) the applications running on the end-device generate frames either according to periodic reporting schemes or in response to any request initiated by the LoRaWAN server, thus regardless of the scheduled radio activity; (ii) frames are generated according to a Poisson process of rate λ frames per unit time; (iii) a frame transmission lasts τ seconds; (iv) the buffer of any end-device has room for at most one frame. Since a frame could be queued up and waiting for its transmission due to any imposed duty cycle limitation, any additional frame generated during this waiting time will be lost. The resulting loss probability is characterized in Sec. 4, while the buffer size impact in mitigating such losses will be investigated in future contributions.

In Accettura et al. (2017), by using geometrical arguments coupled with the inclusion-exclusion principle, the resulting throughput for the radiolocation model in honeycomb-like LoRaWAN deployments (e.g, the one sketched above in Fig. 1) was computed. The difference between the pure throughput and the radiolocation one is that for a frame to be considered in the pure throughput model, it has be to be received correctly by at least one gateway whereas for the radiolocation throughput it has to be received by at least three gateways. Their approach can be used to also obtain a formula for the pure throughput.

According to the work by Accettura et al. (2017), the average throughput Γ^r features applications for radiolocation through multilateration. Indeed, the formulation shown in eq. (1) takes into account that a LoRaWAN frame can be used for radiolocation purposes if and only if it has been received by at least 3 gateways. In addition, such a formulation is computed over a disk with radius R picked randomly onto the considered honeycomb deployment:

$$\Gamma^{r} = p\mu\pi \cdot \left[\left(\frac{2\pi}{\sqrt{3}} - 2 \right) e^{-\frac{(2-p)p\mu\pi}{n} \left(\frac{3}{2} + \frac{\sqrt{3}}{\pi} \right)} + \left(\frac{4\pi}{\sqrt{3}} - 6 \right) e^{-\frac{(2-p)p\mu\pi}{n} \left(\frac{5}{3} + \frac{\sqrt{3}}{\pi} \right)} + \left(-\frac{6\pi}{\sqrt{3}} + 9 \right) e^{-\frac{(2-p)p\mu\pi}{n} \left(\frac{5}{3} + \frac{3\sqrt{3}}{2\pi} \right)} \right].$$
(1)

In the previous expression, n is the number of available channels, while p is the probability that an end-device is transmitting a frame. The value of p is related to the traffic generation rate λ on any end-device, and to the time-on-air of a frame sent by any given end-device to the closest gateways. Assuming that the interarrival times between consecutive frame generations follow an exponential distribution, p is in fact the probability that an end-device starts transmitting a frame on the air in a time interval shorter than τ : $p = 1 - e^{-\tau \lambda}$.

With similar arguments, the average throughput Γ over a disk with radius R picked randomly onto the considered honeycomb deployment and related to the correct reception of a frame by at least one gateway is expressed as follows (its exact derivation will be presented into future works):

$$\Gamma = p\mu\pi \cdot \left[\frac{2\pi}{\sqrt{3}} e^{-\frac{(2-p)p\mu\pi}{n}} + \left(-\frac{4\pi}{\sqrt{3}} + 3 \right) e^{-\frac{(2-p)p\mu\pi}{n} \left(\frac{4}{3} + \frac{\sqrt{3}}{2\pi} \right)} + \left(-\frac{2\pi}{\sqrt{3}} + 3 \right) e^{-\frac{(2-p)p\mu\pi}{n} \left(\frac{5}{3} + \frac{\sqrt{3}}{2\pi} \right)} \\
+ \left(\frac{2\pi}{\sqrt{3}} - 2 \right) e^{-\frac{(2-p)p\mu\pi}{n} \left(\frac{3}{2} + \frac{\sqrt{3}}{\pi} \right)} + \left(\frac{4\pi}{\sqrt{3}} - 6 \right) e^{-\frac{(2-p)p\mu\pi}{n} \left(\frac{5}{3} + \frac{\sqrt{3}}{\pi} \right)} \\
+ \left(-\frac{2\pi}{\sqrt{3}} + 3 \right) e^{-\frac{(2-p)p\mu\pi}{n} \left(\frac{5}{3} + \frac{3\sqrt{3}}{2\pi} \right)} \right].$$
(2)

Since the distance in between any pair of gateways has been set equal to the interference range R, both throughput expressions are independent of R, thus fitting scalable deployments.

3 SIMULATION ENVIRONMENT

In order to validate the model shown so far, city scaled LoRaWAN deployments have been mimicked through an event-based Python simulator, *LoRaWAN-sim*¹, optimized to achieve processing simplicity. Indeed, with the aim of simulating several network conditions, the simulator has been properly developed to accomplish with many deployment settings. For instance, it is possible to control all the transmission parameters of end-devices, while making the deployment configuration easy through the use of modular Python classes. In addition, during each simulator all events are logged to file for post-processing duties.

Given the honeycomb-based displacement of gateways and the resulting independence from the real value of the transmission range R, Fig. 2 pictures an example simulated deployment area large around $10 \times 10R^2$. Gateways are drawn with blue circles, while end-devices are randomly scattered according to a Poisson point process and indicated with red circles.

The models of eqs. (1) and (2) assume that a frame generated by an end-device can collide on gateways closer than R with frames generated by all end-devices closer than 2R. In facts, such models do not take into account that the average performances are expected to be bettered off on the boundaries of finite deployments: a smaller amount of traffic will be produced in such border zones, with consequent less frame collisions. As a consequence, the results considered hereafter for evaluation purposes are those related to end-devices within the gray rectangle A_{sim} (pictured in Fig. 2), whose outline is distant 2R from the outline of the whole simulation area.

¹More information about the simulation environment is present at http://homepages.laas.fr/naccettu.

For each end-device, the number of frames transmitted on the air and received by at least one gateway during the simulation time T_{sim} is registered together with the number of gateways able to decode any frame.

First, the value N is obtained by summing up the number of frames received by at least 1 gateway over all the end-devices falling within the sub-area A_{sim} during T_{sim} . Then, the average occupancy of the radio resources for useful communications Δ is computed as follows:

$$\Delta = \frac{\pi \tau}{\mathcal{A}_{sim} T_{sim}} N,\tag{3}$$

and it represents the estimated throughput for regular communications to be compared against the expected value Γ defined by eq. (2). The portion $\frac{\pi}{A_{sim}}N$ represents the average portion of N packets sent by any end-device belonging to a disk with radius R. Then, the average amount of time of successful transmissions is obtained by multiplying that portion with the time-on-air τ . Normalizing such time to the simulation time T_{sim} , the average occupancy of the radio resources for successful transmissions is achieved.



Figure 2. LoRaWAN deployment example.

Similarly, summing up the number of frames received by at least 3 gateways over all the end-devices falling within the sub-area \mathcal{A}_{sim} during T_{sim} , the value N^r is obtained. This value is indicated as Δ^r , $\Delta^r = \frac{\pi \tau}{\mathcal{A}_{sim} T_{sim}} N^r$, and in fact represents the estimated throughput for radiolocation purposes to be compared against the expected value Γ^r defined by eq. (1).

It is worth pointing out that the developed simulator embeds the possibility of mimicking duty cycle limitations enforced by ETSI (2015) over the 863 – 870 MHz ISM band in Europe (duty cycle limitations are county-specific, and policed by local standardizing bodies). That is, if a LoRaWAN frame is sent by any end-device to the closest gateways for an amount of time of τ seconds, then the transmitting end-device is required to turn off its radio for the following $T_{off} = (1/DC - 1)\tau$ seconds, where DC is the duty cycle limitation expressed as a number belonging to the interval [0, 1].

Such a duty cycle limitation shapes the generated traffic according to the service rate of an M/D/1 queueing system. Indeed, the application running on any end-device generates frames with inter-frame generation time distributed according to an exponential distribution with average rate λ . The service time is deterministic and equal to 1/DC, since the time between two consecutive starts of frame transmission must at least be 1/DC seconds. Moreover, since the buffer can store at most one single packet, the system is an M/D/1/2 queue.

4 VALIDATION RESULTS

The theoretical models presented in Sec. 2 have been validated by results obtained through an extensive simulation campaign with the Python environment presented in Sec. 3. For each run, the simulation duration T_{sim} has been set to be 1 hour long, while the simulation area \mathcal{A}_{sim} is the one pictured in the example of Fig. 2. The LoRaWAN gateways have been configured to operate over 3 channels, according to the default configuration suggested in LoRaWAN regional parameters (2016) that includes 3 channels centered at 868.10, 868.30, and 868.50 MHz (the model presented in Sec. 2 also tailors implementations with more than 3 channels). In addition, the data rate is configured to be the highest among the ones made available with the LoRa modulation. In order to minimize the overhead due to encapsulation duties, each end-device is assumed to generate frames of maximum size. According to Sornin et al. (2016) and to the LoRaWAN regional parameters (2016), when the maximum data rate is used, the maximum frame size has got a physical payload of 240 bytes and an insured application payload of 222 bytes. Using the Semtech specification (2013), the resulting time-on-air τ_{max} for maximum size frames is 368.896 ms long. Since a duty cycle limitation DC = 0.01 is enforced over the 3 aforementioned default channels, each end-device generates a frame every $\frac{\tau_{max}}{DC} = 100\tau_{max}$ seconds on average, according to an exponential distribution with intensity $\lambda = \frac{1}{100\tau_{max}}$.

To picture a general behavior of LoRaWAN networks, simulations have been run by varying the end-device density μ . To show the effect of a duty cycle limitation enforced on end-devices, for each value μ two simulations have been run: in one case, frames generated on any end-device are immediately transmitted; in the other

case, any end-device implements the M/D/1/2 queuing policy described in Sec. 3, that enforces a duty cycle limitation of DC = 0.01. To achieve statistical significance, 10 different simulations have been run for each configuration by feeding the random number generator with 10 different seeds. In facts, the results plotted hereafter represent mean values obtained by averaging over the 10 simulation repetitions.

Blue and red lines in Fig. 3 plot respectively throughput curves related to the correct receptions of Lo-RaWAN frames by at least one and three gateways in the considered honeycomb deployment. As expected, the maximum throughput related to radiolocation applications is achieved for smaller values of the density μ of end-devices per area. This is due to the fact that the probability that at least 3 gateway receive a frame is lower than the one related to the same frame being received by at least 1 gateway. Each of the groups of lines drawn with the same color in Fig. 3 includes 3 plots. Focusing on the throughput related to the correct reception by at least one gateway (blue lines), the dotted line with full circle markers denotes the statistics Δ collected according to eq. (3) and related to simula-



tions run without any duty cycle limitation: it is completely overlapped to the dotted line with circles, that in turn plots the function Γ given by eq. (2). Instead, the thick and continuous line plots the same statistics Δ in the case of application of a duty cycle limitation: it results as stretched along the abscissa, since end-devices will drop some generated packets from their own buffer according to the M/D/1/2 policy, thus reducing the resulting per-device transmission rate.

In this sense, it is worth evaluating also the measured loss rate. Fig. 4 shows black lines related to the measured ratio of frames generated but dropped before being transmitted on the radio. Specifically, the dotted line with thick full circles lying close to the abscissa plots the average number of frames dropped when no duty cycle limitation applies: the values are not zero, since there is the probability that more than 2 frames are generated on an end-device before the end of an ongoing transmission is not null, thus originating very sporadic frame drops due to the M/D/1/2 queuing policy implemented. The thick and continuous black line plots the ratio of packets drop when the duty cycle limitation of 1% on the traffic transmission applies. The two lines do not depend upon the end-device den-



sity μ , because the drop probability is just related to the buffer size of each end-device and is uncorrelated with the number of end-devices. Furthermore, Fig. 4 plots the overall ratio of frames that result useless for monitoring (blue) and radiolocation (red) duties. As matter of fact, frames that result useless for radiolocation purposes because correctly decoded by less than 3 gateways can still be useful for monitoring duties, and this can be noted by the fact that red lines are below blue ones. For each pair of lines of the same color, the dotted line with thick full circles and the continuous and thick line respectively denote the lack and presence of duty cycle limitation. All the lines show that the loss rate increases with increasing end-device density, since more offered traffic produces also more collisions. However, focusing on the blue lines (similar considerations can be done for the case of radiolocation) it is worth noting that the line plotting the loss rate in the case of implementation of a duty cycle policy intersects the line corresponding to the lack of any policy exactly for the same value of μ where the correspondent throughput lines intersect each other in Fig. 3: for higher end-device density, the duty cycle limitation permits higher throughput and less collisions.

5 CONCLUSIONS

LPWANs are going to further extend the scope of the IoT. Among such technologies, LoRaWAN networks enable bidirectional information exchanges over long range communications. Some built-in architectural prop-

erties make LoRaWAN exploitable for enhanced monitoring applications, including radiolocation, geofencing and low-power tracking. This paper timely introduces a probability model able to describe scalable LoRaWAN deployments. To validate the envisaged performances, a Python simulator has been developed and described. It mimics LoRaWAN city scaled deployments, while instrumenting researchers and practitioners with a wide range of parameter settings.

Some preliminary simulation results show that the theoretical model for LoRaWAN throughput in multigateway multi-channel deployments is validated and that realistic LoRaWAN duty cycle policies produce a predictable distortion in the perceived performances. Future works will tackle more advanced features of the network, including mobile nodes and radio interference from other technologies. These results will be validated in real scenarios.

ACKNOWLEDGMENTS

This work has been co-funded by a FEDER-FSE 2014-2020 fund of the Région Midi-Pyrénées and the European Union and also by French Government (program: investment for future) in the project: Smart Services for Connected vehiCles – S2C2.

REFERENCES

- Abramson, N. (1970). THE ALOHA SYSTEM: another alternative for computer communications. In *Proceedings of the November 17-19, 1970, fall joint computer conference*, pp. 281–285. ACM.
- Accettura, N., S. Medjiah, B. Prabhu, and T. Monteil (2017, October). Low Power Radiolocation through Long Range Wide Area Networks: a performance study. In Proc. IEEE International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob 2017).
- Baccelli, F. and B. Błaszczyszyn (2010). Stochastic geometry and wireless networks: Volume i theory. *Foundations and Trends in Networking* 3(34), 249–449.
- Bajović, D., D. Jakovetić, D. Vukobratović, and V. Crnojević (2014, June). Slotted Aloha for networked base stations. In *Proc. IEEE International Conference on Communications Workshops (ICC 2014)*, pp. 520–526.

Bertsekas, D. and R. Gallager (1992). Data Networks, Second Edition. Prentice Hall.

- ETSI (2015, May). ETSI EN 300 220-2 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 2: Harmonized EN covering essential requirements under article 3.2 of the R&TTE Directive. European Telecommunications Standards Institute (ETSI).
- IEEE 802.15.4 Standard (2016, April). IEEE Standard for Low-Rate Wireless Networks. *IEEE Std* 802.15.4-2015 (*Revision of IEEE Std* 802.15.4-2011), 1–709.
- IETF lpwan (2016). IPv6 over Low Power Wide-Area Networks (lpwan) Working Group.
- LoRaWAN regional parameters (2016). LoRaWAN Regional Parameters v1.0.
- Palattella, M., N. Accettura, X. Vilajosana, T. Watteyne, L. A. Grieco, G. Boggia, and M. Dohler (2012). Standardized Protocol Stack for the Internet of (Important) Things. *IEEE Communications Surveys and Tutorials* 15(3), 1389–1406.
- Petäjäjärvi, J., K. Mikhaylov, A. Roivainen, T. Hanninen, and M. Pettissalo (2015, December). On the coverage of LPWANs: range evaluation and channel attenuation model for LoRa technology. In *ITS Telecommunications (ITST), 2015 14th International Conference on*, pp. 55–59.
- Roberts, L. G. (1975, April). ALOHA Packet System with and Without Slots and Capture. SIGCOMM Comput. Commun. Rev. 5(2), 28–42.

Semtech specification (2013, July). SXI272/3/6/7/8 LoRa Modem Design Guide AN1200.13 Revision 1.

- Sornin, N., M. Luis, T. Eirich, T. Kramp, and O. Hersent (2016). LoRaWAN Specification v1.0.2.
- Warneke, B., M. Last, B. Liebowitz, and K. S. J. Pister (2001, January). Smart Dust: communicating with a cubic-millimeter computer. *Computer* 34(1), 44–51.