Li, S., Raza, U., & Khan, A. (2019). How Agile is the Adaptive Data Rate Mechanism of LoRaWAN? In *2018 IEEE Global Communications Conference (GLOBECOM 2018): Proceedings of a meeting held 9-13 December 2018, Abu Dhabi, United Arab Emirates*. [8647469] (IEEE Global Communications Conference (GLOBECOM)). Institute of Electrical and Electronics Engineers (IEEE).
<https://doi.org/10.1109/GLOCOM.2018.8647469>

Peer reviewed version

License (if available):
Other

Link to published version (if available):
[10.1109/GLOCOM.2018.8647469](https://doi.org/10.1109/GLOCOM.2018.8647469)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the accepted author manuscript (AAM). The final published version (version of record) is available online via IEEE at <https://doi.org/10.1109/GLOCOM.2018.8647469> . Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

How Agile is the Adaptive Data Rate Mechanism of LoRaWAN?

Shengyang Li^{*†}, Usman Raza^{*}, and Aftab Khan^{*}

^{*}Toshiba Research Europe Ltd., Telecommunications Research Laboratory, Bristol, UK

[†]University of Bristol, Bristol, UK

Email: sl14470.2014@my.bristol.ac.uk, usman.raza@toshiba-trel.com, aftab.khan@toshiba-trel.com

Abstract—The LoRaWAN based Low Power Wide Area networks aim to provide long-range connectivity to a large number of devices by exploiting limited radio resources. The Adaptive Data Rate (ADR) mechanism controls the assignment of these resources to individual end-devices by a runtime adaptation of their communication parameters when the quality of links inevitably changes over time. This paper provides a detailed performance analysis of the ADR technique presented in the recently released LoRaWAN Specifications (v1.1). We show that the ADR technique lacks the agility to adapt to the changing link conditions, requiring a number of hours to days to converge to a reliable and energy-efficient communication state. As a vital step towards improving this situation, we then change different control knobs or parameters in the ADR technique to observe their effects on the convergence time.

I. INTRODUCTION

The Low Power Wide Area (LPWA) networking technologies [1] are recent outcomes of breakthroughs in communication technologies, as well as miniaturization and decreasing costs of the Internet of Things (IoT) devices. Their fast-paced adoption across all the seven continents aims to provide to tens of billions of devices a range in the order of tens of kilometers at a fraction of the cost and energy consumption of legacy technologies. LoRaWAN has established itself as one of the leading LPWA technologies in the past few years alongside SIGFOX, Ingenu RPMA, and other cellular solutions such as NB-IoT and LTE-M. LoRaWAN is expected to connect a large number of static and mobile end-devices (EDs). Its support of the direct single-hop connection between EDs and the gateways obviates the need for any complex and expensive multi-hop mesh routing schemes.

Nevertheless, LoRaWAN exploits a limited radio bandwidth available in the Sub-GHz part of the industrial, scientific and medical (ISM) band that is also shared among multiple co-existing technologies. In addition to this, the EDs and gateways must respect regional regulations related to the use of this spectrum, which restrict the time-on-air of transmissions. For these very reasons, efficient management of the scarce radio resources is essential for achieving very high *network-wide* performance measured in terms of scale and goodput/reliability of a network. Apart from the problem of limited radio resources, the EDs are often deployed in very far-flung areas in challenging radio environments, resulting in very

high variability in the link quality over time due to multiple different reasons such as obstructions, device mobility and environmental factors [2]. Thus, the responsibility of LoRaWAN stack extends well beyond sharing of the limited resources among a large number of EDs. It must also include intelligent mechanisms capable of *adapting* the communication settings of *individual EDs* to cope up with the link changes and thereby deliver reliable connectivity at all times.

In this paper, we study an *Adaptive Data Rate (ADR)* mechanism that is responsible for the radio resource management and runtime link adaptation for individual EDs in LoRaWAN. This mechanism is recently updated in the LoRaWAN specifications v1.1 [3] and is gradually making its way into commercial deployments. The performance of this mechanism is yet to be fully understood. Therefore, we make a timely contribution in this paper by evaluating the performance of the ADR mechanism as a first step towards exposing different limitations and improving the underlying technique. Unlike most other studies [4], [5], [6], [7], [8] that focus on *network-wide* performance metrics related to scalability, reliability, fairness, and throughput, we focus on performance that the individual EDs receive while running the ADR mechanism. We provide an in-depth study of the agility of the ADR mechanism in adapting the communication settings of EDs in response to link changes.

In this paper, we also provide very first insights into the runtime performance of the official ADR algorithm under dynamic link conditions and various network sizes. Our detailed results reveal that if link conditions change or network size becomes too large, the convergence time of ADR mechanism to a communication setting that provides good reliability and low energy consumption is quite high. A large number of packets are therefore lost, motivating improvements in its design. We then investigate if changing different control knobs, specifically ADR initialization parameters, can improve the convergence rate, and if so to what extent. We also provide a brief discussion of the useful insights gained through this study that can help improve future ADR algorithms.

This paper is organized as follows: Section II presents a very brief primer on LoRa and LoRaWAN immediately prior to the description of the ADR mechanism in Section III. Section IV and Section V present simulation setup, the performance of

Fig. 1: ADR algorithm

the ADR mechanism and the impact of various factors on it. Section VI quantifies the extent to which the agility of the ADR mechanism can be improved by optimizing its different parameters. We then conclude the paper.

II. LoRA AND LoRAWAN

LoRa, a Chirp Spread Spectrum (CSS) technique, is the underlying PHY layer used by LoRaWAN, the upper network stack designed by LoRaWAN Alliance. The CSS technique supports multiple data rates. For a low data rate, it uses a large Spreading Factor (SF) that puts a high level of redundancy and amount of energy in a signal. Therefore, the signal can reach long distances and still retain enough strength to be successfully received. The same explanation holds for a signal sent at a high Transmission Power (TP). From the radio resource management point-of-view, the use of high SF values keeps radio medium busy for a long duration due to low data rates. Therefore, it is desirable to use the lowest possible SF that provides a good link between EDs and network. Different orthogonal SFs enable multiple successful receptions when used to send packets overlapping both in time and frequency channel. In the context of this paper, we restrict our discussion

to SF and TP, the two control knobs adjusted by the ADR algorithm in pursuit of achieving long-distance, reliable, and energy-efficient communication.

LoRaWAN defines the higher layer protocols and the network architecture that enable EDs to directly connect with the gateways using an ALOHA based multiple access scheme over the sub-GHz ISM bands. The gateways are then connected to the network servers that perform device authentication, downlink transmission scheduling, and execution of a part of ADR algorithm among many other important network-level functions. LoRaWAN mechanisms respect the regional regulation related to the use of the sub-GHz ISM spectrum, such as those governing maximum TP and duty cycles. This paper assumes operation of LoRaWAN in Europe where TP and duty cycle are limited to 14 dBm and 1% respectively for the default frequency channels. The SF can be varied from 7 to 12 to adapt both the communication range and data rate. Three device classes are defined based on the application requirements for overall energy-efficiency and downlink communication latency. The most energy-efficient (and typically battery-powered) EDs are the Class A devices that experience

the longest latency to receive downlink messages that are sent by the network only shortly after an uplink transmission. The other device classes provide additional opportunities for receiving downlink messages at the expense of higher energy consumption.

III. ADAPTIVE DATA RATE MECHANISM

The ADR mechanism is part of LoRaWAN specifications. It aims to provide a fairly reliable and battery-friendly connectivity by adapting SF and TP to changes in link conditions. Both EDs and the network play an important role in this process.

If an ED observes that a large number of consecutive uplink transmissions are not followed by a downlink response from the network, it assumes lost connectivity and resolves this issue by gradually stepping up its TP to the maximum before doing the same for SF. These measures gradually improve the robustness of the link. Figure 1a explains the full operation of EDs for adapting their TP and SF according to LoRaWAN Specifications v1.1. The two parameters namely `ADR_ACK_LIMIT` and `ADR_ACK_DELAY` control the number of uplink messages, after which if a downlink response is not received, an ED must increase either TP or SF. The value of these parameters along with the network size, deployment environment, and the amount of link fluctuations, all affect the time to converge to a state where ED is able to successfully re-establish a reliable link to the network. Section V provides a detailed analysis of these aspects.

The EDs adapt communication setting to establish a reliable, but not necessarily an energy-efficient communication with the network. EDs can, however, request the network to step in and monitor the quality of uplink receptions from the recent past. If the link quality calculated over the last N packets is too high compared to the minimum receiver sensitivity threshold, the network decides to reduce SF and/or TP. The new SF and TP values are set such that the expected signal-to-noise ratio of the future packets is above the minimum receiver sensitivity threshold by a pre-configured margin. Reduction in SF and TP would enable faster (high data rate) transmissions that consume less energy. Semtech, the organization that designed LoRa, provides recommendations for implementing the network-side of the ADR algorithm, which is adopted by different operators as well as The Things Network, a popular crowd-sourcing LoRaWAN network. On the network-side, N , the minimum number of received packets that the network require to choose values of TP and SF, significantly affect the agility of the ADR algorithm as highlighted later in Section V.

IV. SIMULATION SETUP

We implement the ADR algorithm in LoRaWANSim [9], a discrete event simulator that already includes LoRa as well as detailed LoRaWAN MAC protocol features including support for downlink traffic, retransmissions and support for Class A devices. Both EDs and the network are made capable of executing their sides of the algorithm shown in Figure 1, one of the contributions of this paper.

TABLE I: Simulation Parameters

Parameter	Value
Communication range	670 m
Average message rate	10 minutes per message
Carrier frequency	g1 sub-band (868.1, 868.3, 868.5 MHz)
Bandwidth	125 kHz
Code rate	4/5
Spreading factor	7 to 12
Transmission power	{2, 5, 8, 11, 14} dBm
Path loss values [4]	$d_0 = 40$ m, $\gamma = 2.08$, $\overline{L}_{pl}(d_0) = 127.41$ dB
Channel variation level	Low to High: $\sigma = \{0, 1.785, 3.57\}$ dB

To evaluate the ADR mechanism, we simulated networks consisting of a gateway with a varying number of randomly distributed EDs that transmit one packet every ten minutes on average. An urban scenario with the same path loss model as used by earlier works [9], [6] is assumed. The EDs are tuned to use the default three central carrier frequencies in the g1 sub-band of the European sub-GHz ISM band, which is subject to a 1% duty cycle limit. The simulation parameters used in following experiments are listed in Table I.

Each experiment accounts for 12 days of simulated time and the reported results are averaged over 30 repetitions. The error bars in all the performance figures represent the standard deviations.

V. PERFORMANCE EVALUATION

We are mainly concerned with how quickly an ED, after experiencing a change in its link quality, converge to a state where it is assigned the right SF and TP values by the network. To quantitatively measure this, we define *convergence time* as the duration from the change in the link quality until when the network receives enough number of packets required to compute the new SF and TP values. We are also interested in the amount of energy consumed by the radio during this period.

Now, we evaluate the impact on the ADR mechanism due to various factors such as network scale, deployment environment, traffic type and link changes.

A. Impact of Network Size

Firstly, we start by looking at how network size affects the runtime performance of the ADR algorithm. For this particular experiment, we initialize networks of varying sizes and let the ADR algorithm run for all the EDs for some time so that the EDs acquire stable values of SF and TP. We then introduce to the networks additional 100 EDs and measure and report their convergence times.

Figure 2a shows that when the network size is increased, the convergence time increases as well from around 200 minutes for a 100-node network to more than 3000 minutes for a 4000-node network. The slow convergence in the large networks is due to a very high contention between a large number of uplink transmissions. This can be validated by Figure 2b that shows that data loss due to collisions increases from approximately 17% to 85% as network scales from 100 to 3000 EDs. Thus,

Fig. 2: Impact of network size on ADR algorithm.

Fig. 4: Packet loss during ADR algorithm

Fig. 3: Impact of channel condition on convergence time of ADR algorithm.

Fig. 5: Impact of change in link quality on convergence time for the ADR algorithm.

it takes more time for the network to receive N transmissions from the EDs required to assign optimal TP and SF values. Furthermore, once the new values of TP and SF are calculated, the network is required to send downlink command messages to inform EDs. This is often not possible for a larger network due to the 1% duty cycle limit on the transmissions from the gateway, resulting in additional delay.

B. Impact of Deployment Environment

Different deployment environments cause different levels of variation in the radio channel. We simulate three different scenarios related to low, medium and high levels of channel variation – consistent with [6]. This is achieved by changing the value of the standard deviation that accounts for the shadowing effect in the log-normal path loss model used in our simulations.

Figure 3 shows the convergence time of the ADR algorithm under three different conditions. When network size is small, a high variation in channel slightly increases the convergence time of ADR algorithm. Whilst one will naturally expect this to happen, higher variation in channel signifies that uplink packets are more likely suffering from losses due to fading. Figure 4 highlights different reasons of packet loss. The total percentage packet loss under highly varying channel exceeds the total percentage packet loss under a less varying channel.

However, collisions happen more often under low variation. High packet loss translates to longer convergence times.

When the network size is large, we can observe that higher variation in a channel does not necessarily lead to an increase in the convergence time. Rather, it actually reduces the convergence time. As shown in Figure 4b, the uplink collisions happen much more for low channel variation. Although uplink fading under high varying channel occurs more often than the one in the low varying channel as expected. But under the low varying channel, the increase in the uplink collisions is so high that the total packet loss outweighs the one under high varying channel. One potential reason is that a high variation in channel introduces randomness in the received signal strength of the uplink packets. This leads to a large difference between signal strengths of overlapping packets, resulting in successful decoding of the strongest signal due to the capture effect. More uplink packets can therefore successfully reach the gateway. Another possible reason is high variation in channel eases the crowded network. High variations in channel cause some uplink packets to not reach the gateway, reducing the contentions for others that do. Thus, the chances of collision reduce at gateway. Overall, the success of uplink packets reaching gateway is more influenced by the number of collisions rather than fading for large network size. This trend is reversed for network size.

Fig. 6: Impact of confirmed uplinks on ADR algorithm

Fig. 8: Impact of ADR_ACK_LIMIT on ADR algorithm

Fig. 7: Impact of number of past frames collected on ADR algorithm

Fig. 9: Impact of ADR_ACK_DELAY on ADR algorithm

C. Impact of Link Changes

LoRaWAN is a technology of choice for various smart city applications that are often deployed in dense urban environments. Wireless links may often degrade or improve in such environments suddenly. Smart parking application is one example where the parking sensors often are obstructed by vehicles causing links to degrade. Now, we simulate such changes in communication environment by altering the mean path loss value of the communication between end devices and gateway. Figure 5a simulates a case when the link quality of an ED degrades due to obstructions or mobility that increase the mean path loss. Figure 5b, on the other end, shows a case where ED improves its link quality by reducing mean path loss.

Figure 5a clearly shows when the link quality degrades, the time required by the ADR algorithm to converge to the right communication parameter setting increases significantly. This overhead mainly comes from the process running on ED to regain connectivity to the gateway. Unfortunately, this process requires EDs to lose sufficient number of sent packets before moving to higher SF or TP values. In Figure 5b, when node has a good link quality, initially ADR takes a slightly longer time, mainly because of channel variation and resulting additional

packets due to loss of uplink packets. However as the link quality continues to improve, gateway continues to receive the packets for computing SF and TP. In our simulations, the convergence time approaches 200 minutes, the minimum time required to receive $N = 20$ packets that are transmitted every 10 minutes on average.

D. Impact of Traffic Type

Some application traffic in LoRaWAN can be of more value than the rest and therefore should be acknowledged by the network. For this purpose, LoRaWAN supports confirmed messages. We now are interested in how the convergence time is affected by the confirmed uplink traffic. Figure 6a shows the effect of different percentage of uplink packets requiring ACK on the convergence time. If an ACK is lost, the frame will be retransmitted after a short time. Retransmissions improve link reliability, enabling the gateway to collect enough number of packets quickly, thus reaching optimal SF and TP earlier compared to the cases when all the messages are unconfirmed.

VI. OPTIMIZATION FOR THE ADR MECHANISM

After learning that the convergence of ADR is slow, we now change different tunable parameters of the ADR algorithm (shown in Table II) to analyze the corresponding effect on its performance. We consider N , the number of packets required