NB-IoT Random Access: Data-driven Analysis and ML-based Enhancements

Giuseppe Caso, Konstantinos Kousias, Özgü Alay, Anna Brunstrom and Marco Neri

***Abstract*—In the context of massive Machine Type Communications (mMTC), the Narrowband Internet of Things (NB-IoT) technology is envisioned to efficiently and reliably deal with massive device connectivity. Hence, it relies on a tailored Random Access (RA) procedure, for which theoretical and empirical analyses are needed for a better understanding and further improvements. This paper presents the first data-driven analysis of NB-IoT RA, exploiting a large scale measurement campaign. We show how the RA procedure and performance are affected by network deployment, radio coverage, and operators' configurations, thus complementing simulation-based investigations, mostly focused on massive connectivity aspects. Comparison with the performance requirements reveals the need for procedure enhancements. Hence, we propose a Machine Learning (ML) approach, and show that RA outcomes are predictable with good accuracy by observing radio conditions. We embed the outcome prediction in a RA enhanced scheme, and show that optimized configurations enable a power consumption reduction of at least 50%. We also make our dataset available for further exploration, toward the discovery of new insights and research perspectives.**

***Index Terms*—Cellular Internet of Things, massive Machine Type Communications, Narrowband Internet of Things, Random Access, Empirical Analysis**

I. INTRODUCTION

Since its introduction by 3rd Generation Partnership Project (3GPP) in Release 13 (Rel-13), the Narrowband Internet of Things (NB-IoT) technology is gaining momentum as a leading solution in the context of Low Power Wide Area Networks (LPWANs) [1]. Along with other technologies, NB-IoT enables IoT services for massive Machine Type Communications (mMTC), including smart city and industrial automation use cases, by exploiting the existing cellular architecture in a low cost and power efficient manner [2]–[6].

In order to cope with the massive access requirement, the NB-IoT design has emphasized the importance of providing reliable connectivity, hence proposing a tailored Random Access (RA) procedure. Even though NB-IoT RA takes its root from Long Term Evolution (LTE) RA, there still exists several peculiarities and further challenges, not only in terms of increased access requests, but also with respect to the environmental scenarios in which many NB-IoT devices are supposed to operate (e.g., deep indoor). For this reason, the research community is increasingly formalizing and analyzing several NB-IoT RA-related aspects. In particular, initial investigations have led to the definition of theoretical performance models, often validated by simulations, which emphasize the massive connectivity aspect [7]–[20] (see §VI for details).

In parallel, many mobile operators are launching NB-IoT networks worldwide [21], making possible to execute field trials and measurement campaigns. These enable a better understanding of the system, including the RA procedure, toward identifying correlations between deployment and performance, and in turn deriving new guidelines for further system enhancements. To this end, data-driven analyses are crucial to understand the complexities of modern communication systems; however, extensive measurement campaigns are scarcely available to researchers, who thus often opt for sub-optimal simulation-based approaches.

Given the above motivations, we provide in this paper the first data-driven analysis and enhancement of NB-IoT RA. To do so, we exploit a large scale measurement campaign conducted in the city of Oslo during 2019, which includes measurements from two mobile operators and corresponding network deployments. To the best of our knowledge, this is the first analysis on NB-IoT operational networks to specifically consider RA procedure and performance across heterogeneous scenarios. The main contributions of this paper are:

- We conduct a thorough analysis on NB-IoT RA, revealing how its performance and outcomes depend on network deployment, radio coverage, and specific configurations adopted by the operators. The results in this paper complement RA studies in the current literature, which mainly emphasize the multiple access aspect and adopt simulation-based analyses;
- We propose a Machine Learning (ML) approach to enhance the RA procedure, demonstrating that a reliable prediction of its outcomes is possible, and can be used to optimize the operations composing the procedure. Our results show that, with the proposed approach, the RA power consumption can be reduced by at least 50%;
- We open-source our dataset [22], which comprises of NB-IoT RA measurements for two Norwegian operators, collected across heterogeneous environmental scenarios. Along with the dataset in [23], which focuses on network deployment and coverage [24], we thus make available a large amount of NB-IoT data for further exploration by the research community.

The rest of the paper is organized as follows. A background on NB-IoT is provided in §II, in terms of technology overview (§II-A) and RA description (§II-B). §III presents the experimental design, executed measurement campaign, and collected dataset. RA performance analysis is reported in §IV,

G. Caso is with Simula Metropolitan CDE.
K. Kousias is with Simula Research Laboratory.
Ö. Alay is with University of Oslo and Simula Metropolitan CDE.
A. Brunstrom is with Karlstad University.
M. Neri is with Rohde&Schwarz.

in which we first study preliminary affecting factors (§IV-A), and then present and discuss RA results (§IV-B). ML-based analysis and enhancement are introduced and discussed in §V. We report related literature in §VI, contrasting with our contributions. We finally conclude the paper in §VII.

II. Background

A. NB-IoT Technology

NB-IoT is a radio interface implemented over the cellular licensed spectrum that leverages the existing architecture [2]. In the following we report the most important features standardized in Rel-13, complemented by the enhancements introduced in Rel-14 and Rel-15 [3] [25].

NB-IoT operates over either a 200 kHz Global System for Mobile Communications (GSM)-like channel or an LTE Physical Resource Block (PRB) of 180 kHz, allowing coexistence with both technologies. It can adopt three operation modes:

- *stand-alone* over a 200 kHz channel in GSM spectrum;
- *in-band* over a single PRB within a set of LTE PRBs;
- *guard-band* over a PRB within a guard band among different sets of LTE PRBs.

After selecting a mode, the operators can provide NB-IoT services via a software upgrade of their infrastructure, i.e., reconfiguring their LTE evolved Node Bs (eNBs) and cells.

Downlink (DL) and Uplink (UL) resources are accessed in Frequency Division Duplex (FDD), but also in Time Division Duplex (TDD) since Rel-15. Orthogonal Frequency Division Multiple Access (OFDMA) is applied in DL, with 15 kHz subcarrier spacing and Cyclic Prefix (CP). The PRB is divided into seven OFDM symbols of twelve subcarriers each, and occupies in time a slot of 0.5 ms. Two slots sum up into a subframe, which is the smallest DL scheduling unit. DL channels and signals include:

- Primary and Secondary Synchronization Signals (NPSS[1] and NSSS), which allow device/cell synchronization;
- Physical Broadcast Channel (NPBCH), which carries Master Information Blocks (MIBs), broadcasting high-level network configurations;
- Physical Downlink Control Channel (NPDCCH), which carries Downlink Control Information (DCI) messages;
- Physical Downlink Shared Channel (NPDSCH), which is used to transmit data and further configurations via System Information Blocks (denoted SIBx, where x identifies a SIB transmitting specific configurations, e.g., SIB2 provides RA-related settings, see §II-B).

Cell-dependent portions of NPBCH, NPDCCH, and NPDSCH are used to transmit the Reference Signal (NRS), which allows to estimate propagation conditions.

Single Carrier Frequency Division Multiple Access (SC-FDMA) is instead applied in UL. The subcarrier spacing can be either 15 kHz or 3.75 kHz, in which case the PRB contains 48 subcarriers and lasts 2 ms. UL channels include:

- Physical Random Access Channel (NPRACH), which allows to initiate the connection toward a cell, i.e., the RA procedure (§II-B);
- Physical Uplink Shared Channel (NPUSCH), which carries data and control information.

In UL, sub-PRB allocation is possible, i.e., less than twelve subcarriers can be allocated. In particular, *single-tone* (single-subcarrier) transmissions are enabled in both spacing modes, while 12, 6, and 3 *multi-tone* transmissions are available for 15 kHz spacing. NPRACH is always 3.75 kHz-spaced and single-toned, while NPUSCH configurations depend on radio conditions and operators' settings.

NB-IoT targets high service reliability and delay-tolerant data exchange mostly in UL. For this reason, advanced modulation and coding are not supported. Rather, Quadrature Phase Shift Keying (QPSK) and Binary Phase Shift Keying (BPSK)/QPSK modulations are adopted on each DL/UL subcarrier, respectively; then, so-called Coverage Enhancement (CE) techniques are used to improve connectivity in harsh environments. On the one hand, a power boost is obtained by narrowing down the signal bandwidth, at the cost of a low data rate. On the other hand, NB-IoT uses repeated transmissions to increase the probability of correct reception. In particular, DL and UL messages can be repeated up to 2048 and 128 times, respectively. As described in §II-B1, the device estimates a so-called Coverage Level (CL) during RA, considering its radio conditions and operator's configurations, and the number of repetitions depends on the CL.

NB-IoT devices operate in *idle* or *connected* modes. While in idle, they trigger the procedures for switching into connected, including cell (re-)selection, RA, and DL paging monitoring. While in connected, they exchange data and continue the monitoring. Finally, aiming at high energy efficiency and long battery lifetime, two energy saving schemes have been introduced, that is, extended Discontinuous Reception (eDRX) and Power Saving Mode (PSM), for which detailed descriptions and analyses are given in [26].

B. NB-IoT Random Access

NB-IoT RA is described in this section, along with the preliminary operations affecting its execution and outcomes.

1) Preliminary operations: Among pre-RA operations, cell (re-)selection and CL estimation play a key role, and are thus described in the following.

a) Cell (re-)selection: Cell selection enables a device to identify, synchronize to, and determine the suitability of a cell. Once the device has decoded the cell identity, it determines its suitability using two SIB1 parameters, that is, minimum required signal strength ($Q_{\mathsf{rx,min}}$) and quality ($Q_{\mathsf{qual,min}}$) [27]. A cell n is considered suitable *iff*:

$$\begin{cases} S_{\mathsf{rx},n} := \mathsf{RSRP}_n - Q_{\mathsf{rx,min}} > 0 \\ S_{\mathsf{qual},n} := \mathsf{RSRQ}_n - Q_{\mathsf{qual,min}} > 0 \end{cases} \tag{1}$$

where $S_{\mathsf{rx},n}$ and $S_{\mathsf{qual},n}$ are signal strength and quality indicators obtained by observing Reference Signal Received Power (RSRP [dBm]) and Quality (RSRQ [dB]) for cell n, i.e., RSRP_n and RSRQ_n in Eq. (1), respectively.

If multiple surrounding cells are suitable, the device selects the strongest in terms of RSRP.

[1] *N* stands for Narrowband in all acronyms.

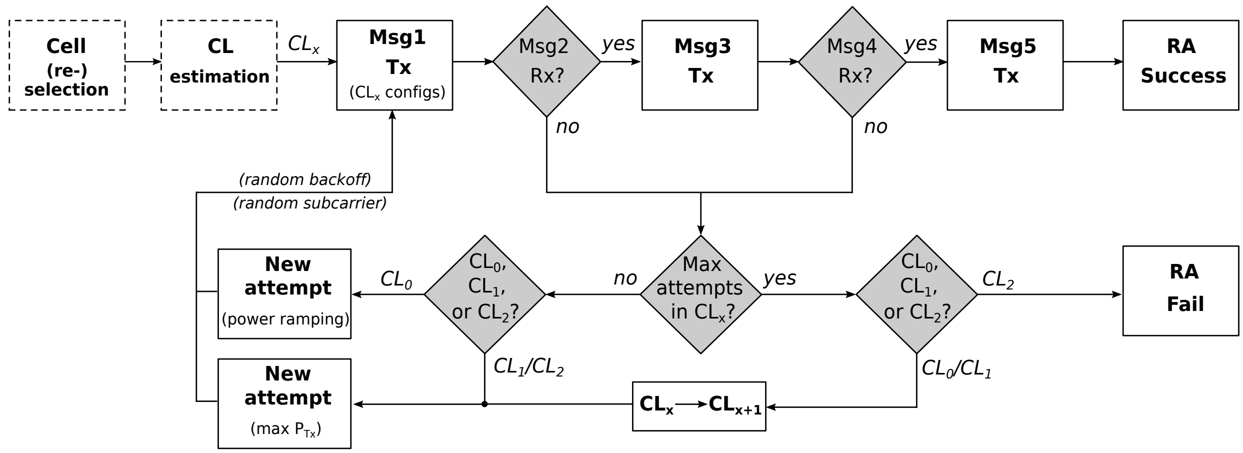

Fig. 1: Block diagram representing NB-IoT RA procedure. Dashed blocks depict pre-RA operations, i.e., cell (re-)selection and CL estimation.

Cell re-selection supports connection mobility during the idle mode. If the device has selected cell n, it uses two SIB3 parameters, $S_{\mathsf{IntraSearchP}}$ and $S_{\mathsf{NonIntraSearch}}$, to decide if a re-selection is needed in the same frequency band (*Intra*) and/or in different bands (*NonIntra*). Re-selection is triggered in both Intra/NonIntra modes *iff*:

$$\begin{cases} S_{\mathsf{rx},n} \leq S_{\mathsf{IntraSearchP}} \\ S_{\mathsf{rx},n} \leq S_{\mathsf{NonIntraSearch}} \end{cases} \tag{2}$$

where $S_{\mathsf{rx},n}$ follows Eq. (1).

In this case, the device monitors several cells; then, it decides to camp on a new cell n' if it is suitable[2], and:

$$\mathsf{RSRP}_{n'} - Q_{\mathsf{off}} > \mathsf{RSRP}_n - Q_{\mathsf{hyst}}, \tag{3}$$

where Q_{hyst} is an hysteresis value preventing ping-pong re-selections and shared in SIB3, while Q_{off} is an offset parameter, specifically applied to NonIntra re-selections and shared in SIB5. In Intra mode, the condition in Eq. (3) should hold for a time interval at least equal to t_{resel}, shared in SIB3. The t_{resel} value for NonIntra mode is instead shared in SIB5.

Aiming at energy savings, stationary devices with high RSRP from the selected cell could avoid re-selections. Rel-14 makes also possible to avoid re-selections throughout a day, if the device observes limited RSRP variation in its current cell. Rel-15 further allows RSRP evaluation on NSSS and NPBCH signals, thus reducing the time for obtaining RSRP values.

b) CL estimation: Once a cell is selected, the device estimates its CL by observing the RSRP, and comparing it with operator-specific thresholds [28]. The standard allows up to two RSRP thresholds, leading to three possible CLs: CL_0 represents LTE-like radio conditions, while CL_1 and CL_2 apply to challenging scenarios, e.g., deep indoor. Assuming two thresholds Th_1 and Th_2 in dB (Th_1, $\mathsf{Th}_2 > 0$, and $\mathsf{Th}_1 > \mathsf{Th}_2$), the device estimates its CL as follows:

$$\mathsf{CL_x} = \begin{cases} \mathsf{CL}_0 & \text{if} \quad \mathsf{RSRP} \geq -140 + \mathsf{Th}_1 \\ \mathsf{CL}_1 & \text{if} \quad -140 + \mathsf{Th}_2 < \mathsf{RSRP} < -140 + \mathsf{Th}_1 \\ \mathsf{CL}_2 & \text{if} \quad \mathsf{RSRP} \leq -140 + \mathsf{Th}_2 \end{cases} \tag{4}$$

where -140 dBm is the receiver sensitivity [28] [29].

2) RA procedure: The RA procedure is depicted in Fig. 1. Once a cell is selected and a CL estimated, the device transmits a *preamble* (*Msg1*) on NPRACH during the first available RA Opportunity (RAO). As clarified later, *Msg1* is transmitted with CL-dependent configurations, and preambles from multiple devices may collide in this phase. If the cell detects *Msg1* of a device, it replies with a Random Access Response (RAR, or *Msg2*) on NPDSCH. With RAR, the device gets the UL grant for transmitting *Msg3* on PUSCH, containing a device identifier (C-RNTI).[3] With *Msg4*, the cell resolves any contention due to preamble collision, transmitting a connection setup command with the C-RNTI of the contention-winning device. The winning device infers the RA success and replies with *Msg5*, finalizing its transition into connected mode and proceeding with data exchange. Adopting a *backoff* scheme, the other devices try to access after waiting for a backoff time, randomly selected in a window between zero and a maximum configurable value. Then, they transmit a new attempt in the next available RAO. Aiming at better handling overloading situations, the *access barring* scheme can be also adopted, where the cell broadcasts an access barring factor that may depend on network conditions. The devices start a new RA if they randomly extract a value lower than this factor. Both schemes are detailed in [31].

We conclude the overview by noticing that the RA procedure is clearly needed in idle mode but, if required, it can be also executed by already connected devices, in order to improve connection reliability while reducing the data rate [2].

Next we further detail some important aspects related to RA, i.e., NPRACH channel, CL adjustment, and power control.

a) NPRACH: As anticipated in the RA description, NPRACH is the time-frequency resource on which the preamble is transmitted. While LTE PRACH is multi-toned, with 1.25 kHz subcarrier spacing and 1.05 MHz bandwidth [32], NPRACH is single-toned, with 3.75 kHz subcarrier spacing and access regulated by Frequency Hopping (FH) [7]. A cell can configure up to 48 NPRACH subcarriers, meaning that there are up to 48 available preambles with unique FH patterns. The devices access NPRACH in multi-channel slotted ALOHA; hence, they may select the same preamble and collide. The preamble consists of four symbol groups, each

[2]Suitability check during re-selection follows Eq. (1), but may adopt different parameters shared in SIB3 [27].

[3]Rel-15 Early Data Transmission (EDT) also allows to transmit (small amount of) data encapsulated in *Msg3* [30].

TABLE I: NB-IoT RA-related terminology adopted in this paper, in alphabetical order.

Term / Acronym	**Description**
$\mathsf{CL_x}$	Coverage Level x, with $\mathsf{x} \in [0, 2]$ (CL estimation follows Eq. (4))
C-RNTI	Cell Radio Network Temporary Identifier, it identifies the RA contention-winning device
Δ_{ramp}	Power ramping parameter for successive preamble attempts in $\mathsf{CL_0}$ (shared in SIB2)
MIB / SIB	Downlink periodic messages carrying network configurations
Msg1–Msg5	Downlink and Uplink messages composing a RA execution
N_{att}	Total number of possible preamble attempts in a RA execution (shared in SIB2)
$N_{\mathsf{att}}^{\mathsf{CL_x}}$	Total number of possible preamble attempts in a RA execution using $\mathsf{CL_x}$ configurations (shared in SIB2)
$N_{\mathsf{rep}}^{\mathsf{CL_x}}$	Total number of repetitions of a preamble attempt using $\mathsf{CL_x}$ configurations (shared in SIB2)
NPDSCH	Downlink channel used for Random Access Response (RAR, or *Msg2*) transmission
NPRACH	Uplink channel used for preamble (*Msg1*) transmission.
NPUSCH	Uplink channel used for request-to-connect (*Msg3*) transmission
NRS	NB-IoT Reference Signal
$P_{\mathsf{tx}}^{\mathsf{max}}$	Maximum device transmission power (20 or 23 dBm for Rel-13, also 14 dBm for Rel-14)
$P_{\mathsf{tg}}^{\mathsf{NPRACH/NPUSCH}}$	Target received power on NPRACH and NPUSCH (used in Eqs. (5) and (7), shared in SIB2)
$P_{\mathsf{tx}}^{\mathsf{NPRACH}}$	Power used to transmit a preamble attempt on NPRACH (defined in Eq. (5))
$P_{\mathsf{tx}}^{\mathsf{NPUSCH}}$	Power used to transmit *Msg3* and data on NPUSCH (defined in Eq. (7))
Q_{hyst}, t_{resel} (Intra mode)	Parameters used for Intra cell re-selection (used in Eq. (3), shared in SIB3)
Q_{off}, t_{resel} (NonIntra mode)	Parameters used for NonIntra cell re-selection (used in Eq. (3), shared in SIB5)
RAO	Random Access Opportunitiy (each time NPRACH resources are available, it depends on NPRACH periodicity)
RSRP, RSRQ	Received Signal Strength Power and Received Signal Strength Quality (their evaluation is standardized by 3GPP)
$S_{\mathsf{IntraSearchP}}$, $S_{\mathsf{NonIntraSearch}}$	Parameters used to trigger cell re-selection (used in Eq. (2), shared in SIB3)
S_{rx}, S_{qual}, $Q_{\mathsf{qual,min}}$, $Q_{\mathsf{rx,min}}$	Parameters used to check cell suitability (used in Eq. (1), $Q_{\mathsf{qual,min}}$ and $Q_{\mathsf{rx,min}}$ are shared in SIB1)
SINR	Signal to Interference plus Noise Ratio (its evaluation is vendor-specific, not standardized by 3GPP)
$\mathsf{Th_1}$, $\mathsf{Th_2}$	Threshold used for CL estimate in Eq. (4)

containing one CP and five symbols. The CP lasts either 66.7 μs or 266.67 μs. Each symbol has a duration of 266.67 μs, and the subcarrier for the first symbol group is randomly chosen from the set of subcarriers allocated to the CL the device is in. The next three symbol groups are transmitted using the FH pattern [7]. NPRACH periodicity (i.e., RAO period) also depends on the CL, reaching up to 2560 ms.

b) CL adjustment and Power control: The CL indicates the resources to use in NPRACH, and the number of repetitions to adopt for transmitting each preamble attempt. The amount of repetitions to use in a CL ($N_{\mathsf{rep}}^{\mathsf{CL_x}}$) is given in SIB2, along with RSRP thresholds for CL estimation and further RA configurations. These include the maximum number of preamble attempts within an entire procedure (N_{att}) and the number of preamble attempts for each CL ($N_{\mathsf{att}}^{\mathsf{CL_x}}$, with $N_{\mathsf{att}}^{\mathsf{CL_x}} \leq N_{\mathsf{att}}$). The device can adjust its initial CL estimate if it is not being able to access after $N_{\mathsf{att}}^{\mathsf{CL_x}}$ attempts. It thus moves into higher CLs, performing other attempts with a higher number of repetitions, until it is able to connect or reaches N_{att}. In $\mathsf{CL_0}$, the device also adopts the following power control to derive the transmission power for the first attempt, denoted $P_{\mathsf{tx}}^{\mathsf{NPRACH}}$ [dBm]:

$$P_{\mathsf{tx}}^{\mathsf{NPRACH}} = \min\{P_{\mathsf{tx}}^{\mathsf{max}}, P_{\mathsf{tg}}^{\mathsf{NPRACH}} + \mathsf{PL}\}, \tag{5}$$

where $P_{\mathsf{tx}}^{\mathsf{max}}$ and $P_{\mathsf{tg}}^{\mathsf{NPRACH}}$ represent maximum transmission and NPRACH target powers, respectively. PL [dB] is the experienced path loss, evaluated as follows:

$$\mathsf{PL} = P_{\mathsf{tx}}^{\mathsf{NRS}} - \mathsf{RSRP}, \tag{6}$$

where $P_{\mathsf{tx}}^{\mathsf{NRS}}$ is the power used by the cell to transmit NRS. If needed, $P_{\mathsf{tx}}^{\mathsf{NPRACH}}$ can be increased in the next $\mathsf{CL_0}$ attempts, adding a power ramping quantity equal to Δ_{ramp} [dB]. $P_{\mathsf{tg}}^{\mathsf{NPRACH}}$, $P_{\mathsf{tx}}^{\mathsf{NRS}}$, and Δ_{ramp} are also signalled in SIB2.

Rel-13 devices do not adopt power control/ramping in $\mathsf{CL_1}$ and $\mathsf{CL_2}$, and always use $P_{\mathsf{tx}}^{\mathsf{max}}$ instead, equal to either 20 or 23 dBm. Rel-14 extends power control/ramping to $\mathsf{CL_1}$, and also introduces a new device power class with $P_{\mathsf{tx}}^{\mathsf{max}} = 14$ dBm.

A similar power control is used to configure the NPUSCH power, denoted $P_{\mathsf{tx}}^{\mathsf{NPUSCH}}$ [dBm], as follows:

$$P_{\mathsf{tx}}^{\mathsf{NPUSCH}} = \min\{P_{\mathsf{tx}}^{\mathsf{max}}, 10\log_{10}(M) + P_{\mathsf{tg}}^{\mathsf{NPUSCH}} + \alpha\mathsf{PL} + \mathsf{c}\}, \tag{7}$$

where M depends on the NPUSCH bandwidth [2, Chapter 7], $0 \leq \alpha \leq 1$ is used to adjust the path loss contribution, and c groups other considered variables, as detailed in [33].

We conclude this section by reporting in Table I the most relevant terminology in the context of NB-IoT RA, selected from the terms used in §II and next sections.

III. Experimental Design

In this section, we present our measurement campaign, providing a description of the adopted hardware and software components, and an overview of the experimental setup and collected dataset.

A. Measurement System

We performed NB-IoT measurements in the area of Oslo, Norway, using the Rohde&Schwarz (R&S) TSMA6 toolkit, along with an Exelonix NarrowBand (NB) USB device and a Global Positioning System (GPS) antenna.

The TSMA6 system integrates a spectrum scanner and a laptop, where the controlling software, named ROMES4, is installed. The spectrum scanner enables passive measurements of all 3GPP mobile radio technologies up to 6 GHz, and supports NB-IoT signal decoding in all operation modes. Besides these functionalities, we leveraged two further TSMA6 features, i.e., a) *Automatic Channel Detection*, for automatically detecting all technologies in the specified spectrum, and b) *Base Transceiver Station (BTS) Position Estimation*, for estimating the position of cells and eNBs. This setup allowed a comprehensive collection of measurements related to network deployment and radio coverage [23] [24].

The Exelonix module is a Qualcomm-based IoT device supporting NB-IoT and LTE-M. We embedded the device with NB-IoT SIM cards of the operators under testing, and connected to TSMA6 via USB. By doing so, we were able to configure and monitor via ROMES4 the device operations, including cell (re-)selection, CL estimation, and RA, and keep track of the Quality of Service (QoS) of active measurements possibly executed after a successful RA. The connection to TSMA6 also enabled the observation of radio conditions from the device perspective, in parallel to the scanner and focusing on the serving operator/cell pair.

B. Measurement Campaign and Dataset

The measurement campaign covered a period of three weeks in summer 2019. We enabled the scanner to perform passive measurements on LTE Band 1, 3, 7, and 20, and corresponding guard bands. By doing so, we detected two LTE operators, denoted as Op1 and Op2, which also provide NB-IoT service in the guard bands of Band 20.

The collection was dissected into multiple sub-campaigns, each characterized by specific features in terms of location, time, and device operations. Considering location and time, we conducted the sub-campaigns in three reference scenarios, i.e., Deep Indoor (DI), for basements and deep enclosed spaces, Indoor (I), for houses and multi-floor buildings, and Outdoor (O), while walking or on public transport. We also extended the measurements over time, and hence the dataset includes morning, afternoon, evening, weekdays, and weekends sub-campaigns. Considering device operations, we exploited the possibility to configure the Exelonix module via ATtention (AT) Commands [34], and created three test cases in ROMES4. For each operator, we run the test cases in the above scenarios, in parallel with the scanner measurements. Considering the focus on RA, we developed a main test case in which the device performed repeated RA executions spaced out by short waiting times. In the other two test cases, the device was required to perform, after a successful RA, either a connectivity test via Internet Control Message Protocol (ICMP) ping, toward the Google Domain Name System (DNS) server located at 8.8.8.8, or a short data upload via File Transfer Protocol (FTP), toward a proprietary server located in Oslo. By doing so, we were able to collect a large number of RA executions, thus enabling a deep inspection of its functioning and performance, reported in §IV.

Overall, we performed 21, 69, and 24 sub-campaigns in DI, I, and O scenarios, respectively. Considering the passive measurements performed by the scanner, each sub-campaign contains parallel collection of NB-IoT and LTE coverage measurements, as well as network deployment data for the operators detected in the monitored bands. We provide in [23] the list of collected attributes and the complete dataset, with anonymized operator-specific information. As regards the active measurements executed through the Exelonix device, each sub-campaign includes anonymized information related to a specific operator's network. The sub-campaigns are almost equally split across operators, and each of them contains several RA-related attributes. These are described in [22], where we also provide the complete dataset and the configurations adopted for the test cases.

IV. Performance Evaluation

In this section, we present the analysis of our measurement campaign. First, we analyze preliminary factors affecting the RA procedure, including network deployment and coverage, along with cell (re-)selection and CL estimation. We then extensively discuss the RA performance.

A. Deployment, Coverage, and Preliminary Operations

In the following, we first analyze NB-IoT network deployment and radio coverage, by leveraging the spectrum scanner measurements. We then move on to cell (re-)selection and CL estimation, hence including the analysis of measurements obtained via the Exelonix module.

1) Network deployment and Radio coverage: Fig. 2a depicts the placement of NB-IoT eNBs for Op1 and Op2 in the Oslo area covered by our measurements, as provided by the TSMA6 BTS Position Estimation functionality. We find that Op1 features a higher number of eNBs compared to Op2 (146 vs. 107), implying a denser infrastructure; moreover, both operators use the existing LTE infrastructure for deploying NB-IoT. On this aspect, considering an average across operators, about 86% of the eNBs supporting LTE have been reconfigured, and now include at least one cell each providing NB-IoT services. For Op1 we detected a few NB-IoT-only eNBs, that could be explained by considering the more penetrating nature of NB-IoT signals compared to LTE.

Since the passive measurements by the scanner do not require any active operation toward cell selection and connection (which are performed by the Exelonix module), we define the operator's coverage in a measurement location as the highest RSRP and Signal to Interference plus Noise Ratio (SINR) values perceived among all the cells detected for that operator. By averaging across all locations in a sub-campaign, we finally obtain the sub-campaign average coverage.

Figs. 2b and 2c depict the sub-campaign average coverage for both operators in a boxplot format, in terms of RSRP (a) and SINR (b), across DI, I, and O scenarios. We find that both RSRP and SINR are significantly lower in DI environments. More specifically, when compared to I, the sub-campaign average RSRPs differ in their averages by 36.36 dB and 35.7 dB for Op1 and Op2, respectively, while the corresponding numbers are 10 dB and 9.88 dB for SINR. The results highlight the negative effect of DI on the signal propagation, which needs to be mitigated via CE techniques. The deviation

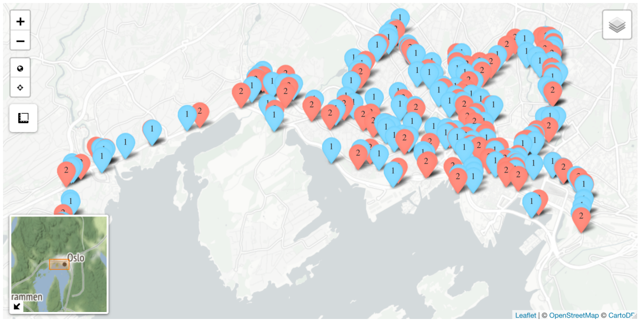

(a)

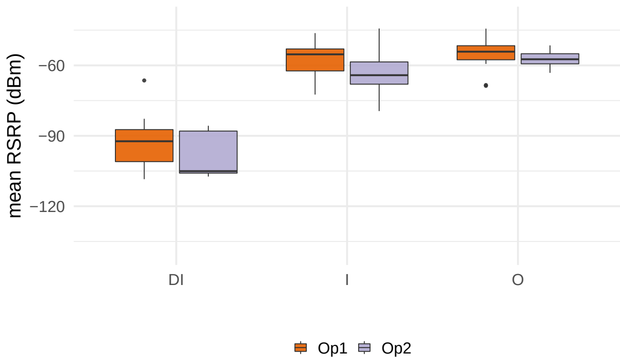

(b)

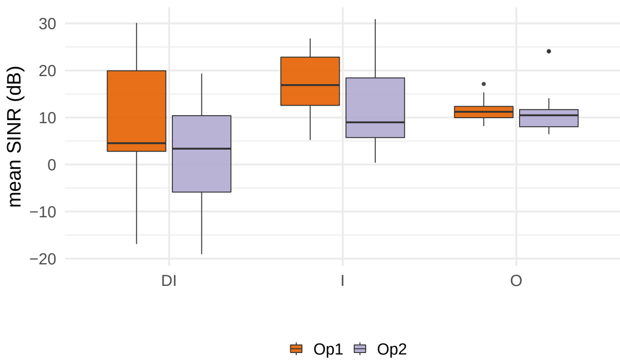

(c)

Fig. 2: Spatial deployment of NB-IoT eNBs for Op1 (blue color, label: 1) and Op2 (red color, label: 2) (a). NB-IoT coverage: sub-campaign average RSRP [dBm] (b) and SINR [dB] (c), grouped by scenario and operator.

between O and I is instead smaller, with an average increase of 1.55 dB (Op1) and 5.75 dB (Op2) for RSRP, and an average decrease of 6.18 dB (Op1) and 0.92 dB (Op2) for SINR. The SINR decrease highlights that the higher heterogeneity and dynamicity of the O scenario (e.g., in terms of mobility) has a direct impact on signal quality rather than signal strength. Across operators, we find that Op1 consistently provides better NB-IoT coverage (i.e., about 4.69 dB and 3.62 dB better on average for RSRP and SINR, respectively), mainly due to its denser infrastructure.

We further validated the above observations by performing two non parametric analysis of variance tests, i.e., Kruskal-Wallis and Dunn's tests, to assess which of the RSRP/SINR distributions have statistically different mean values. In particular, we observed mostly statistically significant difference

TABLE II: Cell (re-)selection parameters for Op1 and Op2. Acronyms and notations from Table I.

Selection Parameter (SIB1)	**Op1**	**Op2**
$Q_{\text{rx,min}}$ [dBm]	-64	-70
$Q_{\text{qual,min}}$ [dBm]	-23	-23
Re-selection Parameter (SIB3)	**Op1**	**Op2**
$S_{\text{IntraSearchP}}$ [dBm]	31	31
$Q_{\text{rx,min}}$ [dBm]	-64	-70
t_{resel} [s]	3	6
$S_{\text{NonIntraSearch}}$ [dBm]	31	10
Q_{hyst} [dB]	3	4

a) between DI and I/O scenarios (for both operators), and b) between operators in I and O scenarios. Due to space limitation, the full set of results is reported in [22].

2) Cell (re)-selection: We report in Table II the parameters adopted by Op1 and Op2 for cell (re-)selection, retrieved by decoding SIB1 and SIB3 messages via TSMA6. We observe that while some parameters set by the operators are similar, there are also significant differences in some other parameters. In particular, with regards to cell suitability (Eq. (1)), Op2 uses a lower value for $Q_{\text{rx,min}}$, meaning that more cells would be considered suitable for connection. In parallel, Op2 tries to slow down the re-selection rate by adopting higher values for t_{resel} and Q_{hyst} (Eq. (3)). We also observe that both operators share in SIB3 a value for $S_{\text{NonIntraSearch}}$, extremely high for Op2. However, we did not detect SIB5 messages, which carry other NonIntra configurations, implying that NonIntra re-selections are not executed. This is in line with the current deployment for both operators, which comprises of a single NB-IoT carrier each in the guard bands of LTE Band 20.

The adopted configurations directly relate to deployment/coverage aspects previously discussed: given its less dense deployment and lower coverage, Op2 enables more (re-)selection opportunities, which may help the devices to find a cell to camp on. However, this solution may lead to excessive cell re-selections, which may not be needed and beneficial for stationary devices, ultimately leading to higher energy consumption. We highlight this aspect by evaluating the statistics of the number of serving cells for DI and I stationary sub-campaigns, as well as for O scenarios. In particular, for each sub-campaign we evaluate the number of cells toward which the Exelonix module has performed RA that resulted in a connection. We find that Op1 has a smaller average value of serving cells than Op2 in both DI (2.44 vs. 3.59) and I (1.84 vs. 2.29), thus confirming that the Op2 deployment/coverage, along with the parameter settings shown in Table II, lead to more cell (re-)selections in stationary scenarios. Due to denser deployment, Op1 has instead more serving cells in O (45 vs. 28.83 on average), where the higher values as compared to DI and I are due to the larger measurement area. In this case, contrarily to the indoor counterpart, the (re-)selections are inherently needed due to device mobility.

3) CL estimation: We now discuss the CL estimation, as described in §II-B1b. As an initial step, we report in Table III the entire set of RA configurations for both operators, retrieved by decoding SIB2 messages. We focus for now

TABLE III: RA parameters for Op1 and Op2, as observed in SIB2. Acronyms and notations from Table I.

Parameter	**Op1**	**Op2**
Th_1, Th_2 [dB]	31, 21	36, 26
Δ_{ramp} [dB]	2	2
$P_{\mathsf{tg}}^{\mathsf{NPRACH}}$ [dBm]	-112	-104
N_{att}	10	10
Parameter (per CL)	**Op1**	**Op2**
$N_{\mathsf{att}}^{\mathsf{CL}_x}$ (a)	4, 4, 2	4, 4, 2
$N_{\mathsf{rep}}^{\mathsf{CL}_x}$ (a)	2, 8, 32	2, 8, 32
NPRACH periodicity(b) [ms]	640	640
NPRACH subcarriers(b)	12	12
Response Window Size(a,c)	pp10, pp5, pp5	pp5, pp5, pp5
MAC Cont. Resol. Timer(a,c)	pp64, pp64, pp64	pp8, pp8, pp8

(a) Values for CL_0, CL_1, and CL_2.

(b) The same value is configured for all CLs.

(c) ppx stands for "x times (N)PDCCH period (pp)".

on the first row of the table, which provides the values for Th_1 and Th_2 adopted in Eq. (4), and discuss the other parameters later. We observe that operators apply different RSRP thresholds, 5 dB more conservative for Op2, which is thus more likely to work at higher CLs. As for cell (re-)selection, we believe that this configuration is related to the operators' deployment/coverage, with Op2 having lower coverage, and thus trying to enhance the connection reliability by operating at higher CLs, ultimately dealing with increased energy consumption and congestion due to repetitions.

We analyze the combined effect of deployment, coverage, and RSRP thresholds in Table IV, which reports, across different scenarios, the ratio of initiating and concluding a RA in the same CL (first three columns per operator) or in different CLs (last three columns per operator). We evaluate the ratio as the number of RA executed in a specific CL or transition, divided by the total number of RA executions. We observe that a generic device operating in Op1 network would likely start and conclude its RA in CL_0, independently on the scenario it is deployed. While this is almost always true in I, transitions to higher CLs are observed in about 5% of occasions in the other scenarios, due to more challenging propagation conditions (DI) and/or mobility and dynamicity effects (O). Due to its different configurations, Op2 works predominantly in CL_1 in DI, where it also experiences transitions to CL_2 (4%). As regards I and O, and similarly to Op1, Op2 mostly starts the RA in CL_0, but requires more transitions to higher CLs (10% on average across scenarios).

As a common trend across operators, we also observe that outdoor scenarios increase the need for transitions to higher CLs, even though better RSRP coverage is observed compared to indoor situations. This result maps with the SINR decrease shown in Fig. 2c, and suggests that in more heterogeneous and dynamic scenarios, the current CL estimation and RA procedures may be sub-optimal, e.g., requiring several attempts before achieving a successful connection. This observation motivates a more thorough study of RA, in terms of its explainability and outcome prediction, which is the focus of the analysis in §V.

B. RA Configuration and Performance

In the following we analyze the RA performance in terms of success vs. fail outcomes, and delay in concluding the procedure. We first highlight how the performance changes across scenarios, and then discuss power-related aspects, by analyzing NPRACH and NPUSCH power control (§II-B2b).

1) RA configurations: Before digging into performance analysis, we finalize the discussion on the RA configurations reported in Table III. Beside RSRP thresholds, the operators adopt similar parameters. As an interesting aspect, we observe that the two operators use different number of attempts and repetitions across CLs, while NPRACH periodicity and subcarriers are the same. The latter settings differ from the assumptions mostly adopted in simulation-based studies (e.g., in [9]). We also highlight different values of $P_{\mathsf{tg}}^{\mathsf{NPRACH}}$ (8 dB higher for Op2), having significant implications on NPRACH power control (cf. §IV-B3). Op1 also allocates larger time windows waiting for *Msg2* (Response Window Size) in CL_0, and for *Msg4* (MAC Contention Resolution Timer).

2) Performance across scenarios: There are two main outcomes of the RA procedure: RA Success and RA Fail, as described in §II-B1 (see Fig. 1). However, a further outcome was registered in our measurements, denoted as RA Abort. This is specifically used by the Qualcomm chipset in the Exelonix module, to pinpoint device/cell misalignments during a RA execution. In particular, majority of aborts were verified when the subcarrier used by the device for transmitting a preamble attempt was wrongly decoded, ultimately causing the RA to stop abruptly. In the following analysis, we do not consider RA Abort, since it is a chipset-specific rather than a standardized RA outcome.

Focusing on RA successes, we dissect them into four sub-levels. By doing so, we aim at emphasizing that, from several operational perspectives including resource allocation and energy consumption, RA successes obtained after the transmission of one preamble attempt only (1^{st} *Attempt*) are inherently different from successes obtained after either performing more attempts with power ramping in CL_0 (*Power Ramping*), or adjusting the initial CL estimate of one or two CLs ($+1$ *CL*, i.e., either $\mathsf{CL}_0 \rightarrow \mathsf{CL}_1$ or $\mathsf{CL}_1 \rightarrow \mathsf{CL}_2$, depending on the initial CL estimate, and $+2$ *CLs*, i.e., $\mathsf{CL}_0 \rightarrow \mathsf{CL}_2$).

Hence, in Table V we evaluate, for each scenario and operator, the ratio of a given outcome by dividing the number of RA with that specific result by the total number of executions. For the successes, this ratio represents an estimate of the achievable success probability, which is a RA key performance indicator (cf. §VI). We observe that the successes are significantly predominant across all scenarios and for both operators, thus hinting at in general reasonably good behaviour of the procedure. Indeed, the success probability constantly is of at least 99% across scenarios, meeting the requirement of having at least 99% success probability over ten RA attempts (note $N_{\mathsf{att}} = 10$ for both operators) [31]. Op1 shows slightly better results compared to Op2; for both, the dynamicity of outdoor scenarios leads to an increase of RA failures.

Overall, we note that the reported results are not likely affected by massive device connectivity, since the measurement period covers initial system testing phases. Hence, they

TABLE IV: CL statistics for Op1 and Op2 across scenarios. For each operator, the first three columns represent the ratio of initiating and concluding a RA procedure in the same CL; the last three columns report the ratio of RA procedures concluded after CL increase.

Scenario	Op1						Op2					
	CL_0	CL_1	CL_2	$CL_0 \rightarrow CL_1$	$CL_1 \rightarrow CL_2$	$CL_0 \rightarrow CL_2$	CL_0	CL_1	CL_2	$CL_0 \rightarrow CL_1$	$CL_1 \rightarrow CL_2$	$CL_0 \rightarrow CL_2$
DI	0.93	0.02	≈ 0	0.03	0.02	≈ 0	0.26	0.69	≈ 0	0.01	0.03	≈ 0
I	0.99	0	0	≈ 0	0	≈ 0	0.91	0	0	0.08	0	≈ 0
O	0.95	0	0	0.025	0	0.025	0.885	≈ 0	0	0.085	0	0.03

TABLE V: RA result statistics for Op1 and Op2.

Scenario	Op1					Op2				
	RA Success				RA Fail	RA Success				RA Fail
	1st Attempt	Power Ramping	+1 CL	+2 CLs		1st Attempt	Power Ramping	+1 CL	+2 CLs	
DI	0.81	0.14	0.04	< 0.01	< 0.01	0.77	0.18	0.04	< 0.01	< 0.01
I	0.99	< 0.01	< 0.01	< 0.01	0	0.88	0.03	0.08	< 0.01	< 0.01
O	0.875	0.075	0.025	0.01	0.015	0.805	0.085	0.08	0.015	0.015

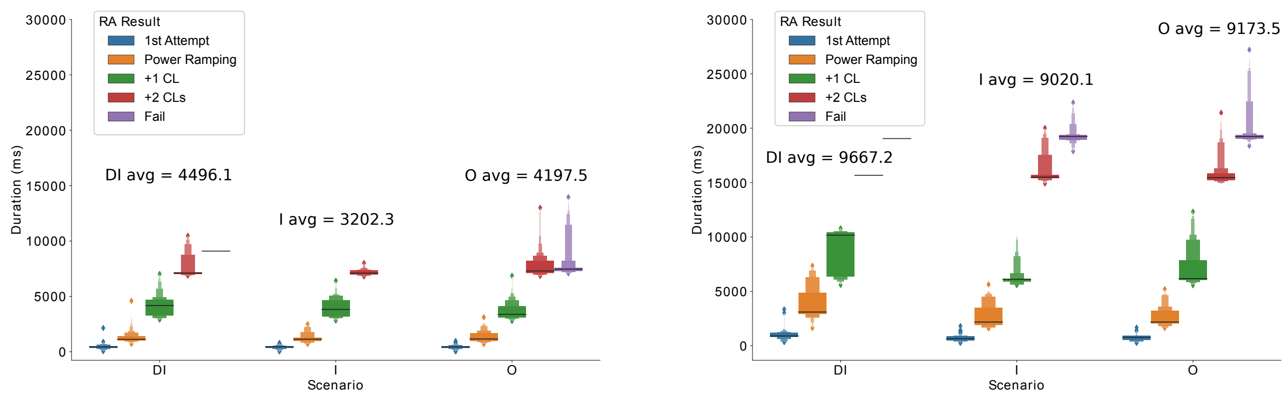

(a) (b)

Fig. 3: Statistics of RA duration for Op1 (a) and Op2 (b), across scenarios and RA results. For each scenario, the numerical indication represents the average value.

provide RA performance bounds to be compared with the performance achievable under massive system usage, for which we plan measurements and analyses for future work. As also highlighted in theoretical and simulation-based studies, we envision that meeting the 99% success probability requirement would be increasingly challenged by massive connectivity, thus calling for adjustments in the RA procedure.

Going deeper on the analysis of different types of successes, Op1 shows 80% of successes after one preamble transmission in DI, with this value increasing in O and peaking at a value of 99% in I. When the first attempt is not successful, the power ramping applicable in the next attempts in CL_0 (+2 dB for up to three more attempts, see Table III) is effective to solve the majority of remaining procedures (about 8% on average across scenarios). Finally, the need for increasing CLs and working at maximum power is experienced in 1-4% of cases. Considering Op2, the success ratio after one attempt is notably decreased, finding a minimum of 77% and a maximum of 88% in DI and I scenarios, respectively. The power ramping in successive attempts still leads to evident benefits, successfully solving 18% of RA procedures in DI, and about 10% across scenarios. When compared to Op1, Op2 also requires more CL increases to achieve connectivity, hinting at the need for working at maximum power, and using more attempts and repetitions. We find these results to be in line with the discussion in [18], where it is observed that a large number of repetitions (as for CL_1 and CL_2) is needed under bad radio conditions (e.g., SINR ≈ -30 dB), while consecutive (unrepeated) RA attempts (also, retransmissions) are sufficient for achieving RA successes with high SINR. In our measurements, we always find SINR > -30 dB, and see that RA is mostly solved with CL_0 attempts, thus requiring a low usage of many repetitions. However, it is worth mentioning that, compared to unrepeated and unramped retransmissions mostly considered in [18], both operators use 2 repetitions for each CL_0 attempt, and apply power ramping.

The results in Table V are complemented by Fig. 3, which reports the statistics of RA duration for Op1 (Fig. 3a) and Op2 (Fig. 3b), split across scenarios and RA results. This is a key performance indicator, representing an estimate of the achievable access delay, often used in RA studies. For both operators, we observe a duration increase when moving from RA successes with one attempt to RA failures, due to the increase of preamble attempts and repetitions. The duration peaks at the maximum value when RA fails, that happens by definition only after the maximum amount of attempts (10 for both operators, see Table III). We also observe that, given a specific result, the RA duration slightly decreases moving from I to O, hinting at a slight effect of propagation conditions. Finally, we highlight that even though presenting similar patterns, operators significantly differ in terms of

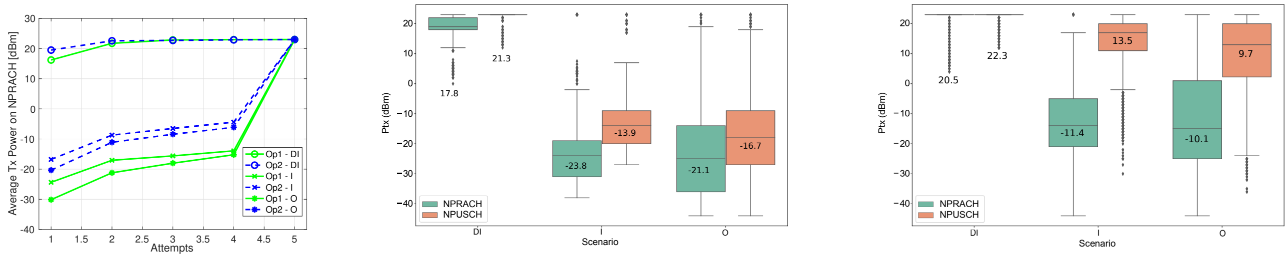

(a) (b) (c)

Fig. 4: NB-IoT power control mechanisms. The effect of NPRACH power control and ramping for both operators and across scenarios is reported in (a). If needed, both operators switch from CL_0 to CL_1 during the 5th preamble attempt (see Table III), where they also adopt $P_{\mathsf{tx}}^{\mathsf{NPRACH}} = P_{\mathsf{tx}}^{\mathsf{max}} = 23$ dBm; the comparison between NPRACH and NPUSCH power control mechanisms is given in (b)(c) for Op1 and Op2, respectively. For each boxplot in (b)(c), the numerical indication represents the average value.

absolute values. This result ties back again with the differences in deployment, coverage, and configurations. Overall, Op2 provides slower access, with a RA duration of above 15 seconds in case of both RA Success after $+2$ CLs increase and Fail, while Op1 is under 10 seconds in most of cases. Such values can be benchmarked by the requirement stated in [35], i.e., a RA should take no longer than 10 seconds to be completed. We see that Op1 mostly meets the constraint, while Op2 has issues when a RA with double CL increase is needed for achieving a success. As also commented for the success probability, massive connectivity will increasingly challenge the achievement of this delay constraint, ultimately requiring RA enhancements. Indeed, the proposals in this and other works (§V-§VI) aim at reducing attempts and repetitions toward faster successful accesses.

3) Power aspects: We now analyze power-related aspects, considering their importance for NB-IoT energy efficiency. As an initial step, we report in Table VI the configurations adopted by the operators for NPRACH and NPUSCH power control.

TABLE VI: NPRACH and NPUSCH power control parameters for Op1 and Op2, as observed in SIB2. Acronyms and notations from Table I and Eq. (7).

Parameter	Op1	Op2
$P_{\mathsf{tg}}^{\mathsf{NPRACH}}$ [dBm]	see TABLE I	
$P_{\mathsf{tg}}^{\mathsf{NPUSCH}}$ [dBm]	-105	-67
α	1	0.7
$\Delta_{\mathsf{NPRACH}}^{(a)}$ [dB]	4	4

$^{(a)}$Included in the term c in Eq. (7).

We already observed in §IV-B1 that Op2 adopts a higher value for $P_{\mathsf{tg}}^{\mathsf{NPRACH}}$, and thus likely uses more power to transmit preamble attempts. We now observe that Op2 also uses a higher value for $P_{\mathsf{tg}}^{\mathsf{NPUSCH}}$, which likely leads to higher power levels to transmit *Msg3* and UL data, even though the adopted α is slightly lower compared to Op1.

These observations are confirmed by the results reported in Fig. 4. Fig. 4a shows the effect of NPRACH power control and ramping during the four preamble attempts in CL_0, in terms of average transmitted power across different scenarios. We observe that both operators mostly work at $P_{\mathsf{tx}}^{\mathsf{max}}$ in DI, with Op1 using slightly less power than Op2. In other scenarios, Op1 clearly exploits its better coverage and less stringent $P_{\mathsf{tg}}^{\mathsf{NPRACH}}$ value, resulting in lower power values compared to Op2. Figs. 4b-4c report the statistics of the transmitted power on both NPRACH and NPUSCH, per operator and across scenarios. We observe that the higher requirements of Op2 on $P_{\mathsf{tg}}^{\mathsf{NPRACH}}$ and $P_{\mathsf{tg}}^{\mathsf{NPUSCH}}$ significantly affect the performance, and lead to higher powers on both channels. In particular, while the difference between NPRACH and NPUSCH powers is limited to 6 dB on average for Op1, it increases to around 17 dB for Op2, due to the significant difference between the target power values.

These results allow to discuss a further aspect toward the deployment and usage of the new device power class standardized in Rel-14, which operates with $P_{\mathsf{tx}}^{\mathsf{max}} = 14$ dBm. The introduction of these devices leads to further heterogeneity and may require the operators to a) rethink their deployment strategies and configurations, and b) communicate to the end-users that different devices may deliver significantly different performance in terms of connectivity and battery lifetime. Given current deployments and configurations, Figs. 4b-4c highlight that the new device power class is not suitable in DI for both operators, as they constantly work above 14 dBm. As regards the other scenarios, we notice that the new power class would likely work in the Op1 network, since NPRACH and NPUSCH powers are below 14 dBm in most of cases; on the contrary, deployment and configuration choices significantly hinder the use of the new power class to Op2 customers; NPUSCH median powers are in fact in the range of 14 dBm, meaning that the new power class is not usable in about 50% of cases, ultimately requiring Op2 to reconfigure its system, e.g., reducing $P_{\mathsf{tg}}^{\mathsf{NPUSCH}}$ value.

V. RA Explainability and Improvement

In this section, we deepen our analysis by employing a ML approach on the collected data. The goal is to reveal the interdependencies between radio conditions, configurations, and RA outcomes, aiming at better understanding the RA procedure, proposing better configurations, and achieving higher efficiency.

In order to do so, we map each RA outcome to a set of *features*, to understand how RA outcomes and features are correlated to each other, and to which extent the latter can be exploited to represent and predict the former. In our analysis, we select as features a set of parameters characterizing both radio conditions and adopted configurations, measured *at the beginning* of each RA execution. The feature set thus comprises of RSRP, RSRQ, and SINR values, as well as initial CL estimate and $P_{\mathsf{tx}}^{\mathsf{NPRACH}}$ adopted for the first preamble attempt.

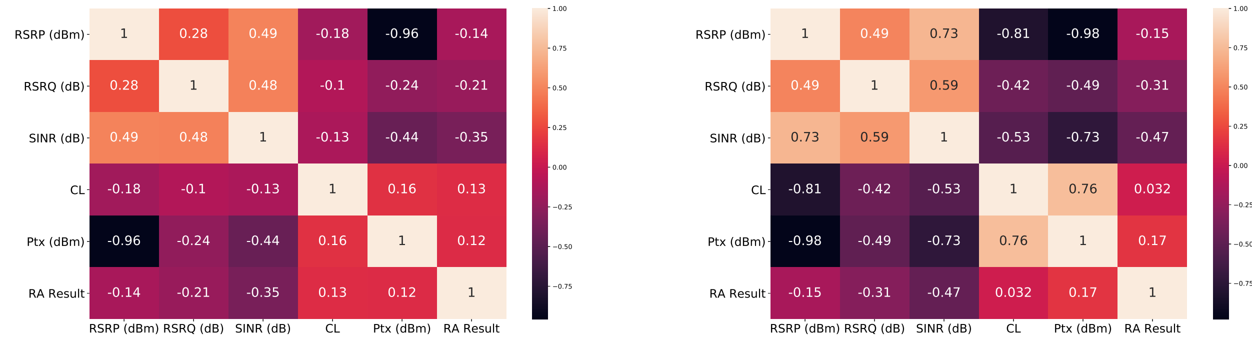

(a) (b)

Fig. 5: Correlation Matrix across features and RA Result for Op1 (a) and Op2 (b). Data from all scenarios (DI, I, and O) are combined together in order to provide an overall overview.

In the following, the initial CL estimate is simply referred to as CL, while $P_{\text{tx}}^{\text{NPRACH}}$ is simplified as P_{tx}.

A. Linear Correlation and Feature Selection

As an initial step, we run a linear correlation and feature selection analysis, showing the results in Fig. 5 and Table VII, respectively.[4] Figure 5 highlights that the three radio parameters (RSRP, RSRQ, and SINR) present the highest correlation coefficients with the RA outcome, across all features and for both operators, thus implying their impact on the RA final result. With respect to CL, we observe low correlation with the radio parameters for Op1, while high values are obtained for Op2, and particularly with RSRP. This can be explained by reminding that the CL feature represents the initial CL estimate, which is derived from RSRP through Eq. (4). However, due to coverage and threshold configurations, the estimate is almost always equal to CL_0 for Op1, ultimately hiding the relationship with RSRP, which is instead clear for Op2. High correlation with RSRP is always observed for P_{tx} which is related to RSRP via Eq. (7).

We then apply two well-known feature selection methods on the feature set, that is, Recursive Feature Elimination (RFE) [36] and Lasso regularization (Lasso) [37], both working while training a linear model between features and RA results.

TABLE VII: Application of RFE and Lasso feature selection algorithms. The initial feature set is $\{$RSRP, RSRQ, SINR, CL, $P_{\text{tx}}\}$.

Scenario	Op1		Op2	
	RFE	Lasso	RFE	Lasso
DI	RSRP	RSRP, RSRQ, SINR	RSRP, SINR	SINR
I	RSRP, SINR	SINR	RSRP, SINR	SINR
O	RSRQ, SINR	SINR	RSRQ, SINR	RSRQ, SINR

Table VII reports the results of the application of RFE and Lasso to our feature set, split by operator and scenario, while training a linear model for the RA Result variable. We observe that both methods highlight the importance of SINR as a feature for predicting the RA outcome. Lasso is in general more selective than RFE, and often leads to selecting SINR as sufficient predictor. We also observe that, in stationary scenarios, RSRP is selected more often than RSRQ as predictor along with SINR, while the opposite happens under mobility. The result implies that, in more dynamic situations, the use of parameters of signal quality rather than signal strength may lead to performance improvement of the RA procedure, as they show stronger relationship with its outcomes. As regards CL and P_{tx}, they are never selected by RFE or Lasso, being both strongly correlated with the radio parameters. From this analysis we can finally conclude that the radio parameters have a direct impact on RA results, while CL and P_{tx} potentially bring limited RA explainability gains.

B. RA Outcome Classification

We then move a step further, casting our RA explainability problem as a classification problem, also considering the discrete nature of the RA Result variable. Our analysis aims at understanding if it would be possible to correctly classify RA outcomes based on the feature observation. From an operational perspective, this would be significantly beneficial, since it provides a methodology for improving the RA efficiency.

To better clarify this aspect, we define in Table VIII an action set that could exploit the RA outcome prediction and improve the procedure. As shown in the table, the prediction of a success after one preamble attempt suggests to rely on the standard procedure, as this already provides the optimal outcome. However, assuming that the prediction is a success after either power ramping or increase of $+1$ or $+2$ CLs, then an efficient action would be to send the initial attempt with a higher power with respect to the one derived from Eq. (5) (e.g., already adding the needed Δ_{ramp}), or directly exploit the configurations of higher CLs. This would in turn avoid the transmission of attempts that are predicted to be unsuccessful.

Similar observations can be done for RA Fail prediction: in this case, it would probably make sense to try the connection only using the most reliable configuration possible (i.e., CL_2 settings), even though this outcome hints that even more drastic adjustments may be needed, e.g., enable the use of more repetitions beyond the limits configured for CL_2.

Holding these observations, we initially run a comparison across possible classifiers, in order to highlight possible gains in using one with respect to another. Recognizing that our

[4]For this analysis, we represent the RA outcomes in a numerical form, with the four types of successes ordered as in Table V and numbered from 1 to 4, while RA Fail is indicated with 5.

TABLE VIII: Possible action set derived from the adoption of the RA outcome prediction scheme, and corresponding effect on the standardized procedure. Reported parameters are defined in Table I; $P_{\text{tx}}^{\text{NPRACH}}$ evaluation follows Eq. (5).

Predicted RA Result	Action	Effect
1st Attempt	**Follow standard procedure**	–
Power Ramping at i^{th} attempt $(i \in [1, N_{\text{att}}^{\text{CL}_0} - 1])$	**Transmit initial attempt with these configurations:** $\begin{cases} N_{\text{rep}} = N_{\text{rep}}^{\text{CL}_0} \\ P_{\text{tx,opt}}^{\text{NPRACH}} = \min\{P_{\text{tx}}^{\text{NPRACH}} + i * \Delta_{\text{ramp}}, P_{\text{tx}}^{\text{max}}\} \end{cases}$	**Avoid Tx of previous i attempts:** $i = 1 \rightarrow$ no Tx of 1st unramped attempt; $i = 2 \rightarrow$ no Tx of 1st unramped and 1st ramped attempt; $i = 3 \rightarrow$ no Tx of 1st unramped and 1st/2nd ramped attempts
+1 CL	**Transmit initial attempt with CL_{x+1} configurations:** $\begin{cases} N_{\text{rep}}^{\text{CL}_{x+1}} \\ P_{\text{tx,opt}}^{\text{NPRACH}} = P_{\text{tx}}^{\text{max}} \end{cases}$	**Avoid Tx of previous $N_{\text{att}}^{\text{CL}_x}$ attempts in CL_x**
+2 CLs	**Transmit initial attempt with CL_2 configurations:** $\begin{cases} N_{\text{rep}}^{\text{CL}_2} \\ P_{\text{tx,opt}}^{\text{NPRACH}} = P_{\text{tx}}^{\text{max}} \end{cases}$	**Avoid Tx of previous $N_{\text{att}}^{\text{CL}_0} + N_{\text{att}}^{\text{CL}_1}$ attempts in CL_0 and CL_1**
Fail	**Follow "+2 CLs"**	**As for "+2 CLs"**

dataset is significantly imbalanced from the RA outcome perspective, we balance it across operators and scenarios adopting the well-known ADASYN method [38] [39]. We then train four classifiers, that is, multi-class Support Vector Machine (SVM) with linear kernel, k-Nearest Neighbors (kNN) with $k = 1$, Decision Tree (DT), and Random Forest (RF). RF is a well-known extension of DT, being an ensemble classifier allowing parallel construction of multiple trees, by resampling with replacement the training data (i.e., *bagging*) [40].

For our analysis, we perform the training/classification 10 times, each time extracting 500 random samples per class from the balanced dataset. Moreover, we adopt RSRP, RSRQ, and SINR as features, due to their importance highlighted in the previous analysis. Finally, we use 50 trees for RF. As visually reported in [22], this is a good yet conservative choice for all the scenarios, since we observe the RF accuracy to converge with about 10-20 trees in DI and I scenarios, while it keeps slightly increasing with more than 20 trees in the O scenario.

We show the obtained average 10-folds cross-validation accuracy for each classifier in Fig. 6, across scenarios and operators. We preliminary observe that the classification performs slightly better on Op2 dataset, except in I scenario, where the task is inherently simpler for Op1, since in this case one class is missing (RA Fail). Moreover, the increased dynamicity in the O scenario challenges the classification, resulting in lower accuracy. Comparing the classifiers, linear SVM shows the worst accuracy, hinting that the faced problem may have some non-linearities to be taken into account. kNN and DT have instead similar performance, significantly better compared to SVM. As regards RF, it slightly improves DT (between 0-5%) in DI and I scenarios, suggesting that DT, which is computationally faster and simpler than RF, can be reliably used in these scenarios. However, the improvement with respect to DT peaks at 9% in O scenario, hinting that the challenges posed by the dynamicity of this scenario can be better handled by more advanced classifiers. Considering the results, we pick RF and perform a deeper analysis aiming at better understanding the impact of using different sets of features on the achievable accuracy.

Figure 7 reports the Out-Of-Bag (OOB) classification accuracy per scenario, operator, and adopted feature set, adopting

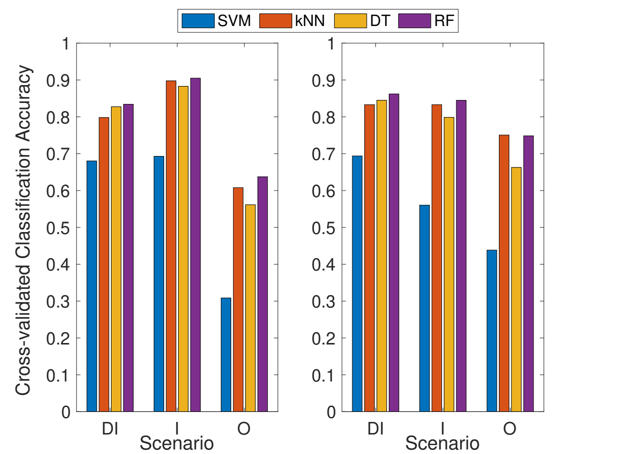

Fig. 6: 10-folds cross-validation accuracy for SVM, kNN ($k = 1$), DT, and RF (50 trees) classifiers across scenarios and operators (Op1: left, Op2: right). The feature set is {RSRP, RSRQ, SINR}.

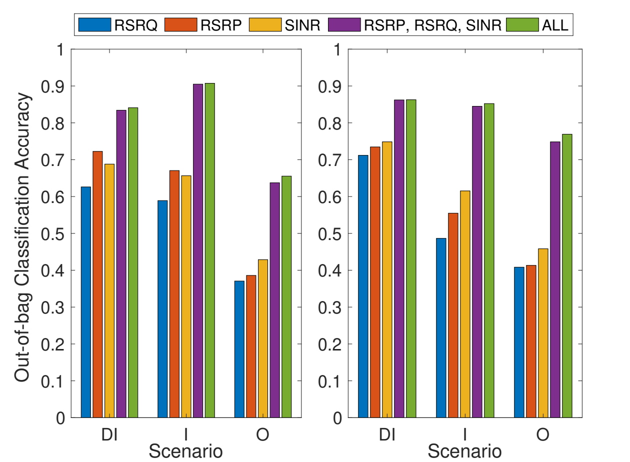

Fig. 7: Out-of-bag classification accuracy in case of DI, I, and O scenarios (Op1: left, Op2: right). RF uses 50 trees.

a RF classifier with 50 trees. In this case, we run the classification with 500 random samples per class, and average over 20 executions. The OOB accuracy is calculated as the average prediction accuracy on the training samples, where each sample is predicted by using the trees trained with a data bootstrap not including that specific sample. It complements

TABLE IX: Power and repetitions saved by adopting the RA outcome prediction scheme and the action set in Table VIII. RA parameter values are assumed as in Table III; $P_{\mathsf{tx}}^{\mathsf{NPRACH}}$ evaluation follows Eq. (5).

RA Result	**Saved power**		**Saved repetitions**
	Case 1: $P_{\mathsf{tx}}^{\mathsf{NPRACH}} + i * 2 < P_{\mathsf{tx}}^{\mathsf{max}} \quad \forall i \in [1,3]$ [(a)]	**Case 2:** $P_{\mathsf{tx}}^{\mathsf{NPRACH}} \geq P_{\mathsf{tx}}^{\mathsf{max}}$	
1st Attempt	–	–	–
Power Ramping at i^{th} attempt [(b)]	$2 * [i * P_{\mathsf{tx}}^{\mathsf{NPRACH}} + i * (i-1)]$	$2 * i * P_{\mathsf{tx}}^{\mathsf{max}}$	$2 * i$
+1 CL	$\mathsf{CL}_0 \rightarrow \mathsf{CL}_1$: $8 * P_{\mathsf{tx}}^{\mathsf{NPRACH}} + 24$ $\mathsf{CL}_1 \rightarrow \mathsf{CL}_2$: $32 * P_{\mathsf{tx}}^{\mathsf{max}}$	$\mathsf{CL}_0 \rightarrow \mathsf{CL}_1$: $8 * P_{\mathsf{tx}}^{\mathsf{max}}$ $\mathsf{CL}_1 \rightarrow \mathsf{CL}_2$: $32 * P_{\mathsf{tx}}^{\mathsf{max}}$	$\mathsf{CL}_0 \rightarrow \mathsf{CL}_1$: 8 $\mathsf{CL}_1 \rightarrow \mathsf{CL}_2$: 32
+2 CLs	$8 * P_{\mathsf{tx}}^{\mathsf{NPRACH}} + 24 + 32 * P_{\mathsf{tx}}^{\mathsf{max}}$	$40 * P_{\mathsf{tx}}^{\mathsf{max}}$	40
Fail	As for "+2 CLs"		

(a) As in Table VIII, $i \in [1, N_{\mathsf{att}}^{\mathsf{CL}_0} - 1]$ and does not include the first unramped RA attempt. From Table III, $N_{\mathsf{att}}^{\mathsf{CL}_0} = 4$ for both operators.
(b) E.g., predicting a RA success at $i = 1$ allows to avoid the transmission of the 1st unramped attempt (repeated $N_{\mathsf{rep}}^{\mathsf{CL}_0} = 2$ times).

to one the OOB error, which is a well-known performance indicator for bagging-based algorithms.

We clearly observe that the accuracy depends on the adopted feature set. In particular, the combination of the three radio parameters significantly improves the performance, compared to the cases in which such parameters are used standalone. On this aspect, we notice that a) SINR standalone provides the best classification results in most of cases compared to RSRP and RSRQ, and b) the use of the entire set of initial features (denoted as ALL in Fig. 7), comprising of CL and P_{tx}, does not bring significant gains compared to the {RSRP, RSRQ, SINR} set, confirming the results of the feature selection analysis. Across operators, we also observe that, independently from the adopted feature set, the classification performs better on the Op2 dataset, apart for the I scenario, due to the same reason discussed for Fig. 6.

For DI and I scenarios, and {RSRP, RSRQ, and SINR} set, the classification accuracy lies between 80-90%, indicating that RA outcomes can be reliably predicted, and in turn used to enhance the standardized procedure. We complement these results in [22], where we report and discuss the confusion matrix for each operator/scenario configuration, showing which RA outcomes are more challenging to predict correctly.

We now assume to embed the outcome prediction in the RA procedure, and analyze the possible advantage of doing so. We propose a centralized implementation, thus considering the operators being able to collect a good amount of executions and outcomes for the standard Rel-13 RA (e.g., via campaigns with dedicated testing devices or by monitoring the devices registered to the network). This allows to train the RF classifier and derive the decision trees for properly selecting one of the optimized RA actions in Table VIII (e.g., *"If RSRP, SINR, and RSRQ are below these thresholds* (derived by the classifier)*, start RA with* $\mathsf{CL}_{\mathsf{x}+1}$ *rather than* CL_{x} *configurations"*, where CL_{x} is derived via Eq. (4)). Then, we propose a slight modification of SIB signaling (e.g., an enhanced SIB2), which would allow each cell to share the mapping between radio conditions and optimized RA actions, ultimately enabling the device to pick the best configuration considering its current status. Backward compatibility can be achieved via a flag indicating the usage of either standard or enhanced procedure. SIB messages can be also easily updated, in case a remapping between conditions and suggested actions is needed (e.g., due to availability of new measurements for classifier retraining).

We report in Table IX the amount of saved power and repetitions when we take a correct outcome prediction and perform the corresponding action proposed in Table VIII. In order to evaluate the power saving, we distinguish upper and lower bound cases, referred to as Case 1 and Case 2. In both cases, the power for transmitting the first attempt is evaluated via Eq. (5); then, in Case 1, such power is low enough so that the possibly needed successive attempts with power ramping are all transmitted with a power lower than $P_{\mathsf{tx}}^{\mathsf{max}}$. On the contrary, in Case 2, Eq. (5) returns a power equal to $P_{\mathsf{tx}}^{\mathsf{max}}$, and hence all successive attempts are also transmitted at $P_{\mathsf{tx}}^{\mathsf{max}}$.

Observing Table IX, we describe in particular the savings achievable when a success after +2 CLs is correctly predicted, since the other cases can be straightforwardly derived in a similar way. In this situation the optimized RA procedure directly starts with CL_2 configurations (Table VIII), thus avoiding $N_{\mathsf{att}}^{\mathsf{CL}_0} = 4$ attempts in CL_0 (each repeated $N_{\mathsf{rep}}^{\mathsf{CL}_0} = 2$ times) and $N_{\mathsf{att}}^{\mathsf{CL}_1} = 4$ attempts in CL_1 (each repeated $N_{\mathsf{rep}}^{\mathsf{CL}_1} = 8$ times). Assuming that the avoided attempts in CL_0 are in either Case 1 or Case 2, it follows that the saved power for not transmitting them (and corresponding repetitions) is either $8 * P_{\mathsf{tx}}^{\mathsf{NPRACH}} + 24$ or $8 * P_{\mathsf{tx}}^{\mathsf{max}}$. Specifically for Case 1, the constant value of 24 dB represents the effect of ramping the power for transmitting the consecutive attempts. Considering $\Delta_{\mathsf{ramp}} = 2$ dB, then the two repetitions of the first attempt are both ramped of 2 dB with respect to the very initial (unramped) preamble, the ones of the second attempt are ramped of 4 dB, and finally the ones of the third attempt are performed with a power ramping of 6 dB each.

Moreover, the saved power for not transmitting the attempts in CL_1 (and corresponding repetitions) is always equal to $32 * P_{\mathsf{tx}}^{\mathsf{max}}$, since they are always transmitted at maximum power. Hence we obtain a saved power equal to $8 * P_{\mathsf{tx}}^{\mathsf{NPRACH}} + 24 + 32 * P_{\mathsf{tx}}^{\mathsf{max}}$ (Case 1) or $40 * P_{\mathsf{tx}}^{\mathsf{max}}$ (Case 2).

In order to evaluate the relative power savings with respect to the standardized procedure, we assume that the RA achieves a success in the first attempt in CL_2, transmitted $N_{\mathsf{rep}}^{\mathsf{CL}_2} = 32$ times at full power. Hence, the standard procedure would transmit a total of 9 attempts, with a global used power of $8 * P_{\mathsf{tx}}^{\mathsf{NPRACH}} + 24 + 64 * P_{\mathsf{tx}}^{\mathsf{max}}$ (Case 1), while the optimized procedure would adopt a unique attempt transmitted with a power of $32 * P_{\mathsf{tx}}^{\mathsf{max}}$. Assuming that the power consumed for CL_0 (and power ramping) attempts is negligible with respect to the power needed for the CL_1 attempts, i.e., $8*(P_{\mathsf{tx}}^{\mathsf{NPRACH}})+$

$24 << 32{*}P_{\mathsf{tx}}^{\mathsf{max}}$, then we conclude that, compared to standard RA, the optimized procedure more than halves the power consumption in Case 1, with a maximum reduction of 55% in Case 2. In terms of repetitions, in both Cases 1 and 2, 32 repetitions are used instead of 72, with a reduction of 55%.

We finally conclude our analysis by highlighting again the accuracy decrease occurring in the O scenario, evidently due to higher dynamicity and in turn unpredictability. The result suggests that other variables, along with more advanced classifiers, should be exploited in this case, e.g., mobility-related information. The analysis of such variables opens the way for future work toward better understanding how to move forward cellular-based mobile IoT deployments and services.

VI. Related Work

In this section, we report and analyze relevant literature in the context of NB-IoT RA, and also emphasize differences and contribution of the present work.

Besides [7], which analyzes NPRACH design rationale, several works have performed analyses of NB-IoT RA. The work in [8] proposes a RA model taking roots from [41] [42], which investigated slotted ALOHA in OFMDA systems. The model considers stationary devices performing RA in three CLs, and estimates success probability and average access delay in terms of number of contending, successful, and collided devices. Simulations show that CL_0 devices have higher success probability than those in CL_1 and CL_2, with these latter also experiencing higher delays. The model is then extended in [9] and [10]. Such works also discuss possible schemes for optimizing RA configurations, aiming at maximizing the success probability under a maximum admissible access delay [9], or assuring fair access to users belonging to different CLs [10]. The parameters being optimized are the number of CLs, N_{att}, and, for each CL, $N_{\mathsf{att}}^{\mathsf{CL_x}}$, backoff windows, and number of subcarriers. The work in [9] confirms that the adoption of three CLs is a reasonable choice. Then, it is observed that if NPRACH periodicity increases with the CLs, optimal RA configurations have a) backoff windows also increasing with the CLs, and b) $N_{\mathsf{att}}^{\mathsf{CL_x}} = N_{\mathsf{att}}$ for each CL. Indeed, it is demonstrated that configuring $N_{\mathsf{att}}^{\mathsf{CL_x}} < N_{\mathsf{att}}$ does not lead to higher success probability. A rule for deciding the CL at which 24 subcarriers should be assigned (out of 48 available, as seen in §II-B2a, and with the other two CLs having 12 subcarriers each) is also given. The same optimization scheme is then used in [10], where a fairness mechanism is added by computing the Jain's index [43] across users in different CLs.

Such optimized settings are interesting but differ from the configurations currently in use in operational networks. E.g., Table III shows that both operators use the same NPRACH periodicity and number of subcarriers across CLs, and adopt $N_{\mathsf{att}}^{\mathsf{CL_x}} < N_{\mathsf{att}}$ for each CL. We also notice that models and strategies in [8]–[10] are purely probabilistic and do not consider practical factors, such as radio propagation, CL estimation, and power control. A similar simulation-based analysis is proposed in [11], focusing on the backoff mechanism. This is modeled as a Markov chain, while a capacity-limited First-In-First-Out (FIFO) queue is used to represent the device data buffer. The work in [12] also introduces a RA model and open-sources the simulator implemented for testing, which is based on LTE-Sim [44] and handles different CLs, attempt counters, and backoff windows.

The RA model in [13] considers the effects of radio propagation and power control (without ramping) on success probability, taking into account the SINR of preambles. The path loss model includes a distance-based loss and identically distributed Rayleigh fading, following [45]. A single CL is however assumed to simplify the analysis. Same assumptions are adopted in [14], which also targets simulation-driven RA performance evaluation. Results show that the success probability decreases as the interference due to contending devices increases. The effect of repetitions on RA reliability is also highlighted. On the one hand, it is observed that repetitions increase the success probability; on the other hand, it is shown that the standard number of repetitions may be insufficient in heavy traffic scenarios, failing to provide the 99% success probability requirement. The works in [15] and [16] extend [13] by considering three CLs. In particular, [16] compares backoff and access barring schemes, showing improvements through their joint use.

As regards RA enhancements, the work in [17] proposes to reduce collisions by slightly sacrificing preamble detection probability, adopting so-called partial preamble transmission. As anticipated in §IV-B, [18] discusses the trade-off between RA repetitions and attempts, assuming a system with one CL. It is shown that repetitions positively affect success probability, access delay, and power consumption in low SINR situations, while retransmissions are sufficient in good conditions.

A RA improvement named TARA is proposed in [19]. TARA is based on Time Alignment (TA), i.e., the estimated time delay between a device and a cell, needed to keep these two synchronized. TARA allows each device to perform two RAs in two consecutive RAOs. A TA-based RA can be executed if the first standard procedure fails. The cell uses a modified *Msg2* in order to broadcast a) the list of TA values estimated for the collided preambles, and b) information on the UL resources allocated to these values. Then, the device tries to match its TA value with one of the candidates. If a match is found, the device continues its RA with a *Msg3* transmitted on the allocated resources. Simulations show that TARA leads to higher success probability and throughput, and lower access delay compared to the standard RA.

A power efficient RA scheme is proposed in [20]. In this case, standard RA is modified so that the device failing a RA attempt in $\mathsf{CL_x}$ is not forced to try $N_{\mathsf{att}}^{\mathsf{CL_x}} - 1$ times more before moving to $\mathsf{CL_{x+1}}$. It can instead re-evaluate RSRP and probabilistically *jump* into another CL. Simulations show higher power efficiency compared to standard RA. This scheme may be considered as a *reactive* version of the one proposed in this paper. Indeed, it includes at least one failed attempt in $\mathsf{CL_x}$ (estimated via Eq. (4)). Our approach is instead *proactive*, since it allows to directly use an optimized configuration, considering previously acquired knowledge and corresponding RA outcome predictions. Moreover, power ramping is neglected in [20], while we keep it in our scheme, since we show it is key for achieving RA successes without increasing repetitions.

Aiming at more general NB-IoT performance analyses, RA is also evaluated in [46] and [47]. On the one hand, [46] models the resource allocation across DL/UL channels, showing the impact of repetitions on device battery lifetime and delay. On the other, [47] focuses on UL performance, including a Markov chain for the backoff mechanism. The path loss model in [48], derived from measurements of 95 macrocells in the US, is used to simulate the radio propagation.

With respect to the above works that are based on theoretical and simulation analyses, this paper provides an empirical analysis on how radio conditions and operators' configurations affect RA procedure and outcomes. The analysis thus enhances and complements the models, which have a stronger emphasis on the impact of multiple devices, and can also be used to configure the simulations with realistic parameters, ultimately verifying the accuracy with respect to real scenarios.

Recent literature is increasingly proposing the use of ML for optimizing NB-IoT systems. Among others, [49] proposes a Multi-Armed Bandit (MAB) framework for enabling distributed dynamic spectrum access across devices, aiming at reducing repetitions and energy consumption while increasing the coverage. A Reinforcement Learning (RL) solution in the form of Q-Learning is instead proposed in [50] for optimizing CL selection and number of repetitions for each device.

A RL-based enhancement is also proposed in [51] and extended in [52] and [53], aiming at optimizing several RA parameters, including the number of preamble attempts and repetitions. The cell performs the optimization via Q-Learning, using prior history in terms of preamble collisions, successful devices, and unused resources. The work in [53] shows that adding a preliminary traffic prediction step (e.g., for predicting the number of preambles sent in the next frame) via a Recurrent Neural Network (RNN) can lead to further access improvements. Clearly, our work is also a ML-based RA enhancement, but exploits the application of offline supervised ML on empirical data rather than online RL. The two enhancements seem to nicely complement each other, as we propose a method for better configuring and adjusting the RA procedure with respect to radio conditions and environmental scenarios, while [51]–[53] propose a scheme for better coping with the negative impact of massive connectivity and collisions.

VII. Conclusion

In this paper we present the first data-driven analysis of NB-IoT RA. By leveraging a large scale measurement campaign, we show the impact of network deployment, radio coverage, and operators' configurations on RA operations and outcomes. In terms of success probability and access delay, we observe that general requirements on such performance indicators are met, but the increasing massive connectivity and scenario heterogeneity will likely challenge these achievements, requiring procedure optimization. In this direction, we propose a ML-based scheme, preliminarily casting the RA explainability problem as a classification task. By doing so, we show that RA outcomes can be predicted with high accuracy by observing radio conditions, i.e., RSRP, SINR, and RSRQ values. The outcome predictor can be then used for driving the devices toward optimized RA configurations, enabling at least 50% power consumption reduction. We also discuss how such scheme could be embedded in a lightweight manner in next NB-IoT releases. For future work, we plan to replicate our measurements in advanced system usage states, in order to quantify the impact of massive connectivity, and test the proposed prediction methodology in more dynamic scenarios, ultimately aiming at possible enhancements.

VIII. Acknowledgment

This work is funded by the EU H2020 research and innovation programme under grant agreement No. 815178 (5GENESIS). The authors wish to thank Leonardo Piciucchi for contributing to the measurement campaign.

References

[1] U. Raza, P. Kulkarni, and M. Sooriyabandara, "Low Power Wide Area Networks: An Overview," *IEEE Commun. Surveys Tuts.*, vol. 19, no. 2, pp. 855–873, 2017.

[2] O. Liberg *et al.*, *Cellular Internet of Things: From Massive Deployments to Critical 5G Applications, 2nd Edition*. Academic Press, 2020.

[3] Y.-P. E. Wang *et al.*, "A Primer on 3GPP Narrowband Internet of Things," *IEEE Commun. Mag.*, vol. 55, no. 3, pp. 117–123, 2017.

[4] W. Yang *et al.*, "Narrowband Wireless Access for Low-Power Massive Internet of Things: A Bandwidth Perspective," *IEEE Wireless Commun.*, vol. 24, no. 3, pp. 138–145, 2017.

[5] M. Chen *et al.*, "Narrow Band Internet of Things," *IEEE Access*, vol. 5, pp. 20 557–20 577, 2017.

[6] J. Xu *et al.*, "Narrowband Internet of Things: Evolutions, Technologies, and Open Issues," *IEEE Internet Things J.*, vol. 5, no. 3, pp. 1449–1462, 2017.

[7] X. Lin, A. Adhikary, and Y.-P. E. Wang, "Random Access Preamble Design and Detection for 3GPP Narrowband IoT Systems," *IEEE Wireless Commun. Lett.*, vol. 5, no. 6, pp. 640–643, 2016.

[8] R. Harwahyu, R.-G. Cheng, and C.-H. Wei, "Investigating the Performance of the Random Access Channel in NB-IoT," in *IEEE Vehicular Technology Conf. (VTC-Fall)*. IEEE, 2017, pp. 1–5.

[9] R. Harwahyu *et al.*, "Optimization of Random Access Channel in NB-IoT," *IEEE Internet Things J.*, vol. 5, no. 1, pp. 391–402, 2017.

[10] ——, "Fair Configuration Scheme for Random Access in NB-IoT with Multiple Coverage Enhancement Levels," *IEEE Trans. Mobile Comput.*, 2019.

[11] Y. Sun *et al.*, "Throughput Modeling and Analysis of Random Access in Narrowband Internet of Things," *IEEE Internet Things J.*, vol. 5, no. 3, pp. 1485–1493, 2017.

[12] S. Martiradonna, G. Piro, and G. Boggia, "On the Evaluation of the NB-IoT Random Access Procedure in Monitoring Infrastructures," *Sensors*, vol. 19, no. 14, p. 3237, 2019.

[13] N. Jiang *et al.*, "RACH Preamble Repetition in NB-IoT Network," *IEEE Commun. Lett.*, vol. 22, no. 6, pp. 1244–1247, 2018.

[14] S. Narayanan *et al.*, "Performance Analysis of NB-IoT Random Access Channel," in *Interactive Mobile Communication, Technologies and Learning*. Springer, 2019, pp. 1118–1128.

[15] Y. Liu *et al.*, "Random Access Performance for Three Coverage Enhancement Groups in NB-IoT Networks," in *IEEE Global Communications Conf. (GLOBECOM)*. IEEE, 2019, pp. 1–6.

[16] ——, "Analysis of Random Access in NB-IoT Networks with Three Coverage Enhancement Groups: A Stochastic Geometry Approach," *arXiv preprint arXiv:2009.06782*, 2020.

[17] T. Kim *et al.*, "An Enhanced Access Reservation Protocol with a Partial Preamble Transmission Mechanism in NB-IoT Systems," *IEEE Commun. Lett.*, vol. 21, no. 10, pp. 2270–2273, 2017.

[18] R. Harwahyu *et al.*, "Repetitions versus Retransmissions: Tradeoff in Configuring NB-IoT Random Access Channels," *IEEE Internet Things J.*, vol. 6, no. 2, pp. 3796–3805, 2019.

[19] J. Zhang, D. Xie, and X. Wang, "TARA: An Efficient Random Access Mechanism for NB-IoT by Exploiting TA Value Difference in Collided Preambles," *IEEE Trans. Mobile Comput.*, 2020.

[20] M. Agiwal, M. K. Maheshwari, and H. Jin, "Power Efficient Random Access for Massive NB-IoT Connectivity," *Sensors*, vol. 19, no. 22, p. 4944, 2019.

[21] "NB-IoT Deployment Guide to Basic Feature Set Requirements," https://www.gsma.com/iot/wp-content/uploads/2019/07/201906-GSMA-NB-IoT-Deployment-Guide-v3.pdf, accessed: May 2020.

[22] https://mosaic-simulamet.com/nbiot-randomaccess, accessed: Nov. 2020.

[23] https://mosaic-simulamet.com/nbiotcoverage, accessed: Nov. 2020.

[24] K. Kousias *et al.*, "Coverage and Deployment Analysis of Narrowband Internet of Things in the Wild," *IEEE Commun. Mag.*, vol. 58, no. 9, pp. 39–45, 2020.

[25] A. Hoglund *et al.*, "Overview of 3GPP Release 14 Enhanced NB-IoT," *IEEE Network*, vol. 31, no. 6, pp. 16–22, 2017.

[26] B. Martinez *et al.*, "Exploring the Performance Boundaries of NB-IoT," *IEEE Internet Things J.*, vol. 6, no. 3, pp. 5702–5712, 2019.

[27] H. Fattah, *5G LTE Narrowband Internet of Things (NB-IoT)*. CRC Press, 2018.

[28] 3GPP, "LTE; Evolved Universal Terrestrial Radio Access (E-UTRA); Requirements for support of radio resource management," TS 36.133, https://www.etsi.org/deliver/etsi_ts/136100_136199/136133/13.03.00_60/ts_136133v130300p.pdf, accessed: May 2020.

[29] ——, "LTE; Evolved Universal Terrestrial Radio Access (E-UTRA); Medium Access Control (MAC) protocol specification," TS 36.321, https://www.etsi.org/deliver/etsi_ts/136300_136399/136321/13.05.00_60/ts_136321v130500p.pdf, accessed: May 2020.

[30] A. Hoglund *et al.*, "3GPP Release 15 Early Data Transmission," *IEEE Commun. Standards Mag.*, vol. 2, no. 2, pp. 90–96, 2018.

[31] 3GPP, "Study on RAN Improvements for Machine-type Communications," TR 37.868, https://www.3gpp.org/DynaReport/FeatureOrStudyItemFile-450015.htm, accessed: May 2020.

[32] ——, "Technical Specification Group Radio Access Network; Evolved Universal Terrestrial Radio Access (E-UTRA); Physical channels and modulation," TS 36.211, https://www.arib.or.jp/english/html/overview/doc/STD-T104v4_10/5_Appendix/Rel13/36/36211-d20.pdf, accessed: May 2020.

[33] ——, "LTE; Evolved Universal Terrestrial Radio Access (E-UTRA); Physical layer procedures," TS 36.213, https://www.etsi.org/deliver/etsi_ts/136200_136299/136213/13.00.00_60/ts_136213v130000p.pdf, accessed: May 2020.

[34] https://www.u-blox.com/sites/default/files/SARA-N2-N3_ATCommands_\%28UBX-16014887\%29.pdf, accessed: Nov. 2020.

[35] 3GPP, "Cellular System Support for Ultra Low Complexity and Low Throughput Internet of Things," TR 45.820, https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=2719, Accessed: Nov. 2020.

[36] I. Guyon *et al.*, "Gene Selection for Cancer Classification using Support Vector Machines," *Machine learning*, vol. 46, no. 1-3, pp. 389–422, 2002.

[37] R. Tibshirani, "Regression Shrinkage and Selection via the Lasso," *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 58, no. 1, pp. 267–288, 1996.

[38] H. He *et al.*, "ADASYN: Adaptive Synthetic Sampling Approach for Imbalanced Learning," in *IEEE Int'l Joint Conf. on Neural Networks (IEEE World Congress on Computational Intelligence)*. IEEE, 2008, pp. 1322–1328.

[39] H. He and E. A. Garcia, "Learning from Imbalanced Data," *IEEE Trans. Knowl. Data Eng.*, vol. 21, no. 9, pp. 1263–1284, 2009.

[40] L. Breiman, "Random Forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.

[41] C.-H. Wei, R.-G. Cheng, and S.-L. Tsao, "Modeling and Estimation of One-Shot Random Access for Finite-User Multichannel Slotted ALOHA Systems," *IEEE Commun. Lett.*, vol. 16, no. 8, pp. 1196–1199, 2012.

[42] C.-H. Wei, G. Bianchi, and R.-G. Cheng, "Modeling and Analysis of Random Access Channels with Bursty Arrivals in OFDMA Wireless Networks," *IEEE Trans. Wireless Commun.*, vol. 14, no. 4, pp. 1940–1953, 2014.

[43] R. K. Jain, D.-M. W. Chiu, and W. R. Hawe, "A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Computer Systems," *Eastern Research Laboratory, Digital Equipment Corporation, Hudson, MA*, 1984.

[44] G. Piro *et al.*, "Simulating LTE Cellular Systems: An Open-Source Framework," *IEEE Trans. Veh. Technol.*, vol. 60, no. 2, pp. 498–513, 2010.

[45] 3GPP, "LTE; Evolved Universal Terrestrial Radio Access (E-UTRA); Physical layer measurements," TS 36.214, https://www.etsi.org/deliver/etsi_ts/136200_136299/136214/11.01.00_60/ts_136214v110100p.pdf, accessed: May 2020.

[46] A. Azari *et al.*, "Latency-Energy Tradeoff based on Channel Scheduling and Repetitions in NB-IoT Systems," in *IEEE Global Communications Conf. (GLOBECOM)*. IEEE, 2018, pp. 1–7.

[47] H. Li *et al.*, "Accurate Performance Modeling of Uplink Transmission in NB-IoT," in *IEEE Int'l Conf. on Parallel and Distributed Systems (ICPADS)*. IEEE, 2018, pp. 910–917.

[48] V. Erceg *et al.*, "An Empirically based Path Loss Model for Wireless Channels in Suburban Environments," *IEEE J. Sel. Areas Commun.*, vol. 17, no. 7, pp. 1205–1211, 1999.

[49] M. Chafii, F. Bader, and J. Palicot, "Enhancing Coverage in Narrow Band-IoT using Machine Learning," in *IEEE Wireless Communications and Networking Conf. (WCNC)*. IEEE, 2018, pp. 1–6.

[50] L.-S. Chen *et al.*, "Adaptive Repetition Scheme with Machine Learning for 3GPP NB-IoT," in *IEEE Pacific Rim Int'l Symp. on Dependable Computing (PRDC)*. IEEE, 2018, pp. 252–256.

[51] N. Jiang *et al.*, "Cooperative Deep Reinforcement Learning for Multiple-Group NB-IoT Networks Optimization," in *IEEE Int'l Conf. on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 8424–8428.

[52] ——, "Reinforcement Learning for Real-Time Optimization in NB-IoT Networks," *IEEE J. Sel. Areas Commun.*, vol. 37, no. 6, pp. 1424–1440, 2019.

[53] N. Jiang, Y. Deng, and A. Nallanathan, "Traffic Prediction and Random Access Control Optimization: Learning and Non-learning based Approaches," *arXiv preprint arXiv:2002.07759*, 2020.

BIOGRAPHIES

Giuseppe Caso [M] is a Postdoctoral Fellow with the Mobile Systems and Analytics (MOSAIC) Department at Simula Metropolitan (SimulaMet). In 2016, he received the Ph.D. degree from Sapienza University of Rome, where he was a Postdoctoral Fellow until 2018. From 2012 to 2018, he has held visiting positions at Leibniz University of Hannover, King's College London, Technical University of Berlin, and Karlstad University. His research interests include cognitive communications, distributed learning, and IoT technologies.

Konstantinos Kousias [S] is a PhD Student at the Faculty of Mathematics and Natural Sciences in University of Oslo affiliated with Simula Research Laboratory. He received his B.Sc. and M.Sc. degrees from the Department of Electrical and Computer Engineering at University of Thessaly in Volos, Greece. His research focuses on the empirical modeling and evaluation of mobile networks using data analytics and artificial intelligence.

Özgü Alay [M] received the B.Sc. and M.Sc. degrees in Electrical and Electronic Engineering from Middle East Technical University, Turkey, and Ph.D. degree in Electrical and Computer Engineering at Tandon School of Engineering at New York University. She is an Associate Professor at University of Oslo, Norway and Head of MOSAIC Department at Simula Metropolitan, Norway. Her research interests include 5G and IoT networks, drone communications, mobile multimedia systems, and multipath protocols. She is author of more than 70 peer-reviewed IEEE and ACM publications and she actively serves on technical boards of major conferences and journals.

Anna Brunstrom [M] received a B.Sc. in Computer Science and Mathematics from Pepperdine University, CA, in 1991, and a M.Sc. and Ph.D. in Computer Science from College of William & Mary, VA, in 1993 and 1996, respectively. She joined the Department of Computer Science at Karlstad University, Sweden, in 1996, where she is currently a Full Professor and Research Manager for the Distributed Systems and Communications Group. Her research interests include Internet architectures and protocols, techniques for low latency and multi-path communications, and performance evaluation of mobile broadband systems including 5G. She has authored/coauthored over 170 international peer-reviewed journal and conference papers.

Marco Neri received the B.Sc. and M.Sc. in Electrical Engineering from Sapienza University of Rome, Italy. He joined Rohde&Schwarz in 2017 as Application Engineer for Mobile Network Testing. His focus is on 5G and cellular IoT testing activities worldwide.