

Empirical Analysis of LoRaWAN Adaptive Data Rate for Mobile Internet of Things Applications

Konstantinos Kousias
Simula Research Laboratory,
Norway
kostas@simula.no

Giuseppe Caso,
Özgü Alay
Simula Metropolitan Center for
Digital Engineering, Norway
{giuseppe,ozgu}@simula.no

Filip Lemic
Internet and Data Lab, University
of Antwerp - imec, Belgium
filip.lemic@uantwerpen.be

ABSTRACT

Built on top of the Long Range (LoRa) physical layer, the LoRa Wide-Area Network (LoRaWAN) protocol has recently emerged as one of the most promising Low-Power Wide-Area Network (LPWAN) technologies, for several Internet of Things (IoT) applications. LoRaWAN introduces the Adaptive Data Rate (ADR) mechanism, aiming to deliver a fair compromise between network performance and system reliability. ADR performs adaptive tuning of communication parameters, e.g., the Spreading Factor (SF), which is used to modulate the transmitted signals. Although the performance of ADR has been explored in conjunction with stationary End-Devices (EDs), little is known about its suitability for mobile IoT applications. In this paper, we investigate the performance of ADR in diverse mobility scenarios by leveraging a large amount of LoRaWAN experimental traces, collected in the urban area of Antwerp, Belgium. Using a data-driven statistical approach, we show that, whilst ADR enhances network reliability and coverage in low mobility settings, its beneficial effects decrease as mobility increases, hence calling for possible improvement and optimization.

CCS CONCEPTS

• **Networks** → **Network performance analysis**; *Mobile networks*; Network mobility.

KEYWORDS

Mobile Internet of Things, Low-Power Wide-Area Networks, Long Range Wide-Area Networks, Adaptive Data Rate

1 INTRODUCTION AND MOTIVATION

Low-Power Wide-Area Networks (LPWANs)¹ are emerging as the norm technology for Internet of Things (IoT) applications requiring long-range, energy-efficient, discontinuous, and low data rate communications [8]. According to the

¹LPWANs can be divided in: I) 3GPP-standards, working in the licensed cellular spectrum, such as Narrowband Internet of Things (NB-IoT) and Long Term Evolution for Machines (LTE-M), and II) technologies operating in the unlicensed spectrum, such as SigFox and Long Range (LoRa), which exploits the LoRa Wide-Area Network (LoRaWAN) protocol stack [4].

Cisco Global Mobile Data Traffic Forecast, LPWAN connections will constitute 14% of wireless connections by the end of 2022². IoT applications such as healthcare, industrial automation, and environmental monitoring, incorporate mobility aspects, in which IoT devices are carried by humans or embedded onto mobile objects [3, 5]. To which extent LPWANs fit mobile IoT applications is not very well explored, to the best of our knowledge, and requires further analysis.

LoRa operates in the 433 and 868 MHz bands in Europe, adopts 125 kHz channels and complies with the European regulations, which limit the emissions by enforcing a 1% duty cycle. The signals are modulated using a Chirp Spread Spectrum (CSS), where the Spreading Factors (SFs) ($SF \in [7, 12]$) indicate the chirp duration. SFs leverage the trade-off between transmission rate, reliability, and coverage, with higher SFs leading to lower data rates but increased coverage. LoRaWAN provides Medium Access Control (MAC) and higher layers functionalities. It works in a star topology, with End-Devices (EDs) communicating with LoRaWAN gateways, which are IP-connected to a central server. Uplink messages can be received by several gateways and repeatedly forwarded to the server.

LoRaWAN introduces Adaptive Data Rate (ADR), a mechanism that aims to optimize performance by dynamically tuning EDs transmission parameters, such as SFs and transmission power. The ED runs device-side ADR, but can also enable network-assisted ADR, in which the server rules the parameters to be adopted in future uplink transmissions. In a nutshell, ADR triggers I) a SF increase (device-side), if the ED is in low coverage, to improve the message delivery in harsh radio conditions, and II) a SF decrease (network-side), if the ED is in good coverage, leading to higher data rates and lower power consumption [6].

Several studies have explored the performance of ADR, mostly by simulations; variations and enhancements of the main algorithm have been also recently proposed [1, 9]. However, extensive analysis in real-world LoRaWAN deployments is currently lacking, including the study of ADR

²<https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white-paper-c11-738429.html>.

performance under mobility scenarios. Besides qualitative recommendations from The Things Network (TTN)³, stating that “ADR should be enabled whenever an end device has sufficiently stable RF conditions”, a thorough quantitative analysis is still missing. An exception can be found in [7], where the impact of mobility on a LoRaWAN ADR-enabled single link is analyzed in both indoor and outdoor scenarios. Results reveal a negative impact in terms of packet loss and delay, but also call for further exploration, given that a single-link deployment cannot address a network-wide analysis.

In this paper, we present results obtained by analyzing the device-side ADR mechanism in a large-scale LoRaWAN deployment in Antwerp, Belgium. In particular, we reveal that under low mobility scenarios, the SF increase leads to increased coverage, hence satisfying its objective; as mobility increases, we observe less beneficial effects, calling for possible advances towards an optimized version.

2 EXPERIMENTAL DESIGN

To disclose the impact of mobility on the ADR performance, we follow a data-driven approach.

Experimental Setup: Several LoRa EDs were mounted on postal trucks, executing fixed routes in the city of Antwerp, Belgium [2]. Each ED was programmed to transmit a LoRa message of either 46 or 51 bytes every 30 seconds, with the transmission power of 14 dBm. We refer to each measurement campaign as a *driving test*.

Dataset Statistics: Given the above setup, the dataset consists of 103236 measurements collected between December '18 and February '19. Each sample represents a LoRa message successfully received and decoded by at least one gateway. The total number of EDs is 15, although messages are not uniformly distributed among them (i.e., more than 50% come from 4 EDs).

Dataset Features: Each sample consists of the following features: an ED identifier (device), the SF ($sf \in [7, 12]$), the message time-on-air ($airtime \in [0.11, 2.30]$ sec), a channel identifier⁴ ($channel \in [1, 8]$), the GPS coordinates of the ED location (latitude and longitude), a hexadecimal representation of the message (payload), and finally, a list of the gateways that successfully received the message. Each gateway is associated with a unique identifier (*id*), a timestamp indicating the receiving time of the message ($rx_time.time$), and power-related indicators, such as the Received Signal Strength Indicator (RSSI) [dBm], Signal to Noise Ratio (SNR) [dB], and Estimated Signal Power (ESP) [dBm]⁵.

³A LoRaWAN-based, IoT platform (<https://www.thethingsnetwork.org/>).

⁴<https://www.thethingsnetwork.org/docs/lorawan/frequency-plans.html>.

⁵ESP is defined as $ESP = RSSI + SNR - 10 * \log_{10}(1 + 10^{(SNR/10)})$.

Table 1: Definition of Mobility classes.

Class ID	Speed Range	No. Samples
Mobility 1	< 4.31 km/h	83094
Mobility 2	4.31 – 11.42 km/h	10586
Mobility 3	11.42 – 32.02 km/h	9556

Feature engineering: Since data are congregated in a single JSON file, and driving test identifiers are not available, we leverage `rx_time.time` and `device` as primary keys to isolate driving tests per day. To discriminate driving tests performed within the same day, we calculate the time distance between two consecutive messages. If the outcome value is higher than a threshold s , we mark the beginning of a new driving test. As the dataset consists of successfully received messages, with no information on the transmission time, selecting s is challenging. High thresholds can result in miss-detection of driving tests performed in close time proximity, while low thresholds can misinterpret a series of *lost* messages with the beginning of a new driving test. We set $s = 450$ sec, equal to the time observed in case 15 messages are lost in a row, that is highly unlikely in our data.

Furthermore, we estimate the average speed under which messages were transmitted by dividing the physical distance between two consecutive transmissions with the elapsed time. To evaluate the physical distance between each pair of consecutive messages, we apply the Haversine formula on the GPS coordinates. After estimating the average speed, we notice that 80% of the messages were transmitted with the EDs moving slower than 4.31 km/h, slightly lower than the average human walking speed (i.e. around 5 km/h)⁶. To better analyze the impact of mobility on the ADR performance, we split the dataset in three classes, based on the estimated speeds, as reported in Table 1.

3 PERFORMANCE EVALUATION

We analyze the ADR performance under the mobility classes in Table 1, by evaluating how the average ESP, as perceived by the gateways, is distributed across SFs. If harsh radio conditions are detected over consecutive messages, device-side ADR will force a SF increase for the next transmissions, aiming to enhance the reliability and allow message delivery in extended coverage. Figure 1 shows the ESP distribution across SFs for each mobility class. We observe that, for all classes, the use of low SFs, e.g., SF = 7, is correctly mapped to relatively high ESP values, while higher SFs match low ESP, improving the coverage. However, the separation between SFs becomes smoother as the mobility becomes higher, highlighting an increasing challenge in selecting the SF to adopt. This also leads to a decrease of the coverage extension benefit.

⁶This is justified by considering the nature of the driving tests; as a matter of fact, postal trucks make multiple stops during a drive, and their routes likely intersect heavily trafficked areas.

Table 2: Pairwise p-values for all combinations of SF increments, across the three mobility scenarios.

	7/8	7/9	7/10	7/11	7/12	8/9	8/10	8/11	8/12	9/10	9/11	9/12	10/11	10/12	11/12
Mobility 1	0	0	0	0	0	$5.2e^{-73}$	$5e^{-167}$	$1.3e^{-172}$	$1.4e^{-121}$	$2.1e^{-27}$	$1.7e^{-48}$	$4.5e^{-44}$	$3.8e^{-8}$	$4.2e^{-13}$	$3.7e^{-3}$
Mobility 2	$1.3e^{-52}$	$6.8e^{-164}$	$5.3e^{-209}$	$4.8e^{-172}$	$2e^{-120}$	e^{-30}	$1.4e^{-61}$	$3.6e^{-69}$	$3.5e^{-63}$	$1.1e^{-8}$	$8.3e^{-19}$	$1.2e^{-26}$	$6.7e^{-5}$	$9.3e^{-13}$	$8.7e^{-5}$
Mobility 3	$1.2e^{-32}$	$2.4e^{-82}$	$2.9e^{-109}$	$1.8e^{-87}$	$3.8e^{-58}$	$1.1e^{-15}$	$9.6e^{-41}$	$8.3e^{-44}$	$1.2e^{-34}$	$1.6e^{-10}$	$2.3e^{-18}$	$1.9e^{-18}$	$3.4e^{-4}$	$3.4e^{-7}$	$2.3e^{-2}$




Figure 1: The distribution of ESP [dBm] as a function of the SFs, for different mobility classes.

As a matter of fact, we observe that the average ESP, when SF = 12, is -126.2 , -126 , and -124 dBm for the three mobility classes, respectively. Moreover the difference between the average ESP in SF = 7 vs. SF = 12 is also shrinking, being 16, 14, and 12 dB for Mobility 1, 2, and 3, respectively⁷.

To quantify to which extent ADR is affected by mobility, we follow a statistical-based approach. One-way Analysis of Variance (ANOVA) is a well-known tool used to assess if there is a statistically significant difference between the means of two or more classes⁸. As the requirements for parametric ANOVA are not fulfilled by our dataset, we select a non-parametric test (Kruskal-Wallis)⁹. Table 2 illustrates the pairwise p-values for all SF increments and mobility classes. We observe a significant difference across the table, implicitly showing that, the SF increase triggered by device-side ADR significantly (and positively) affects the system performance. However, compared to Mobility 1, an increasing trend in the p-values is observed for Mobility 2 and 3, showing that ADR effect keeps vanishing as mobility increases.

4 CONCLUSIONS

In this paper, we studied the performance of the LoRaWAN device-side ADR scheme in diverse mobility scenarios, by leveraging a dataset comprised of LoRaWAN traces collected in the city of Antwerp, Belgium. Results indicate that the benefits of ADR decrease as the ED mobility increases, hence, leaving space for further improvement and optimization.

⁷We do not report the results due to space constraints, but we see the same trend by splitting Mobility 1 in sub-classes on a percentile-based rule.

⁸The null hypothesis is that means of all classes are the same. If p-value is < 0.05 , the null hypothesis is rejected with a 95% Confidence Interval (CI).

⁹We perform pairwise comparisons for each use case with Dunn's Test.

Future work includes a more in-depth analysis of ADR, particularly of the network-side scheme, as well as the design of a smart algorithm that better adapts to different mobility conditions and mobile IoT applications.

5 ACKNOWLEDGMENTS

This work is funded by the EU H2020 research and innovation programme under grant agreement No. 815178 (5GENESIS), and by the Norwegian Research Council project No. 250679 (MEMBRANE).

REFERENCES

- [1] Khaled Q Abdelfadeel, Victor Cionca, and Dirk Pesch. 2018. A fair adaptive data rate algorithm for LoRaWAN. *arXiv preprint arXiv:1801.00522* (2018).
- [2] Michiel Aernouts, Rafael Berkvens, Koen Van Vlaenderen, and Maarten Weyn. 2018. Sigfox and LoRaWAN datasets for fingerprint localization in large urban and rural areas. *Data* 3, 2 (2018), 13.
- [3] Sara Amendola, Rossella Lodato, Sabina Manzari, Cecilia Occhiuzzi, and Gaetano Marrocco. 2014. RFID technology for IoT-based personal healthcare in smart spaces. *IEEE Internet of Things Journal* 1, 2 (2014), 144–152.
- [4] Martin C Bor, John Vidler, and Utz Roedig. 2016. LoRa for the Internet of Things. In *International Conference on Embedded Wireless Systems and Networks (EWSN)*, Vol. 16. 361–366.
- [5] Bin Guo, Zhu Wang, Zhiwen Yu, Yu Wang, Neil Y Yen, Runhe Huang, and Xingshe Zhou. 2015. Mobile crowd sensing and computing: The review of an emerging human-powered sensing paradigm. *ACM Computing Surveys (CSUR)* 48, 1 (2015), 7.
- [6] Shengyang Li, Usman Raza, and Aftab Khan. 2018. How Agile is the Adaptive Data Rate Mechanism of LoRaWAN?. In *2018 IEEE Global Communications Conference (GLOBECOM)*. IEEE, 206–212.
- [7] Dhaval Patel and Myounggyu Won. 2017. Experimental study on low power wide area networks (LPWAN) for mobile Internet of Things. In *2017 IEEE Vehicular Technology Conference (VTC Spring)*. IEEE, 1–5.
- [8] Usman Raza, Parag Kulkarni, and Mahesh Sooriyabandara. 2017. Low power wide area networks: An overview. *IEEE Communications Surveys & Tutorials* 19, 2 (2017), 855–873.
- [9] Mariusz Slabicki et al. 2018. Adaptive configuration of LoRa networks for dense IoT deployments. In *2018 IEEE/IFIP Network Operations and Management Symposium (NOMS)*. IEEE, 1–9.