ARENA: A Data-driven Radio Access Networks Analysis of Football Events

Lanfranco Zanzi, Associate Member, IEEE, Vincenzo Sciancalepore, Senior Member, IEEE, Andres Garcia-Saavedra, Xavier Costa-Pérez, Senior Member, IEEE, Georgios Agapiou, and Hans D. Schotten, Member, IEEE

Abstract—Mass events represent one of the most challenging scenarios for mobile networks because, although their date and time are usually known in advance, the actual demand for resources is difficult to predict due to its dependency on many different factors. Based on data provided by a major European carrier during mass events in a football stadium comprising up to 30.000 people, 16 base station sectors and 1 Km² area, we performed a data-driven analysis of the radio access network infrastructure dynamics during such events.

Given the insights obtained from the analysis, we developed ARENA, a model-free deep learning Radio Access Network (RAN) capacity forecasting solution that, taking as input past network monitoring data and events context information, provides guidance to mobile operators on the expected RAN capacity needed during a future event. Our results, validated against real events contained in the dataset, illustrate the effectiveness of our proposed solution.

Index Terms—Cellular network, Machine Learning, Network monitoring and measurements, Pro-active management.

I. INTRODUCTION

Mass events such as sport events (e.g., football games), religious events (e.g., holy pilgrimages), political events (e.g., demonstrations) or entertainment events (e.g., concerts) are particularly challenging for mobile operators *despite being planned* with months or weeks ahead in most cases [1]. Indeed, operators know *when* a demand surge spawning from such events is coming but they are unable to adapt timely and appropriately [2].

Traditional delay tolerant [3] and information centric [4] approaches offer methods to smooth the traffic volumes during congestion time, e.g., postponing the transmission of delay tolerant traffic outside the busy time windows, or reducing the traffic flowing in the network by a proactive content placement at the edge of the network. However, due to generally scarce radio access resource availability, none of those approaches would avoid mobile traffic congestion in the event premises.

Hans D. Schotten is with Technische Universität Kaiserslautern, Kaiserslautern, Germany. Email: schotten@eit.uni-kl.de.

G. Agapiou is with OTE, Athens, Greece. Email: gaga-piou@oteresearch.gr

To increase the spectrum availability, current approaches imply the on-demand deployment of, e.g., nomadic cells such as Cells on Wheels (CoWs) or Cells on Light Trucks (CoLTs), or fixed small (micro, pico, femto) cells for surge offloading. All these options are rather rudimentary and, needless to say, overly costly. Network slicing [5], which will allow an operator to request further slices of dormant resources (even from other operators) in a more flexible manner, can certainly reduce the cost tied to the hunt for capacity during these events. However, albeit planned ahead in time, the amount of additional resources required to cope with the demand during these situations is largely unpredictable.

Indeed, the amount of network load resulting from these events *depends on contextual features* such as the type of event—different events foster different mobile applications, such as real-time video streaming during concerts; and consumption patterns, such as data avalanches occurring during the breaks of a football match—or the ability of the event to attract attendance, such as the ranking of the matching teams in a football competition. However, although this contextual information is available, the model capturing the relationship between mobile traffic demand and the specific context of a given event *is inherently hard to build because mass events are rare and each is different in nature from one another.*

In this paper, we advocate for the use of model-free deep learning techniques to take up on the challenge. Specifically, we first analyze data obtained from a major operator regarding the cellular coverage in the stadium of an important football team in Greece. It is well understood nowadays that the use of radio spectrum through mobile systems is closely related to human activity [6]. In this way, our analysis not only makes evident the aforementioned challenges, but it also sheds light on the behavioral dynamics of sport fans when attending major football events. Then, we formulate a Bandit Convex Optimization (BCO) model to formalize the problem we are addressing. Our model formalizes some of the aspects that make capacity forecasting during mass events complex, that is, the unknown relationship between user traffic patterns, spectrum usage and the context around the event. Then, to address this problem we design a deep neural network model to estimate the amount of extra spectrum resources required during the event to preserve the quality of service experienced by end users during regular days.

To summarize, the main contributions are the following:

• We provide a quantitative and qualitative analysis of the network statistics during football events taken from a real

L. Zanzi is with NEC Laboratories Europe GmbH, Heidelberg, Germany, and Technische Universität Kaiserslautern, Kaiserslautern, Germany. Email: lanfranco.zanzi@neclab.eu.

V. Sciancalepore, A. Garcia-Saavedra are with NEC Laboratories Europe GmbH., Heidelberg, Germany. Emails: {vincenzo.sciancalepore, andres.garcia.saavedra}@neclab.eu.

X. Costa-Pérez is with NEC Laboratories Europe GmbH, Heidelberg, Germany, and i2CAT Foundation and ICREA, Barcelona, Spain. Email: xavier.costa@neclab.eu.

deployment;

- We formulate our problem by means of a Bandit Convex Optimization (BCO) model;
- We evaluate the performance of different state-of-the-art forecasting models while dealing with real-data;
- We propose a deep learning architecture, namely ARENA that takes as input monitoring metrics from Radio Access Network (RAN) devices as well as other contextual information to assess the extra spectrum resources required during the event;
- We validate the model and finally quantify the expected amount of resources to be proactively deployed in the stadium area to meet the target Quality of Service (QoS).

The rest of the paper is structured as follows. Section II collects the state-of-the-art and related works in the field. Section III discusses the scenario considered in this work. Section IV presents an overview of the dataset of networking measurements and features therein contained, providing insight on the public event dynamics from the mobile network perspective, with focus on the quality of service experienced by the end users. In Section V we model the network dynamics in the stadium by means of bandit optimization theory and propose in Section VI a machine-learning (ML) solution, namely ARENA, to capture the relationship between user activity and network utilization. Section VI-C evaluates ARENA's capabilities to forecast future network dynamics and estimate the additional amount of radio resources to improve the end-user QoS. Finally, Section VII concludes the paper.

Privacy issues: Our research activity does not violate user's privacy rights. The dataset contains only high-level aggregated and anonymous information.

II. RELATED WORK

Mass events (sports, in particular) are gaining substantial attention for analysis in the last few years [7]–[9]. This is certainly a projection of the underlying social attraction to novel applications (and the business therein) that improve the experience for attendees and participants, e.g., virtual/augmented reality (VR/AR) technologies, in the upcoming 5G reign.

An accurate characterization of mobile network channel statistics is of paramount importance to assess the network capacity boundaries. However, as highlighted in [10], stable performances in mobile networks are much more complex to be achieved due to highly variable wireless channel statistics that affect the final service provisioning. The authors of [7] present a detailed analysis of spectator behavior with the goal of designing novel mobile applications for stadium-based sporting events. In their work, they exploit Bluetooth and GPS traces to derive information such as density, location and even travel speeds of crowds of people attending the UEFA world cup final that took place at the City of Manchester Stadium in May 2008. This early study hints at the importance of capturing contextual information to manage mobile services during mass events. A thorough analysis of mobile network performance during the 2013 Superbowl is presented in [8]. The authors conclude that the use of multicast cellular coverage in combination with content caching is paramount to solve the congestion of the uplink LTE channels during super-sized events. The authors of [9] analytically investigate the root causes of cellular network performance drops during crowded events, identifying the random access procedure as the main cause for QoE degradation and suggesting to leverage on device-to-device communication to cluster network access requests as to alleviate the problem. Similarly, [11] provides an in-depth analysis of the subscribers' quality of experience focusing on roaming scenarios.

Mobile traffic forecasting has fostered substantial research effort in its own merit. In [12]-[15] the authors explore mobile traces collected from metropolitan areas to investigate and understand human mobility. In [16], [17] the authors focus on mobile resource utilization of individual services at a national scale. This work sheds the lights on interesting macroscopic properties and provides informative insights of todays' mobile networks, including spatio-temporal traffic patterns as well as the relationship between data volumes and urbanization levels. Moreover, several machine learning mobile traffic prediction solutions have been proposed in the literature. In [18], the authors use deep learning techniques for spatio-temporal modeling and prediction in cellular networks. The authors of [5] apply forecasting of mobile network resource utilization to investigate the network slicing concept in 5G systems. They design a dynamic resource allocation framework accounting for all the network domains and validate their proposal through exhaustive simulation campaigns. Deep models for long-term traffic prediction are also provided in [13], [19], [20]. Of particular interest from these works is how deep learning models substantially outperform traditional predictive methods. Finally, [21] proposes an encoder-decoder neural-network structure to forecast traffic demands on specific services (or network slices) and solve the trade-off between capacity overdimensioning and service requirements guarantees. While this work on traffic forecasting provides insightful understanding on the dynamics of regular-user regular-day when interacting with mobile services, the model capturing such behavior is not valid during rare but massively crowded events.

In the sequel, we will first analyze the case of an important football stadium in Greece and the mobile coverage provided by a major mobile operator during a regular football season. A key observation of our analysis is that the relationship between active users and radio resources during sport events substantially differ from that of a regular day and so it requires special consideration to that given in [5], [16]–[18], [21].

III. PRELIMINARY CONSIDERATIONS

Despite frustrating, the poor network performance achieved in crowded initiatives like concerts or sport matches is commonly accepted by public and professionals attending such events who, after several initial unsuccessful trials, often postpone the upload (or download) of their photos and videos on social media due to network congestion. Similar poor conditions are mostly registered during an event in case of voice calls with a clear dropping rate increase [22]. The problem is well known and accounts for two separate root causes deriving from the need of widely guaranteeing access



Fig. 1: Overview of the stadium area.

to the network, i.e., congestion in the radio access network (RAN) during the random access procedure (uplink traffic) and limited resource availability against a huge data demand surge (downlink traffic).

A. Standardization procedures

During the initial access, the User Equipment (UE) scans the system broadcast information to obtain the main system configuration parameters and synchronize with the network before attempting to attach through the physical random access channel (PRACH). The LTE standard defines two random access procedures: a contention-based and a contention-free. The former is mainly used for initial access, Radio Resource Control (RRC) layer connection (re-)establishment and for uplink data transmission in non-synchronized states or in absence of scheduled resource availability. The latter, less common, is mostly implemented to allow fast handover procedures. The RRC protocol of LTE system includes, among the others, connection establishment and connection release, system information broadcast and radio bearer setup and reconfiguration [23]. In LTE, the base stations or eNBs centralize most of the RRC functionality as well as the resource management and packet scheduling. The set of activities performed by the eNB varies according to the state of the radio connectionwhich can be either active (Connected) or not active (Idle)between the UE and the network. In the RRCIDLE state, the UE passively monitors the system information broadcast by the network without the need to send monitoring reports or mobility updates. Conversely, UEs move into RRC_{CONNECTED} state when performing a call or transmitting data, which increase the monitoring information reports sent to the radio access node in order to keep the session active. During the contention-based procedure, the UE randomly selects one of the $64 - N_{cf}$ orthogonal preamble signatures and the next available subframe for PRACH transmission, where N_{cf} is the number of preamble signatures allocated for the contentionfree procedure (its value can change according to traffic load in the system). After sending the attachment request, the UE monitors the physical downlink control channel (PDCCH) and schedules a timer. In case of simultaneous RACH requests, the eNB uses the preamble sequence to differentiate among users. However, specially in crowded areas, different requests may collide. In this case, the eNB will not be able to decode the requests thereby triggering an exponential backoff procedure in the UE which delays the next access attempt. If no other user selects the same preamble sequence, the eNB is able to decode the request and reply with a Random Access Response (RAR) message through the PDCCH. At this point, the UE sends a RRC connection request message including a temporary identifier and the establishment cause. If accepted by the network, the access procedure terminates, and the UE moves into RRC_{CONNECTED} state so that it is allowed to use the data communication services of the system.

Several optimization steps can be adopted to increase the performance of the RACH procedure. Typical deployments assume collision probability between UEs in the order of 1% with periodic random access occasions distributed every 10 ms. This setup translates into the possibility to handle an offered load of 128 attempts/second over 10MHz bandwidth [24]. Clearly, such a setup leads to low performance when dealing with crowded events. Typical approaches to solve this situation involve the reduction of the access occasion cycle duration and the increase of access opportunities within one frame, however this negatively impacts on the PRACH overhead finally reducing the scheduling opportunities for user data transmissions.

B. Technical challenge

The second issue, data demand surges, is more intuitive and involves reduced rate of transmission opportunities (delay) and small chunks of resources scheduled per user within the Physical Downlink Shared Channel (PDSCH) (low throughput) due to a high number of users in RRC_{CONNECTED} state.

Our work focuses on the latter issue. A simple approach to solve it is to smooth down traffic patterns by delaying load opportunistically to time instances with a less-congested network [3], [8]. However, this approach is only acceptable for delay-tolerant applications and hardly fits with the real-time nature of most use cases motivating 5G applications. The only solution upon such context is to increase the capacity of the system by deploying nomadic cells such as Cells on Wheels (CoWs) or Cells on Light Trucks (CoLTs), or offload traffic to fixed small cells, which effectively increases the capillarity of the network and hence the density of radio access points.

Unfortunately, this approach is overly expensive and slow, incapable of adapting to different traffic patterns occurring during events of different nature or events with different number of attendees. Certainly, network slicing opens the door for novel ways of increasing the capacity to the network opportunistically in a more agile manner (e.g., leasing additional spectrum through a network slice from neighbouring operators)—in the matter of few hours or even minutes. Still, the human relationship with mobile applications and the number of attendees substantially differ from one event to another. Summing up the fact that such massive events are rare (though planned), makes forecasting capacity requirements particularly daunting.



Fig. 2: Temporal evolution of the aggregated downlink (DL) and uplink (UL) traffic volume in the stadium area during a regular day (dashed line) and during the day of a large event (straight line).

IV. DATA ANALYSIS

In this section we present a data-driven analysis of sport events. Our dataset consists of weeks of monitoring data including large sport events taken place in a football stadium during a regular season under the coverage of 6 LTE eNBs, subdivided into 16 sectors, from a major operator in Greece. The set of eNBs covers an area of approximately 1 Km². Such concentration of radio access points is required to accommodate the traffic peaks generated by the almost 30000 people hosted in the stadium during sport events. For the same reason, the base stations are equipped with directive arrays of antennas and support Carrier Aggregation (CA). As shown in Fig. 1, the surrounding district includes both residential neighbourhoods and vehicular streets, which highly characterize the spatio-temporal behaviour of the traces as later detailed in Section IV-B. We do not disclose the location of the eNBs due to privacy matters. Our goal is to characterize the relationship between active users during the events and network usage in contrast to regular days. Therefore, we focus our analysis on the average throughput per active user experienced by the event attendees in the stadium. Throughout the rest of the paper, we will consider it as the key performance indicator (KPI) of the perceived quality of service (QoS) experienced by the end-users. To ease our analysis, we focus on a particular sport match, taken place in February 10th, 2019, which gathered 25097 attendees in

A. Traffic Volume Patterns

Fig. 2 shows the aggregated traffic volume generated within a 24-hour time period of a regular weekday (dashed line) and also during the day of the event (straight line), discerning between uplink (UL) and downlink (DL) traffic with red and blue lines, respectively. The event is scheduled to start at 19.30h and finish at 21.15h. We first observe that both UL and DL patterns are remarkably similar between a regular day and the day of the event up until 3 hours before the beginning of the event and the anomaly data usage persists until around 2 hours after the end of the event. This clearly reflects common human behavior as attendees usually gather in the surroundings of the stadium for social interactions before and after the event. This is evidenced by two volume peaks at around 18h and at around 21.30h. The former gradually vanishes until the event begins, coinciding with crowds moving towards the allotted seats in the stadium, whereas the latter peaks practically immediately after the end of the event and gradually vanishing as crowds move out of the stadium. Interestingly, there is a third peak, concomitant with the break of the event, with two drops in volume during the two halves of the event, which is expected as users are focusing on the game itself. We would like to remark that this pattern repeats across all events of the season with a gain factor dependent on the number of attendees and represents a signature of football games. Different events will leave different footprints and motivate the use of modelfree approaches to forecast capacity requirements.

and aggregated features with a time granularity of 15 minutes.

This ensures that the subscribers' privacy rights are preserved.

Fig. 3 compares the same distributions of traffic volume across different time periods, namely, morning (6-12h), afternoon (12-18h), evening (18-24h) and night (24-6h), as a function of the number of active users (RRC_{CONNECTED}) for both downlink and uplink in a regular weekday (first and second plots) and for the day of the event (third and fourth plots). As expected, the volume distribution during the day of the event is skewed with a long tail concentrating large number

¹While our analysis is based on a single event, we have exhaustively evaluated other 10 different matches finding no relevant performance differences, which makes our approach generalizable.



Fig. 3: Normalized data volume distribution as a function of the number of active users (RRC_{CONNECTED}) for downlink (DL) and uplink (UL) and different periods of the day (morning, afternoon, evening, night).



Fig. 4: Evolution of users activity. Peak and average number of active users (RRC_{CONNECTED}) (outer plot) and peak-to-average ratio (inner plot) during the day of the event and during a regular day.

of users during the afternoon and the evening, precisely the new connections from the event attendees. This allows us to characterize the background traffic volume, generated by less than 250 active connections spawning less than 50% of the total system capacity following a strongly linear relationship. Such linearity is broken during the event for downlink traffic. The average data volume by active users within this regime (above 250 users) is lower than in normal conditions, which suggests poor quality of service experienced by the end users in the stadium.

B. Temporal Distribution of Mobile Users

Fig. 4 compares the activity of active users ($RRC_{CONNECTED}$) in the neighbourhood of the stadium during the day of the event (blue lines) with that of a regular weekday (dark lines), aggregated over all the eNBs in the area. The figure shows the average number of users within each 15-minute time window (dashed lines) and the maximum number of active users within the same window (straight lines). As expected, the number of active connection peaks are highly correlated to the traffic volume time evolution shown in Fig. 2, with the number of concurrent connections ranging up to 600 $\frac{active UEs}{minute}$, over 10 times higher than those during a regular day. The areas highlighted in light blue between straight and dashed lines provide an indication of the variability of the number of active UEs, i.e., the peak-to-average active connections, and



Fig. 6: Evolution of the average downlink (DL) capacity per user (our QoS metric) during the day of the event (outer plot) and a regular day (inner plot). Highlight areas indicate the two periods of time when the sport event is occurring.

evidence a substantial increase in volatility (around 2 times higher peak-to-average connections) during the course of the event. To visualize this, we create an inner plot depicting the peak-to-average ratio in both scenarios (event and regular day). Higher volatility in user activity is observed during the night and early in the morning, which is explained by the low number of net active connections (small variations yield higher peak-to-average ratios). Interestingly, however, the average-topeak ratio increases during the event, which implies that the operator must deal with substantial control plane volatility in addition to data volume surges. It is clear that as soon as the base station reaches the saturation point, the perceived QoE decreases due to limited scheduling of resources. We further investigate this point in the following subsections.

C. Service Degradation

As mentioned above, we conjecture that the skewness of the volume distribution shown in Fig. 3 is due to poor quality of service. To confirm this hypothesis, we take a closer look to the physical spectrum usage, namely the downlink physical resource block (PRB) utilization measured at the eNBs. In this way, Fig. 5 compares the relative PRB utilization as a function of the number of active users during the event (first and third plot) and that of a regular day (second and fourth plot) for the different frequency bands used by the sectors of the eNBs (first two plots) and for different periods of time during the day



Fig. 5: Downlink PRB utilization as a function of the number of active users during the day of the event and during a regular day across three frequency bands managed by the eNB and different time periods (morning, afternoon, evening, night).

6



(a) Overview of different metrics collected by two eNBs in the stadium area during an event day and weekday (Normalized values).

(c) Correlation matrix for different traces collected by a lightly loaded eNB during a weekday (left) and the event (right).

Fig. 7: Temporal correlation of features in different working conditions.

(second two plots). We note that the 1800-MHz band saturates when the number of users grows over 250, which unavoidably translates into QoS degradation for the end users. The time analysis confirms that this event occurs during the time period of the football match.

This is further confirmed in Fig. 6, which compares the average throughput capacity per user over time during the day of the event (outer plot) against a regular weekday (inner plot), considering each of the three bands selected for our study. As expected, individual QoS substantially deteriorates during the sport event, particularly during a 2.5-hour time window before the beginning of the event, during the break of the event and during a 1.5-hour time window after the end of the event. With respect to standard working conditions, the average throughput capacity per user decreases up to 50 times.

D. Inter-Feature Time Correlation

We now study how different features relate to each other and across different neighbouring eNBs from a statistical perspective. To ease the presentation, we focus on two eNBs: an eNB that carries significant traffic volumes during both week and event days, labelled as "High Load eNB" and represented with a blue color palette in Fig. 7; and an eNB that serves less significant amount of traffic, labelled as "Low Load eNB" and represented with a red color palette. Specifically, Fig. 7a shows the temporal dynamics of different features collected by the two eNBs during a regular weekday (dashed lines) and during a match game (shaded area), while Figs. 7b and 7c depicts their correlation matrix. From Fig. 7a, it clearly appears the performance gap with standard operating conditions. Specially for the UL case, during the event both the "High Load" and "Low Load" eNBs present traffic peaks 10 times higher than the traffic volumes achieved in normal conditions. Such UL activity peak can be justified by a sudden increase of social media activity and background cloud synchronization processes [8].

Let us first focus on the left matrices of Figs. 7b and 7c. As expected, since both the eNBs are deployed in the same area and the human activity corresponds to that of a regular weekday, the correlation value (Pearson's r coefficient) between each pair of features is positive and similar in both eNBs. This is true for every couple of features except the pair "DL PRBs-DL Volume", which appears strongly correlated for the highly loaded eNB and only moderately correlated for the lightly loaded eNB. It should be noticed that the latter exhibits a rather flat behavior during normal working conditions. More interesting are the correlation matrices we observe during the match, depicted by the two right-most plots of Figs. 7b and 7c. First, the lightly loaded eNB shows higher correlation values with respect to normal working conditions due to a more dynamic pattern enforced by the attendees in the area. Second and more surprisingly, we observe that most of the correlations available during a weekday in the highly loaded eNB vanish during the match, eventually providing negative value for the couple "Active UEs-DL Volume". The latter represents a counter-intuitive behaviour explainable by the fact that the considered eNB saturates its physical resources (PRBs) during most part of the game distorting the underlying relationship between the features. Together, these observations provide valuable insights not only on the fact that in normal working conditions features such as Active UEs and DL Volume are correlated (which is somewhat expected), but also that

during mass events distortions may occur. This motivates us to design a mechanism that exploits not only the information contained within each time series feature, but also the timedomain correlation that heterogeneous sequences might show among themselves. Moreover, as each eNB exhibits different behaviour, we advocate to model each cell individually taking into account both regular and crowded situations.

V. MODEL DESIGN

In our data analysis, we showed that current network deployments may be insufficient during football events as the unexpected-high density of users can quickly saturate the available network resources. However, providing additional spectrum via increasing the density of radio access points is overly expensive, even with future network slicing mechanisms. Hence, proper mechanisms to accurately account for the required additional resources so that certain QoS targets are satisfied must be in place. Hereafter, we model the problem of network configuration during football events and suggest a new network setting to be dynamically applied by the operator.

A. Notation

We use conventional notation. We let \mathbb{R} and \mathbb{Z} denote the set of real and integer numbers. We use \mathbb{R}_+ , \mathbb{R}^n , and $\mathbb{R}^{m \times n}$ to represent the sets of non-negative real numbers, *n*-dimensional real vectors, and $m \times n$ real matrices, respectively. Vectors are in column form and written in bold font. Matrices are in upper-case font. Subscripts represent an element in a vector and superscripts elements in a sequence. For instance, $\langle \mathbf{x}^{(t)} \rangle$ is a sequence of vectors with $\mathbf{x}^{(t)} = (x_1^{(t)}, \ldots, x_n^{(t)})^T$ being a vector from \mathbb{R}^n . In turn, $x_i^{(t)}$ is the *i*'th component of the *t*'th vector in the sequence, where superscript *T* represents the transpose operator. $\|\vec{x}\|_2$ represents the 2-norm or Euclidean norm of \vec{x} and $\|\vec{x}\|_{\infty}$ its maximum norm $(\max_i |x_i|)$. Finally, $[\cdot]^+$ denotes the projection of a vector onto the non-negative orthant, i.e., $[\vec{x}]^+ = (\max\{0, x_1\}, \ldots, \max\{0, x_n\}), x \in \mathbb{R}^n$.

B. Problem definition

Let us denote the set of base stations in the considered area as $b \in \mathcal{B}$, where $B = |\mathcal{B}|$ is the total amount of base station sectors. Without loss of generality, we consider each sector as an independent base station as each sector shows heterogeneous behaviour and each of them cover a different subset of the users in the area. We assume that time is split into different epochs $e \in E$ with duration T. Let r_{μ}^{e} denote the amount of traffic served for user $u \in \mathcal{U}_{h}^{e}$ during epoch e (bits), where $U_{h}^{e} = |\mathcal{U}_{l}^{e}|$ is the average number of users under the coverage area of base station b connected concurrently within epoch e. Each base station b is provided with a spectrum capacity c_{h}^{e} (number of PRBs), which can be dynamically changed (every epoch e) based on the carrier aggregation policies of the telco operator. Note that in case of seriouslycongested events, the operator may even decide to augment the overall spectrum capacity by placing portable base stations (e.g., Cells on Wheels or Cells on Light Trucks) or leasing or sharing a portion of the available bandwidth from other operators via, e.g., network slicing [25]. We now characterize the satisfaction of individual user connections by formally introducing the QoS metric used in Fig. 6 as the following expression:

$$q_b^e = \frac{\sum_{u \in \mathcal{U}_b^e} r_u^e}{\sum_{u \in \mathcal{U}_b^e} \tau_u^e}, \qquad \forall b \in \mathcal{B}.$$
 (1)

where r_u^e represents the volume of traffic per user under the base station coverage served within the epoch e, and τ_u^e the effective downlink time per user, which accounts for the time when the first part of the PDCP SDU of the downlink buffer was transmitted until the buffer is emptied. It should be noticed that in presence of congestion and fair scheduling of resources, the average active time per user increases as a result of resource contention, thus leading to lower QoS. However, τ_u^e depends on the spectrum capacity covering the event *and* on the contextual data around the event (type of event, expected number of attendees) and the model capturing the behavior of τ_u^e is not known *a priori* and so is q_b^e 's.

Objective. The target is to find the best PRB allocation on different deployed base stations such that the overall cost of such additional spectrum is minimized such that the aggregated QoS measured from all considered base stations is above a certain pre-defined threshold Γ . Of course, our focus is on the amount of spectrum required during mass events, which will typically exceed the capacity of the system and depends on the context of the event itself.

Constraints. We consider two sets of constraints. First, we need to impose that traffic queues are stable, that is, they do not grow infinitely over time, which would lead to dramatic drops in latency performance. We do not impose an instantaneous serving rate greater than the arrival rate; instead we consider that the long-term base station serving rate must be greater or equal to the long-term data arrival rate. Formally, this can be expressed by

$$\sum_{e \in \mathcal{E}} \sum_{u \in \mathcal{U}_b^e} \left(r_u^e - \zeta_u^e(c_b^e) \right) \le 0, \quad \forall b \in \mathcal{B}.$$
 (2)

where ζ_u^e is the average throughput over epoch *e* for user *u*. Second, we shall impose certain threshold on individual QoS performance, which in turn can be succinctly described by

$$\sum_{e \in \mathcal{E}} \left(\Gamma_b^e - q_b^e(c_b^e) \right) \le 0, \quad \forall b \in \mathcal{B}$$
(3)

where $-q_b^e(\cdot)$ is the QoS of base station *b* during epoch *e*. While both functions $q_b^e(\cdot)$ and $\zeta_u^e(\cdot)$ are not known, we can get estimates at the end of each epoch *e*.

Convex functions and decision variables. Telco operators may apply different base station settings, i.e., the number of available spectrum resources (PRBs), and decide to increase the total number of available spectrum in order to tune users' performance. We assume that the functions capturing throughput and QoS performance are both convex with respect to the amount of spectrum or number of selected PRBs (the larger the spectrum, the higher the user throughput, the better the QoS). Note that a trivial solution to satisfy all user requirements, i.e., to keep the QoS at reasonable levels, is to provide an infinitely large amount of spectrum. However, deploying additional spectrum (e.g. via CoWs or network slicing) incurs in a cost δ^e for the operator that we shall minimize.

C. Bandit convex optimization

With the above information, we can mathematically formulate the following optimization problem

$$\min_{\vec{c} \in \mathbb{N}^{B \times E}} \sum_{e \in \mathcal{E}} \sum_{b \in \mathcal{B}} \delta^{e}(c_{b}^{e}) \tag{4}$$
s.t.
$$\sum_{e \in \mathcal{E}} \sum_{u \in \mathcal{U}_{b}^{e}} \left(r_{u}^{e} - \zeta_{u}^{e}(c_{b}^{e}) \right) \leq 0, \quad \forall b \in \mathcal{B};$$

$$\sum_{e \in \mathcal{E}} \left(\Gamma_{b}^{e} - q_{b}^{e}(c_{b}^{e}) \right) \leq 0, \quad \forall b \in \mathcal{B};$$

$$c_{b}^{e} \in \mathbb{R}^{+}, \quad \forall b \in \mathcal{B}, e \in \mathcal{E};$$

where ζ_u^e , q_b^e and δ^e depend on the amount of spectrum (PRBs) c_b^e allotted to base station *b* during epoch *e*. We assume the following cost function

$$\delta^{e}(c_{b}^{e}) = \begin{cases} 0, & \text{if } c_{b}^{e} \leq 1, \\ \alpha (c_{b}^{e})^{\beta}, & \text{if } c_{b}^{e} > 1; \end{cases}$$
(5)

so $\delta(c_b^e)$ captures the fact that base stations might be assigned with PRBs already available within the bandwidth provisioned to the $(c_b^e \leq 1)$ and so it does not incur in extra cost, or additional spectrum is required $(c_b^e > 1)$ with a cost parametrized by α and β .

Recalling the online convex optimization theory [26] and considering an invariant feasible set $C = (c_b^e)$, we can rewrite the above optimization problem with its online Lagrangian version as follows

$$\mathcal{L}_e(\vec{c}^{\ e}, \vec{\lambda}) := \delta^e(\vec{x}) + \vec{\lambda}^T \ \vec{g}^{\ e}(\vec{x}), \tag{6}$$

where $\vec{c}^{e} = (c_{b}^{e})$ is the set of selectable PBRs during epoch e and $\vec{g}^{e}(\vec{c}^{e}) = \sum_{e \in \mathcal{E}} \left(\vec{\Gamma}^{e} - \vec{q}^{e}(\vec{c}^{e}) + \vec{r}^{e} - \vec{\zeta}^{e}(\vec{c}^{e})\right)$. We can then write the recursive \vec{c}^{e+1} given the prime iteration \vec{c}^{e} as the following

$$\vec{c}^{\,e+1} = \arg\min_{\vec{x}\in C} \nabla_{\vec{x}}^T \mathcal{L}^e(\vec{c}^{\,e}, \vec{\lambda}^{\,e})(\vec{x} - \vec{c}^{\,e}) + \frac{1}{2\sigma} ||\vec{x} - \vec{c}^{\,e}||^2, \quad (7)$$

where σ is a predefined constant and $\nabla_{\vec{x}}^T \mathcal{L}^e(\vec{c}^e, \vec{\lambda}^e) = \nabla \delta^e(\vec{c}^e + \nabla \vec{g}^e(\vec{c}^e) \vec{\lambda}^T$. This admits the following solution:

$$\vec{c}^{e+1} = \mathcal{P}_C\left(\vec{c}^e - \sigma \nabla_{\vec{x}} \mathcal{L}^e(\vec{c}^e, \vec{\lambda}^e)\right),\tag{8}$$

where the projector operator is $\mathcal{P}_C(\vec{y}) = \arg\min_{\vec{c} \in C} ||\vec{c} - \vec{y}||^2$, and the dual update as the following

$$\vec{\lambda}^{e+1} = \left[\vec{\lambda}^{e} + \mu \left(\vec{g}^{e}(\vec{c}^{e}) + \nabla^{T} \vec{g}^{e}(\vec{c})^{e}\right)(\vec{c}^{(e+1)} - \vec{c}^{e})\right)\right]^{+}, \quad (9)$$

where $\mu > 0 \in \mathbb{R}^+$ is a step size.

Given that functions $\zeta(\cdot)$ and $q(\cdot)$ are not known, we need to construct a stochastic gradient estimate of unknown functions using the limited value information [27]. In particular, we can evaluate the function at a perturbated point $\vec{x} + \epsilon \vec{u}$ yielding that $\nabla f(\vec{x}) \approx E[\hat{\nabla}^1 f(\vec{x})]$, where $\hat{\nabla}^1 f(\vec{x}) = \frac{d}{\epsilon} f(\vec{x} + \epsilon \vec{u}) \vec{u}$ is the one-point gradient [28]. Therefore, we need to recall the bandit online optimization theory to rewrite Eq. (8) as follows

$$\hat{\boldsymbol{c}}^{e+1} = \mathcal{P}_{(1-\epsilon)} C\left(\hat{\boldsymbol{c}}^e - \sigma \hat{\nabla}_{\vec{c}}^1 \mathcal{L}^e(\vec{c}^e, \vec{\lambda}^e)\right), \tag{10}$$

where the projector operation is performed within a subset of C to ensure feasibility of the perturbed \hat{c}^e . We can then write the dual update operation (Eq. (9)) as the following

$$\vec{\lambda}^{e+1} = \left[\vec{\lambda}^{e} + \mu \left(\vec{g}^{e}(\hat{c}^{e}) + \nabla^{T} \vec{g}^{e}(\hat{c}^{e})\right)(\hat{c}^{(e+1)} - \hat{c}^{e})\right)\right]^{+}, \quad (11)$$

where $\hat{\vec{c}}^e$ is the learning iterate.

The above considerations suggest that the model may fail while trying to approximate functions $\zeta(\cdot)$ and $q(\cdot)$. In addition, we need the exact characterization of the mobility patterns influencing the number of connected user under each base station \mathcal{U}_b^e , for each time epoch e and base station b, as well as the expected traffic demand r_u^e . Unfortunately, all these variables present a strong relationship with respect to multiple real-world variables, making our objective hardly achievable within a reasonable time window. Therefore, in the next section, we propose a machine-learning-based approach to automatically learn and approximate these functions, and provide forecasts on the number of active users as well as their traffic request patterns based on a data-driven approach.

VI. ARENA: DESIGN AND PERFORMANCE

Our solution design, ARENA, requires two different MLbased models to estimate the number of active users and their traffic patterns. On the one hand, an LSTM-based approach allows to predict repetitive time patterns. On the other hand, a Deep-Learning model allows to approximate complex functions, e.g., those related to human mobility and cellular traffic loads. Thus, we first focus on inferring the number of active users in the stadium area, and then we exploit this information to estimate the additional resource utilization necessary to guarantee normal-condition QoE. Finally, we show the performance of ARENA applied to real traffic data.

A. Active users

Spatio-temporal characteristics of mobile traces unveil patterns that might be exploited to perform a forecasting process. In this context, Long Short-Term Memory Recurrent Neural Networks (LSTM-RNNs) [29] have attracted attention given their capacity to capture repetitive schemes within unstructured time series. An LSTM network is composed by a chain of units, which sequentially apply a linear combination of operations on the input data. Such a structure is fundamental to provide the network with the capability to remember useful information gathered from the past data samples and learn long-term trends in the input sequence making LSTM wellsuited to handle mobile data traffic with significant spatiotemporal correlations [16]. LSTMs can be represented as a chain of J units, where each LSTM unit $i \in \mathcal{J}$ consists of 4 main elements, a memory cell and three gates, namely the input I_i , the output O_i and the forget gate F_i . The input and output gates are responsible to insert and read back the data into/from the memory cell itself, while the forget gate



Fig. 8: Comparison of different LSTM-based architectures.

determines how much details should be kept or removed from the unit at each iterative step. The information stored within each unit $j \in \mathcal{J}$ is commonly referred as "state" of the unit, or C_i . To build more complex models, LSTM units are usually stacked and concatenated into consecutive layers, where the output of one layer feeds the input of the subsequent one, and the hidden state h_i of each unit holds the information about previously observed data. During the training phase, the input and the previous hidden state are combined to form a vector, which holds current and past information. Then, the vector passes through several activation functions, namely sigmoid and hyperbolic tangent functions depicted as σ and *tanh* in Fig. 8b, which help regularizing the outcome of the linear combinations, finally adapting weight W and bias b values of each gate. Although the LSTM model has proven powerful in handling temporal correlation, existing work in the literature suggests the use of LSTMs together with Convolutional Neural Network (CNNs) in order to increase the accuracy of the learning models [19]. The key-idea is to exploit not only the information contained within each time series, but also the time-domain correlation heterogeneous sequences might show among themselves, as discussed in Section IV-D. For this purpose, we consider two well-known architectures, namely CNN-LSTM [30] and ConvLSTM [31], and compare their capabilities in predicting the number of RRC_{CONNECTED} devices against the legacy LSTM model within a week time span.

As depicted in Fig. 8, the main difference about the two models is that ConvLSTM embeds the convolution steps into the LSTM unit, while CNN-LSTM models imply the concatenation of the two types of networks sequentially². Both approaches therefore require the input data (and the output of the internal transfer functions) to be shaped as tensors to take advantage of the 2D representations of the network activities and include spatio-temporal information to learn valuable patterns within the dataset. The overall dataset has been split according to a 60/40 ratio for the purposes of training and validation, respectively. We train the models exploiting a moving window of 6 days as input data and inferring on the consecutive one, over a time span of several weeks. Within our experimental setup, we considered several hyperparameter

settings. In particular, we account for a variable number of LSTM cells $J = \{200, 300, 400\}$ and convolutional filters $\Phi = \{64, 128, 256\}$ for each model architecture. The legacy LSTM model is composed of two consecutive layers of J units and a final fully-connected layer that provides the resulting output. The CNN-LSTM model accounts for two convolutional layers characterized by Φ filters each and a kernel size of 3 to learn spatial features, which are followed by an LSTM layer of J units and a final fully-connected layer to learn the correlation in time. With respect to legacy LSTM models, the ConvLSTM architecture neglects the internal dense connection among gates in favour of a convolution operation, thereby reducing the number of model parameters and decreasing the chances of over-fitting. We adopt the mean-squared error (MSE) metric to train the models, and choose the Adam optimizer to optimize the loss function [32]. For the sake of comparison, we additionally include a baseline autoregressive integrated moving average (ARIMA) model characterized by parameters (p, d, q), where p is the order of the autoregressive term, d is the number of differencing required to make the time series stationary, and q is the moving average order [33]. This model requires the input time series to be stationary. Therefore, we applied simple differencing techniques to satisfy this requirement. Additionally, we implemented a grid search algorithm to optimize the ARIMA parameters choice, leading to the following settings (p, d, q) = (5, 1, 0). We resume in Table I the resulting MSE and maximum absolute error measured over the forecasted sequences, averaged over 4 weeks of testing data. As expected, due to the non-stationarity of traffic traces (demonstrated by sudden peaks during the most busy time periods) the ARIMA predictor provides weak performances when dealing with the event time period, resulting in a maximum absolute error of 0.505. From our findings, we can conclude that the CNN-LSTM model with J=200and Φ =64 provides best performances in this scenario. Hence hereafter, we will adopt this model as our active user predictor.

Fig. 9 provides a comparative view of the models exploiting a walk-forward validation. Using such settings, the model is retrained as soon as new observations are made available thus expanding the time window horizon in each training step. Past predictions are then stored and evaluated

 $^{^{2}}$ We refer the reader to state-of-the-art literature such as [30] and [31] for further implementation details.



Fig. 9: Comparison of different LSTM-based forecasting models working on real data from a base station in the stadium area (Normalized values).

against a longer list of observations, leading to more accurate forecasting performances. The different models exhibit good performance in forecasting the number of active users during normal working conditions thanks to the highly repetitive patterns. However, as highlighted by the inner-plot of Fig. 9, the prediction quality deteriorates during the event time period, resulting in an underestimation of the activity peaks. Due to its asymmetry with respect to standard working conditions, it is clear that temporal information alone is not enough to accurately predict the mobile traffic activity in the stadium during public events. In the following, we thus propose a deep learning model that in addition to spatio-temporal information also exploits contextual information to enhance the accuracy of the predictions.

B. Resource utilization

The relationship between QoS and availability of spectrum resources, characterized by functions $\zeta_u^e(c_b^e)$ and $q_b^e(c_b^e)$, is unknown and hard to construct due to the rarity and particularity of mass events. We hence focus on model-free deep learning methods to build our capacity forecasting mechanism.

To this aim, we design a neural network structure per base station, each receiving two sets of inputs pertaining each respective eNB: *i*) a representation of contextual information over the event (such as expected number of attendees, type of the event, etc.), and *ii*) network-specific monitoring data over time (historical time series), as depicted in Fig. 10. In order to reduce the dimensionality of the contextual data yet extract meaningful information we use a standard sparse autoencoder (SAE) [34], [35, Ch.14]. In brief, a SAE consists of two feed-forward neural networks: an encoder (with an output layer of size 1 in our case) and a decoder (with an output layer of size

TABLE I: Error characterization over different settings and models.

	J/Φ	LSTM	CNN-LSTM	ConvLSTM	ARIMA
MSE	64/200	0.0039	0.0032	0.0036	
	128/300	0.0041	0.0035	0.0040	0.055
	256/400	0.0045	0.0039	0.0037	
MAX ABS Error	64/200	0.321	0.339	0.322	
	128/300	0.296	0.322	0.318	0.505
	256/400	0.321	0.328	0.321	1



Fig. 10: Deep neural network architecture used in ARENA.

equal to the dimensions of the input of the encoder). They are trained together so that the reconstructed output of the decoder is as *similar* as possible to the input of the encoder. During exploitation, we simply use the encoder part of our SAE, which effectively compresses the contextual information into a 1-dimensional value. For the purpose of this study, we concentrate on the expected number of attendees to the event as contextual data, and empirically select the CNN-LSTM model as main forecasting module in light of the results discussed in Sec. VI-A. On the other hand, the second set of inputs concerns the historical traces of network measurements introduced in Section IV. In particular, we select the time evolution of average connected users U_{h}^{e} (for base station b during epoch e), volume of downlink traffic as well as our QoS metric q_{b}^{e} . Thus, we train a deep neural network for each base station in the stadium area using 60% of the football events in our dataset and validate the outcome using the remaining 40%.

Fig. 10 depicts the overall architecture of our system, consisting in two neural networks connected in series. The first neural network accounts for two layers of two-dimensional Convolutional Neural Network (2D-CNN) with η and μ filters respectively. The CNN is a class of deep neural network generally used in computer vision for their ability to detect patterns in input images. The workflow assumes each filter to be convolved along width and height of the input image to produce an activation map. Different filters detect distinct features so that a set of activation maps are passed to the next layer of the CNN. We leverage on this and select from the overall measurements matrix M the information related to a specific base station b, i.e., M_b . Then, we split M_b in epochs, or snapshots s_b^e , generating smaller input matrices of size $(F \times T)$, where F is the number of features and T is forecast horizon. We adopt a sliding window over time as data augmentation technique. A normalization function $\mathcal{N}(\cdot)$ is applied at this point to favour the learning process.

The neurons of the 2D-CNN apply a filter $\mathcal{H}(\sum_{e \in \mathcal{E}} \mathbf{I}_e \otimes \mathbf{K} + \mathbf{b}_e)$ where \mathbf{I}_e is the input matrix at epoch *e* (i.e., $\mathbf{I}_e = \mathbf{s}_b^e$), \otimes indicates the 2D convolution operator, **K** is the filter kernel, **b** is a bias vector and $\mathcal{H}(\cdot)$ is a non linear activation function. The significant development of deep learning methods lead to the definition of plethora of different activation functions in the literature [36]. We empirically select the Rectified Linear Units (ReLu) function to overcome the vanishing gradient



Fig. 11: Impact of different hyperparameter settings on ARENA forecasting performances.

problem and allow the models to learn faster. In order to reduce overfitting in favour of generalization, the inner layers of the CNN are interleaved with a Dropout layer. At each training stage, individual nodes and related incoming/outgoing edges are either dropped (with probability p) or kept (with probability 1 - p). This step allows decreasing the co-dependency from adjacent neurons during the training phase. The dropout probability is fixed to 0.2. The output of our CNN consists on a Flatten layer, so we map the input matrices into a fixed-length vector that feeds the second stage of the architecture.

The second neural network consists in three feed-forward layers. In this structure, also called Multi-Layer Perceptrons (MLPs), the neurons of one layer are connected with all the neurons of the adjacent layer to build a finite and acyclic graph. The number of neurons in the first and second layer are fixed to 128 and 64 respectively, while the size of the output layer of the network matches with the number of samples to be predicted T for both output estimations: the active number of users and the PRB utilization during the event. Importantly, the MPLs inherit from standard multi-layer feed-forward networks the capability to approximate continuous n-dimensional functions $f : \mathbb{R}^n \to \mathbb{R}^n$ as stated by the universal approximation theorem [37]. Thus, the output of each MPL layer is a linear combination of weighted real-valued input and a nonlinear activation function, namely $y = \phi(\mathbf{w}^{T}\mathbf{x} + \mathbf{b})$, where **w** denotes the vector of weights, \mathbf{x} is the vector of inputs, \mathbf{b} is the bias and $\phi(\cdot)$ is the non-linear activation function. Also in this case, we adopt the ReLu function for the hidden layers and a linear function for the output layer to retrieve the forecasting samples amplitude. The output of the neural network includes a Tdimensional vector predicting the amount of PRB utilization \vec{c}_b during the T epochs covering the event and a T-dimensional vector with the respective prediction on the number of active users \bar{U}_b . The training phase exploits the adaptive moment estimation (Adam) optimizer, with learning rate 10^{-4} . At each iteration, the model parameters are adjusted to minimize the error between predictions and ground truth measured by means of a loss function. In our model, we exploit the Mean Squared Error (MSE) loss function [38].

The choice of the best hyperparameter settings for ARENA

has been performed by means of an accurate optimization process aiming at reducing the forecasting error. In Fig.11 we investigate the effects of a variable number of convolutional filters η and μ over the two convolutional layers included in our proposed architecture. To deal with the highly heterogeneous traffic patterns, we have individually trained each model exclusively accounting for the traffic traces coming from the respective base station. The methodology to train each model follows the one previously presented. The training phase accounts for 50 epochs and runs over 60% of the overall dataset, while the validation phase encompasses the remaining 40%. Our aim is to deal with mass events within the stadium premises over the event days, therefore, in the upper plot of Fig.11 we depict the maximum forecasting error experienced while estimating the resource consumption of each base station. From the picture, we can notice that for some base stations the resulting performances are poor, regardless the configuration settings of our model. This outcome can be easily explained as follows. Such considered base stations are generally underutilized in regular days. This leads the forecasting models to provide biased estimations of the cell capacity utilization, even in case of significant traffic loads generated during the sport events.

From the picture it can be noticed that an increase in the convolutional layer size yields a general performance degradation, both in terms of MSE and maximum forecasting error. Adding more layers may help to extract more features, but it also increases the number of parameters composing the model that, in turn, augments the chances of overfitted models. Empirically, we have identified the best performance with $\eta = 64$ and $\mu = 128$ settings. Hence, we adopt these values throughout our paper.

With the above neural network structure, our system learns the relationship between active users and the incurred load. Our goal however is to provide additional spectrum capacity such that the same QoS experienced by the users during regular days is preserved during the event. To this aim, we let $\Delta_b^e = \frac{U_b^e}{\overline{U}_b}$ denote a scaling factor, where \overline{U}_b^e is the predicted number of active users during epoch *e* for base station *b* (also marked as the output of our neural network) and \widetilde{U}_b is the number of users required to saturate resources of base station *b* according to a radio usage behavior from a regular day—this can be learned from previous history. The motivation behind is to compensate the expected radio resource usage with respect to the predicted impact that extra active users might have on individual QoS during the event. As a result, we finally devise the amount of spectrum required during the event as

$$\vec{c}_b = \vec{c}_b \cdot \vec{\Delta}_b^T, \qquad \forall b \in \mathcal{B}.$$
(12)

C. Performance evaluation

Hereafter, we evaluate ARENA considering the real traffic traces collected within the stadium neighbourhood. We leverage the neural network capability to approximate the relationship between QoE and PRB utilization as well as the temporal correlation of the measurements (as described in Section V) in order to measure the amount of additional resources required to cope with the traffic demand during



12



Fig. 12: Forecast of the number of active users or $RRC_{CONNECTED}$ (U_b) during a mass event (blue line) and actual evolution (dashed black line) for three sectors chosen for illustration purposes. Forecast on the radio resource usage (blue straight line), actual usage evolution (dashed black line) and estimation of extra radio resource allocation (green line) for the same sectors.

football events. To this aim, we set the time horizon from which our neural network starts making predictions to 2 hours before a real event and then we assess the ability of ARENA to predict the active number of users and spectrum usage after such time horizon. Based on this information, ARENA may take optimal spectrum allocation decisions. A complete snapshot of our validation results is depicted in Fig. 12. On the one hand, Figs. 12(a)-(c) (top three plots) show the evolution over time of the active users for three different sectors chosen for illustration purposes. The black dashed line represents the actual data whereas the blue straight line shows the predicted values. Like before, the area in light red highlights the period of time when the event takes place. Note that both lines overlap before the time horizon of the event as they represent known information from the history. We observe that our neural network structure follows the time evolution of active users during the event remarkably well thereby capturing the particularities of the football match as we have discussed in our analysis in Section IV. On the other hand, Figs. 12(d)-(f) (bottom three plots) show the evolution over time of the downlink spectrum utilization for the same sectors. Similarly, the dashed line represents the actual time series whereas the straight blue line shows the predicted values spawned by our neural network structure. In addition, we show the recommended PRB allocation with a straight green line (c.f. Eq. (12)), obtained varying the input QoS metric as in standard operating conditions. As it can be observed, our prediction on spectrum usage follows closely the real usage until the radio access network capacity limit is reached. In such cases, the recommended PRB allocation value (green line) follows the dynamics of the active users keeping the amount of time wherein extra capacity is required (PRB allocation above

100%) limited, thereby minimizing the incurred cost.

VII. CONCLUSIONS

In this paper we performed a data analysis of the network dynamics during sport events in the neighbourhood of a football stadium. The mobile infrastructure covers a 1-Km² area with 16 base station sectors and is operated by a major European carrier. The events contained traffic peaks generated by about 30.000 people.

Based on the insights obtained, we took up on the mass events RAN capacity forecasting challenge and designed ARENA: a model-free deep learning solution that predicts the expected number of connected users during a future event along with the corresponding RAN spectrum capacity (PRB) needed per sector for providing an average user experience. ARENA takes as input the average number of connected users, downlink traffic volume, downlink throughput and other contextual event information to infer the future traffic loads. Our results show that ARENA *i*) closely predicts the actual number of connected devices during mass events in time and *ii*) that it can provide mobile operators guidance on the actual RAN capacity needed during an event to provide a user experience similar to a regular network situation.

ACKNOWLEDGMENT

The research leading to these results has been supported in part by the H2020 MonB5G Project (grant agreement no. 871780), and in part by the H2020 5G Carmen (grant agreement no. 825012). We thank COSMOTE for providing the network measurements and for the support.

REFERENCES

- [1] A. Abdallah, A. B. MacKenzie, V. Marojevic, J. Kalliovaara, R. Bacchus, A. Riaz, D. Roberson, H. Juhani, and R. Ekman, "Detecting the Impact of Human Mega-events on Spectrum Usage," in 13th IEEE Annual Consumer Communications Networking Conference - CCNC'16, Jan. 2016, pp. 523–529.
- [2] L. Rabieekenari, K. Sayrafian, and J. S. Baras, "Autonomous relocation strategies for cells on wheels in environments with prohibited areas," in *IEEE International Conference on Communications - ICC'17*, May 2017.
- [3] K. Sakai, M. Sun, and W. Ku, "Data-Intensive Routing in Delay-Tolerant Networks," in *IEEE Conference on Computer Communications* - *INFOCOM'19*, Apr. 2019, pp. 2440–2448.
- [4] C. Xu, M. Wang, X. Chen, L. Zhong, and L. A. Grieco, "Optimal Information Centric Caching in 5G Device-to-Device Communications," *IEEE Transactions on Mobile Computing*, vol. 17, no. 9, pp. 2114–2126, 2018.
- [5] J. X. Salvat, L. Zanzi, A. Garcia-Saavedra, V. Sciancalepore, and X. Costa-Pérez, "Overbooking network slices through yield-driven endto-end orchestration," in *Proceedings of the 14th International Conference on emerging Networking EXperiments and Technologies*. ACM CoNEXT, Nov. 2018, pp. 353–365.
- [6] S. D'Oro, F. Restuccia, A. Talamonti, and T. Melodia, "The Slice Is Served: Enforcing Radio Access Network Slicing in Virtualized 5G Systems," in *IEEE Conference on Computer Communications -INFOCOM'19*, Apr. 2019, pp. 442–450.
- [7] A. Morrison, M. Bell, and M. Chalmers, "Visualisation of Spectator Activity at Stadium Events," in *International Conference on Information Visualisation*, Jul. 2009, pp. 220–226.
- [8] J. Erman and K. Ramakrishnan, "Understanding the Super-sized traffic of the Super Bowl," in *Proceedings of the 2013 Conference on Internet Measurement Conference - ACM IMC'13*, 2013, pp. 353–360.
- [9] P. Castagno, V. Mancuso, M. Sereno, and M. A. Marsan, "Why Your Smartphone Doesn't Work in Very Crowded Environments," in *IEEE In*ternational Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMOM), Jun. 2017, pp. 1–9.
- [10] A. S. Khatouni, M. Mellia, M. A. Marsan, S. Alfredsson, J. Karlsson, A. Brunstrom, O. Alay, A. Lutu, C. Midoglu, and V. Mancuso, "Speedtest-Like Measurements in 3G/4G Networks: The MONROE Experience," in 29th International Teletraffic Congress (ITC 29), vol. 1, 2017, pp. 169–177.
- [11] A. M. Mandalari *et al.*, "Experience: Implications of Roaming in Europe," in *International Conference on Mobile Computing and Networking*, ser. MobiCom '18. New York, NY, USA: Association for Computing Machinery, 2018, pp. 179–189.
- [12] D. Zhang, J. Huang, Y. Li, F. Zhang, C. Xu, and T. He, "Exploring Human Mobility with Multi-Source Data at Extremely Large Metropolitan Scales," in *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking - ACM MobiCom '14*, Sept. 2014, pp. 201—212.
- [13] X. Wang, Z. Zhou, F. Xiao, K. Xing, Z. Yang, Y. Liu, and C. Peng, "Spatio-Temporal Analysis and Prediction of Cellular Traffic in Metropolis," *IEEE Transactions on Mobile Computing*, vol. 18, no. 9, pp. 2190–2202, Sep. 2019.
- [14] X. Wang, Z. Zhou, Y. Zhao, X. Zhang, K. Xing, F. Xiao, Z. Yang, and Y. Liu, "Improving Urban Crowd Flow Prediction on Flexible Region Partition," *IEEE Transactions on Mobile Computing*, pp. 1–15, Aug. 2019.
- [15] F. Xu, Y. Li, H. Wang, P. Zhang, and D. Jin, "Understanding Mobile Traffic Patterns of Large Scale Cellular Towers in Urban Environment," *IEEE/ACM Transactions on Networking*, vol. 25, no. 2, pp. 1147–1161, Apr. 2017.
- [16] C. Marquez et al., "Not All Apps Are Created Equal: Analysis of Spatiotemporal Heterogeneity in Nationwide Mobile Service Usage," in Proceedings of the 13th International Conference on emerging Networking EXperiments and Technologies - ACM CoNEXT '17, Dec. 2017, pp. 180–186.
- [17] —, "How Should I Slice My Network? A Multi-Service Empirical Evaluation of Resource Sharing Efficiency," in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking* - ACM MobiCom '18, Oct. 2018, pp. 191—206.
- [18] J. Wang, J. Tang, Z. Xu, Y. Wang, G. Xue, X. Zhang, and D. Yang, "Spatiotemporal Modeling and Prediction in Cellular Networks: A Big Data Enabled Deep Learning Approach," in *IEEE Conference on Computer Communications - INFOCOM'17*, May 2017.

- [19] C. Zhang and P. Patras, "Long-Term Mobile Traffic Forecasting Using Deep Spatio-Temporal Neural Networks," in *Proceedings of the Eighteenth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, ser. Mobihoc '18, 2018, pp. 231–240.
- [20] Z. Pan, Y. Liang, W. Wang, Y. Yu, Y. Zheng, and J. Zhang, "Urban Traffic Prediction from Spatio-Temporal Data Using Deep Meta Learning," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019, pp. 1720—1730.
- [21] D. Bega, M. Gramaglia, M. Fiore, A. Banchs, and X. Costa-Perez, "DeepCog: Cognitive Network Management in Sliced 5G Networks with Deep Learning," in *IEEE Conference on Computer Communications -INFOCOM'19*, Apr. 2019, pp. 280–288.
- [22] M. Z. Shafiq, L. Ji, A. X. Liu, J. Pang, S. Venkataraman, and J. Wang, "A First Look at Cellular Network Performance during Crowded Events," in ACM SIGMETRICS '13, Jun. 2013, pp. 17–28.
- [23] Third Generation Partnership Project (3GPP), "Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Resource Control (RRC) Protocol specification," 3GPP TS 36.331 V15.4.0, Feb. 2019.
- [24] S. Sesia, I. Toufik, and M. Baker, *LTE, The UMTS Long Term Evolution: From Theory to Practice.* Wiley Publishing, 2009.
- [25] V. Sciancalepore, X. Costa-Perez, and A. Banchs, "RL-NSB: Reinforcement Learning-Based 5G Network Slice Broker," *IEEE/ACM Transactions on Networking*, vol. 27, no. 4, pp. 1543–1557, Aug. 2019.
- [26] T. Chen, Q. Ling, and G. B. Giannakis, "An Online Convex Optimization Approach to Proactive Network Resource Allocation," *IEEE Transactions on Signal Processing*, vol. 65, no. 24, pp. 6350–6364, Dec. 2017.
- [27] A. D. Flaxman, A. T. Kalai, and H. B. McMahan, "Online convex optimization in the bandit setting: gradient descent without a gradient," in *Proceedings of ACM SODA*, Jan. 2005, pp. 385–394.
- [28] J. C. Duchi, M. I. Jordan, M. J. Wainwright, and A. Wibisono, "Optimal rates for zero-order convex optimization: The power of two function evaluations," *IEEE Transactions on Information Theory*, vol. 61, no. 5, pp. 2788–2806, May 2015.
- [29] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, pp. 1735—1780, Nov. 1997.
- [30] X. Ma, Z. Dai, Z. He, and Y. Wang, "Learning Traffic as Images: A Deep Convolutional Neural Network for Large-Scale Transportation Network Speed Prediction," in *Sensors*, 2017.
- [31] X. Shi, Z. Chen, H. Wang, D.-Y. Yeung, W.-k. Wong, and W.-c. Woo, "Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting," in *International Conference on Neural Information Processing Systems - NIPS'15*, vol. 1, Dec. 2015, pp. 802– -810.
- [32] S. Vani and T. V. M. Rao, "An Experimental Approach towards the Performance Assessment of Various Optimizers on Convolutional Neural Network," in *International Conference on Trends in Electronics and Informatics - ICOEI'19*, Oct. 2019.
- [33] O. Ozyegen, S. Mohammadjafari, E. Kavurmacioglu, J. Maidens, and A. B. Bener, "Experimental Results on the Impact of Memory in Neural Networks for Spectrum Prediction in Land Mobile Radio Bands," *IEEE Transactions on Cognitive Communications and Networking*, vol. 6, no. 2, pp. 771–782, Jun. 2020.
- [34] B. Liao, J. Zhang, C. Wu, D. McIlwraith, T. Chen, S. Yang, Y. Guo, and F. Wu, "Deep Sequence Learning with Auxiliary Information for Traffic Prediction," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Jun. 2018, pp. 537—546.
- [35] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, http://www.deeplearningbook.org.
- [36] H. K. Vydana and A. K. Vuppala, "Investigative Study of Various Activation Functions for Speech Recognition," in *National Conference* on Communications - NCC'17, Mar. 2017.
- [37] G. Cybenko, "Approximations by Superpositions of Sigmoidal Functions," in *Mathematics of Control, Signals, and Systems*, vol. 2, no. 4, Mar. 1989, pp. 303–314.
- [38] G. Khodabandelou, V. Gauthier, M. Fiore, and M. A. El Yacoubi, "Estimation of Static and Dynamic Urban Populations with Mobile Network Metadata," *IEEE Transactions on Mobile Computing*, pp. 1–14, Sept. 2019.

BIOGRAPHIES



Lanfranco Zanzi received his B.Sc. and M.Sc. in Telecommunication Engineering from Politecnico of Milan (Italy) in 2014 and 2017, respectively. He is currently enrolled as Ph.D. candidate at the Technical University of Kaiserlautern, and works as research scientist at NEC Laboratories Europe. His research interests include network virtualization, machine learning, blockchain, and their applicability to 5G mobile networks in the context of network slicing.



Xavier Costa-Pérez (M'06–SM'18) is Head of Beyond 5G Networks R&D at NEC Laboratories Europe, Scientific Director at the i2Cat R&D Center and Research Professor at ICREA. His team contributes to products roadmap evolution as well as to European Commission R&D collaborative projects and received several awards for successful technology transfers. In addition, the team contributes to related standardization bodies: 3GPP, ETSI NFV, ETSI MEC and IETF. Xavier has been a 5GPPP Technology Board member, served on the Program

Committee of several conferences (including IEEE Greencom, WCNC, and INFOCOM), published at top research venues and holds several patents. He also serves as Editor of IEEE Transactions on Mobile Computing and Transactions on Communications journals. He received both his M.Sc. and Ph.D. degrees in Telecommunications from the Polytechnic University of Catalonia (UPC) in Barcelona and was the recipient of a national award for his Ph.D. thesis.



Vincenzo Sciancalepore (S'11–M'15–SM'19) received his M.Sc. degree in Telecommunications Engineering and Telematics Engineering in 2011 and 2012, respectively, whereas in 2015, he received a double Ph.D. degree. Currently, he is a senior 5G researcher at NEC Laboratories Europe in Heidelberg, focusing his activity on network virtualization and network slicing challenges. He is currently involved in the IEEE Emerging Technologies Committee leading the initiatives on SDN and NFV. He is the Chair of the Emerging Technologies Initiative (ETI)

on Reconfigurable Intelligent Surfaces (RIS). He was also the recipient of the national award for the best Ph.D. thesis in the area of communication technologies (Wireless and Networking) issued by GTTI in 2015.



Georgios Agapiou is Head of Core Research & Network Testing at OTE S.A. and COSMOTE S.A. in Athens. He has worked in several European projects in the area of Mobile and optical Communications, including various IST, STREP, e-ten, Eurescom, FP6, FP7, ESA, H2020, and in more than 15 5G-PPP projects. He is author of more than 180 scientific publications in peer-reviewed journals and conferences, as well as co-author of three technical books. He has also been the Scientific Coordinator of FITCE Greece Association. He received his Diploma

in Electrical Engineering from the University of Louisville, Kentucky, in 1985, and the M.S. and Ph.D. degrees in Electrical Engineering from the Georgia Institute of Technology, in 1987 and 1991, respectively. His main interest and competence include testing of mobile wireless systems and the development of simulation radio systems in the R&D facilities of OTE.



Andres Garcia-Saavedra received the MSc and PhD degrees from the University Carlos III of Madrid (UC3M), in 2010 and 2013, respectively. He then joined the Hamilton Institute, Ireland, as a research fellow till the end of 2014 when he moved to Trinity College Dublin (TCD). Since July 2015, he is a senior researcher with NEC Laboratories Europe. His research interests lie in the application of fundamental mathematics to real-life communications systems and the design and prototype of wireless systems and protocols.



Hans Dieter Schotten received his Diploma and Ph.D. degrees in electrical engineering from Aachen University of Technology RWTH, Germany, in 1990 and 1997, respectively. Since August 2007, he is a full professor and head of the Institute of Wireless Communication and Navigation at Technische Universität Kaiserslautern. Since 2012, he is also scientific director at the German Research Center for Artificial Intelligence, heading the Intelligent Networks department.