

Smooth Handovers via Smoothed Online Learning

Michail Kalntis^{*}, Andra Lutu[†], Jesús Omana Iglesias[†], Fernando A. Kuipers^{*}, George Iosifidis^{*}

^{*}Delft University of Technology, The Netherlands. Emails: {m.kalntis, f.a.kuipers, g.iosifidis}@tudelft.nl

[†]Telefónica Research, Spain. Emails: {andra.lutu, jesusalberto.omana}@telefonica.com

Abstract—With users demanding seamless connectivity, handovers (HOs) have become a fundamental element of cellular networks. However, optimizing HOs is a challenging problem, further exacerbated by the growing complexity of mobile networks. This paper presents the first countrywide study of HO optimization, through the prism of Smoothed Online Learning (SOL). We first analyze an extensive dataset from a commercial mobile network operator (MNO) in Europe with more than 40M users, to understand and reveal important features and performance impacts on HOs. Our findings highlight a correlation between HO failures/delays, and the characteristics of radio cells and end-user devices, showcasing the impact of heterogeneity in mobile networks nowadays. We subsequently model UE-cell associations as dynamic decisions and propose a realistic system model for smooth and accurate HOs that extends existing approaches by (i) incorporating device and cell features on HO optimization, and (ii) eliminating (prior) strong assumptions about requiring future signal measurements and knowledge of end-user mobility. Our algorithm, aligned with the O-RAN paradigm, provides robust dynamic regret guarantees, even in challenging environments, and shows superior performance in multiple scenarios with real-world and synthetic data.

I. INTRODUCTION

Motivation. Mobility management has been a prime consideration for every generation of mobile networks, and occupies a conspicuous position in the agendas of industry and academia [1], [2]. However, as the proposed solutions evolve, so does the problem. On the one hand, users and their devices (user equipments, i.e., UEs) are becoming increasingly mobile and expect higher-data rates and lower-latency services. On the other hand, today's networks are more heterogeneous than ever, encompassing different Radio Access Technology (RAT) generations, multi-tier cells (some with small coverage) / base stations (BSs) operating a wide range of frequencies, and equipment from multiple vendors. This heterogeneity and complexity compound mobility management issues and, in particular, perplex the UE-BS association decisions and the induced handovers (HOs).

HOs are fundamental elements for enabling seamless connectivity for mobile users. Optimizing this procedure is crucial, as inefficient or suboptimal HOs can have a dire twofold impact: on the network-side, they might lead to increased resource consumption and signaling overhead [3]; and on the UE-side, they may cause service interruptions and battery depletion [4], [5]. Thus, latest proposals depart from conventional signal-to-noise ratio (SNR) or signal-to-interference-plus-noise ratio (SINR) based association rules, towards solutions where UE-BS associations are based on network-wide criteria, e.g., aggregate throughput or fairness/load-balancing metrics, while also aiming to reduce HOs whenever possible [6]–[9]. The recent Open Radio Access Network (O-RAN) paradigm [10]

TABLE I: Mobility Dataset Size for European Country.

# of cell sites, # of base stations	> 26k, > 370k
# of UEs measured	≈ 40M
# handovers (daily)	> 1.7B
Measurement duration	7 days (≈ 8 TB/day)

facilitates such holistic approaches and enables central *controllers* to coordinate the network operations dynamically [11], [12], including how the UE associations are implemented.

In this context, the goal of this paper is to design a novel UE association and smooth HO control mechanism that is both *effective* and *robust*, thus overcoming key limitations of the above works (see discussion in Sec. II). The effectiveness alludes to achieving a network utility maximization criterion with guarantees under a wide range of conditions, despite the heterogeneity of UEs and BSs. The robustness, on the other hand, refers to ensuring this performance without the controller requiring access to accurate information about the UEs/network-related parameters, such as their mobility patterns. These two main features of our proposal fill a key gap in the literature and pave the road for the next generation of mobility management solutions.

Methods & Contributions. To understand HOs and their effects in practice, we kick off our study with the analysis of a 1-week capture of all mobility events from a large mobile network operator (MNO) in a European country (see Table I). The analysis of this vast dataset sheds light on the frequency, features, and performance of HOs. For instance, we find that HO failures and delays depend heavily on the RAT generation, and are affected by the UE type. In other words, *HO events vary widely in terms of completion delay*. We also reveal the network complexity in terms of RATs and load in this top-tier operator and discuss the implications of this heterogeneity in terms of HOs. Finally, we measure the effect of HOs on several network and user key performance indicators (KPIs), such as packet loss and throughput. To the best of our knowledge, this is the first *network-side HO optimization and measurement* study that overcomes the limitations of prior UE-side studies, or large network-side measurement-only HO study [13].

Building on these findings, we introduce a realistic system model that captures the impact of the different HOs and mitigates them while maximizing the network throughput. Our model aligns with recent works [6], [7], [9], [11], [12], [14], which we extend substantially by accounting for the network and HO diversity, and importantly, by dropping requirements for access to future SINRs and UE mobility patterns. It is commonly accepted that this is a strict condition that limits the applicability of such solutions. Additionally, we do not assume that the relevant UE/BS parameters are stationary-perturbed since rapid and unpredictable channel fluctuations

are becoming increasingly common in heterogeneous networks and mobile services. In fact, our perturbation model is an adversarial one, where the various random parameters can even be selected by an attacker; still, all results and guarantees hold.

With this in mind, we turn the problem on its head and study HOs through the lens of online convex optimization (OCO) [15]. Namely, we model the UE-BS associations as dynamic decisions that the network controller updates in a time-slotted fashion, where successive (de)associations induce undesirable (sometimes necessary) HOs. We argue a natural framework for this setting is that of *smoothed online learning* (SOL) [16], which maximizes a performance criterion while reducing the decision changes. Using such a framework enables the controller to be oblivious to the SINRs for each UE-BS pair when deciding the associations, and the throughput these will achieve; indeed, it is challenging to predict accurately or know the SINRs over a time window of several msec/sec [17]. Still, following a rigorous analysis, we show that our learning algorithm ensures *sublinear dynamic regret*, i.e., its gap w.r.t. an ideal oracle that has full information about the future diminishes with time [18]. Our model is informed by, and aligned with, the O-RAN paradigm, and the proposed algorithm can be implemented as xApp in *near-real-time* (running time 10msec–1sec) [10]; the reader is kindly referred to [17], [19] for the involved interfaces and detailed steps. Finally, to verify the robustness of our proposed solution, we evaluate its performance on simple and extreme (i.e., adversarial) synthetic scenarios in accordance with related work, as well as in real scenarios using signal measurement from crowdsourced data (i.e., measured on the field).

In summary, the contributions of this work are:

- We present HO and network statistics by using a 1-week dataset from a tier-1 MNO. We highlight the network heterogeneity and identify key factors impacting smooth HOs. We are the first to optimize HOs from the MNO’s perspective, which sets the basis for realistic HO optimization models.
- We model the HO optimization as a *smoothed online learning problem* where the HO delays depend on the RAT and UE type, and assume no prior information for the channels and mobility patterns. This approach departs from related work and its (often simplifying) modeling assumptions.
- We design a scalable (near-real-time) algorithm that achieves sublinear dynamic regret, and we characterize its performance w.r.t. system parameters. We also propose extensions for the case of time-varying HO delays and the case of available (untrusted) forecasting tools.
- We create a simulator using our real, crowdsourced radio signal quality measurements, as well as actual cell and UE information, and evaluate our solution against meaningful benchmarks; e.g., we find up to $\times 79.6$ lower HO cost than previous works, without sacrificing throughput.

II. BACKGROUND & RELATED WORK

Measurements. HOs have mainly been studied using traces from UEs, which are inevitably limited to certain manufacturers [4], services [20], or user types [21], [22]. Our findings on HO duration align with these prior works [4], [21], [23] and offer new insights into the impact of RAT and UE type on

HOs. While several studies have measured HO volume [20], [24], [25], noting that horizontal HOs are more frequent in 5G-Standalone (5G-SA) and 4G, and vertical HOs in 5G-NSA [20], we enhance these findings by revealing the geographic heterogeneity of HOs and measuring their effect on important KPIs. Importantly, our work utilizes a large-scale dataset measuring HOs from the network-side of a top-tier operator with 40M UEs. This contrasts sharply with the few network-side studies, e.g., [5] has less than 0.7% of our UEs. While [13] is the first countrywide study of cellular HOs, it focuses solely on measurements, without providing a HO solution or focusing on the effect of HO failures (HOFs) on KPIs.

Handover Optimization. In user-centric mobility schemes, UE associations (and thus HOs) are often approached as a bandit learning problem where each UE explores which BS offers better throughput [26]–[31]. Conversely, using SINR measurements from all UEs, the network can make more effective centralized decisions. This idea is regaining momentum, but is also more challenging (due to scale, among others) to implement. For instance, a recent thread of studies employs (Deep) Reinforcement Learning to decide HOs or HO rules (e.g., SINR thresholds) [7]–[9], [11], [32], [33]. Despite its modeling appeal, these approaches essentially rely on heuristics. On the other hand, our method comes with performance guarantees, even under adversarial conditions.

Similarly, the joint optimization of throughput and HO delays is considered in [6], [7], [9], [11], [14]. Our goal is very close to these works, but we enrich their model to capture the different delay effects of the various HO types, RATs, and UE types. Also, we drop several of their assumptions and, in particular, their need to have accurate SINR information. Clearly, in today’s volatile and non-stationary network conditions, this assumption is impractical. Finally, some recent papers leverage the ML-provisions of O-RAN, and propose forecasting-assisted HOs [6], [11], [12], [34]. Our solution is also aligned with O-RAN, but unlike these works, it learns on the fly the network/UE parameters and does not require offline training nor the availability of a reliable forecaster. At the same time, when such predictions are available, our algorithm can directly benefit from them to expedite its learning, building on the idea of optimistic learning [35]. For additional discussion on HOs, see [1], [2].

Smoothed Online Learning. In terms of solution, unlike the model-predictive control of [6], the RL strategies in [7], [9], [11], [36], or the heuristics in [14], we approach this problem, for the first time, through the lens of *smoothed online learning* [16], [18], [37]. SOL enriches the online convex optimization toolbox [15], accounting for costs induced by decision changes, and has recently found applications in caching [38], [39], network selection [28], and service deployment [40], among others. Here, the decisions express the UE-BS associations, and the switching cost captures the HO delays in a natural way. Yet, unlike the above works, we leverage the more competitive dynamic regret benchmark, since HOs are unavoidable, and the switching cost models the delay rather than the mere HO count. We also note that SOL should not be confused with mathematical smoothing [41], nor with the stochastic and full-information framework




Fig. 1: Median normalized DL packet loss (left, green y-axis) and normalized average UE throughput (right, orange y-axis) vs binned daily HOFs. Triangles and bold horizontal lines show the mean and median, respectively, in each boxplot.

of decisions with reconfiguration delays [42].

III. DATA COLLECTION & ANALYSIS

Data collection. To show the network’s heterogeneity and pinpoint key factors affecting seamless HOs (and thus, are essential for our system model), we collect passive measurements using commercial tools available within the MNO’s infrastructure. For that, we focus on elements in different RATs, such as the Mobile Management Entity (MME), Mobile Switching Center (MSC), Serving GPRS Support Node (SGSN), Serving Gateway (SGW), and cell sites. The captured data spans the entire country from 24-June-2024 to 30-June-2024 (Table I) and includes the (i) HO timestamp, (ii) HO duration (msec granularity), (iii) anonymized user ID, based on the International Mobile Subscriber/Equipment Identity (IMSI/IMEI), and (iv) source and target BS of the HO along with their respective RATs. At the time this study was conducted, the majority of the traffic for the studied MNO used the 5G-Non-Standalone (NSA) deployment [43]. From a mobility management perspective, 5G-NSA and 4G are identical, since the former relies on the 4G Evolved Packet Core (EPC); thus, we use the term “4G/5G-NSA”.

In addition, we leverage (i) a daily-captured dataset from the MNO that contains cell information such as the longitude, latitude, vendor, and supported RATs, (ii) a daily commercial device catalog, provided by the Global System for Mobile Communications (GSM) Association (GSMA) to examine correlations of user-specific characteristics with HOs, (iii) open datasets published by the official census office in the studied European country to explore the distribution of HOs across the different geographic areas (mainly the 300+ districts defined by the census office), and (iv) commercial crowdsourced data with msec granularity of radio signal measurements to assess the performance of our algorithm.

Effect of HOs on KPIs. We first study how HOs and HOFs affect some important network KPIs, used by the operational team to determine the health of the network and reason about the quality of the service to the end-users. For confidentiality, we normalize each KPI by its median from all days. Also, we discard the outliers, i.e., cells with daily HOFs > 140 (0.1% of data). Fig. 1 depicts the daily HOFs per cell and their impact on downlink (DL) packet loss and user throughput per day. We observe a decrease in normalized average user throughput with increasing HOFs; e.g., (0, 20] HOFs lead to a normalized mean throughput of 1.25k (normalized value, i.e.,




Fig. 2: **Left:** number (norm. by max) of different UE models and their RAT capabilities (up to 2G, 3G, 4G, 5G). **Right:** HOs (norm. by max) per day each of the UE model executes, and to what RAT.

does not mean Kbps or Mbps), which declines by 37.6% to 0.78k for (120, 140] HOFs per day. Similar studies (e.g., [21], [24], [25]) have reported UE throughput but only for specific UE types (smartphones) and manufacturers. Simultaneously, the normalized median DL packet loss rises from a mean of around 0.002 to 0.013 when daily HOFs increase from (0, 20] to (120, 140] respectively. The higher mean compared to the median in each boxplot indicates a few UEs with larger losses compared to the majority.

We verified the findings of Fig. 1 with a generalized linear regression model using as dependent variable the DL packet loss rate and as independent variables the number of HOs and HOFs while controlling for all key covariates; namely, the HO type (inter/intra RAT), district population, BS type/vendor/transmission power, and the district. This ensures the packet loss effects are indeed due to HOs and HOFs, and not to some latent factor. We find that 1% increase of HOs in a cell, increases by 0.02% the packet loss rate, on average, for the served users; and even worse, 1% increase in HOFs increases by 0.6% this loss. A similar model found that, all else being equal (including total uplink/DL volume and physical resource blocks), a HOF increase of 1% reduces by 0.008% the average user throughput. And this drop is more pronounced in cells with few HOFs, as Fig. 1 shows.

Network & UE Heterogeneity. Using the GSMA devices database, we discern the eight most crucial *UE types* (dongle, IoT, feature phone, modem, smartphone, tablet, WLAN router and wearable), and infer their connectivity capabilities (i.e., up to what RAT they support) from their frequency bands. Given the dominance of 4G/5G capable models (98.5%), depicted in Fig. 2, we study the HOs from 4G/5G-NSA to the same or older RATs. In these *HO types* (i.e., Intra 4G/5G-NSA, 4G/5G-NSA to 3G, 4G/5G-NSA to 2G), approximately 94% are Intra 4G/5G-NSA, caused by a wide range of devices. HOs to 3G and 2G hold an important 6%, mainly from smartphones, modems, and tablets, magnifying the heterogeneity in these dimensions as well. As we show in the sequel, the HO delay in the older RATs is $\times 5$ –40 higher. Network heterogeneity is also extensively analyzed in [13].

Fig. 3 and Fig. 4 illustrate the histogram and probability density of different HOs and device types, respectively, w.r.t. the HO delay. From Fig. 3, we observe that the median Intra 4G/5G-NSA HOs require approximately 50msec, whereas HOs to 3G range from 400 to 950msec and to 2G from 750 to 1100msec; thus, especially in the older RATs, the




Fig. 3: Histogram (bars) and distribution (line) of the HO delays for each HO type (all same-colored bars sum to 1).

distribution is significantly more spread. Moreover, from Fig. 4, we deduce the different distributions of the HO delays for each device type. It is interesting to observe the different mean and variances of each UE type; for instance, modems and IoT devices require on average 75msec and 73msec respectively, but around the same number of these models start from 50msec and reach 110msec; smartphones, on the other hand, need 50–62 msec. Consequently, *accounting both for the UE type and HO type is essential when optimizing for the HO delay.*

IV. SYSTEM MODEL & PROBLEM STATEMENT

We consider a heterogeneous cellular network comprising a set \mathcal{J} of J base stations serving a set \mathcal{I} of I users (UEs) in the downlink. Each BS is characterized by an array of features such as its operating frequency, RAT generation (2G-5G), location, etc. Similarly, the set of UEs comprises smartphones, feature phones, IoT devices, and so on. The network is managed centrally by a network controller, in the spirit of recent O-RAN architecture proposals [10]. The system operation is time-slotted where we index the time slots with t and without loss of generality assume the slots have unitary length. These slots refer to the UE association intervals, which subsume other resource scheduling time-slots (e.g., for power control). We study the system for a set \mathcal{T} of T slots. The key metric for the association decisions is the SINR for the signal delivered by BS j to user i :

$$s_{ij}(t) = \frac{p_j \phi_{ij}(t)}{\sum_{k \in \mathcal{B}_j} p_k \phi_{ik}(t) + w_k \sigma^2}, \quad (1)$$

where p_j is the transmit power of BS j , \mathcal{B}_j the set of BSs that operate in the same frequency as j , $\phi_{ij}(t)$ the channel gain (including pathloss, shadowing, and antenna gains), w_k the bandwidth of BS k , and σ^2 the power spectral density. In line with previous works [6], [7], [9], [14], $s_{ij}(t)$ is the average SINR over the slot t (the UEs report multiple SINRs during each slot [6]).

The rate that each user $i \in \mathcal{I}$ associated with BS $j \in \mathcal{J}$ achieves during slot t , can be expressed as:

$$r_{ij}(t) = \frac{c_{ij}(t)(1 - d_{ij}(t))}{y_j(t)}, \quad (2)$$

where $y_j(t)$ is the total number of UEs that BS j needs to serve during t , and $c_{ij}(t) = w_j \log_2(1 + s_{ij}(t))$ is the maximum possible rate for i if it was using BS j exclusively. The rate $r_{ij}(t)$ is discounted by the service disruption time $d_{ij}(t)$ which models any HO delay associated with the assignment of UE i to BS j . Clearly, this delay is negligible if i was already associated with BS j ; is pronounced when the target BS is




Fig. 4: Histogram (bars) and distribution (line) of the HO delays for each UE model (all same-colored bars sum to 1).

of different RAT; and is even larger when there is a HO failure (see Sec. III). We normalize these parameters, i.e., $d_{ij}(t) \leq 1, \forall i, j$, to express the *portion* of slot the UE was not receiving service. Equation (2) assumes the BS resources are allocated fairly across the active users via, e.g., a round robin or a proportional-fair scheduler [44]. As it will become clear, our analysis holds for different scheduling schemes.

We denote with $x_{ij}(t) \in \{0, 1\}$ the association of user i with BS j in the beginning of slot t and define the vector $\mathbf{x}_t = (x_{ij}(t) \in \{0, 1\}, i \in \mathcal{I}, j \in \mathcal{J})$. Then, the problem the network controller wishes to solve can be expressed as:

$$\mathbb{P}: \max_{\{\mathbf{x}_t\}_t} \sum_{t=1}^T \sum_{i=1}^I \sum_{j=1}^J x_{ij}(t) \log r_{ij}(t) \quad (3)$$

$$\text{s.t.} \quad \sum_{j \in \mathcal{J}} x_{ij}(t) = 1, \quad \forall i \in \mathcal{I}, \quad (4)$$

$$y_j(t) = \sum_{i \in \mathcal{I}} x_{ij}(t), \quad \forall j \in \mathcal{J}, \quad (5)$$

$$\mathbf{x}_t \in \{0, 1\}^{I \cdot J}, \quad \forall t \in \mathcal{T}. \quad (6)$$

The logarithmic utility function is selected to maximize the product of rates so as to balance network sum-rate and fairness [6], [7], [45]; constraint (4) ensures each UE is assigned to one BS, and (5) calculates the assigned UEs to each BS.

The solution to this problem at the beginning of horizon T is hindered by several factors. First, at $t = 1$, the controller does not have access to all future SINR values for each UE-BS pair. In fact, the SINRs during each slot t are practically unknown *even at the beginning* of that slot. Secondly, the HO delays $\{d_{ij}(t)\}$ depend on the associations \mathbf{x}_t , but also on the previous associations \mathbf{x}_{t-1} , since these two vectors determine if there is a HO or not; a HO is triggered for user i if $\mathbf{x}_{ij}(t-1) \neq \mathbf{x}_{ij}(t)$. In other words, there is a memory effect in the system, thus the problem cannot be decomposed on per-slot basis. With these challenges in mind, our goal is to design an online association algorithm that is oblivious to these time-varying unknown parameters, and which nevertheless maximizes the throughput while minimizing the HO delays, compared to a meaningful (i.e., competitive) benchmark.

V. ALGORITHM DESIGN AND ANALYSIS

Reformulation & Benchmark. We approach \mathbb{P} as a *smoothed online learning* problem and tackle it via meta-learning based on the *experts framework* [46], [47]. The main idea is to deploy a set of parallel learning algorithms with different learning rates, and a meta-learner that tracks their performance

and discerns on-the-fly the best-performing one. We start by reformulating the objective:

$$\begin{aligned}
f_t(\mathbf{x}_t) &\triangleq \sum_{i=1}^I \sum_{j=1}^J \left[x_{ij}(t) \log c_{ij}(t) - x_{ij}(t) \log y_j(t) \right. \\
&\quad \left. + x_{ij}(t) \log (1 - d_{ij}(t)) \right] \stackrel{(\alpha)}{=} \sum_{i=1}^I \sum_{j=1}^J x_{ij}(t) \log c_{ij}(t) \\
&\quad - \sum_{j=1}^J y_j(t) \log y_j(t) + \sum_{i=1}^I \sum_{j=1}^J x_{ij}(t) \log (1 - d_{ij}(t))
\end{aligned}$$

where (α) follows from $y_j(t) = \sum_{i \in \mathcal{I}} x_{ij}(t), \forall j, t$. The last term corresponds to the performance cost due to HO delays. Extending the rationale of prior works [6], [9], [14], and based on our data analysis, we will capture this cost using:

$$h(\mathbf{x}_t, \mathbf{x}_{t-1}) = -\gamma \|\mathbf{A}(\mathbf{x}_t - \mathbf{x}_{t-1})\| \triangleq -\gamma \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_A$$

where $\mathbf{A} = \text{diag}(a_n > 0, n = 1, \dots, I \cdot J)$ is a positive definite matrix where each element $a_n, n = i \cdot j$ models the delay when UE i is associated (de-associated) to (from) BS j , and $\|\cdot\|_A$ is its induced norm, i.e., $\|\mathbf{x}\|_A^2 = \sum_n a_n x_n^2$ and $\|\mathbf{x}\|_{A^*}^2 = \sum_n x_n^2 / a_n$ [48]. For instance, for $I=3$ and $J=2$, we have: $\|\mathbf{x}_t - \mathbf{x}_{t-1}\|_A^2 = \sum_{i=1}^3 \sum_{j=1}^2 a_{ij} (x_{ij}(t) - x_{ij}(t-1))^2$; if UE $i=1$ moves from BS $j=1$ to $j=2$, it yields a total HO delay $a_{11} + a_{12}$, where a_{11} (a_{12}) is the delay due to the de-association (association) from (to) $j=1$ ($j=2$). The scalarization parameter γ is used to normalize units and prioritize one criterion (i.e., throughput) over the other (i.e., HO delays), based on the preferences of each operator.

This HO model departs from previous works, e.g., [6], [9] and references therein, that merely count the association changes as if all HOs had the same impact on performance. Clearly, the analysis in Sec. III showed this assumption not to be accurate in today's heterogeneous networks. Hence, we opt here instead to modulate the HO costs with parameters reflecting the potentially-different HO time for each UE-BS pair, and additionally with the tunable γ weight. In particular, based on Sec. III, we will be using in our numerical evaluations the BS RAT and UE type as the main features for the HO delays. Finally, we define the decision set:

$$\mathcal{X} = \left\{ \mathbf{x} \in \{0,1\}^{I \cdot J} \mid \sum_{j \in \mathcal{J}} x_{ij} = 1, i \in \mathcal{I} \right\},$$

and its convex hull $\mathcal{X}^c = \text{co}(\mathcal{X})$ that relaxes the integrality, i.e., $\mathbf{x} \in [0,1]^{I \cdot J}$. We will use \mathbf{x} when referring to the discrete associations, and denote with $\mathbf{x}^m \in \mathcal{X}^c$ the respective relaxed vector. Putting these together, we have the next result.

Proposition 1. The throughput and HO cost function is concave on \mathcal{X}^c : $f_t(\mathbf{x}) \triangleq g_t(\mathbf{x}) + h(\mathbf{x}, \mathbf{x}_{t-1}) =$

$$\sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} x_{ij}(t) \log c_{ij}(t) - \sum_{j \in \mathcal{J}} y_j(t) \log y_j(t) - \gamma \|\mathbf{x} - \mathbf{x}_{t-1}\|_A.$$

The concavity of $g_t(\mathbf{x})$ for $\mathbf{x} \in \mathcal{X}^c$ is proved in [6], and it follows that subtracting the A-norm preserves the property.




Fig. 5: Learning mechanism.

The performance of the algorithm will be assessed using the *Expected Dynamic Regret*, defined as:

$$\mathbb{E}[\mathcal{R}_T] \triangleq \sum_{t=1}^T f_t(\mathbf{x}_t^*) - \sum_{t=1}^T \mathbb{E}[f_t(\mathbf{x}_t)], \quad (7)$$

where $\{\mathbf{x}_t\}_t$ are the algorithm decisions, $\{\mathbf{x}_t^*\}_t$ is the benchmark, and the expectation captures any randomization in the algorithm. Our goal is to design an algorithm that ensures this gap diminishes with time, $\lim_{T \rightarrow \infty} \mathbb{E}[\mathcal{R}_T]/T = 0$ for any possible benchmark sequence $\{\mathbf{x}_t^*\}_t$. In other words, we compare our algorithm with the best oracle that has full information at $t=1$ for the SINR and HO delays, for all users and BSs, for the entire horizon T . Clearly, this is a very competitive benchmark, going beyond static and (per-slot) dynamic regret (see discussion in [49]), and thus a sublinear-regret algorithm in this context is highly desirable.

Online Algorithm. The proposed learning mechanism is summarized in Fig. 5. There is a meta-learner that receives suggestions for the association policy from K agents (the *experts*); creates accordingly a weighted policy; and learns gradually how much to trust each expert based on its performance. In turn, each expert learns dynamically the best association strategy, based on feedback from the meta-learner. The experts use different learning steps $\theta = (\theta_k, k \in \mathcal{K})$, that are selected so as to ensure at least one of them will perform optimally w.r.t. the yet-to-be-encountered problem conditions. The details of the mechanism are provided in Algorithm 1 (LDA). At the beginning of each slot, each expert k shares its suggestion \mathbf{x}_t^k , and the meta-learner synthesizes them:

$$\mathbf{x}_t^m = \sum_{k \in \mathcal{K}} w_t^k \mathbf{x}_t^k, \quad (8)$$

where the weights $\mathbf{w}_t = (w_t^k, k \in \mathcal{K})$, with $\mathbf{w}_t^\top \mathbf{1}_K = 1$, is what the meta-learner needs to learn. It follows that if $\mathbf{x}_t^k \in \mathcal{X}^c, \forall k$, then $\mathbf{x}_t \in \mathcal{X}^c$. Next, the meta-learner creates a binary decision vector $\mathbf{x}_t \in \mathcal{X}$ (so as to be implementable) using the quantization routine $Q_{\mathcal{X}}$. For this operation, one can use any unbiased sampling technique, as long as $\mathbb{E}[\mathbf{x}_t] = \mathbf{x}_t^m$. For instance, Madow's sampling [50] ensures this condition which, for the structure of \mathcal{X} , simply picks an element from $\mathbf{x}_i(t)$ with probabilities $\mathbf{x}_i^m(t)$, for each $i \in \mathcal{I}$.

At the end of the slot, the controller observes the system parameters¹ and calculates the gradient $\nabla g_t(\mathbf{x}_t)$, which is sent to all experts. Then, the meta-learner updates the weights:

$$w_{t+1}^k = \frac{w_t^k e^{\beta \ell_t(\mathbf{x}_t^k)}}{\sum_{k \in \mathcal{K}} w_t^k e^{\beta \ell_t(\mathbf{x}_t^k)}} \quad (9)$$

¹Since UEs report their channel gains with all reachable base stations (and not only with the one they are associated), we have a full-information model.

Algorithm 1: Learning Dynamic Associations (LDA)

```

1 Input: Horizon  $T$ ;  $K$  experts with steps  $\{\theta_k\}_{k \in \mathcal{K}}$ ;
   Meta-learning step  $\beta$ 
2 Initialize: set  $w_1^k = \frac{1+(1/K)}{k(k+1)}, \forall k \in \mathcal{K}$ ; draw
    $\mathbf{x}_1^k \in \mathcal{X}, \forall k \in \mathcal{K}$ ; sort  $\theta_1 \leq \theta_2 \leq \dots \leq \theta_K$ ;
3 for  $t = 1, 2, \dots, T$  do
4   Each expert  $k \in \mathcal{K}$  shares its decision  $\mathbf{x}_t^k$ ;
5   The controller synthesizes  $\mathbf{x}_t^m$  using (8);
6   The controller implements  $\mathbf{x}_t = Q_{\mathcal{X}}(\mathbf{x}_t^m)$ ;
7   The controller sends  $\nabla g_t(\mathbf{x}_t)$  to experts;
8   The controller updates its weights using (9);
9   Each expert updates its decision using (11);
end

```

using the surrogate (i.e., partially linearized) loss:

$$\ell_t(\mathbf{x}_t^k) = \langle \nabla g_t(\mathbf{x}_t), \mathbf{x}_t^k - \mathbf{x}_t \rangle - \gamma \|\mathbf{x}_t^k - \mathbf{x}_{t-1}^k\|_A. \quad (10)$$

It is interesting to observe that all experts use the same gradient and not the gradient of their own action. This is possible due to the concavity of g_t , which can be upper bounded by the linearized loss at *any* information point; see also [37].

Finally, the experts perform an online gradient ascent:

$$\mathbf{x}_{t+1}^k = \Pi_{\mathcal{X}}\left(\mathbf{x}_t^k + \theta_k \nabla g_t(\mathbf{x}_t)\right), \quad (11)$$

so as to produce their next suggestion. The number of experts, their rates, initial weights \mathbf{w}_1 , and meta-learning step β , are selected to ensure the regret convergence, as explained next.

Performance Analysis. We first introduce some key parameters that are used below. Defining $a_{\max} = \max_{n \leq I \cdot J} a_n$ we obtain the bounds:

- $\|\mathbf{x} - \mathbf{x}'\|_2 \leq D$, $\|\mathbf{x} - \mathbf{x}'\|_A \leq \sqrt{a_{\max}} D \triangleq D_A$, $\|\mathbf{x} - \mathbf{x}'\|_{A^*} \leq D/\sqrt{a_{\max}} \triangleq D_{A^*}$ with $D = \sqrt{2I}$.
- $\|\nabla g_t(\mathbf{x})\|_2 \leq G$, $\|\nabla g_t(\mathbf{x})\|_A \leq \sqrt{a_{\max}} G \triangleq G_A$, with $G = \sqrt{IJ(\log J + 1/\ln 10)^2}$.

To characterize the performance of Algorithm LDA, we proceed in two steps. First, we bound the regret of the relaxed (continuous) performance and switching cost, with the following Lemma that we prove in the appendix:

Lemma 1. Using the following parameters:

- $K = \lceil \log_2 \sqrt{1+2T} \rceil + 1$.
- $\theta_k = 2^{k-1} \left[\frac{D_A^2}{T(G^2 + 2G_A)} \right]^{1/2}$, $k = 1, \dots, K$.
- $\beta = 1/\sqrt{T\nu}$, with $\nu \triangleq (2GD + D_A)^2(D_A + (1/8))$.

The continuous decisions $\{\mathbf{x}_t^m\}_t$, where $\mathbf{x}_t^m \in \mathcal{X}^c$ ensure:

$$\sum_{t=1}^T f_t(\mathbf{x}_t^*) - \sum_{t=1}^T f_t(\mathbf{x}_t^m) \leq \mathcal{O}\left(\sqrt{T(1+P_T)}\right) \quad (12)$$

with the benchmark's total HO delay $P_T = \sum_{t=1}^T \|\mathbf{x}_t^* - \mathbf{x}_{t-1}^*\|_{A_t}$.

As is expected in dynamic regret, the bound depends on the variability of the benchmark which here, interestingly, has the physical meaning of HO delays. In any case, the meta-learner follows the benchmark and gradually decreases the gap. These decisions refer to a continuous-valued association strategy $(\mathbf{x}_t^m \in \mathcal{X}^c)$ similarly to [6], [9] and others, and can be interpreted as probabilistic associations. Unlike these prior

works, here we make an extra step to provably bound the expected dynamic regret of the implementable *discrete* associations. The observation we utilize is that $\mathbf{0} \preceq \mathbf{x}_t, \mathbf{x}_t^m \preceq \mathbf{1}$, and are related through an unbiased sampling, i.e., $\mathbb{E}[\mathbf{x}_t] = \mathbf{x}_t^m$. Based on this, we are able to obtain the following result, proved in the appendix.

Theorem 1. Algorithm LDA ensures the following bound against *any* benchmark sequence $\{\mathbf{x}_t^*\}_{t=1}^T$:

$$\mathbb{E}[\mathcal{R}_T] \leq \mathcal{O}\left(\sqrt{T(1+P_T)}\right) + G_f T \sqrt{I - (I/J)}$$

where $\|f_t(\mathbf{x})\|_2 \leq G_f, \forall \mathbf{x} \in \mathcal{X}, \forall t$.

Discussion. The theorem shows that the regret of LDA is sublinearly dependent on the oracle's HO delay, which marks its learning capability. There is also an unavoidable non-diminishing error term due to discretization. In fact, this error can be eliminated by normalizing properly the step β and using the doubling trick. Due to lack of space we kindly refer the reader to [51, Lemma 3] for further details. The doubling trick will also eliminate the need to know in advance the horizon T . Besides, even without the step normalization, we find the error to be only 1.1–1.3% of the objective value, when tested on static and volatile scenarios (see Sec. VII).

Moreover, it is important to stress that Algorithm LDA is oblivious to information such as the SINR during the slot, which, in practice, is unknown at the beginning of each slot [17]. This is a key difference of the proposed approach compared to prior works such as [6], [9]. Further, LDA is scalable to the number of UEs and BSs and is amenable to near-real-time execution as it has relatively lightweight operations and thus, can be implemented in O-RAN [17], [19].

VI. EXTENSIONS

Finally, we discuss two key extensions: (i) when HO delays are unknown, where we show that LDA can adapt to them dynamically; and (ii) when there is an ML mechanism that proposes associations based on SINR/mobility forecasts [6], [12], and prove that LDA can seamlessly benefit from them.

Time-varying HO delays. Algorithm LDA does not require knowing in advance the elements of A , i.e., the HO delays. To see this, first, observe that the experts do not use the HO delay when they update their decisions (only the throughput $\nabla g_t(\mathbf{x}_t)$). And secondly, the meta-learner uses the HO delays when it calculates $\ell_t(\mathbf{x}_t^k), \forall k \in \mathcal{K}$ to update the weights \mathbf{w}_t , which takes place *after* the HOs are realized. This flexibility allows to tackle cases where the HO delays for each type of UEs and BSs, change with time. From a technical point of view, instead of the fixed A -norm, LDA can use a time-varying norm $\|\cdot\|_{A_t}$, where each A_t captures the delay that was observed (a posteriori) at each slot t . This information is then used to calculate the weights \mathbf{w}_{t+1} and association \mathbf{x}_{t+1} . The proofs of Lemma 1 and Theorem 1 follow nearly verbatim, with the modification of changing the fixed norm to the time-varying norms, and redefining $P_T = \sum_{t=1}^T \|\mathbf{x}_t^* - \mathbf{x}_{t-1}^*\|_{A_t}$.

Encompassing Forecasters. Unlike our approach that learns from runtime data, several recent works [6], [12] proposed to decide the HOs leveraging forecasters for SINRs. The output of such tools, let us denote it $\{\mathbf{x}_t^p\}_t$, can be very close to $\{\mathbf{x}_t^*\}_t$ if the forecasters are accurate, but very suboptimal




Fig. 6: Static scenario for $\gamma = 20$: (a) average dynamic regret and (b) average obtained objective function f .

otherwise (e.g., if there is a distribution shift). Our framework, on the other hand, can benefit from such tools in a robust fashion, by assessing their accuracy in real time. The meta-learner can include $\{x_t^p\}_t$ as the $(K+1)$ th expert, and assess its performance in real-time so as to discard it if proved inaccurate. And if the forecasting tool is effective, the regret will improve significantly. This can be seen by revisiting the proof of Lemma 1 that utilizes the Hedge algorithm (see [37, Lemma 1]), which bounds the regret of the meta-learner from the best expert (thus, also from the forecaster p) as:

$$\max_{k \in \mathcal{K} \cup p} \sum_{t=1}^T \ell_t(x_t^k) - \sum_{t=1}^T \ell_t(x_t^m) \stackrel{(\alpha)}{=} \sum_{t=1}^T \ell_t(x_t^p) - \sum_{t=1}^T \ell_t(x_t^m) \leq \frac{\beta T c^2}{8} + \frac{1}{\beta} \ln \frac{1}{w_1^k}$$

where $c = 2GD + D_A$ and (α) holds when the forecaster is the best expert. Defining its error $\sum_t \ell_t(x_t^*) - \ell_t(x_t^p) = \epsilon_T$ we get:

$$\mathbb{E}[\mathcal{R}_T] \leq \frac{\beta T c^2}{8} + \frac{1}{\beta} \ln \frac{1}{w_1^k} + \epsilon_T + G_f T \sqrt{I - (I/J)}$$

which is optimized when $\beta = \mathcal{O}(1/\sqrt{T})$ (as in Theorem 1). Comparing this result with Theorem 1, we see that when the forecaster is successful, the overall performance improves by dropping an entire term, and does not depend on P_T ; while, when the forecaster is found to be inaccurate, LDA maintains the previous performance as it relies on a different expert. The idea of combining forecasters with online learning is often referred to as *optimistic learning* [35], [52], and we apply it, for the first time, in the context of SOL with dynamic regret.

VII. PERFORMANCE EVALUATION

LDA undergoes rigorous evaluation across multiple scenarios that encompass both real-world conditions (actual users, cells, and SINRs) and synthetic ones, verifying that its efficacy is broadly applicable. We utilize these scenarios to investigate the algorithm's learning convergence, and compare its performance with different benchmarks in terms of (i) attained dynamic regret, (ii) accumulated objective function, throughput, and HO cost, and (iii) impact of incorporating \mathbf{A} versus using a simple $L2$ norm, or none at all. The employed benchmarks are (i) an LDA-based algorithm using the Euclidean norm in the HO cost, called LDA 2-norm, aiming to highlight the significance of the \mathbf{A} -norm; (ii) an advanced greedy algorithm, Max SINR, that assigns UEs to BSs based on the maximum SINR from the previous slot (the current slot's SINR is




Fig. 7: Volatile scenario for $\gamma = 20$: (a) average dynamic regret and (b) average obtained objective function f .

only known post-association), disregarding HO delays; (iii) an optimal Oracle, used in the definition of dynamic regret in (7), which has complete knowledge of future; (iv) a basic Random algorithm that makes the associations randomly, serving as a minimal benchmark.

Synthetic scenarios. We select two synthetic scenarios that are in line with those in prior work [17], [19], [53], and which we define as follows, for $T = 5k$ slots: (i) *static*: the SINR $s_{ij}(t)$ remains constant across all slots, and (ii) *volatile*: $s_{ij}(t)$ fluctuates every 5 slots within the range of [10, 30]dB [54], encompassing poor to excellent values. In both scenarios, we randomly select bandwidths $w_j \in \{5, 10, 15, 20\}$ MHz [55], while \mathbf{A} takes random values within [0, 1], and $\gamma = 20$ (penalizes more HO cost, without sacrificing throughput). We choose a smaller case study with $I = 100$ UEs and $J = 10$ BS to facilitate the calculation of average regret, as determining the best oracle is computationally intensive. However, it is important to note that the best oracle is not necessary for running LDA. Figs. 6 and 7 show that LDA identifies the best benchmark (diminishing regret) just in solely $T = 5k$ slots. This verifies our claim (see Sec. V) that the error due to discretization is small. Conversely, Max SINR fails to converge, even though the SINR does not change, with the regret remaining constant at 6.5 (exploiting sub-optimal associations). This trend is also evident in the average objective function, where LDA outperforms Max SINR after $t = 1.4k$. As anticipated, Random algorithm performs poorly.

Real-world scenarios. To test LDA in real-world scenarios, we use crowdsourced measurements that contain various signal metrics (e.g., SINR and SNR) along with their precise latitude and longitude information and the concerned BS, at every second. This dataset presents a highly competitive scenario as SINRs can vary arbitrarily, and occasionally exhibit adversarial behavior. The crowdsourced measurements are anonymized, hence we lack specific UE information (e.g., its type or exact mobility pattern). Therefore, \mathbf{A} is chosen from the distribution of our data, as shown in Figs. 2, 3, and 4, ensuring a variety of UE types and HO delays are considered. Moreover, we adopt the Gauss-Markov mobility model [56], in line with prior works [6], with randomness parameter $a = 0.5$ ($a = 0$ being totally random and $a = 1$ being linear motion) and consider velocities in [1, 28] meters per second, and velocity variances in [0, 14]; thus having from pedestrians to fast-moving UEs.

At $t = 1$, 1k UEs are placed randomly on the map at a different location with recorded SINR measurements. At each subsequent slot $t = 2, \dots, T$, for $T = 10k$, each user

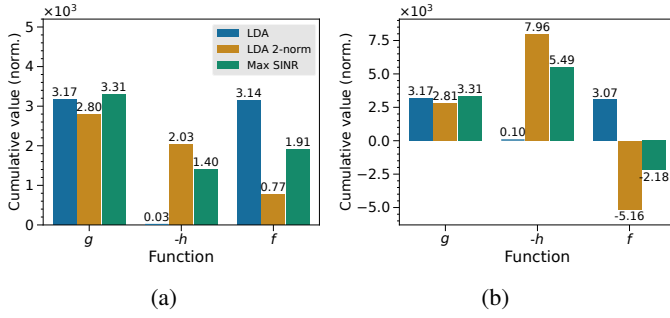


Fig. 8: Comparison of throughput (g), negative HO cost ($-h$), and total (f) values for the real case, when (a) $\gamma = 5$, (b) $\gamma = 20$.

moves in accordance with their individual parameters, to their own new location. To incorporate the signal measurements, we map each new location to the nearest location with available measurements, and continue for the next slot similarly. We utilize measurements from an urban district and choose 25 of the BS with the most data. For the bandwidths of each BS, we use the real ones from our dataset.

We consider two different cases for γ : (i) $\gamma = 5$, and (ii) $\gamma = 20$ (chosen in accordance with throughput and HO cost values). Both cases prioritize high throughput, but the latter penalizes HO costs more. From Fig. 8, we observe the cumulative g values (normalized by the maximum value) are the same for each algorithm, in both $\gamma = 5$ and $\gamma = 20$ case; ensuring no throughput is sacrificed in turn of lower HO delays. At the same time, LDA 2-norm which does not consider the HO delays but solely the number of HOs, achieves $\times 67.7$ and $\times 79.6$ higher HO cost than LDA for $\gamma = 5$ and $\gamma = 20$, respectively. Similarly, although Max SINR achieves 4.4% higher throughput, it does not take into account the HO costs, resulting in 39% (171%) lower total values for $\gamma = 5$ ($\gamma = 20$). Therefore, LDA performs optimally in terms of HO costs, without sacrificing throughput.

To further deep-dive into how LDA responds to γ , we show in Fig. 9 the SINR of two UEs for the first $t = 1k$ slots, and the UE-BS associations for $\gamma = 5$ and $\gamma = 20$. In the former case, we observe that both users change BS more frequently, e.g., for $\gamma = 5$, UE-1 changes 5 times between BS-3 and BS-0 even though the SINR of the latter becomes for a few slots slightly better than the former, while for $\gamma = 20$, UE-1 stays constantly to BS-3. These results show that LDA can be easily adjusted by changing γ , to react more to SINR changes (i.e., maximize throughput), or take more into account the HO cost. The latter is a crucial aspect in mobile networks, where HOs may lead to the ping-pong effect [5].

VIII. CONCLUSIONS

This study addresses the problem of handover (HO) optimization in cellular networks using Smooth Online Learning (SOL). Analyzing countrywide data from a European mobile network with over 40M users and 370k sectors, we identify key correlations between HO failures/delays and the characteristics of radio cells and devices. We develop a dynamic model for UE-to-cell associations that incorporates device and cell features with minimal assumptions. Our proposed algorithm, LDA, aligned with the O-RAN paradigm, demonstrates robust performance and significant improvements in real-world and

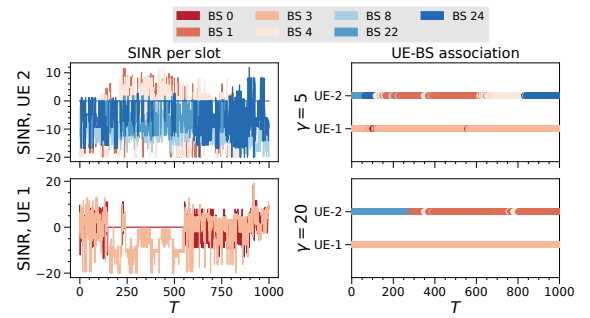


Fig. 9: SINR per slot of two real users (upper and lower left) for $T = 1k$, and the UE-BS associations that LDA chooses, for $\gamma = 5$ (upper right) and $\gamma = 20$ (lower right).

synthetic scenarios. These insights and solutions, based on the MNO's perspective, are crucial for understanding and developing new mechanisms for smooth HOs, providing a foundation for future advancements in network performance.

ACKNOWLEDGMENTS

This work has been supported by (i) the Spanish Ministry of Economic Affairs and Digital Transformation and the European Union – NextGeneration EU, in the framework of the Recovery Plan, Transformation and Resilience (PRTR) through the UNICO I+D 5G SORUS project, and UNICO I+D 5G 2021 refs. number TSI-063000-2021-3, TSI-063000-2021-38, and TSI-063000-2021-52, (ii) the National Growth Fund through the Dutch 6G flagship project “Future Network Services”, and (iii) the European Commission through Grant No. 101139270 (ORIGAMI) and 101192462 (FLECON-6G).

APPENDIX

Proof of Lemma 1. The proof follows the rationale in [18, Th. 4], [37, Th. 3]. First, we bound the regret of each expert w.r.t. benchmark, and then the regret of the meta-learner w.r.t. any expert. W.l.o.g., we set $\gamma = 1$ and define:

$$s_t(\mathbf{x}) = \langle \nabla g_t(\mathbf{x}_t), \mathbf{x} - \mathbf{x}_t \rangle, \text{ and } \bar{\mathbf{x}}_{t+1}^k = \mathbf{x}_t^k + \theta_k \nabla g_t(\mathbf{x}_t).$$

Then, we can write: $s_t(\mathbf{x}_t^*) - s_t(\mathbf{x}_t^k) = \langle \nabla g_t(\mathbf{x}_t), \mathbf{x}_t^* - \mathbf{x}_t^k \rangle$

$$\begin{aligned} &\stackrel{(\alpha)}{\leq} \frac{1}{2\theta_k} (\|\mathbf{x}_t^* - \mathbf{x}_t^k\|_2^2 - \|\bar{\mathbf{x}}_{t+1}^k - \mathbf{x}_t^*\|_2^2) + \frac{\theta_k}{2} G^2 \\ &+ \|\mathbf{x}_{t+1}^k - \mathbf{x}_{t+1}^*\|_2^2 - \|\mathbf{x}_{t+1}^k - \mathbf{x}_t^*\|_2^2 + \frac{\theta_k}{2} G^2 \\ &\stackrel{(\beta)}{=} \frac{1}{2\theta_k} (\|\mathbf{x}_t^k - \mathbf{x}_t^*\|_2^2 - \|\mathbf{x}_{t+1}^k - \mathbf{x}_{t+1}^*\|_2^2 + \\ &+ (\mathbf{x}_{t+1}^k - \mathbf{x}_{t+1}^* + \mathbf{x}_{t+1}^k - \mathbf{x}_t^*)^\top (\mathbf{x}_t^* - \mathbf{x}_{t+1}^*)) + \frac{\theta_k}{2} G^2 \\ &\stackrel{(\gamma)}{\leq} \frac{1}{2\theta_k} (\|\mathbf{x}_t^k - \mathbf{x}_t^*\|_2^2 - \|\mathbf{x}_{t+1}^k - \mathbf{x}_{t+1}^*\|_2^2) + \\ &+ \frac{D_{A^*}}{\theta_k} \|\mathbf{x}_t^* - \mathbf{x}_{t+1}^*\|_A + \frac{\theta_k}{2} G^2. \end{aligned}$$

(α) uses the identity $\langle \mathbf{x}, \mathbf{y} \rangle = (\|\mathbf{x}\|_2^2 + \|\mathbf{y}\|_2^2 - \|\mathbf{x} - \mathbf{y}\|_2^2)/2$, the definition of $\bar{\mathbf{x}}_{t+1}^k$ and $\|\nabla g_t(\mathbf{x}_t)\|_2 \leq G$; (β) uses the projection non-expansiveness, adds/subtracts $\|\mathbf{x}_{t+1}^k - \mathbf{x}_{t+1}^*\|$

and $\|x\|_2^2 - \|y\|_2^2 = (x - y)^\top (x + y)$; and (γ) uses Cauchy-Schwartz, triangle inequality, and D_{A^*} . Telescoping to T and using $\|\mathbf{x}_1^k - \mathbf{x}_1^*\|_A^2 \leq D_A$, gives

$$\sum_{t=1}^T s_t(\mathbf{x}_t^*) - s_t(\mathbf{x}_t^k) \leq \frac{D_A}{2\theta_k} + \frac{D_{A^*}P_T}{\theta_k} + \frac{\theta_k T G^2}{2}.$$

Next, we bound the switching cost of (any) expert k :

$$\begin{aligned} \sum_{t=1}^T \|\mathbf{x}_t^k - \mathbf{x}_{t-1}^k\|_A &= \sum_{t=0}^{T-1} \|\mathbf{x}_{t+1}^k - \mathbf{x}_t^k\|_A \leq \\ \sum_{t=0}^{T-1} \|\bar{\mathbf{x}}_{t+1}^k - \mathbf{x}_t^k\|_A &= \sum_{t=0}^{T-1} \|\mathbf{x}_t^k - \theta_k \nabla g_t(\mathbf{x}_t)\|_A \leq \theta_k T G_A, \end{aligned}$$

and combining with the previous result, we get:

$$\sum_{t=1}^T s_t(\mathbf{x}_t^*) - \sum_{t=1}^T \left(s_t(\mathbf{x}_t^k) - \|\mathbf{x}_t^k - \mathbf{x}_{t-1}^k\|_A \right) \leq \quad (13)$$

$$\frac{D_A^2}{2\theta_k} + \frac{D_{A^*}P_T}{\theta_k} + \theta_k T \left(\frac{G^2}{2} + G_A \right). \quad (14)$$

Next, we bound the gap of the meta-learner from all experts. Following [18, Lem. 3] and using the A -norm:

$$\begin{aligned} \|\mathbf{x}_t^m - \mathbf{x}_{t-1}^m\|_A &= \left\| \sum_{k \in \mathcal{K}} w_t^k \mathbf{x}_t^k - \sum_{k \in \mathcal{K}} w_{t-1}^k \mathbf{x}_{t-1}^k \right\|_A \\ &\leq \left\| \sum_{k \in \mathcal{K}} w_t^k (\mathbf{x}_t^k - \mathbf{x}) - \sum_{k \in \mathcal{K}} w_{t-1}^k (\mathbf{x}_{t-1}^k - \mathbf{x}) \right\|_A \\ &+ \left\| \sum_{k \in \mathcal{K}} w_t^k (\mathbf{x}_{t-1}^k - \mathbf{x}) - \sum_{k \in \mathcal{K}} w_{t-1}^k (\mathbf{x}_{t-1}^k - \mathbf{x}) \right\|_A \\ &\leq \sum_k w_t^k \|\mathbf{x}_t^k - \mathbf{x}_{t-1}^k\|_A + |w_t^k - w_{t-1}^k| \|\mathbf{x}_{t-1}^k - \mathbf{x}\|_A \\ &= \sum_{k \in \mathcal{K}} w_t^k \|\mathbf{x}_t^k - \mathbf{x}_{t-1}^k\|_A + D_A \|\mathbf{w}_t - \mathbf{w}_{t-1}\|_1. \end{aligned}$$

The relative loss of meta-learner is: $\sum_{t=1}^T \left(s_t(\mathbf{x}_t^k) - \|\mathbf{x}_t^k - \mathbf{x}_{t-1}^k\|_A \right) - \sum_{t=1}^T \left(s_t(\mathbf{x}_t^m) - \|\mathbf{x}_t^m - \mathbf{x}_{t-1}^m\|_A \right)$

$$\begin{aligned} &\leq \sum_{t=1}^T \left(\sum_k w_t^k \|\mathbf{x}_t^k - \mathbf{x}_{t-1}^k\|_A + \langle \nabla g_t(\mathbf{x}_t), \mathbf{x}_t^k - \mathbf{x}_t \rangle \right. \\ &\quad \left. - \|\mathbf{x}_t^k - \mathbf{x}_{t-1}^k\|_A \right) + D_A \|\mathbf{w}_t - \mathbf{w}_{t-1}\|_1 = \\ &\sum_{t=1}^T \left(\sum_k w_t^k \ell_t(\mathbf{x}_t^k) - \ell_t(\mathbf{x}_t^k) \right) + D_A \sum_{t=1}^T \|\mathbf{w}_t - \mathbf{w}_{t-1}\|_1. \end{aligned}$$

The first term is bounded noting that $d \leq \ell_t(\mathbf{x}) \leq d + c$, with $d = -DG - D_A$, $c = 2GD + D_A$ and using the Hedge bound [57, Th. 2.2] (see also [37, Lem. 1]):

$$\sum_{t=1}^T \ell_t(\mathbf{x}_t^m) - \min_{k \in \mathcal{K}} \left(\sum_{t=1}^T \ell_t(\mathbf{x}_t^k) + \frac{1}{\beta} \ln \frac{1}{w_1^k} \right) \leq \frac{\beta T c^2}{8}.$$

Using the value of c and the definition of \mathbf{x}_t^m , we get:

$$\begin{aligned} \sum_{t=1}^T \left(\sum_k w_t^k \ell_t(\mathbf{x}_t^k) - \ell_t(\mathbf{x}_t^k) \right) &\leq \\ \frac{1}{\beta} \ln \frac{1}{w_1^k} + \frac{\beta T (2GD + D_A)^2}{8}. \end{aligned} \quad (15)$$

Next, we use the strong convexity of entropic FTRL [58, Lem. 7], as the basis for Hedge, to arrive at:

$$\|\mathbf{w}_t - \mathbf{w}_{t-1}\|_1 \leq \beta \left\| \left[\ell_{t-1}(\mathbf{x}_t^k) \right]_k \right\|_\infty \leq \beta (GD + D_A).$$

Combining the above, we prove that $\forall k \in \mathcal{K}$, it is:

$$\begin{aligned} \sum_{t=1}^T s_t(\mathbf{x}_t^k) - \|\mathbf{x}_t^k - \mathbf{x}_{t-1}^k\|_A - s_t(\mathbf{x}_t^m) + \|\mathbf{x}_t^m - \mathbf{x}_{t-1}^m\|_A &\leq \\ \frac{1}{\beta} \ln \frac{1}{w_1^k} + \beta T \underbrace{\left[(2GD + D_A)^2 (D_A + (1/8)) \right]}_\nu. \end{aligned} \quad (16)$$

and set $\beta = 1/\sqrt{T\nu}$ to balance the RHS (omitting \ln).

Finally, considering the max and min values of P_T , the expert step that minimizes the RHS of (14) lies in:

$$\sqrt{\frac{D_A^2}{T(G^2 + 2G_A)}} \leq \theta^* \leq \sqrt{\frac{D_A^2 + 2D_{A^*}D_A T}{T(G^2 + 2G_A)}} \quad (17)$$

and inspecting the experts' steps $\{\theta_k\}_k$ in Lemma 1, we see that at least one expert step lies in that range. Thus, we obtain:

$$\begin{aligned} \sum_{t=1}^T f_t(\mathbf{x}_t^*) - f_t(\mathbf{x}_t^m) &\leq \sqrt{T} \left[\sqrt{\nu} (1 + \ln(1/w_1^k)) + \right. \\ &\quad \left. (G^2 + 2G_A)^{1/2} (D_A^2 + 2D_{A^*}P_T)^{1/2} \right]. \end{aligned} \quad (18)$$

Proof of Theorem 1. It holds:

$$\begin{aligned} \mathbb{E}[\mathcal{R}_T] &\leq \mathbb{E} \left[\sum_{t=1}^T \left(f_t(\mathbf{x}_t^*) - f_t(\mathbf{x}_t) + f_t(\mathbf{x}_t^m) - f_t(\mathbf{x}_t^m) \right) \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T f_t(\mathbf{x}_t^m) - f_t(\mathbf{x}_t) \right] + \sum_{t=1}^T \left(f_t(\mathbf{x}_t^*) - f_t(\mathbf{x}_t^m) \right). \end{aligned}$$

The second term is bounded by Lemma 1. The first term is the discretization *error*, and we bound it as follows:

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T \left(f_t(\mathbf{x}_t^m) - f_t(\mathbf{x}_t) \right) \right] &\stackrel{(\alpha)}{\leq} \sum_{t=1}^T \left(G_f \mathbb{E} [\|\mathbf{x}_t - \mathbf{x}_t^m\|] \right) \\ &\stackrel{(\beta)}{\leq} G_f \sum_{t=1}^T \sqrt{\mathbb{E} \left[\sum_{i=1}^I \sum_{j=1}^J \left(x_{ij}(t) - \mathbb{E}[x_{ij}(t)] \right)^2 \right]}. \end{aligned}$$

In (α) we used that expectation is a linear operator and the Lipschitz constant of f_t , and in (β) Jensen's inequality and the unbiased sampling. This is the variance of the random binary output of $Q_{\mathcal{X}}$. Since the binary vector is subject to a simplex per user, each sum w.r.t. j is bounded by the variance $1 - (1/J)$, and the overall term by $\mu \triangleq \sqrt{I - (I/J)}$. Using (16):

$$\begin{aligned} \mathbb{E}[\mathcal{R}_T] &= G_f T \mu + \frac{1}{\beta} \ln \frac{1}{w_1^k} + \beta T \nu \\ &\quad + (G^2 + 2G_A)^{1/2} (D_A^2 + 2D_{A^*}P_T)^{1/2}. \end{aligned}$$

The last term can be made sublinear using the step normalization trick of [51, Lem. 3]. Due to lack of space we omit the details of the derivation.

REFERENCES

- [1] H. Tabassum, et al., “Fundamentals of Mobility Aware Performance Characterization of Cellular Networks,” *IEEE Communications Surveys & Tutorials*, vol. 21, no. 3, 2019.
- [2] S. Alraih, et al., “A Survey on Handover Optimization in B5G Networks: Challenges & Solutions,” *IEEE Access*, vol. 11, 2023.
- [3] X. Lin, et al., “Towards Understanding the Fundamentals of Mobility in Cellular Networks,” *IEEE Trans. on Wireless Communications*, vol. 12, no. 4, 2013.
- [4] A. Hassan, et al., “Vivisecting Mobility Management in 5G Cellular Networks,” in *Proc. of ACM SIGCOMM*, 2022.
- [5] D. Zidic, et al., “Analyses of Ping-Pong Handovers in Real 4G Telecommunication Networks,” *Computer Networks*, vol. 227, 5 2023.
- [6] M. Gupta, et al., “Forecaster-Aided User Association and Load Balancing in Multi-Band Mobile Networks,” *IEEE Trans. on Wireless Communications*, vol. 23, no. 5, 2024.
- [7] —, “Load Balancing and Handover Optimization in Multi-band Networks using Deep Reinforcement Learning,” in *Proc. of IEEE GLOBECOM*, 2021.
- [8] Y. Cao, et al., “User Access Control in Open Radio Access Networks: A Federated Deep Reinforcement Learning Approach,” *IEEE Trans. on Wireless Communications*, vol. 21, no. 6, 2022.
- [9] A. Prado, et al., “Enabling Proportionally-Fair Mobility Management With Reinforcement Learning in 5G Networks,” *IEEE JSAC*, vol. 41, no. 6, 2023.
- [10] A. Garcia-Saavedra, et al., “O-RAN: Disrupting the Virtualized RAN Ecosystem,” *IEEE Communications Standards Magazine*, vol. 5, no. 4, 2021.
- [11] A. Lacava, et al., “Programmable and Customized Intelligence for Traffic Steering in 5G Networks Using Open RAN Architectures,” *IEEE Trans. on Mobile Computing*, vol. 23, no. 4, 2024.
- [12] I. Panitsas, et al., “Predictive Handover Strategy in 6G and Beyond: A Deep and Transfer Learning Approach,” in *Proc. of IEEE GLOBECOM*, 2024.
- [13] M. Kalntis, et al., “Through the Telco Lens: A Countrywide Empirical Study of Cellular Handovers,” in *Proc. of the ACM IMC*, 2024.
- [14] J. Choi, et al., “Throughput Estimation Based Distributed Base Station Selection in Heterogeneous Networks,” *IEEE Trans. on Wireless Communications*, vol. 14, no. 11, 2015.
- [15] E. Hazan, “Introduction to Online Convex Optimization,” *Foundations and Trends in Optimization*, vol. 2, no. 3–4, 2016.
- [16] N. Chen, et al., “Smoothed Online Convex Optimization in High Dimensions via Online Balanced Descent,” in *Proc. of COLT*, 2018.
- [17] M. Kalntis, et al., “Adaptive Resource Allocation for Virtualized Base Stations in O-RAN with Online Learning,” *IEEE Trans. on Communications*, 2024.
- [18] L. Zhang, et al., “Revisiting Smoothed Online Learning,” in *Proc. of NeurIPS*, 2021.
- [19] M. Kalntis, et al., “Energy-Aware Scheduling of Virtualized Base Stations in O-RAN with Online Learning,” in *Proc. of IEEE GLOBECOM*, 2022.
- [20] X. Yuan, et al., “Understanding 5G Performance for Real-world Services: a Content Provider’s Perspective,” in *Proc. of ACM SIGCOMM*, 2022.
- [21] M. Ghoshal, et al., “Performance of Cellular Networks on the Wheels,” in *Proc. of the ACM IMC*, 2023.
- [22] Y. Ni, et al., “POLYCORN: Data-driven Cross-layer Multipath Networking for High-speed Railway through Composable Schedulerlets,” in *Proc. of USENIX NSDI* 23, 2023.
- [23] J. Wang, et al., “An Active-Passive Measurement Study of TCP Performance over LTE on High-Speed Rails,” in *Proc. of ACM MobiCom*, 2019.
- [24] A. Narayanan, et al., “A First Look at Commercial 5G Performance on Smartphones,” in *Proc. of The Web Conference*, 2020.
- [25] —, “A Variegated Look at 5G in the Wild: Performance, Power, and QoE Implications,” in *Proc. of ACM SIGCOMM*, 2021.
- [26] C. Shen, et al., “A Non-Stochastic Learning Approach to Energy Efficient Mobility Management,” *IEEE JSAC*, vol. 34, no. 12, 2016.
- [27] —, “A Learning Approach to Frequent Handover Mitigations in 3GPP Mobility Protocols,” in *Proc. of IEEE WCNC*, 2017.
- [28] Y. Zhou, et al., “A Non-Stationary Online Learning Approach to Mobility Management,” *IEEE Trans. on Wireless Communications*, vol. 18, no. 2, 2019.
- [29] Y. Li, et al., “Can Online Learning Increase the Reliability of Extreme Mobility Management?” in *Proc. of IEEE/ACM IWQoS*, 2021.
- [30] L. Sun, et al., “Spatial and Temporal Contextual Multi-Armed Bandit Handovers in Ultra-Dense mmWave Cellular Networks,” *IEEE Trans. on Mobile Computing*, vol. 20, no. 12, 2021.
- [31] S. K. Singh, et al., “User Association in Dense mmWave Networks as Restless Bandits,” *IEEE Trans. on Vehicular Technology*, vol. 71, no. 7, 2022.
- [32] Z. Wang, et al., “Handover Control in Wireless Systems via Asynchronous Multiuser Deep Reinforcement Learning,” *IEEE Internet of Things Journal*, vol. 5, no. 6, 2018.
- [33] Y. Sun, et al., “Efficient Handover Mechanism for Radio Access Network Slicing by Exploiting Distributed Learning,” *IEEE Trans. on Network & Service Management*, vol. 17, no. 4, 2020.
- [34] A. Masri, et al., “Machine-Learning-Based Predictive Handover,” in *Proc. of IFIP/IEEE IM*, 2021.
- [35] A. Rakhlin, et al., “Online Learning with Predictable Sequences,” in *Proc. of COLT*, 2013.
- [36] J. J. Alcaraz, et al., “Model-Based Reinforcement Learning With Kernels for Resource Allocation in RAN Slices,” *IEEE Trans. on Wireless Communications*, vol. 22, no. 1, 2023.
- [37] L. Zhang, et al., “Adaptive Online Learning in Dynamic Environments,” in *Proc. of NeurIPS*, 2018.
- [38] S. Mukhopadhyay, et al., “Online Caching with Optimal Switching Regret,” in *Proc. of IEEE ISIT*, 2021.
- [39] G. Paschos, et al., “Online Convex Optimization for Caching Networks,” *IEEE/ACM Trans. on Networking*, vol. 28, no. 2, 2020.
- [40] J. Steiger, et al., “Constrained Bandit Learning with Switching Costs for Wireless Networks,” in *Proc. of IEEE INFOCOM*, 2023.
- [41] A. Block, et al., “Smoothed Online Learning is as Easy as Statistical Learning,” in *Proc. of COLT*, 2022.
- [42] G. D. Celik, et al., “Scheduling in Networks With Time-Varying Channels and Reconfiguration Delay,” *IEEE/ACM Trans. on Networking*, vol. 23, no. 1, 2015.
- [43] G. Liu, et al., “5G Deployment: Standalone vs. Non-Standalone from the Operator Perspective,” *IEEE Communications Magazine*, vol. 58, no. 11, 2020.
- [44] P. Viswanath, et al., “Opportunistic beamforming using dumb antennas,” *IEEE Trans. Information Theory*, vol. 48, no. 6, 2002.
- [45] R. Srikant, et al., “Communication Networks: An Optimization, Control and Stochastic Networks Perspective,” *Cambridge University Press*, 2014.
- [46] N. Littlestone, et al., “The Weighted Majority Algorithm,” *Proc. of IEEE FOCS*, 1989.
- [47] S. Arora, et al., “The Multiplicative Weights Update Method: a Meta-algorithm and Applications,” *Theory of Computing*, vol. 8, no. 6, 2012.
- [48] A. Beck, “First-Order Methods in Optimization,” *MOS-SIAM Series on Optimization*, 2017.
- [49] N. Cesa-Bianchi, et al., “Online Learning with Switching Costs and Other Adaptive Adversaries,” in *Proc. of NIPS*, 2013.
- [50] W. G. Madow, “On the Theory of Systematic Sampling,” *The Annals of Mathematical Statistics*, vol. 20, no. 3, 1949.
- [51] A. Lesage-Landry, et al., “Online Convex Optimization with Binary Constraints,” *IEEE Trans. on Automatic Control*, vol. 66, no. 12, 2021.
- [52] N. Mhaisen, et al., “Online Caching With no Regret: Optimistic Learning via Recommendations,” *IEEE Trans. on Mobile Computing*, vol. 23, no. 5, 2024.
- [53] J. A. Ayala-Romero, et al., “Bayesian Online Learning for Energy-Aware Resource Orchestration in Virtualized RANs,” in *Proc. of IEEE INFOCOM*, 2021.
- [54] 3GPP, “Requirements for support of radio resource management,” 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 36.133, 2024.
- [55] —, “User Equipment (UE) radio transmission and reception,” 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 36.101, 2024.
- [56] T. Camp, et al., “A Survey of Mobility Models for Ad Hoc Network Research,” *Wireless Communications and Mobile Computing*, vol. 2, no. 5, 2002.
- [57] N. Cesa-Bianchi, et al., “Prediction, Learning, and Games,” *Cambridge University Press*, 2006.
- [58] B. McMahan, “A Survey of Algorithms and Analysis for Adaptive Online Learning,” *Journal of Machine Learning Research*, vol. 18, 2017.