Characterizing and Modeling Mobile Networks User Traffic at Millisecond Level

Pablo Fernández Pérez
pablo.fernandez@imdea.org
IMDEA Networks Institute
Madrid, Spain

Claudio Fiandrino
claudio.fiandrino@imdea.org
IMDEA Networks Institute
Madrid, Spain

Joerg Widmer
joerg.widmer@imdea.org
IMDEA Networks Institute
Madrid, Spain

ABSTRACT

The availability of datasets has been instrumental to drive advances in several disciplines like computer vision, image processing, and natural language processing. However, in the context of mobile traffic, data is often not available because of diverse reasons including data sensitivity, legal considerations and business competition. The lack of dataset availability restrains the research advance at large.

In this paper, we make a twofold contribution. On the one hand, we make available a large dataset of mobile traffic from multiple Base Stations (BSs). The key distinct feature of the dataset is in the nature of the data, which is based on real LTE traffic information decoded from control channel information at the millisecond level. On the other hand, we carry out an in-depth characterization of user traffic and study how widely adopted probability distributions for mobile traffic do apply at short-term scales. Our analysis shows that mobile data traffic exhibits self-similarity and the number of Radio Resource Control (RRC) connected users exhibits a bi-modal distribution. Overall, our contribution key to verify and reproduce research outcomes as well as driving advances of Artificial Intelligence (AI)/Machine Learning (ML) applied to mobile networks.

CCS CONCEPTS

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

WiNTECH'23, October 6, 2023, Madrid, Spain

ACM ISBN 978-1-4503-9990-6/23/10…$15.00
https://doi.org/10.1145/3570361.3613190

KEYWORDS

Mobile networks, network measurements, mobile traffic characterization, datasets.

1 INTRODUCTION

Data analysis has gained momentum in several disciplines over the last few years. The enormous improvements in analytics and predictive modeling coupled with the availability of large datasets have enabled common benchmarking and reproducibility of research outcomes, sparking the growth and advancements in the areas of computer vision, image and natural language processing [7, 18]. Such research advances usually go hand in hand with advances of Artificial Intelligence (AI)/Machine Learning (ML) techniques applied to such areas.

Surprisingly, the landscape of dataset availability in the context of mobile traffic is *fragmented* and *scarce*. Scarcity is mainly attributed to the lack of interest of mobile network operators in publicly releasing such data because of privacy concerns for data sensitivity, legal and regulatory aspects, and to retain a competitive advantage over other operators. Despite the overall scarcity, there exist remarkable initiatives making available mobile traffic data at metropolitan scale at minute-level granularity (e.g., the Telecom Italia Big Data Challenge [4] and NetMob [27]). Reference public datasets providing access to sub-minute scale data are even rare, making imbalanced these two fully independent classes. Indeed, the former class of datasets is useful for optimizing network deployment planning [8] and routing [13], to infer human and economic activities [42],and to better handle crowded events [32]. In contrast, the latter class of datasets is useful for optimizing resource allocation [6], channel sounding [10], congestion control over mobile networks [40] or understanding specific mechanisms like network ID assignment to users [2].

In this paper, we make a two-fold contribution. First, we make publicly available a dataset with millisecond-level information of real LTE Base Stations (BSs) collected over the course of two years in both Madrid city center and sub-metropolitan areas. Second, we carry out an in-depth characterization of traffic at BS- and user-level in terms of traffic load, and inter-transmission times. For this, we study how

widely adopted probability distributions for mobile traffic do apply at short-term scales. Overall, such knowledge is extremely useful to researchers for simulation and experimental (e.g., in large wireless testbeds like Colosseum [28]) evaluation settings. Therefore, our work is positioned to support new research methods and systems in the area of resource allocation for mobile networks enabling result reproducibility.

The expert reader may argue why to focus on releasing a LTE dataset in the 5G era. There are several reasons. First, despite its commercial roll-out dates back to 2019, the vast majority of today's 5G network deployments are Non-StandAlone (NSA) [12, 23, 33], *i.e.*, a 5G NR BS is added as secondary radio access to the existing LTE master BS. Second, technological advances like the introduction of 5G and the consequent introduction of new services and applications, shift in user behaviors and infrastructure improvements are processes that ultimately alter mobile traffic patterns, but at a slow pace. Hence, our dataset given its unique features will still be of help in the above-identified research areas. Second, from a technical standpoint, it is not trivial to create a tool capable of decoding information at such fine-grained scaled and level of detail for 5G. We will elaborate more on this regard in Section 5.

The key insights (denoted as "I") resulting from our analysis can be summarized as follows:

*I*1: We confirm with a recent study an old finding that dates back to the year 2004 [17]: data-plane traffic exhibits self-similarity properties at both individual uplink and downlink components and as a whole. This is different from control-plane mobile network traffic [30]

*I*2: We find that the number of Radio Resource Control (RRC) connected users follows a bi-modal distribution that indicates the presence of circadian cycles resulting in two clusters of different sizes, *i.e.*, users connected during the day and users connected during the night.

We publicly release the anonymized dataset that we use in the analysis at https://git2.networks.imdea.org/wng/madrid-lte-dataset.

2 DATA COLLECTION METHODOLOGY

For our study, we collect a dataset of LTE traffic allocations from multiple BSs located in different areas of Madrid, Spain. We first delve into the background of the tools that can provide such data and next, we describe the characteristics of data collection and processing. Finally, Section 3 will delve into the presentation of the final form of the dataset.

The Tools. Unlike 5G NR [12], in LTE the control channel spans the entire bandwidth and the decoding of the physical downlink control channel (i.e., the PDCCH) contains per-user scheduling information. LTEye [20] is a seminal attempt at performing PDCCH decoding, OWL [5], FALCON [9], and LTESniffer [14] are refinements over the basic mechanisms that improve the decoding performance. NG-SCOPE [39] makes it possible the decoding at multiple BSs at a time.

Our Dataset. For our study, we have used FALCON as LTE passive monitoring tool. We run FALCON on a Linux laptop connected to a USRP B210 to decode the unencrypted information that BSs send to the UEs over the PDCCH channel regarding their traffic allocation at Transmission Time Interval (TTI)-level. During this process, the user's identity is not revealed. Rather, we gather the following information: *(i)* the temporary user ID currently associated with the user (the Radio Network Temporary Identifier (RNTI)), *(ii)* the ID of the frame containing the traffic allocation for each C-RNTI, *(iii)* the associated transport block size (TBS), and *(iv)* information related to the transmission like Modulation and Coding Scheme (MCS) and utilized Physical Resource Block (PRB). This information makes it possible to determine traffic characteristics at the user-level, *i.e.*, the size and duration of per-user traffic bursts and idle times between subsequent traffic bursts. From the collected datawe discard RNTIs reserved for random access (RA-RNTIs with ID 1-960), paging and system notification (P-RNTIs with ID 65534), and broadcast system information (SI-RNTIs with ID 65535).

Being the RNTIs temporary, to identify each user, we analyzed RNTIs like in [2]. The problem with RNTIs being temporal lies in the fact that different users can be identified with the same RNTI because the space from where they are drawn is limited and they are reused over time. To overcome this problem it is necessary to analyze the time elapsed (inter-transmission times) between two consecutive samples with identical RNTI and set a time threshold above which there is an end of the user lifetime. This threshold is set by observing the cumulative distribution functions of inter-transmission times and by selecting a value that is greater than 10 s (common threshold from existing literature) and is before a significant drop in the cumulative distribution function. With this analysis, it is possible to select a threshold that is more realistic for user identification than the commonly used 10 s threshold [2]. After running an extensive analysis on each BS of the whole dataset, we find that common thresholds lie in between the range $[10, 40]$ seconds. For the sake of fair comparison, in this paper, we set the threshold to 10 s.

3 THE FINAL DATASET

The data, collected according to the methodology explained in Section 2 is divided into multiple files one for each different zone, frequency, target, and granularity. Table 1 summarizes the key aspects of the dataset. Depending on the area where the data was collected, we group the data into

Table 1: The dataset in a nutshell

Year	Operator	Zone	Carrier Frequency (MHz)	ID	Total Time (h)	Total Users	Traffic Volume (GB)	Dataset Size (GB)
2020	A	I	796	BS1	247	72 321 636	388 691	11
2020	B	I	1815	BS2	227	23 508 594	242 406	3
2020	A	I	2650	BS3	74	10 493 540	18 403	0.24
2021	B	II	816	BS4	163	14 813 731	579 376	12
2021	C	II	1835	BS5	313	59 463 421	2 378 256	27
2021	A	II	2650	BS6	353	33 650 085	1 308 064	11

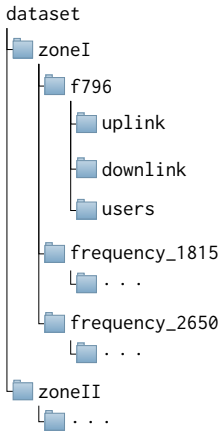

Figure 1: Tree structure of the dataset

two zones, namely Zone I and Zone II. Zone I corresponds to BSs located within the diameter of the orbital highway M-30 circling all the central districts of Madrid municipality. In contrast, zone II contains data about BSs located outside the orbital highway, *i.e.*, in sub-urban areas. Zone I features BSs with different carriers frequencies: 796 MHz (10 MHz channel bandwidth), 1815 MHz (20 MHz channel bandwidth), and 2650 MHz (20 MHz channel bandwidth). The same applies for zone II, whose BSs feature the following carriers frequencies: 816 MHz (10 MHz channel bandwidth), 1835 MHz (20 MHz channel bandwidth), and 2650 MHz (20 MHz channel bandwidth). BSs with carrier frequencies 796 MHz and 2650 MHz belong to operator A, carrier frequencies 796 MHz and 2650 MHz to operator B, and finally carrier frequencies 796 MHz to operator C. Note that A, B, and C are the three major operators in Spain.

The dataset records information at the millisecond level about: timestamp (in unix format), traffic direction (uplink or downlink), RNTI, and Tranport Block Size (TBS). Such fine-grained information makes it possible to estimate, among others, the inter-transmission times and traffic load for each user, and the number of users connected at BS level. The above definitions ground on the fact that we consider RRC connected users, *i.e.*, users that are capable of transmitting and receiving traffic from the BS. Note that being connected at RRC level does not necessarily mean that the users transmit or receive traffic in a given TTI. Figure 1 shows the high-level description of the structured set of folders in which the dataset is organized. These folders contain the cleaned and granularity-reduced data.

4 ANALYSIS AND INSIGHTS

In this Section, we provide the complete analysis of our dataset that encompasses the study of mobile traffic when observed at BS level (Section 4.1), and at the user level (Section 4.2). Finally, we study how the number of RRC connected users varies over time (Section 4.3).

4.1 Traffic at BS Level

We now characterize the main features of the mobile traffic at BS level by analyzing the total amount of traffic observed and the breakdown into downlink and uplink traffic.

Data Filtering. Before delving into the discussion, we make the following observation. Figure 2 shows Cumulative Distribution Functions (CDFs) that indicate the fraction of traffic consumed by a given fraction of users. We observe a consistent pattern across the BSs except for BS3. Specifically, in 5 out of 6 BSs, less than 10% of the users transmit 95% of the traffic. As for BS3, this fraction changes and 40% of the users transmits 95% of the traffic. Based on such observation, in the remainder of the paper we focus on the fraction of users that transmit most of the traffic. Note that at the time of performing trace analysis of data collected from LTE DCI decoding tools, it is highly recommended to perform such preliminary data cleaning.

Measuring Traffic Burstiness. An important concept often overlooked when analyzing mobile network traffic traces is self-similarity. Self-similarity is a phenomenon first noticed by the hydrologist Harold Edwin Hurst [16] and later mathematically defined by Benoit B. Mandelbrot [26], and consists in finding structural similarities across all different scales in the data. In mobile traffic modeling, the degree of self-similarity can be used to measure the burstiness of the data [21]. Measuring burstiness is important: bursty traffic strains network resources such as bandwidth, which is often limited in wireless systems pre-mmwave era, and poses a burden to designing effective resource allocation schemes.

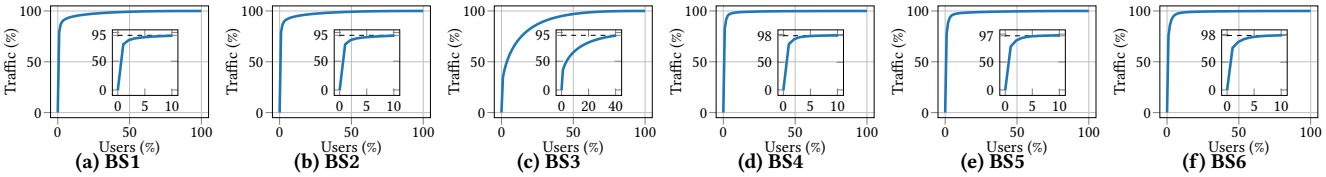

(a) BS1 (b) BS2 (c) BS3 (d) BS4 (e) BS5 (f) BS6

Figure 2: The amount of traffic consumed by users

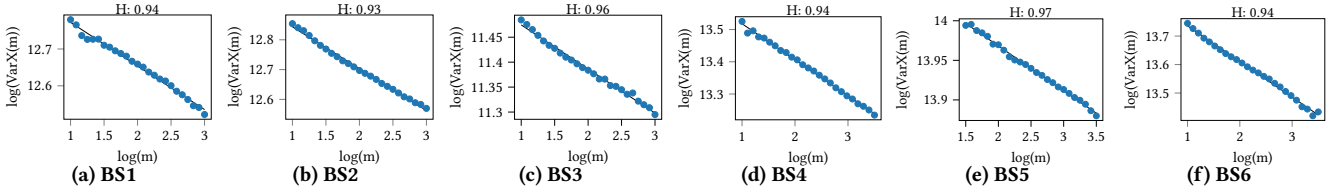

(a) BS1 (b) BS2 (c) BS3 (d) BS4 (e) BS5 (f) BS6

Figure 3: Time-variance plots with the Hurst parameter computed H for combined uplink and downlink traffic load for six different BSs

The Hurst parameter H is commonly used to measure the degree of self-similarity. This parameter is estimated by plotting time-variance plots and by calculating the slope β of the slowing decay variance along different aggregation group sizes m [41]. H is then calculated using the expression: $H = 1 - \beta/2$. A Hurst parameter greater than 0.5 and close to 1 implies that the sample variance will asymptotically decrease at a slower rate than the sample size (long-range dependence - positive correlation, persistent behavior), unlike what happens in stochastic processes like traditional Poisson modeling process where H tends to values around 0.5 and the sample variance decreases exponentially concerning the sample size (no time range dependence - no correlation, randomness). Finally, values of H lower than 0.5 and close to 0 indicate that the process alternates between high and low values (short range dependence - negative correlation, anti-persistent behavior).

This poses a specific problem in traditional models because they systematically fail to reflect this statistical property which is commonly known as the Hurst phenomenon. Asserting self-similarity in mobile traffic networks means that for modeling, researchers should take into account a number of self-similar stochastic models to better reflect the reality of network traffic.(e.g., [34]). To help facilitate traffic modeling at BS level, we report the H parameter for the different BSs in our dataset by aggregating traffic at the second level. Figure 3 shows time-variance plots computed on our dataset which makes it possible to extract the Hurst parameter H. We observe that the degree of self-similarity is very high and close to 1. This means that mobile network traffic is statistically similar for time windows of different sizes as opposed to stochastic processes and therefore the burstiness is similar at different time scales. Practically, this means that modeling traffic with such a function over a horizon of 30 seconds or 10 minutes would be statistically correct.

Table 2: Degree of self-similarity of uplink and downlink traffic

BS ID	Uplink (Hurst)	Downlink (Hurst)
BS1	0.92	0.94
BS2	0.81	0.93
BS3	0.92	0.95
BS4	0.87	0.94
BS5	0.78	0.97
BS6	0.89	0.95

We now ask ourselves if the self-similarity property also applies to the individual traffic components, uplink, and downlink. Table 2 confirms the property and shows that it is more evident for downlink traffic than for the uplink. We attribute this behavior to upper-layer protocols like TCP that introduce mechanisms like congestion control that lead to higher burstiness than the downlink counterpart.

4.2 Traffic at User Level

If observed with the microscopic lens at millisecond resolution, mobile traffic is highly sparse. Figure 4 illustrates the sparsity of traffic for a 10 seconds traffic snapshot for BS2. At the top of Figure 4 we show the PRB allocation (normalized to 100, the number of PRB with 20 MHz channel bandwidth) for each active UE at the TTI level (1 ms). Note that BSs prefer to assign all resources to a single user in one TTI, which explains the yellow dots [3]. At the bottom of Figure 4, we show the corresponding load of the BS. Figure 5 shows a 5 s snapshot of traffic generated by one specific UE. We observe that traffic is highly sporadic and is therefore well described

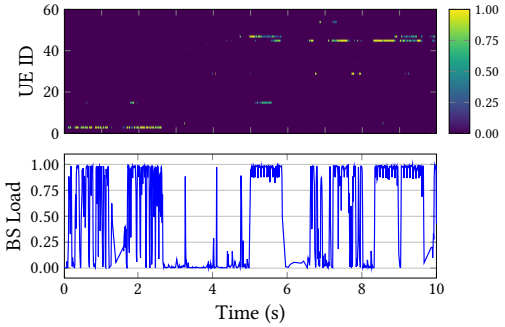

Figure 4: Mobile traffic sparsity during 10 s of a 20 MHz channel BS.

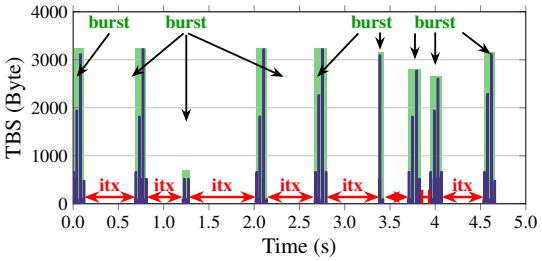

Figure 5: Snapshot of 5 s of traffic generated by a UE. The example illustrates how per-user traffic is sporadic and irregular. Markers highlight bursts and inter-transmission times ("itx" in Figure).

through two components. The *bursts*, measured in Bytes, define the amount of traffic computed summing all TBS of subsequent data transmissions. The *inter-transmission times*, measured in milliseconds, define the time between two consecutive bursts with no data transmission in the uplink or downlink direction.

Having set the stage by defining *bursts* and *inter-transmission*, we now delve into the characteristics of each component separately in Section 4.2.1 and Section 4.2.2 respectively.

4.2.1 Traffic Load. Figure 6 shows the overall distribution of traffic bursts for all the users of three different BS. Specifically, we choose to present one BS per operator to better highlight the difference and similarities we find with this analysis. More concretely, we plot on the axis-x the burst defined in terms of TBS sum (in other words, all the traffic destined to a specific RNTI in a 10 ms frame), and on the axis-y, the number of occurrences a given burst is found in the user lifetime (in other words, the duration of the burst). Finally, we count the number of users that have been seen transmitting a given burst for so many occurrences and render this information with colors. The resulting plot provides therefore a unique fingerprint of users' traffic activities in the corresponding BS. From the Figure, we can observe that very few users transmit long bursts and when this occurs, these are bursts that carry little amount of traffic. The vast majority of the users (lighter colors in the bottom part of the Figure) are usually associated with shortly-lived bursts regardless of the burst size, which may be the consequence of the mix of highly diverse applications utilized. The differences between the BSs mainly lie in the areas of the dark clouds on the top part of Figure 6: we attribute this difference to the joint effect of the traffic nature itself and to how each operator performs resource allocation.

We now break down the above analysis for uplink and downlink traffic and show the result in Figure 7 for the BS4 solely. As we expect, the uplink traffic is lighter and this becomes clear in Figure 7(b) as the maximum burst size is significantly lower than the downlink counterpart. Also, note that the long bursts that carry a little amount of traffic are mainly generated by a few users transmitting in the uplink direction (see the top left part of the figure).

4.2.2 Inter-transmission Times. For inter-transmission times, we perform a similar analysis to traffic load. Specifically, we show in Figure 8 the inter-transmission times of two different BSs. Specifically, we plot on the axis-x the inter-transmission times, and on the axis-y the number of occurrences a given inter-transmission time is found in the user lifetime. Finally, we render with colors the number of users that have been seen idle for a certain duration for so many occurrences. As expected, we do not observe any inter-transmission time higher than 10 seconds, that is the value of the threshold set for RNTI expiration in this work. From the figure, we can observe that the vast majority of the cases (light colors) appear in the bottom left part of the figure, which indicates that the vast majority of the inter-transmission times are very short.

4.3 Estimated Number of Connected Users

With the methodology described in Section 2, we can identify the number of users that are active at RRC layer. As we will better describe later in Section 5, such information is particularly relevant to analyze the root cause of signaling storms in core networks. We show in Figure 9 a 5 minute snapshot of the number of RRC connected users for BS4 when looking at different rolling averages. Increasing the rolling average duration smooths out noise and allows to identification of patterns that may be exploited by ML.

By analyzing the distribution of RRC connected users, we find that it can be described as a bimodal Gaussian distribution. This distribution is the sum of two Gaussian distributions and can be mathematically expressed as:

$$\sum_{i=1}^{2} A_i e^{-\frac{1}{2}(\frac{x-\mu_i}{\sigma_i})^2}, \quad (1)$$

where A is the height of the curve's peak, μ is the mean and σ is the standard deviation.

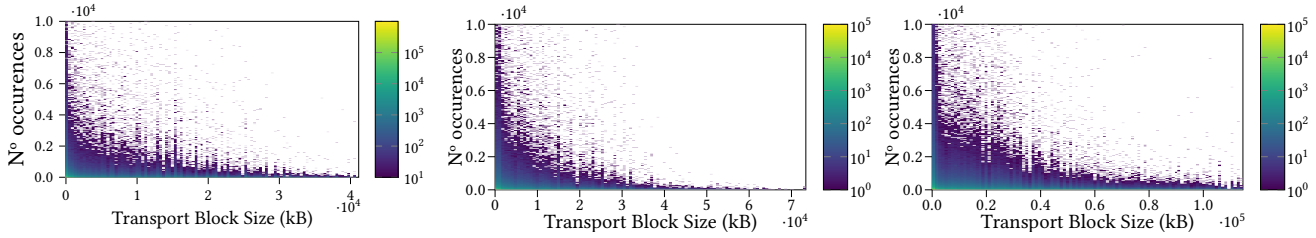

(a) BS1 (b) BS4 (c) BS5

Figure 6: Distribution of user-level traffic across different BSs

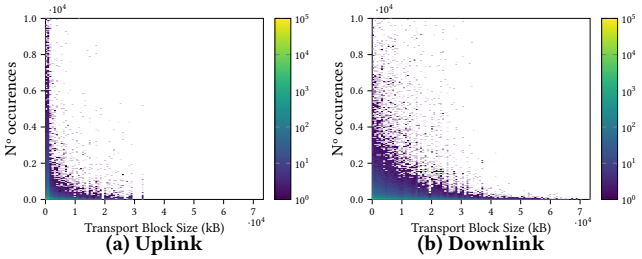

(a) Uplink (b) Downlink

Figure 7: Breakdown of uplink and downlink user-level traffic for BS4

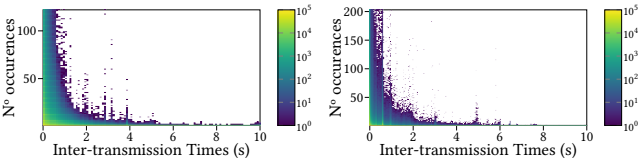

(a) BS1 (b) BS5

Figure 8: Inter-transmission times for BS4

We apply curve fitting to the BSs of our dataset. Curve fitting is computed as the least square error linear regression of the parameters appearing in (1) to fit the histogram of the RRC-connected users computed on the raw time series. To increase convergence, it is usually recommended to set the initial parameters close to the final parameters and this can be done by visually inspecting the histogram bimodal shape (e.g., μ_i and A_i are the x-axis point and the y-axis point respectively at the peak of the Gaussian bell i). Table 3 reports all the parameters for fitting the distributions of RRC connected users.

Figure 10 shows that a bimodal distribution for BS4 (on the left of the Figure) and BS6 (on the right) can be observed. Note that this particular distribution can be due to the Circadian cycles that result in two clusters with different numbers of users during the day and at night. Such information can be leveraged by a ML solution that forecasts the future number of users in the network.

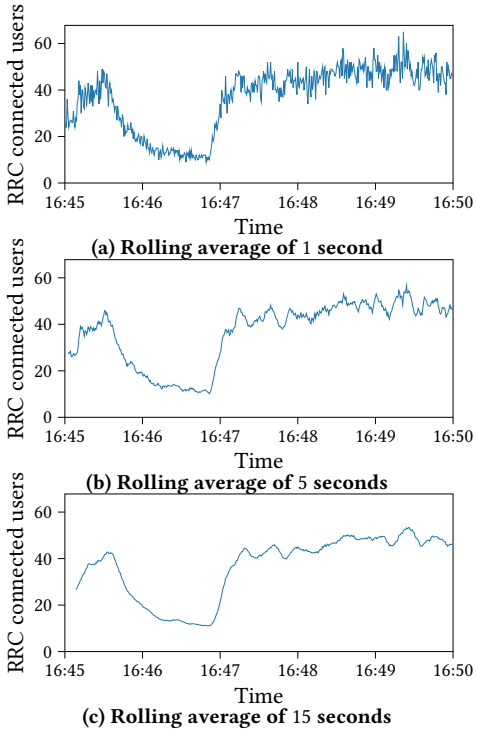

(a) Rolling average of 1 second

(b) Rolling average of 5 seconds

(c) Rolling average of 15 seconds

Figure 9: A 5 minute snapshot of BS4

Table 3: Bimodal RRC users distribution parameters

BS ID	μ_1	σ_1	A_1	μ_2	σ_2	A_2
BS1	2.88	2.37	3008.43	59.51	43.31	7920.10
BS2	7.93	1.88	6732.34	24.50	17.05	19129.48
BS3	5	0.05	2000.00	39.53	12.22	8067.96
BS4	9.41	5.28	17659.29	35.05	10.22	14117.49
BS5	49.99	12.85	31621.99	108.81	17.70	1823.15
BS6	16.45	4.27	29527.18	29.86	9.36	40123.96

5 DISCUSSION

In this Section, we elaborate on potential applications of our datasets to diverse research areas and briefly discuss challenges to overcome for a possible extension to 5G NR.

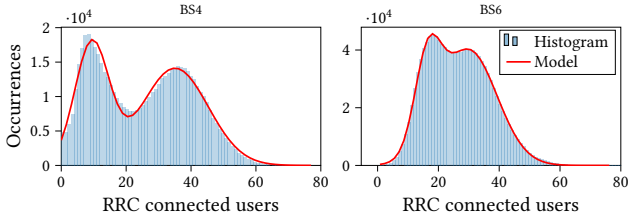

Figure 10: Bi-modal distribution

Resource Allocation. Harnessing the insights that can be obtained at the mobile phone level (*e.g.*,, via MobileInsight [22]) like many recent measurement studies do to characterize 5G performance [12, 15, 23] does not make it possible to infer inner mechanisms of BSs like scheduling. This only becomes possible by observing the traffic from the BS standpoint like mobile decoders do [5, 9, 14]. Regarding resource allocation, our dataset can contribute to advance the research at the Radio Access Network (RAN) and at the core network levels at the same time.

At the RAN level, past research has focused on anticipatory resource allocation at individual BSs [6], or in the presence of carrier aggregation [24]. This problem is becoming more and more important given that 5G networks introduce an additional layer of complexity regarding how scheduling is performed because of the introduction of network slicing and millimeter-wave technology, among other factors. It has also been shown that, to a certain degree, it is possible to infer the class of applications and applications themselves from coarse TBS information [29].

At the core network level, past research has utilized information like traffic volumes for CPU allocation of Virtual Network Function (VNF) like the Access and Mobility Function (AMF) [1]. One of the bottlenecks of the AMF lies in the signaling procedures handled at a user level. A simple use registration, for example, triggers signaling from the BS to the AMF, in turn from the AMF to the Authentication Service Function (AUSF) with at least two rounds of back and forth for a one-shot successful registration and, finally, from the AMF to the BS [19]. Therefore, forecasting the number of users that can potentially be connected at RRC level becomes important. Note that such information is also relevant for other core network functions like the Location Management Function (LMF) [37].

Applications to ML at Large. The present dataset has the potential to be highly useful for several forecasting activities where ML finds applicability. The use of a common dataset for benchmarking is key to improve model design and provision model explanations under the EXplainable Artificial Intelligence (XAI) umbrella [11]. A lack of explainability may lead to poor model design, which has been proved detrimental in the presence of adversarial attacks [31].

We have already discussed above regarding past research on anticipatory networking [6] for resource allocation and, as well, for scheduling pilot signals in a scalable fashion to ultimately improve the quality of channel estimation [10]. Mobile traffic load forecasting can be delivered via traditional ML models focusing on individual BSs or multiple BSs at a time via federated learning [36]. In Section 4.3, we have shown that the number of RRC connected users exhibits a bi-modal distribution that is amenable to forecasting. This would be certainly of interest for VNF resource allocation. Further, past research has also shown the capability of performing anomaly detection [35, 38].

Extension to 5G NR. To the present date, to the best of our knowledge, it does not exist tool able to perform DCI decoding for 5G NR as possible for LTE. The main reason lies in the way 5G NR improves over LTE. First, part of the information required for decoding is delivered via encrypted channels, unlike LTE. Second, such information is not delivered over the entire channel bandwidth like in LTE, but in sub-parts (known as bandwidth parts) to augment flexibility and increase capacity. A preliminary attempt to build such a 5G NR sniffer allows to retrieval RNTIs and basic information for traffic analysis, but not yet a fully-fledged decoder [25].

6 CONCLUSIONS

In this paper, we present and analyze a large dataset of mobile traffic data whose key distinct feature is like the data, which is based on real LTE traffic information decoded from control channel information at the millisecond level. The dataset contains traces at BS level from different mobile network operators that we make anonymous. We carry out an in-depth characterization of mobile traffic at BS and user levels and study how widely adopted probability distributions for mobile traffic do apply at such short-term scales. Our analysis shows that mobile data traffic exhibits self-similarity and that the number of RRC connected users exhibit a bi-modal distribution making our contribution key to verifying and reproducing research outcomes as well as to driving advances of AI/ML applied to mobile networks.

ACKNOWLEDGMENT

This work is partially supported by the Spanish Ministry of Science and Innovation through the Juan de la Cierva grant IJC2019-039885-I, and grant PID2021-128250NB-I00 ("bRAIN"). P. Fernández received funding from the EU--NextGenerationEU and SEPE/PRTR called "Programa Investigo" (grant 2022-C23.I01.P03.S0020-0000038).

REFERENCES

[1] 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 (2018), 42–49.

[2] G. Attanasio, C. Fiandrino, M. Fiore, J. Widmer, and et al. 2022. In-depth study of RNTI management in mobile networks: Allocation strategies and implications on data trace analysis. *Computer Networks* 219 (2022), 109428.

[3] A. Balasingam, M. Bansal, R. Misra, and et al. 2019. Detecting If LTE is the Bottleneck with BurstTracker. In *Proc. of ACM MobiCom*. 1–15.

[4] G. Barlacchi, M. De Nadai, R. Larcher, and et al. 2015. A multi-source dataset of urban life in the city of Milan and the Province of Trentino. *Scientific data* (2015).

[5] N. Bui and J. Widmer. 2016. OWL: A Reliable Online Watcher for LTE Control Channel Measurements. In *Proc. of ACM Workshop AllThings-Cellular*. 25–30.

[6] N. Bui and J. Widmer. 2018. Data-Driven Evaluation of Anticipatory Networking in LTE Networks. *IEEE Transactions on Mobile Computing* 17, 10 (2018), 2252–2265.

[7] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).

[8] P. Di Francesco, F. Malandrino, and L. A. DaSilva. 2018. Assembling and Using a Cellular Dataset for Mobile Network Analysis and Planning. *IEEE Transactions on Big Data* 4, 4 (2018), 614–620.

[9] 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.

[10] C. Fiandrino, G. Attanasio, M. Fiore, and J. Widmer. 2021. Traffic-Driven Sounding Reference Signal ResourceAllocation in (Beyond) 5G Networks. In *Proc. of IEEE SECON*. 1–9.

[11] C. Fiandrino, G. Attanasio, M. Fiore, and J. Widmer. 2022. Toward native explainable and robust AI in 6G networks: Current state, challenges and road ahead. *Computer Communications* 193 (2022), 47–52. https://doi.org/10.1016/j.comcom.2022.06.036

[12] C. Fiandrino, D. Juárez Martínez-Villanueva, and J. Widmer. 2022. Uncovering 5G Performance on Public Transit Systems with an App-Based Measurement Study. In *Proc. of ACM MSWiM*. 65–73.

[13] 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).

[14] Tuan D. H. and et al. 2023. LTESniffer: An Open-source LTE Downlink/Uplink Eavesdropper. In *Proc. of ACM ACM Conference on Security and Privacy in Wireless and Mobile Networks*. 25–30.

[15] A. Hassan, A. Narayanan, A. Zhang, and et al. 2022. Vivisecting Mobility Management in 5G Cellular Networks. In *Proc. of ACM SIGCOMM*. 86–100.

[16] H. E. Hurst. 1951. Long-term storage capacity of reservoirs. *Transactions of the American society of civil engineers* 116, 1 (1951), 770–799.

[17] R. Kalden and S. Ibrahim. 2004. Searching for self-similarity in GPRS. In *Passive and Active Network Measurement*. Springer, 83–92.

[18] A. Krizhevsky, I. Sutskever, and G. E. Hinton. 2017. ImageNet classification with deep convolutional neural networks. *Commun. ACM* 60, 6 (2017), 84–90.

[19] A. Kumar, P. Naik, S. Patki, P. Chaudhary, and M. Vutukuru. 2022. Evaluating Network Stacks for the Virtualized Mobile Packet Core. In *Proc. of ACM APNet*. 72–79.

[20] S. Kumar, E. Hamed, D. Katabi, and L. Erran Li. 2014. LTE Radio Analytics Made Easy and Accessible. In *Proc. of ACM SIGCOMM*. 211–222.

[21] Will E. Leland, Murad S. Taqqu, Walter Willinger, and Daniel V. Wilson. 1993. On the Self-Similar Nature of Ethernet Traffic. In *Proc. of ACM SIGCOMM*. 183–193. https://doi.org/10.1145/166237.166255

[22] Y. Li, C. Peng, Z. Zhang, Z. Tan, and et al. 2021. Experience: A Five-Year Retrospective of MobileInsight. In *Proc. of the ACM MobiCom*. 28–41.

[23] Y. Liu and C. Peng. 2023. A Close Look at 5G in the Wild: Unrealized Potentials and Implications. In *Proc. of IEEE INFOCOM*. 1–10.

[24] Norbert Ludant, Nicola Bui, Ana García Armada, and Joerg Widmer. 2017. Data-driven performance evaluation of carrier aggregation in LTE-Advanced. In *Proc. of IEEE PIMRC*. 1–6.

[25] N. Ludant, P. Robyns, and G. Noubir. 2022. From 5G Sniffing to Harvesting Leakages of Privacy-Preserving Messengers. In *Proc. of IEEE Symposium on Security and Privacy*. 1919–1934.

[26] B. B. Mandelbrot and J. W. Van Ness. 1968. Fractional Brownian motions, fractional noises and applications. *SIAM review* 10, 4 (1968), 422–437.

[27] O. E. Martínez-Durive, S. Mishra, C. Ziemlicki, S. Rubrichi, Z. Smoreda, and M. Fiore. 2023. The NetMob23 Dataset: A High-resolution Multi-region Service-level Mobile Data Traffic Cartography. arXiv:2305.06933 [cs.NI]

[28] T. Melodia, S. Basagni, Kaushik R. Chowdhury, and et al. 2021. Colosseum, the World's Largest Wireless Network Emulator. In *Proc. of ACM MobiCom*. 860–861.

[29] 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.

[30] J. Meng, J. Huang, Y. C. Hu, and et al. 2022. Characterizing and Modeling Control-Plane Traffic for Mobile Core Network. arXiv:2212.13248 [cs.NI]

[31] S. Moghadas Gholian, C. Fiandrino, A. Collet, G. Attanasio, M. Fiore, and J. Widmer. 2023. Spotting Deep Neural Network Vulnerabilities in Mobile Traffic Forecasting with an Explainable AI Lens. In *IEEE INFOCOM*.

[32] A. Noulas, C. Mascolo, and E. Frias-Martinez. 2013. Exploiting Foursquare and Cellular Data to Infer User Activity in Urban Environments. In *Proc. of IEEE MDM*, Vol. 1. 167–176.

[33] P. Parastar, A. Lutu, Ö Alay, G. Caso, and D. Perino. 2023. Spotlight on 5G: Performance, Device Evolution and Challenges from a Mobile Operator Perspective. In *Proc. of IEEE INFOCOM*. 1–10.

[34] V. Paxson. 1997. Fast, approximate synthesis of fractional Gaussian noise for generating self-similar network traffic. *ACM SIGCOMM Computer Communication Review* 27, 5 (1997), 5–18.

[35] A. Pelati, M. Meo, and P. Dini. 2022. Traffic Anomaly Detection Using Deep Semi-Supervised Learning at the Mobile Edge. *IEEE Transactions on Vehicular Technology* 71, 8 (2022), 8919–8932.

[36] H. P. Phyu, R. Stanica, and D. Naboulsi. 2023. Multi-Slice Privacy-Aware Traffic Forecasting at RAN Level: A Scalable Federated-Learning Approach. *IEEE Transactions on Network and Service Management* (2023), 1–1.

[37] A. Pinto, G. Santaromita, C. Fiandrino, and et al. 2022. Characterizing Location Management Function Performance in 5G Core Networks. In *Proc. of IEEE NFV-SDN*. 66–71.

[38] 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.

[39] Y. Xie and K. Jamieson. 2022. NG-Scope: Fine-Grained Telemetry for NextG Cellular Networks. *Proc. ACM Meas. Anal. Comput. Syst.* 6, 1 (Feb 2022).

[40] Y. Xie, F. Yi, and K. Jamieson. 2020. PBE-CC: Congestion Control via Endpoint-Centric, Physical-Layer Bandwidth Measurements. In *Proc. of ACM SIGCOMM*. 451–464.

[41] HF Zhang, YT Shu, and Oliver Yang. 1997. Estimation of Hurst parameter by variance-time plots. In *1997 IEEE Pacific Rim Conference on Communications, Computers and Signal Processing, PACRIM. 10 Years Networking the Pacific Rim, 1987-1997*, Vol. 2. IEEE, 883–886.

[42] M. Zhang, H. Fu, Y. Li, and S. Chen. 2019. Understanding Urban Dynamics From Massive Mobile Traffic Data. *IEEE Transactions on Big Data* 5, 2 (2019), 266–278.