

Evaluation of LoRa Performance in a City-wide Testbed: Experimentation Insights and Findings

Giannis Kazdaridis^{*†}, Stratos Keranidis[†], Polychronis Symeonidis[†],
Panagiotis Tzimotoudis[†], Ioannis Zographopoulos[†], Panagiotis Skrimponis[†]
and Thanasis Korakis^{*†}

[†]Department of Electrical and Computer Engineering, University of Thessaly, Greece

^{*}Centre for Research and Technology Hellas, CERTH, Greece

{iokazdarid,efkerani,posymeon,tzimotou,zografop,skrimpon,korakis}@uth.gr

ABSTRACT

In this paper, we present a LoRa based city-scale testbed that employs several sensing devices scattered across the urban area to characterize air quality and weather conditions in real-time. The installation is augmented through a custom link quality evaluation framework that continuously monitors the Packet Delivery Ratio versus RSSI relation to characterize the performance of LoRa standard under realistic conditions. Experimental results, collected over a period of 2 months, efficiently analyze LoRa's performance across a wide range of protocol configurations. Finally, we also present in-lab experiments that characterize the efficiency of LoRa modules in terms of power and energy efficiency per bit, along with valuable insights aimed at the development of energy efficient protocol improvements.

KEYWORDS

LoRa Technology, Testbed Experimentation, LoRa Performance Evaluation, Energy Efficiency

1 INTRODUCTION

Air pollution is a crucial environmental situation that drastically impacts on public health, global environment, and worldwide economy. The World Health Organization confirms that air pollution is now the world's largest single

The research leading to these results has received funding by GSRT, under the act of "HELIX-National Infrastructures for Research", MIS No 5002781.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

WiNTECH'19, October 25, 2019, Los Cabos, Mexico

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6931-2/19/10...\$15.00

<https://doi.org/10.1145/3349623.3355474>

environmental health risk, responsible for numerous premature deaths [1], while long-term exposition to atmospheric pollutants cause health disorders [2]. All the above highlight the urgent need for air pollution monitoring facilities, towards providing meaningful data to environmental protection agencies to better understand the effects and protect the public. Latest advancements in embedded system technologies, facilitate the development of low-cost and small-size devices, capable of addressing the challenge for large-scale, dense deployments to monitor the pattern and presence of harmful pollutants. To date, several research institutes have deployed pilot city-scale facilities to effectively provide real-time measurements of the atmospheric pollutants [3], [4].

Evidently, a major concern in such installations is the backbone connectivity, as it has to cope with the large volume and the distribution of the deployed sensors. The recently introduced IoT Low-Power WAN (LPWAN) technologies, such as NB-IoT, LoRa, LoRaWAN and SigFox appear to be attractive candidates to deal with the nature of the aforementioned scenario. All these technologies use a physical layer that trades throughput performance, supporting from hundreds of bits to a few kilobits per second, to provide long range communication in the order of tens of kilometers. LoRa [5] is one prominent technology of this type, employing *Chirp Spread Spectrum (CSS)* modulation. This standard allows the configuration of four critical parameters, *Bandwidth (BW)*, *Spreading Factor (SF)*, *Coding Rate (CR)* and *Transmit Power (TX_P)*. Despite the huge potential that the LoRa standard exhibits, it has been insufficiently researched and studied. Several works [6–8], study the performance of LoRa and LoRaWAN technologies, providing some first results in the behavior of this technology.

In this paper we present a custom city-scale monitoring installation, reporting air quality and weather conditions in real-time. The framework consists of sensing devices that are distributed across the city area of Volos, Greece, while they employ the LoRa standard to propagate their measurements. The key contributions of our work follow:

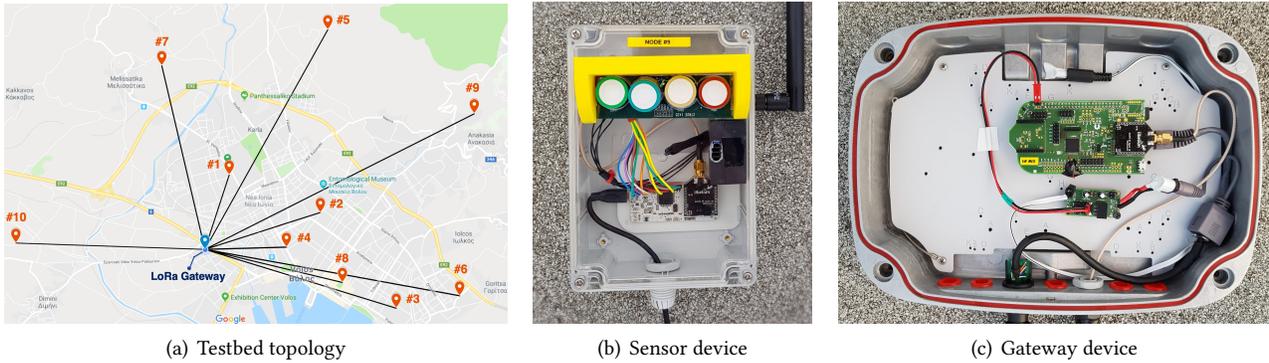


Figure 1: Testbed topology and devices

- development of a real-time air quality and weather conditions evaluation testbed setup in city-scale, employing custom sensor nodes and gateway devices
- thorough analysis of PDR versus RSSI performance exploiting measurements over a period of 2 months
- detailed in-lab measurements characterizing the energy efficiency performance of the LoRa chipset, across a wide range of protocol configurations along with meaningful insights and findings

The remainder of the paper is organized as follows. Section 2 reviews related work. System Implementation and developed components are described in Section 3, along with Experimentation Parameters configuration in Section 4. Section 5 characterizes the obtained power consumption results, while experimentation evaluation is demonstrated in Section 6. Finally, Section 7 concludes the paper.

2 RELATED WORK

In the domain of LoRa performance as a communication protocol, several existing works [9–13] provide measurements regarding the coverage the capacity, the delay and the throughput performance of LoRa protocol. In [6] authors measure the performance of LoRa protocol in an urban environment, highlighting its variable behavior delivering a packet error rate between 3 and 90 %, even when using high transmission power settings. [14] illustrates an experimental study of the impact of the LoRa transmission parameters selection on the communication performance. An indoor setup of two nodes is used along with a developed link probing scheme to evaluate the optimal parameters for the packet transmissions. However, the setup used can not lead to safe conclusions since it is based only on two nodes located in a specific environment.

Considering testbed platforms, [8, 15] provide framework tools to evaluate LoRa technology. Specifically, *OpenChirp* [15], demonstrates custom LoRa-enabled sensing nodes and Gateways that utilize the Semtech’s *SX1276* chipset. The firmware is built on top of the open-source *TI-RTOS* and users

can register and define the parameters they wish to follow. Similarly, *WiSH-WaIT* [8], is a framework for controllable and reproducible experiments. A custom interface allows the parameterization of LoRa settings, such as *TX*, *SF*, *BW* and *Error Coding Rate*. The sensor nodes feature again the *SX1276* and the paper also includes findings regarding the correlation of PDR versus the spreading factor used as well as *PDR* versus *TX_p*.

We remark the work in [7] that provides several useful findings based on realistic experimentation. Authors proved that LoRa is capable of communicating over 10 km links in *LOS* conditions, while in *NLOS* is severely affected by obstructions and vegetation. Moreover, authors quantify the importance of the *SF* parameter over the rest ones and urges the selection of the lowest possible *SF* whenever possible. In our work, we employ a city-scale experimental setup to obtain meaningful data in order to characterize the performance of LoRa links. More specifically, we analyze the obtained results and correlate the obtained *PDR* with the *RSSI*, while we also suggest the optimal rate configurations to be used for each different link scenario.

In the field of energy consumption characterization of LoRa protocol, there are several works [14, 16] that are based on data-sheet power data, while some other [7, 15, 17] employ measurement tools to obtain power consumption. All works aim at modeling the energy performance of LoRa to characterize the lifetime of the under consideration setups. We distinguish the work in [7], in which authors correlate the obtained power consumption data versus the selected transmission power setting and estimate the lifetime of four different nodes that utilize the Semtech’s *SX1276* and the *RFM96* LoRa chipsets. In our work we focus on detailing the instantaneous current draw for each possible combination, providing also meaningful findings, while we also calculate the energy efficiency per bit in both transmission and reception phase, across different payload sizes to aid the development of energy efficient rate adaptation algorithms.

ID	GW Dist. (km)	Altitude (m)	LOS/NLOS	Sensors	Bytes
GW	-	23	-	-	-
1	1.3	43	NLOS	3	24
2	1.98	30	NLOS	3	24
3	3.25	32	NLOS	1	14
4	1.33	25	NLOS	1	14
5	4.12	179	LOS	3	24
6	4.22	55	NLOS	3	24
7	3.11	128	LOS	7	44
8	2.32	35	NLOS	7	44
9	4.92	170	LOS	7	44
10	3.10	8	NLOS	3	24

Table 1: Testbed Node Characteristics

3 SYSTEM ARCHITECTURE & IMPLEMENTATION

In this section, we present the components and the architecture of our testbed that consists of one LoRa Gateway (*GW*) that is deployed on the rooftop of our University's premises and ten Sensing nodes that are scattered across the city of Volos, Greece. The end nodes feature various air quality probes, characterizing gas concentrations, dust particle concentration, along with outdoor air temperature and humidity modules. Fig. 1(a) illustrates the testbed topology, while the characteristics of each individual sensor device are presented in Table 1. We clearly observe that the testbed offers a wide range of wireless links that differentiate in terms of communication distance, elevation difference and *LOS/NLOS* conditions. For each node, we also note the amount of payload bytes transmitted when transferring the respective number of sensed parameters, which includes 5 Bytes per environmental parameter plus 9 Bytes of own protocol overhead, while it does not include the LoRa *preamble*, *header* and *CRC*, as specified in [5]. Below we present the custom developed *GW* and sensing devices, as well as the software toolkit for obtaining LoRa performance measurements.

LoRa Gateway: The *GW* is based on the BeagleBone Black Rev. C board [18] which is a low-cost, embedded platform characterized by sufficient processing power capabilities (*ARM Cortex-A8 Processor at 1 GHz with 512 MB RAM*) and several I/O pins, while it also embeds an Ethernet port for communicating with the backbone network. On top of BeagleBone we attach a custom-designed PCB board (cape) that integrates our electronic components and features a slot for connecting a LoRa transceiver. More specifically, the cape features the MK20DX128VLH5 micro-processor which is interfaced with the LoRa transceiver responsible to run all the requisite libraries for its operation. The micro-processor software is also responsible of implementing a polling mechanism to acquire measurements from the sensor nodes and also initiate the link quality experiments. For the LoRa transceiver we employ the SX1272 [19] chipset manufactured by Semtech, which is paired with a 4.5 dBi antenna.

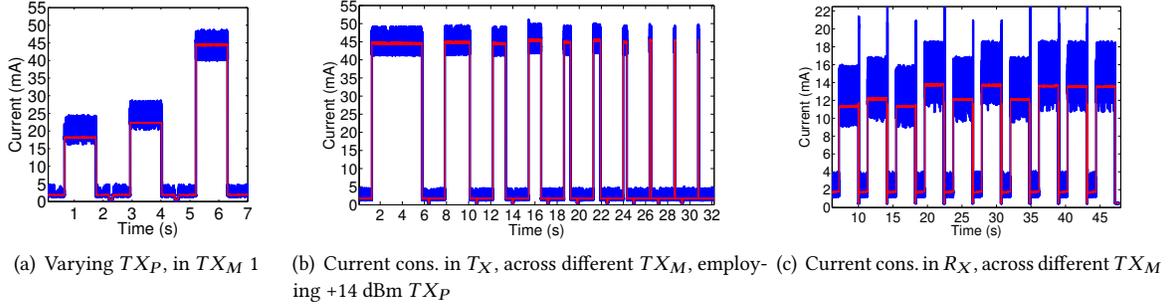
TX_M	BW	SF	Data Rate (bps)	Sensitivity (dBm)	Airtime (msec)
1	125	12	293	-137	1155.07
2	250	12	586	-135	577.53
3	125	10	977	-133	329.72
4	500	12	1172	-129	288.76
5	250	10	1954	-130	164.86
6	500	11	2149	-128	164.86
7	250	9	3516	-128	82.43
8	500	9	7032	-122	41.21
9	500	8	12500	-119	23.16
10	500	7	21875	-116	11.58

Table 2: LoRa Protocol specifications

Fig. 1(c) illustrates the developed gateway device deployed inside a waterproof aluminum case.

LoRa Sensing Devices: The core module of our sensing devices is again the MK20DX128VLH5 [20]. The 32-bit *ARM Cortex-M4 CPU*, clocks up to 96 MHz and supports ultra-low power sleeping modes to allow for power efficient operations. The node employs the same LoRa interface (*SX1272*) as the gateway device, attached to a 4.5 dBi omni-directional antenna. Each node is equipped with a varying set of sensors, as illustrated in Table 1, featuring 1, 3 or 7 sensors in total. The sensors utilized are gas concentrations probes (*NO2, SO2, CO, NH3, O3*), dust particle concentration *PM2.5, PM10* and a temperature & humidity module. The micro-processor hosts a program that implements the sensor reading and collection of measurements, and also the communication between itself and the *SX1272* module. Fig.1(b) presents one sensor device attached with four gas concentration probes and a dust particle module.

Software Toolkit: Through the deployed list of devices, we continuously collect detailed air quality measurements, while also evaluate the performance of the employed LoRa communication links. The overall procedure is orchestrated at the *GW* level, which periodically polls the different sensing nodes that continuously remain in receive mode. Sensor polling is executed in a Round-robin fashion, instructing sensors to transmit collected air quality measurements (with an interval of 10 minutes) or to participate in a link quality evaluation experiment (executed twice per day). The two procedures are executed independently of each other, taking advantage of the employed polling approach, which enables us to avoid the impact of overlapping transmission interference, when performing link quality experiments. Upon receiving a request for experiment, each sensing node transmits 10 consecutive packets at the *data rate*, *payload length* and *transmission power* that has been determined by the *GW*. The gateway sends the experiment settings embedded to the experiment request packet. The per-node link quality of the network is determined through the *Packet Delivery Ratio (PDR)* of the aforementioned experiments.


Figure 2: Obtained Current Consumption measurements

4 EXPERIMENTATION PARAMETERS

In this section we detail the procedure we are following to evaluate the performance of LoRa links. In essence, we are testing, adopting a probing fashion, all possible LoRa configurations to generate a rich set of measurements. More specifically, we vary the *Transmission Power* (TX_P), the *PHY* layer *data rate* and the frame *payload* size at each transmitter node. Starting with the (TX_P), we configure three different levels (0, +7, +14 dBm). Through this configuration, we result in 30 different link combinations, as illustrated in terms of uplink RSSI in Fig. 7. Considering the parameters affecting the resulting *PHY* layer *data rate*, we configure the full list of available settings of (BW , SF and CR), by using the list of combinations presented in Table 2. For simplicity, we refer to these combinations as *Transmission Modes* (TX_M) from 1 to 10. We specifically use the aforementioned list of combinations, considering Semtech’s recommendation of using SF settings of $SF7$ to $SF12$ and BW 125, 250, and 500 kHz [19], so as to ensure acceptable transmission distance and data rate trade-off, since both SF and BW affect propagation duration and data rate. Notably, in all combinations the CR option is set to 4/5. Additionally, Table 2 lists the resulting sensitivity and indicative airtime for each TX_M , using a reference payload size of 10 Bytes. The specified *data rate* combinations in Table 2 are also followed by the Libelium [21] as well as by the [22] if consider only CR of 4/5. Finally, regarding the varying payload sizes, we configure the values of 14, 24 and 44 Bytes, which correspond to the actual number of payload bytes transmitted when collecting air quality measurements from 1, 3 and 7 air sensors accordingly. In addition, we also use the payload size of 255 Bytes for stress testing the LoRa protocol, as this is the maximum supported value.

5 LORA POWER CONSUMPTION EVALUATION

In this section, we characterize the power consumption profile of the LoRa SX1272 chipset, under various protocol configurations. The experiments are conducted in the lab, employing a set of a transmitter and a receiver. To measure

TX_M	TX_P (mA) +0 dBm	TX_P (mA) +7 dBm	TX_P (mA) +14 dBm	RX_P (mA)
1	18.96	22.13	44.64	11.24
2	19.27	22.41	44.93	12.26
3	18.96	22.13	44.64	11.24
4	19.84	23.02	45.51	13.86
5	19.23	22.46	44.97	12.26
6	19.84	23.01	45.54	13.86
7	19.20	22.45	44.93	12.26
8	19.78	22.99	45.51	13.86
9	19.72	22.90	45.46	13.86
10	19.72	22.90	45.47	13.86

Table 3: Instantaneous Power Draw for SX1272 chipset

power draw, we place a small shunt resistor (1Ω) between the sensor’s regulated power supply (3.3 Volts) and the SX1272’s input supply pin. The voltage drop across this resistor is proportional to the current draw of the LoRa interface, which we measure with our high-end power consumption monitoring tool, presented in [23, 24]. The measurements were obtained with 100 MHz sampling speed at resolution of 16 Bits that provides the required level of detail. To characterize specific events, such as packet transmission or reception, we isolate the respective part of the measurements and calculate the average value.

5.1 Instantaneous current consumption

In the first experiment, we characterize the instantaneous current consumption profile of the SX1272 LoRa chipset, during frame transmission and reception under various TX_P and TX_M settings. Specifically, we configure the LoRa chipset in all available configurations, firstly in transmission mode and then in reception mode and we isolate the respective events in order to measure their instantaneous current draw.

The results of the first experiment are illustrated in Fig. 2(a) where the TX_P increases from 0 dBm to +7 dBm and +14 dBm, while using the fixed TX_M setting of the first rate configuration and a payload size of 10 Bytes. Similarly, in Fig. 2(b) we set the TX_P to +14 dBm and vary the TX_M to all the available rate configurations from 1 to 10. In this plot we use a fixed

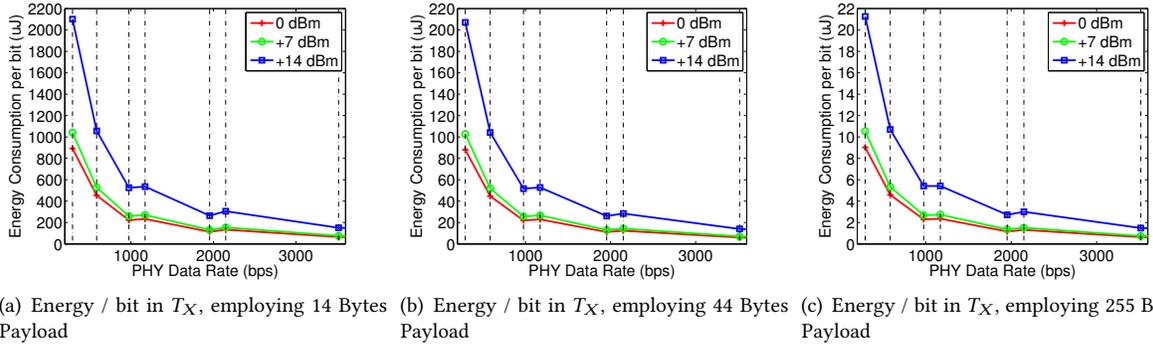


Figure 3: Energy Efficiency per bit measurements in T_X mode across different Payload sizes

payload size of 100 Bytes and we clearly observe that the relation between the duration of each frame propagated and the transmission speed of each TX_M is inversely proportional as expected. We repeat the same experiment but this time we configure the LoRa chipset as a receiver and we measure the instantaneous power draw, as illustrated in Fig. 2(c). We can observe the variation in instantaneous power consumption in each different TX_M . Notably, the red line illustrated in Fig. 2(a), Fig. 2(b) and Fig. 2(c) represents the average power draw calculated as a moving mean over a window of 100 values in an effort to provide a more representative indication to derive safe results, since the obtained measurements present high variations. Finally, the sudden drops from idle state to zero between each packet transmission or other spikes illustrated are related to the re-configuration of the LoRa parameters executed prior to each transmission or reception.

The analytic results across all tested parameters are listed in Table 3. We clearly observe that switching from the $+14$ dBm TX_P (44.64 mA) to the $+7$ dBm TX_P (22.13 mA), we can reduce current consumption by 50.43 %, while marginal decrease of 14.33 % is obtained when switching from $+7$ dBm to 0 dBm TX_P . Both observations are made for the first TX_M configuration. Considering the whole set of TX_M configurations the reduction varies at the same levels, specifically in the case of switching from the $+14$ dBm to the $+7$ dBm TX_P it varies from 50.43 % (TX_M : 1) to 49.42 % (TX_M : 4), while in the case of $+7$ dBm to 0 dBm TX_P we obtain a variation from 14.45 % (TX_M : 7) to 13.78 % (TX_M : 6). It is worth noting that the differentiation among the different TX_P levels, stems from the fact that a Power Amplifier (PA) circuit block is engaged in order to amplify the propagated signal, which in some cases drains the most of the total power [25]. The latter, suggests that it is preferable to use $+7$ dBm TX_P instead of 0 dBm, to achieve sufficiently higher PDR at almost the same energy cost in scenarios where required higher link budget.

The obtained results also highlight the impact of TX_M parameter in power consumption under the same TX_P . Relating consumption data with BW and SF parameters of Table 2, it

is made clear that the BW parameter plays significant role in the resulting draw, while SF configurations merely impact the obtained current consumption. We observe that TX_M configurations with fixed BW present roughly stable power draw. For instance when employing BW of 125 KHz (which is realized in TX_M 1 and 3) the power consumption remains exactly the same under the same TX_P setting. Whilst, in the rest BW configurations of 250 KHz and 500 KHz the power consumption variation under the same BW and TX_P settings does never exceed 1 %, which further highlights that the BW parameter totally affects the resulted power draw, while SF setting plays minor role. Notably, the highest obtained draw variation under the same TX_P is 4.47 % for the 0 dBm TX_P , 4.01 % for the $+7$ dBm and 2.03 % for the $+14$ dBm. It is worth noting that the adjacent TX_M 5 and 6 present significant draw variation despite the fact they present marginal difference in their attained data rates, 1.954 kbps and 2.149 kbps respectively. Specifically, they present 3.08 % current increase.

Identical findings can be drawn when using reception mode (RX_P), where the impact of BW is even more highlighted. More specifically, during reception, we note a decrease of 11.54 % when switching from 500 KHz BW (TX_M : 4) to 250 KHz (TX_M : 2) and the remarkable variation of 18.9 % when switching from 500 KHz (TX_M : 4) to 125 KHz (TX_M : 1). Notably, the same variations in TX_M during transmission in the max TX_P provides marginal power consumption decrease of up to 2.03%. This clearly shows that the power draw of the reception mode, is even more affected by the employed BW configuration, which is also noticeable by Fig. 2(b).

5.2 Energy Efficiency per Bit

As frame transmission duration is monotonically related to the configured data rate, it is important to quantify energy efficiency in terms of energy consumption per transmitted bit of information (E_B), as analyzed in our previous works [26, 27], which focus on analyzing the efficiency of the WiFi protocol. We calculate E_B , expressed in *Joules/bit*, as the

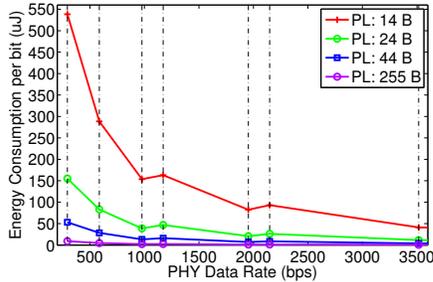


Figure 4: Energy Efficiency per bit measurements in R_X mode across different Payload sizes

division of the power consumption values collected for each different TX_M and TX_P by the exact *data rate* value.

In Fig. 3, we plot the obtained transmission E_B across the available *data rate* configurations and zoom in on the first 7 TX_M settings. We plot the E_B for the payload sizes of 14, 44 and 255 Bytes, while we omit the case of 24 Bytes for space reasons. Notably, the dashed lines on the plots represent the respective *data rates* of each TX_M . We notice that higher *data rate* settings always result in lower E_B , which is mainly due to the decreased duration of frame transmissions. The only two cases that present different behavior are the adjacent pairs 3, 4 and 5, 6. We observe that the calculated E_B in TX_M 4 is marginally higher than that of TX_M 3, which is also true for TX_M 6 and TX_M 5. This happens since, TX_M 4 and 6 present higher draw than 3 and 5 respectively, while the achieved *data rates* are nearly equal. This finding suggests skipping TX_M 3 and 5, and proceeding to TX_M 4 and 6 respectively, when considering energy efficient rate adaptation algorithms. Moreover, the highest observed variation is up to 15.12 % when comparing the pair of TX_M 5 and 6, realized in the +14 dBm setting when transmitting packets of 14 Bytes payload and 9.98 % and 11 % for the cases of 44 and 255 Bytes respectively. Regarding the comparison of the calculated E_B among the illustrated payload sizes we observe an extreme variation. More specifically, when comparing the E_B of 14 Bytes payload with the one of 44 Bytes under the TX_M 1 the reduction is 90.14 %. Whilst, the difference is even bigger when comparing the 14 Bytes with the 255 Bytes under the TX_M 1, with reduction reaching up to 98.99 %.

Similarly, we plot the E_B under varying payload sizes in the reception phase, in Fig. 4, by zooming in on the first 7 TX_M . The obtained measurements present identical behavior for TX_M 3, 4 and 5, 6 as in the transmission phase. To note that the variation of the E_B for the pairs TX_M 3, 4 and 5, 6 is even higher than the ones in the transmission phase. Namely, it is 27.33 %, 19.17 % and 11.35 % for the cases of 14, 44 and 255 Bytes respectively, when considering the pair TX_M 5 and 6. Thus, we further highlight the urge to omit the TX_M 3 and 5, and select the TX_M 4 and 6 when designing power efficient rate adaptation algorithms. Direct comparison of the E_B among payload sizes of 14, 24 and 44 Bytes,

during reception, presents remarkable energy efficiency improvement of 71 % and 90 % respectively. The variation is even higher when comparing payloads of 14 and 255 Bytes, reaching up to 99.78 %. This finding, highlights the fact that measurement aggregation techniques can achieve extremely improved energy efficiency in LoRa setups.

Assuming an energy experiment based on our deployed network, in which the sensor nodes propagate their measurements once per hour, we can estimate the lifetime of the network on typical batteries and compare it with the same experiment when employing aggregation mechanisms. In a scenario where a node is propagating 14 Bytes of information using the maximum TX_P setting, we can calculate that the LoRa chipset will consume roughly 140 mAh of the total power budget in a year, if we assume ideal conditions in sleep mode and wake-up operation. Considering a power budget of 400 mAh dedicated only for LoRa transmissions the module's lifetime will be approximately 2.87 years. When employing an aggregation mechanism, grouping several individual measurement frames into a single 100 Bytes packet the consumed power will be 62 mAh in a year, which suggests lifetime of roughly 6.38 years. The latter implies an increase of 122 % in lifetime, while aggregation using 255 Bytes frame will lead to an increase of 160 % in lifetime.

6 EXPERIMENTAL LORA PERFORMANCE EVALUATION

This section presents a thorough analysis of LoRa performance across a wide range of conducted testbed experiments, under uplink communication initiated from by each one of the 10 nodes and destined to the GW. The list of varying parameters (Table 2) corresponds to a total of 1200 experimental combinations that are conducted every hour, corresponding to 24K frames transmitted per day. The resulting large dataset includes measurements collected over a period of 2 months. In order to effectively visualize the wide set of collected measurements, we decided to group the node and TX_P combinations, in terms of the resulting RSSI per link, thus LoRa performance is evaluated in terms of PDR and spans the RSSI range between -102 and -137 dBm.

Figures 5(a) and 5(b) illustrate the PDR performance versus uplink RSSI and TX_M , across the 2 different tested payloads. The solid lines correspond to the different TX_M (Table 2), while the dashed lines refer to the corresponding sensitivity thresholds, as calculated employing payload of 10 bytes [19] compared to the ones calculated through our realistic experiments (14 Bytes). Considering the results across increasing payload sizes, we clearly see that the obtained sensitivity thresholds tend to increase (at maximum by 1 dB per step), which also verifies the validity of the experimental results. In the extreme cases employing the TX_M 1 and 10, we observe the remarkable performance of LoRa achieve approximately

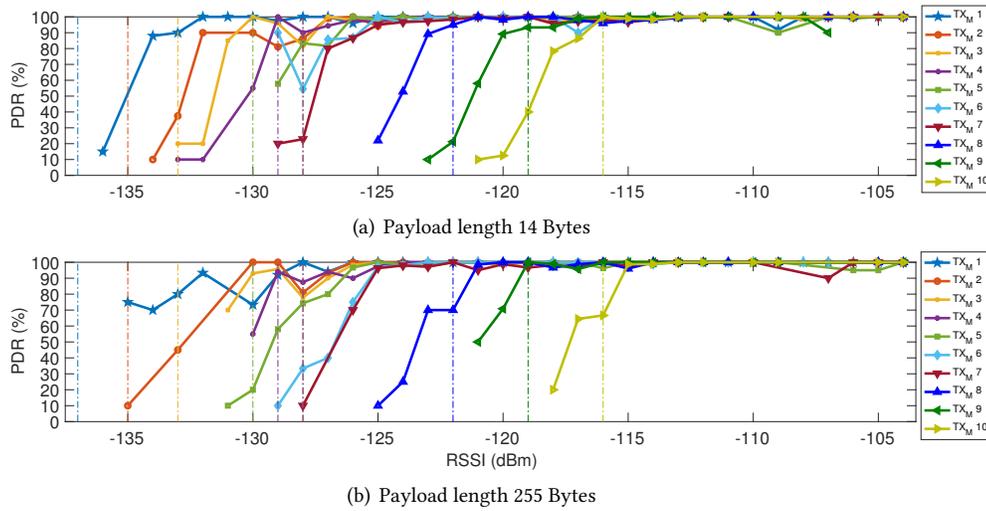


Figure 5: LoRa PDR performance versus uplink RSSI and TX_M

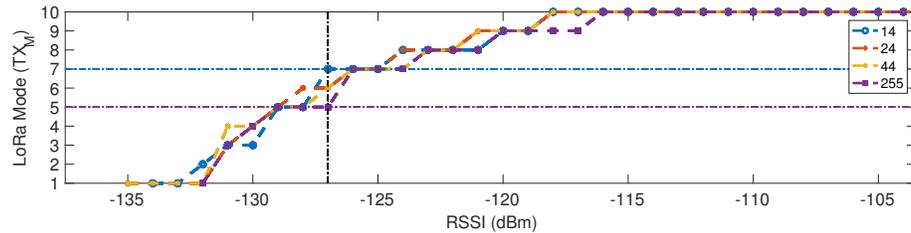


Figure 6: LoRa mode versus uplink RSSI and payload length

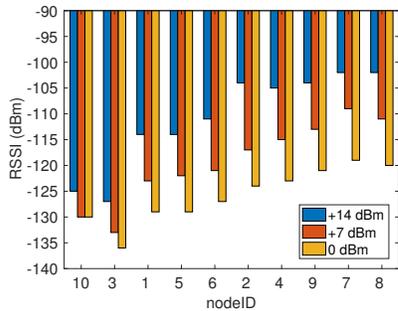


Figure 7: Uplink RSSI per LoRa link across different TX_P levels (0, +7, +14 dBm)

90% PDR at -134 dBm and -118 dBm RSSI for transmissions at TX_M 1 and 10 respectively, under realistic city-scale experiments. Focusing on the maximum payload case, we can observe that even -115 dBm LoRa links are able to transfer frames of 255 Bytes at TX_M 10 with approximately 100% PDR that is able to achieve 21.89 Kbps and can support the transferring of low-quality images or even video streams in real-time. Relating this insight with Fig. 7, we notice that 8 out of the 10 considered links are able to support such demanding applications in the maximum TX_P setting.

In the rest of this section, we attempt to define the optimal mapping of RSSI to TX_M , when considering the 30 different link scenarios (10 nodes in 3 different TX_P levels),

as depicted in Fig. 7. Following a heuristic approach commonly employed in rate adaptation mechanisms, we identify the best PHY layer setting, as the maximum setting able to maintain higher than 90% of PDR. Fig. 8 illustrates the mapping of TX_M to the 30 link types, considering the aforementioned rate selection principle. We clearly see that most of the tested links for the 0 dBm TX_P level, are able to maintain higher than 90% of PDR for TX_M higher than 4 for RSSI levels higher than -129 dBm. On the other hand, considering the maximum TX_P setting (+14 dBm), we notice that most links can maintain the TX_M of 7 for RSSI levels higher than -114 dBm. Only 6 out of 30 tested links, generated from nodes 3 and 10, correspond to link quality lower than -127 dBm and require the configuration of lower TX_M settings. The sum of collected results highlight the potential of LoRa to achieve stable communication performance for data rates much higher than the basic one, across a wide range of low quality links.

Finally, in Fig. 6, we illustrate the detailed mapping of RSSI to TX_M , considering the varying frame payload parameter as well. It is made clear that the transmission of larger frames, requires the increase of TX_M , by up to two levels, as in the case of -127 dBm, which corresponds to results collected from nodes 3 and 6 under TX_P accordingly. Concluding, we

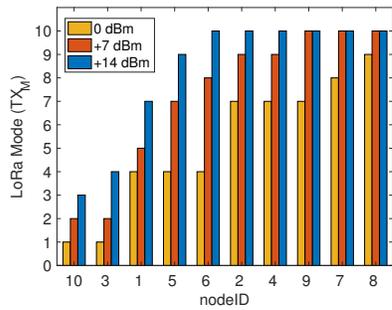


Figure 8: Minimum calculated LoRa mode per link, across different TX_p levels, for achieving at least 90% PDR

remark that the wide range of collected results highlights the importance of employing automated rate adaptation for LoRa communication that jointly considers the varying RSSI and payload parameters.

7 CONCLUSIONS & FUTURE WORK

In this paper we presented a city-scale monitoring LoRa-based infrastructure that provides real-time air and link quality indications. Furthermore, we characterize the energy efficiency of the utilized LoRa chipset by conducting in-depth lab experiments. We conclude that the LoRa standard is a very attractive option for deploying energy and cost efficient *IoT* applications in city-scale environments. Specifically, our findings indicate that rate adaptation algorithms and packet aggregation techniques can drastically benefit the energy efficiency of the LoRa protocol, when considering battery powered applications. Moreover, our city-scale experiments suggest that LoRa can support demanding applications achieving impressively high PDR at extremely low RSSI conditions, even when using high data rates. Our future directions include the expansion and the rearrangement of our testbed, so as to extend the range of collected data, specifically of low RSSI links and to characterize the performance of the setup across more advanced metrics, such as throughput and latency.

REFERENCES

- [1] World Health Organization. "7 million premature deaths annually linked to air pollution", Geneva, 2014.
- [2] David Mage, Guntis Ozolins, Peter Peterson, Anthony Webster, Rudi Orthofer, Veerle Vandeweerd, and Michael Gwynne. Urban air pollution in megacities of the world. *Atmospheric Environment. Supercities: Environment Quality and Sustainable Development*.
- [3] D. Hasenfratz, O. Saukh, C. Walser, C. Hueglin, M. Fierz, and L. Thiele. Pushing the spatio-temporal resolution limit of urban air pollution maps. In *2014 IEEE Int. Conf. on PerCom*, pages 69–77, March 2014.
- [4] Luis Sanchez, Luis Muoz, Jose Antonio Galache, Pablo Sotres, Juan R. Santana, Veronica Gutierrez, Rajiv Ramdhany, Alex Gluhak, Srdjan Krco, Evangelos Theodoridis, and Dennis Pfisterer. Smartsantander: Iot experimentation over a smart city testbed. *Computer Networks*, 61, 2014. Special issue on Future Internet Testbeds - Part I.
- [5] Lora Alliance. <https://goo.gl/DOR0CS>, [Accessed 27-June-2019].

- [6] T. Petri, M. Goessens, L. Nuaymi, L. Toutain, and A. Pelov. Measurements, performance and analysis of lora fabian, a real-world implementation of lpwan. In *2016 IEEE 27th Annual Int. Symp. on Personal, Indoor, and Mobile Radio Communications (PIMRC)*, 2016.
- [7] JC. Liando, A. Gamage, AW. Tengourtius, and M. Li. Known and unknown facts of lora: Experiences from a large-scale measurement study. *ACM Transactions on Sensor Networks (TOSN)*, 2019.
- [8] A. Duda, Q. Lone, E. Duble, F. Rousseau, I. Moerman, and S. Gianoulis. Wish-walt: A framework for controllable and reproducible lora testbeds. In *IEEE PIMRC*, 2018.
- [9] M. Bor, JE. Vidler, and U. Roedig. Lora for the internet of things. 2016.
- [10] Norbert Blenn and Fernando A. Kuipers. Lorawan in the wild: Measurements from the things network. *CoRR*, 2017.
- [11] A. Augustin, J. Yi, T. Clausen, and W. Townsley. A study of lora: Long range & low power networks for the internet of things. *Sensors*, 2016.
- [12] KE. Nolan, W. Guibene, and MY. Kelly. An evaluation of low power wide area network technologies for the internet of things. In *2016 IWCMC*, 2016.
- [13] J. Petjrvi, K. Mikhaylov, M. Pettissalo, J. Janhunen, and J. Iinatti. Performance of a low-power wide-area network based on lora technology: Doppler robustness, scalability, and coverage. *Int. Journal of Distributed Sensor Networks*, 2017.
- [14] M. Bor and U. Roedig. Lora transmission parameter selection. In *2017 13th Int. Conf. on DCOSS*, 2017.
- [15] A. Dongare, C. Hesling, K. Bhatia, A. Balanuta, RL. Pereira, B. Iannucci, and A. Rowe. Openchirp: A low-power wide-area networking architecture. In *2017 IEEE Int. Conf. on PerCom Workshops*, 2017.
- [16] MS. Mahmoud and AAH. Mohamad. A study of efficient power consumption wireless communication techniques/modules for internet of things (iot) applications. 2016.
- [17] L. Casals, B. Mir, R. Vidal, and C. Gomez. Modeling the energy performance of lorawan. *Sensors*, 2017.
- [18] BeagleBone Black. <https://goo.gl/jkfCcl>, [Accessed 27-June-2019].
- [19] SX1272/73 860 MHz to 1020 MHz Low Power Long Range Transceiver. Rev. 3.1 March 2017. <http://tiny.cc/tw8w8y>, [Accessed 27-June-2019].
- [20] Freescale MK20DX128VLH5 microprocessor. <https://goo.gl/eoT8xx>, [Accessed 27-June-2019].
- [21] LoRa SX1272 Networking Guide Libelium's. <http://tinyurl.com/y6nrb877>, [Accessed 27-June-2019].
- [22] M. Cattani, C. Boano, and K. Rmer. An experimental evaluation of the reliability of lora long-range low-power wireless communication. *Journal of Sensor and Actuator Networks*, 2017.
- [23] G. Kazdaridis, S. Keranidis, P. Symeonidis, P. S. Dias, P. Gongalves, B. Loureiro, P. Gjanci, and C. Petrioli. Everun: Enabling power consumption monitoring in underwater networking platforms. In *Proc. of the 11th Workshop on WiNTECH*, pages 83–90. ACM, 2017.
- [24] G. Kazdaridis, I. Zographopoulos, P. Symeonidis, P. Skrimponis, T. Korakis, and L. Tassiulas. In-situ power consumption meter for sensor networks supporting extreme dynamic range. In *Proc. of the 11th Workshop on WiNTECH*, 2017.
- [25] R. Khereddine, L. Abdallah, E. Simeu, S. Mir, and F. Cenni. Adaptive logical control of rfin performances for efficient energy consumption. In *IFIP/IEEE Int. Conf. on Very Large Scale Integration-System on a Chip*. Springer, 2010.
- [26] S. Keranidis, G. Kazdaridis, N. Makris, T. Korakis, I. Koutsopoulos, and L. Tassiulas. Experimental evaluation and comparative study on energy efficiency of the evolving ieee 802.11 standards. In *Proceedings of the 5th Int. Conf. on Future Energy Systems*, e-Energy '14, 2014.
- [27] S. Keranidis, G. Kazdaridis, V. Passas, T. Korakis, I. Koutsopoulos, and L. Tassiulas. Online energy consumption monitoring of wireless testbed infrastructure through the nitos emf framework. In *Proc. of the 8th ACM int. workshop on WiNTECH*. ACM, 2013.