

Characterizing RNTI Allocation and Management in Mobile Networks

Giulia Attanasio
IMDEA Networks Institute
University Carlos III
Madrid, Spain
giulia.attanasio@imdea.org

Claudio Fiandrino, Marco Fiore, Joerg
Widmer
IMDEA Networks Institute
Madrid, Spain
firstname.lastname@imdea.org

ABSTRACT

The characterization of user behavior is key to perform traffic analysis, modeling and optimization of network components and protocols. This is especially true for 5G and beyond networks that will heavily rely on machine learning for network optimization. In Base Station (BS) traffic traces, users are uniquely identified by a Radio Network Temporary Identifier (RNTI) assigned to them. RNTIs are not bound to a user but are reused upon expiration of an inactivity timer, whose duration is operator dependent. This implies that, over time, multiple users can be mapped to the same RNTI ID in diverse ways. Hence, when using real-world radio access measurement traces for traffic analysis, distinguishing individual users within the RNTI space is a non-trivial task. In this paper, we provide the first in-depth study of the RNTI allocation process and shed light not only on the setting of the inactivity timer, but also on the relationship of the RNTI allocation scheme and the user characteristics. For this, we collect a large dataset of mobile traffic from multiple BSs of several mobile network operators. The analysis of the decoded control messages of the BS unveils that the RNTI allocation process changes over time depending on the BSs observed load and time of day. We also observe that the RNTI expiration threshold is on the order of minutes, and demonstrate how using thresholds around 10 s that are reported in the vast majority of the literature can bias subsequent analyses. Overall, our work provides an important step towards dependable mobile network trace analysis, and lays more solid foundations to research relying on traffic traces for data-driven analysis and simulation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MSWiM '21, November 22–26, 2021, Alicante, Spain

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-9077-4/21/11...\$15.00

<https://doi.org/10.1145/3479239.3485685>

CCS CONCEPTS

• **Networks** → **Network dynamics**; **Network experimentation**; **Network dynamics**.

KEYWORDS

5G, LTE, RNTI expiration threshold, RNTI allocation

ACM Reference Format:

Giulia Attanasio and Claudio Fiandrino, Marco Fiore, Joerg Widmer. 2021. Characterizing RNTI Allocation and Management in Mobile Networks. In *Proceedings of the 24th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM '21), November 22–26, 2021, Alicante, Spain*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3479239.3485685>

1 INTRODUCTION

5G subscriptions are forecast to reach 3.5 billion globally by the end of 2026 [9]. 5G and beyond mobile networks are expected to support a wide range of demanding use cases (*e.g.*, Augmented Reality, 3D video, industry automation, smart grid and smart cities) and to serve unprecedented traffic demands. For this, (beyond) 5G networks should evolve to quickly adapt to changes in the network state and traffic [19]. For this reason, these systems are becoming increasingly complex, heterogeneous, dynamic and dense, which makes it hard to correctly model their behavior [12].

To advance mobile network technologies and components, data-driven approaches are becoming more and more popular. Examples range from optimization for robustness to mobility (*i.e.*, self-configuration of handover parameters) [13], slicing [6], automatic resource scaling of 5G core virtualized network functions [5], down to the physical layer to optimize beam management of radios operating at millimeter-wave frequencies [22]. The characterization of the user behavior is a key aspect to properly analyze traffic traces that are used for data-driven simulations.

In mobile networks, a BS (*i.e.*, the gNB in 5G NR or the eNB in LTE) identifies users by means of temporary and permanent identifiers [4]. The International Mobile Equipment Identity (IMEI) and the International Mobile Subscriber Identity (IMSI) are permanent and identify respectively the

device and the mobile subscriber. Upon authentication with IMSI, an encrypted channel is set up and from that moment on the terminal uses temporary identifiers like Temporary Mobile Subscriber Identity (TMSI) for communication with the core network, while the RNTI identifies the User Equipment (UE)¹ within the radio cell. RNTIs themselves do not reveal the user’s identity and expire after the user is inactive for a certain period of time. In practice, the user inactivity is monitored with a timer and when the inactivity timer exceeds a given threshold, the user-RNTI association is released. Then, the same RNTI can be re-assigned to another user. The 3GPP standards define the process, but leave the specific implementation of the timer, threshold and RNTI value selection to the operators. This makes reverse engineering single users from RNTI sequences in measurements data a complex task. At the same time, the incorrect identification of user sessions has clear consequences on traffic analysis: counting multiple users as one, or vice versa, results in incorrect estimates of traffic distribution, user activity, and resource allocation, which affects modeling and simulations and potentially biases the conclusions.

In this paper, we provide a first in-depth study of the RNTI allocation process. We use the well-known FALCON [10] LTE passive monitoring tool, and collect traffic traces from 8 different BSs of three major European network operators (see Fig. 1 for an example of the measurement setup) during a 6 months measurement campaign. Our analysis of the collected traces sheds light on a number of aspects related to RNTI allocation and management, including insights on the actual values of inactivity thresholds and on the allocation schemes adopted by different operators. We demonstrate that different RNTI allocation schemes may be used at a same BS, and investigate correlations between the user behavior and the employed RNTI allocation scheme. Overall, our results help to better understand and analyze BS measurement data, and pave the way for more dependable research on data-driven modeling and simulations of (beyond) 5G mobile network components.

The major contributions of this work are as follows.

- We provide a methodology to correctly set the RNTI inactivity threshold for trace analysis, to properly identify the life cycle of users connected to a BS. In contrast to past research that commonly sets the inactivity threshold to ~10 s, our analysis shows that actual values for the threshold are in the 1–10 minutes range.
- We characterize the RNTI allocation process adopted by multiple network operators, and find that different operators implement different RNTI allocation processes, that may change over time at the same BS.
- We characterize user traffic patterns given the specific

¹In this paper, we use the terms UE and *user* interchangeably.

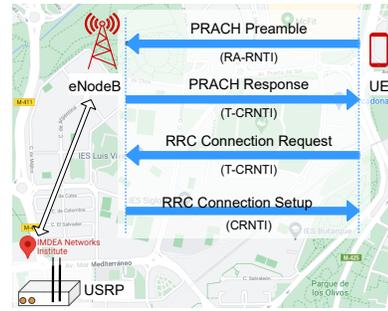


Figure 1: Example of the suburban area measurement setup and RNTI allocation procedure during RA

RNTI allocation process. We find a correlation between the given RNTI allocation scheme and frame occupancy in terms of physical resource block usage.

The rest of the paper is organized as follows. Sec. 2 provides background on user identification in mobile networks. Sec. 3 presents our real-world LTE traffic dataset, while Sec. 4 delves into the RNTI analysis. Sec. 5 evaluates the impact of incorrect inactivity thresholds. Sec. 6 provides a discussion of the research areas that may benefit from our results. Finally, Sec. 7 draws conclusions.

2 TEMPORARY USER IDENTIFIERS IN MOBILE NETWORKS

In order to get service from the network, a mobile phone that is just turned on has to perform an Initial Access procedure, which consists of three steps: cell search, system information extraction and random access procedure. Once these steps are completed, the user switches from IDLE to CONNECTED mode and can now communicate with the BS over scheduled channels. Specifically, after the random access procedure is completed, the UE is mapped to a Cell Radio Network Temporary Identifier value (C-RNTI). A new RA procedure is required if the user moves between BSs during handovers, or if the user remains inactive within the same cell for a time that is not explicitly defined by the standards.

The random access procedure through which a C-RNTI is assigned to an incoming user is summarized in Fig. 1. The UE sends a Physical Random Access Channel (PRACH) Preamble signal to the BS to initiate the random access. The PRACH Preamble contains the Random Access RNTI (RA-RNTI) which is specified by the UE according to the sending subframe. The BS replies to the PRACH Preamble by sending a PRACH Response. This response contains the temporary C-RNTI (TC-RNTI) that will be used by the UE for the RRC connection request. If the connection is successfully set up, the TC-RNTI is promoted to C-RNTI. This value will be then used to uniquely identify UE traffic within the cell. In the rest of this paper, we will refer to C-RNTI simply as RNTI.

Table 1: RNTI Assignment in LTE

Value (hexa-decimal)	Type of RNTI
0000	N/A
0001-0960	RA-RNTI, C-RNTI, Semi-Persistent Scheduling C-RNTI, TC-RNTI, eIMTA-RNTI, TPC-PUCCH-RNTI, TPC-PUSCH-RNTI, SL-RNTI, G-RNTI, SL-V-RNTI, UL Semi-Persistent Scheduling V-RNTI, SL Semi-Persistent Scheduling V-RNTI, SRS-TPC-RNTI, and AUL C-RNTI
0961-FFF3	C-RNTI, Semi-Persistent Scheduling C-RNTI, eIMTA-RNTI, TC-RNTI, TPC-PUCCH-RNTI, TPC-PUSCH-RNTI, SL RNTI, G-RNTI, SL-V-RNTI, UL Semi-Persistent Scheduling V-RNTI, SL Semi-Persistent Scheduling V-RNTI, SRS-TPC-RNTI and AUL C-RNTI
FFF4-FFF8; FFF9	Reserved for future use; SI-RNTI
FFFA; FFFB; FFFC	SC-N-RNTI; SC-RNTI; CC-RNTI
FFFD; FFFE; FFFF	M-RNTI; P-RNTI; SI-RNTI

Table 2: RNTI Assignment in 5G NR

Value (hexa-decimal)	Type of RNTI
0000	N/A
0001-FFEF	RA-RNTI, TC-RNTI, C-RNTI, MCS-C-RNTI, CS-RNTI, TPC-PUCCH-RNTI, TPC-PUSCH-RNTI, TPC-SRS-RNTI, INT-RNTI, SFI-RNTI, and SP-CSI-RNTI
FFF0-FFFD; FFFE; FFFF	Reserved for future use; P-RNTI; SI-RNTI

According to 3GPP standards [1], the pool of available RNTI IDs is limited and ranges from 0x1 to 0xFFFF. A subset of them is not allocated to any UE explicitly, but reserved for specific messages such as paging (P-RNTI), system information (SI-RNTI) and power control purposes (TPC-RNTI). Tab. 1 provides the complete list of LTE RNTI values. For 5G NR, 3GPP specifications [3] introduce new types of RNTI to support novel uses such as the SFI-RNTI and the MCS-C-RNTI (see Tab. 2 for a complete list). Specifically, the first one is generally assigned to a group of users and used for the notification of slot format information. The second enables for both downlink and uplink the usage of an alternative MCS table for the scheduling of packets with high reliability. For our analysis over LTE, we consider the RNTI associated to actual UE transmissions within the Physical Downlink Control Channel (PDCCH), *i.e.*, the ones whose values are between 0x3D and 0xFFFF3.

Once an RNTI is assigned to a UE, it is maintained until the user remains inactive for a certain amount of time. In that case, the corresponding RNTI ID value is released and it can be reassigned to a new user. Thus, if a UE whose RNTI was released wants to reconnect to the BS, it needs to perform a new access procedure to obtain a new RNTI. Existing works set the inactivity timer to values that are close to the System Frame Number cycle. This counter, incrementally assigned to each LTE frame, repeats itself after 1024 frames, *i.e.*, 10.24 s. Specifically, other research works sets the inactivity timer to 10 s [14], 10.15 s [27], 10.24 s [7] and 11.57 s [15]. Another work [21] observed that an LTE smartphone connected to a major US operator maintained the same RNTI ID for over 4 hours but does not mention whether the smartphone was in fact exchanging traffic with the BS or not.

Table 3: Collected dataset

BS	Operator	Channel bandwidth	Location	Total collected time
1	Operator 1	10 MHz	inner city	2 months
2	Operator 1	20 MHz	inner city	2 months
3	Operator 1	20 MHz	suburban area	1 month
4	Operator 2	10 MHz	inner city	2 months
5	Operator 2	20 MHz	inner city	2 months
6	Operator 2	10 MHz	suburban area	1 month
7	Operator 3	20 MHz	suburban area	1 month
8	Operator 3	20 MHz	city center	1 month

The few studies that investigated RNTI allocation schemes mainly focused on the blind decoding mechanism performed by the UE at the receiver side. In order to decode the Downlink Control Information (DCI) message transmitted through the PDCCH, the UE will use its RNTI to scramble the Cyclic Redundancy Check part of the radio channel messages over a set of possible control channel elements. Within such set, the PDCCH carries information about the resource allocation for each user. In order to increase control channel element utilization several allocation schemes have been proposed, among them a table-based scheme in [8] and preamble power estimation based approach in [23].

Ultimately, no previous work analyzes the RNTI expiration threshold or the RNTI allocation process adopted by different network operators. Also, no characterization of traffic patterns according to the used allocation scheme is presently available. Our study closes these important gaps.

3 MEASUREMENT DATASET

This section presents the methodology for the data collection, the dataset of raw DCI information from different BSs and operators and the pre-processing required for our study.

For our analysis, we consider real LTE traffic information gathered from the PDCCH through FALCON [10]. The data was collected for over 6 months from 8 different BSs of 3 major European operators. During this period we monitored up to two different BSs at the same time. All the considered BSs are located on different sites and measurements were collected in the city center, in the suburban area and in the inner city. Tab. 3 summarizes the characteristics of our dataset. The gathered PDCCH data contains per user scheduling information in both uplink and downlink directions. Within each frame and transmission time interval, FALCON reports per user specific resource allocation in terms of PRB, TBS and MCS. The collected raw traces from each monitored BS are then grouped per unique registered RNTI. For our study, we collect traffic traces running FALCON with two setups, an Intel Core CPU at 3.6 GHz connected via 10 Gigabit Ethernet to an X310 USRP and an Intel Core CPU at 3.4 GHz connected via USB 3.0 to a B210 USRP (located at the USRP site in Fig. 1). Both PCs run Linux Ubuntu 20.04. We do not observe any difference in terms of DCI decoding success rate with the two setups.

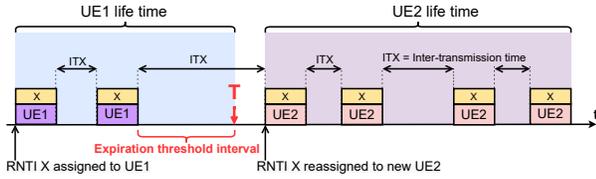


Figure 2: Diagram illustrating inter-transmission times (ITX) for the same RNTI ID (X). If the inter-transmission time exceeds the expiration threshold T , the RNTI ID can be reassigned to a new user.

Since we are only interested in user traffic, we discard RNTI ID values that are not UE-related. Specifically, we remove the ones associated to paging and system information (0xFFFF4 to 0xFFFFF) and to the PRACH channel (0x1 to 0x3), and only include the RNTI values from actual UE transmissions, *i.e.*, from 0x3D to 0xFFFF3, see Tab. 1.

4 RNTI ANALYSIS

Our RNTI analysis has two objectives. First, we investigate how RNTIs are allocated and reassigned to identify the RNTI timer expiration thresholds. Second, we characterize the RNTI behavior across different BSs and operators to analyze how the RNTI allocation varies over daily and weekly cycles.

4.1 Threshold Detection

The data traffic associated to an RNTI is generated by different UEs that take the same identifier in different time periods (see Sec. 1). Since the set of available RNTI ID is limited, the BS reassigned identifiers that have been inactive for a certain period to new users. We denote as *user life time* the time elapsed between the first and the last transmission of a given user. We define the inter-transmission time as the time elapsed among two consecutive transmissions for the same RNTI. Then, if the inter-transmission time is smaller than the *expiration threshold* T , we consider that the consecutive transmissions belong to the same UE; if it is instead greater than T , we consider the transmissions to be from different UEs, as shown in Fig. 2.

In order to verify if users associated to the same RNTI ID can be correctly discriminated through the commonly adopted threshold of 10 s, we plot the probability density function (PDF) of the inter-transmission times for each registered RNTI. Fig. 3 presents the results for all the monitored operators and BSs. For simplicity, we report only downlink traffic given that is predominant in mobile networks, but our approach is also valid for uplink traffic. In fact, we confirm that the fraction of uplink traffic is tiny with respect to the downlink traffic. The results show inter-transmission times from 0 to 60 s and, for each plot, we show an inset in the range from 2 to 30 s.

From Fig. 3 we can observe that, across all BSs, the highest number of occurrences, normalized so that the area under the histogram integrates to 1, is always registered for values that are smaller or equal than 1 s, *i.e.*, 0.85 for Operator 1 BS 4, 0.98 for Operator 3 BS 7 and 0.96 for Operator 1 BS 3.

The threshold T determines the end of the user life time. As soon as the user remains inactive for a period of time that exceeds T , its RNTI becomes available for reassignment. Within its life time, a user transmits $k \in \mathbb{N}$ times. For each new RNTI reassignment, we expect to see $k - 1$ inter-transmission time values smaller than T and one inter-transmission time value that exceeds T . Considering all active users, the probability of observing inter-transmission times greater than T (since a new user arrives at that time and happens to be assigned the same RNTI) should be much smaller than the probability of observing values smaller than T (where the same user is still active). Thus, we expect that the number of registered occurrences to drop sharply for values larger than T .

Fig. 3 shows that the number of occurrences within the interval 2-30 s is two order of magnitude smaller with respect to the one registered within the interval 0-1 s. However, no significant drop can be observed in the interval 10-11 s, where according to existing literature, the threshold should be set.

Our aim is that to identify the proper threshold for correct user life time discrimination. However, according to the results presented in Fig. 3, there is no indication that the commonly adopted threshold of 10 s is correct. In fact, an indication of the correct threshold setting can be derived from a significant variation of the inter-transmission time values. A change in behavior can be appreciated by looking at the survival function (complementary CDF) of consecutive RNTI transmissions. We evaluate the inter-transmission times considering one week of continuous traffic. Fig. 4 shows the results for each of the monitored BSs.

A suitable threshold for all BSs and operators falls within the range from 30 s to 100 s. In fact, from Fig. 4(b), we observe four main stages. First, there is a rapid decay phase (red background) that ranges from 1 to 3 s for all the considered BSs. Most inter-transmission times are in this range, *i.e.*, the majority of consecutive user transmissions are within 3 s. The rapid decay is followed by a smoothly decreasing stage (green background) that spans over a larger inter-transmission time interval. After this stage, we can identify a leveling out period (gray background), followed by the tail of the distribution (blue background) for which the number of observed values drops significantly.

For the threshold setting we are interested in the change of behavior that occurs between the second and the third phase, *i.e.*, the smoothly decreasing stage and the leveling out period. In fact, from the beginning of the leveling out period up to the tail of the distribution few inter-transmission times values can be observed as the survival function remains

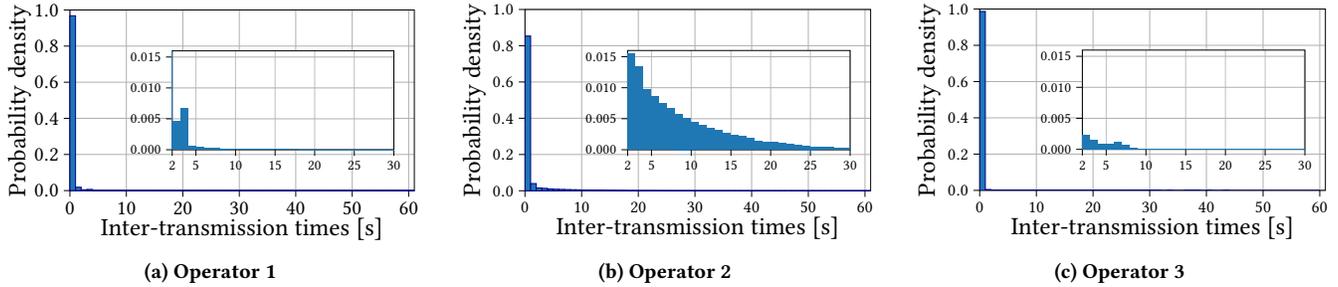


Figure 3: PDF of inter-transmission times for different operators for values smaller than 60 s.

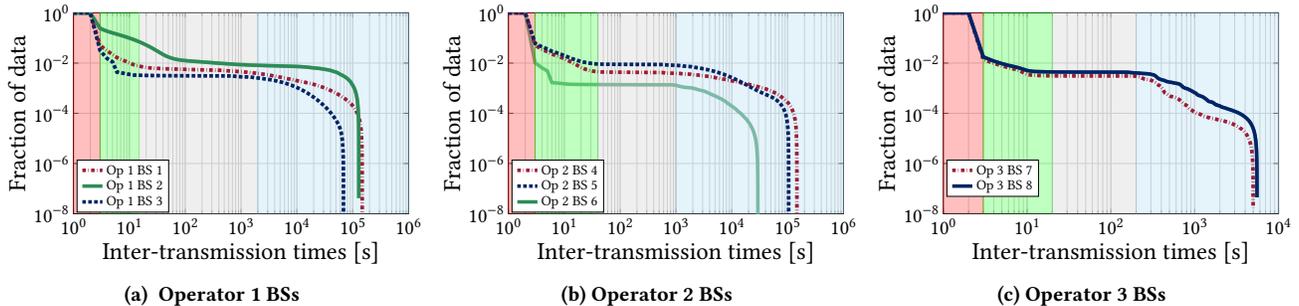


Figure 4: Survival function of the inter-transmission times with a rapid decay phase (red), a smoothly decreasing stage (green), the leveling out period (gray) and distribution tail (blue).

almost constant. Hence, smaller inter-transmission times within the rapid decay and the smoothly decreasing stage identify sessions of the same user. Thus, a possible threshold can be set right after the knee point between the smoothly decreasing stage and the leveling out period. We can neither consider the tail of the distribution since the number of registered observations is very sparse.

Furthermore, the significant change that occurs between the rapid decay and the smoothly decreasing stage would not represent a reasonable interval for the threshold setting. In this case the threshold would be 3 s. However, it would not make sense for the BS to set such a small value since also data continuity should be guaranteed to connected users. In fact, such a small value would trigger too many re-establishment procedures leading the system to collapse, a phenomenon known as signalling storm [24]. Thus, in the first place, we set the expiration threshold to 60 s for BS 1, 2, 4, 5, and to 30 s for BS 3, 6, 7 and 8. We adjust the estimated expiration threshold for all the considered BSs, according to the methodology described in Sec 4.4.

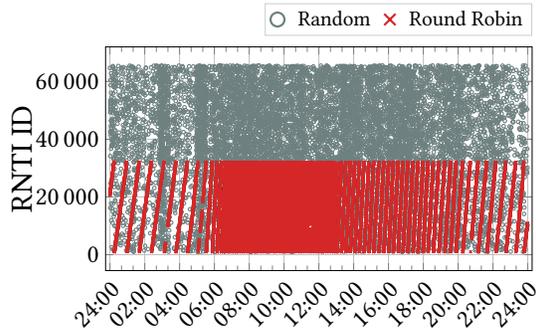
Takeaway message. The threshold for the RNTI expiration should be set to much higher values than those around 10 s proposed in the literature. A reasonable threshold value is the beginning of the leveling out period identified through the survival function of the inter-transmission times (Fig. 4).

4.2 RNTI Allocation Schemes

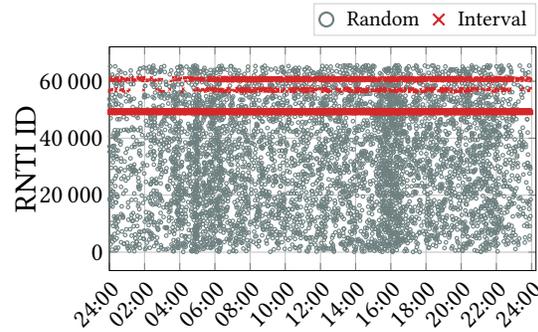
We now investigate how the monitored BSs assign RNTIs. For this purpose, we set the inactivity expiration threshold to 60 s for BS 1, 2, 4 and 5 and to 30 s for BS 3 and 6 according to the results presented in Sec. 4.1.

Fig. 5(a) presents, over a period of 24 hours, the allocation schemes adopted by Operator 1 BS 2. From this analysis, it emerges that the BSs from Operator 1 and 2 use two RNTI allocation strategies combined, namely random (gray circles) and round robin (red crosses along slopes). Specifically, we found that the RNTI IDs whose value is greater than 32060 (hereafter defined as 1-headed, according to the first digit of their binary representation), are always allocated randomly. Instead, the RNTI IDs that are smaller than 32060 (hereafter defined as 0-headed), can be allocated either randomly or according to a round robin scheme.

The BSs from Operator 3 instead, use a random allocation strategy (gray circles) and an interval-based scheme (red crosses along horizontal lines) as show in Fig. 5(b). In particular, the interval-based scheme allocates randomly a subset of RNTI IDs values. However, those values are allocated per hour 15 times more frequently than the other random allocated ones (IDs not belonging to the interval). For both Operator 3 BSs, we identified three intervals, *i.e.*, low, medium and high according to the range that they cover. With the sole exception of one interval of BS 7, all the other intervals cover



(a) Random (gray) and round robin (red) allocation strategies for Operator 1 BS 3 over 24 hours.



(b) Random (gray) and interval (red) allocation strategies for Operator 3 BS 8 over 24 hours.

Figure 5: Allocation strategies over time

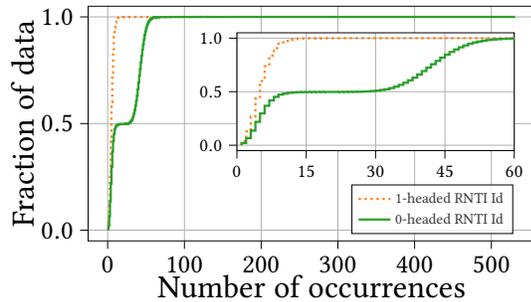


Figure 6: Occurrences of 0-headed and 1-headed RNTI IDs for Operator 1 BS 1

a range of 1400 RNTI ID values. Specifically, the interval that ranges from RNTI ID 48600 to 50000, is common to both BSs. The other two are instead different: 22000-24000, 3000-4400 for BS 7 and 56200-57600, 60000-61400 for BS 8.

Motivated by the fact that operators can allocate RNTI to incoming users following different strategies, we investigate the behavior of 0-headed and 1-headed RNTI with respect to the random/round robin allocation scheme. Thus, we evaluate the CDF of the number of occurrences for each registered RNTI over one week. This analysis is performed on the traffic traces collected from Operator 1 and 2 only, since Operator 3 does not allocate RNTI according to a round robin strategy. The results presented for Operator 1 BS 4 in Fig. 6 are consistent across all BSs from Operator 1 and 2, but not presented here for space reasons. For both operators, the 95% of 1-headed RNTI IDs, *i.e.*, the ones associated to random periods only, are concentrated within the interval that ranges from 1 to 10. From the same figure, the analysis of 0-headed RNTI IDs reveals the existence of two groups. The first group registers a significant smaller number of occurrences with respect to the second one. Specifically, the first one ranges from 1 to 15, while the second from 30 up to 531. By analyzing the RNTI IDs of both groups, we discover that the first one only comprises even IDs, while the second

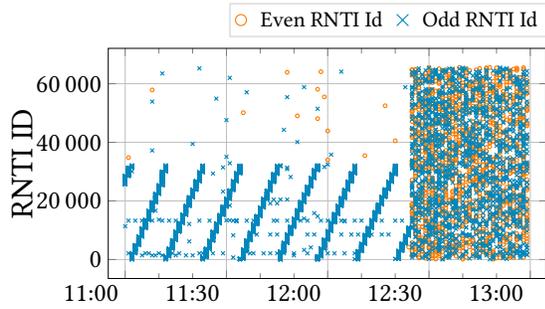
only comprises odd IDs values. Motivated by this finding, in the next Section we characterize odd and even RNTI IDs.

Takeaway message. Operators adopt different RNTI allocation strategies. Specifically, they combine random allocation with either round robin or interval-based schemes.

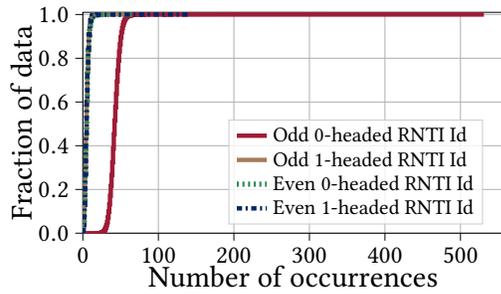
4.3 RNTI Characterization

We plot the inter-transmission times of odd and even RNTI. Fig. 7(a) presents the results for Operator 1 BS 4 over a period of 2 hours. We find that 1-headed RNTIs are reserved for random allocation only, no matter if they are even or odd. Instead, 0-headed RNTI are used for either random and round robin allocation. Specifically, only odd 0-headed RNTI are used by the BS for the round robin mechanism. This result holds across all monitored Operator 1 and Operator 2 BSs. Fig. 7(b) illustrates the CDF of the number of occurrences of even/odd 0-headed and 1-headed RNTI IDs over one week of traffic data from Operator 1 BS 3. We observe that the number of occurrences is smaller for odd 1-headed and for both 0-headed and 1-headed even RNTI IDs with respect to odd 0-headed ones. The highest number of occurrences, *i.e.*, > 100 is registered for spurious 0-headed odd RNTI ID (see Sec. 4.4). As stated in Sec. 4.2, Operator 3 combines a random and an interval scheme allocation. The same analysis reveals that Operator 3 does not allocate even and odd RNTI differently within random and interval periods.

We now analyze whether RNTI assigned according to different strategies (random, round robin, and interval-based) have different traffic characteristics. Furthermore, we investigate the relationship between the traffic load (that exhibits patterns, *e.g.*, day/night) registered at the BS and the RNTI allocation scheme that is used. From the analysis of Operator 3 traffic traces, we find that the medium interval for both BS is less used during night periods and early mornings, within a range that approximately lasts from 23:00 to 6:00. Fig. 8(a), shows that as soon as the medium interval (red) is progressively less used, a drop in the BS load is registered too. The



(a) Odd RNTI IDs can be assigned both randomly and in round robin (linear slopes), whereas even RNTI IDs are only assigned randomly.



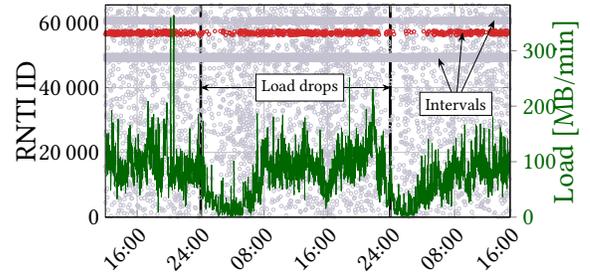
(b) Occurrences of odd and even 0-headed and 1-headed RNTI IDs

Figure 7: Odd and even RNTI ID analysis

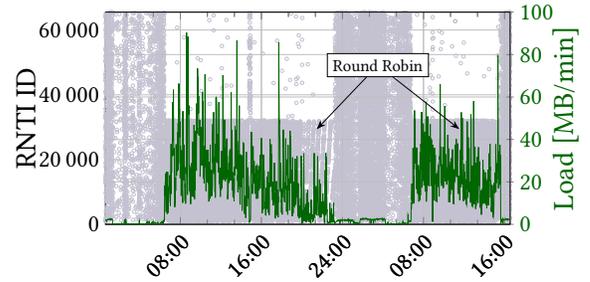
first load drop of 97 % occurs at 24:00 and a second one of 94 % is registered around 24:00 of the next day. In both cases, as soon as the load grows over approximately 80 MB, RNTI IDs start to be assigned again to that interval. Our conjecture is that the BS changes the allocation to random for that interval whenever it detects a low level of traffic volume per period. The same considerations hold for Operator 2 BS 4 as shown in Fig. 8(b). In this case the BS swaps the allocation from random to round robin as soon as the load grows above 20 MB (08:00 of both consecutive days). In the same way, the allocation swaps from round robin to random as soon as the load drops below 15 MB (24:00 and 16:00).

Thus, we consider frame occupancy in terms of PRB utilization which is recommended as the evaluation criteria of load [29] and depends on the application in use and on the amount of resources available. PRB utilization is defined as the ratio between the number of PRBs used for transmission and the total number of PRB of the BS, in each transmission time interval [2]. By considering one week of continuous traffic, we evaluate the CDF of PRB utilization. Fig. 9 presents the results for Operator 1 BS 3 and Operator 3 BS 8.

Fig. 9(a) shows similar results between odd 1-headed RNTI IDs, even RNTI IDs and Operator 3 RNTI IDs that are allocated randomly. In this case, the 90 % of the IDs registers a frame occupancy smaller than 10 %. Instead, when we look



(a) Operator 3 BS 8 (30 hours). The interval RNTI allocation is highlighted in red.



(b) Operator 2 BS 4 (40 hours)

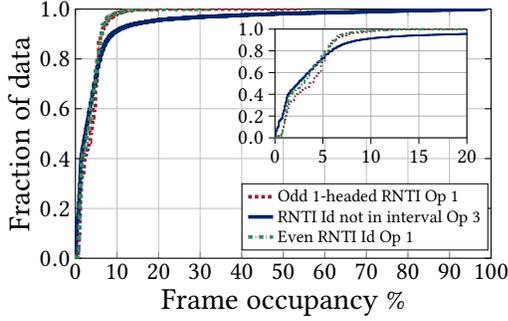
Figure 8: Inter-transmission times and traffic load for different RNTI allocation schemes

at odd 0-headed RNTI IDs and to the interval assigned ones, see Fig. 9(b), it emerges that a higher frame occupancy is achieved by both with respect to previous cases. From this result it emerges that when the random and round robin allocation schemes are used by the BS, a higher number of PRB resources is consumed by UEs with odd RNTI. The same result holds for Operator 3. In this case, we register a higher frame occupancy for the interval-allocated RNTI with respect to the ones that are not interval-allocated.

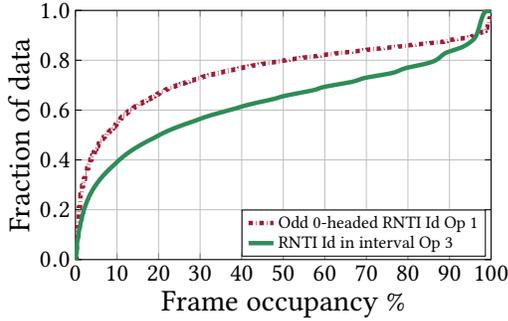
Takeaway message. Our analysis shows that specific ranges of RNTI IDs are reserved for different allocation strategies, *i.e.*, random, round robin, or interval-based. In addition, round robin and interval-based allocations are only observed during high-load time periods. Finally, users with round robin or interval-based RNTI allocation exhibit higher frame occupancy than user with RNTIs that were randomly assigned.

4.4 Threshold Tuning

We now perform an in-depth study on the micro-level behavior of RNTI allocation. This analysis reveals that the threshold T can be tuned further. By analyzing our traces, we notice that some RNTIs are first allocated in round robin fashion and then obtain subsequent allocations over random periods. We define these as spurious cases and Fig. 10(a) exemplifies this behavior for three RNTI IDs of BS 4 of Operator 2. The frequency through which these allocations occur



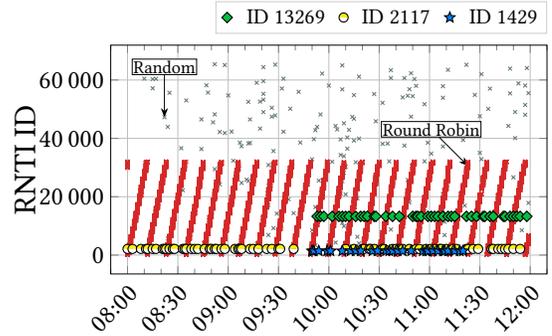
(a) RNTI assigned by random allocation only. Inset over 0-20% interval.



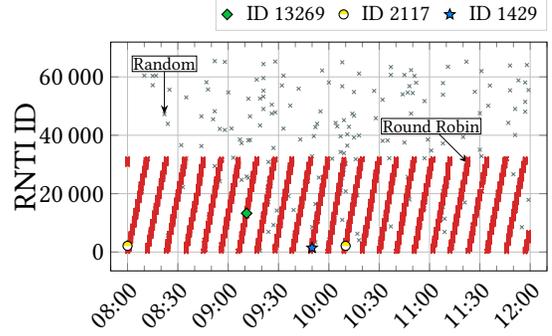
(b) Round robin and interval

Figure 9: Pairwise comparison of (a) random and (b) round robin vs interval for operators 1 and 3 with different RNTI allocation combinations

indicate that such transmissions actually belong to the same UE (since the probability of this occurring frequently at random is vanishingly small). Thus, the RNTI is preserved for an inactivity period which is larger than the one that we set in our first analysis, *i.e.*, 30 s or 60 s according to the monitored BS. We refine the expiration threshold per BS by removing these random spurious RNTI allocations (see Fig. 10(b)). For this set of RNTIs, a real allocation to a new incoming user happens when the next round robin period occurs. For each spurious RNTI from the set, we evaluate the minimum threshold T' that allows to filter out all random RNTI allocations and preserve only the round robin ones. Finally, among all the T' obtained for each spurious RNTI, we take the minimum one. This analysis leads to $T' = 9$ min. **Takeaway message.** Whenever both random and round robin allocation schemes are used by the BS, it is possible to further refine T . We identify spurious RNTIs that are first allocated through round robin and then are reallocated by the random scheme. By eliminating random reallocations and preserving real round robin ones, we observe that the threshold can be extended up to tens of minutes.



(a) RNTI expiration threshold of 60 s. Three spurious cases are highlighted, *i.e.*, RNTI IDs 2117, 1429, 13269.



(b) RNTI threshold of 9 min. Spurious random allocations disappear while round robin allocations are preserved.

Figure 10: Threshold tuning procedure

5 EVALUATION

We now investigate the impact of the incorrect setting of the expiration threshold on user characteristics.

We present the results for Operator 2 BS 1 over a period of 24 hours. Specifically, we analyze how the number of extracted users and per-user load change when considering T_R and T_L . We denote as T_R our refined threshold, *i.e.*, 9 min (see Sec. 4.4) and as T_L the 10 s threshold used in the literature (see Sec. 2). Using T_L underestimates the user life time, *i.e.*, it assumes the RNTI was released while the user is in fact still active with the same RNTI, and the remaining time assigned to a new user that in reality does not exist. This has two main implications i) the loss of data for the underestimated user and ii) the shift of such data to a fake created user. In fact, we observe that the number of extracted users is significantly higher for T_L than for T_R , *i.e.*, 107982 vs. 18170 respectively.

Fig. 11(a) presents over 1 min bins the number of occurrences of users duration extracted through T_L and T_R . The figure shows that for both thresholds the highest number of occurrences is registered for users whose duration is smaller than 1 min. In the considered dataset, 40% of users have an estimated life time smaller than 1 min, and this number grows to 98% when T_L is considered. The figure also shows

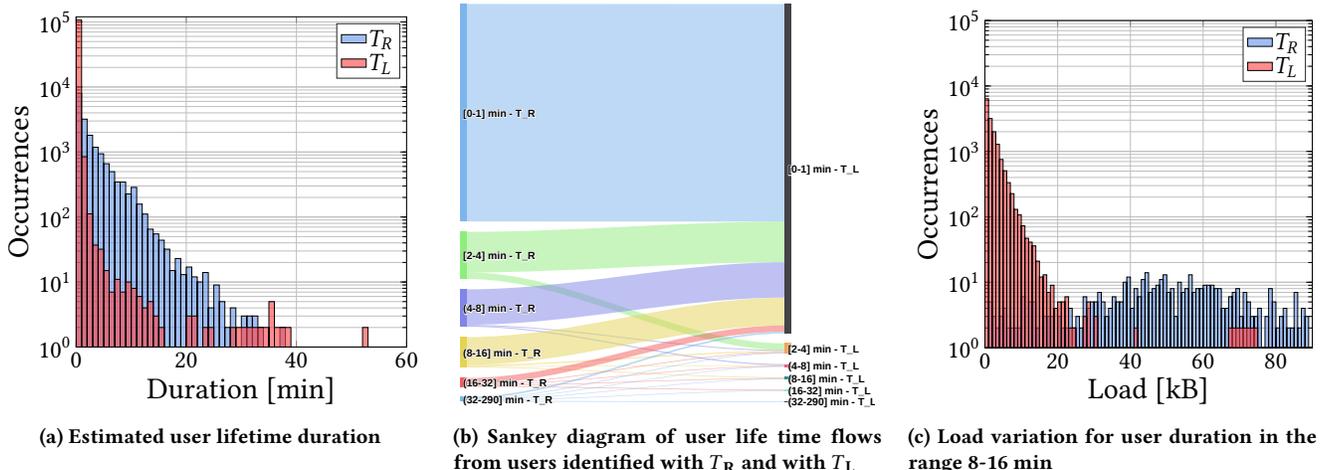


Figure 11: Implication of setting correct (T_R) and incorrect (T_L) RNTI expiration thresholds

that as soon as the larger user durations when using T_R are considered, the corresponding number of occurrences in bin (0-1) min drops by two orders of magnitude, since through T_L many fake users are created in precisely that range.

The Sankey diagram in Fig. 11(b) illustrates the user life time reallocation, where the width of the edge flow is proportional to the number of occurrences that are shifted from larger to smaller intervals when using T_L instead of the correct value T_R . We can observe that all users' life time that ranges from 2 min up to the maximum are mostly allocated to non-existent users with a duration of (0-1) min. Instead, only a small fraction of users preserves the same life time duration from T_R to T_L .

An incorrect choice of threshold has also an impact on the considered per-user load. Fig. 11(c) considers user durations that fall within the bin (8-16) min and analyzes their characteristics in terms of per user-load. The average load of T_R extracted users is significantly higher than T_L extracted ones, *i.e.*, 830 kB vs. 32 kB. The figure shows a zoom over the interval 0 - 120 kB. We observe that when T_L is considered, 99% of the extracted users have a load which is smaller than 20 kB. Instead, the actual load when T_R is used as expiration threshold shows higher values ranging from 25 kB to 120 kB with a median value of 60 kB.

Takeaway message. An incorrect setting of the expiration threshold affects the number of extracted users and per-user load. When T_L instead of T_R is used, non-existent users are created, that have a smaller load value than the T_R extracted ones. These errors can have substantial impact on the subsequent trace analyses, and may bias the conclusions.

6 DISCUSSION

Our analysis can serve a wide range of applications in resource management, traffic forecasting and security.

Resource allocation. *The resource allocation process at the BS side can benefit from the correct identification of the user life time over the network. Possible applications range from scheduling optimization for video streaming to allocation of pilot symbols.* For example, in [28], the number of active C-RNTIs is used to identify traffic variations and thus to detect urban anomalies. By correctly identifying the number of active UEs at a certain time period it is possible to discriminate normal and anomalous network behavior. Furthermore, by leveraging ms-level resource allocation patterns extracted from LTE, PERCEIVE in [18], aims at uplink throughput predictions to improve video streaming experience. TRADER [11] defines an efficient policy for aperiodic Sounding Reference Signal scheduling, leveraging per user traffic forecasts at ms level.

Traffic forecasting. *The characterization of RNTI IDs in Sec. 4.3 can improve the predictive analysis of cellular networks.* For example in [27], the authors use the traffic information derived from the DCI to predict mobile traffic up to the next 150 ms. A neural network can benefit from the information about the adopted allocation scheme, *e.g.*, a change from one allocation scheme to another can be tied to a load variation (see Fig. 8a) over time in order to optimize the predictions. Furthermore, in [26], the RNTI ID is used to track the data flow of a certain user and then to learn a classifier on the applications executed at the connected mobile terminals. In particular, for some applications, *e.g.*, video streaming, *a higher frame occupancy is desirable*, and thus such type of traffic should be linked to certain RNTI allocation schemes, *i.e.*, round robin or interval allocation.

Security. *The identification of the user life time allows to determine the user permanence within a cell and thus to expose the user to privacy issues.* In fact, several works show that the RNTI can be mapped to the user end terminal, exposing an adversary to track and localize users within a cell [17, 25].

Furthermore, the information on the websites visited and types of application ran by the user can be extracted by decoding the RNTI from the DCI, allowing an adversary to run passive and active fingerprinting attacks [16, 20].

7 CONCLUSIONS

In this paper we provided an extensive study of the RNTI allocation mechanism by collecting and analyzing a real LTE traffic dataset from different BSs and operators. We proposed a complete methodology for setting the RNTI expiration threshold for trace analysis. Furthermore, according to the allocation scheme adopted by the BS, we characterized RNTI IDs in terms of frame occupancy and traffic load. This is an important step towards improving analysis and use of traffic traces gathered from operational mobile networks.

Our results reveal that the operators combine different allocation strategies (*e.g.*, random, round robin and interval) and which strategy is used varies according to the traffic load over night/day periods.

ACKNOWLEDGMENTS

This work is partially supported by a Juan de la Cierva grant from the Spanish Ministry of Science and Innovation (IJC2019-039885-I), by the Atracción de Talento Investigador grant number 2019-T1/TIC-16037 NetSense, funded by the Comunidad de Madrid, and by the Madrid Regional Government through the TAPIR-CM program (S2018/TCS-4496).

REFERENCES

- [1] 3GPP TS 36.321. 2017. *LTE Medium Access Control (MAC) protocol specification*. Technical Report 36.321. Version 14.4.0.
- [2] 3GPP TS 38.314. 2009. *LTE Evolved Universal Terrestrial Radio Access Network (E-UTRAN)*; Technical Report. Version 8.1.0.
- [3] 3GPP TS 38.321. 2019. *5G NR Medium Access Control (MAC) protocol specification*. Technical Report 38.321. Version 15.6.0.
- [4] Sassan Ahmadi. 2010. *Mobile WiMAX: A systems approach to understanding IEEE 802.16 m radio access technology*. Academic Press.
- [5] I. Alawe, A. Ksentini, Y. Hadjadj-Aoul, and P. Bertin. 2018. Improving Traffic Forecasting for 5G Core Network Scalability: A Machine Learning Approach. *IEEE Network* 32, 6 (Nov 2018), 42–49. <https://doi.org/10.1109/MNET.2018.1800104>
- [6] D. Bega, M. Gramaglia, M. Fiore, A. Banchs, and X. Costa-Pérez. 2020. DeepCog: Optimizing Resource Provisioning in Network Slicing With AI-Based Capacity Forecasting. *IEEE JSAC* 38, 2 (2020), 361–376.
- [7] N. Bui and J. Widmer. 2016. OWL: A Reliable Online Watcher for LTE Control Channel Measurements. In *Proc. of ACM Workshop on All Things Cellular: Operations, Applications and Challenges*. 25–30. <https://doi.org/10.1145/2980055.2980057>
- [8] X. Chen, W. Yang, C. Xu, and Y. Kim. 2012. RNTI Allocation Schemes for User Equipments in LTE System. In *Proc. of IEEE WiCOM*. 1–4.
- [9] Ericsson. 2020. Ericsson Mobility Report. <https://www.ericsson.com/4adc87/assets/local/mobility-report/documents/2020/november-2020-ericsson-mobility-report.pdf>.
- [10] R. Falkenberg and C. Wietfeld. 2019. FALCON: An accurate real-time monitor for client-based mobile network data analytics. In *Proc. of IEEE GLOBECOM*. 1–7.
- [11] C. Fiandrino, G. Attanasio, M. Fiore, and J. Widmer. 2021. Traffic-Driven Sounding Reference Signal Resource Allocation in (Beyond) 5G Networks. In *Proc. of IEEE SECON*. 1–9. <https://doi.org/10.1109/SECON52354.2021.9491611>
- [12] C. Fiandrino, C. Zhang, P. Patras, A. Banchs, and J. Widmer. 2020. A Machine Learning-based Framework for Optimizing the Operation of Future Networks. *IEEE Communications Magazine* 58, 6 (2020), 20–25. <https://doi.org/10.1109/MCOM.001.1900601>
- [13] C. Gijón, M. Toril, S. Luna-Ramírez, and M. Luisa Mari-Altozano. 2019. A Data-Driven Traffic Steering Algorithm for Optimizing User Experience in Multi-Tier LTE Networks. *IEEE Transactions on Vehicular Technology* 68, 10 (2019), 9414–9424. <https://doi.org/10.1109/TVT.2019.2933068>
- [14] S. Hailu and M. Säily. 2017. Hybrid paging and location tracking scheme for inactive 5G UEs. In *Proc. of IEEE EuCNC*. 1–6. <https://doi.org/10.1109/EuCNC.2017.7980730>
- [15] J. Huang, F. Qian, A. Gerber, Z. Morley Mao, S. Sen, and O. Spatscheck. 2012. A close examination of performance and power characteristics of 4G LTE networks. In *Proc. of ACM MobiSys*. 225–238.
- [16] K. Kohls, D. Rupperecht, T. Holz, and C. Pöpper. 2019. Lost traffic encryption: fingerprinting LTE/4G traffic on layer two. In *Proc. of ACM WiSec*. 249–260.
- [17] S. Kumar, E. Hamed, D. Katabi, and L. Erran Li. 2014. LTE radio analytics made easy and accessible. *Proc. of ACM SIGCOMM* 44, 4 (2014), 211–222.
- [18] J. Lee, S. Lee, J. Lee, S. D. Sathyanarayana, et al. 2020. PERCEIVE: deep learning-based cellular uplink prediction using real-time scheduling patterns. In *Proc. of ACM MobiSys*. 377–390.
- [19] F. Malandrino and C.-F. Chiasserini. 2019. 5G Traffic Forecasting: If Verticals and Mobile Operators Cooperate. In *Proc. of IEEE WONS*. 79–82.
- [20] F. Meneghello, M. Rossi, and N. Bui. 2020. Smartphone identification via passive traffic fingerprinting: A sequence-to-sequence learning approach. *IEEE Network* 34, 2 (2020), 112–120.
- [21] R. Piqueras Jover. 2016. LTE security, protocol exploits and location tracking experimentation with low-cost software radio. *arXiv:1607.05171* (2016).
- [22] M. Polese, F. Restuccia, and T. Melodia. 2021. DeepBeam: Deep Waveform Learning for Coordination-Free Beam Management in mmWave Networks. *Proc. of ACM MobiHoc* (2021).
- [23] C. Ramesh, A. Jafar, and G. Prasad. 2016. An effective RNTI allocation scheme for user equipment in LTE systems. In *Proc. of WiSPNET*. 1737–1740. <https://doi.org/10.1109/WiSPNET.2016.7566436>
- [24] M. T. Raza, D. Kim, K.-H. Kim, S. Lu, and M. Gerla. 2017. Rethinking LTE network functions virtualization. In *Proc. of IEEE ICNP*. 1–10.
- [25] D. Rupperecht, K. Kohls, T. Holz, and C. Pöpper. 2019. Breaking LTE on layer two. In *Proc. of IEEE SP*. 1121–1136.
- [26] H. D. Trinh, Á. Fernández G., L. Giupponi, M. Rossi, and P. Dini. 2021. Mobile Traffic Classification Through Physical Control Channel Fingerprinting: A Deep Learning Approach. *IEEE Transactions on Network and Service Management* 18, 2 (2021), 1946–1961. <https://doi.org/10.1109/TNSM.2020.3028197>
- [27] H. D. Trinh, L. Giupponi, and P. Dini. 2018. Mobile Traffic Prediction from Raw Data Using LSTM Networks. In *Proc. of IEEE PIMRC*. 1827–1832. <https://doi.org/10.1109/PIMRC.2018.8581000>
- [28] H. D. Trinh, L. Giupponi, and P. Dini. 2019. Urban anomaly detection by processing mobile traffic traces with LSTM neural networks. In *Proc. of IEEE SECON*. 1–8.
- [29] H. Zhang, L. Qiu, X. and Meng, and X. Zhang. 2010. Design of distributed and autonomic load balancing for self-organization LTE. In *Proc. of IEEE VTC Fall*. 1–5.