



HAL
open science

Dynamic Rate and Channel Selection in Cognitive Radio Systems

Richard Combes, Alexandre Proutière

► **To cite this version:**

Richard Combes, Alexandre Proutière. Dynamic Rate and Channel Selection in Cognitive Radio Systems. *IEEE Journal on Selected Areas in Communications*, 2015, 33 (5), pp.910-921. 10.1109/JSAC.2014.2361084 . hal-01101352

HAL Id: hal-01101352

<https://hal.science/hal-01101352v1>

Submitted on 9 Apr 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Dynamic Rate and Channel Selection in Cognitive Radio Systems

R. Combes and A. Proutiere
KTH, The Royal Institute of Technology

Abstract—In this paper, we investigate dynamic channel and rate selection in cognitive radio systems which exploit a large number of channels free from primary users. In such systems, transmitters may rapidly change the selected (channel, rate) pair to opportunistically learn and track the pair offering the highest throughput. We formulate the problem of sequential channel and rate selection as an online optimization problem, and show its equivalence to a *structured Multi-Armed Bandit* problem. The structure stems from inherent properties of the achieved throughput as a function of the selected channel and rate. We derive fundamental performance limits satisfied by *any* channel and rate adaptation algorithm, and propose algorithms that achieve (or approach) these limits. In turn, the proposed algorithms optimally exploit the inherent structure of the throughput. We illustrate the efficiency of our algorithms using both test-bed and simulation experiments, in both stationary and non-stationary radio environments. In stationary environments, the packet successful transmission probabilities at the various channel and rate pairs do not evolve over time, whereas in non-stationary environments, they may evolve. In practical scenarios, the proposed algorithms are able to track the best channel and rate quite accurately without the need of any explicit measurement and feedback of the quality of the various channels.

I. INTRODUCTION

In cognitive radio systems, radio devices may access a potentially large number of frequency bands or channels. An example of such systems are those exploiting “white space” spectrum, the unused part of the TV/UHF spectrum (unallocated or not used locally). The FCC 2008 ruling allowed unlicensed devices to use parts of this spectrum, provided that devices can detect primary users (TV transmitters and wireless microphones). As a part of the 2010 ruling [1], FCC mandates the use of a geolocation database to identify which frequencies are free from primary users. By querying the geolocation database, we are guaranteed to obtain a set of channels free from primary transmitters and we avoid the difficult problem of sensing primary users.

We consider systems exploiting channels known to be free from primary users. For the transmission of each

packet, transmitters can select a coding rate from a finite predefined set (as in 802.11 systems for example) and a channel from the set of available channels. The outcome of a packet transmission is random, and the probabilities of successfully transmitting a packet using the various (channel, rate) pairs are a priori unknown at the transmitter; they need to be learnt based on trial and error. These probabilities can vary significantly and randomly over time and across channels; they also strongly depend on the chosen coding rate. As a consequence, tracking the best (channel, coding rate) pair for transmission may greatly improve the system performance. In this paper, we aim at designing sequential channel and coding rate selection schemes that efficiently track the best available channel and the corresponding coding rate.

As shown in previous works, see e.g. [2], [3], RSSI (Receive Signal Strength Indicator) is a poor predictor of channel quality, and hence of the packet successful transmission probabilities. In OFDM systems for example, this stems from the fact that RSSI does not report the individual signal strength experienced on the various sub-carriers. In order to accurately estimate the quality of a wide-band channel, more sophisticated techniques with specific hardware are needed [4], [5]. But these techniques are not typically supported in current commercial radio hardware. Instead, we need to infer the quality of each channel at each transmission rate through probing. Here we consider 802.11-like systems, where the only feedback sent from the receiver to the transmitter is whether a data packet has been successfully received or not. Hence by probing, we mean that several real data packets have to be sent on each channel and at each rate to construct a reliable estimate of the channel quality. In the design of channel and rate selection schemes, we then face a classical exploration vs. exploitation trade-off problem. We need to exploit the (channel, rate) pair that has offered the best throughput so far, whilst constantly exploring other pairs in case one of them is actually optimal.

We rigorously formulate the design of the optimal sequential (channel, rate) selection algorithms as an online stochastic optimization problem. In this problem,

the objective is to maximize the number of packets successfully sent over a finite time horizon. We show that this problem reduces to a Multi-Armed Bandit (MAB) problem [6]. In MAB problems, a decision maker sequentially selects an action (also called an “arm”), and observes the corresponding reward. Rewards of a given arm are random variables with unknown distribution. The objective is to design sequential action selection strategies that maximize the expected reward over a given time horizon. These strategies have to achieve an optimal trade-off between exploitation (actions that have provided high rewards so far have to be selected) and exploration (sub-optimal actions have to be chosen so as to learn their average rewards). For our (channel, rate) selection problem, the various arms correspond to the decisions available at the transmitter to send packets, i.e., an arm corresponds to a channel and a coding rate. When a (channel, rate) pair is selected for a packet transmission, the reward is equal to 1 if the transmission is successful, and equal to 0 otherwise. The average successful packet transmission probabilities at the various (channel, rate) pairs are unknown, and have to be learnt.

The MAB problem corresponding to the design of channel and rate selection mechanisms is referred to as a *structured* MAB problem. It differs from classical MAB problems. (i) First, the rewards associated with the various rates on a given channel are stochastically correlated, i.e., the outcomes of transmissions at different rates are not independent: for example, if a transmission at a high rate is successful, it would be also successful at lower rates. (ii) Then, the average throughputs achieved at various rates exhibit natural structural properties. For a given channel, the throughput is an unimodal function of the selected rate. (iii) In addition, most often, on all channels, the packet successful transmission probabilities are close to 1 at low rates, and abruptly decrease to 0 as the rate increases. This additional structure, referred to as *graphical unimodality*, allows us to predict the outcomes of transmissions on various channels. As we demonstrate, correlations and (graphical) unimodality are instrumental in the design of channel and rate selection mechanisms, and can be exploited to learn and track the best (channel, rate) pair quickly and efficiently. Finally, note that most MAB problems consider stationary environments, which, for our problem, means that the successful packet transmission probabilities for the different (channel, rate) pairs do not vary over time. In practice, the transmitter faces a non-stationary environment as these probabilities could evolve over time. We consider both stationary and non-stationary radio environments.

In the case of stationary environments, we derive

an upper bound of the expected reward that can be achieved in structured MAB problems. This provides a fundamental performance limit that *any* (channel, rate) selection algorithm cannot exceed. This limit quantifies the inevitable performance loss due to the need to explore sub-optimal (channel, rate) pairs. It also indicates the performance gains that can be achieved by devising schemes that optimally exploit the correlations and the structural properties of the MAB problem. We present sequential (channel, rate) selection algorithms that optimally exploit the structural properties of the problem: for our algorithms, we prove that the performance loss due to the need to explore sub-optimal (channel, rate) pairs does not depend on the number of available rates. We also extend our algorithms to non-stationary radio environments. Finally, we evaluate the performance of the proposed algorithms using an office white-space testbed operating in the 500MHz-600MHz band, and simulation experiments.

To our knowledge, the problem of sequential channel and rate selection has only been investigated in [3], where heuristic algorithms have been developed. In contrast, we formulate and solve this problem rigorously, i.e., we provide fundamental upper performance bounds satisfied under any channel and rate selection algorithm and design optimal algorithms that match these bounds.

Contributions.

- We formulate the design of (channel, rate) selection algorithms as an online optimization problem, and establish its equivalence to a structured MAB problem.
- We derive a performance upper bound satisfied by any (channel, rate) selection algorithm, depending on the assumptions made on the structure of the problem – three scenarios with increasing structure are considered: 1. no structural assumption is made; 2. the throughput on each channel is a unimodal function of the rate; 3. the throughput is a graphically unimodal function of the channel and rate. We also quantify the performance gains that one may achieve by exploiting the structural properties of the problem.
- We propose three (channel, rate) selection algorithms, one for each of the above scenarios, and analyze their performance in stationary radio environments. We prove that our algorithms optimally exploit the structural properties of the throughputs.
- We briefly discuss the extensions of our algorithms to non-stationary radio environments.
- Finally, we evaluate the performance of our algorithms using simulation experiments. To this aim, we use artificially generated traces, as well as traces

extracted from a white space test-bed operating in the UHF radio spectrum .

Paper organization. The next section is devoted to the related work. In Sections III and IV, we formulate the problem of sequential selection of channel and rate selection as a *structured* bandit problem. Section V presents fundamental upper performance bounds for this problem. In Section VI, asymptotically optimal algorithms are proposed. Section VII deals with non-stationary radio environments. Section VIII presents numerical experiments to evaluate the performance of our algorithms.

II. RELATED WORK

First observe that the joint channel and rate selection problem is considerably more difficult than detecting channels with no primary users as considered in a lot of recent works, see e.g. [7], [8], [9], [10], [11], [12], [13]. In some of these papers, a MAB framework has been used to design primary user detection algorithms. The presence or the absence of primary users just means that a channel is either good or bad. When selecting both channel and rate, the dimension of the problem becomes larger, and there are multiple and numerous possible channel states. Primary users are not considered in our work, as we assume that transmitters can use a geolocation database to get a list of channels free from primary users [1].

It should also be observed that most of the work on dynamic spectrum access considers stationary radio environments. In [9], [10] for example, the authors use classical stochastic control techniques (Markov Decision Processes) to sequentially select a channel for transmission. The underlying assumption is that the environment is stationary, i.e., the packet successful transmission probabilities do not evolve over time. In this paper, both stationary and non-stationary radio environments are explored. Test-bed experiments actually suggest that the environment is non-stationary in practice, even in networks where nodes do not move such as indoor offices, see e.g. [3].

Our problem resembles the rate adaptation problem in 802.11 systems, see e.g. [14], [15], [16]. But again, our problem has one additional dimension (a channel has to be selected): in turn, the number of available decisions at the transmitter is much larger than in 802.11 systems where only the rate has to be chosen. Rate adaptation algorithms are not applicable when the channel can also be selected for each packet transmission. This is due to the fact that the transmitter does not continuously monitor the same channel (as in 802.11 systems), and has to switch channels often to discover the best (channel, rate) pair as rapidly as possible.

There is an abundant literature on MAB problems, and engineers have applied these problems to dynamic spectrum access [8], [10], [11], [17]. Most existing theoretical results, see [18] for a recent survey, are concerned with *unstructured* MAB problems, i.e., problems where the average reward associated with the various decisions are not related. For this kind of problems, Lai and Robbins [6] derived an asymptotic lower bound on regret and also designed optimal sequential decision algorithms. When the average rewards are structured (as this is the case for our problem), the design of optimal decision algorithms is more challenging, see e.g. [18]. Non-stationary environments have not been extensively studied in the bandit literature: Most often unstructured MAB only are analyzed, see [19], [20], [21].

To our knowledge, the only work dealing with joint (channel, rate) selection is [3]. However there, the structural properties of the corresponding MAB problem had not been identified, and the authors only proposed algorithms based on heuristics. This contrasts with the present work: we rigorously determine fundamental limits satisfied by any (channel, rate) adaptation algorithm, and propose algorithms approaching these limits.

III. MODELS

We consider a single link (a transmitter-receiver pair). At time 0, the link becomes active and the transmitter starts sending packets to the receiver. For each packet, the transmitter selects a channel from a finite set $\mathcal{C} = \{1, \dots, C\}$, and a coding and modulation scheme from a finite set $\mathcal{K} = \{1, \dots, K\}$. The transmission rate when using coding and modulation scheme k is denoted by r_k and we define the set of rates $\mathcal{R} = \{r_k, k \in \mathcal{K}\}$. \mathcal{R} is ordered, i.e., $r_1 < r_2 < \dots < r_K$. After a packet is sent, the transmitter is informed of whether the transmission has been successful. Based on the observed past transmission successes and failures at the various channels and rates, the transmitter has to select a channel and rate pair for the next packet transmission. Let Π denote the set of all possible sequential (channel, rate) selection schemes. Packets are assumed to be of equal size, and without loss of generality, for any k , the duration of a packet transmission at rate r_k is $1/r_k$.

A. Channel models

For the i -th packet transmission on channel c at rate r_k , a binary random variable $X_{ck}(i)$ represents the success ($X_{ck}(i) = 1$) or failure ($X_{ck}(i) = 0$) of the transmission. $\mathbb{E}[X_{ck}(i)]$ refers to as the packet successful transmission probability on channel c at rate r_k (i.e., it is the packet reception rate). We consider

both stationary and non-stationary radio environments. In stationary environments, the success transmission probabilities on the various channels and at different rates do not evolve over time. This arises when the system considered is static (in particular, the transmitter and receiver do not move). In non-stationary environments, success transmission probabilities can evolve over time. Unless otherwise specified, we consider stationary radio environments. Non-stationary environments are treated in Section VII.

We assume that $X_{ck}(i)$, $i = 1, 2, \dots$, are independent and identically distributed, and we denote by θ_{ck} the success transmission probability on channel c at rate r_k : $\theta_{ck} = \mathbb{E}[X_{ck}(i)]$. We verified that the i.i.d. assumption holds in our test-bed and simulation framework. Denote by (c^*, k^*) the optimal (channel, rate) pair, i.e., $(c^*, k^*) \in \arg \max_{c,k} r_k \theta_{ck}$. To simplify the exposition and the notation, we assume that the optimal (channel, rate) pair is unique, i.e., $r_{k^*} \theta_{c^* k^*} > r_k \theta_{ck}$, for all $(c, k) \neq (c^*, k^*)$. Our analysis can be extended in an easy way to scenarios where the optimal channel and rate pair is not unique. Note however that scenarios where different channel and rate pairs yield exactly the same throughput should happen very rarely in practice. We further introduce, for any channel c , the optimal rate $r_{k_c^*}$, i.e., $(c, k_c^*) \in \arg \max_k r_k \theta_{ck}$. Again for simplicity, we assume that on any channel, the optimal rate is unique: $r_{k_c^*} \theta_{ck_c^*} > r_k \theta_{ck}$, for all $k \neq k_c^*$. The throughput achieved using (channel, rate) pair (c, k) is denoted by $\mu_{ck} = r_k \theta_{ck}$. The maximum throughput on channel c is $\mu_c^* = \mu_{ck_c^*}$, and the throughput achieved using the optimal (channel, rate) pair is $\mu^* = \mu_{c^* k^*} = \mu_{c^* k^*}$.

Although we do not account for the presence of primary users in this work, we could actually model scenarios where on each channel c , primary users occupy the channel with some fixed probability ζ_c , and in an i.i.d. manner across time. Indeed in such scenarios, we just need to replace θ_{ck} by $(1 - \zeta_c) \theta_{ck}$. If the primary users occupy channels for long periods of time (not in an i.i.d. manner), the analysis would be significantly more challenging. This kind of situations is investigated in [9] for example.

B. Structural properties

The successful transmission probabilities $\theta = (\theta_{ck}, c \in \mathcal{C}, k \in \mathcal{K})$ are initially unknown at the transmitter, and have to be learnt. When the number of (channel, rate) pairs grows large, learning the best pair for transmission then becomes really challenging. Fortunately, the outcomes of transmissions using the various (channel, rate) pairs exhibit structural properties

that can be exploited to speed up the learning process. To emphasize the importance of exploiting the structural properties, we consider three scenarios with increasing structure.

1) *Scenario 1 – No structure*: If no structural assumptions are made regarding the successful transmission probabilities, then $\theta \in [0, 1]^{C \times K}$. In such scenarios, we will show that the performance loss due to the need to explore sub-optimal (channel, rate) pairs scales linearly with the number of channels and rates.

2) *Scenario 2 – Unimodality*: First observe that the successes and failures of transmissions on a given channel at various rates are statistically correlated. Indeed, if a transmission is successful at a high rate, it has to be successful at a lower rate. Similarly, if a low-rate transmission fails, then transmitting at a higher rate would also fail. Formally this means that for any channel c , $\theta_c = (\theta_{c1}, \dots, \theta_{cK}) \in \mathcal{T}$, where $\mathcal{T} = \{\eta \in [0, 1]^K : \eta_1 \geq \dots \geq \eta_K\}$. Then, in practice, it has been observed (and this is confirmed in our numerical experiments) that the throughput achieved on a given channel is a unimodal function of the transmission rate, see e.g. [5], [16]. In other words, for any channel c , $\theta_c \in \mathcal{U}$, where $\mathcal{U} = \{\eta \in [0, 1]^K : \exists k_1, r_1 \eta_1 < \dots < r_{k_1} \eta_{k_1}, r_{k_1} \eta_{k_1} > r_{k_1+1} \eta_{k_1+1} > \dots > r_K \eta_K\}$. In summary in Scenario 2, for any channel c , $\theta_c \in \mathcal{T} \cap \mathcal{U}$.

3) *Scenario 3 – Graphical unimodality*: We further observe (see Section VIII) that on a given channel, the throughput first grows linearly with the rate (the successful transmission probability is close to 1), and then abruptly decreases to 0. This observation has been made in earlier work, see [14] (the author refers to this scenario as the *steep throughput* scenario), [5]. This knowledge can be exploited to build a relationship between the throughputs achieved on various channels. Indeed, for example, the throughputs observed on two different channels are roughly identical in their growth phase (when the rates are low and the success probabilities are close to 1). To exploit this observation, we remark that if it holds, the throughput is a *graphically unimodal* function of the (channel, rate) pair as defined below.

We first construct a directed graph $G = (V, E)$ whose vertices correspond to the (channel, rate) pairs. When $(d, d') \in E$, we say that the decision d' is a neighbor of decision d , and we define $\mathcal{N}(d) = \{d' \in V : (d, d') \in E\}$ as the set of neighbors of d . The throughput or average reward of decision $d = (c, k)$ is denoted by $\mu_d = r_k \theta_{ck}$. Graphical unimodality expresses the fact that when the optimal decision is $d^* = (c^*, k^*)$, then for any $d \in V$, there exists a path in G from d to d^* along which the expected reward is strictly increasing. In other

words there is no *local* maximum in terms of expected reward except at d^* . The notion of locality is defined through that of neighborhood $\mathcal{N}(d), d \in V$. Formally, $\theta \in \mathcal{U}_G$, where \mathcal{U}_G is the set of successful transmission probabilities $\theta \in [0, 1]^{C \times K}$ such that, if $d^* = (c^*, k^*) \in \arg \max_{(c,k)} r_k \theta_{ck}$, for any $d = (c, k) \in V$, there exists a path $(d_0 = d, d_1, \dots, d_p = d^*)$ in G such that for any $i = 1, \dots, p$, $\mu_{d_i} > \mu_{d_{i-1}}$.

Let us now complete the construction of graph G . The set of edges E is: $((c, k), (c, k - 1)), ((c, k), (c, k + 1))$ and $((c, k), (c', k)), ((c, k), (c', k + 1))$ for all (channel, rate) pair (c, k) , and all c' . An example of such a graph G is presented in Figure 1 – for 2 channels and 4 rates. When the above observation made on θ holds (steep scenario as defined in [14]), it is easy to check that the throughput is a graphically unimodal function (w.r.t. graph G) of the channel and rate. In all practical cases, beyond the steep throughput scenario, we have actually observed that the graph G as constructed above had enough edges to guarantee the graphical unimodality of the throughput, see Section VIII.




Fig. 1

EXAMPLE OF A GRAPH PROVIDING UNIMODALITY OF THE THROUGHPUT.

In summary, in Scenario 3, we assume that $\theta \in \mathcal{T}^C \cap \mathcal{U}_G$. Note that there is more structure in Scenario 3 than in Scenario 2: if $\theta \in \mathcal{T}^C \cap \mathcal{U}_G$, then for any $c, \theta_c \in \mathcal{T} \cap \mathcal{U}$.

Remark 1: Note that the sharp decrease of the throughput when the rate increases may not hold in some scenarios as observed in several papers. The sharp transition of the reception rate motivates the use of graphical unimodality, but the latter is more general, and may hold even in absence of this sharp transition. Actually, we are free to design the graph G , and adapt its topology depending on the specificity of the radio environment so as to get graphical unimodality.

Remark 2: The notion of graphical unimodality is generic. Our approach consists in constructing a minimal graph G combining structural properties satisfied by the throughput as a function of the (channel, rate) pair, and results from experiments run off-line, so that the throughput is graphically unimodal w.r.t. G . It should be observed that in MIMO systems (e.g. as in 802.11n and subsequent standards), the throughput is no longer a unimodal function of the rate (due to the different

available MIMO modes, with single or multiple streams). The proposed framework can be adapted to account for the various MIMO modes. To this aim, the set of decisions would correspond to the set of all possible (channel, MIMO mode, rate) triplets, and the graph G would be constructed so that the throughput is a graphically unimodal function of these triplets.

IV. OBJECTIVES AND MULTI-ARMED BANDITS

Our goal is to devise a sequential (channel, rate) selection scheme that maximizes the number of packets successfully transmitted over a finite time horizon. Such a design can be formulated as an online stochastic optimization problem. The choice of the time horizon, denoted by T , is not really important as long as during time interval T , a large number of packets can be sent – so that inferring the success transmission probabilities efficiently is possible.

Consider a rate adaption scheme $\pi \in \Pi$ that selects (channel, rate) pair $(c^\pi(t), k^\pi(t))$ for the t -th packet transmission. The number of packets $\gamma^\pi(T)$ that have been successfully sent under algorithm π up to time T is: $\gamma^\pi(T) = \sum_{c,k} \sum_{i=1}^{s_{ck}^\pi(T)} X_{ck}(i)$, where $s_{ck}^\pi(T)$ is the number of transmission attempts on channel c at rate r_k before time T . The $s_{ck}(T)$'s are random variables (since the rates selected under π depend on the past random successes and failures), and satisfy the following constraint:

$$\sum_{c,k} s_{ck}^\pi(T) \times \frac{1}{r_k} \leq T.$$

Wald's lemma implies that the expected number of packets successfully sent up to time T is: $\mathbb{E}[\gamma^\pi(T)] = \sum_{c,k} \mathbb{E}[s_{c,k}^\pi(T)] \theta_{ck}$. Thus, our objective is to design an algorithm solving the following online stochastic optimization problem:

$$\begin{aligned} \max_{\pi \in \Pi} \quad & \sum_{c,k} \mathbb{E}[s_{c,k}^\pi(T)] \theta_{ck}, \\ \text{s.t.} \quad & \sum_{c,k} s_{c,k}^\pi(T) \times \frac{1}{r_k} \leq T, \quad \forall c, k, s_{c,k}^\pi(T) \in \mathbb{N}. \end{aligned} \quad (1)$$

A. An equivalent Multi-Armed Bandit (MAB) problem

Next we show that the above online stochastic optimization problem is equivalent to a Multi-Armed Bandit (MAB) problem.

1) *An alternative system:* Without loss of generality, we assume that time can be divided into slots whose durations are such that for any k , the time it takes to transmit one packet at rate r_k corresponds to an integer number of slots. Under this convention, the optimization

problem (1) can be written as:

$$\begin{aligned} \max_{\pi \in \Pi} \quad & \sum_{c,k} \mathbb{E}[t_{ck}^\pi(T)] r_k \theta_{ck}, \\ \text{s.t.} \quad & \sum_{c,k} t_{ck}^\pi(T) \leq T, \\ & \forall c, k, t_{ck}^\pi(T) \in \frac{1}{r_k} \mathbb{N} = \left\{ \frac{u}{r_k}, u \in \mathbb{N} \right\}, \end{aligned} \quad (2)$$

where $t_{ck}^\pi(T) = s_{ck}^\pi(T)/r_k$ represents the amount of time (in slots) that the transmitter spends, before T , on sending packets on channel c at rate r_k . The constraint $t_{ck}^\pi(T) \in \frac{1}{r_k} \mathbb{N}$ indicates that when a rate is selected, this rate selection remains the same for the next $1/r_k$ slots. By relaxing this constraint, we obtain an optimization problem corresponding to a MAB problem. Indeed, consider now an alternative system where rate selection is made *every* slot. If at any given slot, (channel, rate) pair (c, k) is selected for the i -th times, then if $X_{ck}(i) = 1$, the transmitter successfully sends r_k bits in this slot, and if $X_{ck}(i) = 0$, then no bit are received. A (channel, rate) selection algorithm then decides in each slot which (channel, rate) pair to use. There is a natural mapping between rate selection algorithms in the original system and in the alternative system: let $\pi \in \Pi$, if for the t -th packet transmission, rate r_k is selected under π in the original system, then π selects the same rate r_k in the t -th slot.

For the alternative system, the objective is to design $\pi \in \Pi$ solving the following optimization problem, which can be interpreted as a relaxation of (2).

$$\begin{aligned} \max_{\pi \in \Pi} \quad & \sum_{c,k} \mathbb{E}[t_{ck}^\pi(T)] r_k \theta_{ck}, \\ \text{s.t.} \quad & \sum_{c,k} t_{ck}^\pi(T) \leq T, \\ & \forall c, k, t_{ck}^\pi(T) \in \mathbb{N}. \end{aligned} \quad (3)$$

The above optimization problem corresponds to a MAB problem, where in each slot a decision is taken (i.e., a channel and a rate are selected), and where when (c, k) is chosen, the obtained reward is r_k with probability θ_{ck} and 0 with probability $1 - \theta_{ck}$.

2) *Regrets*: We quantify the performance of an algorithm $\pi \in \Pi$ in both original and alternative systems through the notion of *regret*. The regret up to slot T compares the performance of π to that achieved by an algorithm always selecting the best (channel, rate) pair. If the parameter $\theta = (\theta_{ck}, c, k)$ was known, then in both systems, it would be optimal to always select (channel, rate) pair (c^*, k^*) . The regret of algorithm π up to time slot T in the original system is then defined by:

$$R_1^\pi(T) = \theta_{c^*k^*} \lceil r_{k^*} T \rceil - \sum_{c,k} \theta_{ck} \mathbb{E}[s_{ck}^\pi(T)],$$

where $\lceil x \rceil$ denotes the largest integer smaller than x .

The regret of algorithm π up to time slot T in the

alternative system is similarly defined by:

$$R^\pi(T) = \theta_{c^*k^*} r_{k^*} T - \sum_{c,k} \theta_{ck} r_k \mathbb{E}[t_{ck}^\pi(T)].$$

3) *Asymptotic equivalence*: In the next section, we show that an asymptotic lower bound for the regret $R^\pi(T)$ is of the form $c(\theta) \log(T)$ where $c(\theta)$ is a strictly positive constant that we can explicitly characterize. It means that for all $\pi \in \Pi$, $\liminf_{T \rightarrow \infty} R^\pi(T)/\log(T) \geq c(\theta)$. It can be also shown that there exists an algorithm $\pi^* \in \Pi$ that actually achieves this lower bound in the alternative system, in the sense that $\limsup_{T \rightarrow \infty} R^{\pi^*}(T)/\log(T) \leq c(\theta)$. In such a case, we say that π^* is asymptotically optimal. The following proposition states that actually, the same lower bound is valid in the original system, and that any asymptotically optimal algorithm in the alternative system is also asymptotically optimal in the original system.

Proposition 1: Let $\pi \in \Pi$. For any $\beta > 0$, we have:

$$\liminf_{T \rightarrow \infty} \frac{R^\pi(T)}{\log(T)} \geq \beta \implies \liminf_{T \rightarrow \infty} \frac{R_1^\pi(T)}{\log(T)} \geq \beta,$$

and

$$\limsup_{T \rightarrow \infty} \frac{R^\pi(T)}{\log(T)} \leq \beta \implies \limsup_{T \rightarrow \infty} \frac{R_1^\pi(T)}{\log(T)} \leq \beta.$$

Proof. Let $T > 0$. By time T , we know that there have been at least $\lfloor Tr_1 \rfloor$ transmissions, but no more than $\lceil Tr_K \rceil$. Also observe that both regrets R^π and R_1^π are increasing functions of time. We deduce that:

$$R^\pi(\lfloor Tr_1 \rfloor) \leq R_1^\pi(T) \leq R^\pi(\lceil Tr_K \rceil).$$

Now

$$\begin{aligned} \liminf_{T \rightarrow \infty} \frac{R_1^\pi(T)}{\log(T)} &\geq \liminf_{T \rightarrow \infty} \frac{R^\pi(\lfloor Tr_1 \rfloor)}{\log(T)} \\ &= \liminf_{T \rightarrow \infty} \frac{R^\pi(\lfloor Tr_1 \rfloor)}{\log(\lfloor Tr_1 \rfloor)} \geq \beta. \end{aligned}$$

The second statement can be derived similarly. \square

B. MAB problems

Instead of trying to solve (1), we rather focus on analyzing the MAB problem (3). We know that optimal algorithms for (3) will also be optimal for the original problem. The nature of the MAB problem (3) depends on the structural assumption made on the successful transmission probabilities θ . In absence of such assumption (Scenario 1), we get a classical MAB problem where the rewards provided by all decisions are independent. In Scenarios 2 and 3, we get a *structured* MAB problem, as we know a priori that θ belongs to a structured set, which helps learning the best (channel, rate) pair. Next

we summarize the MAB problems obtained in the three different scenarios.

We have a set $\{1, \dots, C\} \times \{1, \dots, K\}$ of possible decisions (i.e., (channel, rate) pairs). If decision (c, k) is taken for the i -th time, we receive a reward $r_k X_{ck}(i)$. ($X_{ck}(i), i = 1, 2, \dots$) are i.i.d. with Bernoulli distribution with mean θ_{ck} . The objective is to design a decision scheme minimizing the regret $R^\pi(T)$ over all possible algorithms $\pi \in \Pi$. The three MAB problems differ depending on the structural assumptions made on θ .

Unstructured MAB (P_I). No assumption is made on θ : $\theta \in [0, 1]^{C \times K}$.

Structured MAB (P_U). We assume that $\theta_c \in \mathcal{T} \cap \mathcal{U}$ for all channel c .

Structured MAB (P_{GU}). We assume that $\theta \in \mathcal{T}^C \cap \mathcal{U}_G$.

V. REGRET LOWER BOUNDS

In this section, we derive an asymptotic (as T grows large) lower bound of the regret $R^\pi(T)$ satisfied by any algorithm $\pi \in \Pi$ in the three MAB bandit problems (P_I), (P_U), and (P_{GU}). These lower bounds provide insightful theoretical performance limits satisfied by any (channel, rate) selection scheme. By comparing the lower bounds derived for the three problems, we also quantify the performance gains that can be achieved by smartly exploiting the (a priori) known structure.

A. Unstructured MAB (P_I)

The regret lower bound for MAB problem (P_I) can be derived using the direct technique used by Lai and Robbins [6]. Note that the only difference between (P_I) and the classical MAB problems [6] lies in the fact that in (P_I), we know that the average reward of decision (c, k) is of the form $r_k \theta_{ck}$ for known r_k . The analysis of (P_I) is then similar to that of classical bandit problems.

We first introduce the notion of uniformly good algorithms. An algorithm π is uniformly good, if for all parameters θ , for any $\alpha > 0$, we have¹: $\mathbb{E}[t_{ck}^\pi(T)] = o(T^\alpha)$, $\forall (c, k) \neq (c^*, k^*)$, where $t_{ck}^\pi(T)$ is the number of times (channel, rate) pair (c, k) has been chosen up to the T -th decision, and (c^*, k^*) is the optimal channel and rate pair (it depends on θ). Uniformly good algorithms exist as we shall see later on.

Let $N = \{k : \mu^* \leq r_k\}$ – note that N depends on θ . There exists k_0 such that $N = \{k_0, \dots, K\}$, with the convention $k_0 = K + 1$ if $N = \emptyset$. Note that if $k < k_0$, then for any channel c , $r_{k^*} \theta_{ck^*} > r_k$, which means that even if all transmissions at rate r_k on channel c were

successful, i.e., $\theta_{ck} = 1$, rate r_k would be sub-optimal. Hence, there is no need to select rate r_k to discover this fact, since by only selecting rate r_{k^*} on channel c^* , we get to know whether $r_{k^*} \theta_{c^* k^*} > r_k \geq r_k \theta_k$.

Finally, we introduce the Kullback-Leibler (KL) divergence, a well-known measure for dissimilarity between two distributions. When we compare two Bernoulli distributions with respective averages p and q , the KL divergence is: $I(p, q) = p \log \frac{p}{q} + (1-p) \log \frac{1-p}{1-q}$.

Theorem 1: Let $\pi \in \Pi$ be a uniformly good rate selection algorithm for MAB problem (P_I). We have: $\liminf_{T \rightarrow \infty} \frac{R^\pi(T)}{\log(T)} \geq c_I(\theta)$, where

$$c_I(\theta) = \sum_{k=k_0: k \neq k^*}^K \frac{\mu^* - r_k \theta_{c^* k}}{I(\theta_{c^* k}, \frac{\mu^*}{r_k})} + \sum_{c \neq c^*} \sum_{k=k_0}^K \frac{\mu^* - r_k \theta_{ck}}{I(\theta_{ck}, \frac{\mu^*}{r_k})}.$$

The proof of the previous theorem is similar to that of the regret lower bound in [6], and is omitted here. In view of this result, if we do not exploit structural properties of the problem, then the regret of any algorithm scales at least as $CK \log(T)$. Hence, when the number of channels and rates grow large, no algorithm is able to learn the best (channel, rate) pair rapidly and efficiently.

B. Structured MAB (P_U)

To derive a regret lower bound for MAB problem (P_U), we need to introduce additional notations. We define $M = N \cap \{k^* - 1, k^* + 1\}$. For any channel c , let $N_c = \{k : \mu_c^* \leq r_k\}$, and k_{0c} such that $N_c = \{k_{0c}, \dots, K\}$, with the convention $k_{0c} = K + 1$ if $N_c = \emptyset$. Observe that for any $c \neq c^*$, $k_{0c} \leq k_0$. Let $M_c = N_c \cap \{k_c^* - 1, k_c^* + 1\}$. The following theorem is proved in [22]. Due to space limitations, all the proofs are presented in [22].

Theorem 2: Let $\pi \in \Pi$ be a uniformly good (channel, rate) selection algorithm for MAB problem (P_U). We have: $\liminf_{T \rightarrow \infty} \frac{R^\pi(T)}{\log(T)} \geq c_U(\theta)$, where $c_U(\theta)$ is the optimal value of the following optimization problem:

$$\begin{aligned} & \inf_{\alpha_{ck} \geq 0, \forall k, c} \sum_{(c, k) \neq (c^*, k^*)} \alpha_{ck} (\mu^* - \mu_{ck}) \\ & \text{s.t. } \forall k \in M, \alpha_{c^* k} I(\theta_{c^* k}, \frac{\mu^*}{r_k}) \geq 1, \\ & \forall c \neq c^* : k_c^* \geq k_0, \alpha_{c k_c^*} I(\theta_{c k_c^*}, \frac{\mu^*}{r_{k_c^*}}) \geq 1, \text{ and} \\ & \forall k \geq k_0, k \neq k^*, \inf_{\lambda_c \in C_k} \sum_l \alpha_{cl} I(\theta_{cl}, \lambda_{cl}) \geq 1, \end{aligned}$$

where $C_k = \{\lambda_c \in \mathcal{U} \cap \mathcal{T} : r_k \lambda_{ck} > \mu^*\}$.

The above theorem does not provide a fully explicit regret lower bound. In particular, it remains unclear how

¹ $f(T) = o(g(T))$ means that $\lim_{T \rightarrow \infty} f(T)/g(T) = 0$.

this lower bound scales with the numbers of rates and channels. In the following theorem, we further exploit the structural properties of the MAB problem (P_U) to show that $c_U(\theta)$ scales at most linearly with the number of channels, and does not scale with the number of rates.

Theorem 3: We have $c_U(\theta) \leq c'_U(\theta)$ where

$$\begin{aligned} c'_U(\theta) &= \sum_{k \in M} \frac{\mu^* - \mu_{c^*k}}{I(\theta_{c^*k}, \frac{\mu^*}{r_k})} \\ &+ \sum_{c \neq c^*} \left[\frac{\mu^* - \mu_{ck_c^*}}{\min\{I(\theta_{ck_c^*}, \frac{\mu^*}{r_{k_c^*}}), I(\theta_{ck_c^*}, \theta_{ck_c^*} - \frac{\delta_c}{r_{k_c^*}})\}} \right] \\ &+ \sum_{k \in M_c} \frac{\mu^* - \mu_{ck}}{I(\theta_{ck}, \theta_{ck} + \frac{\delta_c}{r_k})}. \end{aligned} \quad (4)$$

and

$$\delta_c = \min_{k \in \{k_c^* - 1, k_c^* + 1\}} (\mu_{ck_c^*} - \mu_{ck}) / 2.$$

In particular, $c'_U(\theta)$ is proportional to the number of channels and independent of the number of rates.

From the above analysis, we conclude that the minimum regret for the MAB problem (P_U) scales at most as $3C \log(T)$. Hence we expect that exploiting the structure of the problem (the fact that $\theta_c \in \mathcal{T} \cap \mathcal{U}$ for any channel c) may significantly improve the system performance. Indeed we expect a regret that does not depend on the number of available rates. In the next section, we design an algorithm with such a regret.

C. Structured MAB problem (P_{GU})

Graphical unimodal bandit problems have been recently studied in [23], [24]. A regret lower bound is derived in [24]. The only difference between our graphically unimodal MAB problem and those considered in [24] is that we consider directed graphs, but the analysis is similar. We use here the notation introduced in §III-B.3, and recall that $N = \{k : \mu^* \leq r_k\}$. For any (c, k) , we define $\mathcal{N}'(c, k) = \mathcal{N}(c, k) \cap N$. $\mathcal{N}'(c, k)$ is the set of (channel, rate) pairs that are neighbors of vertex (c, k) , and that need to be explored if one wants to know whether they provide better throughput than (c, k) .

Theorem 4: [24] Let $\pi \in \Pi$ be a uniformly good (channel, rate) selection algorithm for MAB problem (P_{GU}). We have:

$$\liminf_{T \rightarrow +\infty} \frac{R^\pi(T)}{\log(T)} \geq c_{GU}(\theta), \quad (5)$$

where

$$c_{GU}(\theta) = \sum_{(c,k) \in \mathcal{N}'(c^*, k^*)} \frac{\mu^* - \mu_{ck}}{I(\theta_{ck}, \frac{\mu^*}{r_k})}.$$

In view of the above theorem, for the MAB problem (P_{GU}), the minimum regret scales as $\gamma \log(T)$, where γ

is γ is the maximum node degree in the graph G . Note that for our graph G , $\gamma \leq 2C$. Hence, by exploiting the graphical unimodal structure, we may expect to design algorithms whose regret does not depend on the number of available rates. In the next section, an algorithm whose regret matches the lower bound of Theorem 4 is proposed.

In this section, we have shown that the regret lower bound can be significantly improved when structural assumptions are made, i.e., $c_{GU}(\theta) \leq c_U(\theta) \leq c_I(\theta)$. By exploiting the structure of the problem, we may actually design algorithms whose regrets does not depend on the number of available rates. Such algorithms do not exist when the structure is not exploited (see Theorem 1).

VI. ALGORITHMS

In this section, we present algorithms for the three MAB problems (P_I), (P_U), and (P_{GU}), and analyze their regrets. For the two structured MAB problems, the proposed algorithms exhibit a regret that does not depend on the number of available rates.

A. The KL-UCB algorithm for MAB problem (P_I)

Classical unstructured bandit problems have been extensively studied in the past, and numerous efficient algorithms have been proposed. We build on this previous work, and present a simple extension of KL-UCB algorithm [21] to the MAB problem (P_I). This algorithm does not exploit any structural properties, and is asymptotically optimal: its regret matches the lower bound derived in Theorem 1.

Under the KL-UCB algorithm, each (channel, rate) pair (c, k) is associated with an index $q_{ck}(n)$ for the $(n+1)$ -th packet transmission:

$$q_{ck}(n) = \max\{q \in [0, r_k] : t_{ck}(n) I\left(\frac{\hat{\mu}_{ck}(n)}{r_k}, \frac{q}{r_k}\right) \leq \log(n) + 3 \log \log(n)\},$$

where $t_{ck}(n)$ denotes the number of times (c, k) has been selected up to the n -th transmission, and

$$\hat{\mu}_{ck}(n) = \frac{1}{t_{ck}(n)} \sum_{i=1}^{t_{ck}(n)} r_k X_{ck}(i),$$

is the empirical throughput or reward of (channel, rate) pair (c, k) up to the n -th transmission. The algorithm selects the (channel, rate) pair with highest index:

Algorithm 1 KL-UCB

For $n = 0, \dots, CK - 1$ (initialization): for the $(n+1)$ -th transmission, select (channel, rate) pair $(c, k)(n+1) =$

$(c' + 1, k' + 1)$ where $n = Kc' + k'$, $k' \in \{0, \dots, K - 1\}$. For $n \geq CK$, for the $(n + 1)$ -th transmission, select $(c, k)(n + 1)$ where $(c, k)(n + 1) \in \arg \max_{(c, k)} q_{ck}(n)$.

The rationale behind the design of KL-UCB is the same as that of the classical UCB algorithm. We construct an index for each (channel, rate) pair, which in turn constitutes an upper confidence bound of the corresponding throughput. By selecting the pair with the highest upper confidence bound, we force the exploration of suboptimal pairs if the latter have not been explored enough (in such a case, the upper confidence bound of a suboptimal pair can be higher than that of the optimal pair). KL-UCB is designed so that the number of times a suboptimal pair is selected matches the optimal number of times it is explored in the regret lower bound. In fact, KL-UCB is known to be asymptotically optimal in classical bandit problems [21]. It can be easily established that its extension is also optimal for the problem (P_I) :

Theorem 5: For any $\theta \in [0, 1]^{C \times K}$, the regret of the $\pi = \text{KL-UCB}$ algorithm satisfies:

$$\limsup_{T \rightarrow \infty} \frac{R^\pi(T)}{\log(T)} \leq c_I(\theta),$$

In particular, the regret under KL-UCB scales linearly with the numbers of channels and rates. When the later become large, the performance of KL-UCB can be quite poor.

Note that one may actually derive finite-time upper bounds on the regret of KL-UCB, as done in [21]. KL-UCB is asymptotically optimal, but also provides good performance over a finite time horizon.

Finally, regarding the computational complexity of implementing KL-UCB, note that we just require to maintain an index for each pair, which requires a number of operations that scales as CK after each transmission. The comparison between the various indexes can be done with $CK \log(CK)$ operations.

B. The CRS-T algorithm for MAB problem (P_U)

Next, we present CRS-T (Channel and Rate Sampling with Tests), an algorithm that exploits the structure of the MAB problem (P_U) , i.e., the fact that on each channel, the throughput is a unimodal function of the rate. To describe our algorithm, we introduce the following notations. After the n -th transmission, the rate with the highest average empirical throughput on channel c is referred to as the *leader* on channel c , and is $l_c(n) = \arg \max_k \hat{\mu}_{ck}(n)$. The *global leader* $l(n)$ is the (channel, rate) pair with highest average empirical throughput: $l(n) = \arg \max_{(c, k)} \hat{\mu}_{ck}(n)$.

Similar to the KL-UCB index $q_{ck}(n)$, we define the lower confidence bound $\underline{q}_{ck}(n)$ as:

$$\underline{q}_{ck}(n) = \min\{q \in [0, r_k] : t_{ck}(n) I\left(\frac{\hat{\mu}_{ck}(n)}{r_k}, \frac{q}{r_k}\right) \leq \log(n) + 3 \log \log(n)\},$$

We introduce a statistical test, which will be used to assess whether the leader on channel c , $l_c(n)$, provides a larger reward than its neighbors $l_c(n) - 1$, $l_c(n) + 1$ on the same channel. Define the test for channel c at time n through $U_c(n)$:

$$U_c(n) = \mathbf{1}\{\underline{q}_{cl_c(n)}(n) \geq \sup_{k: |k - l_c(n)| = 1} q_{ck}(n)\}$$

The test can be interpreted as follows. $U_c(n) = 1$ means that $l_c(n)$ is better than its neighbors with high probability, and $U_c(n) = 0$ means that we do not have enough samples to determine whether $l_c(n)$ is better than its neighbors. After the n -th packet transmission, we define $\mathcal{U}(n) = \{c : U_c(n) = 0\}$ the set of channels for which we cannot determine whether the leader $l_c(n)$ corresponds to the best rate k_c^* on this channel.

The sequential decisions under the CRS-T algorithm are based on the indexes of the various (channel, rate) pairs, and can be easily implemented. The index $b_k(n)$ of decision (c, k) for the $(n + 1)$ -th packet transmission is:

$$b_{ck}(n) = q_{ck}(n) \mathbf{1}\{k = l_c(n)\},$$

where $q_{ck}(n)$ is the index used in the KL-UCB algorithm. Note that the index of decision (c, k) is equal to 0 if k is not the leader on channel c . The pseudo-code for CRS-T is given below (each time the decision is ambiguous, ties are broken arbitrarily).

Algorithm 2 CRS-T

For $n = 0, \dots, CK - 1$ (initialization): for the $(n + 1)$ -th transmission, select (channel, rate) pair $(c, k)(n + 1) = (c' + 1, k' + 1)$ where $n = Kc' + k'$, $k' \in \{0, \dots, K - 1\}$. For $n \geq CK$: for the $(n + 1)$ -th transmission, select $(c, k)(n + 1)$ where

- if $\mathcal{U}(n) \neq \emptyset$, then $c(n + 1) \in \mathcal{U}(n)$ and $k(n + 1) \in \arg \min_{k': |k' - l_{c(n+1)}(n)| \leq 1} t_{c(n+1)k'}(n)$;
 - else $(c, k)(n + 1) \in \arg \max_{c', k'} b_{c'k'}(n)$.
-

The design of the CRS-T algorithm is motivated by the following objectives: (1) For all channels, we need to play the leader and all its neighbours until we can determine with high probability that the leader $l_c(n)$ is the best rate k_c^* . (2) Once we have determined that $l_c(n)$ is k_c^* for all channels, then we play the leader of the channel with the largest index, i.e we apply KL-UCB restricted to the set of leaders.

Define $\Delta = \min_{(c,k)} \min_{k':|k'-k|=1} |\mu_{ck} - \mu_{ck'}|$ the minimal separation between two neighboring rates on any channel and $\tilde{\mu}_c = (\mu_{ck_c^*} + \max_{|k-k_c^*|=1} \mu_{ck})/2$. The next theorem, proved in [22], provides a finite time upper bound on the regret under CRS-T. The asymptotic regret $\log(T)c^{\text{CRS-T}}(\theta)$ scales linearly with the number of channels, but is independent of the number of available rates. In particular, CRS-T exploits the structure of the MAB problem (P_U).

Theorem 6: For any θ such that for all $c, \theta_c \in \mathcal{T} \cap \mathcal{U}$, and for all $\epsilon > 0$, the regret of the CRS-T algorithm satisfies:

$$R^{\text{CRS-T}}(T) \leq (1 + \epsilon)c^{\text{CRS-T}}(\theta) \log(T) + \Gamma_\epsilon(\Delta^{-2} + \log(\log(T))),$$

where $\Gamma_\epsilon > 0$ depends on ϵ, C, K and \mathcal{R} but not on θ , and

$$c^{\text{CRS-T}}(\theta) = \sum_{c=1}^C \tau_c^{-1} \sum_{k:|k-k_c^*|\leq 1} (\mu^* - \mu_{ck}),$$

with

$$\tau_c = \min \left(\min_{k:|k-k_c^*|\leq 1} I(\theta_{ck}, \tilde{\mu}_c/r_k), I(\theta_{ck_c^*}, \mu^*/r_{k_c^*}) \right)$$

The regret under CRS-T does not depend on the number of available rates, and hence exploits (at least asymptotically) the unimodal structure of (P_U). The computational complexity of CRS-T is similar to that of KL-UCB because it essentially requires to maintain the indexes of the various channel and rate pairs: it scales linearly with CK (up to a logarithmic factor).

C. The KL-UCB-U algorithm for MAB problem (P_{GU})

Finally, we present KL-UCB-U, an algorithm for MAB problem (P_{GU}). KL-UCB-U is a natural extension of an algorithm proposed in [24] for graphically unimodal bandits with undirected graphs. This algorithm is asymptotically optimal (its regret matches the lower bound derived in Theorem 4).

Recall that the global leader is denoted by $l(n)$ before the $(n+1)$ -th transmission. We introduce $v_{(c,k)}(n)$ the number of times that (channel, rate) pair (c, k) has been the global leader up to the n -th transmission: $v_{(c,k)}(n) = \sum_{n'=1}^n \mathbf{1}\{l(n') = (c, k)\}$. The index associated with decision (c, k) before the $(n+1)$ -th transmission is:

$$b_{ck}(n) = \max \left\{ q \in [0, r_k] : t_{ck}(n) I\left(\frac{\hat{\mu}_{ck}(n)}{r_k}, \frac{q}{r_k}\right) \leq \log(v_{l(n)}(n)) + 3 \log(\log(v_{l(n)}(n))) \right\},$$

For the $(n+1)$ -th transmission, KL-UCB-U selects the (channel, rate) pair in the neighborhood of the leader with maximum index. Ties are broken arbitrarily.

Algorithm 3 KL-UCB-U

For $n = 0, \dots, CK - 1$ (initialization): for the $(n+1)$ -th transmission, select (channel, rate) pair $(c, k)(n+1) = (c'+1, k'+1)$ where $n = Kc' + k', k' \in \{0, \dots, K-1\}$. For $n \geq CK$: for the $(n+1)$ -th transmission, select $(c, k)(n+1)$ where:

$$(c, k)(n+1) = \begin{cases} l(n) & \text{if } (v_{l(n)}(n) - 1)/\gamma \in \mathbb{N}, \\ \arg \max_{(c,k) \in \mathcal{N}(l(n))} b_{ck}(n) & \text{otherwise.} \end{cases}$$

Remember that γ is the maximum number neighbors in G of a given (channel, rate) pair. The KL-UCB-U algorithm periodically selects the leader to make sure that the latter is often selected. The design of KL-UCB-U is based on the lower regret bound derived for the MAB problem (P_{GU}). This lower bound implies that an optimal algorithm explores suboptimal (channel, rate) pairs a number of times that scales with $\log(T)$ only for pairs that are neighbours of the the optimal pair in the graph G . Hence in KL-UCB-U, the exploration is restricted to the neighbours of the current leader in G . As in [24], we can establish that KL-UCB-U is asymptotically optimal:

Theorem 7: For any $\theta \in \mathcal{T}^C \cap \mathcal{U}_G$, the regret of $\pi = \text{KL-UCB-U}$ satisfies:

$$\limsup_{T \rightarrow \infty} \frac{R^\pi(T)}{\log(T)} \leq c_{GU}(\theta),$$

In particular, KL-UCB-U optimally exploits the structure of MAB problem (P_{GU}). In turn, if the throughput is a graphically unimodal function of the (channel, rate) pair, then KL-UCB-U asymptotically outperforms any other algorithm, and in particular CRS-T, an algorithm designed to exploit the unimodal structure per channel only.

Note that a finite-time regret analysis for KL-UCB-U is possible as shown in [24]. KL-UCB-U is asymptotically optimal, but also provides good performance over a finite time horizon. Finally, again, the computational complexity of KL-UCB-U is similar to that of KL-UCB.

VII. NON-STATIONARY RADIO ENVIRONMENTS

In practice, channel conditions may be non-stationary, i.e., the success probabilities at various (channel, rate) pair could evolve over time. In many situations, the evolution over time is rather slow – refer to [3] and

to Section V for test-bed measurements. These slow variations allow us to devise (channel,rate) adaptation schemes that efficiently track the best (channel,rate) pair for transmission.

We assume that for all (c, k) pairs, the transmissions outcomes $X_{ck}(n)$, $n = 1, 2, \dots$ are independent, with expectation $\theta_{ck}(n) = \mathbb{E}[X_{ck}(n)]$. At time n we define the throughput of (c, k) $\mu_{ck}(n) = r_k \theta_{ck}(n)$, the best throughput $\mu^*(n) = \max_{c,k} \mu_{ck}(n)$ and the optimal decision $(c^*, r^*)(n) = \arg \max_{c,k} \mu_{ck}(n)$.

Any algorithm designed for stationary radio environments can readily be extended to non-stationary environments. These extensions are obtained by replacing empirical averages by averages over a sliding time window. Let $\tau \geq 1$ denote the sliding window size, and define the empirical reward $\hat{\mu}_{ck}(n)$ as:

$$\hat{\mu}_{ck}^\tau(n) = \frac{r_k}{t_{ck}^\tau(n)} \sum_{n'=n-\tau+1}^n X_{ck}(n') \mathbf{1}\{(c, k)(n') = (c, k)\},$$

where

$$t_{ck}^\tau(n) = \sum_{n'=n-\tau+1}^n \mathbf{1}\{(c, k)(n') = (c, k)\},$$

with the convention $\hat{\mu}_{ck}^\tau(n) = 0$ if $t_{ck}^\tau(n) = 0$. We also define the upper confidence index of (channel, rate) pair (c, k) as:

$$q_{ck}^\tau(n) = \max\{q \in [0, r_k] : I\left(\frac{\hat{\mu}_{ck}^\tau(n)}{r_k}, \frac{q}{r_k}\right) \leq \log(\tau) + 3 \log(\log(\tau))\}.$$

We define sliding window variants of the algorithms presented in Section VI by replacing $t_{ck}(n)$ by $t_{ck}^\tau(n)$, $\hat{\mu}_{ck}(n)$ by $\hat{\mu}_{ck}^\tau(n)$ and $q_{ck}(n)$ by $q_{ck}^\tau(n)$. For instance, KL-UCB with sliding window is the algorithm which selects $(c, k)(n) \in \arg \max_{c,k} q_{ck}^\tau(n)$.

In [25], the authors show that algorithms with sliding windows efficiently track the best decision over time provided that the environment evolves relatively slowly. This is confirmed in [24], where the performance of algorithms similar to KL-UCB and KL-UCB-U with sliding window is analyzed. Due to space limitation, we skip this analysis; refer to [24] for more details.

The way the window size τ should be chosen in practice is dictated by the following remarks. First, τ should be relatively small compared to the time it takes for the packet successful transmission probabilities to evolve. This ensures that the algorithms with sliding window track the best channel and rate pair. Then, τ should be sufficiently large so that the throughput of the various channel and rate pairs could be estimated with precision using samples collected in an interval of time of duration

τ . Typically 10 or 20 packets sent using the same channel and rate pair is enough. These two requirements for τ go in opposite directions, and clearly when the packet successful transmission probabilities evolve rapidly, an appropriate design of τ is not possible. In practice, τ may be tuned by just observing the way these probabilities evolve (via off-line preliminary experiments). It is not necessary to adapt τ over time if the radio environment does not change (say for example, users stay inside a building). However, in other scenarios, we may need to adapt τ .

VIII. NUMERICAL EXPERIMENTS


In this section we numerically illustrate the performance of the proposed (channel, rate) selection algorithms. To this aim, we conduct simulation experiments where our algorithms are tested against channel quality traces that are either extracted from a test-bed [3], or artificially generated. In the latter case, we generate traces based on a widely used statistical model for radio propagation and on a mapping between channel quality and probability of packet successful transmission on a given (channel,rate) pair [5].

A. Traces extracted form a test-bed

In this subsection we present trace-driven experiments using the test-bed described in [3]. The test-bed is based on a SDR platform (Lyrtech SFF-SDR), and is located in an indoor office. The PHY layer is OFDM, as in 802.11a/g/n. There are 3 available rates $\{4.5, 6, 6.75\}$ Mbps corresponding to QPSK modulation with respective coding rates $\{1/2, 2/3, 3/4\}$. We consider 5 channels in the UHF band centred at $\{510, 530, 550, 580, 600\}$ Mhz . The bandwidth of each channel is 10 Mhz, and the packet size is 1500 bytes. The trace duration is 600s.

The traces are collected as follows. The transmitter transmits 10 packets at each rate on a given channel before moving to the next channel. Each measurement round (where each (channel, rate) pair is probed) hence consists of the transmission of $10 \times 5 \times 3$ packets, and lasts less than 1s. This means we sample a given channel at a given rate once every second. In each round and for each channel, we calculate, by averaging over 10 packets, the RSSI, the successful packet transmission probabilities and the goodputs at the 3 different rates.

In Fig.2, we plot the best decision (c^*, k^*) as a function of time. the radio environment is non-stationary, and the optimal decision remains constant for several seconds. Since a packet transmission lasts about 1ms, the packet successful transmission probabilities for various




TEST-BED: BEST (CHANNEL, RATE) PAIR $(c^*, k^*)(t)$ AS A FUNCTION OF TIME.

decisions stay constant for thousands of packet transmissions. Therefore we have quite a lot of statistical information to find the best decision. Furthermore the window size used in the tested algorithms should be of the order of a few seconds – we fix it to 2s.

In Fig.3, we plot the throughput under KL-UCB and KL-UCB-U algorithms. For the sake of comparison, we also plot $\mu^*(t)$ the throughput of an Oracle algorithm that always selects the optimal decision. We also plot the throughput obtained by choosing the best static (channel, rate) pair, computed offline. We observe that selecting the best static pair is clearly sub-optimal, so that adaptive algorithms can lead to a large gain in throughput. Both decision algorithms, KL-UCB and KL-UCB-U, manage to closely follow the best (channel, rate) pair. KL-UCB-U provides a throughput equal to 95% of that obtained under the Oracle algorithm, whereas the throughput under KL-UCB is equal to 90% of that of the Oracle algorithm. There is not a huge performance gap between KL-UCB and KL-UCB-U because there are few available rates, $K = 3$. Hence KL-UCB explores $C \times K = 15$ (channel,rate) pairs, while KL-UCB-U explores (in the worse case) $2C + 1 = 11$ pairs. We will show that increasing the number of available rates K makes this difference significantly larger.

B. Artificial traces

We also present numerical results based on a widely used statistical model for radio propagation. Namely, we assume that the channel is a multi-path Rayleigh fading channel. When a signal is transmitted, several delayed



TEST-BED: THROUGHPUTS OF THE VARIOUS ALGORITHMS AS A FUNCTION OF TIME.

copies of this signal are received and the amplitude and phase of each delayed copy is an independent Rayleigh fading process. We use Jakes' model to simulate Rayleigh fading with user speed set to match the time variability of the test-bed trace presented in VIII-A. This corresponds to static users such as laptops in an office environment. The expected power of each delayed path is chosen according to the field measurements presented in [26].

We assume that OFDM is used, and the mapping between the strength of received signal on each sub-carrier and the probability of successful transmission is calculated by the method presented in [5]. We consider 5 channels with bandwidth 20 Mhz in the 2.4 GHz band centred at $\{2.4, 2.41, 2.42, 2.43, 2.44\}$ GHz, respectively. Each channel has 52 sub-carriers and the packet size is 1500 bytes. We consider 8 available rates: $\{6, 13, 19.5, 26, 39, 52, 58.5, 65\}$ Mbps, and a transmitter-receiver pair with an average SNR of 20 dB. The trace length is 600 seconds.

We first consider stationary environments, so that a snapshot of the success probabilities for all (channel,rate) pairs is drawn and kept constant throughout the simulation. Fig.4 shows the packet successful transmission probabilities and throughputs of different (channel,rate) pairs. As announced, graphical unimodality holds: the throughput on each channel is a unimodal function of the rate, and given the optimal rate k_c^* on sub-optimal channel $c \neq c^*$, there exists another channel $c' \neq c$ such that either $\mu_{c',k_c^*} > \mu_{c,k_c^*}$ or $\mu_{c',k_c^*+1} > \mu_{c,k_c^*}$. Graphical unimodality results from the fact that we are in a steep environment as defined in [14]. Fig.5 presents

r_k	6	13	19.5	26	39	52	58.5	65
$\theta_{1,k}$	1	1	1	1	1	0.2	0	0
$\theta_{2,k}$	1	1	1	1	1	1	0.7	0.1
$\theta_{3,k}$	1	1	1	1	1	0.6	0	0
$\theta_{4,k}$	0	0	0	0	0	0	0	0
$\theta_{5,k}$	1	1	0.8	0.2	0	0	0	0
$\mu_{1,k}$	6	13	19.5	26	39	13	0	0
$\mu_{2,k}$	6	13	19.5	26	39	52	41	8
$\mu_{3,k}$	6	13	19.5	26	39	29	0	0
$\mu_{4,k}$	0	0	0	0	0	0	0	0
$\mu_{5,k}$	6	13	16	6	0	0	0	0

Fig. 4

SIMULATION: PACKET SUCCESSFUL TRANSMISSION PROBABILITIES AND THROUGHPUTS AT DIFFERENT (CHANNEL,RATE) PAIRS IN A STATIONARY ENVIRONMENT.




Fig. 5

SIMULATION: REGRET OF DIFFERENT DECISION RULES AS A FUNCTION OF TIME IN A STATIONARY ENVIRONMENT.

the regret of KL-UCB and KL-UCB-U as a function of time. KL-UCB-U beats KL-UCB and for large time horizons the regret under KL-UCB-U is roughly half of that under KL-UCB. Hence exploiting the graphical unimodal structure significantly helps.

We now turn to non-stationary environments. As in VIII-A, we present the best pair as a function of time and the throughputs of different algorithms in Fig.6. Again KL-UCB-U beats KL-UCB.

For both the test-bed and simulation, the performance of KL-UCB-U is rather impressive: its throughput is at least 95% of that of the Oracle, without knowing the throughputs of the various (channel,rate) pairs beforehand. This shows that given a good decision rule, the

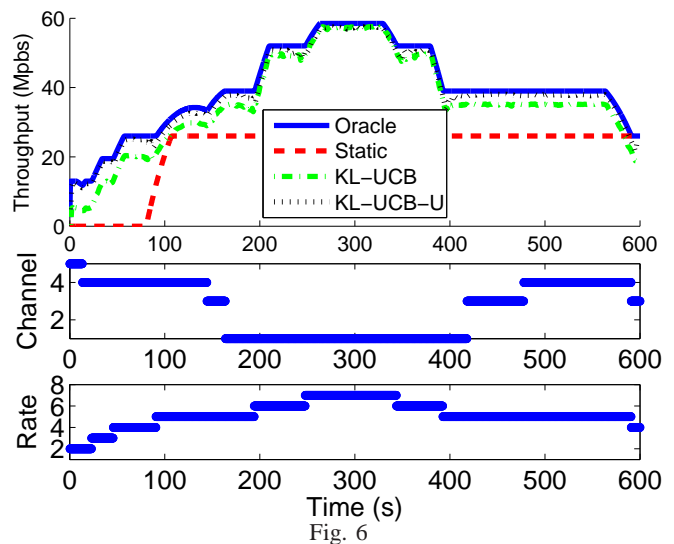


Fig. 6

SIMULATION: THROUGHPUTS OF THE VARIOUS ALGORITHMS (ABOVE) AND THE BEST PAIR (c^*, k^*) (BELOW) AS A FUNCTION OF TIME IN A NON-STATIONARY ENVIRONMENT – LOW VARIATION SPEED.

Speed	$\times 1$	$\times 20$	$\times 100$
Static	52 %	45 %	43 %
KL-UCB	90 %	83 %	57 %
KL-UCB-U	96 %	91 %	79 %
Oracle	100 %	100 %	100 %

Fig. 7

IMPACT OF THE SPEED OF VARIATION OF THE SUCCESSFUL TRANSMISSION PROBABILITIES ON PERFORMANCE IN A NON-STATIONARY ENVIRONMENT.

selection of channel and rate can be done solely based on ACK/NACK feedback with excellent performance. This is critical for real-world systems because feedback of channel measurements is problematic in practice both in terms of delay and overhead.

So far, in non-stationary environments, the packet successful transmission probabilities were evolving slowly. Next we vary the speed at which they evolve, by artificially accelerating our traces by a factor 20 and 100. Results are presented in Fig. 7. At all speeds, KL-UCB-U beats KL-UCB, and the performance gap between the two algorithms increases with the speed. When the environment changes faster, the performance of KL-UCB becomes poor, as the algorithm needs to explore all (channel, rate) pairs, and cannot track the best pair. On the contrary, KL-UCB-U exploits the structure and explores less, which makes its performance more robust.

IX. CONCLUSION

In this paper, we have addressed the problem of joint channel and rate adaptation in cognitive radio systems. We have shown that the problem is equivalent to a structured MAB problem, where the structure stems from inherent properties of the throughput as a function of the selected channel and rate. For several assumptions on this structure, we have derived fundamental performance limits satisfied by any sequential (channel, rate) selection algorithm. For each structure type, we have also proposed algorithms which are either close or achieve these limits. Finally we have assessed the efficiency of the proposed algorithms through trace-driven experiments and simulations. The two key insights from our results are: (a) The channel and rate adaptation problem has a strong structure. This structure can be exploited to devise algorithms whose performance does not depend on the number of available rates, and is close to that of an Oracle algorithm that perfectly knows the packet successful transmission probabilities at any available (channel, rate) pair. (b) There exist readily implementable algorithms which allow almost perfect channel and rate selection without the need of any measurement and explicit feedback of the quality of the various channels.

REFERENCES

- [1] FCC, "Second memorandum opinion and order, FCC 10-174," 2010.
- [2] J. Camp and E. Knightly, "Modulation rate adaptation in urban and vehicular environments: cross-layer implementation and experimental evaluation," in *Proceedings of ACM MOBICOM*, 2008.
- [3] B. Radunovic, A. Proutiere, D. Gunawardena, and P. Key, "Dynamic channel, rate selection and scheduling for white spaces," in *Proceedings of ACM CoNEXT*, 2011.
- [4] S. Sen, N. Santhapuri, R. R. Choudhury, and S. Nelakuditi, "Accurate: Constellation based rate estimation in wireless networks," in *Proceedings of USENIX NSDI*, 2010.
- [5] D. Halperin, W. Hu, A. Sheth, and D. Wetherall, "Predictable 802.11 packet delivery from wireless channel measurements," in *Proceedings of ACM SIGCOMM*, 2010.
- [6] T. Lai and H. Robbins, "Asymptotically efficient adaptive allocation rules," *Advances in Applied Mathematics*, vol. 6, no. 1, pp. 4–2, 1985.
- [7] H. Kim and K. G. Shin, "In-band spectrum sensing in cognitive radio networks: Energy detection or feature detection?" in *Proceedings of ACM MOBICOM*, 2008.
- [8] A. Motamedi and A. Bahai, "Optimal channel selection for spectrum-agile low-power wireless packet switched networks in unlicensed band," *EURASIP J. Wireless Comm. and Networking*, 2008.
- [9] K. Liu, Q. Zhao, and B. Krishnamachari, "Dynamic multichannel access with imperfect channel state detection," *IEEE Trans. Signal Process.*, vol. 58, no. 5, pp. 2795–2808, 2010.
- [10] L. Lai, H. Jiang, and H. V. Poor, "Medium access in cognitive radio networks: A competitive multi-armed bandit framework," in *Proceedings of ASILOMAR*, 2008.
- [11] L. Lai, H. El Gamal, H. Jiang, and H. Poor, "Cognitive medium access: Exploration, exploitation, and competition," *IEEE Transactions on Mobile Computing*, vol. 10, no. 2, pp. 239–253, 2011.
- [12] Q. Zhao, L. Tong, A. Swami, and Y. Chen, "Decentralized cognitive MAC for opportunistic spectrum access in ad hoc networks: A POMDP framework," *Selected Areas in Communications, IEEE Journal on*, vol. 25, no. 3, pp. 589–600, 2007.
- [13] S. Ahmad, M. Liu, T. Javidi, Q. Zhao, and B. Krishnamachari, "Optimality of myopic sensing in multichannel opportunistic access," *IEEE Transactions on Information Theory*, vol. 55, no. 9, pp. 4040–4050, 2009.
- [14] J. Bicket, "Bit-rate selection in wireless networks," Ph.D. dissertation, Massachusetts Institute of Technology, 2005.
- [15] S. H. Y. Wong, H. Yang, S. Lu, and V. Bharghavan, "Robust rate adaptation for 802.11 wireless networks," in *Proceedings of ACM MOBICOM*, 2006.
- [16] L. Deek, E. Garcia-Villegas, E. Belding, S.-J. Lee, and K. Almeroth, "Joint rate and channel width adaptation in 802.11 MIMO wireless networks," in *Proceedings of IEEE Secon*, 2013.
- [17] C. Tekin and M. Liu, "Online learning in opportunistic spectrum access: A restless bandit approach," in *Proceedings of IEEE INFOCOM*, 2011.
- [18] S. Bubeck and N. Cesa-Bianchi, "Regret analysis of stochastic and nonstochastic multi-armed bandit problems," *Foundations and Trends in Machine Learning*, vol. 5, no. 1, 2012.
- [19] L. Kocsis and C. Szepesvári, "Discounted UCB," in *Proceedings of the 23rd PASCAL Challenges Workshop*, 2006.
- [20] J. Y. Yu and S. Mannor, "Piecewise-stationary bandit problems with side observations," in *Proceedings of ICML*, 2009.
- [21] A. Garivier and O. Cappé, "The KL-UCB algorithm for bounded stochastic bandits and beyond," in *Proceedings of COLT*, 2011.
- [22] R. Combes and A. Proutiere, "Dynamic rate and channel selection in cognitive radio systems," Technical Report. [Online]. Available: <http://arxiv.org/abs/1402.5666>
- [23] J. Y. Yu and S. Mannor, "Unimodal bandits," in *Proceedings of ICML*, 2011.
- [24] R. Combes and A. Proutiere, "Unimodal bandits: Regret lower bounds and optimal algorithms," in *Proceedings of ICML*, 2014.
- [25] A. Garivier and E. Moulines, "On upper-confidence bound policies for non-stationary bandit problems," 2008. [Online]. Available: <http://arxiv.org/abs/0805.3415>
- [26] S. Sen, B. Radunovic, J. Lee, and K. H. Kim, "CSpy: Finding the best quality channel without probing," in *Proceedings of ACM MOBICOM*, 2013.
- [27] T. L. Graves and T. L. Lai, "Asymptotically efficient adaptive choice of control laws in controlled markov chains," *SIAM J. Control and Optimization*, vol. 35, no. 3, pp. 715–743, 1997.
- [28] R. Combes and A. Proutiere, "Unimodal bandits: Regret lower bounds and optimal algorithms," Technical Report, 2014. [Online]. Available: people.kth.se/~alepro/pdf/tr-icml2014.pdf

APPENDIX

PROOF OF THEOREM 2

We derive here the regret lower bounds for the MAB problem (P_U). To this aim, we apply the techniques used by Graves and Lai [27] to investigate efficient adaptive decision rules in controlled Markov chains. We recall here their general framework. Consider a controlled Markov chain $(X_t)_{t \geq 0}$ on a finite state space \mathcal{S} with a

control set U . The transition probabilities given control $u \in U$ are parametrized by θ taking values in a compact metric space Θ : the probability to move from state x to state y given the control u and the parameter θ is $p(x, y; u, \theta)$. The parameter θ is not known. The decision maker is provided with a finite set of stationary control laws $G = \{g_1, \dots, g_K\}$ where each control law g_j is a mapping from \mathcal{S} to U : when control law g_j is applied in state x , the applied control is $u = g_j(x)$. It is assumed that if the decision maker always selects the same control law g the Markov chain is then irreducible with stationary distribution π_θ^g . Now the reward obtained when applying control u in state x is denoted by $r(x, u)$, so that the expected reward achieved under control law g is: $\mu_\theta(g) = \sum_x r(x, g(x))\pi_\theta^g(x)$. There is an optimal control law given θ whose expected reward is denoted $\mu_\theta^* \in \arg \max_{g \in G} \mu_\theta(g)$. Now the objective of the decision maker is to sequentially control laws so as to maximize the expected reward up to a given time horizon T . As for MAB problems, the performance of a decision scheme can be quantified through the notion of regret which compares the expected reward to that obtained by always applying the optimal control law.

We now apply the above framework to our MAB problem. For (P_U) , for all c , the parameter θ_c takes values in $\mathcal{T} \cap \mathcal{U}$. The Markov chain has values in $\mathcal{S} = \{0, r_1, \dots, r_K\}$. The set of control laws is $G = \{1, \dots, C\} \times \{1, \dots, K\}$. These laws are constant, in the sense that the control applied by control law (c, k) does not depend on the state of the Markov chain, and corresponds to selecting (channel, rate) pair (c, k) . The transition probabilities are given as follows: for all $x, y \in \mathcal{S}$,

$$p(x, y; (c, k), \theta) = p(y; (c, k), \theta) = \begin{cases} \theta_{ck}, & \text{if } y = r_k, \\ 1 - \theta_{ck}, & \text{if } y = 0. \end{cases}$$

Finally, the reward $r(x, (c, k))$ does not depend on the state and is equal to $r_k \theta_{ck}$, which is also the expected reward obtained by always using control law (c, k) .

We now fix θ : $\forall c, \theta_c \in \mathcal{T} \cap \mathcal{U}$. Define $I^{(c,k)}(\theta, \lambda) = I(\theta_{ck}, \lambda_{ck})$ for any (c, k) . Further define the set $B(\theta)$ consisting of all *bad* parameters λ : $\forall c, \lambda_c \in \mathcal{T} \cap \mathcal{U}$ such that (c^*, k^*) is not optimal under parameter λ , but which are statistically *indistinguishable* from θ :

$$B(\theta) = \left\{ \lambda : \forall c, \lambda_c \in \mathcal{T} \cap \mathcal{U}, \right. \\ \left. \lambda_{c^*k^*} = \theta_{c^*k^*}, \max_{(c,k)} r_k \lambda_{ck} > \mu^* \right\},$$

$B(\theta)$ can be written as the union of sets $B_{ck}(\theta)$ defined as:

$$B_{ck}(\theta) = \{ \lambda \in B(\theta) : r_k \lambda_{ck} > r_{k^*} \lambda_{c^*k^*} \}.$$

Note that $B_{ck}(\theta) = \emptyset$ if $r_k < r_{k^*} \theta_{c^*k^*}$, hence if $k \notin N$.

By applying Theorem 1 in [27], we know that $c_U(\theta)$ is the minimal value of the following LP:

$$\min \sum_{c,k} \alpha_{ck} (\mu^* - r_k \theta_{ck}) \quad (6)$$

$$\text{s.t.} \quad \inf_{\lambda \in B(\theta)} \sum_{(c,k) \neq (c^*, k^*)} \alpha_{ck} I^{(c,k)}(\theta, \lambda) \geq 1, \quad (7)$$

$$\alpha_{ck} \geq 0, \quad \forall (c, k). \quad (8)$$

Next we detail the constraints (7). These constraints are equivalent to:

$$\inf_{\lambda \in B_{c^*k}(\theta)} \sum_{(c,l) \neq (c^*, k^*)} \alpha_{cl} I^{(c,l)}(\theta, \lambda) \geq 1, \quad \forall k \neq k^* \quad (9)$$

$$\inf_{\lambda \in B_{c^*k^*}(\theta)} \sum_{(c',l) \neq (c^*, k^*)} \alpha_{c'l} I^{(c',l)}(\theta, \lambda) \geq 1, \quad \forall c \neq c^* \quad (10)$$

$$\inf_{\lambda \in B_{ck}(\theta)} \sum_{(c',l) \neq (c^*, k^*)} \alpha_{c'l} I^{(c',l)}(\theta, \lambda) \geq 1, \quad \forall c \neq c^*, \forall k \neq k^*. \quad (11)$$

Constraint (9). We prove that (9) is equivalent to:

$$\min_{k \in M} \alpha_{c^*k} I(\theta_{c^*k}, \frac{\mu^*}{r_k}) \geq 1. \quad (12)$$

Observe that if $k < k_0$ (i.e., if $k \notin N$), then $B_{c^*k}(\theta) = \emptyset$. Let $k \in N$ with $k \neq k^*$. Without loss of generality assume that $k > k^*$. We prove that:

$$\inf_{\lambda \in B_{c^*k}(\theta)} \sum_{(c,l) \neq (c^*, k^*)} \alpha_{cl} I^{(cl)}(\theta, \lambda) = \sum_{l=k^*+1}^k \alpha_{c^*l} I(\theta_{c^*l}, \frac{\mu^*}{r_l}). \quad (13)$$

This is due to the fact that we can always choose $\lambda_{cl} = \theta_{cl}$ for all $c \neq c^*$, and to the following two observations:

- for all $\lambda \in B_{c^*k}(\theta)$, we have $\lambda_{c^*k^*} r_{k^*} = \theta_{c^*k^*} r_{k^*}$ and $\lambda_{c^*k} r_k > \lambda_{c^*k^*} r_{k^*}$, which using the unimodality of λ , implies that for any $l \in \{k^*, \dots, k\}$, $\lambda_{c^*l} r_l \geq \theta_{c^*k^*} r_{k^*}$. Hence:

$$\sum_{l \neq k^*} \alpha_{c^*l} I^{(c^*,l)}(\theta, \lambda) \geq \sum_{l=k^*+1}^k \alpha_{c^*l} I(\theta_l, \frac{\mu^*}{r_l}).$$

- For $\epsilon > 0$, define λ_ϵ as follows: for all $l \in \{k^*, \dots, k\}$, $\lambda_{c^*l} = (1 + (l - k^*)\epsilon) \frac{\mu^*}{r_l}$, and for all $l \notin \{k^*, \dots, k\}$, $\lambda_{c^*l} = \theta_{c^*l}$. By construction, $\lambda_\epsilon \in B_{c^*k}(\theta)$, and

$$\lim_{\epsilon \rightarrow 0} \sum_{l \neq k^*} \alpha_{c^*l} I^{(c^*,l)}(\theta, \lambda_\epsilon) = \sum_{l=k^*+1}^k \alpha_{c^*l} I(\theta_l, \frac{\mu^*}{r_l}).$$

From (13), we deduce that constraints (7) are equivalent to (12) (indeed, only the constraints related to $k \in M$ are really active, and for $k \in M$, (7) is equivalent to $\alpha_{c^*k} I(\theta_{c^*k}, \frac{\mu^*}{r_k}) \geq 1$).

Constraint (10). Note that if $k_c^* < k_0$, then $B_{c^*k_c^*}(\theta) = \emptyset$.

Assume that $k_c^* \geq k_0$. When $\lambda \in B_{ck_c^*}(\theta)$, the optimal (channel, rate) pair under λ is (c, k_c^*) . This implies that $r_{k_c^*} \lambda_{ck_c^*} \geq \mu^*$, and so:

$$\sum_{(c', l) \neq (c^*, k^*)} \alpha_{c'l} I^{(c', l)}(\theta, \lambda) \geq \alpha_{ck_c^*} I(\theta_{ck_c^*}, \frac{\mu^*}{r_{k_c^*}}).$$

Now select λ_ϵ as follows: $\lambda_{c'k} = \theta_{c'k}$ for all $(c', k) \neq (c, k_c^*)$, and $\lambda_{ck_c^*} = \mu^*/r_{k_c^*} + \epsilon$. Then $\lambda_\epsilon \in B_{ck_c^*}(\theta)$ for all $\epsilon > 0$, and

$$\lim_{\epsilon \rightarrow 0} \sum_{(c', l) \neq (c^*, k^*)} \alpha_{c'l} I^{(c', l)}(\theta, \lambda) = \alpha_{ck_c^*} I(\theta_{ck_c^*}, \frac{\mu^*}{r_{k_c^*}}).$$

We conclude that (10) is equivalent to:

$$\forall c \neq c^*, \alpha_{ck_c^*} I(\theta_{ck_c^*}, \frac{\mu^*}{r_{k_c^*}}) \geq 1_{k_c^* \geq k_0}.$$

Constraint (11). For $k < k_0$, $B_{ck}(\theta) = \emptyset$. Assume that $\tilde{k} \geq k_0$ and $\tilde{k} \neq k_c^*$. Then $\sum_{(c', l) \neq (c^*, k^*)} \alpha_{c'l} I^{(c', l)}(\theta, \lambda)$ is minimized over $B_{ck}(\theta)$ when for all $c' \neq c$ and all k , $\lambda_{c'k} = \theta_{c'k}$, and is actually equal to: $\inf_{\lambda_c \in C_k} \sum_l \alpha_{cl} I(\theta_{cl}, \lambda_{cl})$. Unfortunately, the above optimization problem cannot be further reduced. \square

PROOF OF THEOREM 3

For $\lambda: \forall c, \theta_c \in \mathcal{T} \cap \mathcal{U}$ and $\alpha \in \mathbb{R}_+^{C \times K}$, we define:

$$D(\theta, \lambda, \alpha) = \sum_{(c, k)} \alpha_{ck} I(\theta_{ck}, \lambda_{ck}).$$

As defined previously, the ‘‘bad parameter set’’ is:

$$B(\theta) = \{\lambda \in \mathcal{T} \cap \mathcal{U} : \lambda_{c^*k^*} = \theta_{c^*k^*}, \max_{(c, k)} r_k \lambda_{ck} > \mu^*\}.$$

Further define the set:

$$\mathcal{C} = \{\alpha \in \mathbb{R}_+^{C \times K} : \inf_{\lambda \in B(\theta)} D(\theta, \lambda, \alpha) \geq 1\}.$$

$c_U(\theta)$ in Theorem 2 is the solution to a minimization problem over \mathcal{C} . An upper bound of $c_U(\theta)$ is obtained by choosing $\alpha \in \mathcal{C}$ and by computing the value of the objective function at α which is $\sum_{c, k} \alpha_{ck} (\mu^* - \mu_{ck})$. We prove that if we define α as:

- $\alpha_{c^*k} = 1/I(\theta_{c^*k}, \frac{\mu^*}{r_k})$ if $k \in M$,
- $\alpha_{ck_c^*} = (\min\{I(\theta_{ck_c^*}, \frac{\mu^*}{r_{k_c^*}}), I(\theta_{ck_c^*}, \theta_{ck_c^*} - \frac{\delta_c}{r_{k_c^*}})\})^{-1}$ if $c \neq c^*$,
- $\alpha_{ck} = 1/I(\theta_{ck}, \theta_{ck} + \frac{\delta_c}{r_k})$ if $c \neq c^*$ and $k \in M_c$,
- $\alpha_{ck} = 0$ if $c \neq c^*$ and $k \notin M_c \cup \{k_c^*\}$.

then $\alpha \in \mathcal{C}$. To do so, we use the following decomposition: $B(\theta) = \cup_{(c, k) \neq (c^*, k^*)} B_{ck}(\theta)$ where

$$B_{ck}(\theta) = \{\lambda \in B(\theta) : (c, k) \in \arg \max_{(c', k')} r_{k'} \lambda_{c'k'}\}.$$

(i) If $\lambda \in \cup_{k \neq k^*} B_{c^*k}(\theta)$. Since $\lambda \in \cup_{k \neq k^*} B_{c^*k}(\theta)$, we

have $\theta_{c^*k^*} = \lambda_{c^*k^*}$. Since $k \mapsto r_k \lambda_{c^*k}$ is unimodal, and $k^* \notin \arg \max_k r_k \lambda_{c^*k}$, then there must exist a neighbour k' of k^* such that $r_{k'} \lambda_{c^*k'} \geq r_{k^*} \lambda_{c^*k^*} = r_{k^*} \theta_{c^*k^*} = \mu^*$. Hence $\lambda_{c^*k'} \geq \mu^*/r_{k'}$. Using the monotonicity of the KL divergence:

$$\begin{aligned} D(\theta, \lambda, \alpha) &\geq \alpha_{c^*k'} I(\theta_{c^*k'}, \lambda_{c^*k'}) \\ &\geq \alpha_{c^*k'} I(\theta_{c^*k'}, \mu^*/r_{k'}) \\ &\geq 1. \end{aligned}$$

(ii) If $\lambda \in \cup_k B_{ck}(\theta)$, $c \neq c^*$. Under parameter λ , let

$\tilde{k} = \arg \max_k r_k \lambda_{ck}$ be the optimal rate for channel c . We further consider two cases depending on whether \tilde{k} is equal to k_c^* .

Case (a): $\tilde{k} = k_c^*$. Then $\lambda \in B_{ck_c^*}(\theta)$, and we have $r_{k_c^*} \lambda_{ck_c^*} \geq \mu^*$. Hence:

$$\begin{aligned} D(\theta, \lambda, \alpha) &\geq \alpha_{ck_c^*} I(\theta_{ck_c^*}, \lambda_{ck_c^*}) \\ &\geq \alpha_{ck_c^*} I(\theta_{ck_c^*}, \mu^*/r_{k_c^*}) \\ &\geq 1. \end{aligned}$$

Case (b): $\tilde{k} \neq k_c^*$. Since $k \mapsto r_k \lambda_{ck}$ is unimodal and $k_c^* \neq \arg \max_k r_k \lambda_{ck}$, there must exist a neighbour k' of k_c^* such that $r_{k'} \lambda_{ck'} \geq r_{k_c^*} \lambda_{ck_c^*}$. Since $k \mapsto r_k \theta_{ck}$ is unimodal and $k_c^* = \arg \max_k r_k \theta_{ck}$, we have $r_{k_c^*} \theta_{ck_c^*} \geq r_{k'} \theta_{ck'}$. Therefore:

$$\begin{aligned} \max(r_{k_c^*} |\lambda_{ck_c^*} - \theta_{ck_c^*}|, r_{k'} |\lambda_{ck'} - \theta_{ck'}|) \\ \geq (r_{k_c^*} \theta_{ck_c^*} - r_{k'} \theta_{ck'})/2 \geq \delta_c. \end{aligned}$$

To establish the above inequality, we have used the fact that for all $a, b > 0$ and for all $x \in \mathbb{R}$, $\max(|x|, |x + a + b|) \geq (a + b)/2$, and have applied this result for $x = r_{k_c^*} (\lambda_{ck_c^*} - \theta_{ck_c^*})$, $a = r_{k'} \lambda_{ck'} - r_{k_c^*} \lambda_{ck_c^*}$, and $b = r_{k_c^*} \theta_{ck_c^*} - r_{k'} \theta_{ck'}$. We have shown that:

- either (b1) $r_{k_c^*} |\lambda_{ck_c^*} - \theta_{ck_c^*}| \geq \delta_c$;
- or (b2) $r_{k'} |\lambda_{ck'} - \theta_{ck'}| \geq \delta_c$.

If (b1) holds, then either $\lambda_{ck_c^*} \leq \theta_{ck_c^*} - \delta_c/r_{k_c^*}$ or $\lambda_{ck_c^*} \geq \theta_{ck_c^*} + \delta_c/r_{k_c^*}$. In the latter case, we have:

$$\begin{aligned} \lambda_{ck'} &\geq \frac{r_{k_c^*}}{r_{k'}} \lambda_{ck_c^*} \geq \frac{r_{k_c^*}}{r_{k'}} \theta_{ck_c^*} + \frac{\delta_c}{r_{k'}} \\ &\geq \theta_{ck'} + \frac{\delta_c}{r_{k'}}. \end{aligned}$$

If (b2) holds, then either $\lambda_{ck'} \geq \theta_{ck'} + \delta_c/r_{k'}$, or $\lambda_{ck'} \leq \theta_{ck'} - \delta_c/r_{k'}$. In the latter case, we have:

$$\begin{aligned} \lambda_{ck_c^*} &\leq \frac{r_{k'}}{r_{k_c^*}} \lambda_{ck'} \leq \frac{r_{k'}}{r_{k_c^*}} \theta_{ck'} - \frac{\delta_c}{r_{k_c^*}} \\ &\leq \theta_{ck_c^*} - \frac{\delta_c}{r_{k_c^*}}. \end{aligned}$$

In both cases (b1) and (b2), we have proved that either

$\lambda_{ck_c^*} \leq \theta_{ck_c^*} - \delta_c/r_{k_c^*}$ or $\lambda_{ck'} \geq \theta_{ck'} + \delta_c/r_{k'}$. Finally:

$$\begin{aligned} D(\theta, \lambda, \alpha) &\geq \alpha_{ck_c^*} I(\theta_{ck_c^*}, \lambda_{ck_c^*}) + \alpha_{ck'} I(\theta_{ck'}, \lambda_{ck'}) \\ &\geq \max\left\{\alpha_{ck_c^*} I(\theta_{ck_c^*}, \theta_{ck_c^*} - \frac{\delta_c}{r_{k_c^*}}), \right. \\ &\quad \left. \alpha_{ck'} I(\theta_{ck'}, \theta_{ck'} + \frac{\delta_c}{r_{k'}})\right\} \\ &\geq 1. \end{aligned}$$

We have proved that $\inf_{\lambda \in B(\theta)} D(\theta, \lambda, \alpha) \geq 1$, and thus $\alpha \in \mathcal{C}$. Now for our choice of α , the value of the objective function of the optimization problem in Theorem 2 is $c'_u(\theta)$. We conclude that $c_U(\theta) \leq c'_U(\theta)$. \square

PROOF OF THEOREM 6

We first decompose the regret as the sum of the regret due to exploring the neighbourhood of the optimal rate on each channel, and the other (channel,rate) pairs. By definition of the regret:

$$\begin{aligned} R^\pi(T) &= \sum_{(c,k)} (\mu^* - \mu_{ck}) \mathbb{E}[t_{ck}(T)] \\ &= \sum_c \sum_{|k-k_c^*| \leq 1} (\mu^* - \mu_{ck}) \mathbb{E}[t_{ck}(T)] \\ &\quad + \sum_c \sum_{|k-k_c^*| > 1} (\mu^* - \mu_{ck}) \mathbb{E}[t_{ck}(T)]. \end{aligned}$$

Therefore, to prove the theorem, it is sufficient to prove that for all channels c and all $\epsilon > 0$:

$$\mathbb{E}[t_{ck}(T)] \leq (1 + \epsilon) \log(T) \tau_c^{-1} + O(\log(\log(T)) + \Delta^{-2}),$$

if $|k - k_c^*| \leq 1$ and:

$$\mathbb{E}[t_{ck}(T)] \leq O(\log(\log(T)) + \Delta^{-2}).$$

otherwise. To ease notation we define $f(n) = \log(n) + 3 \log(\log(n))$, and the empirical success probability of (channel,rate) pair (c, k) as $\hat{\theta}_{ck}(n) = \hat{\mu}_{ck}(n)/\mu_{ck}$.

Define A the set of instants where there exists at least a (channel, rate) pair (c, k) such that either its index (i.e its upper confidence bound) $q_{ck}(n)$ under-estimates its expected value μ_{ck} , or its lower-confidence bound $\underline{q}_{ck}(n)$ over-estimates its expected value μ_{ck} :

$$\begin{aligned} A^1 &= \cup_{(c',k')} \{1 \leq n \leq T : q_{c'k'}(n) < \mu_{c'k'}\}, \\ A^2 &= \cup_{(c',k')} \{1 \leq n \leq T : \underline{q}_{c'k'}(n) > \mu_{c'k'}\}, \\ A &= A^1 \cup A^2. \end{aligned}$$

Consider (c, k) and $\epsilon > 0$ both fixed and define $\tilde{t}_{ck} = (1 + \epsilon) f(T) \tau_c^{-1}$ if $|k - k_c^*| \leq 1$ and $\tilde{t}_{ck} = 0$ otherwise.

Further define the sets of instants:

$$\begin{aligned} B_{ck} &= \{1 \leq n \leq T : (c, k)(n) = (c, k), t_{ck}(n) < \tilde{t}_{ck}\} \\ C_{ck} &= \{1 \leq n \leq T : n \notin A, (c, k)(n) = (c, k), \\ &\quad l_c(n) \neq k_c^*\} \\ D_{ck} &= \{1 \leq n \leq T : n \notin A, (c, k)(n) = (c, k), \\ &\quad l_c(n) = k_c^*, t_{ck}(n) \geq \tilde{t}_{ck}\}. \end{aligned}$$

and we have that:

$$\mathbb{E}[t_{ck}(T)] \leq \mathbb{E}[|A|] + \mathbb{E}[|B_{ck}|] + \mathbb{E}[|C_{ck}|] + \mathbb{E}[|D_{ck}|].$$

At all instants $n \in B_{ck}$, $t_{ck}(n)$ is incremented, therefore $|B_{ck}| \leq \tilde{t}_{ck}$. From the above inequality we deduce:

$$\mathbb{E}[t_{ck}(T)] \leq \mathbb{E}[|A|] + \tilde{t}_{ck} + \mathbb{E}[|C_{ck}|] + \mathbb{E}[|D_{ck}|].$$

To prove the announced result, it is sufficient to show that the expected size of A , B_{ck} and D_{ck} are at most of order $O(\Delta^{-2} + \log(\log(T)))$.

We will prove the following upper bounds:

- (i) $\mathbb{E}[|A|] = O(\log(\log(T)))$
- (ii) $\mathbb{E}[|C_{ck}|] = O(\Delta^{-2})$
- (iii) $\mathbb{E}[|D_{ck}|] = O(1)$.

(i) $\mathbb{E}[|A|] = O(\log(\log(T)))$

To upper bound the expected size of A , let us prove that:

$$A \subset \cup_{(c,k)} \{1 \leq n \leq T : t_{ck}(n) I(\hat{\theta}_{ck}(n), \theta_{ck}) \geq f(n)\}. \quad (14)$$

We recall the definition of index $q_{ck}(n)$:

$$q_{ck}(n) = \sup\{q \in [0, r_k], t_{ck}(n) I(\hat{\theta}_{ck}(n), q/r_k) \leq f(n)\}.$$

Since $q \mapsto I(p, q)$ is increasing for $q \geq p$, if $q_{ck}(n) < \mu_{ck}$ then:

$$t_{ck}(n) I(\hat{\theta}_{ck}(n), \mu_{ck}/r_k) = t_{ck}(n) I(\hat{\theta}_{ck}(n), \theta_{ck}) \geq f(n).$$

Using the same reasoning for the lower confidence bound $\underline{q}_{ck}(n)$:

$$\underline{q}_{ck}(n) = \inf\{q \in [0, r_k], t_{ck}(n) I(\hat{\theta}_{ck}(n), q/r_k) \leq f(n)\},$$

and since $q \mapsto I(p, q)$ is decreasing for $q \leq p$, if $\underline{q}_{ck}(n) > \mu_{ck}$ then $t_{ck}(n) I(\hat{\theta}_{ck}(n), \theta_{ck}) \geq f(n)$, so that equation (14) is valid.

Applying [21][Theorem 10] we have that:

$$\mathbb{P}[t_{ck}(n) I(\hat{\theta}_{ck}(n), \theta_{ck}) \geq f(n)] \leq e \frac{[\log(n) f(n)]}{n(\log(n))^3}.$$

Using a union bound we obtain the announced in-

equality:

$$\begin{aligned} \mathbb{E}[|A|] &\leq \sum_{n=1}^T \sum_{(c,k)} \mathbb{P}[t_{ck}(n)I(\hat{\theta}_{ck}(n), \theta_{ck}) \geq f(n)] \\ &\leq CK e \sum_{n=1}^T \frac{[\log(n)f(n)]}{n(\log(n))^3} \\ &\leq O(\log(\log(T))). \end{aligned}$$

$$\boxed{\text{(ii)} \mathbb{E}[|C_{ck}|] = O(\Delta^{-2})}$$

Let us decompose C_{ck} depending on the index of the leader:

$$\begin{aligned} C_{ck} &= \cup_{k': |k-k'| \leq 1, k' \neq k_c^*} C_{ckk'}, \\ C_{ckk'} &= \{1 \leq n \leq T : n \notin A, (c, k)(n) = (c, k), \\ &\quad l_c(n) = k'\}, \end{aligned}$$

the set of instants $n \notin A$ where (c, k) is selected and k' is the leader on channel c .

Fix $k' \neq k_c^*$ and consider $n \in C_{ckk'}$. There exists \tilde{k} such that $|\tilde{k} - k'| = 1$ and $\mu_{c\tilde{k}} \geq \mu_{ck'} + \Delta$ since $k' \neq k_c^*$ and $k \mapsto \mu_{ck}$ is unimodal for all c . Let us prove that we must have $U_c(n) = 0$. Assume that $U_c(n) = 1$ so that $\underline{q}_{ck'}(n) \geq \max_{k'': |k''-k'|=1} q_{ck''}(n) \geq q_{c\tilde{k}}(n)$. Since $n \notin A$ we have: $q_{c\tilde{k}}(n) \geq \mu_{\tilde{k}} > \mu_{k'} \geq \underline{q}_{ck'}(n)$, a contradiction.

Define $s = \sum_{n'=1}^n \mathbf{1}\{n' \in C_{ckk'}\}$ the number of instants in $C_{ckk'}$ between 1 and n . Since for all $n \in C_{ckk'}$ we have $U_c(n) = 0$, we must have $(c, k)(n) \in \arg \min_{k'': |k''-k'| \leq 1} t_{ck''}(n)$ so that $t_{ck''}(n) \geq s/3$ for all k'' such that $|k'' - k'| \leq 1$. Since $\mu_{\tilde{k}} \geq \mu_{k'} + \Delta$ and $\hat{\mu}_{k'}(n) \geq \hat{\mu}_{\tilde{k}}(n)$, we must have either $|\hat{\mu}_{k'}(n) - \mu_{k'}| \geq \Delta/2$ or $|\hat{\mu}_{\tilde{k}}(n) - \mu_{\tilde{k}}| \geq \Delta/2$.

So $t_{ck''}(n) \geq s/3$ for all k'' such that $|k' - k''| \leq 1$ and $\max_{k'': |k''-k'| \leq 1} |\hat{\mu}_{k''}(n) - \mu_{k''}(n)| \geq \Delta/2$, and applying Lemma 1, we obtain that $\mathbb{E}[|C|] = O(\Delta^{-2})$.

$$\boxed{\text{(iii)} \mathbb{E}[|D_{ck}|] = O(1)}$$

First it is noted that if $|k - k_c^*| > 1$, then $D_{ck} = \emptyset$, since $n \in D$ implies $l_c(n) = k_c^*$ and by design of CRS-T: $|k(n) - k_c^*| = |k(n) - l_{c(n)}(n)| \leq 1$.

Now consider k such that $|k - k_c^*| \leq 1$ and decompose D_{ck} depending on the value of the test $U_c(n) \in \{0, 1\}$:

$$\begin{aligned} D_{ck} &= D_{ck}^0 \cup D_{ck}^1, \\ D_{ck}^0 &= \{1 \leq n \leq T : n \in D_{ck}, U_c(n) = 0\}, \\ D_{ck}^1 &= \{1 \leq n \leq T : n \in D_{ck}, U_c(n) = 1\}. \end{aligned}$$

Consider $n \in D_{ck}^1$. Since $l_c(n) = k_c^*$ and $U_c(n) = 1$, by design of CRS-T we have $U_{c'}(n) = 1$ for all channels c' otherwise c is not selected. By the same reasoning as above, since $n \notin A$, $U_{c'}(n) = 1$ for all c' implies $l_{c'}(n) = k_c^*$ for all c' . Also, by design of CRS-T $k = k_c^*$

otherwise (c, k) is not selected. Since (c, k) is selected $q_{ck_c^*}(n) \geq q_{c^*l_{c^*}}(n) = q_{c^*k^*}(n) \geq \mu^*$ since $n \notin A$. By definition of the index q , $q_{ck_c^*}(n) \geq \mu^*$ implies:

$$t_k(n)I(\hat{\theta}_{ck}(n), \mu^*/r_k) \leq f(n) \leq f(T)$$

since $q \mapsto I(p, q)$ is increasing for $q \geq p$ and $n \mapsto f(n)$ is increasing. By definition of τ_c we have:

$$\begin{aligned} t_k(n) &\geq \tilde{t}_{ck} \\ &= (1 + \epsilon)f(T)\tau_c^{-1} \\ &\geq (1 + \epsilon)f(T)/I(\mu_{ck}, \mu^*/r_k), \end{aligned}$$

so that:

$$I(\hat{\theta}_{ck}(n), \mu^*/r_k) \leq I(\theta_{ck}, \mu^*/r_k)/(1 + \epsilon).$$

Therefore, since $p \mapsto I(p, q)$ is decreasing for $p \leq q$ there exists $\eta > 0$ (depending on ϵ) such that:

$$|\hat{\theta}_{ck}(n) - \theta_{ck}| \geq \eta.$$

Therefore:

$$D_{ck}^1 \subset \{(c, k)(n) = (c, k), |\hat{\theta}_{ck}(n) - \theta_{ck}| \geq \eta\}$$

so that $\mathbb{E}[D_{ck}^1] = O(\eta^{-2})$ applying Lemma 1 once again.

We turn to D_{ck}^0 . Since $U_c(n) = 0$, we have $\underline{q}_{ck_c^*}(n) \leq \max_{k': |k'-k_c^*|=1} q_{ck'}(n)$ so that either (a) $\underline{q}_{ck_c^*}(n) \leq \tilde{\mu}_c$ or (b) $\max_{k': |k'-k_c^*|=1} q_{ck'}(n) \geq \tilde{\mu}_c$.

Case (a): if $\underline{q}_{ck_c^*}(n) \leq \tilde{\mu}_c$ we have:

$$t_{ck_c^*}(n)I(\hat{\theta}_{ck_c^*}(n), \tilde{\mu}_c/r_{k_c^*}) \leq f(n) \leq f(T),$$

and since $t_{ck_c^*}(n) \geq (1 + \epsilon)f(T)/I(\theta_{ck_c^*}, \tilde{\mu}_c/r_{k_c^*})$:

$$I(\hat{\theta}_{ck_c^*}(n), \tilde{\mu}_c/r_{k_c^*}) \leq I(\theta_{ck_c^*}, \tilde{\mu}_c/r_{k_c^*})/(1 + \epsilon). \quad (15)$$

Case (b): consider k' such that $|k' - k_c^*| = 1$, if $q_{ck'}(n) \geq \tilde{\mu}_c$, then:

$$t_{ck'}(n)I(\hat{\theta}_{ck'}(n), \tilde{\mu}_c/r_{k'}) \leq f(n) \leq f(T),$$

and since $t_{ck'}(n) \geq (1 + \epsilon)f(T)/I(\theta_{ck'}, \tilde{\mu}_c/r_{k'})$:

$$I(\hat{\theta}_{ck'}(n), \tilde{\mu}_c/r_{k'}) \leq I(\theta_{ck'}(n), \tilde{\mu}_c/r_{k'})/(1 + \epsilon). \quad (16)$$

Putting (15) and (16) together, we deduce that there must exist η such that:

$$\sup_{k': |k'-k_c^*| \leq 1} |\hat{\theta}_{ck'}(n) - \theta_{ck'}| \geq \eta,$$

and therefore:

$$\begin{aligned} D_{ck}^0 \subset \{(c, k)(n) = \arg \min_{k': |k'-k_c^*| \leq 1} t_{ck'}(n), \\ \sup_{k': |k'-k_c^*| \leq 1} |\hat{\theta}_{ck'}(n) - \theta_{ck'}| \geq \eta\}, \end{aligned}$$

and using Lemma 1 a third time we have that $\mathbb{E}[|D_{ck}^0|] = O(\eta^{-2})$ which concludes the proof. \square

A DEVIATION RESULT

The following result proven in [28][Lemma 2.2] is reproduced here for completeness.

Lemma 1 ([28]): Let $\epsilon > 0$. Consider $(X(t))_{t \geq 0}$ i.i.d. random variables in $[0, 1]$ with common expectation μ . Define \mathcal{F}_n the σ -algebra generated by $(X(t))_{1 \leq t \leq n}$. Consider a random variable $B_t \in \{0, 1\}$ such that B_t is \mathcal{F}_{t-1} measurable for all $t \geq 0$, and define $t(n) = \sum_{t=1}^n B_t$ and $\hat{\mu}(n) = (1/t(n)) \sum_{t=1}^n B_t X_t$. Let $\Lambda \subset \mathbb{N}$ be a (random) set of instants. Assume that there exists a sequence of (random) sets $(\Lambda(s))_{s \geq 1}$ such that (i) $\Lambda \subset \cup_{s \geq 1} \Lambda(s)$, (ii) for all $s \geq 1$ and all $n \in \Lambda(s)$, $t(n) \geq \epsilon s$, (iii) $|\Lambda(s)| \leq 1$, and (iv) the event $n \in \Lambda(s)$ is \mathcal{F}_n -measurable. Then for all $\delta > 0$:

$$\mathbb{E}\left[\sum_{n \geq 1} \mathbf{1}\{n \in \Lambda, |\hat{\mu}(n) - \mu| > \delta\}\right] \leq \frac{1}{\epsilon \delta^2}.$$