

Guidelines for an Energy Efficient Tuning of the NB-IoT Stack

Muhammad Tahir Abbas, Johan Eklund, Karl-Johan Grinnemo, Anna Brunstrom, Stefan Alfredsson
Karlstad University, Karlstad, Sweden

{tahir.abbas, johan.eklund, karl-johan.grinnemo, anna.brunstrom, stefan.alfredsson}@kau.se

Özgü Alay

University of Oslo and Simula Metropolitan, Oslo, Norway
ozgua@ifi.uio.no

Sándor Katona, Gergely Seres, Bela Rathonyi
Ericsson AB,

{sandor.katona, gergely.seres, bela.rathonyi}@ericsson.com

Abstract—In this paper, we study the energy consumption of Narrowband IoT devices. The paper suggests that key to saving energy for NB-IoT devices is the usage of full Discontinuous Reception (DRX), including the use of connected-mode DRX (cDRX): In some cases, cDRX reduced the energy consumption over a 10-year period with as much as 50%. However, the paper also suggests that tunable parameters, such as the inactivity timer, do have a significant impact. On the basis of our findings, guidelines are provided on how to tune the NB-IoT device so that it meets the target of the 3GPP, i.e., a 5-Wh battery should last for at least 10 years. It is further evident from our results that the energy consumption is largely dependent on the intensity and burstiness of the traffic, and thus could be significantly reduced if data is sent in bursts with less intensity, irrespective of cDRX support.

I. INTRODUCTION

Cellular Internet of Things (CIoT) delivers connectivity on a massive scale, and, in doing so, enables several key IoT applications. Some of these applications include smart metering to help control operation and to reduce energy consumption; asset monitoring and smart logistics to enhance distribution efficiency; smart environmental monitoring to reduce city pollution and improve efficiency in agriculture; and, smart wearables to measure health parameters, and making constant health monitoring possible without unnecessarily constrain patient's lives, etc. Due to its versatility, we have witnessed an exponential increase in the number of CIoT connections worldwide since 2015, with over 1 billion CIoT connections in 2020, and a forecast of roughly 5 billion connections in 2025 [4].

The two dominating CIoT technologies are Long-Term Evolution for Machine-Type Communications (LTE-M) and Narrowband IoT (NB-IoT), which are projected to account for more than half of all CIoT connections in 2025 [4]. Of these two technologies, LTE-M targets use cases such as connected vehicles, wearable devices, and trackers of goods, i.e., use cases where quite large data rates (up to 1 Mbps) matter. On the other hand, NB-IoT targets use cases such as smart meters and monitoring sensors, i.e., use cases where device cost

and complexity, flexible and efficient radio spectrum usage, and good coverage in challenging radio conditions are of paramount importance.

Both LTE-M and NB-IoT target battery-powered devices and are designed for mass deployment; in fact, they are designed to support 100,000 or more devices per cell. Since replacing batteries for such a large number of devices would not only be cumbersome, but also excessively costly, battery lifetime is a key concern for both these technologies. This paper focuses on NB-IoT and the great challenge to reconcile a widespread coverage in radio-challenging environments with a low energy consumption. Particularly, to provide a battery lifetime of more than 10 years with a battery capacity of 5 Wh under these circumstances.

To meet the energy consumption targets, NB-IoT extends and refines the energy-saving mechanisms already existing in 4G/LTE. Devices can initiate the connection in a so-called Active mode for an extended period of time with the introduction of extended Discontinuous Reception (eDRX) cycle. This paper studies in a simulation campaign the impact of pertinent, tunable parameters of the protocol stack, e.g., the RRC inactivity timer where the device stay connected with eNodeB, the Active (T3324) timer during which the device is not connected and stays in idle state, the CoAP retransmission timer which triggers a data retransmission after timeout, the eDRX parameters that listens to paging occasions, etc., as well as traffic properties, on the NB-IoT device energy consumption. On the basis of this study, it provides configuration guidelines. Although there are some previous works on the energy consumption of NB-IoT devices, to the best of our knowledge, this paper is among the first to actually provide guidelines on how to configure the NB-IoT protocol stack in order to save energy.

The remainder of this paper is organized as follows. Section II provides a brief overview of NB-IoT with a focus on those parts of the technology that matter the most for the energy consumption. In Section III, the simulation environment is discussed including the modeling of the NB-IoT

protocol stack and the considered network topology. Next, in Section IV, the energy consumption of an NB-IoT device is evaluated, and we study how energy consumption is impacted by several tunable parameters and traffic properties. Section V provides configuration guidelines for an NB-IoT device and Section VI surveys related work. The paper concludes in Section VII with a brief summary of the paper and a short discussion of future work.

II. OVERVIEW OF NB-IoT

NB-IoT is a 3GPP radio technology standard introduced in Release 13 [21] that is characterized by an improved indoor coverage, support for a massive number of low-throughput devices, low delay sensitivity, ultra-low device cost, and a low device energy consumption. NB-IoT is an LTE-based cellular radio access technology that provides Low Power Wide Area (LPWA) connectivity in licensed spectrum. It can be deployed in-band, within a LTE cell, or guard band, outside of the LTE cell in the guard band, or standalone, in a dedicated spectrum. Additionally, through the use of a supporting technology, MulteFire [12], NB-IoT is able to operate in unlicensed spectrum.

Although there are several competing LPWA technologies to NB-IoT, e.g., Bluetooth, ZigBee, SigFox, LoRa, they mostly operate in the unlicensed spectrum, and have their weaknesses in terms of availability, Quality-of-Service (QoS), and network security. The strengths of the CIoT approach include its worldwide availability, its guaranteed QoS, its support for carrier-class network security, and its strong ecosystem of service providers and device manufacturers.




Fig. 1: NB-IoT network architecture. The filled arrow depicts the route of traffic when CP CIoT EPS optimization is used, and the dashed arrow when UP CIoT EPS optimization is used.

NB-IoT normally exchanges traffic with an application server, in the *user plane* using the Serving Gateway (SGW) and Packet Gateway (PGW) in the Evolved Packet System (EPS). This is shown by the dashed line in Figure 1. To reduce energy, 3GPP specifies an optimization, the so-called User Plane (UP) CIoT EPS optimization, in this communication path, where the RRC connections can be suspended and resumed. As follows from the solid line in Figure 1, 3GPP also specifies an optimization where CIoT traffic can be sent over Internet Protocol (IP) in the *control plane*, similar to the Short Message Service (SMS). This optimization, a.k.a. the CP CIoT EPS optimization, avoids triggering the user plane radio bearer establishment, thus saving energy.

As the data reaches the E-UTRAN Node B (eNodeB), it is forwarded to the Mobility Management Entity (MME), where it can be forwarded to the SGW and routed the same way as the user plane would. Alternatively, it can be routed via the Service Capability Exposure Function (SCEF) in the control plane, avoiding the user plane altogether. Eventually, the data is forwarded to the application server.

The user plane optimization is optional for NB-IoT devices to support. Most commonly the control plane optimization is deployed, and it is therefore the optimization considered in this paper.¹

A. Protocol Stack

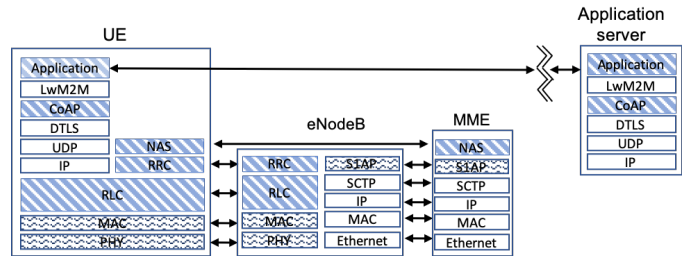


Fig. 2: The protocol stacks at the NB-IoT device, the eNodeB, and the MME, for the CP CIoT EPS optimization. In our simulation environment, the striped patterned, dark blue-colored layers are more or less completely implemented; the wavy patterned, light blue-colored layers are partially implemented; and, the white-colored layers are essentially stubs.

Figure 2 illustrates the protocol stacks at the NB-IoT device and at pertinent nodes in the EPS when NB-IoT traffic is transferred over the control plane, as prescribed by the CP CIoT EPS optimization. As follows, the radio stack is more or less a functionally reduced version of the LTE protocol stack. On the physical layer, NB-IoT employs the same radio technologies as legacy LTE with OFDM in the downlink and SC-FDMA in the uplink. However, the resource scheduling differs in the uplink from LTE in that it uses the subcarrier instead of the physical resource block as the scheduling unit, something which makes it scale better with the number of connected devices. On the Medium access control (MAC) layer, NB-IoT introduces features that specifically target small-data transmission, including RRC-connection suspend/resume and multi-carrier support. To keep the complexity of the MAC layer down, and to simplify its implementation, NB-IoT only supports one (3GPP Rel-13) or two (3GPP Rel-14) Hybrid Automatic Repeat Request (HARQ) processes in both downlink and uplink. Moreover, in 3GPP Rel-13, it could be noted that NB-IoT only supports contention-based random access.

The upper-layer protocols of the radio stack work essentially the same way as in legacy LTE, however, a bit simplified:

¹To keep the cost down, several NB-IoT devices do not even support UP CIoT EPS. However, LTE-M devices, not being equally cost sensitive, must support UP CIoT EPS.

handover is not supported. Thus, when an NB-IoT device moves out of the coverage area of the serving cell, it will experience a radio-link failure, and, it will perform cell re-selection and resume data transfer either by re-establishing the RRC connection or setting up a new one.

The lower protocols of an IP stack with support for NB-IoT are essentially the same as for IP stacks in enterprise networks, however, the upper layers differ a bit from these stacks. In particular, since verbose, web-based application-layer protocols are too heavy for NB-IoT devices, new lightweight protocols that are better suited for large numbers of constrained nodes are used, e.g., the Message Queuing Telemetry Transport (MQTT) [6] or the Lightweight Machine-to-Machine (LwM2M) [13] protocols. The MQTT protocol is a publish-/subscribe-based protocol that primarily targets telemetry applications, while LwM2M is a more general protocol with support for several types of applications. In particular, the LwM2M protocol features a modern architectural design based on Representational State Transfer (REST). It offers a data model that is both extensible and which is designed with performance and the constraints of IoT devices in mind.

The LwM2M protocol is tightly coupled with the IETF Constrained Application Protocol (CoAP) [16], and primarily uses this protocol as the underlying transfer protocol. The CoAP protocol may run in reliable as well as in unreliable mode. When reliable transfer is employed, end-to-end reliability is achieved through the retransmission of lost messages. These retransmissions are managed by a CoAP retransmission timer. To secure the communication between an NB-IoT device and an Application Server, CoAP runs over the Datagram Transport Layer Security (DTLS) [15] protocol.

B. Energy Saving Features

The RRC layer controls communications between an NB-IoT device and an eNodeB at the radio interface. There are two RRC states: RRC_CONNECTED and RRC_IDLE. If an RRC connection exists between the NB-IoT device and the eNodeB, the device is in RRC_CONNECTED. If no RRC connection exists between the device and the eNodeB, the device is in RRC_IDLE. In RRC_CONNECTED, uplink and downlink application-layer data transmissions are possible. The time a device stays in RRC_CONNECTED is usually governed by a timer, the RRC inactivity timer, whose value is controlled by the eNodeB. Each time traffic is sent or received in RRC_CONNECTED, the RRC inactivity timer is reset and the time in this state is extended. When the RRC inactivity timer expires, the RRC connection between the device and the eNodeB is released and the device enters RRC_IDLE. The NB-IoT device may also transmit a Release Assistance Indication (RAI) message to either the MME (3GPP Rel-13) or the eNodeB (3GPP Rel-14) to request the RRC connection to be released. This paper does not consider the use of RAI, and thus, for our purposes, the time in RRC_CONNECTED is completely governed by the RRC Inactivity Timer.

To optimize its energy consumption, NB-IoT uses several energy-saving features including the Power Saving Mode

(PSM) [19], the connected-mode Discontinuous Reception mechanism (cDRX) [20], the idle-mode Discontinuous Reception mechanism (iDRX) [20], and the extended Discontinuous Reception (eDRX) [20], [23]. Figure 3 illustrates how cDRX, iDRX, eDRX and PSM interplay to significantly reduce the energy consumption of an NB-IoT device.

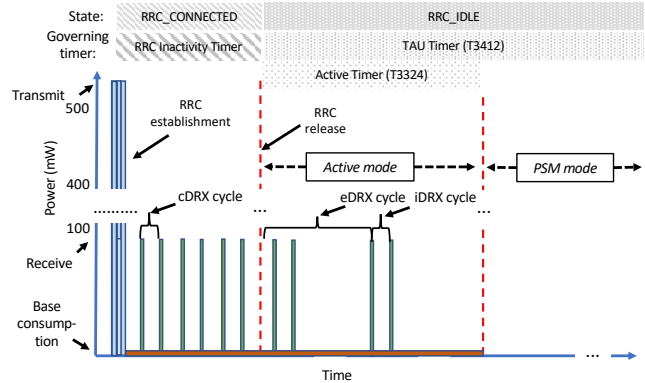


Fig. 3: Power consumption of an NB-IoT device during the RRC_CONNECTED and RRC_IDLE states. Note how the cDRX, iDRX, eDRX, and PSM energy-saving mechanisms help to reduce the power consumption, and thus assist to conserve energy.

The frequency with which an NB-IoT device checks for paging messages is defined by its Discontinuous Reception (DRX) cycle. As follows from Figure 3, the eDRX mechanism is similar to DRX but with less frequent listening periods; the length of the listening periods in eDRX is governed by the Paging Time Window (PTW). In between these listening cycles, an NB-IoT device enters a sleep state to save energy. In the PSM mode, an NB-IoT device turns off its radio and make a transition into a deep sleep or hibernation state. In this state, a device stays registered with the network, however, only resumes a connection provided there is any data from the application layer, or it needs to perform a Tracking Area Update (TAU) operation to advertise its location to the network.

When an NB-IoT device registers with the cellular network, it indicates its preferred eDRX, and PSM parameter settings. As part of accepting a device into the network, the network configures the device with the parameter settings to be used. Two of those parameters being configured during a device registration are T3324 and T3412. The T3324, a.k.a. active timer, determines the duration during which the device remains reachable for a device-terminated transaction in RRC_IDLE: the so-called active mode. In NB-IoT, the active timer can be set as high as 310 hours. The T3412 or periodic Tracking Area Update (TAU) timer specifies the time interval between Tracking Area updates, i.e., how often an NB-IoT device notifies its availability to the cellular network, in RRC_IDLE. The NB-IoT device negotiates the value of the TAU and active timers with the MME during the TAU procedure. Although, the TAU timer predates NB-IoT, its range is extended in NB-

IoT to 413 days.

During the RRC_CONNECTED state, a device may employ the cDRX mechanism. In those cases cDRX is used, the device cycles between a listening mode, when the device actively listens for paging occasions, and a sleep mode, during which the radio unit of the NB-IoT device is switched off. The length of these cycles is determined by the length of the cDRX cycle, and the time spent in listening mode is determined by a listening-time parameter. Since cDRX adds quite some complexity to an NB-IoT device, it is not yet generally deployed. In view of this, we have in our simulations covered both scenarios with full DRX support, i.e., scenarios in which the NB-IoT supports cDRX, as well as scenarios with partial DRX support, in which the NB-IoT device lacks support for cDRX, and instead stays in continuous listening mode.

When an RRC connection is released, an NB-IoT device ends up in the active mode in RRC_IDLE. During the active mode, eDRX, an extension of the legacy IDLE mode DRX, is employed. The length of an eDRX cycle may vary between 20.48 s and 2.91 hours. As depicted in Figure 3, an eDRX cycle starts with a listening period, comprising a train of listening occasions, followed by a comparatively long dormant period. The frequency of listening occasions during the listening period is governed by the length of the iDRX cycle. When the active timer expires, the NB-IoT device enters the PSM mode: a long sleep cycle, during which the device is considered not reachable by the network.

III. SIMULATION SETUP

To study the energy consumption of NB-IoT devices, we have developed an NB-IoT simulation environment in the OMNeT++ simulation framework – a discrete-event simulation framework written in C++ [24]. Our simulation environment models NB-IoT Release 13, and supports the NB-IoT protocol stacks as depicted in Figure 2. The striped patterned, dark blue-colored layers have been more or less completely implemented as specified by 3GPP in our simulation environment. Notably, the CoAP layer in our simulation environment has support for reliable transfer, with confirmable messages being retransmitted using a default timeout and exponential back-off between retransmissions; the Non-Access Stratum (NAS) and RRC layers implement non-access stratum and access-stratum control signaling; and, the Radio Link Control (RLC) layer has Automatic Repeat reQuest (ARQ) support, and is able to offer an ordered and reliable transfer in Acknowledged Mode (AM). The wavy patterned, light blue-colored layers have been partially implemented in our simulation environment, and implement those parts deemed essential to properly model the energy consumption of an NB-IoT device, e.g., the Hybrid Automatic Repeat reQuest (HARQ) mechanism in the MAC layer, the DRX mechanisms in the MAC and PHY layers, etc. Finally, the white-colored layers are only implemented in our simulation environment to the extent that they are able to properly communicate with their adjacent upper and lower layers.

In our simulations, we considered a scenario in which an NB-IoT device generated reliable, uplink traffic in bursts of one or several messages. The burstiness of the traffic was controlled by the NB-IoT application on the basis of its requirements. The uplink traffic was routed from the NB-IoT device, via the MME, SGW, and PGW, before it finally reached an application server. During the device attach and paging procedures, the MME communicated location updates to the Home Subscriber Server (HSS). The studied scenario is depicted in Figure 4, and the corresponding simulation parameters are listed in Table I.

Each simulation was repeated 30 times, and each simulation run covered 10 years of NB-IoT device operation. The energy consumption was measured in terms of average device energy consumption over a 10-year period along with a 95% confidence interval.




Fig. 4: The NB-IoT network scenario studied in our evaluation of the energy consumption of an NB-IoT device.

Parameter	Values
PSM power consumption	15 μ W
RRC_IDLE power consumption	3 mW
Transmission power	545 mW
Reception power	90 mW
Maximum output power	+23 dBm [22]
Transmission time	5 ms
Listening time (cDRX & eDRX)	20 ms
DRX inactivity timer (cDRX)	42.67 ms
Packet size	100 bytes
CIoT-Uu link delay	10 ms
S1-MME link delay	1 ms
S11 link delay	1 ms
S5 link delay	1 ms
S6a link delay	1 ms
SGi link delay	10 ms
CoAP retransmission timer	2 s
RRC inactivity timer	20 s
Active timer	30 s
cDRX/iDRX cycle	1.28 s
eDRX cycle	20.48 s
PTW	2.56 s
HARQ block error rate (BLER)	10 %
PGW-AS packet loss	1 %

TABLE I: Simulation parameters.

IV. EVALUATION

As mentioned, 3GPP prescribes a battery lifetime of at least 10 years with a battery capacity of 5Wh. To get an appreciation for how challenging this target indeed is, we should keep in mind that the energy required to barely keep

an NB-IoT device alive in PSM mode for 10 years is about 1.2 Wh. Since NB-IoT networks with both partial and full DRX support, i.e., scenarios where cDRX is not used and scenarios where it is used, have been deployed, both scenarios are considered.

Our first set of simulations in Section IV-A looks at the impact of traffic properties and packet loss on the energy consumption of an NB-IoT device, i.e., factors that are often not possible to configure. Due to page restrictions, these simulations only cover scenarios with full DRX support. We prioritize full DRX support since it is encouraged by 3GPP, and since it is gradually becoming commonplace. Our second set of simulations in Section IV-B, considers the impact of pertinent protocol-stack parameters, i.e., tunable parameters, on device energy consumption. This set of simulations covers both scenarios with partial and full DRX support.

A. Traffic Properties and Packet Loss

This section considers how the intensity and burstiness of traffic generated by a generic application, i.e., an application that is not tied to a particular use case, impact the energy consumption of an NB-IoT device in a scenario with full DRX support. Moreover, the section studies how deteriorating conditions and Block Error Rate (BLER) over the radio channel (CIoT-Uu) and packet losses over the Internet (PGW-AS) influence energy consumption for these devices.

1) *Traffic Intensity*: To study the effects of different traffic intensities on the energy consumption for an NB-IoT device, we consider the device energy consumption over 10 years for three traffic intensities: 1 message per 3 hours, 1 message per 4 hours, and 1 message per 5 hours. Moreover, we consider the energy consumption with increasing degrees of contention on the radio interface, resulting in increasing random access delays during the device attachment. Figure 5 presents our results.

We observe that provided we do not have extreme random access delays, the 3GPP 5-Wh target is reachable with the lowest studied traffic intensity. Still, the target is not within reach for the highest studied traffic intensity, and only within reach with the mid traffic intensity provided the radio interface contention is limited, i.e., the random access delay is reasonably low. We note that our results align with those of Lauridsen et al. [8], which point out in their work that the lifetime of an NB-IoT device is often a bit shorter than that predicted by 3GPP [22].

2) *Traffic Burstiness*: To observe the effects of traffic burstiness on the energy consumption of an NB-IoT device, we carry out a series of simulations in which the same number of messages was sent over a 10-year period but in bursts of 1, 2 and 4 messages. On the basis of Section IV-A1, the number of messages sent corresponds to a traffic intensity of 1 message per 4 hours. The results from our simulations are presented in Figure 6; the graph in Figure 6a shows how the energy consumption varies with the random access delay for traffic with different burstiness, and the graph in Figure 6b complements Figure 6a and presents the number of RRC connection

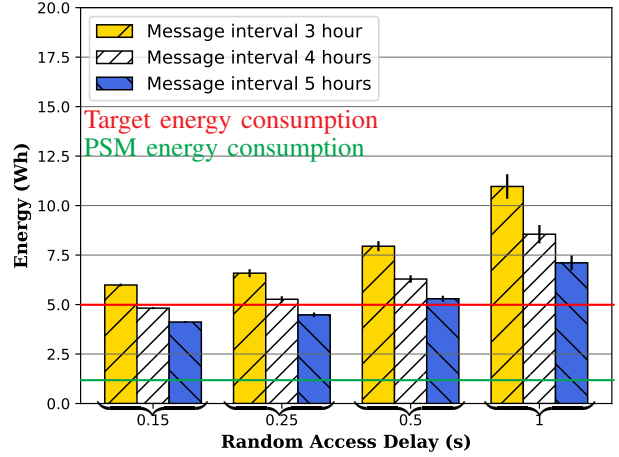


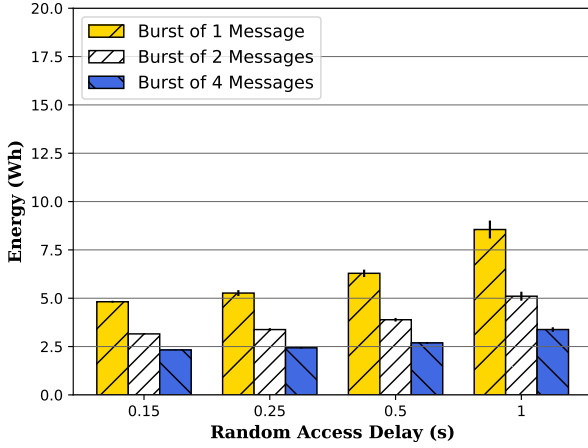
Fig. 5: Energy consumption for an NB-IoT device with full DRX support vs. random access delay for NB-IoT traffic generated with different traffic intensities.

establishments carried out in the simulations in Figure 6a. As follows, the device energy consumption is significantly reduced when the same amount of messages is sent in bursts as compared to being sent separately. It also follows that this is true irrespective of the experienced random access delay. Instead, it seems that the device energy consumption correlates closely with the number of RRC connection establishments.

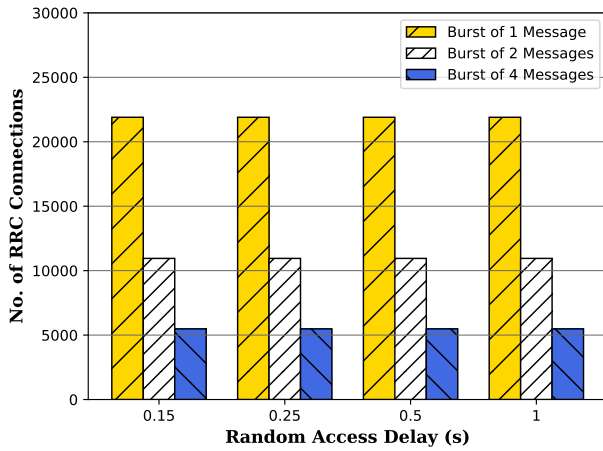
3) *Packet Loss*: As follows from Table I, we introduced a 10% BLER at the MAC layer over the radio channel (CIoT-Uu) and a 1% packet loss over the Internet (PGW-AS) as default values. To study how block errors over the radio channel and packet losses over the Internet impact the energy consumption of an NB-IoT device, we extend our previous simulations by also considering BLERs of 5% and 20% on the radio channel, and packet-loss rates of 3% and 5% over the Internet. Our simulations suggest that neither block errors on the radio channel nor packet losses over the Internet have any significant impact on the device energy consumption. On the radio channel, block errors are mostly recovered by the MAC HARQ mechanism, and are cheap in terms of energy consumption. The packet losses on the Internet path between the PGW and the AS are fewer than those over the radio channel, and are recovered by retransmissions at the CoAP layer. Since the RRC Inactivity timer is 20s – much larger than the end-to-end round-trip times (RTTs) experienced by the NB-IoT device – the radio link between the device and eNodeB is still up when the CoAP retransmission timer expires. Thus, these retransmissions do not involve any RRC connection re-establishments, and are fairly inexpensive in terms of energy consumption.

B. Protocol Stack Parameters

This section considers the influence on energy consumption of a number of parameters in the protocol stack of an NB-IoT device. Since it is believed that cDRX has a significant impact



(a) Energy consumption vs. random access delay.



(b) Number of RRC connection establishments vs. random access delay.

Fig. 6: Energy consumption and number of RRC Connections for an NB-IoT device with full DRX support for different burst sizes.

on the energy consumption of an NB-IoT device which only transfers a small amount of data infrequently, we differentiate between scenarios with full and partial DRX support.

1) *Full DRX Support*: First, we study the impact of the protocol-stack parameters, the RRC inactivity timer, the active timer, the CoAP retransmission timer, and the eDRX cycle length, on an NB-IoT device in scenarios where cDRX is used.

a) *RRC Inactivity Timer, Active Timer, and eDRX*:

Figure 7 shows how the device energy consumption varies with increasing random access delays and with different settings of the RRC inactivity timer. As follows, the NB-IoT device energy consumption increases significantly with an increased random access delay. It is also evident that provided the RRC inactivity timer is enabled, i.e., is set to a value larger than

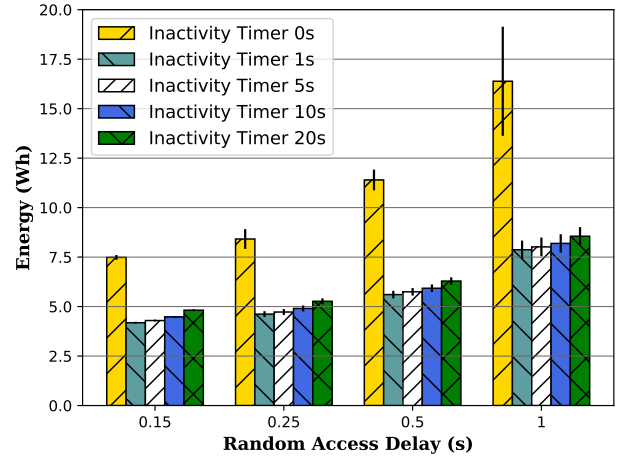


Fig. 7: Energy consumption for an NB-IoT device with full DRX support vs. random access delay for different RRC inactivity timers.

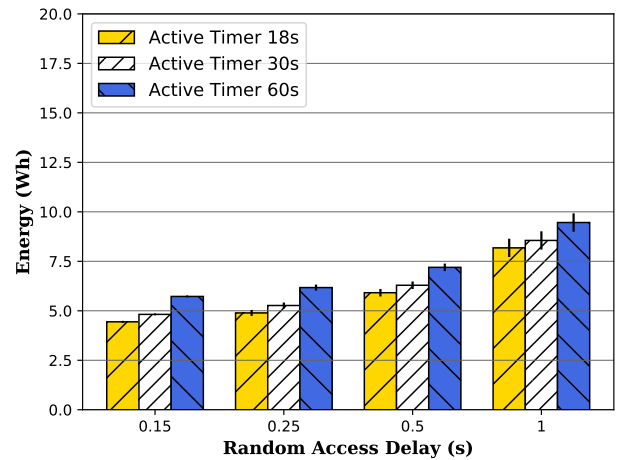


Fig. 8: Energy consumption for an NB-IoT device with full DRX support vs. random access delay for different active timers.

0s, varying the timer between 1s and 20s has a limited impact on the device energy consumption. The reason being that in none of these cases, the RRC inactivity timer expires prematurely, i.e., before outstanding CoAP messages have been acknowledged, and thus no extra, energy-costly RRC connection establishments take place. Instead, the only reason we observe an increase in energy consumption when the RRC inactivity timer is increased, is due to the fact that the device then stays longer in the RRC_CONNECTED state. Still, it is obvious that this result will not hold true if, for some reason, the end-to-end RTT experienced by the NB-IoT device drastically increases.

As follows from Figure 8, also the active timer has a limited impact on the device energy consumption, a result which is

rather expected, and only confirms that eDRX cycles conserve energy. However, we observe that setting the active timer to excessively large values do have a significant impact on the device energy consumption. Still, the length of the eDRX cycle seems only to marginally influence the device energy consumption. In fact, when we run the simulations in Figure 8 with eDRX cycles of lengths, 10.24 s, 20.48 s, and 40.96 s, and with an active timer of 18 s, 30 s and 60 s, respectively, we only observe a slight increase in energy consumption over a 10-year period.




Fig. 9: Energy consumption for an NB-IoT device with full DRX support vs. RRC inactivity timer for different CoAP retransmission timers with an end-to-end RTT of 8 s.

b) CoAP Retransmission Timer: To study the effect of the CoAP retransmission timer and its interaction with the RRC inactivity timer on the energy consumption of an NB-IoT device, we increase the end-to-end RTT experienced by the NB-IoT device to 8 s, i.e., a very long, but still feasible, end-to-end RTT [10]. Figure 9 shows how the CoAP retransmission timer impacts the device energy consumption when the RRC inactivity timer is gradually increased up to 20 s. As follows, provided the RRC inactivity timer is configured large enough to prevent the RRC connection between the NB-IoT device and eNodeB to prematurely expire and provoke extra RRC connection establishments, its impact on the device energy consumption is rather small.

Next, let us consider the simulations with an RRC inactivity timer configured so to avoid extra RRC connection establishments, i.e., those with a timer configured to 5 s or larger. Also RRC inactivity timer of 5 s is efficient for longer RTTs as uplink delay is higher than a downlink delay due to connection setup procedure and further the RRC inactivity timer starts after the transmission of a message. We note that although the higher end-to-end RTT results in some CoAP retransmissions, the low packet-loss rate of 1% over the PGW-AS Internet path, and the fact that the RRC connection is never prematurely torn down, make these come cheap in terms of energy consumption.

That is, the CoAP retransmission timer has only a small impact on the device energy consumption under these circumstances.

2) Partial DRX Support: Next, we consider how the impact on energy consumption of the same protocol-stack parameters as were studied in Section IV-B1, i.e., in scenarios with cDRX support, changes when we consider scenarios in which cDRX is not used. In these scenarios, an NB-IoT device will continuously monitor the Physical Downlink Control Channel (PDCCH) in RRC_CONNECTED.




Fig. 10: Energy consumption for an NB-IoT device with partial DRX support vs. random access delay for different RRC inactivity timers.

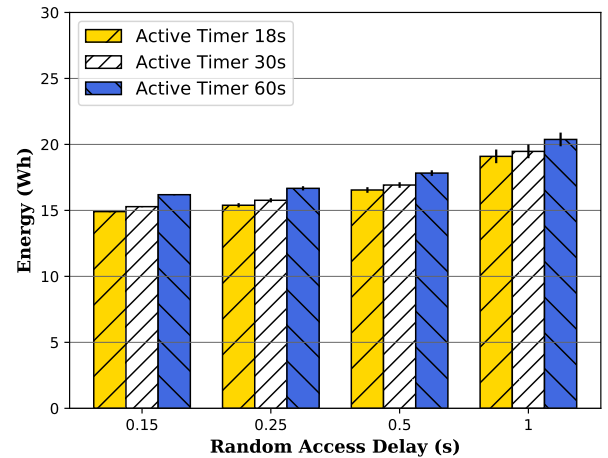


Fig. 11: Energy consumption for an NB-IoT device with partial DRX support vs. random access delay for different active timers.

a) RRC Inactivity Timer, Active Timer, and eDRX:

Figure 10 illustrates how the device energy consumption varies with increasing random access delays and with different settings on the RRC inactivity timer when cDRX is not used.

We compare these results with the ones obtained with full DRX support, i.e., the results in Figure 7, and observe that keeping the RRC inactivity timer enabled is still important to conserve energy, however, in contrast to before, it is this time also important not to set the timer to a too large value; increasing the value of the RRC inactivity timer beyond what is needed to prevent a premature expiration of the RRC inactivity timer, implies that the NB-IoT device listens longer than necessary on the radio channel, something that impacts heavily on its energy consumption. We also observe that as before the device energy consumption increases significantly with an increased random access delay.

Next, let us consider the impact of the active timer and eDRX cycle on the energy consumption of an NB-IoT device when cDRX is not used. Figure 11 corresponds with Figure 8, and shows how the active timer influences the energy consumption of an NB-IoT device when cDRX is not used and the RRC inactivity timer is 20 s. Evidently, cDRX has little to no impact on the way the active timer impacts the device energy consumption. Moreover, it seems that its influence over the device energy consumption does not change with an increasing random access delay. Similarly, the length of the eDRX cycle seems to have the same impact on the energy consumption of an NB-IoT device, irrespective of support for cDRX: Again, when we run the simulations in Figure 11 with eDRX cycles of lengths, 10.24 s, 20.48 s, and 40.96 s, the device energy consumption only marginally goes up with increasing eDRX-cycle lengths.

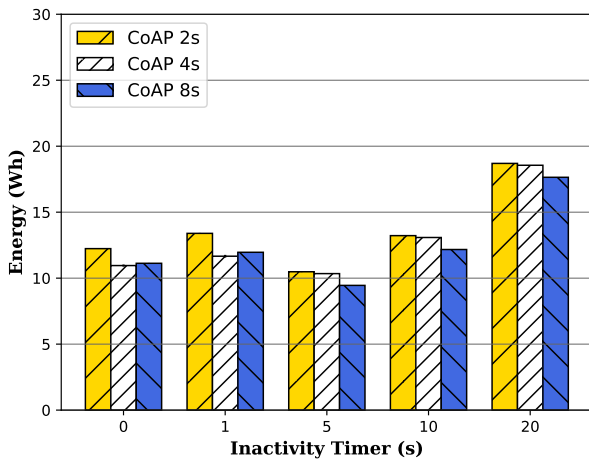


Fig. 12: Energy consumption for an NB-IoT device with partial DRX support vs. RRC inactivity timer for different CoAP retransmission timers with an end-to-end RTT of 8 s.

b) *CoAP Retransmission Timer*: Figure 12 shows, in the same way as Figure 9, how the CoAP retransmission timer and the RRC inactivity timer impact the energy consumption of an NB-IoT device, however, this time when cDRX is not used. As before, we consider an end-to-end RTT of 8 s. As follows, cDRX does not seem to change the effect of the CoAP

retransmission timer on the device energy consumption, and our simulations suggest that this finding holds true irrespective of the value of the RRC inactivity timer. Moreover, as was previously observed (cf. Figure 10), a proper configuration of the RRC inactivity timer is particularly important when cDRX is not used. Still, in our simulations, the configuration of the RRC inactivity timer does not seem to change the effect of the CoAP retransmission timer on the device energy consumption.

V. CONFIGURATION GUIDELINES

This section tries to summarize our findings from Section IV by formulating a number of configuration guidelines for NB-IoT devices. Still, it should be recognized that configuring an NB-IoT device to conserve energy and extend its battery lifetime is not always an easy task, not least since it depends on several interrelated factors, many of which are not even controllable. With this and the limitations pointed out in Section III in mind, Section V-A offers guidelines for scenarios with full DRX support and Section V-B provides guidelines for scenarios with partial DRX support.

A. Full DRX Support

1) *Traffic Properties*: We have found that provided the traffic intensity is kept low – our simulations suggest 1 packet every four hours or less – the energy consumption of an NB-IoT device in scenarios with full DRX support could be kept low enough to enable 10 years of operation. Still, other parameters also need to be properly configured. Moreover, our simulations suggest that sending packets in a burst is indeed efficient. However, there is still a limitation on how much energy that can be reduced. In particular, we have to keep in mind that an NB-IoT device consumes around 1.2-Wh/10 years even when it camps in the PSM state.

2) *RRC Inactivity Timer*: It is always efficient to keep the radio active during a reliable message exchange, i.e., when we expect transmitted messages to be acknowledged by the receiving end. However, keeping the device connected after it has received an acknowledgement from the application server is inefficient. A proper setting of the RRC inactivity timer seems to be the length of the expected end-to-end RTT. In our simulations that meant setting the RRC inactivity timer to a value between 1 s and 10 s. Still, it should be noted that letting an NB-IoT device stay a bit longer than the expected end-to-end RTT in RRC_CONNECTED, when it experiences long and varying RTTs, could conserve energy, since it otherwise runs the risk of having to re-establish the connection in order to receive outstanding CoAP acknowledgements.

3) *Active Timer*: The active timer is enabled when the NB-IoT device enters the RRC_IDLE state, i.e., where eDRX makes sure the device energy consumption is kept low, thus the configuration of the active timer is not crucial. To this end, we recommend an active timer that is at least set long enough to cover two eDRX listening periods.

4) *CoAP Retransmission Timer*: The CoAP retransmission timer has an impact on the energy consumption of an NB-IoT device in those cases it is configured shorter than the expected

end-to-end RTT. In those cases, it generates several spurious retransmissions. Still, it only has a significant impact on the energy consumption when also the RRC inactivity timer is set too short. Under these circumstances, several of the spurious CoAP retransmissions trigger RRC connection establishments. In view of this, we suggest setting the CoAP retransmission timer to a value a bit larger than the expected end-to-end RTT.

B. Partial DRX Support

1) *RRC Inactivity Timer*: Our simulations suggest that the configuration of the RRC inactivity timer becomes much more important when an NB-IoT device does not use cDRX. Still, the same recommendation as when cDRX is used holds, i.e., it should be set to the length of the expected end-to-end RTT. On the other hand when cDRX is not used, the cost in terms of consumed energy of setting the RTT inactivity timer to a value much larger than the expected end-to-end RTT in order to prevent premature RRC connection releases is non-negligible.

2) *Active Timer and CoAP Retransmission Timer*: The active timer and the CoAP retransmission timer seem to have more or less the same impact on the energy consumption of an NB-IoT device independent of the usage of cDRX, so the same recommendations hold true.

VI. RELATED WORK

This section surveys related work in three broad categories: computer simulations, theoretical and mathematical approaches, and device measurements in LTE and 5G Networks.

A. Computer Simulations

As part of their proposal for a prediction-based energy saving mechanism, Prediction-based energy saving mechanism (PBESM), Jinseong Lee and Jaiyong Lee [9] simulate the impact of the NB-IoT energy-saving mechanisms, including PSM and eDRX, as well as the lower parts of the NB-IoT protocol stack, on the energy consumption of a device. Moreover, Sultania *et al.* [17], [18], Soussi *et al.* [3], and Martiradonna *et al.* [11] simulate the NB-IoT energy-saving mechanisms, however, they omit the upper parts of the NB-IoT stack. Our simulation study extends these studies by also considering the upper parts of the NB-IoT protocol stack, by taking into account the impact of different traffic patterns on the NB-IoT energy consumption, and, not least, by providing a more detailed modeling of the RRC signaling procedures.

B. Theoretical and Mathematical Approaches

Several works have theoretically studied the NB-IoT energy-saving mechanisms. Early work was conducted by Zhou *et al.* [25], and by Hawar Ramazanalil and Alexey Vinel [14]. Both these works propose fairly comprehensive Markov models over the Discontinuous Reception (DRX) mechanism in LTE/LTE-Advance that are also largely applicable for NB-IoT. More recently, Maldonado *et al.* [1], [2] proposed a Markov model for the operation modes of an NB-IoT device, i.e., RRC_CONNECTED/IDLE, Inactive, and Communication, and, on the basis of this model, studied the lifetime of

a battery in an NB-IoT device under different circumstances. Also Sultania *et al.* [17], Soussi *et al.* [3], and Jörke *et al.* [7] have mathematically modeled the NB-IoT energy-saving mechanisms and their impact on the energy consumption.

C. Device Measurements in LTE and 5G Networks

There are a few publicly reported energy-consumption measurements on NB-IoT devices: Lauridsen *et al.* [8] and Hertlein *et al.* [5]. Lauridsen *et al.* conducted their measurements in a controlled, emulated environment, and estimated the battery lifetime of the NB-IoT device through energy consumption and battery lifetime models [8]. Moreover, Hertlein *et al.* carried out their measurements in a testbed environment, however, in contrast to Lauridsen *et al.*, they considered the energy consumption of an NB-IoT device in different transmission modes, e.g., RRC_CONNECTED and RRC_IDLE modes, rather than trying to predict its operation lifetime. Andres-Maldonado *et al.* [1] evaluated the energy consumption for NB-IoT devices for small data transmissions by CP CIoT EPS optimization in LTE from 3GPP Release 13. Results showed a battery life of more than 2 years in bad coverage, and up to 8 years with good coverage. When Non-Access Stratum (NAS) RAI is investigated from the Release 13, the results showed that keeping the device in DRX will decrease the battery lifetime. NAS RAI enables the MME to know if there is more pending traffic and thus makes it possible for an NB-IoT device to minimize the time it spends in RRC_CONNECTED. In that way, NAS RAI is able to increase the battery lifetime by up to 78% compared to a normal coverage.

Our simulation study complements these measurement studies by providing a more extensive treatment of the impact of pertinent parameters in NB-IoT on energy consumption —not least tunable parameters of the NB-IoT protocol stack, and by explicitly aim for guidelines on how to energy efficiently tune the NB-IoT protocol stack.

VII. CONCLUSION

This paper studies the energy consumption of an NB-IoT device, and how the energy consumption is influenced by several tunable and non-tunable factors, including traffic intensity and burstiness, packet loss, the cDRX, iDRX, and eDRX configurations, the RRC inactivity timer, the CoAP retransmission timer, etc. Evidently, full DRX support, i.e., using cDRX is key for an NB-IoT device to conserve energy. Moreover, the tuning of the RRC inactivity timer has a significant impact on the device energy consumption, irrespective of DRX support, i.e., irrespective of whether cDRX is used or not. In contrast, our study suggests that the configuration of the CoAP retransmission timer and the eDRX cycle have limited impact on the device energy consumption, and this independent of the use of cDRX. Furthermore, our study indicates that several, many times, non-tunable factors, e.g., radio conditions, traffic intensity, and traffic burstiness, have a significant impact on the energy consumption of an NB-IoT device. Thus, the placement of the device and the design of the

application running on the device could be equally important as tuning the NB-IoT protocol stack, and sometimes even more important. On the basis of the study presented in this paper, guidelines are provided on how to configure an NB-IoT device in a way that conserves energy and prolongs the lifetime of its battery. Guidelines are provided for cases with both full and partial DRX support, i.e., for cases when cDRX is and is not used. As part of our future work, we intend to compare the energy consumption for an NB-IoT device with full DRX support, with a device with partial DRX support, but which supports RAI. We also intend to complement our simulations with live energy measurements in an NB-IoT testbed and in commercial cellular networks.

ACKNOWLEDGEMENT

This research has received funding from the Swedish Knowledge Foundation's research profile, High Quality Networked Services in a Mobile World (HITS), and from European Union's Horizon 2020 research and innovation programme under grant agreement No. 815178 (5GENESIS). The views expressed are solely those of the authors.

REFERENCES

- [1] P. Andres-Maldonado, P. Ameigeiras, J. Prados-Garzon, J. Navarro-Ortiz, and J. M. Lopez-Soler, "Narrowband IoT Data Transmission Procedures for Massive Machine-Type Communications," *IEEE Network*, vol. 31, no. 6, pp. 8–15, Nov. 2017.
- [2] P. Andres-Maldonado, P. Ameigeiras, J. Prados-Garzon, J. J. Ramos-Munoz, and J. M. Lopez-Soler, "Optimized LTE Data Transmission Procedures for IoT: Device Side Energy Consumption Analysis," in *2017 IEEE International Conference on Communications Workshops (ICC Workshops)*, May 2017, pp. 540–545.
- [3] M. El Soussi, P. Zand, F. Pasveer, and G. Dolmans, "Evaluating the Performance of eMTC and NB-IoT for Smart City Applications," in *2018 IEEE International Conference on Communications (ICC)*, May 2018, pp. 1–7.
- [4] Ericsson AB, "Ericsson Mobility Report," 2020. [Online]. Available: <https://www.ericsson.com/49da93/assets/local/mobility-report/documents/2020/june2020-ericsson-mobility-report.pdf>
- [5] M. Hertlein, S. Breun, G. Cappel, A. Schwarzmeier, F. Lurz, R. Weigel, and G. Fischer, "Evaluation of Cellular Standards for Low Data Rate Applications Regarding Power Consumption and Timing Parameters," in *2018 IEEE Radio and Wireless Symposium (RWS)*, Jan. 2018, pp. 217–219.
- [6] ISO/IEC 20922, "Information technology Message Queuing Telemetry Transport (MQTT)," 2016. [Online]. Available: <https://standards.iso.org/ittf/PubliclyAvailableStandards/index.html>
- [7] P. Jörke, R. Falkenberg, and C. Wietfeld, "Power Consumption Analysis of NB-IoT and eMTC in Challenging Smart City Environments," in *2018 IEEE Globecom Workshops*, Dec. 2018, pp. 1–6.
- [12] MulteFire Alliance, "MulteFire Release 1.1 Specification," 2 2019. [Online]. Available: <https://www.multefire.org/release-1-1-specification-download/>
- [8] M. Lauridsen, R. Krigslund, M. Rohr, and G. Madueno, "An Empirical NB-IoT Power Consumption Model for Battery Lifetime Estimation," in *2018 IEEE 87th Vehicular Technology Conference (VTC Spring)*, Jun. 2018, pp. 1–5.
- [9] Lee, Jinseong and Lee, Jaiyong, "Prediction-based Energy Saving Mechanism in 3GPP NB-IoT Networks," *Sensors*, vol. 17, p. 2008, 9 2017.
- [10] B. Martinez, F. Adelantado, A. Bartoli, and X. Vilajosana, "Exploring the performance boundaries of NB-IoT," *IEEE Internet of Things Journal*, vol. 6, no. 3, pp. 5702–5712, 2019.
- [11] S. Martiradonna, A. Grassi, G. Piro, L. A. Grieco, and G. Boggia, "An Open Source Platform for Exploring NB-IoT System Performance," in *European Wireless 2018; 24th European Wireless Conference*, May 2018, pp. 1–6.
- [13] *OMA Lightweight Machine to Machine protocol Version 1.2*, Open Mobile Alliance, Feb. 2019, Available online at www.openmobilealliance.org.
- [14] H. Ramazanali and A. Vinel, "Performance Evaluation of LTE/LTE-A DRX: A Markovian Approach," *IEEE Internet of Things Journal*, vol. 3, no. 3, pp. 386–397, Jun. 2016.
- [15] E. Rescorla and N. Modadugu, "Datagram Transport Layer Security Version 1.2," Internet Requests for Comments, RFC Editor, RFC 6347, Jan. 2012, <http://www.rfc-editor.org/rfc/rfc6347.txt>. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc6347.txt>
- [16] Z. Shelby, K. Hartke, and C. Bormann, "The Constrained Application Protocol (CoAP)," Internet Requests for Comments, RFC Editor, RFC 7252, Jun. 2014, <http://www.rfc-editor.org/rfc/rfc7252.txt>. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc7252.txt>
- [17] A. K. Sultania, P. Zand, C. Blondia, and J. Famaey, "Energy Modeling and Evaluation of NB-IoT with PSM and eDRX," in *2018 IEEE Globecom Workshops (GC Wkshps)*, Dec. 2018, pp. 1–7.
- [18] A. K. Sultania, C. Delgado, and J. Famaey, "Implementation of nb-iot power saving schemes in ns-3," in *Proceedings of the 2019 Workshop on Next-Generation Wireless with ns-3*, 2019, pp. 5–8.
- [19] *Non-Access-Stratum (NAS) Protocol for Evolved Packet System (EPS)*, Technical Report 3rd Generation Partnership Project (3GPP), 2008, Technical Report TS 24.301.
- [20] *Evolved Universal Terrestrial Radio Access (E-UTRA); Medium Access Control (MAC) Protocol Specification*, Technical Report 3rd Generation Partnership Project (3GPP), 2012, Technical Report TS 36.321 Version 8.12.0 Release 8.
- [21] *Cellular System Support for Ultra-low Complexity and Low Throughput Internet of Things (CIoT)*, Technical Report 3rd Generation Partnership Project (3GPP), 2015, 3GPP TS 45.820 Version 13.1.0.
- [22] *NB-LTE - Battery Lifetime Evaluation*, Technical Report 3rd Generation Partnership Project (3GPP), 2015, RP-151393.
- [23] *Radio Resource Control (RRC); Protocol specification*, Technical Report 3rd Generation Partnership Project (3GPP), 2017, 3GPP TS 36.331 Version 15.8.0.
- [24] A. Varga and R. Hornig, "An Overview of the OMNeT++ Simulation Environment," in *1st International ICST Conference on Simulation Tools and Techniques for Communications, Networks and Systems*. ICST, 2010.
- [25] K. Zhou, N. Nikaein, and T. Spyropoulos, "LTE/LTE-A Discontinuous Reception Modeling for Machine Type Communications," *IEEE Wireless Communications Letters*, vol. 2, no. 1, pp. 102–105, Feb. 2013.