Sleeping Experts in Wireless Networks*

Johannes Dams[1], Martin Hoefer[2], and Thomas Kesselheim[3]

[1] Dept. of Computer Science, RWTH Aachen University, Germany
dams@cs.rwth-aachen.de
[2] Max-Planck-Institut für Informatik and Saarland University, Germany
mhoefer@mpi-inf.mpg.de
[3] Dept. of Computer Science, Cornell University, United States
kesselheim@cs.cornell.edu

Abstract We consider capacity maximization algorithms for wireless networks with changing availabilities of spectrum. There are n sender-receiver pairs (called *links*) and k channels. We consider an iterative round-based scenario, where in each round the set of channels available to each link changes. Each link independently decides about access to one available channel in order to implement a successful transmission. Transmissions are subject to interference and noise, and we use a general approach based on affectance to define which attempts are successful. This includes recently popular interference models based on SINR.
Our main result is that efficient distributed algorithms from sleeping-expert regret learning can be used to obtain constant-factor approximations if channel availability is stochastic and independently distributed among links. In general, sublinear approximation factors cannot be obtained without the assumption of stochastic independence among links. A direct application of the no-external regret property is not sufficient to guarantee small approximation factors.

1 Introduction

One of the most important problems in the development of wireless networks is to overcome spectrum scarcity resulting from the static allocation schemes currently used by national regulators. This poses a variety of important regulatory and, in particular, algorithmic challenges. The idea is that licensed *primary users* open up their spectrum bands temporarily in local areas where it is unused. This creates spectrum opportunities for *secondary users* and results in much more efficient usage. A prominent approach that is currently discussed in industry is based on a database that records which channels are currently available for secondary usage in which areas. Primary users announce whether the channel is available to secondary users via this database organized by regulatory

* This work has been supported by DFG through Cluster of Excellence MMCI, UMIC Research Centre at RWTH Aachen University, and grant Ho 3831/3-1 and it has been supported by a fellowship within the Postdoc-Programme of the German Academic Exchange Service (DAAD).

authorities. In this case, secondary users obtain information about the channels available to them querying the database and then decide independently about channel access.

In this paper, we study an underlying algorithmic problem in this scenario and analyze the performance of distributed regret-based learning algorithms. In our model, there are k channels and n secondary users represented by *links*, i.e., by sender-receiver pairs located in a metric space. We consider a round-based approach, where in each round the set of channels available to each link can change, e.g., due to changing usage of the licensed primary users. Each link gets informed about the channels available to him and then decides about making a transmission attempt on an available channel. Transmissions are subject to interference and noise, and the success of a transmission attempt depends on conflicts defined using an interference model. Instead of relying on a particular model, we use a general approach to define conflicts based on a notion of affectance. This approach encompasses a variety of graph-based interference models, like disk graphs or the protocol model, as well as more realistic models based on the signal-to-interference-plus-noise-ratio (SINR).

We consider distributed learning algorithms that are executed for each link independently. The algorithms receive as input in each round the set of available channels and, in case they decide to transmit, a binary feedback if the transmission was successful or not. In particular, they do not need to know the exact SINR or whether and which other links made a (successful) transmission attempt. The decision for transmission follows an evaluation based on a natural utility function, which rewards previous successful transmissions and punishes failed attempts. Each no-regret algorithm aims at optimizing these utilities in a unilateral fashion, and therefore this scenario also has connections to game-theory.

While each link uses a no-regret algorithm to optimize his own successful transmissions, the obvious overall goal is *capacity maximization*, i.e., to maximize the total number of successful transmissions in the system. Our main result is that if all links use algorithms that satisfy a no-regret property resulting from a sleeping expert learning model [3], the number of successful transmissions converges to a constant-factor approximation for capacity maximization. For example, the surprisingly simple protocol of [12] can be used to obtain this result with high probability after a polynomial number of rounds. The analysis is based on a novel formulation of distributed capacity maximization using linear programming duality.

All our algorithms require channel availabilities to be stochastic and independently distributed for each link. This includes as a special case also the natural deterministic variant, where each link has a subset of available channels that does not change over time. We show that independence of the availability distributions among links is necessary, because the no-regret properties do not suffice to guarantee similar bounds when distributions are correlated. In addition, we show that a direct application of no-external regret as in previous work [6] does not provide similar results.

1.1 Contribution and Related Work

Capacity maximization, i.e., the task of maximizing the number of simultaneous transmissions, has been a prominent algorithmic problem over the last decade, e.g., in graph-based interference models [7, 16, 18]. With the seminal work of Moscibroda and Wattenhofer [15] attention has shifted to more realistic settings based on signal-to-interference-plus-noise-ratio (SINR).

We consider no-regret learning algorithms to solve capacity maximization with stochastic channel availabilities. As our main result, we show in Section 3 that no-ordering-regret algorithms converge to a constant-factor approximation for capacity maximization if availabilities are drawn independently at random for each link. Our analysis is based on a conflict graph representation of the interference model and, in particular, on a notion of a C-independence. C-independence turns out to be a key parameter for the performance of no-ordering-regret algorithms in this setting. If channel availabilities are stochastically independent for each link, the Sleeping-follow-the-perturbed-leader algorithm of [12] guarantees polynomial convergence time. This also holds when for each single link the availabilities of the different channels are arbitrarily correlated.

In contrast to this result for no-ordering-regret, we observe in Section 4 that for a direct application of the simpler no-external-regret condition, the successful transmissions can on average still be a factor of $\Omega(k)$ or $\Omega(n)$ smaller than in the optimum, where k is the number of channels and n the number of links. In addition, we highlight that without independence of channel availabilities among different links, there exist examples where even the no-ordering-regret property guarantees only a $\Omega(n)$-factor, as well.

Our main result is shown using a novel technique to analyze the performance of regret learning algorithms based on linear programming. This approach extends related works on capacity maximization on a single channel with uniform powers in SINR [1, 9] and Rayleigh-fading models [5], for which no-external regret algorithms are known to converge to constant-factor approximations [2,5,6]. Closest to our approach is our companion paper [4], in which we introduce a general framework based on the LP technique to study no-external regret learning with adversarial jamming on a single channel. The jammer yields more restrictive feedback, as availability of the channel remains unknown. Instead, an unavailable channel yields the same feedback to the link as being unsuccessful because of interference.

In this paper, we make a first step towards capacity maximization with multiple channels. For static availability, it is easy to see that previous results on no-external regret learning in [2, 5] directly extend to multiple channels. When we consider varying availabilities, however, multiple channels represent a significant complication. For a single channel and availabilities, our LP technique can be applied using no-external regret algorithms and the more challenging jammer feedback [4]. The main idea is to repeat a chosen action sufficiently long in order to obtain a "representative" feedback. For multiple channels, a similar approach is unlikely to work as we must learn on some channels while others are unavailable. This changes the regret and feedback conditions, and the connection between regret, feedback, and optimal solution value becomes more intricate to establish. In this paper, we resort to a stronger notion of no-ordering regret and use stochastic independence assumptions among links to obtain a constant-factor approximation. Omitting the independence assumptions constitutes a major open problem. Still, our work is a strong indication that efficient capacity maximization with availabilities and multiple channels is achievable in practice.

Action availabilities are the subject of a recent line of literature in online learning [12–14]. The actions of a game (or the experts in the learning setting) are not always available, and availability is based on adversarial decisions or random coin flips. The stochastic availabilities in our setting are similarly defined, and we use no-regret learning algorithms from the sleeping-experts setting to design a protocol for capacity maximization. Designing learning algorithms for sleeping expert settings started with Blum [3] and Freund et al. [8]. For definition of regret many works in this area do not use the best single strategy in hindsight. Instead, they resort to the best ordering of actions in hindsight [12–14], where unavailable actions can be accounted for. This has led to the design of multiple no-ordering-regret learning algorithms for the sleeping-experts setting. We will use ordering regret in our analysis as well.

2 Formal Description

2.1 General Problem Statement

We assume our network to consist of a set V of n wireless links $\ell_v = (s_v, r_v)$ for $v \in V$, each consisting of a sender and a respective receiver. We denote the set of channels by K and the number of channels by k. In each step, the availability of a channel $\kappa \in K$ to a link ℓ_v is the result of a random trial. We will assume throughout that the distributions for the random availabilities are independent among the links. However, among the channels of a link the availabilities can be arbitrarily correlated in our model. We justify this assumption by giving a lower bound, where we assume channel availabilities correlated among the links, in Section 4.

In a specific time step t and a specific link ℓ_v, some subset of channels is available. For any subset of channels M we define $p_{v,M}$ to be the probability that at least one channel out of M is available to ℓ_v. Then the random variable $P_{v,M}^{(t)}$ is defined to be 1 if and only if at least one channel out of the set M is available to link ℓ_v in time slot t. Let $p_{\min} = \min \{p_{v,M} \mid v \in V, M \neq \emptyset, p_{v,M} > 0\}$ be the minimal probability of a channel to be available. We can define $p_{\min}$ such that $p_{\min} > 0$ as channels which are never available to ℓ_v can be removed from consideration because they are neither available to the optimum nor the algorithm.

In each step t, each link first gets to know the outcome of the random trial and his set of available channels. It then has to decide whether and on which channel to send. Thus, a link can either not attempt transmission or transmit on one chosen channel. Success of transmissions can be defined in various ways, e.g., using the SINR model. In fact, our proofs rely on a more general condition which is also fulfilled by other interference models, e.g., based on bounded-independence graphs like unit-disk graphs [17].

In particular, we formally rely on conflict graphs to model interference (see, e.g., [11]). A *conflict graph* is a directed graph $G = (V, E)$ consisting of the links as vertices and weights $b_v(w)$ for any edge $(v, w) \in E$. We assume the weights to be defined such that a link ℓ_w can transmit successfully if and only if $\sum_{v \in L} b_v(w) \leq 1$, where L is the set of other links transmitting. We consider the conflict graph to be the same for all channels. Actually, this is only to simplify notation. It is easy to observe that our proofs also hold when the conflict graph is different in different channels. A subset of links is called *feasible* (on a channel) if all links in this set can transmit simultaneously (i.e., fullfil the condition above on that channel). The overall goal in this setting is to do *capacity maximization* in every single time step. That is to select for every time step depending on the availabilities a maximal cardinality subset of links and one available channel for those link such the sets of links are feasible on their respective channel.

We define the following notion of *C-independence* inspired by [2] as a key parameter to identify the connection between the specific interference model and the performance of our algorithm.

Definition 1 (cf. [2]). *A conflict graph is called* C-independent *if for any feasible set* L *there exists a subset* $L' \subset L$ *with* $|L'| = \Omega(|L|)$ *and* $\sum_{v \in L'} b_u(v) \leq C$ *for all* $u \in V$.

C-independence generalizes the bounded-independence property popular in the distributed computing literature. To embed the SINR model into this framework, let us outline how we can construct such a conflict graph. Let ϕ_v be the transmission power of link ℓ_v. Success of transmissions is defined in the SINR model as follows. Each sending link w emits a signal from sender s_w. This signal is received by receiver r_v at a strength of $\frac{\phi_w}{d_{w,v}^\alpha}$, where $d_{w,v}$ is the distance from sender s_w to receiver r_v and $\alpha > 0$ the path-loss exponent. The receiver r_v can successfully decode the signal transmitted by its sender s_v, if the SINR is above a certain threshold β. Using a constant $\nu \geq 0$ to denote ambient noise, the SINR condition formally reads

$$\frac{\frac{\phi_v}{d_{v,v}^\alpha}}{\sum_{w \neq v} \frac{\phi_w}{d_{w,v}^\alpha} + \nu} \geq \beta \ .$$

To turn this condition into appropriate edge weights of a conflict graph, we can use the notion of *affectance* as a measure of interference. It was defined for the SINR model in [10] as follows.

Definition 2. *The* affectance $a(w, v)$ *of link* ℓ_v *caused by another link* ℓ_w *is*

$$a(w, v) = \min \left\{ 1, \beta \frac{\frac{\phi_w}{d_{w,v}^\alpha}}{\frac{\phi_v}{d_{v,v}^\alpha} - \beta\nu} \right\} \ .$$

If all links use the same uniform power for transmission, this results in C-independence with a constant C. This was proven by Ásgeirsson and Mitra [2]. Using affectance it is straightforward to construct the corresponding conflict graph by simply setting weights $b_u(v) = a(u, v)$.

For simplicity we will assume that the conflict graphs satisfy C-independence for constant C throughout the paper. Nevertheless, losing a factor of C in the approximation guarantee our main theorem on the performance of regret learning can be directly generalized to arbitrary conflict graphs.

2.2 No-Regret Learning

We apply no-regret learning algorithms to solve capacity maximization. The links independently decide in every time slot whether and on which channel to transmit using appropriate learning algorithms. Every algorithm adjusts its decisions based on the outcome of its previous decisions. To measure the quality of an outcome every link i uses an utility function $u_i(a_i, a_{-i})$ depending on action a_i chosen by player i and a_{-i}, the vector of actions of all other players. Throughout this paper we define the utility of a link i as follows. This utility function was already used for a single channel case where the channel is always available by Andrews and Dinitz [1] and later by Ásgeirsson and Mitra [2].

$$u_i(a_i, a_{-i}) = \begin{cases} 1 & \text{if } i \text{ transmits successfully,} \\ -1 & \text{if } i \text{ attempts and the transmission fails,} \\ 0 & \text{otherwise.} \end{cases}$$

This utility reflects that the best a link can achieve in one time slot is successful transmission, for which is rewarded with a utility of 1. The worst that can happen is an unsuccessful attempt, which is penalized by a utility of -1. This strikes a balance between reducing interference on other links (when not being successful) and increasing the number of transmissions (when being successful).

For our utility functions we can consider different notions of regret. The easiest notion is external regret given as follows.

Definition 3. *Let $a^{(1)}, \ldots, a^{(T)}$ be a sequence of action vectors. The* external regret *of this sequence for link i is defined by*

$$\max_{a'_i \in \mathcal{A}} \sum_{t=1}^{T} u_i(a'_i, a_{-i}^{(t)}) - \sum_{t=1}^{T} u_i(a_i^{(t)}, a_{-i}^{(t)}) \ ,$$

where $\mathcal{A}$ denotes the set of actions.

This notion of regret can only be used if the utilities are defined also for actions that are not available in a time slot. We assume that choosing an unavailable channel is equivalent to choosing not to send at all, which is an action that we assume to be always available. This allows to directly apply no-external-regret algorithms in our scenario.

In addition, let us also consider a different notion of regret from sleeping experts learning. This notion of regret is introduced by Kanade et al. [12].

Definition 4 (Kanade et al. [12]). *Let $a^{(1)}, \dots, a^{(T)}$ be a sequence of action vectors. The* ordering regret *of this sequence for link i in the sleeping experts setting is defined as*

$$\max_{\sigma \in S_{\mathcal{A}}} \mathbb{E}\left[\sum_{t=1}^{T} u_i(\sigma(\mathcal{A}^{(t)}), a_{-i}^{(t)})\right] - \sum_{t=1}^{T} u_i(a_i^{(t)}, a_{-i}^{(t)}) \enspace ,$$

where the expectation is over the random availabilities. Here, $\mathcal{A}$ denotes the set of actions, $S_{\mathcal{A}}$ the set of all permutations on $\mathcal{A}$, and $\sigma(\mathcal{A}^{(t)})$ the action ordered topmost in σ of the actions available in time slot t.

In contrast to external regret, ordering regret does not measure the utility difference to the best action in hindsight but to the utility resulting from the best ordering in hindsight. The utility for an ordering is computed by assuming that in every step the topmost available action in the ordering is played. Additionally, the expectation over the availabilities is considered for comparison. Note that we do not consider the expectation in this definition to be taken over the random choices of the algorithm as, e.g., in [14]. Considering the expectation this way is possible due to the stochastic independence assumption. Thus, we can keep the choices of other players and also their availabilities fixed and just take the expectation over the availabilities of one player i.

An infinite sequence of actions or an algorithm has the *no-external regret property* if external regret grows in $o(T)$. We analogously define the *no-ordering-regret property*. Throughout this paper we will, whenever it is clear from context, use regret as a synonym for either ordering regret or external regret.

3 Convergence with No-Ordering-Regret Learning

In this section we show our main result that using no-ordering-regret algorithms the number of successful transmissions converges to a constant-factor approximation of the optimal capacity. As discussed before, the optimum is different in different time slots depending on available channels. Let us denote by $OPT_\kappa^{(t)}$ the set of links transmitting on channel κ in time slot t in the optimal solution. Thus, we compare the number of successful transmissions of the no-regret algorithms to the empirical average capacity of all optima, i.e., $\overline{|OPT|} = \frac{1}{T}\sum_t \sum_{\kappa \in K} |OPT_\kappa^{(t)}|$.

For simplicity we assume that conflict graphs are $O(1)$-independent and highlight the places at which the factor C comes into play if this assumption does not hold. Note that, in particular, conflict graphs resulting from the SINR model under uniform power yield constant C [2].

Theorem 1. *If all links use no-ordering-regret algorithms, the average number of successful transmissions becomes a constant-factor approximation for capacity maximization after a number of time steps polynomial in n and linear in k with high probability, i.e., with probability at least $1 - \frac{1}{n^c}$ for any constant c. More generally, in C-independent conflict graphs the same result holds for convergence to an $O(C)$-approximation.*

While the overall approach is in the spirit of previous work, our setting is quite different and the notion of regret also differs. Similar to [2,6], our analysis starts with the observation that a constant fraction of all transmission attempts are successful. Afterwards, we combine this with the result that the number of transmission attempts is in $\Omega\left(\overline{|OPT|}\right)$. Especially the proof of this latter statement in Lemma 2 below needs more advanced techniques. Together both statements prove our theorem.

In the remainder of this section, we denote the fraction of time slots in which link ℓ_v transmits on channel κ by $q_{v,\kappa}$. The sum over the channels is denoted by $q_v = \sum_\kappa q_{v,\kappa}$. The fraction of time slots in which link ℓ_v transmits successfully is denoted by $w_{v,\kappa}$, and the sum over channels by w_v. In the following, we will denote the fraction of all time steps in which link ℓ_v (no matter whether it attempted to transmit) would not be able to be successful on channel κ by $f_{v,\kappa}$ no matter if κ was actually available to link ℓ_v. Throughout this section, we assume that for each link the ordering regret after T time slots is at most $\varepsilon \cdot T$.

First of all, let us bound the number of successful transmissions by the number of transmission attemps.

Lemma 1 (cf. [2,5]). *It holds* $w_v \leq q_v \leq 2 \cdot w_v + \varepsilon$ *and* $\sum_v w_v \leq \sum_v q_v \leq 2 \cdot \sum_v w_v + \varepsilon n$.

Proof. The first inequality follows by definition. For the second inequality, we use the fact that for each link v the average regret is at most ε. Therefore, not to send at all can increase the average utility per step by at most ε. Formally this means $(q_v - w_v) - w_v = q_v - 2w_v \leq \varepsilon$. Taking the sum over all v we get $\sum_v q_v - 2 \cdot \sum_v w_v \leq \varepsilon n$. This yields the claim. □

Lemma 1 shows that the number of successful transmissions and the number of transmission attempts only differ by a constant factor. Together with the following lemma this proves Theorem 1.

Lemma 2. *Every sequence of length* $T \in \Omega\left(\frac{1}{p_{\min}}(\ln n + k)\right)$ *with ordering regret at most* $\varepsilon \cdot T < \frac{1}{4n} \cdot T$ *yields* $\sum_v q_v = \Omega\left(\overline{|OPT|}\right)$ *with high probability.*

To prove Lemma 2, we use a primal-dual approach using an appropriately defined linear program. Recall that $p_{\min} = \min\{p_{v,M} \mid v \in V, M \neq \emptyset, p_{v,M} > 0\}$ is the minimal availability probability of all the channels.

Let us start by showing in Lemma 3 that for a number of time slots $T \in \Omega\left(\frac{1}{p_{\min}}(\ln n + k)\right)$ with high probability the empirical fraction of slots $\bar{P}_{v,M}$ in which at least one channel out of M was available to link ℓ_v is close to the probability $p_{v,M}$, for every set of channels M and every link ℓ_v. Afterwards, we will use this result to draw a connection between transmission attempts, availabilities, and experienced affectances in Lemma 4 finally proving Lemma 2.

Lemma 3. *After a number of time steps* $T \in \Omega\left(\frac{1}{p_{\min}} \cdot (\ln n + k)\right)$ *it holds* $|\bar{P}_{v,M} - p_{v,M}| \leq \frac{1}{2}\bar{P}_{v,M}$ *for all sets of channels* M *and all links* ℓ_v *with high probability.*

Proof. Consider the random variable $P^{(t)}_{v,M} \in \{0,1\}$ indicating whether any channel of the set M is available for link ℓ_v in time slot t. Let $Y = \sum_t P^{(t)}_{v,M}$. Thus, we need $|\mathbb{E}(Y) - Y| < \frac{1}{2}Y$ to hold, because this directly yields $|\bar{P}_{v,M} - p_{v,M}| \leq \frac{1}{2}\bar{P}_{v,M}$ by division with T. Equivalently we need $\frac{1}{2}Y < \mathbb{E}(Y) < \frac{3}{2}Y$ to hold.

As the channel availabilities are drawn independently in every time slot, we can apply a Chernoff bound. This yields $\mathbf{Pr}\left[Y \geq (1+\delta)\mathbb{E}(Y)\right] \leq \exp\left(\frac{-\delta^2}{3}\mathbb{E}(Y)\right)$ and $\mathbf{Pr}\left[Y \leq (1-\delta)\mathbb{E}(Y)\right] \leq \exp\left(\frac{-\delta^2}{2}\mathbb{E}(Y)\right)$ for every $\delta \in [0,1]$.

Using this we get $\mathbf{Pr}\left[Y \geq 2\mathbb{E}(Y)\right] \leq \exp\left(-\frac{1}{3}\mathbb{E}(Y)\right)$ and $\mathbf{Pr}\left[Y \leq \frac{2}{3}\mathbb{E}(Y)\right] \leq \exp\left(-\frac{1}{18}\mathbb{E}(Y)\right)$. With a union bound, the probability that $|\bar{P}_{v,M} - p_{v,M}| \leq \frac{1}{2}\bar{P}_{v,M}$ does not hold for a particular set M is

$$\mathbf{Pr}\left[|\bar{P}_{v,M} - p_{v,M}| > \frac{1}{2}\bar{P}_{v,M}\right] \leq \exp\left(-\frac{p_{v,M}T}{3}\right) + \exp\left(-\frac{p_{v,M}T}{18}\right) \quad .$$

This is at most $2 \cdot \exp\left(-\frac{1}{18}p_{v,M}T\right)$. Applying another union bound yields

$$\sum_{v \in V} \sum_{M \subseteq K} \mathbf{Pr}\left[|\bar{P}_{v,M} - p_{v,M}| > \frac{1}{2}\bar{P}_{v,M}\right] \leq 2^k n \cdot 2 \cdot \exp\left(-\frac{1}{18}p_{\min} \cdot T\right) \quad .$$

Setting $T \geq \frac{18}{p_{\min}}\left((c+1)\ln n + (k+1)\cdot \ln 2\right)$ shows that the probability that for any arbitrary set of channels M the property $|\bar{P}_{v,M} - p_{v,M}| \leq \frac{1}{2}\bar{P}_{v,M}$ does not hold is at most n^{-c}. $\square$

Consider the set of channels with a low congestion where a link will be unsuccessful in a small fraction of time slots. For these channels we will show that the number of transmission attempts yields an upper bound on the availabilities. This fact will be used in the proof of Lemma 2.

Lemma 4. *Let M be any set of channels such that for every channel $\kappa \in M$ it holds $f_{v,\kappa} \leq \frac{1}{4}$. If regret is at most ε and $|\bar{P}_{v,M} - p_{v,M}| \leq \frac{1}{2}\bar{P}_{v,M}$, then it follows*

$$4 \sum_{\kappa \in K} q_{v,\kappa} + 4\varepsilon \geq \bar{P}_{v,M} \quad .$$

Proof. The expected utility of the best ordering in hindsight is obviously at least as high as the expected utility of the ordering in which all $\kappa \in M$ are ordered above the action 'not sending' followed by all other channels.

First, we consider just one channel. On any channel $\kappa \in M$ link ℓ_v is not successful in an $f_{v,\kappa}$-fraction of all time steps. That leaves $T \cdot (1 - f_{v,\kappa})$ time steps possibly successful each yielding a utility of $+1$ for choosing κ if it was available. Choosing κ in contrast also yields -1 as a utility in $T \cdot f_{v,\kappa}$ time steps if κ was available. Thus ordering only κ before not sending yields a total expected utility of $p_{v,\kappa}\left(T \cdot (1 - f_{v,\kappa}) - T \cdot f_{v,\kappa}\right)$. We extend this argument to the set M such that it depends on $p_{v,M}$ instead of $p_{v,\kappa}$ in the following way.

For any ordering with all $\kappa \in M$ ordered above not sending (and all other channels below), we get in expectation at least the utility of the worst channel

$\kappa \in M$ if any channel of M is available. This only holds due to the independence of the availabilities between different links as we can fix the actions of other links. This way, considering the expectation over ℓ_v's own availabilities we yield at least $\min_\kappa ((1 - f_{v,\kappa}) - f_{v,\kappa})$ for time steps where any channel in M is available. For the expected utility of the best ordering in hindsight this yields

$$\max_{\sigma \in S_{\mathcal{A}}} \mathbb{E}\left[\sum_{t=1}^{T} u_v(\sigma(\mathcal{A}^{(t)}), a_{-v}^{(t)})\right] \geq \min_\kappa ((1 - f_{v,\kappa}) - f_{v,\kappa}) \, p_{v,M} \cdot T \enspace .$$

Note that as discussed above we can only bound the expected utility by that of one channel due to the availabilities of channels between links being stochastically independent. Otherwise those could be correlated in such a way that the expected unsuccessful time steps are not at most $T \cdot \max_\kappa f_{v,\kappa} p_{v,M}$ but could be worse. This is due to correlation, for example, being able to force all interference of other links on a channel (even if it occurs in few time steps in total) occur in available time steps only.

This yields $\frac{1}{T} \max_{\sigma \in M_{\mathcal{A}}} \mathbb{E}\left[\sum_{t=1}^{T} u_v(\sigma(\mathcal{A}^{(t)}), a_{-v})\right] \geq \left(\frac{3}{4} - \frac{1}{4}\right) p_{v,M} = \frac{1}{2} p_{v,M}$. Using $|\bar{P}_{v,M} - p_{v,M}| \leq \frac{1}{2}\bar{P}_{v,M}$ we can easily bound this from below by $\frac{1}{4}\bar{P}_{v,M}$. With the fact that the regret is at most ε and that the utility is at most q_v we get $\frac{1}{4}\bar{P}_{v,M} \leq q_v + \varepsilon$. □

This connection between the availability of a set, its interference, and the actions played now allows us to prove Lemma 2.

Proof (Proof of Lemma 2). Recall the definition of C-independence. Note that the conditions given in Definition 1 can be transfered for each channel κ from the single time steps to all time steps by averaging as follows. Let $OPT_\kappa'^{(t)}$ be L' out of Definition 1 when setting $L = OPT_\kappa^{(t)}$ yielding $|OPT_\kappa'^{(t)}| \geq \Omega\left(|OPT_\kappa^{(t)}|\right)$ and $\sum_{v \in OPT_\kappa'^{(t)}} b_u(v) \leq C$ for every time step t. By averaging over all time steps this is

$$\frac{1}{T}\sum_t |OPT_\kappa'^{(t)}| \geq \Omega\left(\frac{1}{T}\sum_t |OPT_\kappa^{(t)}|\right) \quad \text{and} \quad \frac{1}{T}\sum_t \sum_{v \in OPT_\kappa'^{(t)}} b_u(v) \leq C \enspace .$$

As C-independence holds in the given network for any feasible set on each channel, it also holds in this averaged variant for the optimum on each channel.

We will prove our lemma with the following primal-dual approach. For the following primal LP we will essentially consider the optimum averaged over all time steps and utilize C-independence. The above result is this way useful to show feasibility of the primal solution.

$$
\begin{aligned}
\text{Max.} & \sum_{v \in V} \sum_{\kappa \in K} x_{v,\kappa} \\
\text{s.t.} & \sum_{v \in V} b_u(v) x_{v,\kappa} \le C \qquad \forall u \in V, \kappa \in \mathcal{K} \\
& \sum_{\kappa \in M} x_{v,\kappa} \le \bar{P}_{v,M} \ \forall v \in V, M \subseteq K \\
& x_v \ge 0 \qquad \forall v \in V
\end{aligned}
$$

Observe that $x_{v,\kappa} = \frac{\left|\left\{t \middle| v \in OPT'^{(t)}_{\kappa}\right\}\right|}{T}$ represents a feasible solution to this LP. The first constraint is fulfilled as C-independence is fulfilled for every single time slot. The second constraint is fulfilled due to the fact that at most one channel is used at a time. Thus, we get $\sum_{v \in V} \sum_{\kappa \in K} x_{v,k} \ge \Omega\left(\overline{|OPT|}\right)$ by the definition of C-independence for the single-slot optima.

Constructing the dual to this primal LP yields

$$
\begin{aligned}
\text{Min.} & \sum_{v \in V} \sum_{\kappa \in K} C \cdot y_{v,k} + \sum_{v \in V} \sum_{M \subseteq K} \bar{P}_{v,M} \cdot z_{v,M} \\
\text{s.t.} & \sum_{u \in V} b_u(v) y_{u,\kappa} + \sum_{M : \kappa \in M} z_{v,M} \ge 1 \ \forall v \in V, \kappa \in K \\
& y_{v,\kappa}, z_{v,M} \ge 0 \ \forall v \in V, \kappa \in K, M \subseteq K
\end{aligned}
$$

We construct the following dual solution that gives an upper bound to the solution of the primal LP. Let $M_v = \left\{\kappa \in K \;\middle|\; f_{v,\kappa} \le \frac{1}{4}\right\}$, where $f_{v,\kappa}$ again denotes the fraction of all time steps in which link ℓ_v would not be able to transmit successfully on channel κ. So M_v represents the set of channels with low congestion. We set $y_{v,\kappa} = 4 \cdot q_{v,\kappa}$, $z_{v,M_v} = 1$, and $z_{v,S} = 0$ for all $S \neq M_v$.

First, let us observe that this is a feasible solution and the constraints are fulfilled. Recall the definition of $f_{v,\kappa}$ being the fraction of time steps ℓ_v would have been unsuccessful on channel κ no matter whether the channel was available to ℓ_v. Thus, for any channel κ in which $f_{v,\kappa} \ge \frac{1}{4}$, it holds $\sum_{u \in V} b_u(v) q_{u,\kappa} \ge \frac{1}{4}$. So $\sum_u b_u(v) \cdot y_{u,\kappa} \ge 1$ with the chosen $y_{u,\kappa}$. For the other case with $f_{v,\kappa} < \frac{1}{4}$ we set $z_{v,M_v} = 1$ and by definition $\kappa \in M_v$. Therefore, the constraint is fulfilled.

Using Lemma 4 leads to an upper bound on the objective function of the dual LP of

$$
\sum_v \left(4Cq_v + \bar{P}_{v,M_v}\right) \le \sum_v \left(4Cq_v + 4q_v\right) + 4\varepsilon n \ .
$$

Combined with the primal LP this yields $\sum_v q_v = \Omega\left(\overline{|OPT|}\right)$ for $\varepsilon < \frac{1}{4n}$. In particular, for arbitrary C-independence the last derivation obviously implies $C \cdot \sum_v q_v = \Omega\left(\overline{|OPT|}\right)$ and directly yields an approximation factor in $O(C)$. □

We have seen that after a number of time slots linear in k and logarithmic in n a sequence with low regret converges to a constant-factor approximation with high probability. Additionally, this time bound depends on the minimal probability of the availabilities $p_{\min}$. It is clear that a similar parameter must occur in the convergence time as links may not learn in time slots in which they have no channel available at all.

Theorem 1 and Lemma 2 show that the no-ordering-regret property allows to converge to constant-factor approximations. The algorithms in [12] have this property, which allows to directly use them for capacity maximization in our scenario. The sleeping-follow-the-perturbed-leader algorithm of [12] yields an ordering regret of at most $\sqrt{T \log k}$ in expectation after T time slots. While this algorithm runs in the full-information model getting feedback also for actions not chosen, Kanade et al. also propose an algorithm yielding low regret roughly the size $(k \cdot T \cdot \log T)^{4/5}$ in the partial-information model, where only feedback for chosen actions is given. To reach $\varepsilon < \frac{1}{4n}$ we therefore need only an additional factor polynomial in n for the number of time slots.

4 Lower Bounds

In this section, we show that a direct application of no-external-regret algorithms does not necessarily yield a constant-factor approximation. In fact, we will give an example that shows approximation factors in $\Omega(k)$ and $\Omega(n)$. Note that these factors can already be reached by algorithms where just one channel is utilized or just one link transmits, respectively. Additionally, we show that our assumption of stochastic independence in the availabilities among links is necessary. All our lower bound constructions can trivially be embedded into 1-independent conflict graphs. Thus, they establish linear lower bounds even in cases, where no-ordering-regret obtains constant-factor approximations.

Theorem 2. *For every number of channels k there is an instance such that for a sequence yielding 0 external regret the number of successful transmissions is at least a factor of k smaller than in an optimal schedule.*

Proof. Let us assume that all n links can be successful simultaneously on every channel. This allows us to consider only a single link. We first consider a sequence of deterministic availabilities in which channel κ is available in time slots t with $(t \mod k+1) = \kappa$. Here there is a 0-external-regret sequence in which exactly one channel is chosen. The link will transmit only in every k-th time step, choosing exactly one channel. In contrast, in the optimum the link can simply choose another channel in every single time step. This yields the factor of k.

To reproduce the same arguments with stochastic availability, we set the probabilities for each channel availability to $\frac{1}{k}$. This yields the same structure. Again, if the link chooses only a single channel for transmission, it will encounter vanishing external regret as in the long run all channels have the same availability and success. However, it will only transmit in an $\frac{1}{k}$-fraction of all time slots. In contrast, in expectation in every time slot there is at least one channel available. This implies that in the long run a factor of k. □

This result also implies an $\Omega(n)$ bound by setting $k = n$. Therefore, using no-external regret in this direct way does not imply a constant-factor approximation.

Corollary 1. *For every number of links n there is a network such that for a sequence yielding* 0 *external regret the number of successful transmissions is at least a factor of n smaller than in an optimal schedule.*

In contrast to directly applying the no-external-regret property, one might consider using multiple such no-external-regret algorithms. It is an interesting open problem if this allows to establish similar properties as for the sleeping-follow-the-perturbed-leader algorithm leading to a constant-factor approximation.

In the previous sections, we have assumed that the channel availabilities of different links are independent. We will use a similar example as in the proof above to see that this assumption is necessary to achieve convergence to a constant-factor approximation, even for no-ordering-regret algorithms.

Theorem 3. *For every number of links n there exists a network with correlated availabilities such that for a sequence yielding* 0 *ordering regret the number of successful transmissions is at least a factor of n smaller than in an optimal schedule.*

Proof. Suppose there is only one channel. We construct the network as follows. No pair of links can transmit simultaneously on the channel. This can easily be achieved be placing links (almost) in the same location and constructing the interference appropriately.

The channel is either available for all n links simultaneously or for only one single link ℓ_v with $v \in \{2, \ldots, n\}$. The probability for each of these n cases is $\frac{1}{n}$.

We construct a 0-ordering-regret sequence by scheduling link ℓ_1 to send whenever the channel is available to him. All other links choose not to send at all. The dependence of the availabilities implies that the expected utility of the best response in hindsight for all links $\ell_2, \ldots, \ell_n$ becomes 0 because, in the long run, for each of these links every second available slot is occupied by ℓ_1.

In contrast, in the optimum letting every link $\ell_2, \ldots, \ell_n$ transmit when the channel is available to him alone yields a successful transmission in every time slot. This proves the theorem. □

The one-or-all structure of availabilities used in the proof of Theorem 3 can still occur with a very low probability if we do not assume correlation and instead let the channel be available to each link independently with probability $\frac{1}{n}$. In this case, however, the transmission choices in the proof of Theorem 3 do not yield 0 ordering regret.

With a slight adjustment of the one-or-all structure, it is possible to show even slightly stronger lower bounds close to $3n/2$. We proved our positive results under the assumption that availabilities of links are independent and encounter no correlation at all. In contrast, the lower bound in Theorem 3 heavily relies on correlation. It is an interesting open problem to characterize influence of correlation of availability distributions on the performance of no-regret learning algorithms (e.g., when correlation results from a locality structure of primary and secondary users).

References

1. M. Andrews and M. Dinitz. Maximizing capacity in arbitrary wireless networks in the SINR model: Complexity and game theory. In *Proc. 28th IEEE Conf. Computer Communications (INFOCOM)*, pages 1332–1340, 2009.
2. E. I. Asgeirsson and P. Mitra. On a game theoretic approach to capacity maximization in wireless networks. In *Proc. 30th IEEE Conf. Computer Communications (INFOCOM)*, pages 3029–3037, 2011.
3. A. Blum. Empirical support for winnow and weighted-majority algorithms: Results on a calendar scheduling domain. *Machine Learning*, 26(1):5–23, 1997.
4. J. Dams, M. Hoefer, and T. Kesselheim. Jamming-resistant learning in wireless networks. *CoRR*, abs/1307.5290, 2013.
5. J. Dams, M. Hoefer, and T. Kesselheim. Scheduling in wireless networks with Rayleigh-fading interference. In *Proc. 24th Symp. Parallelism in Algorithms and Architectures (SPAA)*, pages 327–335, 2012.
6. M. Dinitz. Distributed algorithms for approximating wireless network capacity. In *Proc. 29th IEEE Conf. Computer Communications (INFOCOM)*, pages 1397–1405, 2010.
7. T. Erlebach, K. Jansen, and E. Seidel. Polynomial-time approximation schemes for geometric graphs. *SIAM J. Comput.*, 34(6):1302–1323, 2005.
8. Y. Freund, R. Schapire, Y. Singer, and M. Warmuth. Using and combining predictors that specialize. In *Proc. 29th Symp. Theory of Computing (STOC)*, pages 334–343, 1997.
9. O. Goussevskaia, M. Halldórsson, R. Wattenhofer, and E. Welzl. Capacity of arbitrary wireless networks. In *Proc. 28th IEEE Conf. Computer Communications (INFOCOM)*, pages 1872–1880, 2009.
10. M. Halldórsson and R. Wattenhofer. Wireless communication is in APX. In *Proc. 36th Intl. Coll. Automata, Languages and Programming (ICALP)*, volume 1, pages 525–536, 2009.
11. M. Hoefer, T. Kesselheim, and B. Vöcking. Approximation algorithms for secondary spectrum auctions. In *Proc. 23rd Symp. Parallelism in Algorithms and Architectures (SPAA)*, pages 177–186, 2011.
12. V. Kanade, B. McMahan, and B. Bryan. Sleeping experts and bandits with stochastic action availability and adversarial rewards. *J. Machine Learning Res. - Proc. Track*, 5:272–279, 2009.
13. V. Kanade and T. Steinke. Learning hurdles for sleeping experts. In *Proc. 3rd Symp. Innovations in Theoret. Computer Science (ITCS)*, pages 11–18, 2012.
14. R. Kleinberg, A. Niculescu-Mizil, and Y. Sharma. Regret bounds for sleeping experts and bandits. *Machine Learning*, 80(2–3):245–272, 2010.
15. T. Moscibroda and R. Wattenhofer. The complexity of connectivity in wireless networks. In *Proc. 25th IEEE Conf. Computer Communications (INFOCOM)*, pages 1–13, 2006.
16. T. Nieberg, J. Hurink, and W. Kern. Approximation schemes for wireless networks. *ACM Trans. Algorithms*, 4(4), 2008.
17. J. Schneider and R. Wattenhofer. Coloring unstructured wireless multi-hop networks. In *Proc. 28th Symp. Principles of Distrib. Comput. (PODC)*, pages 210–219, 2009.
18. J. Schneider and R. Wattenhofer. An optimal maximal independent set algorithm for bounded-independence graphs. *Distributed Computing*, 22(5–6):349–361, 2010.