

No-Regret Learning in Wireless Networks

Von der Fakultät für Mathematik, Informatik und Naturwissenschaften
der RWTH Aachen University zur Erlangung des akademischen Grades
eines Doktors der Naturwissenschaften genehmigte Dissertation

vorgelegt von

Diplom-Informatiker

Johannes Dams

aus Geldern

Berichter: *Privatdozent Dr. Martin Hoefer*
Universitätsprofessor Dr. Martin Grohe

Tag der mündlichen Prüfung: *16. Oktober 2014*

Diese Dissertation ist auf den Internetseiten der Hochschulbibliothek online verfügbar.

Abstract

The defining property of wireless networks is that communications are not done over a wired connection. In contrast, radio spectrum is used which is subject to a variety of influences. Especially, interference from other concurrent transmissions can cause communication attempts to be unsuccessful. Wireless devices have to operate and perform transmissions despite this interference.

Algorithmic considerations often assume a centralized coordination of devices. This coordination behavior can typically not be assumed, as wireless networks are inherently distributed. Even if a central authority existed, the necessary coordination would introduce additional traffic and interference. In this thesis, we consider distributed adaptive algorithms that only rely on very little information to maximize the number of successful transmissions. We give provable guarantees for the performance of these adaptive protocols.

We model a wireless network to consist of sender-receiver pairs denoted as links whose transmissions interfere based on a conflict graph. This model generalizes previous interference models like unit-disk graphs or ones based on the signal-to-interference-plus-noise ratio (SINR).

We consider no-regret-learning approaches known from game theory. These algorithms adapt transmission probabilities based on utilities, which represent feedback on the success of previous transmission attempts. Our approach extends previous considerations of no-regret techniques for capacity maximization in wireless networks.

In the first part of this thesis, we consider the *capacity-maximization problem*. The task is to choose a maximal-cardinality set of links that can simultaneously transmit successfully. We identify key properties and use them to introduce a flexible proof template for the analysis of adaptive no-regret learning algorithms. This proof template is applied to settings with adversarial jammers, multiple receivers, and links being allowed to leave the network almost arbitrarily. Settings with multiple channels and channel availabilities are also considered by applying similar techniques.

We extend previous works by introducing concurrent networks or malicious devices, which are modeled as adversarial jammers. These jammers are limited and can render all transmissions in a certain fraction of time steps unsuccessful. We show that no-regret learning algorithms are able to converge with approximation factors depending on parameters of the jammer. Assuming these parameters to be known to the algorithm, they even achieve constant-factor approximations when jammers tightly fulfill their limitations or against a stochastic jammer.

The proof template yields the same guarantees when applied to further generalized scenarios. Here, senders do not strive to transmit data to one receiver but to a set of receivers. At the expense of $O(\log n)$ in the approximation guarantees, we also omit the standard as-

sumption from related works that links stay in the network indefinitely.

As the wireless spectrum becomes more and more utilized, regulatory authorities devise plans to let primary users with the exclusive right to use a certain part of the spectrum open up their spectrum for secondary usage. To consider this scenario, we extend the capacity-maximization setting to multiple channels and stochastic channel availabilities. We utilize different notions of regret to achieve constant-factor approximations if the availabilities are uncorrelated among links. Otherwise the guarantees depend on the degree of correlation.

Many wireless devices are actually capable of adjusting their transmission power. Thus, in the second part of this thesis, we consider the *power-control problem*, where we strive to find a minimal power assignment such that all links can transmit successfully. We analyze a well-known fixed-point approach of Foschini and Miljanic in terms of convergence time and relate this to no-regret learning also converging to the optimal power assignment. For no-regret learning algorithms we can even guarantee a certain fraction of time steps to be successful.

Zusammenfassung

Eine bestimmende Eigenschaft von drahtlosen Netzwerken ist die kabellos per Funk stattfindende Kommunikation. Das Funkspektrum unterliegt hierbei einer Vielzahl von Einflüssen. Insbesondere Interferenz gleichzeitiger Übertragungen kann zu erfolgloser Kommunikation führen. Geräte müssen trotz dieser Interferenz agieren.

Die algorithmische Betrachtung von Funknetzwerken impliziert häufig zentralisiert koordinierte Geräte. Diese kann man typischerweise nicht annehmen, da Funknetzwerke von Natur aus dezentralisiert sind. Sogar wenn eine zentrale Instanz existiert, würde die Koordination zu zusätzlichen Übertragungen und Interferenz führen. In dieser Dissertation betrachten wir verteilte adaptive Algorithmen, die nur auf wenige Informationen angewiesen sind, um die Anzahl gleichzeitig erfolgreicher Übertragungen zu maximieren. Wir zeigen beweisbare Garantien für diese adaptiven Protokolle.

Wir modellieren ein Funknetzwerk als Sender-Empfänger-Paare oder Links, deren Übertragungen sich auf einem Konflikt-Graphen basierend beeinflussen. Dieses Interferenzmodell verallgemeinert vorherige Modelle wie Unit-Disk-Graphen und das Signal-zu-Interferenz-plus-Rausch-Verhältnis (signal-to-interference-plus-noise ratio – SINR).

Wir betrachten No-Regret-Algorithmen bekannt aus der Spieltheorie. Diese Algorithmen passen ihre Übertragungswahrscheinlichkeit basierend auf Nutzenfunktionen an, welche den Erfolg vorheriger Übertragungen darstellen. Unser Ansatz erweitert frühere Betrachtungen von Kapazitätsmaximierung in Funknetzwerken mit No-Regret-Techniken.

Im ersten Teil dieser Arbeit betrachten wir das *Capacity-Maximization Problem*, in dem eine Menge maximaler Kardinalität von Links zu wählen ist, sodass Links gleichzeitig erfolgreich übertragen können. Wir identifizieren relevante Eigenschaften und erstellen eine Beweisvorlage für die Analyse von No-Regret-Algorithmen für Capacity-Maximization. Die Beweisvorlage wenden wir auf Szenarien mit böartigen Störsendern, mehreren Empfängern und Links, die das Netzwerk wieder verlassen dürfen, an. Mehrere Kanäle und Kanalverfügbarkeiten werden unter Anwendung ähnlicher Techniken ebenfalls betrachten.

Wir erweitern frühere Arbeiten durch konkurrierende Netzwerke und böartige Störsender, die wir als Gegenspieler modellieren. Diese Störsender sind eingeschränkt und können nur in einem bestimmten Anteil der Zeitschritte Übertragungen stören. Wir zeigen, dass adaptive No-Regret-Algorithmen konvergieren und die erreichten Approximationsfaktoren von den Parametern des Störsenders abhängen. Angenommen diese Parameter sind bekannt, dann erreichen die Algorithmen sogar konstante Approximationsfaktoren, wenn Störsender genau ihre Beschränkungen erfüllen oder gegen stochastische Störsender.

Der Beweis liefert die selben Garantien für weiter verallgemeinerte Szenarien, in denen Sender nicht einen Empfänger erreichen wollen sondern eine Menge von Empfängern. Wir verzichten auch auf die Standardannahme verwandter Arbeiten, dass Links unbegrenzt im

Netzwerk verbleiben. Wir erneuern die Beweisvorlage mit einem Verlust von $O(\log n)$ in der Approximationsgüte.

Weil das Funkspektrum mehr und mehr ausgelastet ist, haben Regulierungsbehörden Pläne ausgearbeitet, dass Primärnutzer, welche das exklusive Recht besitzen bestimmte Frequenzen oder Kanäle zu nutzen, ihr Spektrum für Sekundärnutzer öffnen können. Für dieses Szenario erweitern wir Capacity-Maximization durch mehrere Kanäle und stochastische Verfügbarkeit dieser Kanäle. Wir nutzen verschiedene Arten von Regret um konstante Approximationsfaktoren zu erreichen, wenn die Verfügbarkeit für verschiedene Links unkorreliert ist. Andernfalls hängt die Garantie von der Korrelation ab.

Tatsächlich sind viele Funksender fähig, ihre Sendeleistung zu beeinflussen. Daher betrachten wir im zweiten Teil dieser Dissertation das *Power-Control Problem*, in dem Links die minimale Energie finden sollen, so dass alle Links erfolgreich übertragen können. Wir analysieren einen bekannten Fixpunkt-Ansatz von Foschini und Miljanic im Hinblick auf die Konvergenzzeit und beziehen dies auf No-Regret-Sequenzen, die ebenfalls zur optimalen Energiezuweisung konvergieren. Für No-Regret-Algorithmen können wir sogar einen bestimmten Anteil erfolgreicher Übertragungen garantieren.

Acknowledgments

Many people were part of the process of creating this thesis. First of all, I want to thank my PhD adviser Martin Hoefer. He supported my research and offered helpful advice and a lot of good ideas and inspirations. I am also grateful for having the opportunity to start this PhD project under his advice in the first place.

I also want to thank Martin Grohe for being the second adviser of my thesis. I want to express my gratitude for being able to finish my thesis.

I am really grateful, that Berthold Vöcking made it possible to work at his chair at RWTH Aachen University and that he took the time to coadvise my work and thesis. I was able to strengthen my scientific abilities under his guidance and advice.

Large parts of this thesis are based on results developed in collaboration with Martin Hoefer and Thomas Kesselheim. I really do want to express my thanks to both of them for many in-depth discussions, which held fruitful and inspiring insights.

Furthermore, I would like to thank Oliver Göbel, Klaus Radke, Andreas Tönnis, and Lisa Wagner for proofreading this thesis. Their feedback helped a lot to finish and polish this thesis.

Besides these people, all other former and current members of the algorithms and complexity group at the computer science department of RWTH Aachen University provided an environment, which helped in doing focused and creative research. I wish to thank all of them for the honor and pleasure to work with them.

This work has been supported by DFG through UMIC Research Centre, RWTH Aachen University, and grant Ho 3831/3-1.

Contents

1	Introduction	1
1.1	Network and Interference Model	3
1.1.1	Signal-to-Interference-plus-Noise Ratio and the Physical Model	5
1.1.2	Conflict Graphs	6
1.2	Capacity Maximization	8
1.3	Power Control	9
1.4	No-Regret Learning Algorithms	10
1.5	Outline and Results Overview	12
1.5.1	Bibliographic Notes	14
1.6	Related Work	14
2	Capacity Maximization under Adversarial Jamming	19
2.1	Results	20
2.2	No-External-Regret Algorithms	20
2.3	General Framework	22
2.4	Jamming-Resistant Capacity Maximization	28
2.4.1	Model Description	28
2.4.2	$(T', 1 - \delta)$ -bounded Adversary	29
2.4.3	Unknown T'	31
2.4.4	Unknown δ	34
2.4.5	Stochastic Adversary	35
2.5	Simulation Results	38
2.6	Summary and Further Considerations	41
3	Extensions of the Proof Template for Jamming-Resistant Learning	43
3.1	Results	43
3.2	Multi-Receiver Setting	44
3.2.1	Model Description	44
3.2.2	Approximation Guarantees	46
3.3	Joining and Leaving Links	49
3.3.1	Model Description	49
3.3.2	Approximation Guarantees	50
4	Scheduling with Multiple Channels and Availabilities	55
4.1	Results	56
4.2	Availability and Multi-Channel Model	56

4.3	No-Availability-Ordering Regret	59
4.4	Convergence Results	62
4.4.1	Approximation Factors	62
4.4.2	Uncorrelated Availabilities	68
4.4.3	Lower Bounds	72
4.5	Summary	75
5	Power Control for Capacity Maximization	77
5.1	Results	77
5.2	Power-Control Model	78
5.3	Foschini-Miljanic Iteration	79
5.3.1	Convergence Time	80
5.4	Utilizing No-Regret Learning	83
5.5	Computing No-Swap-Regret Sequences	87
5.6	Convergence of No-Swap-Regret Sequences	89
6	Conclusion and Further Work	97

CHAPTER 1

Introduction

During the last decade, *wireless networks* increasingly gained importance and seemed to become ubiquitous in today's daily life. Most people in our modern society use one or the other kind of wireless communication frequently. The wide range of applications spans over mobile phones to sensor networks, WiFi at home, and many other settings. In sensor networks a large number of devices need to coordinate to achieve a common task like monitoring forest fires etc. Other applications offer means for personal communication like mobile phones and WiFi. Thus, for most individuals today, wireless networks are omnipresent. This way, wireless networks are highly influential for our daily life.

Due to a rising number of applications the number of *protocols* used in networks increases, too. These protocols have to operate in parallel spatially close together and can often not directly coordinate their transmissions.

Wireless networks have to communicate using certain frequencies of the *radio spectrum* as *channels*, which leads to the frequency spectrum becoming more and more utilized. The protocols utilized here need to coordinate the usage of the frequency spectrum, as concurrent transmissions interfere possibly rendering transmissions unsuccessful. For this coordination many protocols use additional transmissions or *acknowledgements* indicating to the sender whether a transmission attempt was successful.

Depending on the actual network, the use case, and regulatory restrictions devices use specific protocols sticking to a specific part of the spectrum and frequency bands. A large portion of the possible spectrum is licensed by regulatory authorities like FCC, OfCom, or the German Bundesnetzagentur for specific use cases, e.g., for mobile phone communication. This way a company can have the exclusive rights to use parts of the spectrum. These license owners are often referred to as *primary users*. Thus, parts of the spectrum can only be used by primary users like telecommunication companies owning the respective license.

Other parts of the radio spectrum are not exclusively licensed by the regulatory authorities and usable for a wide range of applications. Especially devices using this spectrum suffer from having to coordinate the spectrum usage. The spectrum available for these purposes is referred to as the *ISM-band* (Industrial, Scientific, and Medical Band). This part of the spectrum is not only used for scientific purposes but also for widespread applications like WiFi, bluetooth, etc. All these protocols have to deal with interference from other devices using other protocols, and they have to cope with scarce and highly utilized spectrum. Coordination under these influences is one of the major algorithmic challenges in this research domain.

We consider wireless networks to consist of sender-receiver pairs denoted as *links*, where every sender strives to transmit to a respective receiver. This communication critically depends on transmission attempts being successful and yields certain limits on wireless networks.

Wireless networks have to utilize the available spectrum as efficiently as possible. To do so, one crucial task is to decide which transmissions should be conducted. In general, not all transmissions can be executed successfully at the same time. Here, it is typically assumed that the time is divided into discrete *time slots*. The goal often considered in wireless networks is to maximize the number of simultaneous transmissions per time slot. This task is often referred to as *capacity maximization*.

In wireless networks, transmissions are subject to interference from other transmitting links, ambient noise, and other additional influences. Due to the mutual *interference*, transmitting links have to coordinate their transmissions. This was already analyzed in early research considering wireless networks, e.g., in the slotted ALOHA protocol [Abr70] by re-transmitting with an exponential backoff.

The mutual influences of transmitting links and the influences due to ambient noise can be modeled in various different ways. Some of those models are based on graphs [JPPQ05, SW10, HKV11] modeling the conflicts due to concurrent transmissions. Another model often used is the *signal-to-interference-plus-noise ratio (SINR)* [MW06]. Here, the “success” of a transmission is based on the sum of interference from other links. The received power of any sender at a receiver decreases as a power law of the distance between the sender and the receiver. The SINR contrasts the received power of the intended transmission with the power received from other senders and ambient noise. If the SINR is larger than a given threshold, a transmission is considered successful.

Some strategies used to coordinate wireless links utilize centralized approaches to decide which links should transmit in which time slot [GWHW09, GOW07, GK00, HW09]. Practically speaking, these approaches afford wireless devices to have a connection to a centralized authority for coordination. Communication to the central instance introduces additional traffic into the network, and it makes the approach inflexible for application to large and changing networks. Interference also increases this way. Especially in local urban or industrial environments it becomes more and more necessary to implement a large number of links, and thus it is important to communicate as efficiently as possible.

More flexible approaches can also ensure guarantees on approximation factors in distributed settings. This contrasts the performance of the algorithm to that of an optimal solution in terms of capacity maximization. Considering *distributed algorithms*, which can be executed on every single sender, is necessary as wireless networks are inherently decentralized. In particular, *wireless ad-hoc networks*, where links connect to a network without using any kind of predefined connections, cannot rely on a centralized authority or a connection to one to coordinate their communications. Many modern wireless networks are distributed and, thus, it is intuitive to solve algorithmic problems in these networks in a distributed fashion using the least additional communication possible.

In this thesis, we consider adaptive learning algorithms to solve these challenges assuming only very restricted feedback. That is, links are able to adjust their transmission behavior based only on feedback whether their previous transmission attempts were successful. Thus, links learn over time whether to transmit or not. This approach has clear advantages to previously considered ones by introducing less overhead communications for coordination and still guaranteeing good approximation factors for capacity maximization as we show later

in this thesis.

In large-scale networks and ad-hoc networks it can simply be infeasible to give links more than very restricted information about their transmissions or about the network's structure and size. This way, our approach extends and follows a similar thread as previous works [Din10, AM11]. We extend previous considerations concerning adaptive distributed learning algorithms like no-regret learning to more intricate settings. A no-regret learning algorithm iteratively adapts its probability to choose one of multiple possible actions based only on feedback for previously chosen actions.

We start by giving a quite general proof template based on linear programming, we extend the simple setting of one network consisting of n links to a setting which allows concurrent networks not following the same protocol. We even consider adversarial jammers, which are adversaries that render transmissions unsuccessful. By assuming an adversary we take the fact into account that environments can change rapidly due to co-existing networks.

We also omit the assumption that a sender strives to transmit to one receiver by introducing a multi-receiver setting. Another extension allows links to leave the network again. For both of these settings the same proof technique is used as for the jammer setting above.

In real-world wireless networks most devices are actually capable of deciding on which of multiple frequency bands to transmit. This can be used to further increase the utilization of available spectrum. Besides this, legislative institutions all over the world start to temporarily open up licensed spectrum for usage by secondary users [FNK⁺11, Nek09]. That is, primary users owning the exclusive right to use a certain part of the spectrum can allow secondary users to utilize the spectrum whenever the primary user has no need to transmit. Assuming links are able to choose among different channels brings the considered capacity-maximization setting closer to practical settings, where links can utilize such secondary spectrum.

While choosing which links transmit is reasonable for doing capacity maximization, it is not the only measure applicable to optimize wireless networks. We also consider power control to adjust the interference transmitting links exert on other links. In capacity maximization it is typically assumed that links use a fixed power level for transmissions. Using power control, links can reduce the power they use still guaranteeing successful reception. Reducing the transmission powers reduces interference exerted on other links. This reduction of the transmission power also lowers the energy consumption necessary for successful communication. Thus, battery-driven devices can reduce drainage of their batteries. Thereby the number of links transmitting successfully in the same time slot can be significantly increased using power control.

1.1 Network and Interference Model

Communication in wireless networks is subject to interference from concurrent transmissions. That is, other simultaneous transmission attempts influence the reception of a specific transmission at a receiver. This is one of the main differences to wired network models, where transmissions are often assumed to interfere in a *one-or-none* fashion, i.e., as soon as more than one sender emits a signal all transmissions fail. In wireless networks transmissions interfere in a more localized fashion. In general the received signal strength from nearby senders is stronger than the one from senders far away. This propagation and fading

of the signal can be modeled in various ways and can be incorporated in different models to depict the influence one sender can have on other transmissions.

There exist several approaches to model this typical characteristic of wireless networks. We consider some of them throughout this section concluding with a rather general one, on which most of the results in this thesis are based. Before doing so, we need to define the general structure of our wireless network model.

We assume that a *wireless network* is defined as set V of n links. A *link*, denoted by $\ell_v = (s_v, r_v)$, consists of a *sender* s_v and a *receiver* r_v . Senders and receivers are placed in a metric space inducing a *distance* $d(s_v, r_v)$ also denoted as the *length* of the link. In most parts of this thesis, we apply a slightly more general notion to model the mutual interference of links.

Transmissions are performed by senders with the goal that the respective receiver can successfully receive the transmission. Based on the considered interference model these transmissions are either successful or unsuccessful. We also call a transmission *conflict-free* if it is successful at the receiver. On the other hand, we define that a conflict occurs at the receiver and the transmission is unsuccessful based on transmissions of other links and possibly additional ambient noise. Formally, this definition actually depends on the considered interference model.

Note that we abandon the standard notion of links as sender-receiver pairs in parts of this work. Here we define a link to consist of a sender and multiple receivers. The notion of *conflict-free receivers* becomes very useful in this context.

Interference models are used to define the influence concurrent transmissions have on each other. A very simple one exhibits the one-or-none fashion mentioned above. This is often also denoted as a *single-hop network model* [Gal85, Chl01]. While this model is very simple, there exists a variety of more intricate models considering more localized and more practical versions of interference.

One property exhibited in other interference models is that signals in wireless networks degrade over distance. While they can be received with a high signal strength near to the sender, the signal strength typically decreases for receivers far away. There exist different approaches capturing this property in the interference model.

One way to model interference in a very localized way is the *disk-graph model* (see for example [MW05, SW09]). Here, senders and receivers are assumed to be placed in the Euclidean plane. Each sender is depicted by a disk. In the *unit-disk-graph model*, all disks are assumed to have the same radius. Successful transmissions are defined as follows. A transmission can be *received successfully* at a point of the plane – or a receiver situated at this point is *conflict-free* – if the only disk of transmitting senders covering this point is the one of the respective sender. Thus, this models very harshly that signals do not travel indefinitely far through space and cannot be received at distances too far away from the sender. Nevertheless for every single point of the Euclidean plane the one-or-none structure of the single-hop model remains.

One straightforward generalization of the disk-graph model to general metric spaces is the *protocol model* [GK00, JPPQ05]. A transmission from sender s_v to receiver r_v is *successful* if $d(s_v, r_v) \leq (1 + \Delta)d(s_k, r_v)$ holds for all other transmitting senders s_k , where $d(s_i, r_j)$ denotes the distance between sender s_i and receiver r_j and Δ is a constant parameter. Intuitively, a transmission of a link ℓ_v is successful if the length of this link is at most a factor $(1 + \Delta)$ larger than the distance between any other transmitting sender and the receiver. This way, the success of transmission attempts is subject to a pairwise condition of the links.

Although the models above capture the locality of signal propagation, both these models

lack of capturing the fact that signals can propagate arbitrary far with decreasing strength. In practical networks one crucial property is that the signals from different links interfering with a transmission add up. Both these disadvantages of the above models can be resolved by using the *physical model* based on the *signal-to-interference-plus-noise ratio* (SINR) described below. Another property captured by the physical model is that senders can choose different power levels for their transmissions.

A further generalization is the *conflict-graph model* considered afterwards. It generalizes the SINR and the simpler settings mentioned above.

1.1.1 Signal-to-Interference-plus-Noise Ratio and the Physical Model

Many related works, especially those considering the more practical engineering perspective [Rap96], use the *physical model* to model interference in a realistic way. This interference model relies on the *signal-to-interference-plus-noise ratio* (SINR) measuring the quality of the signal. The SINR is the ratio between the strength of the received signal of the intended transmission to the signal strength of other transmissions plus ambient noise. This way, it measures how strong the received signal is compared to interference. A stronger received signal allows a better decoding of the transmitted message.

In the physical model, the received signal strength depends on the power used by the sender and the distance between the sender and the receiver. It decreases as a power law of the distance. Let p_v be the *transmission power* used by sender s_v . Formally, the received power at receiver r_w then is $\frac{p_v}{d(s_v, r_w)^\alpha}$, where $d(s_v, r_w)$ denotes the distance in the metric space between s_v and r_w . The parameter $\alpha \in \mathbb{R}_{\geq 0}$ is a constant *path-loss exponent*. It describes how strong the signal strength degrades over distance. In practical settings this constant is often assumed to be between 2 and 6 [RM92, Rap96]. Nevertheless, our considerations hold for any $\alpha > 0$. One can exchange the distance part of this definition by a general *gain* value $g_{w,v}$. This results in a received power of $g_{w,v} \cdot p_v$ by using $g_{w,v} = \frac{1}{d(s_v, r_w)^\alpha}$.

Note that additionally to the interference of other links, a constant *ambient noise* $\nu \in \mathbb{R}_{\geq 0}$ affects the transmissions. Having defined how the signal propagates, we can now define the signal-to-interference-plus-noise ratio formally. Using the notion of SINR we are able to define under which circumstances a transmission is considered successful in the physical model.

Definition 1.1. Let $L \subseteq V$ be a set of links in a metric space transmitting simultaneously. Let p_v be the power assignment for link ℓ_v and ν the ambient noise. The SINR condition for link ℓ_v is defined by

$$\beta_v \leq \frac{\frac{p_v}{d(s_v, r_v)^\alpha}}{\sum_{\substack{\ell_u \in L \\ u \neq v}} \frac{p_u}{d(s_u, r_v)^\alpha} + \nu} ,$$

where β_v is the SINR threshold of link ℓ_v . The right-hand side is the signal-to-interference-plus-noise ratio (SINR).

A transmission is considered *successful* or *conflict-free* if the SINR condition is fulfilled. That is, the SINR is at least as large as a fixed threshold β_v . Although it is often assumed that this threshold is the same for all links, our approaches also work more generally when each link has its own threshold.

Using this model we do not account for the actual quality of the transmission but only for whether the transmission is successful. In some related settings, the reached SINR can be associated with the data rate achieved by a transmission. This is directly related to the *Shannon-capacity* and the length of packets to transmit [Har28, Tel99].

Some more advanced interference models based on the SINR model also capture additional effects due to the environment. One example is the *Rayleigh-fading model*, where the actually SINR achieved by a transmission attempt is drawn from an exponential distribution with the expected value at the deterministic SINR. The algorithmic problem of capacity maximization analyzed in this thesis was also considered under Rayleigh-fading conditions [DHK12b].

1.1.2 Conflict Graphs

Most parts of this thesis rely on a more general notion of interference than the SINR condition described above. In this section, we generalize and unify the interference models considered above. Using such a general notion of interference in our proofs enables us to derive results in a more general and more flexible way. Our more general framework for interference encompasses different interference models, including the models discussed above or other models based on bounded-independence graphs [SW10].

We model interference based on *conflict graphs*. Intuitively, the unit-disk-graph model and the protocol model can be represented as a directed graph. The graph consists of one vertex per link. An edge from the vertex of link ℓ_v to the vertex of ℓ_w exists if the transmission of ℓ_v renders the transmission of ℓ_w unsuccessful. For unit-disk graphs and the protocol model this construction is straightforward.

To further generalize this kind of conflict graph, we consider an *edge-weighted* version also considered in [HKV11]. In the simple unweighted graph model above, we can identify the set of links being able to transmit simultaneously with an *independent set*, which consists of vertices not being connected via edges in the graph. For our edge-weighted version we extend this notion.

Formally, a *conflict graph* is a directed graph $G = (V, E)$, where the vertices V are identified with the links, and each edge $(u, v) \in E$ is associated with a weight $b_u(v)$. Given a set $L \subseteq V$ of transmitting links, a link ℓ_v can transmit *successfully* (the receiver is conflict-free) if the incoming edge weights of other transmitting links is at most 1.

The set of links which are successful simultaneously can be identified with a weighted version of an independent set. We call such a set being able to transmit successfully in the same time step *feasible*.

Definition 1.2. A set of links L is feasible if the transmissions of all links $\ell_v \in L$ are successful, which means that

$$\sum_{\substack{\ell_u \in L \\ u \neq v}} b_u(v) \leq 1 ,$$

is fulfilled for all links $\ell_v \in L$.

We use C-independence as a key parameter in our proofs. This notion allows to formalize the connection between the underlying interference model used to define the edge weights and the used algorithm.

Definition 1.3. A conflict graph is called *C-independent* if for any feasible set of links L there exists a subset $L' \subseteq L$ with $|L'| \geq \frac{1}{2} \cdot |L|$ and $\sum_{\ell_u \in L', u \neq v} b_u(v) \leq C$ for all $u \in V$, where $|L|$ and $|L'|$ denote the number of transmitting links in these sets.

C-independence constitutes formally that the influence of any transmitting sender on a feasible set is bounded. Intuitively, this property holds in metric spaces under SINR interference with *uniform transmission powers*, where all senders use the same transmission power, as the interference from one sender decreases over distance. The feasible set of links itself needs certain distances between senders and receivers of different links to be successful.

Note that we set the constant in the definition above to be $\frac{1}{2}$ simply for ease. In fact, any constant can be used here. Thus, we can exchange the condition on the size of the set L' by $|L'| \geq z \cdot |L|$ for $z > 0$ being a constant independent of the actual conflict graph (but possibly depending on the interference model). This can also be achieved by utilizing the C-independence definition recursively. This way, C-independence generalizes the bounded-independence property popular in the distributed computing literature [SW10].

While disk graphs and the protocol model can directly be transferred to the conflict graph model, we need some more in-depth consideration to transfer the SINR model here. How to embed the SINR model into the conflict graph was also considered by [HKV11] using the notion of affectance [HW09] defined below.

Definition 1.4 ([HW09, HHMW13]). The affectance $a(w, v)$ of link ℓ_v caused by another link ℓ_w is

$$a(w, v) = \min \left\{ 1, \beta \frac{\frac{p_w}{d(s_w, r_v)^\alpha}}{\frac{p_v}{d(s_v, r_v)^\alpha} - \beta v} \right\}.$$

To use affectance for the construction of the conflict graph, we simply use $b_u(v) = a(u, v)$. In the definition, the SINR threshold is assumed to be β for all links. One can easily use the same construction exchanging β with β_v to consider individual thresholds.

With this the success condition of the conflict graph becomes the SINR condition. This yields for a feasible set L

$$\sum_{\substack{\ell_u \in L \\ u \neq v}} b_u(v) \leq 1 \Rightarrow \frac{p_v}{d(v, v)^\alpha} \geq \beta \left(\sum_{\substack{\ell_u \in L \\ u \neq v}} \frac{p_u}{d(u, v)^\alpha} + v \right).$$

If the distances in the SINR model are based on senders and receivers being placed in a metric space and uniform powers are used, we can see that this construction provides conflict graphs, that are $O(1)$ -independent. The constant is independent of the actual network and conflict graph. This is due to considerations from Ásgeirsson and Mitra [AM11, Lemma 8, Lemma 11].

The considered conflict graph model is actually more general than a union of the protocol and the physical model. For example, the weights or the affectance can be defined using adjusted gain values based on measurements instead of $g_{v,w} = \frac{1}{d(s_v, r_w)^\alpha}$ of the SINR model. This kind of interference model was analyzed by Bodlaender and Halldórsson [BH14] and by Gudmundsdottir et al. [GÁB⁺14] still guaranteeing similar C-independence. Thus, our approaches can be applied to these settings.

1.2 Capacity Maximization

One of the probably most important combinatorial optimization problems in wireless networks is to maximize the capacity of the network, i.e., to maximize the number of links being able to transmit simultaneously in the same time slot. In networks under interference conditions as described in Section 1.1, we want to decide which subset of all links should transmit, such that all these links transmit successfully. Throughout this thesis, mostly this kind of problem is considered in wireless networks based on the weighted conflict-graph model.

Here, the *capacity-maximization problem* is defined as follows. Given a network consisting of a set V of n links, the task is to find the subset of links $L \subseteq V$ maximal in its cardinality $|L|$ being able to transmit simultaneously, i.e., guaranteeing that all receivers of links in L are conflict-free. In other words, we want to find a *feasible one-shot schedule* L of the links. We mainly focus on the conflict-graph-based interference to model the interdependence of links. Thus, we have to select the maximum cardinality subset L of links such that for every link in L the summed up incoming edge weight from other links in L is at most 1, i.e., $\sum_{\ell_u \in L, u \neq v} b_u(v) \leq 1$ for all $v \in L$.

When considering capacity maximization in the SINR model, we assume that links use *uniform* or *fixed powers* in most parts of this thesis. Thus the transmission powers are fixed to the same value for all links when constructing the weights of the conflict graph. We consider the different optimization problem described in the next section utilizing different power levels and even optimizing them.

Our analyses are done in terms of worst-case scenarios, where no additional assumptions about the network are made. In this context capacity maximization is known to be an *NP-hard* problem [GOW07]. Thus, we cannot expect to have efficient algorithms to find the optimal solution (unless $P = NP$). Here, we even consider distributed algorithms in which links do not coordinate their behavior via direct communication.

We introduce randomized algorithms based on game-theoretic assumptions to solve this problem. In game theory the main underlying assumption is that *agents* or *users* behave selfishly given a certain set of possible actions and their outcome in terms of a utility. Over time these algorithms adjust their transmission probability, and the overall network converges. We will show that these learning dynamics even admit a constant-factor approximation after having run for a time long enough to adapt to the network. Here, the *approximation guarantee* or *factor* is given by the ratio of the optimal capacity of the network to the number of successful transmissions of the algorithm.

When considering the unweighted version of conflict graphs, maximization of the network's capacity directly reflects the *maximum-independent-set problem*. The set of links of an optimal solution would then only consist of vertices in the conflict graph such that no two neighboring vertices in the conflict graph are chosen. This also means that the chosen links in the network can transmit with receivers being conflict-free.

For our notion of weighted conflict graphs, the capacity maximization in the same way represents a generalization of the maximum-independent-set problem to edge-weighted graphs, where the incoming edge weights from links in the set is at most 1.

1.3 Power Control

While we consider the capacity-maximization problem defined above throughout most parts of this thesis, we focus on a problem of a different kind in Chapter 5.

In wireless networks, senders can use *power control* to attempt transmissions at different power levels. Many modern devices utilized in wireless networks have the capability to switch between different power levels. Being able to adapt the powers used in the network can dramatically increase the capacity of the network. Links being successful with a fixed power can reduce their power without becoming unsuccessful. This way, the overall interference in the network can be reduced, and possibly additional links can transmit successfully.

Concerning the SINR or physical model, it is obvious that the SINR of a link is directly associated with the power this link uses for transmission (see Definition 1.1). The chosen transmission power influences the SINR linearly. Thus, increasing the power results in an increased SINR and decreasing the power reduces the achieved SINR with interference staying fixed. With power control links being successful can achieve this by exactly fulfilling the SINR constraint.

Power control can reduce interference in the network but it also enables devices to improve *battery life*, which is the time until the battery of a device is drained. Wireless devices being battery driven can reduce the power usage by choosing the smallest power value necessary for successful transmission. As batteries only offer a restricted amount of energy, spending less energy allows devices to stay functional longer, and this is clearly beneficial for the overall performance of a network.

Our considerations mainly assume powers to be fixed and, thus, links use powers in an oblivious way without considering the structure of the network. Concerning the SINR model we basically assume a uniform power setting, where all links use the same power level for all transmissions. This is also assumed by [HHMW13]. Other oblivious power assignments can also be incorporated into the conflict-graph weights using the respective affectance in Definition 1.4. These are power settings that only depend on the respective link. Such power settings can depend on the distance between the sender and the receiver of the links. Two such examples are the linear power assignment, where the power is chosen linear in $d(s_v, r_v)^\alpha$, or the square-root power assignment, where the power is proportional to $\sqrt{d(s_v, r_v)^\alpha}$. Concerning capacity maximization, it is known that even the optimal set of links transmitting under uniform powers is only an $O(\log \Delta)$ -factor away from the optimal feasible set of links under arbitrary powers [AM11]. The parameter Δ denotes the ratio between the length of the longest and the shortest link. For square-root power assignments this factor becomes $O(\log \log \Delta)$ [HHMW13].

Nevertheless, choosing powers in such an oblivious way has some disadvantages. One is the loss in the approximation factors mentioned above and another one is the fact that links have to know the distance between their sender and receiver.

To overcome these disadvantages, we can define the *power-control problem* as follows. Consider a wireless network with n links in a metric space with interference based on the SINR model. We assume that there exists a power assignment such that all links are feasible. The task is to find the minimal power assignment rendering all links successful.

This problem can actually be solved efficiently in a centralized way by solving an appropriate linear program. In wireless networks it is often infeasible to apply such a centralized

solution to the links, as the solution needs to be communicated to the links and also the network needs to be known to the algorithm. Thus, as before, we consider the power control problem in a distributed environment, where links have to decide on their transmission powers for themselves.

Firstly, we assume links get to know their reached SINR and this way links are able to use the distributed approach introduced by Foschini and Miljanich [FM92]. Afterwards, we again consider adaptive learning algorithms which only need to know whether previous transmissions were successful. For the first part, we assume that unconstrained powers can be chosen. For the second part, we assume a *maximal power* exists and only powers $p_v \in [0, p_{\max}]$ can be used.

1.4 No-Regret Learning Algorithms

In this thesis, we solve the considered algorithmic problems from a game-theoretic perspective. We assume links to decide in a distributed fashion striving to optimize the transmission behavior in a selfish way.

Links want to transmit successfully, and they do not want to waste energy by transmitting unsuccessfully. Thus, links selfishly choose whether to transmit or not. This kind of selfishness is often considered in game-theoretic analyses. Here, each link is represented by an *agent* – also called *user*, or *player* – choosing which of multiple actions to play. In our base setting these actions are either to transmit or not. This way, links iteratively play a game against other players and the environment. Over time they adjust their behavior to the network and other players.

To do so, we need to model capacity maximization as a *repeated game*. This is done by defining a suitable *utility function* $u_i^{(t)}(a_i^{(t)}, a_{-i}^{(t)})$ depending on the action $a_i^{(t)}$ chosen by player i (belonging to link ℓ_i) and a vector of actions chosen by other players, which is denoted by $a_{-i}^{(t)}$. Here, actions $a_i^{(t)}$ in every time step t are chosen from an *action space* denoted by $\mathcal{A}$. The definition of the actual utility and the action space depends on the settings considered in the subsequent chapters.

Given such a utility function determined by the environment and other players, we can define this to be a game played where the players want to maximize the overall summed up utility gained over time.

This game and the utility function considered in later chapters is highly related to the one used by Andrews and Dinitz [AD09] and others [Din10, AM11]. It rewards any successful transmission with a utility of +1 and penalizes an unsuccessful one with a utility of -1. Not transmitting at all guarantees a utility of 0. This captures the intuition described above by punishing failed transmission attempts and rewarding successful ones.

When considering power control in Chapter 5, we assume a different utility function. This utility depends on the power used by the link to achieve successful transmissions.

One approach often used to analyze such games is based on equilibrium concepts like the famous Nash equilibrium [NRTV07]. Intuitively, in an equilibrium no player wants to change the strategy unilaterally assuming all other player strategies are fixed. For capacity maximization, analyses of these concepts can for example be seen in [DKV11, AD09].

We examine a different solution concept. It originates from adaptive protocols, whose behavior converges to stable patterns of play. One such intuitive concept is *best-response*

dynamics, where players choose the action that offered the most utility in the previous time step. Note that this does not guarantee the no-regret property used in most parts of our work, which is defined below.

We consider a more intricate kind of dynamic protocol, namely *no-regret-learning dynamics* to provide properties, which we utilize within our proofs. There exists a variety of notions of regret from which we consider different ones depending on the actual setting.

A rather general definition of regret is the following notion of Φ -regret, which can also be found in [BM07, NRTV07, GJM11].

Note that we use the notation $u_i(a_i^{(t)}, a_{-i}^{(t)})$ in the definition below to emphasize the dependence of the utility function on the actions chosen by other players (denoted by $a_{-i}^{(t)}$). We consider an *action vector* $a^{(t)}$ for a time step t consisting of the chosen actions of all players. This definition differs from the ones used in related game-theoretic works by explicitly introducing actions of other players into the definition instead of considering them to be part of the environment or an adaptive adversary defining the utility.

Definition 1.5 ([BM07, NRTV07, GJM11]). *Let Φ be a set of measurable functions such that each $\phi \in \Phi$ is a map $\phi: \mathcal{A} \rightarrow \mathcal{A}$. Given a sequence of action vectors $a^{(1)}, \dots, a^{(T)}$, the Φ -regret encountered by player i is*

$$R_i^\Phi(T) = \sup_{\phi \in \Phi} \sum_{t=1}^T u_i(\phi(a_i^{(t)}), a_{-i}^{(t)}) - u_i(a_i^{(t)}, a_{-i}^{(t)}) ,$$

where u_i is the utility function.

An algorithm computing an action sequence with $R_i^\Phi(T) \in o(T)$ exhibits the no- Φ -regret property and is called a no- Φ -regret algorithm.

Intuitively, regret measures the difference in the summed up utility between a sequence of actions computed by an algorithm and a best choice of actions in hindsight. Depending on the set Φ , the best choice in hindsight can be restricted to a best single action or a certain mapping from chosen strategies to other actions.

The set Φ can be considered as a parameter defining different notions of regret. For example, we can consider Φ to be the set of all constant functions, such that it consists of functions $\phi_{a'}$ for every $a' \in \mathcal{A}$ with $\phi_{a'}(a) = a'$ for all $a \in \mathcal{A}$. This yields the notion of *external regret*, for which we give an alternative formal definition in Section 2.2. The notion of *swap regret* considered and defined in Section 5.4 can also be expressed in terms of Φ -regret by choosing Φ to be the set of all measurable functions. For more details on these considerations, we refer to the respective chapters and to [BM07, GGM08].

Depending on the notion of regret, we can define suitable *equilibrium concepts* to which the sequence of actions played by a no-regret algorithm converges (see for example [BM07, GGM08]). While we stick to *sequences of actions with low regret* in most parts of this thesis, we move to equilibria in Chapter 5. We actually use the notion of *correlated equilibria* associated with no-swap-regret algorithms.

Algorithms yielding a no-regret property work in a randomized fashion, for example, by choosing the actions with a probability proportional to weights, which are updated based on utilities seen for previous actions. Those algorithms get to know the utilities in an online fashion and have to adapt to these. One such example is the *Randomized Weighted Majority algorithm* depicted in Algorithm 1 in Section 2.2. In particular, algorithms offering regret bounds with high probability are helpful here. We say an algorithm has a no-regret property *with high probability (whp)* if it holds with probability at least $1 - \frac{1}{n^c}$ for any constant c .

While some of these no-regret algorithms need to know the utilities of all actions in every time step no matter whether these were played, some algorithms only need to observe the utilities of actions taken by the algorithm. The first ones follow the *full-information model*, assuming that utilities for all actions are known in a time step. This is a strong assumption as the necessary information might not be available in realistic settings. The latter one considers the *partial-information model*, which is often also denoted as *bandit model*. The partial-information model is clearly a more suitable assumption in wireless networks, as it allows to only rely on the feedback for chosen actions. This feedback constitutes for example whether transmission attempts were successful.

We apply different notion of regret and respective no-regret algorithms using utility functions tailored to the setting for capacity maximization and power control throughout this thesis.

1.5 Outline and Results Overview

In this thesis, we consider the problems defined in Section 1.2 and Section 1.3. The general outline of the thesis also follows the same order by first considering the described capacity maximization in Chapters 2 to 4. This constitutes the main part of the thesis and the results span over various different settings. Chapter 5 complements these considerations by analyzing distributed approaches for power control. In these chapters the respective notions of regret are introduced and the performance of no-regret learning algorithms in the context of capacity maximization and power control are analyzed.

In Chapter 2, we begin by formally introducing no-external regret. Afterwards, we establish two major properties, which intuitively are inherent in good approximation algorithms applicable to our distributed online setting. These are the (γ, ϵ) -successful and (η, ϵ) -blocking properties. These properties identify connections between the number of transmission attempts and successful transmissions and between the number of transmissions and interference from other links. Using this, we are able to give a general proof template in Section 2.3. This constitutes one of the main results of this chapter.

The other results of Chapter 2 are derived by using this template. We are able to consider capacity maximization under additional interference conditions subject to an adversarial jammer. That is, besides interference defined using a conflict graph, transmissions can be rendered unsuccessful by an adversary. Allowing this kind of additional influences on the network makes it more difficult to adjust to the network in a distributed fashion. Thus, throughout Section 2.4 we consider different restrictions on the jammer and different knowledge of algorithms about the jammer. We assume the conflict graphs to be C -independent and the external regret of the used algorithms to be polynomially small in the number of links. Using this, we show that, when parameters are public knowledge, no-external-regret learning algorithms can be used to obtain an $O(\frac{C}{\delta})$ -approximation against an adversary, which jams at most $\delta \cdot T'$ time steps of any time window of T' time steps. If one of the parameters is not known, we can still guarantee an $O(\frac{C}{\delta^2})$ -approximation factor. For a stochastic jammer an $O(C)$ -bound is shown to hold with high probability in the long run. These results are complemented with simulations showing an astonishingly close and fast convergence towards the optimal capacity.

We extend these results in Chapter 3. Firstly, we consider a scenario where links do not

consist of sender-receiver pairs but of senders with multiple receivers. In this multi-receiver setting, we assume that senders want to transmit to all of their receivers, to at least one, or that the utility of a link depends on the number of its conflict-free receivers. For the last of these settings, we show that no-external regret applied here cannot guarantee non-trivial approximation factors. For the other multi-receiver settings, we are able to show that up to a constant factor the same parameters and factors hold as used in Chapter 2.

A second extension to the proof template seen before is given in Section 3.3. Here, we reestablish the general result how to apply the proof template in a setting where links are allowed to leave the network again once they achieved low regret. Again assuming (γ, ϵ) -successful and (η, ϵ) -blocking properties, we prove the same general results as before roughly losing a factor of $O(\log n)$ in the approximation guarantees. As both results in Chapter 3 utilize the same properties as in Chapter 2, we are able to ensure that all previous results also hold in the respective multi-receiver settings, and with a slight loss when links can join and leave almost arbitrarily.

In Chapter 4, we consider a different generalization of the capacity-maximization setting. We assume that links can utilize different channels, which are subject to stochastic availabilities. In this setting, algorithms do not only need to decide whether to transmit but also which of the available channels to use. This models scenarios in which a primary user, that has the exclusive right to use a certain part of the wireless spectrum, can allow secondary users to utilize the communication channels in some time steps.

We apply no-availability-ordering-regret learning algorithms tuned to the availability setting. This guarantees approximation factors depending on a newly introduced measure, which extends C -independence to the correlation of availabilities among links. We prove similar bounds for uncorrelated channel availabilities. Here, slightly different notions of regret can be applied. These results are accompanied by lower bounds in Section 4.4.3. We show that arbitrary correlation among the channel availabilities of links can lead to an $\Omega(n)$ -factor for the used notion of regret.

The part of this thesis dealing with power control is situated in Chapter 5. There, we consider the power-control problem, where links have to find a power level such that all links can transmit successfully in the same time slot. To consider power control at all, we resort to the SINR interference model instead of using conflict graphs.

We begin by considering a distributed iterative approach proposed by Foschini and Miljanic [FM92]. The first result here is that this approach converges to the fixed point of the iteration up to a factor $(1 - \delta)$ within $O(n \log n \log \delta)$ time steps for a network with n links. This fixed point coincides with the optimal power assignment. The biggest disadvantage of this iterative protocol is that links need to know their achieved SINR and that there exist runs of the algorithm where no link transmits successfully.

To overcome the drawbacks of the iterative approach by Foschini and Miljanic, we consider no-regret-learning techniques in Section 5.4 and the following ones. We introduce a suitable no-regret approach and show that the fraction of rounds with successful transmissions converges to 1. Our approach is based on swap regret and only needs to know whether previous transmissions were successful to compute a suitable utility function.

This thesis is concluded in Chapter 6 with some comments about possible future research questions, which extend the considered problems.

1.5.1 Bibliographic Notes

Most results presented throughout this thesis have previously been published as joint work.

In detail, the results concerning the jamming-resistant learning in Chapter 2 and parts of the extensions to the considered proof template in Chapter 3 (Section 3.3) have been published in the following paper and its preprint or fullversion.

- [DHK14] Johannes Dams, Martin Hoefer, and Thomas Kesselheim. Jamming-resistant learning in wireless networks. In *Proceedings of the 41st International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 447–458. Springer, 2014

The results assuming links consisting of a sender and multiple receivers can also be found in the preprint of the above paper.

- [DHK13a] Johannes Dams, Martin Hoefer, and Thomas Kesselheim. Jamming-resistant learning in wireless networks. *CoRR*, abs/1307.5290, 2013

Our analysis concerning capacity maximization under multiple channels from Chapter 4 have partially been published in the paper mentioned below. Especially the results concerning uncorrelated availabilities of the channels considered in Section 4.4.2 and the lower bounds considered in Section 4.4.3 can be found in the following paper.

- [DHK13b] Johannes Dams, Martin Hoefer, and Thomas Kesselheim. Sleeping experts in wireless networks. In *Proceedings of the 27th International Symposium on Distributed Computing (DISC)*, pages 344–357. Springer, 2013

We previously published our results concerning power control and the results in Chapter 5 in the following two papers.

- [DHK11] Johannes Dams, Martin Hoefer, and Thomas Kesselheim. Convergence time of power-control dynamics. In *Proceedings of the 38th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 2, pages 637–649. Springer, 2011
- [DHK12a] Johannes Dams, Martin Hoefer, and Thomas Kesselheim. Convergence time of power-control dynamics. *IEEE Journal on Selected Areas in Communications (J-SAC)*, 30(11):2231–2237, 2012

1.6 Related Work

In this section, we give an overview over related works concerning wireless networks. We focus our survey on results being related to the algorithmic problems considered throughout this thesis. Thus, we focus on related work with special regard on capacity maximization and power control in the context of game-theoretic learning dynamics. Nevertheless, we try to give a slightly more general insight by also considering non-game-theoretic results here.

Wireless networks and the algorithmic problems arising in those networks have been subject to many works in the past. Starting with analyses considering the well-known and well-studied ALOHA protocol [Abr70] especially in single-hop networks [Abr70, SK83, ABAJ04] the attention shifted to more intricate interference models like the SINR model since

the works of Moscibroda and Wattenhofer [MW06] and of Gupta and Kumar [GK00]. They considered the problem of maximizing the sum of distances of links being feasible simultaneously in both the unit-disk graph and the SINR model.

Nevertheless, a lot of recent research still sticks to the general idea of synchronized time steps also considered in slotted ALOHA (see for example [Din10, AM11]).

Many topics considered in these works follow the question of how to coordinate in wireless networks. For the single-hop networks considered in these early works the tasks are mainly variants of collision or conflict resolution similar to slotted ALOHA [Chl01, Gal85], where links retransmit with a certain probability after unsuccessful transmissions. Other tasks considered there are information-theoretic approaches like coding [Gal85]. Later extensions of the interference model also considered more intricate problems. These coordination tasks range from capacity maximization, which is highly related to the aforementioned conflict resolution, with and without power control to routing of packets via multiple hops [BMJ⁺98, KKT04], broadcasting a message in the networks [KLN⁺10], or the usage of secondary spectrum [HKV11, HBH06a].

In the single-hop model, it is generally assumed that any concurrent transmission results in interference making any reception of transmissions impossible. This is generalized to more realistic interference models like the protocol model [GK00] and the SINR model [Rap96, MW06]. Several works consider the differences between these two models [GK00, IRK09]. Recent works even extend this to more intricate models where transmitters and senders can even utilize directed antennas [SS03, SBH14].

One of the algorithmic problems based on these models analyzed quite often is the capacity-maximization problem also considered in this thesis. Past research started with scheduling problems similar to capacity maximization considered in the protocol model and went on to the SINR model and further more challenging problem definitions.

In the protocol model, scheduling links is mainly considered in terms of *latency minimization*. That is, instead of finding the maximal feasible subset of all links, one strives to find a partition into a minimal number of sets of links such that each set is feasible. More intuitively speaking, for latency minimization we want to schedule the links allowing feasible transmissions minimizing the time until all links transmitted successfully at least once.

Concerning the graph-based protocol model, latency minimization coincides with the algorithmic problem of finding a minimal coloring in the graph. As this task is NP-complete in general, approximation algorithms were considered [RL93, MW05]. Capacity maximization under graph-based interference models was also analyzed in various works [EJS05, NHK08, SW10].

In the more realistic SINR model both latency minimization and capacity maximization were considered. In [KV10] latency minimization was considered in the SINR model for fixed power assignments. While this was mainly considered using centralized approaches, capacity maximization was also considered in terms of distributed and adaptive learning algorithms. In general, when dealing with centralized algorithms one can easily utilize algorithms for capacity maximization to construct algorithms for latency minimization losing a factor of $O(\log n)$ in the approximation guarantee [GWHW09].

A lot of works were published considering capacity maximization in the SINR model. Centralized algorithms were proposed that provide a constant-factor approximation for the capacity under different additional assumptions [HM12, GWHW09, GHW13]. There are also approaches for the online case, where links arrive one by one [FGHV13, GHK⁺14].

As wireless networks are inherently distributed it is essential for links to consider prob-

lems like capacity maximization in terms of distributed algorithms. Especially works considering techniques from game theory got more attention throughout recent years in this context [Din10, AM11, DP13].

Game-theoretic techniques were considered frequently for capacity maximization under SINR conditions. Those range from analyses concerning Nash equilibria to no-regret-learning dynamics. Capacity maximization can be modeled as a game where links are players having to choose a transmission probability as action [MW01a, BJ10, DKV11]. Here, a (pure) *Nash equilibrium* constitutes a vector of actions or a state where no player wants to unilaterally deviate from this state to one other action. These states exist under certain conditions, and the quality of these states was analyzed in terms of a suitable utility function. Multi-channel environments similar to our considerations in Chapter 4 were for example considered in [NC06].

Other game-theoretic approaches consider capacity maximization using no-regret learning algorithms [Din10, AM11]. We consider some of these works in more detail in the other chapters. These approaches exist for both the unit-disk-graph model and the SINR model. In summary, capacity maximization is modeled as a game and no-regret learning algorithms adjust to the experienced interference and the network structure over time.

Generalizations of the capacity-maximization problem were also considered in recent works. One of which is the jammer scenario we also consider in Chapter 2. This is considered by [ARS08, RSSZ10, RSSZ11, RSSZ13] assuming similar restrictions as we do on the jammer under the single-hop model and the unit-disk-graph model. This approach was also extended to the SINR setting [ORS⁺13]. Possibly the most distinctive difference to the setting in this thesis is that these works are not link-centered but assume that senders want to reach any arbitrary other device. Thus, we clearly extend this by considering adaptive regret learning in a link-centered scenario. Another recent approach by Gilbert et al. [GKP⁺14] considered jamming in a single-hop network when a message needs to be broadcasted to all nodes.

Other extensions to the capacity-maximization problem assume adversarial instances by considering the dual-graph model in which edges of a conflict graph based on the protocol model are determined by an adversary [KLN⁺10, GLN13, GMW08]. In these papers, senders strive to transmit a message successfully to multiple devices doing a broadcast to all devices either in the whole network or locally.

Modern wireless devices can exploit the ability to adjust their transmission powers to increase the achieved capacity of a network. This allows to transmit with the minimal power possible still yielding successful transmissions. This obviously reduces the interference acting on other links in the SINR model. So power control can clearly increase the overall capacity of the network.

Simpler considerations assume links to use preset oblivious power schemes independent of other links like the square-root power assignment considered in [FKRV09, HM11, HHMW13]. These *oblivious power assignments* define a choice of powers for each link depending only on the distance between sender and receiver of the respective link.

A short literature survey concerning more intricate power-control topics can be found in [SK11]. Lots of the related works rely on the same model as considered by Foschini and Miljanic [FM92]. The existence of a power assignment rendering all links feasible is assumed. Here, the task is to find this power assignment [Yat95, YH95, HY98, KG05]. Also other algorithmic approaches than the fixed-point approaches were considered (see for example [LPPP11]).

Other approaches consider game-theoretic solutions for the power-control problem. Similar to the ones dealing with capacity maximization, these can be divided into some considering equilibria as a solution concept [KG05, JH98, ABSA02, AFB⁺08] and others considering adaptive dynamics. These started with converging best-response dynamics like the iteration introduced by Foschini and Miljanic [FM92] (see also [HGGB04, WB01, GZY95]). Nash equilibria in similar settings were also considered [Hei06, HBH06b, AD09]. Some of these works extend the model in terms of throughput via utility functions, which logarithmically depend on the reached SINR [HBH06b, MW01a, MW01b, FMGS02, MCPS06]. This well-studied *Shannon capacity* is often considered in practice [Har28, Tel99]. These considerations involve continuous utility functions. This allows to apply different equilibrium concepts or algorithms. Thus, these concepts clearly differ from our approach, where links have to beat an SINR threshold to transmit successfully.

It is a natural extension to combine power control directly with scheduling (i.e., deciding which links transmit) to do capacity maximization. This combination is known to be NP-hard [AD09]. Thus, it was considered by quite a few papers throughout the recent years [KAY10, BCP95, Kes11, Kes12, Hal12]. All of these approaches divide the task into two separate phases for scheduling and power control. ElBatt and Ephremides [EE04] consider this scenario not only for arbitrary link-based networks but also for networks with base stations and terminals. Thus, these apply to more practice-related setting.

The related works of Kesselheim [Kes11, Kes12] consider power control and capacity maximization that achieve a constant-factor approximation for fading metrics. These works allow flexible data rates using individual SINR thresholds in a centralized way.

Besides these settings, further problems related to capacity maximization were considered in literature. These range from extensions of the power-control setting to incorporate interference cancellation [JJ09] to applications of coding theory and coding schemes [KYR06] or works considering transmission queues and how the queue length stabilizes [ÁHM12]. Other considered topics are multicast [EG10], i.e., transmitting to multiple receivers, or the negative effects of new techniques like power control or interference cancellation on stable states considering selfish agents [DP13]. These topics are not the only ones considered in this context, but we do not strive to give an exhaustive overview here.

CHAPTER 2

Capacity Maximization under Adversarial Jamming

We consider using adaptive learning dynamics to optimize the capacity achieved by a network of n links. In this chapter, we apply *no-external-regret-learning techniques* to show convergence to a transmission behavior, which yields a constant-factor approximation of the maximal capacity for $O(1)$ -independent conflict graphs.

To begin, we introduce the notion of external regret and introduce algorithms with a no-external-regret property. These algorithms update a probability distribution over possible actions (i.e., attempting a transmission or not) based on a utility function as feedback for previously chosen actions. The no-external-regret property guarantees that given this utility function the gained utility for actually chosen actions converges to the utility of the best single action in hindsight. Introducing algorithms exhibiting this property, we consider a quite general framework based on linear programming to analyze these adaptive algorithms for capacity maximization.

This framework constitutes a proof template generalizing previous results and techniques considered when analyzing no-regret learning algorithms for capacity maximization in wireless networks. We apply our framework to no-regret algorithms used in a setting where transmissions are not only subject to interference but also an adversarial jammer rendering transmissions unsuccessful in a certain amount of time steps.

In applied settings, wireless networks are actually subject to influences of other sources than links in the same network. These influences range from concurrent networks and malicious jammers even to interference from other hardware or appliances like, for example, microwave ovens. Concurrent networks might not follow the same protocol as the considered network and this way do not coordinate with the network. We model these kinds of external interference by assuming an adversarial jammer similar to related work [RSSZ10, RSSZ11, ARS08].

Later in Chapter 3, we then show that the proof template considered here can also be applied to a wider range of settings. Those include settings where a sender does not only try to transmit to one receiver but multiple ones and a setting where links can leave the network again once their own regret is low.

2.1 Results

To use our general proof template, as a first result we identify two properties intuitively inherent to algorithms with good approximation factors. These are namely the (γ, ϵ) -*successful* and (η, ϵ) -*blocking* properties. Depending on these parameters we derive bounds on the capacity achieved by algorithms yielding low average external regret in scenarios with C -independent conflict graphs.

Using this result, we apply the proof template to different settings where a jammer renders transmissions unsuccessful. Clearly, it is necessary to limit the abilities of jammers, as otherwise jammers can simply make all transmission attempts unsuccessful. We consider different assumptions on how the jamming adversary can be constrained. We prove a constant-factor approximation guarantee for constant-independence or an $O(C)$ -factor for general C -independence assuming that the adversary jams exactly $(1 - \delta) \cdot T'$ time steps of any time window of T' time steps. If the jammer does at most but not exactly destroy $(1 - \delta) \cdot T'$ time steps, we lose a factor of δ in the guarantees. When the parameters bounding the jammer are not known, we introduce algorithms with an additional loss in the approximation guarantees of δ each for unknown T' and unknown δ . We also consider a stochastic adversary against which our algorithms achieve a constant approximation factor.

To conclude, we also draw a line from these theoretical results to more applied results by conducting simulations. Our algorithms perform favourably in these simulations. They show that the factors and guarantees seem to be even better in realistic settings than expected from the previous theoretical results in worst-case analysis.

2.2 No-External-Regret Algorithms

We consider algorithms for links to adjust to the network structure and the interference over time. In total we want to maximize the capacity, which is the number of successful transmissions. Our algorithms for capacity maximization are based on no-regret-learning techniques. As described before, links decide independently in every time slot whether to transmit or not using an appropriate learning algorithm.

Our assumption that algorithms do not have any additional information about the structure of the network requires algorithms to adjust their behavior based only on the outcome of previous decisions. This outcome is either a successful transmission or an unsuccessful one. To measure the quality of the outcome of an action, a suitable *utility function* $u_i^{(t)}(a_i^{(t)}, a_{-i}^{(t)})$ is used depending on action $a_i^{(t)}$ chosen by player i in time step t and on actions $a_{-i}^{(t)}$ chosen by other players in t .

In the setting considered throughout this chapter, there are only two possible actions in each time step – sending or not sending. The detailed problem setting is defined in Section 2.3, and we use a utility function depending on the details of the respective setting. The utility functions $u_i^{(t)}$, which is defined in the subsequent sections, strikes a balance between interference minimization and throughput maximization for each link. We also account for different kinds of adversarial jamming.

The adaptive algorithms, which are considered here, are learning algorithms yielding a game-theoretic property called no-external regret. Given the respective setup with appropri-

ate utility functions, we assume links apply arbitrary such no-regret learning algorithms that minimize a notion called external regret. The *external regret* for an algorithm or a sequence of chosen actions is defined below.

In terms of Φ -regret introduced in Section 1.4, we can define the set Φ to the set of constant functions. Thus, the functions map all actions to only one action as the best action in hindsight. Alternatively we can use the following definition.

Definition 2.1. Let $a^{(1)}, \dots, a^{(T)}$ be a sequence of action vectors. The external regret of this sequence for link ℓ_i is defined by

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

where $\mathcal{A}$ denotes the set of possible actions. An algorithm has the no-external-regret property or is a no-external-regret algorithm if the external regret of the computed sequence of actions grows in $o(T)$.

Note that for simplicity we also speak of regret instead of external regret throughout this and the next chapter. The utility in Definition 2.1 can be considered to be given by an adaptive online adversary possibly depending on previous decisions of link ℓ_i . In fact, in the capacity-maximization setting it is determined by other links and their transmission decision, which by itself might depend on previous actions of link ℓ_i . Thus, we use $u_i^{(t)}(a_i, a_{-i})$ as a notation for the utility, where a_{-i} denotes the actions chosen by all links but ℓ_i .

Algorithms like the famous and surprisingly simple *Randomized Weighted Majority algorithm* (RWM) by Littlestone and Warmuth [LW94] or the variant called *Hedge* [FS97] guarantee this no-regret property by updating a probability distribution over the actions without actually calculating the regret. This algorithm is depicted in Algorithm 1.

Algorithm 1: Randomized Weighted Majority [LW94, FS97]

```

Initialize weights  $w_a = 1$  for all actions  $a \in \mathcal{A}$ ;
foreach  $t \in T$  do
     $W = \sum_{a \in \mathcal{A}} w_a$ ;
    Choose action  $a \in \mathcal{A}$  with probability  $p_a = \frac{w_a}{W}$ ;
    foreach  $a \in \mathcal{A}$  do
        Observe loss  $l^t(a)$ ;
         $w_a = w_a \cdot (1 - \rho)^{l^t(a)}$ ;

```

The way the algorithm is stated here uses the notion of loss $l^t(a) \in [0, 1]$ instead of utility for an action $a \in \mathcal{A}$ in time step $t \in T$, where $\mathcal{A}$ is the set of all available actions and $\rho \in [0, \frac{1}{2})$ is a suitable chosen parameter. The loss can easily be constructed from a utility by multiplying with -1 and scaling appropriately.

Choosing the parameter ρ appropriately allows the Randomized Weighted Majority algorithm to have the no-regret property. A detailed proof of the following lemma can be found in [NRTV07, FS97].

Lemma 2.1 ([NRTV07]). Setting $\rho = \sqrt{\frac{\ln |\mathcal{A}|}{T}}$ Algorithm 1 has expected external-regret at most $\sqrt{T \ln |\mathcal{A}|}$.

Note that to implement this lemma, links need to know the total time T it will run beforehand. This might be impossible in many settings. In these cases, one can drop back to the “guess and double” approach considered for example by Blum [Blu98]. Guessing a certain T and doubling this guess once the time step T is reached still yields a no-regret guarantee.

While the famous RWM algorithm needs feedback also for actions not chosen in a time step, it seems useful to consider algorithms only using feedback for actually chosen actions in our setting. There exist several such *partial-feedback* or *bandit-model* no-regret algorithms [ACBFS02, Aue03, MB04].

One such algorithm is introduced by Auer et al. [ACBFS02]. It even yields a no-regret guarantee with high probability as given in the lemma below. The algorithm works similarly as RWM by updating a probability distribution in a multiplicative weights fashion. It is applicable in the bandit model, where after each round the weights are only updated for the action chosen by the algorithm.

Lemma 2.2 ([ACBFS02]). *Algorithm Exp3.P by Auer et al. yields regret in $O\left(\sqrt{T|\mathcal{A}|\ln\left(\frac{|\mathcal{A}|T}{c}\right)}\right)$ for any $c > 0$ with probability $1 - c$.*

This lemma can be used to get a high probability bound on the regret. Having this bound holding for an algorithm allows us to use this algorithm on each link altogether giving a high probability bound on the time needed until all links have low regret.

We use the no-external-regret property in a rather general way in the following sections to derive bounds on the performance of algorithms for capacity maximization. There is no explicit restriction to a specific algorithm as long as the used algorithm provides low regret. Different links are thus also not assumed to use the same algorithm.

2.3 General Framework

In this section, we introduce a general framework usable to analyze regret-learning dynamics for capacity maximization. Though we do not stick to this framework in all later chapters, this general framework gives a first insight into our analysis techniques. We show that this framework can be used in different settings and that already known results can also be reestablished using our novel proof template. This way, our approach unifies and extends previous analyses of simpler problem variants.

Depending on the actual setting, we define appropriate utility functions to adapt no-external-regret learning algorithms to the respective setting. A central idea in our construction is not to let the algorithm update its transmission probabilities in every time step. To overcome some of the difficulties introduced in the more intricate settings we will use algorithms updating the probability distribution after an interval of multiple time steps instead of learning after every time step. We call such an interval of k time steps a *phase*, where k will be chosen appropriately in the respective settings. A link is assumed to stick to the same chosen action throughout the whole phase observing the outcome of its decision in every time step. So a link will either transmit for a phase of k consecutive time steps or not throughout this phase. This way, we adapt no-regret learning algorithms such that one round (update step) of the algorithm coincides with a phase and not with a single time step. By $\mathcal{R}_v$ we denote the *set of phases* of link ℓ_v .

For our learning algorithm we consider a phase in which the link transmits to be either successful or not. This depends on the fraction of successful transmissions in this phase. It is *successful* if link ℓ_v attempted transmission throughout the phase and a fraction $\mu \in (0, 1]$ of time steps within the phase have been successful. We use μ as a parameter in our proof to address specific settings below. In this and the following chapters, we use the same proof template. Consider a computed sequence of actions. Let q_v denote the fraction of phases in which ℓ_v attempted transmission and w_v the fraction of successful phases.

We assume the conflict graphs of the considered wireless networks to be C-independent. As the first step, we identify two properties indicating the relation between attempted and successful transmissions and between attempted transmissions and interference from other links. These properties depend on parameters, which differ depending on the considered setting. Our general proof uses the notions of (γ, ϵ) -successful and (η, ϵ) -blocking algorithms and C-independence to establish an approximation guarantee depending on these parameters γ , η , and C. It also relies on the parameter ϵ , which denotes the external regret averaged over the phases (instead of over the time steps – c.f. Section 2.2). Where suitable, we speak of ϵ_v to emphasize that this regret might be different for different links.

Let $\epsilon = \sum_{v \in V} \epsilon_v$ be the *summed up average regret per phase*. Now we start by defining the properties.

Definition 2.2. A sequence of actions is (γ, ϵ) -successful for a link ℓ_v if it holds

$$\frac{1}{\gamma} \cdot (2w_v + \epsilon) \geq q_v .$$

This represents the relation between attempted and successful transmissions. Both this and the (η, ϵ) -blocking below capture the intuition of how a good approximation algorithm might behave. Being (γ, ϵ) -successful implies that a certain fraction of phases with attempted transmissions in a computed sequence of actions must be successful. It seems intuitively obvious that for a “good algorithm” a certain fraction of its transmission attempts is successful. This property roughly states that an $\left(\frac{2}{\gamma}\right)$ -fraction of all transmission attempts is successful. Otherwise the algorithm would have decided not to transmit. In subsequent sections we will see that the no-external-regret property can be used to guarantee this property.

In a very similar way, we have (η, ϵ) -blocking constituting a second intuitive property for good approximation algorithms. This somehow mirrors being (γ, ϵ) -successful.

Definition 2.3. A sequence of actions is (η, ϵ) -blocking for a link ℓ_v if with $q_v \leq \frac{1}{4}\eta - \epsilon$ we have for the fraction of unsuccessful phases due to other links f_v (independent of whether ℓ_v transmits)

$$f_v \geq \frac{1}{4}\eta \quad \text{and} \quad \sum_{u \in V} b_u(v) q_u \geq \frac{1}{8}\eta .$$

The notion of (η, ϵ) -blocking allows to obtain a bound on the incoming edge weights from other transmitting links. Intuitively speaking, an algorithm sending seldom would have only done so if it would have been unsuccessful often. Otherwise it would have had the opportunity to exploit low interference for successful transmission. Thus, every link that rarely attempts transmission must have experienced a lot of interference. To model this property, f_v in the definition above is the fraction of unsuccessful phases not restricted to those phases in which ℓ_v transmits.

Note that the definition is slightly different to the η -blocking definition in previous publications. This allows to prove a bound on the approximation factor depending on the sum of regrets instead of demanding every link to have a low regret.

We observe that both these properties can be achieved by no-external-regret learning even in non-trivial settings, for example, with jammers later in this chapter. The main result of this section is the following theorem.

Theorem 2.1. *Suppose an algorithm computes a sequence of actions which is (η, ϵ_v) -blocking and (γ, ϵ_v) -successful for all links ℓ_v .*

For this sequence of actions, the average number of successful transmissions yields

$$\sum_{v \in V} w_v \in \Omega\left(\frac{\mu\gamma\eta}{C} |\overline{OPT}| - \mu \cdot \varepsilon\right),$$

where $|\overline{OPT}|$ denotes the average optimum (or the single-slot optimum) and $\varepsilon = \sum_{v \in V} \epsilon_v$.

Here, $|\overline{OPT}|$ can denote the single-slot optimum, which assumes the optimum to be the same in all time steps. To consider settings with changing optima, $|\overline{OPT}|$ can instead refer to the average optimum over all time steps $\frac{1}{T} \sum_t |\overline{OPT}^{(t)}|$, where $\overline{OPT}^{(t)}$ is the maximum-cardinality set of links transmitting successfully in time step t . It follows from this theorem, that if all links have regret low enough (i.e., $\epsilon_v \leq \frac{\gamma\eta}{C_n}$) the approximation ratio of the no-regret sequence is in

$$O\left(\frac{C}{\mu \cdot \gamma \cdot \eta}\right).$$

To prove this theorem, we use a primal-dual approach. That is, we consider a *primal linear program (LP)* and its *dual* such that the optimal capacity is related to a feasible solution of the primal LP and the regret sequence to a feasible solution of the dual LP. This way the value of the objective function of the dual is an upper bound for this value of the primal. This allows us to relate the optimal capacity to the regret sequence.

We construct a primal LP suitable for this proof. In subsequent chapters, we will use different LPs to comply to the respective settings considered. For a primal LP, its dual LP can be constructed. A primal LP consists of an objective function to be maximized based on variables which are subject to constraints. An assignment to the variables is *feasible* if the constraints of the LP are fulfilled. The objective value of any feasible solution of the primal is bound from above by the objective value of any feasible solution of the dual LP. This fact will be useful throughout our proof. Using a scheduling with optimal capacity and the one resulting from the no-regret learning algorithm as solutions of the respective LP yields the claim.

Proof of Theorem 2.1. The primal LP is given below. See also [HM12] for an example, where such an LP is used within an algorithm directly. Here, we use this LP only for the analysis.

In the linear program, the objective is to maximize the number of successful transmissions. Thus, we sum over the variables x_v as the objective function and relate x_v to the fraction of time steps in which link ℓ_v transmits in the optimal solution. Obviously a link can at most transmit in all time steps, which leads to the second constraint. The first constraint is the C-independence for every link ℓ_u introduced in Definition 1.3 and is this way fulfilled for a feasible set of links transmitting simultaneously.

Note that the linear program is actually not directly tailored to finding a feasible set of links with maximal cardinality. Instead of directly constituting the interference condition, it considers the weaker C -independence condition. Nevertheless, the optimal value of the objective function of the LP is at least the optimal capacity.

Recall that interference from link ℓ_u exerted on link ℓ_v is defined via a conflict graph with edge weights $b_u(v)$.

$$\begin{aligned} \max. \quad & \sum_{v \in V} x_v \\ \text{s.t.} \quad & \sum_{v \in V} b_u(v) \cdot x_v \leq C \quad \forall u \in V \\ & x_v \leq 1 \quad \forall v \in V \\ & x_v \geq 0 \quad \forall v \in V \end{aligned}$$

Let OPT' denote the set L' for $L = OPT$ from Definition 1.3 of C -independence. Here, OPT is the set of links transmitting in the optimum. Thus, it holds $|L'| \geq \frac{1}{2} \cdot |L|$, where $\frac{1}{2}$ can be substituted with any constant independent of the actual network. Besides this it holds $\sum_{v \in L'} b_u(v) \leq C$ for all $u \in V$.

We set x_v to represent the single-slot optimum $x_v = 1$ if link ℓ_v is in OPT' and $x_v = 0$ otherwise. This ensures that the constraints of the primal LP are fulfilled.

In a similar way, we can use the average optimum instead of a single-slot one by defining OPT_t to be the set of links transmitting to achieve the optimal capacity in time step t . Let OPT'_t then be defined as L' for $L = OPT_t$ from the C -independence definition.

Let T be the number of time steps. Note that, for simplicity we write $t \in T$ for time steps instead of $t \in \{1, \dots, T\}$. We set $x_v = \frac{|\{t \in T \mid \ell_v \in OPT'_t\}|}{T}$. So x_v represents the fraction of time steps in which ℓ_v is in the optimum of all time steps or L' , respectively. This averaging choice of x_v results in

$$\sum_{v \in V} b_u(v) x_v = \frac{1}{T} \sum_{t=1}^T \sum_{v \in OPT'_t} b_u(v) .$$

As C -independence holds for all single OPT_t we have $\sum_v x_v \geq \frac{1}{2} \cdot \frac{1}{T} \sum_t |OPT_t|$ giving a lower bound on the primal objective function. We also conclude

$$\frac{1}{T} \sum_{t=1}^T \sum_{v \in OPT'_t} b_u(v) \leq \frac{1}{T} \sum_{t=1}^T C = C .$$

So we have constructed feasible solutions for the primal LP based on the single-slot optimum or the average optimum.

Now we consider the dual LP and use the average number of transmissions in an action sequence with average-regret ϵ_v for link ℓ_v to find a feasible solution. In total, this then leads to an upper bound on the solution of the primal LP proving the theorem.

$$\begin{aligned} \min. \quad & \sum_{v \in V} C \cdot y_v + \sum_{v \in V} z_v \\ \text{s.t.} \quad & \sum_{u \in V} b_u(v) y_u + z_v \geq 1 \quad \forall v \in V \\ & y_v, z_v \geq 0 \quad \forall v \in V \end{aligned}$$

Roughly speaking, the objective function of the dual LP is the sum over dual variables, which we relate to the number of transmissions in the no-regret sequence.

To construct a feasible solution for the dual LP we set $y_v = \frac{1}{\eta} \cdot 8q_v$ and $z_v = \frac{1}{\eta} \cdot 4(q_v + \epsilon_v)$.

Note that in the first constraint at least one of both terms on the left hand side needs to be large enough to fulfill the constraint. So, the solution needs to strike a balance between interference from other links and the own transmissions of link ℓ_v .

With this choice of variables, the value of the objective function is in $O\left(\sum_v \frac{q_v}{\eta}\right)$ for small ϵ_v . Later on we will relate this to the fraction of successful transmissions. First of all, we check whether this solution fulfills the constraints.

The choices for y_v and z_v are clearly greater or equal to zero. To see that this solution fulfills the other constraint we consider two cases. If $q_v \geq \frac{1}{4}\eta - \epsilon_v$, it directly follows that $z_v \geq 1$ and fulfills the constraint.

We consider the case of $q_v < \frac{1}{4}\eta - \epsilon_v$ now. By (η, ϵ) -blocking in Definition 2.3 we directly get the fraction of unsuccessful phases $f_v \geq \frac{1}{4}\eta$. Remember that f_v does not only consider phases in which ℓ_v actually tries to transmit but all phases. This directly yields $\sum_{u \in V} b_u(v)q_u \geq \frac{1}{8}\eta$. Thus, the constraint is also fulfilled in this case.

It is now left to consider the objective functions of the primal and dual LP. The feasible solutions directly yield $\sum_{v \in V} x_v \leq \sum_{v \in V} C \cdot y_v + \sum_{v \in V} z_v$. That is

$$\sum_{v \in V} \frac{|\{t \in T \mid \ell_v \in OPT'_t\}|}{T} \leq \sum_{v \in V} C \cdot \frac{1}{\eta} \cdot 8q_v + \sum_{v \in V} \frac{4}{\eta} q_v + \sum_{v \in V} \frac{4}{\eta} \epsilon_v .$$

We can use the (γ, ϵ) -successful property in Definition 2.2 to relate the fraction of phases with transmission attempts q_v to the fraction of successful phases w_v . We get

$$\begin{aligned} \sum_{v \in V} \frac{|\{t \in T \mid \ell_v \in OPT'_t\}|}{T} &\leq \sum_{v \in V} C \cdot \frac{12}{\eta} \cdot \frac{2}{\gamma} (w_v + \epsilon_v) + \sum_{v \in V} \frac{4}{\eta} \epsilon_v \\ &\leq \sum_{v \in V} C \cdot \frac{12}{\eta} \cdot \frac{2}{\gamma} w_v + C \cdot \frac{16}{\eta} \cdot \frac{2}{\gamma} \epsilon . \end{aligned}$$

Remember that a phase is a consecutive interval of k time steps. Note that $k = 1$ is used in some settings to learn in single time slots. A phase of length k is defined to be considered successful if at least $\mu \cdot k$ of its time steps have a successful transmission. This allows us to relate the number of successful phases to the number of successful time steps losing an additional factor of μ .

$$\sum_{v \in V} \frac{|\{t \in T \mid \ell_v \in OPT'_t\}|}{T} \leq \sum_{v \in V} C \cdot \frac{24}{\eta \gamma \mu} \sum_t \frac{w_v^{(t)}}{T} + C \cdot \frac{32}{\eta \gamma} \epsilon ,$$

where $w_v^{(t)} = 1$ if link ℓ_v transmits successfully in time step t and $w_v^{(t)} = 0$ otherwise.

As already discussed, applying the definition of C -independence for the optimum as a feasible set of links results in $\sum_{v \in V} |\{t \in T \mid \ell_v \in OPT'_t\}| \geq \frac{1}{2} \cdot \sum_{v \in V} |\{t \in T \mid \ell_v \in OPT_t\}|$. In total we get the claimed approximation guarantee or even an approximation guarantee in $O\left(\frac{C}{\eta \gamma \mu}\right)$ for $\epsilon_v < \frac{1}{4\eta} \eta \gamma$ for all links ℓ_v with respect to the primal optimum (when at least one link can transmit successfully).

Again note that both the single-slot optimum and the average optimum can be used to construct the solution of the primal LP. Thus, we have proven the theorem. $\square$

While we use this proof template in the following sections to prove bounds on no-regret learning for wireless capacity maximization, we can also use it to reestablish the results proven by Ásgeirsson and Mitra [AM11]. This gives us a first intuition on how general this proof technique can be applied. Nevertheless, in some of the settings considered in later chapters, we need to apply more specialized proofs, to accommodate for the complexity of the respective setting, without complicating this general framework and losing the simplicity and intuition at this point.

Ásgeirsson and Mitra considered capacity maximization in wireless networks under SINR fading. Their model did not include any further influences like jamming adversaries, multiple receivers, or joining and leaving links as done in the next sections and chapters or even multiple channels as considered in Chapter 4.

To reestablish their result we simply need to argue that the setting is $(1, \epsilon)$ -successful and $(1, \epsilon)$ -blocking. The $(1, \epsilon)$ -successful property directly reflects Lemma 8 of [Din10] (or Lemma 6 of [AM11]). Though those lemmas are formulated in terms of the summed fraction of successful transmissions $\sum_v w_v$ and the summed fraction of transmission attempts $\sum_v q_v$, their proof directly proves the action sequence to be $(1, \epsilon)$ -successful by considering “not sending” as the possible best action in hindsight. Lemma 7 of [AM11] directly shows the $(1, \epsilon)$ -blocking property when considering the definition of the conflict graph based on the SINR model.

Corollary 2.1. *Any sequence with regret ϵ is $(1, \epsilon)$ -successful and $(1, \epsilon)$ -blocking for capacity maximization in the scenario considered by Dinitz [Din10] and Ásgeirsson and Mitra [AM11].*

This way we directly reestablished the results known for capacity maximization in this base setting already proven by Ásgeirsson and Mitra [AM11]. This result constitutes the following corollary, which we can directly derive from Corollary 2.1 and Theorem 2.1.

Corollary 2.2 ([AM11]). *No-external-regret learning converges to a constant-factor approximation for capacity maximization in the SINR interference model with uniform powers.*

Deriving this previously known result reusing our proof template shows the flexibility of this LP-based proof technique. We use the general result in Theorem 2.1 and the notion of (γ, ϵ) -successful and (η, ϵ) -blocking algorithms to apply no-regret learning to different settings in the following section and the next chapters. Those range from introducing adversarial jammers into the general capacity-maximization setting to considering joining and leaving links or multiple receivers.

In later chapters, we switch to slightly different proofs abstracting from this first proof template. This is necessary to adapt to the settings considered there.

Note that Ásgeirsson and Mitra [AM11] also argue that using uniform powers like our algorithm still guarantees an $O(\log \Delta)$ -approximation for general power assignments under general metrics, where Δ is the ratio between the lengths of the longest and shortest links in the network. They actually show that the optimum under arbitrary powers is at most a $(\log \Delta)$ -factor better than the optimum under uniform powers. It is easy to realize that we can easily apply this to our results, because it also holds for the single optima considered. Thus, for arbitrary powers we lose a factor of $\log \Delta$ in the approximation factors proven in the following sections and chapters using our proof template.

2.4 Jamming-Resistant Capacity Maximization

We have established a general way to consider approximation guarantees for capacity maximization in wireless networks in the previous section. This allowed to reestablish previously known results using this proof template. Now, we focus on a more intricate setting also applying the same template.

We assume that concurrent networks, which do not follow the same protocol, or other (possibly adversarial) influences are modeled using the adversarial jammer defined below. Then, we apply the proof template to different variants of this jammer setting, and we derive suitable approximation guarantees.

2.4.1 Model Description

To model the jamming behavior in wireless networks, we rely on *jammer models* considered in previous works [RSSZ10, RSSZ11, ARS08]. Similarly as in these previous works we assume that there is an adversary rendering transmission attempts unsuccessful. It is clear that an adversarial jammer without any restrictions would simply render all transmissions unsuccessful. Therefore, we restrict the jammer by preventing it from blocking all time steps and making communication impossible.

In general, a jammer decides on a per-time-step-basis whether to destroy transmissions in the specific time step. We distinguish between the following variants of such adversaries.

- A (*global*) $(T', 1 - \delta)$ -*bounded* adversary can jam at most a $(1 - \delta)$ -fraction of the time steps in any time window of length T' or larger.
- We also consider the special case of a (*global*) $(T', 1 - \delta)$ -*exact* adversary, which exactly jams an $(1 - \delta)$ -fraction of any time window of length T' .
- As a third variant, we treat a (*global*) *stochastic* adversary, where we assume any time step to be independently jammed with a probability $1 - \delta$.

Whereas these adversaries jam the channel globally for all links, we also consider *individual* adversaries being able to block each link individually. This leads to similar definitions of *individual* $(T', 1 - \delta)$ -*bounded*, *individual* $(T', 1 - \delta)$ -*exact* and *individual stochastic* adversaries. The restrictions the jammer has to obey are the same on the type and number of jammed time slots for each link, but the jammer decides individually for each link if a slot is jammed. The respective $(1 - \delta)T'$ bounds hold for each link individually. Thus, the individual jammers are clearly stronger than the corresponding global ones. Note that the random trials of the individual stochastic adversary can be correlated between links but are assumed to be independent between time steps.

When the (individual) adversary jams a time slot, every attempted transmission (of the jammed link) in this time slot becomes unsuccessful no matter whether or not it would have been successful based on the interference model. As discussed earlier, we assume that links receive only success or failure of their own transmissions as information. And due to this, they cannot distinguish whether a transmission failed due to adversarial jamming or interference from other transmissions. Thus, a protocol has to base the decisions about transmission only on the feedback of success or failure of previous time steps.

2.4.2 $(T', 1 - \delta)$ -bounded Adversary

We will now apply the framework seen in Section 2.3 to a setting with adversarial jamming behavior. As a starting point, we assume that links know the bounds on the jammer being T' and $1 - \delta$ for the $(T', 1 - \delta)$ -bounded adversary. Knowing these parameters enables the algorithms to adjust their behavior to the jammer in an efficient way. This results in an $O(\frac{C}{\delta})$ -approximation against the bounded jammer in general. In the special case of global exact jammers no-regret learning algorithms even guarantee a constant-factor approximation.

To apply the framework, we need to utilize and adapt the no-external-regret algorithms to have provable bounds on the (γ, ϵ) -successful and (η, ϵ) -blocking properties. This way, we bound the parameters γ and η .

As described earlier, we use no-external-regret algorithms not learning in every time step but after a phase of a certain number of time steps. When T' is known to every link, algorithms can clearly first observe the outcome of multiple time steps to better adjust to the jammer. We set the length of a phase to $k = T'$. This allows to see the outcome of transmissions in a time window in which the jammer is restricted before adjusting the transmission probabilities.

The utility function applied after a phase is defined via $\mu = \frac{1}{2}\delta$. This ensures that a phase is considered successful iff at least $\frac{1}{2}\delta T'$ time steps are successful. Thus, at least half of the time steps guaranteed to be unjammed need to be successful. For a successful phase a utility of $+1$ is applied and for an unsuccessful one -1 . This utility function is inspired by [Din10]. It penalizes transmitting unsuccessfully and rewards successful transmissions.

Let $w_u^R = \frac{|\{t \in R \mid \ell_u \text{ transmits successfully}\}|}{|R|}$ be the fraction of successful transmissions during phase R . Then the utility function can be formalized as follows.

$$u_i^{(R)}(a_i, a_{-i}) = \begin{cases} 1 & \text{if } i \text{ transmits and } w_u^R \geq \frac{1}{2}\delta \\ -1 & \text{if } i \text{ transmits and } w_u^R < \frac{1}{2}\delta \\ 0 & \text{otherwise.} \end{cases}$$

Here a_i represents the decision whether to transmit throughout phase R . Recall, that a_{-i} denotes the actions (i.e., transmission decisions) of other links.

Using no-external-regret learning as explained above yields an approximation guarantee in $O(\frac{1}{\delta})$ for individual bounded jammers and a constant one against global exact jammers for constant C -independence. For general C -independence these bounds are a factor C worse. Firstly, we will prove the following lemma bounding the parameters used in the proof template.

Lemma 2.3 (cf. [Din10, AM11, DHK12b]). *Every action sequence with average regret per phase ϵ with respect to the utility above is $(1, \epsilon)$ -successful.*

The general proof idea follows the same thread as used in related work [AM11, Din10]. We compare the overall utility reached by the learning algorithm (or action sequence) to the overall utility reachable by sending throughout all time steps or not sending at all. As the difference between the algorithm's utility and the utility of one action in hindsight is bounded by the external regret, we can use this to find a suitable bound on the parameters.

In this and the next proof, we will use the following notation. Let $Q_v^R = 1$ if link ℓ_v transmits in phase R of the no-regret sequence and $Q_v^R = 0$ otherwise. Similarly we define

$W_v^R = 1$ if phase R is considered a successful phase, i.e., the fraction of successful transmissions throughout phase R is $w_u^R \geq \frac{1}{2}\delta$, and $W_v^R = 0$ otherwise.

In our proofs, we consider the fraction of phases in which ℓ_v transmits and the fraction of successful phases. Formally, that is $q_v = \frac{1}{|\mathcal{R}_v|} \sum_{R \in \mathcal{R}_v} Q_v^R$ and $w_v = \frac{1}{|\mathcal{R}_v|} \sum_{R \in \mathcal{R}_v} W_v^R$ respectively.

Proof of Lemma 2.3. The utility gained by the algorithm is $|\mathcal{R}_v|(w_v - (q_v - w_v))$, as there are $w_v|\mathcal{R}_v|$ successful phases and $(q_v - w_v)|\mathcal{R}_v|$ unsuccessful ones with transmission attempts. For comparison we consider the utility of not sending at all, which is always 0.

The difference between those two values is a lower bound on the regret ϵ of the no-external-regret algorithm. In average over the phases this then is $(q_v - w_v) - w_v \leq \epsilon$. By simple calculus, we get $2w_v + \epsilon \geq q_v$ directly proving the action sequence to be $(1, \epsilon)$ -successful as claimed. $\square$

In a similar fashion we will now consider (η, ϵ) -blocking and prove the following lemma.

Lemma 2.4. *Every action sequence with average regret ϵ with respect to the utility above is $(1, \epsilon)$ -blocking.*

Proof. We now prove the $(1, \epsilon)$ -blocking property. Here, we compare the utility of the algorithm with the utility gained when transmitting in all phases.

The $(1, \epsilon)$ -blocking property is equivalent to the following formulation. For $\sum_{R \in \mathcal{R}_v} Q_v^R < \frac{1}{4}|\mathcal{R}_v| - \epsilon|\mathcal{R}_v|$ we have $\sum_{R \in \mathcal{R}_v} F_v^R \geq \frac{1}{4}|\mathcal{R}_v|$ and $\sum_{u \in V} \sum_{R \in \mathcal{R}_u} b_u(v) Q_u^R \geq \frac{1}{8}|\mathcal{R}_v|$, where the sum $\sum_{R \in \mathcal{R}_v} F_v^R$ denotes the number of phases with $w_v^R < \frac{1}{2}\delta$ when ℓ_v would have transmitted. Thus $F_v^R = 1$ iff the phase would have been unsuccessful if ℓ_v would have transmitted.

If link ℓ_v transmitted in all time steps, the utility would have been $(|\mathcal{R}_v| - \sum_{R \in \mathcal{R}_v} F_v^R) - \sum_{R \in \mathcal{R}_v} F_v^R$. The no-regret algorithm in contrast gains at most a utility of $\sum_{R \in \mathcal{R}_v} Q_v^R$ as the link can be successful at most in these phases. Thus, the regret ϵ yields a bound of

$$\left(|\mathcal{R}_v| - \sum_{R \in \mathcal{R}_v} F_v^R \right) - \sum_{R \in \mathcal{R}_v} F_v^R - \sum_{R \in \mathcal{R}_v} Q_v^R \leq \epsilon \cdot |\mathcal{R}_v| .$$

Let us now assume that $\sum_{R \in \mathcal{R}_v} F_v^R < \frac{1}{4}|\mathcal{R}_v|$. We will see that this results in a contradiction with $\sum_{R \in \mathcal{R}_v} Q_v^R < \frac{1}{4}|\mathcal{R}_v| - \epsilon|\mathcal{R}_v|$. Plugging in both $\sum_{R \in \mathcal{R}_v} F_v^R$ and $\sum_{R \in \mathcal{R}_v} Q_v^R$ yields

$$\begin{aligned} & \left(|\mathcal{R}_v| - \sum_{R \in \mathcal{R}_v} F_v^R \right) - \sum_{R \in \mathcal{R}_v} F_v^R - \sum_{R \in \mathcal{R}_v} Q_v^R \leq \epsilon|\mathcal{R}_v| \\ \Leftrightarrow & \left(|\mathcal{R}_v| - \frac{1}{4}|\mathcal{R}_v| \right) - \frac{1}{4}|\mathcal{R}_v| - \frac{1}{4}|\mathcal{R}_v| + \epsilon|\mathcal{R}_v| \leq \epsilon|\mathcal{R}_v| \\ \Leftrightarrow & \frac{3}{4}|\mathcal{R}_v| - \frac{1}{4}|\mathcal{R}_v| - \frac{1}{4}|\mathcal{R}_v| = \frac{1}{4}|\mathcal{R}_v| \leq 0 . \end{aligned}$$

This is a contradiction and proves $\sum_{R \in \mathcal{R}_v} F_v^R > \frac{1}{4}|\mathcal{R}_v|$.

It is directly obvious that for an unsuccessful phase the number of time steps unsuccessful in this phase is at least $\mu \cdot k = \frac{1}{2}\delta T'$ and the sum of incoming edge weights from transmitting links of the conflict graph is at least 1 in every such time step. While this would

result in having at least $\frac{1}{2}\delta$ average incoming edge weights from transmitting links, we will now actually argue that the sum of $b_u(v)$ from transmitting links on link ℓ_v is higher.

Assuming that every link uses the phase-based construction given in this section, it is obvious that all links ℓ_u transmitting in one time step also transmit throughout one whole phase of their algorithm. Note that we do not assume phases of different links to be synchronized. Nevertheless, one phase of link ℓ_v overlaps with at most two phases of any link ℓ_u and vice versa. These two arguments together prove that the summed incoming edge weights from transmitting links for a link ℓ_v , which results from the overlapping phases is

$$\sum_{R \in \mathcal{R}_v} \sum_{\substack{R' \in \mathcal{R}_u \\ R' \cap R \neq \emptyset}} b_u(v) Q_u^{R'} = \sum_{R' \in \mathcal{R}_u} \sum_{\substack{R \in \mathcal{R}_v \\ R' \cap R \neq \emptyset}} b_u(v) Q_u^{R'} \leq 2 \cdot \sum_{R' \in \mathcal{R}_u} b_u(v) Q_u^{R'}.$$

Having $\sum_{R \in \mathcal{R}_v} F_v^R > \frac{1}{4} |\mathcal{R}_v|$ we can bound the incoming edge weights from overlapping phases of other links and we get

$$\sum_{R \in \mathcal{R}_v} \sum_{\substack{R' \in \mathcal{R}_u \\ R' \cap R \neq \emptyset}} b_u(v) Q_u^{R'} > \frac{1}{4} |\mathcal{R}_v|.$$

By the arguments given before this results in $\sum_{R' \in \mathcal{R}_u} b_u(v) Q_u^{R'} \geq \frac{1}{8} |\mathcal{R}_v|$, which directly proves $(1, \epsilon)$ -blocking for our algorithm. $\square$

Having chosen $\mu = \frac{1}{2}\delta$ we can use Theorem 2.1 to prove convergence to an $O(\frac{C}{\delta})$ of the average optimum and also of the single-slot optimum. In the given jammer scenario it is easy to see that this directly yields the same approximation factor against individual $(T', 1 - \delta)$ -bounded adversaries. Assuming a global jammer every time slot is either jammed (and thus the optimal capacity is 0) or the single-slot optimum is optimal. This way we see that the average optimum against a global $(T', 1 - \delta)$ -exact adversary is a δ -portion of the single-slot optimum, because only a δ -fraction of time slots exhibits an optimal capacity greater than 0. Let $|\overline{OPT}|$ again denote the average optimal capacity.

Theorem 2.2. *Consider a sequence of action vectors with average regret per phase of ϵ_v for links ℓ_v and $\epsilon = \sum_{v \in V} \epsilon_v$. The number of successful transmissions per time step is in $\Omega\left(\frac{\delta}{C} \cdot |\overline{OPT}| - \delta\epsilon\right)$ against individual $(T', 1 - \delta)$ -bounded adversaries and in $\Omega\left(\frac{1}{C} |\overline{OPT}| - \delta\epsilon\right)$ against global $(T', 1 - \delta)$ -exact adversaries.*

This theorem directly follows from plugging Lemmas 2.3 and 2.4 into Theorem 2.1. Note that for constant C -independence (e.g., when considering the SINR model), this even yields a constant-factor approximation against global-exact adversaries.

2.4.3 Unknown T'

In the previous section, we applied our proof template to settings with $(T', 1 - \delta)$ -bounded jammers. This way we showed an approximation factor in $O(\frac{C}{\delta})$ for C -independent conflict graphs. The utility function used there to apply no-external-regret algorithms needs to know δ and T' to apply the idea of using phases of length $k = T'$. In practical settings this might not be the case. It is still crucial that a jammer is restricted and thus not able to destroy

communications in all time steps to preserve provable performance. An algorithm might actually not know these bounds. Thus it is a natural extension of the model to consider that links do not know the parameters restricting the jammer.

At first, we assume that T' is unknown. In Section 2.4.4 we will then also consider the setting where δ is not known to the algorithm. We design an algorithm and its utility function capturing the jammer without considering T' in the phase length. Thus we set $k = 1$ and this way links learn and adjust the transmission probability in every time step.

The following utility function can then be used to again find guarantees on the (γ, ϵ) -successful and (η, ϵ) -blocking properties for suitable parameters γ and η .

$$u_i^{(t)}(a_i, a_{-i}) = \begin{cases} 1 & \text{if } \ell_i \text{ transmits successfully in } t \\ -\frac{\delta_v}{2-\delta_v} & \text{if } \ell_i \text{ transmits unsuccessfully in } t \\ 0 & \text{otherwise} \end{cases}$$

Compared to the utility function used in the previous section, this utility allows to account for the jammed time steps in the utility function instead of using suitable phase lengths. The unsuccessful transmissions possibly due to jamming are considered by reducing the penalty, which a link gets for unsuccessful transmission attempts.

This utility function yields the following theorem provable with similar techniques seen in the section above. Note that to apply this utility function, the algorithm uses $\mu = 1$ by differentiating between successful and unsuccessful transmissions in every time slot.

Theorem 2.3. *Every sequence of action vectors with average regret per phase of ϵ_v for links ℓ_v and with $\epsilon = \sum_{v \in V} \epsilon_v$ yields the number of successful transmissions per time step in $\Omega\left(\frac{\delta^2}{C} \cdot |\text{OPT}| - \epsilon\right)$ against individual $(T', 1 - \delta)$ -bounded adversaries and in $\Omega\left(\frac{\delta}{C} |\text{OPT}| - \epsilon\right)$ against global $(T', 1 - \delta)$ -exact adversaries.*

We will now prove Lemma 2.5 and 2.6 below providing the $(\frac{\delta}{2}, \epsilon)$ -successful and (δ, ϵ) -blocking properties. Theorem 2.3 then directly follows from Theorem 2.1.

As we consider $k = 1$ in our construction here, we cannot account for the jammer as easily as in the proof of Theorem 2.2 before. Here, it might have a quite significant effect on the algorithm if the jammer actually jams less than an $(1 - \delta)$ -fraction of time steps of a time window. That is, an algorithm might consider unsuccessful time steps to be unsuccessful due to the jammer and this way underestimates the interference from links. Thus we consider δ'_v to be the fraction of time steps, in which the adversary really jams. Against an $(T', 1 - \delta)$ -exact adversary it holds $\delta'_v = \delta$ and for bounded adversaries we have $\delta'_v \geq \delta$. Let again as before $q_v = \frac{|\{t | \ell_v \text{ transmits in } t\}|}{T}$ be the fraction of time steps in which ℓ_v transmits and $w_v = \frac{|\{t | \ell_v \text{ transmits successfully in } t\}|}{T}$ the fraction of time steps with successful transmissions of ℓ_v .

Lemma 2.5. *Every sequence of actions with average external regret ϵ for the given utility is $(\frac{\delta}{2}, \epsilon)$ -successful.*

Proof. To prove that a sequence is $(\frac{\delta}{2}, \epsilon)$ -successful, we consider the utility reached by the algorithm. That is $w_u - (q_u - w_u) \cdot \frac{\delta_v}{2-\delta_v}$. In total, we will show that this yields

$$q_u \leq \frac{4}{\delta} w_u + \epsilon \frac{2}{\delta}.$$

The difference between the utility of not sending throughout all time steps and the utility of the algorithm is

$$(q_u - w_u) \frac{\delta_v}{2 - \delta_v} - w_u \leq \epsilon .$$

The inequality can be easily rearranged to get

$$q_u \frac{\delta_v}{2 - \delta_v} \leq \left(1 + \frac{\delta_v}{2 - \delta_v}\right) w_u + \epsilon \leq 2w_u + \epsilon ,$$

and thus we have proven $q_u \frac{\delta_v}{2} \leq 2w_u + \epsilon$. $\square$

Now we consider (η, ϵ) -blocking and show that an action sequence with low regret fulfills the conditions as needed to prove Theorem 2.3. That is, we show that the (δ, ϵ) -blocking property holds.

Lemma 2.6. *Every sequence of actions with average regret ϵ for the given utility is (δ, ϵ) -blocking.*

Proof. Let f_v again be the fraction of all time steps in which ℓ_v would have been unsuccessful due to other links transmitting. Having ℓ_v transmit in all time steps would therefore result in a fraction of $\delta'_v - f_v$ successful transmission out of all transmissions and in a fraction of $1 - \delta'_v + f_v$ unsuccessful ones. This yields $-(1 - \delta'_v) \frac{\delta}{2 - \delta} - f_v \frac{\delta}{2 - \delta} + (\delta'_v - f_v)$ as the utility of transmitting in all time steps. The utility gathered by the algorithm is again bounded by q_v and thus yields

$$-(1 - \delta'_v) \frac{\delta}{2 - \delta} - f_v \frac{\delta}{2 - \delta} + (\delta'_v - f_v) - q_v \leq \epsilon .$$

Note that $-(1 - \delta) \frac{\delta}{2 - \delta} = \frac{1}{2} \frac{\delta^2}{2 - \delta} - \frac{1}{2} \delta$ holds by easy calculus. As by definition $(1 - \delta'_v) \leq (1 - \delta)$, this implies

$$\frac{1}{2} \frac{\delta^2}{2 - \delta} - \frac{1}{2} \delta - f_v \frac{\delta}{2 - \delta} + (\delta'_v - f_v) - q_v \leq \epsilon .$$

For contradiction we plug in the assumption that the number of interfered time steps is low, which means $f_v < \frac{1}{4} \delta$. This yields

$$\frac{1}{2} \frac{\delta^2}{2 - \delta} - \frac{1}{2} \delta - \frac{1}{4} \frac{\delta^2}{2 - \delta} + \delta'_v - \frac{1}{4} \delta - q_v \leq \epsilon .$$

Transmitting only in a low fraction of all time steps $q_v < \frac{1}{4} \delta - \epsilon$ and having $\delta'_v \geq \delta$ shows

$$\frac{1}{4} \frac{\delta^2}{2 - \delta} + \frac{1}{4} \delta - \frac{1}{4} \delta + \epsilon \leq \epsilon_v .$$

This directly results in the contradiction necessary for our proof. Then the summed up conflict-graph weights are

$$\sum_{u \in V} b_u(v) q_u \geq \frac{1}{4} \delta \geq \frac{1}{8} \delta .$$

This concludes the proof. $\square$

2.4.4 Unknown δ

We have just seen that the assumption of knowing T' can safely be omitted. The proven approximation guarantees only lose a δ -factor. We now consider what happens if the other parameter δ , on which our algorithm relies, is unknown. The approach described here holds for both cases considered above – the case where T' is known and the case where it is not known. Thus we consider links to either know T' – and only not know δ – or not to know both.

To overcome the additional difficulties introduced by not having any information about the jammer's bounds we resort to a coordinated and synchronized approach here. Though we still assume that all links learn and behave independently, we assume that they all join the network at the same time step and run the algorithm in a synchronized fashion. This allows us to cope with this scenario by just adding an additional δ factor to the approximation guarantee.

We adjust the algorithms for known δ from Section 2.4.2 and 2.4.3 used by links to run in a coordinated way such that different values for δ can be tested by all links simultaneously. The general idea is that different δ values are tested, each in a sufficient large number of time steps. Thus the correct δ actually used by the jammer is tested in a large enough fraction of all time steps. In this fraction of time steps, the algorithms behave as if they “play” against an exact adversary. Thus, we can get the approximation factor from the bounds against $(T', 1 - \delta)$ -exact jammers in the sections before.

More formally, we assume links to execute the phases of the learning algorithm in a synchronized fashion. The algorithms start by considering $\delta = \frac{1}{2}$ in half of all time steps. In a quarter of all time steps links consider $\delta = \frac{1}{4}$ and so forth – they consider any $\delta = \frac{1}{2^h}$ in an $\frac{1}{2^h}$ -fraction of all time steps for any $h \in \mathbb{N}$ with $h > 0$. It is easy to see that one can arrange the phases to accomplish this scheme of testing different δ . The algorithms can use $\delta = \frac{1}{2}$ in every second phase, $\delta = \frac{1}{4}$ every second phase of the remaining ones and so on. Considering algorithms having same phase length for different δ like our algorithms do allows this construction easily. Here, we stick to the phase length as given by the respective algorithm in Section 2.4.2 and 2.4.3. Thus generally speaking, our construction uses $\delta = \frac{1}{2^h}$ in a $\frac{1}{2^h}$ -fraction of all time steps.

For each such δ , links have to provide a no-external-regret sequence independent of the other δ -values. Observe that the regret bounds used above are achieved after a polynomial number of phases with high probability using an appropriate no-external-regret algorithm. Thus, we need to consider only a polynomial number of no-regret sequences until convergence, and the algorithm only needs to keep track of a polynomial number of simultaneous algorithm runs.

In fact, the algorithms considered before yield the proven approximation factors after time polynomial in δ . Thus, the time necessary in total is only an additional δ -factor larger. After T time steps it is clear that the smallest δ tested yet is the one being only used in 1 time step, thus, being $\delta = \frac{1}{T}$. By the construction we would have used $\log(T)$ different independent no-regret sequences until then.

In phases where up to a factor of 2 the correct δ is assumed the approximation factors for exact adversaries from the sections before also hold up to a factor of 2. Having an approximation bound in $O(C)$ with C -independent conflict graphs against global $(T', 1 - \delta)$ -exact adversaries for known T' and $O(\frac{C}{\delta})$ for unknown T' , we can directly conclude a bound on the approximation guarantee of the described approach with unknown δ . Due to the fact

that we only guess the correct δ up to a factor of 2 within a δ -fraction of all time steps the following theorem follows.

Theorem 2.4. *Let $\varepsilon = \sum_{v \in V} \epsilon_v$ with average regret ϵ_v for links ℓ_v . There exists synchronized algorithms that yield the number of successful transmissions per time slot being in $\Omega\left(\frac{\delta}{C} \cdot |\overline{OPT}| - \delta\varepsilon\right)$ against any (T', δ) -bounded adversary without knowledge of δ , and $\Omega\left(\frac{\delta^2}{C} \cdot |\overline{OPT}| - \varepsilon\right)$ without knowledge of δ and T' .*

Note that the running time increases by a factor of $\frac{1}{\delta}$ over the asynchronous case, as we need the regret to be sufficiently low in the phases with the correct assumption on δ .

2.4.5 Stochastic Adversary

In this section, we assume jamming behavior different to the adversarial jammers considered before by considering stochastic jamming. To stay consistent with the already discussed notation we assume that the jammer jams a time step with probability $1 - \delta$ in a global fashion. To deal with this kind of adversary, we slightly modify the algorithm seen in Section 2.4.2. Roughly speaking, we lengthen phases such that the stochastic jammer resembles an $(T', 1 - \delta)$ -exact adversary.

More formally, we extend the phases such that with high probability only a small fraction of phases is considered unsuccessful due to jammed time slots. Besides this, we argue that the optimum also converges to the single-slot optimum. We answer the question after how many time slots jamming behavior is close enough to the expected behavior to guarantee a constant approximation factor.

Note that the (γ, ϵ) -successful property only depends on the considered utility function and is independent of the jammer. Thus, the same parameter γ still holds when the algorithms from previous sections are applied here.

As a first step we show that the (η, ϵ) -blocking property can be transferred from the results considered before to this stochastic setting given that the algorithms run a sufficiently long time.

In the remainder of this section, we will bound the probability that an algorithm considers a phase unsuccessful due to interference although the phase is unsuccessful due to jamming. Afterwards, we will see that the (η, ϵ) -blocking property stays valid with the same parameter besides a constant after a sufficient amount of time steps. After a certain amount of time, the actual optimum approaches the expected one. These insights are then combined to get an approximation guarantee.

Let p_k be the probability that a phase of length k is considered unsuccessful due to the stochastic jammer even if it would not have done so for an exact jammer.

Recall that a phase is considered successful if the fraction of successful time steps during this phase is at least $\mu \cdot k$. Thus, the algorithm considers a phase unsuccessful due to the stochastic jammer if the jammer destroys more than $(1 - \delta + \mu) \cdot k$ time steps. Using this, we prove the following lemma.

Lemma 2.7. *Let $\mu < \delta$. Then for $k \geq 1$ it holds*

$$p_k \leq \exp\left(-\frac{\left(\frac{\mu}{\delta}\right)^2 \delta k}{2}\right).$$

Proof. Let $X^{(t)} \in \{0, 1\}$ be a random variable indicating whether time step t is not jammed for link ℓ_v . Then we consider $Y = \sum_{t \in R} X^{(t)}$ the number of unjammed time steps for an arbitrary phase $R \in \mathcal{R}_v$. We get

$$p_k = \Pr[Y < (\delta - \mu) \cdot k] = \Pr\left[Y < \left(1 - \frac{\mu}{\delta}\right) \mathbf{E}[Y]\right].$$

Here the equality follows from the expected value of Y being $\mathbf{E}[Y] = \delta \cdot k$.

As the stochastic adversary draws its jamming in every time step stochastically independent, we can apply Chernoff bounds (see for example [DP98]) on the probability to bound p_k from above.

$$\Pr\left[Y < \left(1 - \frac{\mu}{\delta}\right) \mathbf{E}[Y]\right] \leq \exp\left(-\frac{\left(\frac{\mu}{\delta}\right)^2}{2} \delta \cdot k\right).$$

□

Having bounded p_k we now bound the number of phases until the (η, ϵ) -blocking property also holds for the stochastic setting. This bound will obviously depends on p_k . We say that this bound holds *with high probability (whp)* when it holds with a probability of $1 - \frac{1}{n^c}$ for any constant c .

Lemma 2.8. *Consider an algorithm that computes a sequence of actions which is (η, ϵ) -blocking against an $(T', 1 - \delta)$ -exact adversary. After $T \geq \frac{\max\{p_k, 1 - p_k\}}{\eta^2} \cdot 8^2 \cdot 3 \cdot c \cdot \ln(n) + \ln(n)$ phases, the computed sequence is $(\frac{\eta}{2}, \epsilon)$ -blocking against a stochastic adversary with probability at least $1 - \frac{1}{n^c}$.*

Proof. Remember that (η, ϵ) -blocking yields a bound on the unsuccessful phases of $\frac{1}{4}\eta$ for links sending rarely (see Definition 2.3). To consider $(\frac{\eta}{2}, \epsilon)$ -blocking we will bound probability that $\frac{1}{8}\eta$ phases more than expected are unsuccessful due to the adversary. Note that the expected behavior of the stochastic adversary matches that of an exact one, and so we can prove the $(\frac{\eta}{2}, \epsilon)$ -blocking property.

Let $X_v^{(R)} \in \{0, 1\}$ be a random variable for every phase R and link ℓ_v indicating whether R is not unsuccessful due to the jammer. This way $Y_v = \sum_{R \in \mathcal{R}_v} X_v^{(R)}$ is the number of unjammed phases. The number of phases unsuccessful due to the jammer then is given by $\bar{Y}_v = T - Y_v$.

As mentioned above, we will bound $\Pr[Y_v \geq \mathbf{E}[Y_v] - \frac{1}{8}\eta T]$.

We first consider the case $\eta < 8p_k$, which yields

$$\Pr\left[Y_v \geq \mathbf{E}[Y_v] - \frac{1}{8}\eta T\right] = 1 - \Pr\left[Y_v < \mathbf{E}[Y_v] - \frac{1}{8}\eta T\right].$$

This evaluates to $1 - \Pr\left[Y_v < \left(1 - \frac{1}{8 \cdot p_k} \eta\right) \mathbf{E}[Y_v]\right]$.

Using Chernoff bounds we get

$$\Pr\left[Y_v < \left(1 - \frac{1}{8 \cdot p_k} \eta\right) \mathbf{E}[Y_v]\right] \leq \exp\left(-\frac{\left(\frac{\eta}{8p_k}\right)^2}{2} p_k \cdot T\right).$$

Similarly, the second case $\eta \geq 8p_k$ results in

$$\Pr\left[\bar{Y}_v > \left(1 + \frac{1}{8 \cdot (1 - p_k)}\eta\right)\mathbf{E}[\bar{Y}_v]\right] \leq \exp\left(-\frac{\left(\frac{\eta}{8(1-p_k)}\right)^2}{3}(1 - p_k) \cdot T\right).$$

In both cases, we can apply a union bound over all links and get an upper bound on the respective probabilities of

$$n \cdot \exp\left(-\frac{\eta^2}{3 \cdot 8^2 \cdot \max\{p_k, 1 - p_k\}} \cdot T\right).$$

Setting $T \geq \frac{\max\{p_k, 1 - p_k\}}{\eta^2} \cdot 8^2 \cdot 3 \cdot (c + 1) \cdot \ln(n)$ directly yields that the number of phases unsuccessful due to the jammer exceeding the expectation is at most $\frac{1}{8}\eta$ with high probability. Thus, any $(1, \epsilon)$ -blocking algorithm against the exact jammer would provide the $(\frac{\eta}{2}, \epsilon)$ -blocking property here. $\square$

For a global jammer, the optimal capacity in the stochastic setting might actually be better than the optimal capacity against an exact-adversary. So we would like to bound the number of time steps until the actual optimum converges to the optimum against an exact-adversary. We will prove that with high probability the stochastic optimum is not to much larger after a polynomial number of time steps. With this fact we can easily conclude the approximation factor.

Lemma 2.9. *After $T \geq \frac{8^2}{3\delta} \cdot c \cdot \ln(n)$ time steps, the optimum against the stochastic adversary is at most $\frac{9}{8}$ of the optimum against an exact adversary with probability $1 - \frac{1}{n^c}$.*

Proof. In the same way as in the proof of Lemma 2.8 we consider $X^{(t)} \in \{0, 1\}$ to indicate whether time step t is not jammed. This way $Y = \sum_{t \in R, R \in \mathcal{R}_v} X^{(t)}$ is again the number of unjammed time steps. To bound the probability that more than $\frac{9}{8}$ of the expected time steps are not jammed, we take a look at

$$\Pr\left[Y \leq \left(1 + \frac{1}{8}\right)\mathbf{E}[Y]\right] = 1 - \Pr\left[Y > \left(1 + \frac{1}{8}\right)\mathbf{E}[Y]\right].$$

Using Chernoff bounds in a similar way as before we can directly bound this to be

$$1 - \Pr\left[Y > \left(1 + \frac{1}{8}\right)\mathbf{E}[Y]\right] \geq 1 - \exp\left(-\frac{1}{8^2 \cdot 3}\delta \cdot T\right).$$

Setting $T \geq \frac{8^2}{3\delta} \cdot c \cdot \ln(n)$ gives the bound on the probability as requested by the lemma. Having this way only $\frac{1}{8}$ more unjammed time steps against an stochastic jammer compared to an exact adversary, we conclude that the average optimum is at most a factor of $\frac{9}{8}$ higher than the optimum against an exact adversary. This proves the lemma. $\square$

Summing up, we have seen that after a time $T \in O\left(\frac{\max\{p_k, 1 - p_k\}}{\eta^2} \ln(n)\right)$ any algorithm being η -blocking against an exact adversary is again $(O(\eta), \epsilon)$ -blocking and after a time $T \in O\left(\frac{\ln(n)}{\delta}\right)$ the optimum is close to the optimum against an exact adversary. In total, we can conclude the following corollary. We can simply set the length of each phase to be $k \geq \frac{2}{\delta} \cdot \ln(8)$ to use our bound on p_k from Lemma 2.8 showing $p_k \leq \exp\left(-\frac{\mu^2}{\delta^2} \ln(8)\right) \leq \frac{1}{8}$.

Corollary 2.3. *With high probability, by setting $k = \frac{2}{\delta} \cdot \ln(8)$ the algorithm in Section 2.4.2 yields a number of successful transmissions in $\Omega\left(\frac{1}{C}|\overline{OPT}| - \delta\epsilon\right)$ after $T \in O(\ln(n))$ phases against a (global) stochastic adversary. Here, $\epsilon = \sum_v \epsilon_v$ for average regrets ϵ_v for links ℓ_v .*

This corollary can easily be shown by setting $\mu = \frac{\delta}{2}$ and using the $(1, \epsilon)$ -blocking algorithm from Section 2.3 for known T' . This k implies $1 - p_k > \frac{7}{8}$ by Lemma 2.7, and in Lemma 2.8 we use this. Together with Lemma 2.9 this proves the claim.

The total runtime necessary for this result due to phase length and number of phases is in $O\left(\frac{\ln(n)}{\delta}\right)$. The no-external-regret algorithms as presented in Section 2.2 typically need a polynomial number of learning steps to provide low regret as used for the results in Section 2.2. Thus, we end up with a run time polynomial in the number of links n and δ .

Applying the same arguments to the algorithm of Section 2.4.3 yields a slightly worse bound.

2.5 Simulation Results

In this section, we draw a line from the theoretical results shown in the previous section to a more practical point of view. We simulate wireless networks with jamming utilizing no-regret learning algorithms. Although we consider the scenario with jammer here, similar behavior can also be seen in simpler settings, as considered by Ásgeirsson and Mitra [AM11].

In the simulations of Ásgeirsson and Mitra [AM11] the no-external-regret learning algorithm in a random network shows a quick convergence to a stable solution within the first 20 to 40 iterations. This is clearly less than one would expect from the theoretical analysis. In our simulations here, we consider a similar setting extended by a jammer. Our simulations show a very similar behavior by converging quickly close to the optimum.

We conduct simulations with randomly generated networks under SINR-interference. This way, we show that our proposed approach exhibits a good convergence towards the optimum. Not only the time until the no-regret-sequence stabilizes but also the distance to the optimum seem smaller than one might expect from our theoretical results. It is especially promising that the constant factors used in our proofs seem to be negligible in these simulations. This implies that no-regret learning might actually even be near optimal with a far better approximation factor than seen in our theoretical analysis.

We consider global and individual stochastic jammers. The no-regret learning algorithm considered is a variant of the Randomized Weighted Majority algorithm [LW94], which is also depicted in Algorithm 1. The algorithm uses transmission probabilities proportional to weights which are updated based on feedback of previous actions. Instead of using a fixed ρ as assumed in Section 2.2, we adapt the parameter ρ over time to account for the unknown total number of time steps. This way we basically follow a “guess and double” approach.

This adapted algorithm works as follows. The weights are initialized with 1 and multiplied by $(1 - \rho)^{l_a \cdot k}$ in every time step, where $l_0 = 0.5$ is the loss of not sending and the loss of sending being $l_1 = 0$ for a successful phase and $l_1 = 1$ for an unsuccessful phase. To increase the effect of learning we multiplied the loss by the length of the phases before updating the weights. These losses correspond to the utility function used in Section 2.4.2. The length of a phase is set to $k = \frac{6}{\delta}$ as given in Corollary 2.3. The factor ρ starts with $\sqrt{0.5}$ and is multiplied




Figure 2.1: Number of successful transmissions for no-regret learning over time against global and individual jammer.

by $\sqrt{0.5}$ every time the number of time steps is increased above the next power of 2. This reflects the typical “guess and double” approach as also considered by Blum [Blu98].

We conduct our simulations on networks constructed randomly, which consists of 200 links with receivers placed randomly on a 1000×1000 plane. Note that the placement is not restricted to the grid points of the plane. Senders are placed with a random angle and within a random distance between 0 and 100 near their respective receiver. The SINR parameters are $\alpha = 2.1$, $\beta_v = 1.1$ for all links, and $\nu = 4 \cdot 10^{-7}$. The transmission power of all senders is set to 2. Placing links on an Euclidean plane we used the SINR model to directly get a conflict graph with constant-independence as considered in our theoretical analysis.

This construction provides networks where interference from other links is the main reason for unsuccessful transmissions. Nevertheless, placing senders not too far away from the respective receiver still yields a sufficient possibility to achieve successful transmissions. To simulate links joining in different time steps, we let each link start its algorithm at a random time step during the first phase (i.e. during the first k time steps). This ensures that the algorithms do not coordinate quicker due to being already synchronized initially.

Figure 2.1 depicts the number of successful transmissions for one run of the algorithm against a global stochastic adversary and an individual stochastic adversary with $\delta = 0.8$. To simplify comparison we additionally plotted the size of the single-slot optimum without jammer $|OPT|$ and the expected average optimum $\delta \cdot |OPT|$ against the global jammer. For the case of individual jammers we plotted the average optimum $\sum_{v \in V} \frac{|\{t \in T | \ell_v \in OPT'_t\}|}{T}$ of the considered run. This is also the reference value of our theoretical proof of Corollary 2.3 before.

Our theoretical results imply that no-regret algorithms converge to a constant-factor approximation and this is also visible in simulations. Our simulations show that considering the case of a global jammer the algorithm approaches the actual single-slot optimum in un-jammed time steps depicted by dots in Figure 2.1. For time steps where the jammer prohibits any communication, the number of successful transmissions is 0. The dots on the x-axis rep-




Figure 2.2: Number of successful transmissions using no-regret learning over time for different values of δ . The optimal capacity is depicted by a dotted line.

resent no throughput in jammed time slots. The run is not averaged over time and, thus, lies above the expected optimum.

Besides some fluctuations in the beginning the algorithm stabilizes during the first 50 time steps and eventually even reaches the optimum in single time steps afterwards. So the algorithm achieves a reasonable throughput very early on. This result conforms to the simulation results presented by Ásgeirsson and Mitra [AM11] showing a quite similar behavior for no-regret learning without jammer. While Figure 2.1 depicts the run of our algorithm only by one example, Figure 2.2 averages over multiple runs and shows the same general behavior.

Figure 2.1 also shows that no-regret learning in the case of an individual jammer shows basically the same trend though it underlies higher fluctuation. This is due to the optimum changing in every time step. It is reasonable that these fluctuation will remain as they are introduced due to the jammer. Nevertheless, the algorithm shows a clear tendency towards stabilizing in spite of the mentioned fluctuations and approaching the average optimum. The algorithm also outperforms the average optimum in single time steps and clearly converges closely to the average optimum.

Both, for the global and the individual jammer it is clearly visible that the sequence of actions with diminishing regret converges closely to the optimum with an approximation factor closer to 1 than one would anticipate given the previous theoretical proofs. Thus, we consider these results to be very promising for the application of no-regret learning algorithms to real world networks.

The construction of the algorithm depends on δ by fixing the phase length accordingly and using the proposed utility functions. Knowing δ is not only useful in the construction, but as discussed before, it is crucial for the performance of the algorithms to know the correct parameter δ . Using a δ twice as high as the δ used by the jammer enables the jammer to trick the algorithm into believing that there is a lot of interference from other links. Thus, not transmitting can be the preferred action for all links allowing no performance at all.

We conducted simulations to investigate the actual effect of using the wrong δ in our algorithm. This is depicted in Figure 2.2. Here, we use δ to denote the parameter used by the algorithm and $\delta' = 0.35$ to denote the one actually used by the global jammer.

The simulations were executed over 10 networks, which are constructed in the same way as above and 1000 different random seeds for the jammer. The plots depict the average success in those runs. The average in a time step is taken only over runs where the channel was not jammed. This emphasizes the difference between the different δ . We iterate over $\delta \in \{0.2, 0.35, 0.6, 0.7, 0.9\}$. All other parameters stay as before. This way, our theoretical results again imply a constant-factor approximation when $\delta = \delta'$ is used by the algorithm.

Figure 2.2 shows that assuming a low δ makes the algorithm converge slower, although it still converges close to the optimum. This holds also due to the choice of the length of a phase k depending on δ . Phases correspond to the ridges visible in the plots, which makes it obvious that the links adapt their transmission behavior. The ridges are emphasized by averaging over different runs with the same phase length. Nevertheless a $\delta \leq \delta'$ still allows to approach the optimum, while more time is needed to reach a good approximation. Roughly speaking, the algorithm considers longer phases and overestimates the adversary still allowing to adjust to the real interference. Using $\delta = 0.2$ seems to show exactly the same adjustments as $\delta = \delta'$ stretched in the time dimension.

Surprisingly, by assuming $\delta > \delta'$ the performance does not seem to suffer severely. For $\delta = 0.6$ the algorithm still converges slowly to a reasonable approximation. As δ reaches $2\delta'$ this changes as expected. The algorithms tend to converge to not sending as the unsuccessful time steps due to the jammer overrule the interference. The adversary obviously tricks the algorithm into believing that there is much interference, and this way the algorithms of all links reduce their transmission probability. This results in a drop of performance as expected. In total, no-regret learning seems quite robust against using the wrong δ . In practical terms, our algorithmic approach is still applicable when only a rough guess of δ' is known and the dependence of the performance on using the exact δ is smaller than assumed.

2.6 Summary and Further Considerations

Throughout this chapter, we have considered a rather general proof template based on two intuitive properties of distributed approximation algorithms. We have then considered a setting in which transmissions can be rendered unsuccessful due to an adversary. This jammer is subject to restrictions on the time steps it can render unsuccessful. Depending on whether the parameters restricting the jammer are known to links, we derived different bounds on the performance of no-external-regret learning algorithms using suitable utility functions. We also considered a stochastic jammer.

We contrasted the theoretical results by simulations in Section 2.5 showing a fast convergence closely to the optimal capacity. In conclusion, no-regret learning can be used to successfully tackle capacity maximization with jamming in both theory and simulations. In simulations the constant factors inherent in our theoretical analysis appear negligible. The algorithms converge in reasonable time also being quite low compared to the time polynomial in the number of links one might expect from our analysis. Besides this, simulations imply that assuming wrong parameters δ different from the ones used by the jammer is not as bad as expected and that performance of no-regret learning remains robust in this case. It

nevertheless remains open whether these properties can also be observed in more intricate jamming settings discussed below.

While we considered a jammer being restricted in the number of time steps, in which it is allowed to render transmissions unsuccessful, other related work considers a jammer addressing the SINR setting directly. This is done by introducing a power-based jammer having a certain amount of power available to increase the ambient noise. Recently Ogierman et al. [ORS⁺13] considered such a specialized jammer. In contrast to our work for a general class of interference models, this work targets the SINR model rather specifically – the adversary has a budget of power to influence ambient noise. Besides this, the goal of transmitting is not to reach a specific receiver but any arbitrary one. That means it is not link-centered.

The algorithm proposed by Ogierman et al. to tackle the power-based jamming updates the transmission probability of senders in a very similar fashion as RWM by a kind of exponential backoff. Nevertheless, it uses additional information by being able to receive transmissions from other links.

Concerning our setting it is an interesting open question whether this kind of further feedback can be used to improve the algorithms based on no-regret learning. Especially, it seems to be a fruitful direction to consider power-based jamming with this additional feedback, as it is not clear how one could utilize learning approaches without this kind of setting here. In general, one can argue that power-based restricted jammers are more realistic than the ones we considered. So the question of how to utilize no-regret learning under these conditions is an interesting further research direction.

It is not trivial and not clear whether the algorithms can be utilized with the very restricted feedback considered here and whether provable guarantees against more intricate jammer settings like the power-based ones can be obtained. Obviously, our algorithmic approach cannot be applied straightforward, as we cannot account for jammers in the utility functions as easily as before.

For other scenarios, which are not restricted to the jammer scenario, we can nevertheless use the considered proof template and extend it. We will see that senders transmitting to multiple receivers and scenarios, where links can leave the network almost arbitrarily, can be handled using the same general idea in the following chapter.

CHAPTER 3

Extensions of the Proof Template for Jamming-Resistant Learning

The proof template considered in the previous chapter cannot only be applied to the jamming scenario considered there. We assumed adversarial jammers, which can render a certain fraction of time steps unsuccessful. This setting can be extended to more challenging settings using the same notions of (γ, ϵ) -successful and (η, ϵ) -blocking (see Definition 2.2 and 2.3). These extensions do not only broaden the jammer setting from the previous chapter but also extend scenarios considered in previous works.

In this chapter, we utilize the general ideas and the proof template of the previous section for two further settings. Firstly, we show that the same template of Theorem 2.1 can be applied to links consisting of one sender associated with *multiple receivers*. Secondly, we consider that links do not need to stay within the network once they joined. This was an assumption underlying the previous considerations and also previous work on regret learning for capacity maximization.

Both results considered in this chapter can be combined with the jammer scenario. In fact, the results can be applied in any scenario in which we have a bound on the (γ, ϵ) -successful and (η, ϵ) -blocking properties.

In the *multi-receiver setting*, we consider three different scenarios. The first one assumes that each sender strives to reach all of its receivers. Another scenario is considered where a transmission is successful whenever any of its receivers are reached successfully, and a third one assumes that the success or more precise the utility of a transmission attempt depends on the number of links receiving successfully.

3.1 Results

In total, we see that the proof template can be applied more generally than only in the considered jamming-resistant learning from Chapter 2. We show that with multiple receivers per sender we can reestablish the same parameters γ and η of the (γ, ϵ) -successful and (η, ϵ) -blocking properties up to a constant. This holds when a transmission attempt is considered successful if either all of the respective receivers receive the transmission successfully or if at least one receiver is conflict-free. This way, we prove the same approximation bounds given

in the previous chapter. When the utility depends on the number of receivers reached by a sender, we see that the properties do not hold and that no-external-regret learning does not suffice to get good approximation guarantees.

Similarly, we reestablish the proof template in a setting, where links can join and leave the network arbitrarily once they have low regret themselves. Here, we prove convergence to approximation factors in $O\left(\left(\log n + \log\left(\frac{1}{\eta}\right)\right) \frac{C}{\mu \cdot \gamma \cdot \eta}\right)$. This roughly increases the previously known bounds by a factor of $\log n$ over the case when the links are assumed to stay within the network.

In fact, the extension considered here is not only applicable to the setting with jammers. Recall the definitions of the properties from Section 2.3. We can use the results here for any setting that is (η, ϵ) -blocking and (γ, ϵ) -successful. Thus also the general setting of Dinitz [Din10] and Ásgeirsson and Mitra [AM11] is easily extended to multiple receivers and links leaving the network almost arbitrarily.

3.2 Multi-Receiver Setting

Extending the proof template considered for jamming-resistant learning from the previous chapter, we will start with a scenario where a sender is not only associated with one receiver but with multiple ones. Thus, each link does not consist of a sender-receiver pair but of a sender and a set of receivers.

This extends previous considerations to a more practical setting, as wireless networks are essentially multi-receiver networks by nature. Signals propagate through space in any direction. Thus, a sender can try to utilize this property of wireless communication to reach multiple receivers simultaneously with the same transmission. This kind of multicast setting, where a specific set of receivers should receive a transmission and an anycast setting, where an arbitrary of the possible receivers should be reached, will be analyzed in this section.

3.2.1 Model Description

When considering multiple receivers one also has to consider how to define a successful transmission, as now a transmission might be considered successful if all of the multiple receivers get the message or any of them. The utility of a transmission might even depend on the number of receivers receiving the transmission successfully.

More formally, we extend the previously considered network in the following way. The *network* now consists of n senders s_v , each coupled with a *set of receivers* $\{r_{v,1}, \dots, r_{v,k}\}$. Note that we do not make any assumption on how this set of receivers looks like or what size it has. Concerning the success of a transmission attempt we apply the following three kinds of settings.

Again we assume that a receiver can receive from its sender successfully if there is no conflict. That is, there is not too much interference from other senders, i.e., the summed up conflict-graph weights are low enough for the transmission to be successful. This results in the following settings to define the circumstances for a transmission attempt to be considered successful.

- In the *to-all setting*, a transmission for link ℓ_v is successful iff all of its receivers are conflict-free.
- The *to-one setting* considers a transmission for link ℓ_v successful, iff at least one of its receivers is conflict-free.
- The utility of a link in the *to-many setting* is linear in the number of receivers that are conflict-free.

For the first two settings, the to-all and the to-one setting, we assume that the no-external-regret-learning technique is applied as discussed before in Section 2.2. We assume that the utility functions basically stay the same. Only the notion of success is altered. For the last one we assume the utility to actually depend on the number of conflict-free receivers. It can, for example, depend linearly on the number of conflict-free receivers.

In real world wireless settings, it is a natural task to disseminate information to a set of multiple receivers. This task is often referred to as *multicast*. In wireless sensor networks one might even only want to distribute information to any arbitrary of these multiple receivers, which is also called an *anycast scenario*.

From a global point of view we can identify the task of capacity maximization with different global objectives for the network. In a to-all setting we maximize the number of transmissions of senders such that all receivers can receive the transmission successfully. In contrast to this, the to-one setting clearly strikes at the total number of senders reaching at least one of their receivers. This is sender-centered as it does not matter how many receivers actually receive the transmissions. The to-many setting is obviously more receiver-centered, as capacity maximization here can be identified with maximizing the number of receivers successfully receiving a transmission from their respective sender.

It is clearly visible that these three settings are inherently different. Especially the last one exhibits a different point of view. We will in fact see that no-external-regret learning applied in our simple approach is not suitable for optimization in the to-many setting.

We will show that having an algorithm that is (γ, ϵ) -successful and (η, ϵ) -blocking in the single receiver setting, we can easily transfer these properties to the to-one and to-all settings. Thus our proof template from Section 2.3 can be applied here.

We assume the same restricted feedback for our algorithms as before, i.e., links only get to know whether previous transmissions were successful or not.

Having considered only single-receiver settings until now, we use the same basic notion of C -independence here. In fact, we define C -independence for the multi-receiver setting by simply demanding C -independence for every single-receiver conflict graph under the considered interference model. That is, C -independence holds for each conflict graph constructed from every combination between senders and one of their receivers.

Definition 3.1. *A multi-receiver setting is C -independent if every conflict graph resulting from the combination of every sender with one of its receivers is C -independent.*

Though this seems to be a quite harsh condition, it is fulfilled for an interference model if this interference model provides C -independent conflict graphs for the single receiver settings in general. So for the SINR-based interference model with uniform powers we can again claim that we have $O(1)$ -independence as it can directly be transferred.

In the rest of this section, we apply no-external-regret learning algorithms from the previous section without any major alteration to the multi-receiver setting. To do so, we stick to

the same *utility functions* used before. We keep the general structure penalizing unsuccessful transmissions or phases and rewarding successful ones with the same values used before. For example, for the setting without jammer the utility is $+1$ as a reward, -1 as a penalty, and 0 if the link does not transmit. We only need to redefine under what circumstances a transmission attempt or a phase is considered successful or unsuccessful, respectively. Thus considering the to-all setting, a transmission attempt of sender s_v is successful if and only if all of its receivers $r_{v,i}$ are conflict free. In the to-one setting a transmission is defined to be successful iff at least one of its receivers is conflict free. For the to-many setting we clearly need to adjust the utility function to be linearly depending on the number of successful transmissions. The parameters μ and k used by the algorithm in Chapter 2 stay the same as before. Here, μ again denotes the fraction of time steps with a successful transmission in a phase needed for the algorithm to consider the phase successful and k denoting the length or number of time steps of every phase.

3.2.2 Approximation Guarantees

We will now consider the settings with multiple receivers described above, where each sender wants to reach one of its receivers or all of them. For those we will prove that the same approximation guarantees hold as before. To do so, we show that any algorithm guaranteeing the (γ, ϵ) -successful and (η, ϵ) -blocking properties in single-receiver settings also fulfills these properties for multiple receivers with the same parameters γ and η up to a constant.

Proposition 3.1. *An algorithm that computes a sequence of actions that is (γ, ϵ) -successful in a single-receiver setting also computes a sequence that is (γ, ϵ) -successful in the to-one and to-all multi-receiver settings.*

Proposition 3.1 follows straightforward from the definition of success. The utility functions stay the same, and our property only relies on the utility and average regret ϵ . As the definition of success in the phases is independent of this, we can directly conclude that the algorithm behaves in the same manner as before maintaining that the algorithm is (γ, ϵ) -successful.

Let q_v again denote the fraction of phases with transmissions and w_v the fraction of phases with successful transmissions. Thus, an algorithm having $\frac{1}{\gamma}(2w_v + \epsilon) \geq q_v$ before with the same utility function and the redefined success still fulfills this property.

To show the same for (η, ϵ) -blocking and transfer it to the multi-receiver setting, we have to argue a bit more in detail. Using the C-independence will enable us to construct weights for a conflict graph being C-independent in the single-receiver sense and this way ensure the (η, ϵ) -blocking property. These conflict-graph weights will be compatible with and depend on a specific run of the algorithm and sustain the implications about transmitting links.

Lemma 3.1. *There exists a C-independent conflict graph such that every algorithm that computes an (η, ϵ) -blocking sequence of actions in a single-receiver setting also computes an (η, ϵ) -blocking sequence in the to-one and to-all multi-receiver settings.*

Proof. We will construct conflict-graph weights based on a specific conflict graph in each time step. Then we will average over those weights for different time steps to prove that (η, ϵ) -blocking holds given these weights. The corresponding weights are denoted by $b_u^{(t)}(v)$

and averaging over all steps $t \in T$ yields $\bar{b}_u(v) = \frac{1}{T} \sum_{t \in T} b_u^{(t)}(v)$. In every time step the conflict graph is C-independent, so an average conflict graph with averaged weights is also C-independent.

Thus, in each of the settings we basically average over different single-receiver conflict graphs. These conflict graphs are chosen differently in the to-one and the to-all setting. They depend on which receivers are conflict free or which are not conflict-free, respectively.

At first, we consider the *to-one setting*. For the to-one setting the chosen weights $b_u^{(t)}(v)$ depend on the receivers receiving successfully in the optimal schedule. This choice is obviously independent of the time step t . Thus the average weight used in our proof is $\bar{b}_u(v) = b_u^{(t)}(v)$.

When considering a scenario, where the optimum differs in some time steps like in the jammer setting considered before, we can use $\bar{b}_u(v) = \frac{1}{T} \sum_t b_u^{(t)}(v)$.

For each link $\ell_v \in OPT$ or $\ell_v \in OPT_t$, respectively, to define $b_u^{(t)}(v)$ we consider an arbitrary receiver that is conflict-free in the optimum. For other links $\ell_v \notin OPT$ we arbitrarily choose any of its receivers. An unsuccessful transmission of the link is unsuccessful to all of its receivers. Thus it is also not conflict-free at the chosen receiver.

With this conflict-graph weight, having an unsuccessful transmission of link ℓ_v in time step t means that especially the transmission cannot be received successfully of the chosen receiver and thus $\sum_{u \text{ transmitting in } t} \bar{b}_u(v) \geq 1$. As the algorithm is (η, ϵ) -blocking in the single-receiver setting where exactly the chosen receiver is the single receiver, our choice of conflict-graph weights yields

$$\sum_t \sum_{u \text{ transmitting in } t} b_u^{(t)}(v) \geq \frac{1}{8} \eta \cdot T.$$

Thus, we directly obtain the (η, ϵ) -blocking property in the to-one setting.

For the *to-all setting* we choose weights not depending on the optimum as before but depending on the sequence of actions calculated by the algorithm. In fact, it depends on the receivers being in conflict during the run of the algorithm. Note that this construction only seems more restricted due to its dependence on the actual run of the algorithm. This does not matter, as this construction holds for all runs of the algorithm and as our proofs consider an already finished run of the algorithm. This construction is viable as long as we can still prove the necessary properties, which we will do in the following.

If the transmission of the sender of ℓ_v is not received by receiver $r_{v,i}$ in time step t , we use the pair $(s_v, r_{v,i})$ to construct $b_u^{(t)}(v)$ simply using the respective single-receiver conflict-graph weights. If multiple receivers do not receive the transmission, we choose an arbitrary of these receivers. For the time steps in which ℓ_v is successful we set $b_u^{(t)}(v)$ to the average $b_u^{(t')}(v)$ of all unsuccessful time steps t' . This way we ensure that the average $b_u^{(t)}(v)$ over all time steps is the same as the average over all unsuccessful time steps. For a sender that is always successful, we can easily choose an arbitrary of its receivers.

It is easy to see that C-independence also holds for this conflict graph. For any feasible set of links L we can construct L' as in Definition 1.3, and it holds $\sum_{v \in L'_t} \bar{b}_u(v) \leq \sum_{v \in L'_t} \max_{r_{v,i}} b_u^{r_{v,i}}(v)$ for all $u \in V$. Here $b_u^{r_{v,i}}(v)$ denotes the conflict-graph weight of sender s_u to the i -th receiver of link ℓ_v . Thus, it refers to a specific single-receiver conflict graph, where the receiver of link ℓ_v is $r_{v,i}$. As C-independence is fulfilled for every receiver of ℓ_v , it also holds for the receivers with $\max_{r_{v,i}} b_u^{r_{v,i}}(v)$. This directly shows $\sum_{v \in L'_t} \max_{r_{v,i}} b_u^{r_{v,i}}(v) \leq C$

for all $u \in V$ in the constructed conflict graph. Averaging over different time steps still guarantees C -independence.

We will now show that this construction for the to-all setting is still (η, ϵ) -blocking using the (η, ϵ) -blocking property from the single-receiver setting. We know that in a time step which is unsuccessful due to other links, we have $\sum_u b_u^{(t)}(v) > 1$ by definition. Thus, by (η, ϵ) -blocking, $f_v \geq \frac{1}{4}\eta$ for small q_v and summing over unsuccessful time steps yields

$$\sum_{\substack{t \in T \\ \ell_v \text{ unsuccessful}}} \sum_{u \text{ transmits in } t} b_u^{(t)}(v) \geq \frac{1}{8}\eta T .$$

Having $\bar{b}_u(v)$ defined to be the average of $b_u^{(t)}(v)$ in steps unsuccessful for ℓ_v we get

$$\sum_{\substack{t \in T \\ \ell_v \text{ unsuccessful}}} \sum_{u \text{ transmits in } t} \bar{b}_u(v) \geq \frac{1}{8}\eta T .$$

Easy calculus results in $\sum_u \bar{b}_u(v) \cdot |\{t \in T \mid u \text{ transmits in } t\}| \geq \frac{1}{8}\eta T$ proving (η, ϵ) -blocking. This concludes the proof of the lemma. $\square$

We have shown that the (η, ϵ) -blocking and (γ, ϵ) -successful properties also hold in the to-one and to-all settings by Proposition 3.1 and Lemma 3.1. This way, it is straightforward that our proofs from the previous sections can be directly applied using the constructed weights $\bar{b}_u(v)$.

Corollary 3.1. *Theorems 2.2 to 3.1 and Corollaries 2.1 and 2.3 hold in the respective multi-receiver to-one and to-all settings.*

Besides applying the results to the jamming scenario considered in Chapter 2, we can also apply the multi-receiver to-one and to-all settings to the result from Ásgeirsson and Mitra [AM11] in Corollary 2.2. This proves the following corollary.

Corollary 3.2. *No-external-regret learning in the multi-receiver to-one and to-all settings converges to a constant-factor approximation for capacity maximization in the SINR interference model with uniform powers.*

While the constructions in the to-one and the to-all settings were quite straightforward, the to-many setting, where utility depends on the number of conflict-free receivers, does not yield such a construction. Thus, we cannot transfer the results that easily. To be precise, we show a negative result for this setting. There exists an instance and a no-regret sequence with an approximation factor linear in the maximum number of receivers per sender. This problem arises as a sender does not get feedback on how many receivers of other senders are blocked by its transmission attempts.

Lemma 3.2. *In the to-many setting there exists an instance such that every sequence of action vectors with 0 regret yields an approximation factor linear in the maximum number of receivers per sender.*

Proof. The network for the lower bound consists of two links – one with sender s_1 and receivers $r_{1,1}$ to $r_{1,k}$ and one with sender s_2 and receiver $r_{2,1}$. The conflicts are constructed such that receivers $r_{1,i}$ can only be conflict-free if s_2 decides not to transmit. The second link can

always transmit successfully. This can be implemented in the SINR model by simply putting all $r_{1,j}$ close together and closer to s_2 together with $r_{2,1}$ than to s_1 .

In every no-regret sequence, s_2 is transmitting almost all the time, and for 0 regret s_2 is always transmitting. For the average regret being close to 0 the other link s_1 almost never transmits, as it can not have a successful transmission when s_2 transmits. This implies a total objective function value of 1. It is clear that the optimal schedule would choose s_1 for transmission and s_2 for not transmitting at all. Thus, s_1 reaches k receivers in the optimum. Note that s_2 would only reach 1 if it transmits. The optimal capacity being k directly proves the claim. $\square$

In total, we can extend the previously considered setting using the proof template in Chapter 2 to multiple receivers. This can be done in a straightforward way when all receivers or when at least one of the receivers should be able to successfully receive the transmission. If the utility depends on the number of conflict-free receivers, the same cannot be shown and the regret-learning approach is shown to fail.

3.3 *Joining and Leaving Links*

Besides considering the extension to multi-receiver settings in the section above, we will now consider a setting where links are allowed to *join and leave* the network almost arbitrarily.

While the approach considered in Chapter 2 and Section 3.2 above in general does not assume links to join the network simultaneously, it assumes links to stay in the network. That is, a link will not stop behaving according to the algorithm until the end of time once it started to transmit.

Now we assume that links only have to stay in the network until their own regret is low enough. For this setting, we will construct a similar proof template as in Chapter 2. This result can be combined with the previous ones concerning the jammer scenario, the multiple receivers, and previous work by Ásgeirsson and Mitra. Thus, we extend the previous works and the previously known bounds while allowing links to leave the network.

3.3.1 *Model Description*

Obviously, it is not a realistic assumption that links stay within the network indefinitely. So we relax this assumption made in the previous sections here and consider links being able to leave the network earlier. Besides this we also allow links to join and leave the network almost arbitrarily. However, they are assumed to stay until their *own* regret is low. For this we prove convergence to an $O\left(C \cdot \frac{\log(n)}{\mu \cdot \gamma \cdot \eta}\right)$ -approximation for an (γ, ϵ) -successful and (η, ϵ) -blocking algorithm as defined in Section 2.3.

This way, we can apply all results observed in the previous chapter and the previous section to the setting with joining and leaving links by basically losing a factor of $\log n$. This also includes the standard setting from Ásgeirsson and Mitra as discussed before.

More formally, each link ℓ_v comes with an interval of phases $\mathcal{R}_v$. During these phases the link is present in the network being able to transmit and observe the outcome from its actions. We assume that outside of its time interval a link cannot observe, learn, or transmit. The average regret of an algorithm is now considered with respect to the phases in $\mathcal{R}_v$. In the

following section we will adjust the theorem and proof as seen in Section 2.3 to this setting to get a very similar approximation guarantee.

3.3.2 Approximation Guarantees

During this section, we will reestablish the general result from Section 2.3 used as a template in Sections 2.4 to 3.2 losing only a factor of $O(\log n)$. To do so, we will prove the following theorem. This proof needs some more advanced considerations compared to the one of Theorem 2.1 to be applicable to joining and leaving links.

Theorem 3.1. *Consider an action sequence, which is (η, ϵ_v) -blocking, (γ, ϵ_v) -successful with average regrets ϵ_v for links ℓ_v , and has at least μ successful time steps in each phase considered successful in a C -independent conflict graph. For this sequence of actions, the number of successful transmissions yields*

$$\sum_{v \in V} \sum_t w_v^{(t)} \in \Omega \left(\frac{\mu \gamma \eta}{\left(\log n + \log \left(\frac{1}{\eta} \right) \right) C} \sum_t |OPT_t| - \mu \cdot \epsilon \cdot T \right),$$

where OPT_t denotes the optimal set of links transmitting in time step t or the single-slot optimum and $\epsilon = \sum_{v \in V} \epsilon_v$.

To prove this, we again use a primal-dual approach as in Section 2.3 for the proof of Theorem 2.1. To adjust to the joining and leaving links, we use a different linear program formulation. This LP is more complex and introduces additional factors in the approximation guarantee.

In terms of approximation guarantees, if the regret is low enough, ϵ converges to 0 and this theorem implies a factor in

$$O \left(C \cdot \frac{\log n + \log \left(\frac{1}{\eta} \right)}{\mu \gamma \eta} \right).$$

Proof. The general proof concept follows the same thread as the one of Theorem 2.1 using a linear-programming approach. Here, we cannot simply consider an average over all time steps when constructing the LPs, as different links might be in the system in different time intervals. As a primal LP we use the following one, where OPT'_t is defined as L' out of Definition 1.3 with $L = OPT_t$ being the optimal schedule in time step t . Recall, that $b_u(v)$ are the conflict-graph weights modelling the interference. In the remainder of this proof, we will for simplicity denote the set of time steps by T .

$$\begin{aligned} \max. \quad & \sum_{v \in V} \sum_{\substack{t \in T \\ v \in OPT'_t}} x_{v,t} \\ \text{s.t.} \quad & \sum_{v \in OPT'_t} b_u(v) x_{v,t} \leq C \quad \forall u \in V, t \in T \\ & x_{v,t} \leq 1 \quad \forall t \in T, v \in OPT'_t \\ & x_{v,t} \geq 0 \quad \forall t \in T, v \in OPT'_t \end{aligned}$$

The linear program only differs from the one used before by considering the summed up number of transmissions. This way, the constraints need to be considered for every time step, and the objective function basically sums over all time steps in which link ℓ_v is in the optimum.

We consider a variable $x_{v,t}$ only to be existing for time steps in which $v \in OPT'_t$, and we set $x_{v,t} = 1$ if $v \in OPT'_t$. This is a feasible solution as C -independence holds in every time step fulfilling the first constraint. The other constraints are also fulfilled.

The dual LP is

$$\begin{aligned} \min. \quad & \sum_{v \in V} \sum_{t \in T} C \cdot y_{v,t} + \sum_{t \in T} \sum_{v \in OPT'_t} z_{v,t} \\ \text{s.t.} \quad & \sum_{u \in V} b_u(v) y_{u,t} + z_{v,t} \geq 1 \quad \forall v \in V, t \in T \text{ with } v \in OPT'_t \\ & y_{v,t}, z_{v,t} \geq 0 \quad \forall v \in V, t \in T \text{ with } v \in OPT'_t. \end{aligned}$$

To construct a dual solution we need more detailed considerations. We assume the phases of any link ℓ_v to be numbered such that $R_v^{(i)}$ denotes the i -th phase of link ℓ_v . Let $i_v(t)$ denote the number of the phase containing time step t , formally, $t \in R_v^{(i_v(t))}$. Considering the maximum number of transmissions of multiple consecutive learning intervals we can construct a feasible solution for the dual LP in a similar way as before in the proof of Theorem 2.1. Let $Q_v^R = 1$ if the link ℓ_v transmits in phase $R \in \mathcal{R}_v$ and $Q_v^R = 0$ otherwise. Thus, we define

$$\mathcal{J}_v = \{\lceil \log |\mathcal{R}_v| \rceil - \log 16n + \log \eta, \dots, \lceil \log |\mathcal{R}_v| \rceil + \log 16n\}$$

to construct the solution for the dual LP by setting

$$\begin{aligned} z_{v,t} &= 8 \cdot \frac{1}{\eta} \cdot \left(\max_{j \in \mathcal{J}_v} \frac{1}{2^j} \sum_{i'=i_v(t)-2^j}^{i_v(t)+2^j} Q_v^{R_v^{(i')}} + \frac{\epsilon_v}{2} \right) \\ y_{v,t} &= 32 \cdot \frac{1}{\eta} \cdot \max_{j \in \mathcal{J}_v} \frac{1}{2^j} \sum_{i'=i_v(t)-2^j}^{i_v(t)+2^j} Q_v^{R_v^{(i')}}. \end{aligned}$$

For the case of $z_{v,t} \geq 1$ the solution is clearly feasible. To show that the solution is feasible, we only have to analyze the case of $z_{v,t} \leq 1$. We now consider some $v \in V$ and t such that $v \in OPT'_t$. We know that if $z_{v,t} \leq 1$ for all $j \in \mathcal{J}_v$ we have

$$\frac{1}{2^j} \sum_{i'=i_v(t)-2^j}^{i_v(t)+2^j} Q_v^{R_v^{(i')}} \leq \frac{1}{8} \eta - \frac{\epsilon_v}{2}.$$

For $j = \lceil \log |\mathcal{R}_v| \rceil$ this yields

$$\frac{1}{2 \cdot |\mathcal{R}_v|} \sum_{i'=i_v(t)-|\mathcal{R}_v|}^{i_v(t)+|\mathcal{R}_v|} Q_v^{R_v^{(i')}} \leq \frac{1}{8} \eta - \frac{\epsilon_v}{2}.$$

By easy calculus we get

$$\frac{1}{|\mathcal{R}_v|} \sum_{i'=i_v(t)-|\mathcal{R}_v|}^{i_v(t)+|\mathcal{R}_v|} Q_v^{R_v^{(i')}} \leq \frac{1}{4} \eta - \epsilon_v.$$

The left side of the inequality is at least the fraction of phases in which ℓ_v chooses to transmit as the sum goes at least over all phases in $\mathcal{R}_v$. Thus, we can bound the fraction of all phases in which ℓ_v transmits by $\frac{1}{4}\eta - \epsilon_v$.

In the following, let $R_u \cap \mathcal{R}_v$ denote the set of time steps that phase R_u of link ℓ_u and the phases of ℓ_v share. In the same way $t \in \mathcal{R}_v$ is a time step in any phase of ℓ_v . Using (η, ϵ) -blocking we conclude that the summed up weights $b_u(v)$ of other links transmitting in phases of ℓ_v is at least $\frac{1}{8}\eta$. Formally, we can express this conclusion by

$$\frac{1}{|\mathcal{R}_v|} \sum_u \sum_{\substack{R_u \in \mathcal{R}_u \\ R_u \cap \mathcal{R}_v \neq \emptyset}} b_u(v) Q_u^{R_u} \geq \frac{1}{8}\eta .$$

Reordering the sums yields

$$\sum_u b_u(v) \frac{1}{|\mathcal{R}_v|} \sum_{\substack{R_u \in \mathcal{R}_u \\ R_u \cap \mathcal{R}_v \neq \emptyset}} Q_u^{R_u} \geq \frac{1}{8}\eta .$$

We will prove that it follows that the $y_{u,t}$ are large enough to fulfill the constraint. To do so, we now drop all links ℓ_u with which ℓ_v shares only few phases from the consideration. We will show that these phases can constitute only a minor part of the interference. The part considered to be dropped is

$$\sum_{\substack{u \in V \\ |\mathcal{R}_v \cap \mathcal{R}_u| < \frac{1}{16n}\eta |\mathcal{R}_v|}} b_u(v) \frac{1}{|\mathcal{R}_v|} \sum_{\substack{R_u \in \mathcal{R}_u \\ R_u \cap \mathcal{R}_v \neq \emptyset}} Q_u^{R_u} .$$

This can be bounded using $Q_u^{R_u} \in [0, 1]$ by

$$\sum_{\substack{u \in V \\ |\mathcal{R}_v \cap \mathcal{R}_u| < \frac{1}{16n}\eta |\mathcal{R}_v|}} b_u(v) \frac{1}{|\mathcal{R}_v|} |\mathcal{R}_v \cap \mathcal{R}_u| \leq n \cdot \frac{1}{16n}\eta = \frac{1}{16}\eta .$$

Therefore, we have

$$\sum_{\substack{u \in V \\ |\mathcal{R}_v \cap \mathcal{R}_u| \geq \frac{1}{16n}\eta |\mathcal{R}_v|}} b_u(v) \frac{1}{|\mathcal{R}_v|} \sum_{\substack{R_u \in \mathcal{R}_u \\ R_u \cap \mathcal{R}_v \neq \emptyset}} Q_u^{R_u^{(i')}} \geq \frac{1}{16}\eta .$$

We can assume $|\mathcal{R}_v| \geq \frac{1}{16n} \max_{u \in V} |\mathcal{R}_u|$ for all links $\ell_v \in OPT_t'$ for any $t \in T$ without loss of generality, as otherwise it can only make up for an $(\frac{1}{16n})$ -th fraction of the average optimum.

This yields as a bound on the experienced interference of ℓ_v

$$\sum_{\substack{u \in V \\ \frac{1}{16n}\eta |\mathcal{R}_v| \leq |\mathcal{R}_u| \leq 16n |\mathcal{R}_v|}} b_u(v) \frac{1}{|\mathcal{R}_v|} \sum_{\substack{R_u \in \mathcal{R}_u \\ R_u \cap \mathcal{R}_v \neq \emptyset}} Q_u^{R_u^{(i')}} \geq \frac{1}{16}\eta .$$

Let j be such that $2^{j-1} \leq |\mathcal{R}_v| \leq 2^j$. Then we have for all $u \in V$ with $\frac{1}{16n}\eta |\mathcal{R}_v| \leq |\mathcal{R}_u| \leq$

$16n|\mathcal{R}_v|$ and for $t \in \mathcal{R}_v$ a bound on $y_{u,t}$ being

$$\begin{aligned} y_{u,t} &\geq 32 \frac{1}{\eta} \cdot \frac{1}{2^j} \sum_{i'=i_u(t)-2^j}^{i_u(t)+2^j} Q_u^{R_u^{(i')}} \geq 32 \frac{1}{\eta} \cdot \frac{1}{2^j} \sum_{\substack{R_u \in \mathcal{R}_u \\ R_u \cap \mathcal{R}_v \neq \emptyset}} Q_u^{R_u^{(i')}} \\ &\geq 32 \frac{1}{\eta} \cdot \frac{1}{2 \cdot |\mathcal{R}_v|} \sum_{\substack{R_u \in \mathcal{R}_u \\ R_u \cap \mathcal{R}_v \neq \emptyset}} Q_u^{R_u^{(i')}}. \end{aligned}$$

We combine this with the experienced interference of link ℓ_v deduced above, that is

$$\begin{aligned} \sum_{u \in V} b_u(v) y_{u,t} &\geq \sum_{u \in V} b_u(v) 32 \cdot \frac{1}{2 \cdot \eta \cdot |\mathcal{R}_v|} \sum_{\substack{I_u \in \mathcal{R}_u \\ R_u \cap \mathcal{R}_v \neq \emptyset}} Q_u^{R_u^{(i')}} \\ &\geq \frac{16}{\eta} \cdot \sum_{u \in V} b_u(v) \frac{1}{|\mathcal{R}_v|} \sum_{\substack{I_u \in \mathcal{R}_u \\ R_u \cap \mathcal{R}_v \neq \emptyset}} Q_u^{R_u^{(i')}} \\ &\geq \frac{16}{\eta} \cdot \frac{1}{16} \eta \geq 1. \end{aligned}$$

This way, $\sum_{u \in V} b_u(v) y_{u,t} \geq 1$ fulfills the constraint. Thus the solution is feasible. It is left to show that the objective function of the dual LP compared to that of the primal LP results in the claimed approximation factor.

For the objective function of the dual LP we have

$$\begin{aligned} \sum_{v \in V} \sum_t (C y_{v,t} + z_{v,t}) &\leq 40 \cdot \frac{C}{\eta} \sum_{v \in V} \sum_t \left(\max_{j \in \mathcal{J}_v} \frac{1}{2^j} \sum_{i'=i_v(t)-2^j}^{i_v(t)+2^j} Q_v^{R_v^{(i')}} + \epsilon_v \right) \\ &\leq 40 \cdot \frac{C}{\eta} \sum_{v \in V} \left(\sum_{j=\lceil \log |\mathcal{R}_v| \rceil - \log 16n + \log \eta}^{\lceil \log |\mathcal{R}_v| \rceil + \log 16n} \frac{1}{2^j} \sum_t \sum_{i'=i_v(t)-2^j}^{i_v(t)+2^j} Q_v^{R_v^{(i')}} + \epsilon_v T \right) \\ &= 40 \cdot \frac{C}{\eta} \left(\sum_{v \in V} \sum_{j=\lceil \log |\mathcal{R}_v| \rceil - \log 16n + \log \eta}^{\lceil \log |\mathcal{R}_v| \rceil + \log 16n} \frac{1}{2^j} \sum_t (2 \cdot 2^j + 1) Q_v^{R_v^{(i(t))}} + \epsilon T \right) \\ &\leq 40 \cdot \frac{C}{\eta} \left(\sum_{v \in V} (2 \cdot \log(16n) - \log \eta) \cdot 3 \cdot \sum_t Q_v^{R_v^{(i(t))}} + \epsilon T \right). \end{aligned}$$

The last inequality holds due to $\frac{2 \cdot 2^j + 1}{2^j} \leq 3$. This directly yields

$$\begin{aligned} \sum_{v \in V} \sum_t (C y_{v,t} + z_{v,t}) &\leq O \left((\log n - \log \eta) \cdot \frac{C}{\eta} \left(\sum_{v \in V} \sum_t Q_v^{R_v^{(i(t))}} + \epsilon T \right) \right) \\ &= O \left((\log n - \log \eta) \cdot \frac{C}{\eta} \left(\sum_{v \in V} \sum_{R_v \in \mathcal{R}_v} Q_v^{R_v} \cdot k + \epsilon T \right) \right). \end{aligned}$$

Using that the action sequence is (γ, ϵ) -successful, we get an upper bound on the objective value depending on the number of successful transmissions yielding

$$O \left(\left(\log n + \log \frac{1}{\eta} \right) \frac{C}{\eta} \left(\sum_{v \in V} \sum_t \frac{2}{\gamma \mu} \cdot w_v^{(t)} + \frac{1}{\gamma} \epsilon T \right) \right).$$

The comparison of this to the objective function of the primal LP yields

$$\begin{aligned} \sum_{v \in V} \sum_{\substack{t \in \mathcal{R}_v \\ v \in OPT'_t}} x_{v,t} &= \sum_{t \in T} |OPT'_t| \\ &\leq O\left(\left(\log n + \log\left(\frac{1}{\eta}\right)\right) \frac{C}{\mu \cdot \gamma \cdot \eta} \cdot \left(\sum_{v \in V} \sum_t w_v^{(t)} + \mu \epsilon T\right)\right), \end{aligned}$$

which in total shows the claim concluding the proof. $\square$

Using this theorem we are able to transfer the general ideas seen in Chapter 2 and the previous section to the setting, where we allow links to join and to leave the network again. In fact, all approximation guarantees for all settings analyzed in those sections can be transferred to this case. It also allows to introduce joining and leaving links to previous results like the one from Ásgeirsson and Mitra [AM11]. It is easy to see that their approach can be incorporated into our notions using $\delta = 1$ and $T' = 1$. This way extending their setting we have no-regret learning yielding an $O(\log n)$ -approximation. In general the guarantees increase by a factor of $O\left(\log n + \log \frac{1}{\eta}\right)$.

In total, we have seen throughout this chapter that the proof template considered in Chapter 2 can be further generalized. We are able to do this for both a setting with multiple receivers per sender and a setting where links are allowed to leave the network. Besides this, these results can be combined. Also, the approximation guarantees of the jammer setting can be transferred accordingly losing a factor of $O(\log n)$ here.

CHAPTER 4

Scheduling with Multiple Channels and Availabilities

In this chapter, we will further generalize the considered applications of no-regret learning algorithms for capacity maximization. We will focus on scenarios with multiple channels, where links do not only have to decide whether to transmit but also which channel to use. These channels are assumed to be available or unavailable to the links based on a stochastic process.

Recently, a variety of important regulatory and, in particular, algorithmic challenges in wireless networks were posed [FNK⁺11, Nek09]. These target spectrum licensed to a *primary user* for exclusive usage and how to utilize this for secondary users. Primary users legally have the right to exclusively use a specific part of the wireless spectrum. A general idea to better utilize scarce spectrum is to temporarily open up the spectrum bands of primary users in local areas where it is unused. This creates spectrum opportunities for *secondary users*, which can then use these available channels, allowing a more efficient usage of spectrum. The idea has been introduced by different regulatory authorities like OfCom and FCC.

A prominent approach currently discussed in industry is based on a database recording which channels are currently available for secondary usage in which spatial areas [FNK⁺11, Nek09]. Primary users announce via this database whether the channel is available to be used by secondary users. Secondary users obtain information about the channels available to them by querying the database. Then they decide independently about channel access. In real world networks a link might query this database via some kind of primary communication mean, for example the mobile-phone network. Then links can coordinate among themselves based on the gained information about channel availabilities.

We consider the underlying algorithmic problem of capacity maximization in this kind of setting. We will apply a notion of regret different to the one considered in previous chapters. This notion is known from the game-theoretic setting of *sleeping experts*, where actions can be unavailable [Blu97, FSSW97].

Obviously, the availabilities of channels will be correlated among links to a certain extend. This can for example be due to the locality of links. Multiple links can be in the same spatial region, where a specific primary user has the exclusive right to use certain spectrum. To consider this correlation, we introduce the notion of *group-correlated availabilities* and analyze the algorithmic approach based on this notion.

4.1 Results

Our notion of regret slightly differs from similar notions considered in related work [KMB09, KNMS10, KS12]. Thus, we will start with our first result showing how to utilize existing external-regret algorithms to yield *no-availability-ordering regret* tuned to the availability setting. We will exemplarily give bounds in $o(T)$ on the regret obtained by the algorithm. Besides this, we will also consider the partial information model, where links only get to know the outcome of actions actually chosen. This way, the algorithms are applicable in a wider range of practical settings.

Using these algorithms and the given regret bounds, we will prove approximation guarantees depending on a newly introduced measure capturing the correlation of the availabilities among links. Additionally, we give similar bounds if the channel availabilities are uncorrelated. In this case, we can even utilize slightly different notions of regret already considered in [KMB09] and thus we can apply a wider range of algorithms.

These results are accompanied by lower bounds in the last section of this chapter. We prove that for arbitrary correlation among the availabilities of links, we can only guarantee an $\Omega(n)$ -approximation. This is a trivial bound and can be reached by letting just one link in the network transmit. This underlines the relevance of restricting the correlation of channel availabilities among links, as considered in the following section.

4.2 Availability and Multi-Channel Model

While we stick to the same general network model with a set of n links V as before, we will extend this model to a multi-channel environment. Instead of assuming links to coordinate the usage of spectrum on one channel, we assume that k different channels are available. Interference influences communications per channel and does not influence the interference on other channels. To model capacity maximization in this setting, we assume that all links can in general choose an arbitrary of all available channels. Thus, a link is now not only forced to decide whether to transmit or not, but in case of transmission it also has to decide which channel to use.

Besides this, we assume that channels are not available throughout all time slots. This allows us to model primary users occupying the channel and having the exclusive right to use the channel. Whenever the primary user does not use the channel, it can allow secondary users to use the spectrum. Thus, channels become available to the links in some time steps.

We denote the set of channels by $\mathcal{K}$ and the number of channels by k . In each step, the availability of a channel $\kappa \in \mathcal{K}$ to a link ℓ_v is the result of a random trial. Algorithms like *Sleeping-Follow-the-Perturbed-Leader (SFPL)* from Kanade et al. [KMB09] yield a notion of no-regret property sufficient to show convergence to a constant-factor approximation if the availabilities of channels are independent among links. We will consider this notion of regret in Section 4.4.2. We will extend this to a setting, where the random availabilities can be correlated to some extent as described below.

We assume that a primary user influences a fixed set of links, which can be defined based on locality in the network. Based on this, the channel is available or unavailable to links. The availabilities can this way be correlated depending on how they are influenced by primary users. Formally, we model this assumption as described below.

For each channel κ let G_κ be a set of sets of links $G_{\kappa,1}, \dots, G_{\kappa,m} \subseteq V$ called *groups*. We assume the distributions for the random availabilities are independent among the groups. However, among the channels the availabilities can be arbitrarily correlated in our model. On a channel κ for each group $G_{\kappa,i}$ the group is blocked on this channel, e.g., a primary user occupies the channel with a probability $1 - p_{G_{\kappa,i}}$. A channel is available to link ℓ_v , such that ℓ_v can transmit on the channel if it is not occupied for all groups G with $\ell_v \in G$. Thus, the probability that the channel is available to ℓ_v is $p_{v,\kappa} = \prod_{G \in G_\kappa, \ell_v \in G} p_G$. This way, availabilities are correlated for different links using those groups. We call this setting a *group-correlated setting*.

To express how availabilities among links are correlated, let $p_{u,v,\kappa} = \prod_{G \in G_\kappa, \ell_v, \ell_u \in G} p_G$ denote the probability that for links ℓ_v and ℓ_u the groups both links share are available. The set of time steps in which channel κ is available to link ℓ_v is denoted by $I_{v,\kappa}$ and the set of time steps in which groups shared by ℓ_v and ℓ_u are not occupied is analogously denoted by the set $I_{u,v,\kappa}$.

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 $P_{v,M}^{(t)} \in \{0, 1\}$ 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 its set of available channels. That is, it gets to know the outcome of a random trial given the stochastic process described above. 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.

Our approximation bounds depend on the factor C for C -independent conflict graphs as already considered in Section 1.1.2. Note that, in particular, conflict graphs resulting from the SINR model under uniform power have constant C [AM11]. This then guarantees a constant-factor approximation for uncorrelated availabilities as shown in Section 4.4.2.

To account for the correlation among availabilities of different links, we introduce the notion of (C, D) -correlation-independence extending the C -independence assumed before. We consider a partition of links into sets each being C -independent with scaled weights. This serves as a measure for the correlation. Note that this measure is actually directly tuned to our proof approach. Thus, it is unclear how it could server as a general metric to measure the correlation.

Throughout this chapter, we use $OPT_\kappa^{(t)}$ to denote the set of links using channel κ in time step t in the optimal solution. As this forms a feasible set being able to transmit simultaneously, we are able to apply the C -independence definition to this set.

Definition 4.1. A correlation setting is called (C, D) -correlation-independent if for every link ℓ_u and every channel $\kappa \in \mathcal{K}$ there exist a set $\mathcal{D}_{u,\kappa}$ of at most D sets of links $E_{u,\kappa,1}, \dots, E_{u,\kappa,D} \subseteq V$ with $E_{u,\kappa,1} \cup E_{u,\kappa,2} \cup \dots \cup E_{u,\kappa,D} = V$ such that for each feasible set of links within a set $E_{u,\kappa,i}$ the following variant of the C -independence condition is fulfilled. That is, for each link ℓ_u and each set $E \in \mathcal{D}_{u,\kappa}$ it holds

$$\sum_{\ell_v \in E} b_u(v) \cdot \frac{|I_{v,\kappa}|}{|I_{u,v,\kappa}|} x_{v,\kappa} \leq C ,$$

where $x_{v,\kappa}$ is the fraction of time steps in which ℓ_v is in L' as defined in the C -independence definition (Definition 1.3) with $L = \text{OPT}_\kappa^{(t)} \cap E$.

One could obviously also consider this definition in expectation. This changes the condition to $\sum_{\ell_v \in E} b_u(v) \cdot \frac{p_{v,\kappa}}{p_{u,v,\kappa}} x_{v,\kappa} \leq C$. Using this notion can be handy in some settings to prove approximation guarantees. In our setting, we stick to the definition using the actual values.

For different settings the availabilities can be correlated differently among links depending on the way links are grouped. In Section 4.4.3 we will even prove a lower bound in $\Omega(n)$ on the approximation factor achieved by an action sequence with zero regret for group correlation. Thus, harsh correlation yields that the no-availability-ordering regret defined in the following section does not suffice to guarantee any non-trivial approximation factor. For settings with further restrictions on the group correlation, we will prove better approximation factors.

We will consider the following restriction on the group correlation. We structure the groups and links suitably to restrict the correlation of availabilities. Our proofs later hold for general (C, D) -correlation-independence. The restriction makes the parameter D more manageable and more obviously depending on the considered structure.

Definition 4.2. For every channel $\kappa \in \mathcal{K}$ let T_κ be a rooted tree such that there exists a bijection f between the groups in G_κ and the vertices of T_κ and a surjection g from links in V to vertices of the tree T_κ .

A group-correlated setting is called *laminar* if for every link ℓ_v trees as described above exist, such that $\ell_v \in G$ if and only if the vertex $f(G)$ is on the path from $g(\ell_v)$ to the root (including both).

The maximal height of these trees is called the *height of the laminar group correlation*.

Note that the setting, where the availabilities of any channel κ are *stochastically independent* among the links, is a special case of a laminar setting with trees of height 2. Here, the root is associated with a group being always unoccupied and thus available for the links to transmit with probability 1. The tree consists of n leaves and the function g is constructed to be a bijection from links to the leaves. The root of the tree is not associated with any link. This constitutes a laminar correlation setting. The groups associated with the leaves are independent and so are the availabilities among the links.

Considering laminar correlation-settings, we will prove that these are (C, h) -correlation-independent, where h is the height of the laminar group correlation. This yields an approximation factor depending on this height h as we will see in Section 4.4.1.

Lemma 4.1. Consider a network of n links with C -independent conflict graph on k channels and laminar group-correlated channel availabilities with height h . This setting is (C, h) -correlation-independent.

Proof. To show (C, h) -correlation-independence, we have to show that for every link ℓ_u and channel κ we can partition the set of links into h sets in $\mathcal{D}_{u,\kappa}$, such that every set exhibits a property similar to C -independence being

$$\sum_{\ell_v \in E} b_u(v) \cdot \frac{|I_{v,\kappa}|}{|I_{u,v,\kappa}|} x_{v,\kappa} \leq C ,$$

where $x_{v,\kappa}$ is the fraction of time steps in which ℓ_v is in L' from the C -independence definition with $L = \text{OPT}_\kappa^{(t)} \cap E$ for a set of links $E \in \mathcal{D}_{u,\kappa}$. For a link ℓ_u we can construct this partition easily based on the tree structure.

Fix the tree for any channel κ and link ℓ_u . For every vertex z on the path between $g(\ell_u)$ and the root of the tree, let T_z be the subtree rooted at z and T'_z the subtree rooted at the child of z with $g(\ell_u) \in T'_z$. Here, $g(\ell_u) \in T'_z$ if the vertex $g(\ell_u)$ is a vertex of the tree T'_z . We define $E_z = \{\ell_v \in V | g(\ell_v) \in T_z, g(\ell_v) \notin T'_z\}$ as links with the vertices whose path from the root branches from the one of $g(\ell_u)$ at vertex z . Those are obviously at most h groups. It is left to show that all these sets fulfill the property of the definition.

Consider one set E_z and the respective L' . Note that the time steps in which all groups $f^{-1}(y)$ for all y on the way from z to the root are not occupied are exactly the ones not occupied in time steps in $I_{u,v,\kappa}$. Let $I_{Y_z} = I_{u,v,\kappa}$. As ℓ_v is in all these groups, it holds

$$\sum_{\ell_v \in E_z} b_u(v) \cdot \frac{|I_{v,\kappa}|}{|I_{u,v,\kappa}|} x_{v,\kappa} \leq \sum_{\ell_v \in E_z} b_u(v) \cdot \frac{T}{|I_{u,v,\kappa}|} \frac{1}{T} \sum_{t=1}^T x_{v,\kappa}^{(t)} = \frac{1}{|I_{Y_z}|} \sum_{t \in I_{Y_z}} \sum_{\ell_v \in E_z} b_u(v) x_{v,\kappa}^{(t)}.$$

As $\sum_{\ell_v \in E_z} b_u(v) x_{v,\kappa}^{(t)} \leq C$ for all time steps t by C -independence, the average over the $t \in I_{Y_z}$ yields

$$\sum_{\ell_v \in E_z} b_u(v) \cdot \frac{|I_{v,\kappa}|}{|I_{u,v,\kappa}|} x_{v,\kappa} \leq C.$$

This shows that we can find a partition of links into h sets such that every set fulfills the scaled C -independence condition. This results in (C, h) -correlated independence. $\square$

4.3 No-Availability-Ordering Regret

Given the correlation model above, we define a notion of regret, which is highly related to the regret used by Kleinberg et al. [KNMS10] and Kanade et al. [KMB09]. Similar to them, we consider the best action in hindsight to be defined via the *best ordering of actions in hindsight*. That is, we compare the summed utilities of the algorithm to the utilities the best ordering would have achieved. This assumes that always the topmost action of the available ones is chosen taking the ordering into account. The main difference is that we do not consider this in terms of its expectation in our later proofs. While Kanade et al. consider the expectation only over the utilities of the best ordering in hindsight, Kleinberg et al. consider the overall regret in expectation. Beside this, Kanade et al. assume the rewards or utility of one time step to be independent of the actual availability. As this does not hold here, we need to resort to the approach considered by Kleinberg et al.

Let $\mathcal{A}^t$ be the *set of actions available* to the algorithm in time step t . The set $\mathcal{A}^t$ consists of one action for each available channel representing transmission and an action for not transmitting at all. For an *ordering* of actions $\sigma \in S_{\mathcal{A}}$ and a time step t we define $\sigma(\mathcal{A}^t) = a$ to be the action a ordered *topmost* in σ of the available actions. Here $\mathcal{A}$ is the set of all actions and $S_{\mathcal{A}}$ is the set of all permutations on $\mathcal{A}$. We consider the availabilities fixed in hindsight. We also consider the utilities – and thus the choices of other players – fixed in hindsight and thus only depending on the actions actually chosen by a link in previous time steps.

Let $u_i^{(t)}(a_i, a_{-i})$ be the *utility* (or reward) the algorithm gets for choosing an action $a_i \in \mathcal{A}^t$ in time step t given all choices of the environment, for example other players or links. The notation $u_i^{(t)}(a_i, a_{-i})$ emphasizes the dependence on actions chosen by other players denoted by a_{-i} . Again an *action vector* consists of actions a_i and a_{-i} .

Definition 4.3. A sequence of action vectors $a^{(1)}, \dots, a^{(T)}$ yields availability-ordering regret $\epsilon \cdot T$ for player i if

$$\max_{\sigma \in S_{\mathcal{A}}} \left(\sum_t u_i^{(t)}(\sigma(\mathcal{A}^t), a_{-i}^{(t)}) \right) - \sum_t u_i^{(t)}(a_i^{(t)}, a_{-i}^{(t)}) \leq \epsilon \cdot T .$$

Note that this definition differs from the one of Kanade et al. only by the fact that the best ordering of actions in hindsight is not considered in expectation over the availabilities. In contrast, the best ordering in hindsight is here considered with the experienced availabilities.

Assuming low regret here is seemingly a weaker assumption than assuming low regret based on Kanade et al.'s definition. Nevertheless, the algorithms proposed by Kanade et al. assume either a full-information setting observing the utilities of an action also in time steps when it is unavailable, or a partial-information setting. Especially, the first one is not applicable here.

Kleinberg et al. [KNMS10] define the regret in basically the same way as we do in the definition above. They consider it in terms of the expectation over the algorithms random choices. To get bounds on this regret, they use the transformation of no-external-regret algorithms as described below for the full-information setting. They also use a similar transformation for the partial-information setting. In these cases, the number of actions actually used by the no-external-regret algorithm is given by the number of permutations of the original actions. This increases the guaranteed regret.

Considering no-external-regret algorithms under adversarial rewards (see Section 2.2), it is clear that those can also be adapted to yield no-availability-ordering regret. For this we extend the set of actions A' such that it consists of all permutations over the actions. Thus, we can use the same algorithms as considered in Section 2.2. To do so, the utility is redefined in terms of the ordering.

For a permutation σ_i^t chosen as action the utility is $u_i(\sigma_i^t(\mathcal{A}^t), a_{-i}^{(t)})$. By Definition 2.1, a sequence of permutations with external regret $\epsilon \cdot T$ in this setting directly results in the definition of low availability-ordering regret, which is

$$\max_{\sigma \in A'} \left(\sum_t u_i^{(t)}(\sigma(\mathcal{A}^t), a_{-i}^{(t)}) \right) - \sum_t u_i^{(t)}(\sigma_i^{(t)}(\mathcal{A}^t), a_{-i}^{(t)}) \leq \epsilon \cdot T .$$

This way, we can easily guarantee the low availability-ordering regret and even get bounds on the regret holding with high probability, i.e., with probability $1 - \frac{1}{n^c}$ for any constant c , using the respective known results (c.f., [LW94, KV05, HP05, ACBFS02]). On the downside, this increases the number of actions exponentially and also the regret, as it typically depends on this number.

The algorithms observe the outcome (i.e., utility values) of either all actions – even of actions unavailable – or only the chosen action. While the bandit model is quite suitable for our availability setting, it is not clear how to define the utility of actions being unavailable to the algorithm in the full-information model. Thus, we start by considering an algorithm given by Kleinberg et al. [KNMS10] in a *semi-partial-information model*, where we assume utilities to be known only for available actions. This complies with the construction for external regret. The algorithms given by Kanade et al. [KMB09] in contrast are either adapted directly to the full-information setting or to the partial-information (or bandit) model as introduced in Section 1.4.

Kleinberg et al. used the approach described above for no-external-regret algorithms considering the set of all permutations as actions. Though this increases the number of actions exponentially, one can consider no-external-regret algorithms having external regret depending logarithmically on the number of actions. This then still guarantees regret polynomial in the number of actions.

Kleinberg et al. apply this approach to a variant of the Randomized Weighted Majority algorithm discussed in Section 2.2, which was introduced by [FS97].

Lemma 4.2 (Theorem 13 in [KNMS10]). *In the semi-partial-information model, a variant of RWM (Algorithm 1) yields in expectation availability-ordering regret of*

$$\mathbf{E} \left[\max_{\sigma \in S_{\mathcal{A}}} \left(\sum_t u_i^{(t)}(\sigma(\mathcal{A}^t)) \right) - \sum_t u_i^{(t)}(\sigma_i^{(t)}(\mathcal{A}^t)) \right] \in O \left(\sqrt{T |\mathcal{A}| \log |\mathcal{A}|} \right).$$

As discussed before, it is more reasonable in many scenarios to consider algorithms for the partial-information model, as an algorithm might not be able to know the outcome of actions not taken. Thus, after having constructed this semi-partial algorithm, we now show that applying very similar techniques yields regret in $o(T)$ in expectation and even with high probability after a sufficient amount of time.

In a very similar fashion Kleinberg et al. also applied the approach to an algorithm for the partial-information setting. They used the algorithm *Exp4* introduced by Auer et al. in [ACBFS02]. Again Kleinberg et al. considered an increased action set extended to all permutations with some additional tweaks to use *Exp4*.

In total, the considerations of Kleinberg et al. give the following guarantee.

Lemma 4.3 (Theorem 15 in [KNMS10]). *In the partial-information model, the *Exp4* algorithm (see [ACBFS02]) yields in expectation availability-ordering regret of*

$$\mathbf{E} \left[\max_{\sigma \in S_{\mathcal{A}}} \left(\sum_t u_i^{(t)}(\sigma(\mathcal{A}^t)) \right) - \sum_t u_i^{(t)}(\sigma_i^{(t)}(\mathcal{A}^t)) \right] \in O \left(|\mathcal{A}| \cdot \sqrt{T \log |\mathcal{A}|} \right).$$

Note that these algorithms make it necessary to update a number of variables or weights exponential in the original number of actions. This exponentially increases the time necessary for one learning step.

While this gives us a bound on the regret in expectation, one can also apply the described approach to no-regret learning algorithms and their high-probability bounds. This way, we can easily extend those high-probability bounds to availability-ordering regret. The *Exp3.P* algorithm used in Lemma 2.2 with probability $1 - c$ actually guarantees availability-ordering regret in

$$O \left(\sqrt{T \cdot 2^{|\mathcal{A}|} \cdot \log \left(2^{|\mathcal{A}|} \frac{T}{c} \right)} \right).$$

Although this bound depends exponentially on the number of actions $|\mathcal{A}|$, it converges to a low regret over time.

We have seen that one can also utilize no-external-regret algorithms in the sleeping-experts setting. However, those have higher regret in terms of the number of actions than in the standard external regret for which they were designed. As there exist algorithms providing low external regret guarantees with high probability, we are able to restore this kind of result when applying these algorithms in our setting.

The high probability bound enables us to consider all links to have low regret whp. Note that our results later on do actually not need every link to have low regret. Instead an assumption, which is easier to achieve, is considered. We assume the *summed up regret* of all links to be low.

4.4 Convergence Results

In this section, we consider algorithms guaranteeing no-availability-ordering regret and show that the number of successful transmissions converges to a constant-factor approximation of the optimal capacity. This result needs to assume links being in separate groups, i.e., the availabilities of channels among links being uncorrelated.

For the more challenging setting of group-correlated availabilities we get an approximation factor depending on how the availabilities are actually correlated. At first we consider our notion of no-availability-ordering regret and in Section 4.4.2 later the notion of no-ordering regret as also considered by Kanade et al. [KMB09].

In fact, the worst-case correlation can even be so harsh that our regret-based algorithm can only reach a trivial approximation guarantee in $\Omega(n)$. The latter result is considered in Section 4.4.3.

4.4.1 Approximation Factors

Now we prove the upper bound on the approximation factors guaranteed by algorithms offering the no-availability-ordering-regret property under group-correlated availabilities. As discussed before, the optimum is different in different time slots depending on available channels. Let us denote the set of links transmitting in the optimal solution on channel κ in time slot t by $OPT_\kappa^{(t)}$. Thus, we compare the number of successful transmissions of the no-regret algorithms to the empirical average capacity of all optima, which is

$$|\overline{OPT}| = \frac{1}{T} \sum_t \sum_{\kappa \in K} |OPT_\kappa^{(t)}|.$$

We prove approximation factors relying on the (C, D) -correlation-independence as a measure for the correlation. Using this notion of independence we prove the theorem stated below.

We use the same notations as introduced earlier in Section 2.3. We denote the fraction of all time steps in which link ℓ_v transmits by q_v and the fraction of time steps with a successful transmission of link ℓ_v by w_v . Analogously we define $q_{v,\kappa}$ and $w_{v,\kappa}$ to be the respective fractions restricted to one channel $\kappa \in K$ and we also consider variables $q_{v,\kappa}^{(t)} \in \{0, 1\}$, where $q_{v,\kappa}^{(t)} = 1$ iff ℓ_v transmits on channel κ in time step t .

Let ϵ_v be the availability-ordering regret of link ℓ_v averaged over the time steps and let $\epsilon = \sum_{v \in V} \epsilon_v$. Action sequences yielding the no-availability-ordering-regret property guarantee ϵ converging to 0.

Theorem 4.1. *A sequence of action vectors of length $T \in \Omega(n^2(\ln(n^2) + \ln(k)))$ with availability-ordering regret ϵ_v for links ℓ_v resulting from a (C, D) -correlation-independent algorithms yields*

$$\sum_v w_v \in \Omega\left(\frac{1}{C \cdot D} \cdot |\overline{OPT}| - \epsilon\right)$$

with high probability $1 - \frac{1}{n}$, where $\overline{OPT}$ denotes the average optimum and $\varepsilon = \sum_{v \in V} \epsilon_v$.

Note that this theorem also directly implies an $O(C \cdot D)$ -approximation factor when $\varepsilon < 1$. So, considering a slightly stronger assumption by assuming $\epsilon_v \leq \frac{1}{n}$ for all links guarantees this approximation factor. This holds with high probability after a number of time steps polynomial in the network size whenever the no-availability-ordering regret decreases polynomial in the number of time steps with high probability.

While the overall approach is in the spirit of previous works and the previous Sections 2.3 and 3.3, our setting is quite different and the notion of regret also differs. Similar as before, our analysis starts with the observation that a constant fraction of all transmission attempts is successful. Afterwards, we combine this with the result that the number of transmission attempts is in $\Omega(\overline{OPT})$. Especially the proof of this latter statement in the proof of Theorem 4.1 below needs more advanced techniques. To combine these results, we also need to bound the difference between the expected and the actually experienced interference. Together these statements prove our theorem.

First of all, let us bound the number of successful transmissions by the number of transmission attempts. The proof directly follows the same thread as the one seen for Lemma 2.3 and others in Chapter 2 and 3. We extend the arguments seen there to the setting with multiple channels.

Lemma 4.4 (cf. [Din10, AM11, DHK12b]). *It holds $w_v \leq q_v \leq 2 \cdot w_v + \epsilon_v$ and $\sum_v w_v \leq \sum_v q_v \leq 2 \cdot \sum_v w_v + \varepsilon$.*

Proof. The left-hand side of the first inequality follows by definition. For the right-hand side, we use the fact that for each link ℓ_v the average regret is at most ϵ_v . Therefore, not to send at all can increase the average utility per step by at most ϵ_v . Ordering “not sending” above all channels guarantees at least a reward of 0. Thus, formally

$$\max_{\sigma} \left(\sum_t \sum_{a \in \sigma(A^t)} u_i(a, a_{-i}^t) \right) \geq 0 .$$

This way, we get

$$-\sum_t u_i(a_{-i}^t, a_i) \leq \epsilon_v \cdot T ,$$

which can be evaluated to $-w_v - (q_v - w_v) \leq \epsilon_v$.

Simple calculus directly yields $q_v \leq 2w_v + \epsilon_v$. Taking the sum over all v we get $\sum_v q_v - 2 \cdot \sum_v w_v \leq \varepsilon$. This proves the claim. $\square$

Lemma 4.4 shows that the number of successful transmissions and the number of transmission attempts only differ by a constant factor. The property considered in this lemma matches the (γ, ϵ) -successful property used in Section 2.3.

To prove Theorem 4.1, we use a primal-dual approach using an appropriately defined linear program. Though this general approach is similar to the one used in Section 2.3, we need to account for some more intricate details here.

We define $f'_{v,\kappa}$ to be the fraction of time steps in which link ℓ_v would not have been successful if it had transmitted. Formally it holds

$$f'_{v,\kappa} = \frac{1}{|I_{v,\kappa}|} \sum_{t \in I_{v,\kappa}} \sum_{u \in V, u \neq v} b_u(v) q_{u,\kappa}^{(t)} ,$$

where $I_{v,\kappa}$ is the set of time steps in which channel κ is available to link ℓ_v . We now start with some consideration concerning the stochastic nature of the model. The stochastic process considered here effectively yields a bound on a random variable similar to the expected interference from other links. Let $F'_{v,\kappa}$ be the interference assuming that the $q_{u,\kappa}^{(t)}$ are fixed. Thus,

$$F'_{v,\kappa} = \frac{1}{T} \sum_{u \in V, u \neq v} \sum_{t=1}^T b_u(v) q_{u,\kappa}^{(t)} X_{u,t,\kappa}^* ,$$

for random variables $X_{u,t,\kappa}^* \in \{0, 1\}$ with $X_{u,t,\kappa}^* = 1$ iff the channel is available to link ℓ_v given that it is available to link ℓ_u . It holds $\Pr[X_{u,t,\kappa}^* = 1] = \frac{|I_{v,\kappa}|}{|I_{u,v,\kappa}|}$. Note that those random variables are actually not independent among different u but among the time steps t .

Due to fixing $q_{u,\kappa}^{(t)}$ and the size of $I_{v,\kappa}$ as observed in the run of the algorithm, the expectation of this random variable is $\mathbf{E}[F'_{v,\kappa}] = \mathbf{E}[f'_{v,\kappa} \mid |I_{v,\kappa}|, q_{u,\kappa}^{(t)} \text{ for all } \ell_u \text{ and } t]$. With this we prove the following lemma stating that the actual experienced interference $f'_{v,\kappa}$ does not differ too much from the expected interference based on $F'_{v,\kappa}$.

Lemma 4.5. *Assume $T \in O(n^2 \cdot (\ln(n^2) + \ln(k)))$. It holds with high probability that for all $\ell_v \in V$ and all channels $\kappa \in K$ with $\mathbf{E}[F'_{v,\kappa}] \leq \frac{1}{8}$ we have*

$$f'_{v,\kappa} \leq \frac{1}{4} .$$

Proof. To prove this lemma, we bound the probability $\Pr[f'_{v,\kappa} - \mathbf{E}[F'_{v,\kappa}] \geq \frac{1}{8}]$. With this it is straightforward that assuming $f'_{v,\kappa} - \mathbf{E}[F'_{v,\kappa}] \leq \frac{1}{8}$ together with $\mathbf{E}[F'_{v,\kappa}] \leq \frac{1}{8}$ we get $f'_{v,\kappa} \leq \frac{1}{4}$.

Let $f'_{u,v,\kappa} = \frac{1}{|I_{v,\kappa}|} \sum_{t \in I_v} b_u(v) q_u^{(t)}$ and $F'_{u,v,\kappa} = \frac{1}{T} \sum_t b_u(v) q_u^{(t)} X_{u,t,\kappa}^*$. We now consider

$$\Pr\left[f'_{u,v,\kappa} - \mathbf{E}[F'_{u,v,\kappa}] \geq \frac{1}{8n}\right] = \Pr\left[f'_{u,v,\kappa} - \mathbf{E}[f'_{u,v,\kappa} \mid |I_{v,\kappa}|, q_{u,\kappa}^{(t)} \text{ for all } \ell_u \text{ and } t] \geq \frac{1}{8n}\right] .$$

This holds as again $\mathbf{E}[F'_{u,v,\kappa}] = \mathbf{E}[f'_{u,v,\kappa} \mid |I_{v,\kappa}|, q_{u,\kappa}^{(t)} \text{ for all } \ell_u \text{ and } t]$. We can reinterpret the random process leading to $f'_{u,v,\kappa}$ by using the fact that the fixed variables can be considered to be given by an adversary only depending on previous time steps.

Consider the random variable $Y_t = \sum_{t'=1}^t b_u(v) q_u^{(t')} X_{u,t',\kappa}^* - \mathbf{E}\left[\sum_{t'=1}^t b_u(v) q_u^{(t')} X_{u,t',\kappa}^* \mid H_t\right]$, where H_t denotes the history of the algorithms run. The sequence of Y_t constitutes a martingale with $|Y_t - Y_{t-1}| \leq 1$. Thus, we can apply Hoeffding bounds even when the coefficients are fixed in this way and get

$$\Pr\left[\frac{1}{T} \sum_t b_u(v) q_u^{(t)} X_{u,t,\kappa}^* - \mathbf{E}\left[\frac{1}{T} \sum_t b_u(v) q_u^{(t)} X_{u,t,\kappa}^*\right] \geq \frac{1}{8n}\right] \leq \exp\left(-2 \frac{T}{8^2 n^2}\right) .$$

Using a union bound over the different links yields

$$\Pr\left[\forall u \in V : \frac{1}{T} \sum_t b_u(v) q_u^{(t)} X_{u,t,\kappa}^* \geq \mathbf{E}\left[\frac{1}{T} \sum_t b_u(v) q_u^{(t)} X_{u,t,\kappa}^*\right] + \frac{1}{8n}\right] \leq n \cdot \exp\left(-2 \frac{T}{8^2 n^2}\right) .$$

Having $\frac{1}{T} \sum_t b_u(v) q_u^{(t)} X_{u,t,\kappa}^* - \mathbf{E} \left[\frac{1}{T} \sum_t b_u(v) q_u^{(t)} X_{u,t,\kappa}^* \right] \geq \frac{1}{8n}$ for all u directly shows

$$\frac{1}{T} \sum_u \sum_t b_u(v) q_u^{(t)} X_{u,t,\kappa}^* - \mathbf{E} \left[\frac{1}{T} \sum_u \sum_t b_u(v) q_u^{(t)} X_{u,t,\kappa}^* \right] \geq \frac{1}{8} .$$

Thus, taking again the union bound over all links and over all channels we get

$$\Pr \left[\forall v \in V, \kappa \in K : f'_{v,\kappa} - \mathbf{E} [F'_{v,\kappa}] \geq \frac{1}{8} \right] \leq n^2 \cdot k \cdot \exp \left(-2 \frac{T}{8^2 n^2} \right) .$$

Setting $T \geq \frac{8^2 n^2}{2} ((c+1) \cdot \ln(n^2) + \ln(k))$ yields a probability of at least $1 - \frac{1}{n^c}$. Setting $\mathbf{E} [F'_{v,\kappa}] \leq \frac{1}{8}$ then results in $f'_{v,\kappa} \leq \frac{1}{4}$. $\square$

We now consider the set of channels with *low congestion*, where a link would be unsuccessful in a small fraction of time slots. For these channels we show that the number of transmission attempts is an upper bound on the availabilities. We then use this result in the final proof of Theorem 4.1.

Lemma 4.6. *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 the regret of link ℓ_v is at most ϵ_v it holds*

$$2 \sum_{\kappa \in K} q_{v,\kappa} + 2\epsilon_v \geq \bar{P}_{v,M} .$$

Proof. The utility of the best ordering in hindsight is obviously at least as high as the utility of any other ordering. Thus, it is also at least as high as the utility of the ordering in which all $\kappa \in M$ are ordered above the action ‘not sending’. All other channels are ordered below ‘not sending’.

To illustrate the general thread of arguments at first, we consider just one channel. On any channel $\kappa \in M$ link ℓ_v is not successful in at most an $f'_{v,\kappa}$ -fraction of time steps in which channel κ is available. This holds, due to the summed up interference being an upper bound on the number of unsuccessful time slots. That leaves $T \cdot (1 - f'_{v,\kappa})$ time steps possibly successful each having a utility of +1 for choosing κ if it was available. Choosing κ in contrast also has -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 utility of at least $\bar{P}_{v,\kappa} \cdot (T \cdot (1 - f'_{v,\kappa}) - T \cdot f'_{v,\kappa})$. We extend this argument to the set M in the following way.

For any ordering with all $\kappa \in M$ ordered above “not sending” and all other channels below, we get at least the utility of the worst channel $\kappa \in M$ if any channel of M is available. This way, we get at least $\min_{\kappa} ((1 - f'_{v,\kappa}) - f'_{v,\kappa})$ for time steps where any channel in M is available. For the utility of the best ordering in hindsight this is

$$\max_{\sigma \in S_A} \sum_{t=1}^T \sum_{a=\sigma(\mathcal{A}^t)} u_v(a, a_{-v}^{(t)}) \geq \min_{\kappa} ((1 - f'_{v,\kappa}) - f'_{v,\kappa}) \bar{P}_{v,M} \cdot T . \quad (4.1)$$

We use $f'_{v,\kappa} \leq \frac{1}{4}$, which results in

$$\frac{1}{T} \max_{\sigma \in S_A} \sum_{t=1}^T \sum_{a=\sigma(\mathcal{A}^t)} u_v(a, a_{-v}^{(t)}) \geq \left(\frac{3}{4} - \frac{1}{4} \right) \bar{P}_{v,M} = \frac{1}{2} \bar{P}_{v,M} .$$

Using the regret being at most ϵ_v and the utility being at most q_v we get $\frac{1}{2} \bar{P}_{v,M} \leq q_v + \epsilon_v$. $\square$

This connection between the availability of a set, its interference, and the actions played now allows us to prove Theorem 4.1. We use a very similar linear-programming-based proof as done in Section 2.3. The LPs need some modifications to accustom for the multi-channel environment.

Proof of Theorem 4.1. We will prove our theorem with the following primal-dual approach. For a primal linear program, the set of links that yields the optimal capacity is a viable solution. We can construct a dual linear program, for which the no-regret sequence is a solution. By definition of linear programs, the objective function of any feasible solution to the dual linear program is an upper bound on the value of the objective function of any feasible solution to the primal linear program. This fact will be used for the proof of the approximation guarantee given in the theorem.

For the following primal LP we essentially consider the average optimum and utilize (C, D) -independence.

$$\begin{aligned}
 \max. \quad & \sum_{v \in V} \sum_{\kappa \in \mathcal{K}} x_{v,\kappa} \\
 \text{s.t.} \quad & \sum_{v \in V} \frac{|I_v|}{|I_{u,v,\kappa}|} b_u(v) x_{v,\kappa} \leq C \cdot D \quad \forall u \in V, \kappa \in \mathcal{K} \\
 & \sum_{\kappa \in M} x_{v,\kappa} \leq \bar{P}_{v,M} \quad \forall v \in V, M \subseteq \mathcal{K} \\
 & x_{v,\kappa} \geq 0 \quad \forall v \in V, \kappa \in \mathcal{K}
 \end{aligned}$$

The LP mainly differs from the one considered in Section 2.3 by incorporating multiple channels. The variable $x_{v,\kappa}$ represents the fraction of all time steps where link ℓ_v transmits on channel κ in the optimum. So the objective function is the average number of simultaneous transmissions per time step. The first constraint reflects (C, D) -correlation-independence. The second constrains the variable $x_{v,\kappa}$ summed up for channels κ in any set $M \subseteq \mathcal{K}$ to be at most the fraction of time steps in which a channel out of M was available. This clearly needs to be fulfilled for any feasible schedule for capacity maximization.

Observe that $x_{v,\kappa} = \frac{|\{t \mid v \in \text{OPT}'^{(t)}_{\kappa}\}|}{T}$ represents a feasible solution to this LP. The first constraint is fulfilled as the scaled C -independence is fulfilled for any ℓ_u and every set L in any set of the disjoint partition $\mathcal{D}_{u,\kappa}$ in every single time slot. As there are at most D such sets, the constraint is fulfilled. The second constraint is fulfilled due to the fact that at most one channel is used at a time. Thus, we get $\sum_{v \in V} \sum_{\kappa \in \mathcal{K}} x_{v,\kappa} \geq \Omega(|\text{OPT}'|)$ by the definition of C -independence for the single-slot optimum.

Constructing the dual to this primal LP yields

$$\begin{aligned}
 \min. \quad & \sum_{v \in V} \sum_{\kappa \in \mathcal{K}} CD \cdot y_{v,\kappa} + \sum_{v \in V} \sum_{M \subseteq \mathcal{K}} \bar{P}_{v,M} \cdot z_{v,M} \\
 \text{s.t.} \quad & \sum_{u \in V} \frac{|I_v|}{|I_{u,v,\kappa}|} b_u(v) y_{u,\kappa} + \sum_{M: \kappa \in M} z_{v,M} \geq 1 \quad \forall v \in V, \kappa \in \mathcal{K} \\
 & y_{v,\kappa}, z_{v,M} \geq 0 \quad \forall v \in V, \kappa \in \mathcal{K}, M \subseteq \mathcal{K}
 \end{aligned}$$

We construct the following dual solution that gives an upper bound to the solution of the primal LP. Let $M_v = \{\kappa \in \mathcal{K} \mid f'_{v,\kappa} \leq \frac{1}{4}\}$, 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 $y_{v,\kappa} = 8 \cdot q_{v,\kappa}$, $z_{v,M_v} = 1$, 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. As all $b_u(v) \geq 0$ it is easy to see that we get

$$\sum_{u \in V} \frac{|I_v|}{|I_{u,v,\kappa}|} b_u(v) y_{u,\kappa} = 8 \cdot \sum_{u \in V} \frac{|I_v|}{|I_{u,v,\kappa}|} b_u(v) q_{u,\kappa} = 8 \cdot \sum_{u \in V} \frac{|I_v|}{|I_{u,v,\kappa}|} \frac{1}{T} b_u(v) \sum_{t \in T} q_{u,\kappa}^{(t)}.$$

This directly yields

$$8 \cdot \frac{1}{T} \sum_{t \in T} \sum_{u \in V} b_u(v) q_{u,\kappa}^{(t)} \frac{|I_v|}{|I_{u,v,\kappa}|} = 8 \cdot \mathbf{E}[F'_{v,\kappa}].$$

Any channel κ with $\mathbf{E}[F'_{v,\kappa}] \geq \frac{1}{8}$ guarantees $\sum_{u \in V} \frac{|I_{v,\kappa}|}{|I_{u,v,\kappa}|} b_u(v) q_{u,\kappa} \geq \frac{1}{8}$. So the constraint is fulfilled with the chosen $y_{u,\kappa}$.

For the other case with $\mathbf{E}[F'_{v,\kappa}] < \frac{1}{8}$ we have Lemma 4.5 to conclude that $f'_{v,\kappa} < \frac{1}{4}$ holds with high probability. Thus, $z_{v,M_v} = 1$, where by definition $\kappa \in M_v$, fulfills the constraint as necessary.

Using Lemma 4.6 leads to an upper bound on the objective function of the dual LP of

$$\sum_v (8CDq_v + \bar{P}_{v,M}) \leq \sum_v (8CDq_v + 2q_v + 2\epsilon_v).$$

Combined with the primal LP this is $C \cdot D \cdot \sum_v (q_v + 2\epsilon_v) = \Omega(|\overline{OPT}|)$. To relate this bound to the fraction of successful transmissions, we can use Lemma 4.4 directly yielding $C \cdot D \cdot \sum_v (w_v + \epsilon_v) = \Omega(|\overline{OPT}|)$. In particular, for arbitrary constant independence the last derivation obviously implies $\sum_v w_v = \Omega(|\overline{OPT}| - \epsilon)$ and directly yields a constant approximation factor. $\square$

The theorem shows that after a number of time slots logarithmic in k and polynomial in n , a sequence with low regret converges to a constant-factor approximation with high probability for constant (C, D) -independence. For an algorithm to actually converge, one also has to consider the time needed until the action sequence of the algorithm has low regret. That is, for example, an average regret being below $\frac{1}{2n}$ for all links or the summed up regret of all links being below 1.

Combining the insights of Theorem 4.1 with Lemma 4.1 we can directly conclude the following corollary and Corollary 4.2 in the next section.

Corollary 4.1. *Consider a network of n links with C -independent conflict graph and laminar group-correlated channel availabilities of k channels with a height of h .*

The sequence of action vectors of length $T \in \Omega(n^2(\ln(n^2) + \ln k))$ with availability-ordering regret ϵ_v for all links ℓ_v yields

$$\sum_v w_v \in \Omega\left(\frac{1}{C \cdot h} \cdot |\overline{OPT}| - \epsilon\right)$$

with high probability, where $\epsilon = \sum_v \epsilon_v$.

The algorithmic approaches presented in Section 4.3 provide an availability-ordering regret of at most $O\left((k+1)\sqrt{T \log(k+1)}\right)$ in expectation in the partial-information setting.

This is even slightly better in the semi-partial-information model. To reach for example $\epsilon_v < \frac{1}{2n}$ we therefore expect to need only a number of time steps polynomial in n and in the number of channels k .

This shows that no-regret-learning approaches can be used to guarantee a number of successful transmissions being only a constant factor from the optimal number in a constant-independent network after a polynomial time.

4.4.2 Uncorrelated Availabilities

Algorithms yielding no-availability-ordering regret as proposed in Section 4.3 can also be applied when the availabilities are uncorrelated among different links. The previous result even guarantees these to converge to an $O(C)$ -approximation. This is easy to see as the setting can be embedded into a laminar setting by having one group with probability 1 not occupied by the primary user and a separate group for every link with the respective probability. This laminar correlation setting has height 2 and results in the following corollary.

Corollary 4.2. *Consider a network of n links with C -independent conflict graph and group-correlated channels with availabilities independent among the links.*

The sequence of action vectors of length $T \in \Omega(n^2(\ln(n^2) + \ln k))$ with availability-ordering regrets ϵ_v for links ℓ_v and $\epsilon = \sum_{v \in V} \epsilon_v$ yields

$$\sum_v w_v = \Omega\left(\frac{1}{C} \cdot |\overline{OPT}| - \epsilon\right)$$

with high probability.

As argued in Section 4.3 we use the kind of regret introduced above to compensate for correlation of the availabilities. Thus, we are not able to use the kind of regret introduced by Kanade et al. [KMB09] there. We can nevertheless use this notion of *ordering regret* when the availabilities are not correlated among the links. The following notion of regret is very closely related to the availability-ordering regret considered above.

Definition 4.4 (Kanade et al. [KMB09]). *Let $a^{(1)}, \dots, a^{(T)}$ be a sequence of action vectors. The ordering regret of this sequence for player i is defined as*

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

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 .

Similar as before, an algorithm computing an action sequence with ordering regret in $o(T)$ is called a *no-ordering-regret algorithm*.

Kanade et al. give an algorithm having diminishing ordering regret for an increasing number time steps. This algorithm is a variant of the *Follow the Perturbed Leader algorithm* (FPL) by Kalai and Vempala [KV05]. In the full-information setting, their algorithm *Sleeping Follow the Perturbed Leader* (SFPL) maintains a list of total rewards (or utilities) observed for the actions. In each time step, the algorithm orders the actions based on this summed up rewards and some *perturbation* based on a parameter μ and chooses the action of the available ones ordered topmost in this list.

Lemma 4.7 (Lemma 3 in [KMB09]). *SFPL in expectation has ordering-regret in*

$$O\left(\mu(k+1)T + \frac{1}{\mu}\right)$$

when the pairwise differences between all utilities are bounded by a constant.

Choosing the parameter μ appropriately depending on T guarantees the no-regret property. For cases where the total number of time steps is not known in advance this can still be achieved by standard techniques like the doubling trick. Starting with $T = 1$ and doubling the used T whenever T time steps elapsed ensures that the no-regret property can still be achieved.

It is obvious that other variants of FPL could also be used in the construction of Kanade et al.'s SFPL algorithm. This then yields slightly different regret bounds and can even enable us to avoid the doubling trick for unknown T . The variant introduced by Kalai and Vempala [KV05] as FPL^* for example uses exponentially distributed variables for the perturbation. Applying this variant, Hutter and Poland [HP05] show that for T unknown in advance we do not need to use the doubling trick but can choose the parameters based on the time elapsed until then.

Another algorithm given by Kanade et al. applicable in the full-information model even provides a bound of the expected regret of $\sqrt{T \log(k+1)}$.

Based on SFPL, Kanade et al. also introduce a partial-information algorithm $BSFPL$ with the following regret bound.

Lemma 4.8 (Theorem 2 in [KMB09]). *BSFPL with suitable parameters yields ordering-regret in*

$$O\left((k+1)^{\frac{4}{5}} T^{\frac{4}{5}} \log T\right)$$

in expectation.

Considering an algorithm having ordering-regret $\epsilon \cdot T$, Lemma 4.4 holds accordingly without any adjustment of the proof. To find an equivalent result to Lemma 4.6 we have to find a bound on the time until availabilities are close to the expected ones. Recall that $p_{\min} = \min\{p_{v,M} \mid v \in V, M \neq \emptyset, p_{v,M} > 0\}$ is the minimal probability of the availabilities of all the channels.

Let us now show in Lemma 4.9 that for a number of time slots $T \in \Omega\left(\frac{\ln n}{p_{\min}}\right)$ 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 use this result to draw a connection between transmission attempts, availabilities, and experienced affectances in Lemma 4.10 finally proving convergence to a constant-factor approximation.

Lemma 4.9. *After a number of time steps $T \in \Omega\left(\frac{\ln n}{p_{\min}}\right)$ it holds*

$$|\bar{P}_{v,M_v} - p_{v,M_v}| \leq \frac{1}{2} \bar{P}_{v,M_v}$$

for a (fixed) set of channels M_v for all links ℓ_v with probability $1 - \frac{1}{n^c}$ for any constant c .

Proof. Consider the random variable $P_{v,M_v}^{(t)} \in \{0,1\}$ indicating whether any channel of the set M_v is available for link ℓ_v in time slot t . Let $Y = \sum_t P_{v,M_v}^{(t)}$. Thus, we need $|\mathbb{E}(Y) - Y| < \frac{1}{2}Y$ to hold, because this directly implies $|\bar{P}_{v,M_v} - p_{v,M_v}| \leq \frac{1}{2}\bar{P}_{v,M_v}$. 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 (see for example [DP98]). This directly yields $\Pr[Y \geq (1 + \delta)\mathbb{E}(Y)] \leq \exp\left(-\frac{\delta^2}{3}\mathbb{E}(Y)\right)$ and $\Pr[Y \leq (1 - \delta)\mathbb{E}(Y)] \leq \exp\left(-\frac{\delta^2}{2}\mathbb{E}(Y)\right)$ for every $\delta \in [0, 1]$.

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

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

This is at most $2 \cdot \exp(-\frac{1}{18}p_{v,M_v}T)$. Applying another union bound to guarantee that this holds for all links (and the respective sets M_v) results in

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

Setting $T \geq \frac{18}{p_{\min}}((c+1) \ln n \cdot \ln 2)$ results in the probability for any arbitrary set of channels M to not have $|\bar{P}_{v,M} - p_{v,M}| \leq \frac{1}{2}\bar{P}_{v,M}$ is at most n^{-c} . $\square$

Using this result, we are able to basically reestablish the results in Lemma 4.6 from the previous section. As the proof is almost the same as the one of the previous section, we will only consider the differences to this proof here.

Let $f_{v,\kappa}$ be the experienced fraction of all time steps, in which link ℓ_v would not have been able to transmit successfully. Note that $f_{v,\kappa}$ is not restricted to time steps in which the channel κ was actually available to ℓ_v for transmission.

Lemma 4.10. *For every link ℓ_v let M_v be a set of channels such that every channel $\kappa \in M$ yields $f_{v,\kappa} \leq \frac{1}{4}$. If the regret is at most ϵ_v and the number of time steps $T \in \Omega\left(\frac{\ln n}{p_{\min}}\right)$, then with high probability*

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

Proof. The proof almost exactly follows the one of Lemma 4.6. Up to Equation 4.1 the same arguments hold. To consider the ordering regret here, we have to exchange Equation 4.1 by the one below.

This holds due to the independence of the availabilities between different links, as we can fix the actions of other links. For the expected utility of the best ordering in hindsight we get

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

This yields

$$\frac{1}{T} \max_{\sigma \in S_{\mathcal{A}}} \mathbb{E}\left[\sum_{t=1}^T u_v(\sigma(\mathcal{A}^{(t)}), a_{-v}^{(t)})\right] \geq \left(\frac{3}{4} - \frac{1}{4}\right) p_{v,M} = \frac{1}{2} p_{v,M}. \quad (4.2)$$

Using $|\bar{P}_{v,M} - p_{v,M}| \leq \frac{1}{2}\bar{P}_{v,M}$ from Lemma 4.9, we can easily bound Equation 4.2 from below by $\frac{1}{4}\bar{P}_{v,M}$. With the fact that the regret is at most ϵ_v and that the utility is at most q_v we get $\frac{1}{4}\bar{P}_{v,M} \leq q_v + \epsilon_v$.

Note that 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 more than $T \cdot \max_{\kappa} f_{v,\kappa} \cdot p_{v,M}$. This is due to correlation, for example, being able to force all interference of other links on a channel to occur only in available time steps even if this represents only a in few time steps in total. $\square$

In total we can now consider the approximation factors guaranteed by no-ordering-regret algorithms. This proof again highly resembles the proof of Theorem 4.1 in the previous section. Thus, we will again focus on the differences.

Theorem 4.2. *Consider a network of n links with a C -independent conflict graph. Every sequence of length $T \in \Omega\left(\frac{\ln n}{p_{\min}}\right)$ with ordering regrets ϵ_v for links ℓ_v yields with high probability*

$$\sum_v w_v \in \Omega\left(\frac{|\overline{OPT}|}{C} - \frac{\varepsilon}{C}\right),$$

where $\varepsilon = \sum_{v \in V} \epsilon_v$.

Proof. The proof follows the same primal-dual scheme as the one of Theorem 4.1. Nevertheless, we consider a slightly different linear program, which allows us to utilize the fact that the availabilities among the links are assumed to be independent.

Recall Definition 1.3 of C -independence. Note that the conditions are defined for a feasible set of links in a single time step. They can be transferred for each channel κ to all time steps by averaging as follows. Let $OPT_{\kappa}'^{(t)}$ be L' from Definition 1.3 when setting $L = OPT_{\kappa}^{(t)}$ yielding $|OPT_{\kappa}'^{(t)}| \geq \Omega(|OPT_{\kappa}^{(t)}|)$ and $\sum_{v \in OPT_{\kappa}'^{(t)}} b_u(v) \leq C$ for each time step t . By averaging over all time steps we get

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

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 use a primal-dual approach again to prove the theorem which is only slightly different to the one used in the proof of Theorem 4.1. Thus, we consider the optimum averaged over all time steps again and utilize C -independence. We exchange the second constraint in the LP used in the proof of Theorem 4.1 for the (C, D) -correlation-independence with the C -independence. The considerations above are useful to show feasibility of the primal solution.

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

Observe that the averaged optimum $x_{v,\kappa} = \frac{|\{t \mid v \in OPT'_\kappa(t)\}|}{T}$ again represents a feasible solution for this LP and we get $\sum_{v \in V} \sum_{\kappa \in K} x_{v,\kappa} \geq \Omega(|OPT|)$ by definition.

The dual LP to this primal LP is

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

The dual solution can again be constructed as follows. Let $M_v = \{\kappa \in K \mid f_{v,\kappa} \leq \frac{1}{4}\}$ represent 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 $S \neq M_v$.

This is a feasible solution. 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 $f_{v,\kappa} \geq \frac{1}{4}$ we get $\sum_{u \in V} b_u(v) q_{u,\kappa} \geq \frac{1}{4}$. Hence, $\sum_u b_u(v) \cdot y_{u,\kappa} \geq 1$ 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 4.4 and Lemma 4.10 for the set M_v leads to an upper bound on the objective function of the dual LP of

$$\sum_v (4Cq_v + \bar{P}_{v,M_v}) \leq \sum_v (4Cq_v + 4q_v + 4\epsilon_v) \leq \sum_v (8Cw_v + 8w_v) + 4\epsilon.$$

Combined with the primal LP this results in $\sum_v w_v \in \Omega\left(\frac{|OPT|}{C} - \frac{\epsilon}{C}\right)$. $\square$

This result shows that not only the algorithm derived by utilizing no-external-regret algorithms but also other algorithms like the ones proposed by Kanade et al. can be applied to general multi-channel capacity maximization. It also emphasizes how general the proof technique can be applied to analyze no-regret algorithms for capacity maximization.

4.4.3 Lower Bounds

In this section, we begin by showing that a direct application of no-external-regret algorithms does not necessarily yield a constant-factor approximation. This is not surprising considering that external-regret algorithms applied directly without modification of the action space can only choose one channel.

In fact, we give an example of a sequence of action vectors with 0 external regret that shows approximation factors in $\Omega(k)$ and $\Omega(n)$. These factors can already be reached by algorithms utilizing just one channel or letting just one link transmit, respectively.

Additionally, we show that action sequences under group-correlated availabilities with zero availability-ordering regret can result in approximation factors as bad as $\Omega(n)$. The same example also shows that for 0 ordering regret the assumption of stochastic independence of the availabilities among links considered in Section 4.4.2 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 without correlation obtains constant-factor approximations.

Let the set of actions consist of one action per channel and an additional action representing “not transmitting”. Links have to choose on which channel to transmit or not to transmit at all.

Theorem 4.3. *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 simultaneously successful on every channel. For any arbitrary link we consider a sequence of deterministic availabilities in which channel κ is available in time slots t with $(t \bmod k + 1) = \kappa$. It is easy to see that there is a 0-external-regret sequence in which exactly one channel is chosen. Thus, the link transmits only in every k -th time step. In contrast, in the optimum the link can simply choose another channel in every single time step. This way, the factor is k .

This argument does not need the sequence to be fixed but can also be reproduced with stochastic availability. This is clear when we allow the availabilities to be correlated among the channels of a link. If this is not the case, we can 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 transmits only 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 a factor of k in the long run. $\square$

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

Corollary 4.3. *For every number of links n there is a network, such that for every sequence of actions 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, we have seen that one can utilize no-external-regret algorithms by increasing the action space exponentially. It is nevertheless an interesting open problem if even more intricate approaches allow to establish similar properties as for these algorithms leading to a constant-factor approximation without the disadvantage of increased regret as discussed in Section 4.3.

In the previous sections, we have assumed group-correlated availabilities. Our results mainly consider laminar group-correlation and show constant-factor approximations only for constant heights. This implies a limited independence of the availabilities among links.

We show that for group-correlated availabilities, we can find a setting where each link is in at most h groups and a 0-availability-ordering-regret sequence actually yields at best an $O(h)$ -approximation. This result complements the result shown in Corollary 4.1. We use a similar example here as in the proof above.

Theorem 4.4. *There are instances of a group-correlated setting with n links with at most $h \leq n$ groups and a sequence of actions yielding 0 ordering regret for every link, such that the number of successful transmissions is at least a factor of h smaller than in an optimal schedule.*




Figure 4.1: Example of an optimal schedule (circles) and a no-regret sequence (crosses) for an availability structure with $h = n = 5$ used in the proof of Theorem 4.4

Proof. We assume a setting with only one channel and construct the network as follows. No pair of links can transmit simultaneously on the channel. This can easily be achieved by placing links (almost) in the same location and constructing the interference and the conflict graph appropriately.

We partition the links into $h + 1$ sets $S_0, \dots, S_h$, where at least one link is in each set. The channel is either available for all h sets simultaneously or for only one set of links. Let $G_1, \dots, G_h$ be the groups of the group-correlation and the probability for each group to be unoccupied is $\frac{1}{2}$. The links in set S_0 are in all h groups. The links of every other set S_i are in all groups but G_i .

This way, for every set of links S_i with $i \neq 0$ the probability that the channel is also available to links in S_0 when it is available to S_i is $\frac{1}{2}$. Due to this choice of groups the channel can only be available to all links simultaneously, to exactly one set of links, or to no link at all. Note that we can safely omit time steps where the channel is unavailable to all links. Then, it holds that in an $\frac{1}{h}$ -fraction of the time the channel is available to all links and in the rest to exactly one set of links.

This structure of channel availabilities is depicted in Figure 4.1 with $h = n = 5$. It is easy to see that in every n -th time step the channel is available to all links. In these time steps one link (e.g., link ℓ_0) is in the optimum and one link transmits in the no-regret sequence (here, also ℓ_0). In the remaining time steps only one link at a time is able to use the channel. In the optimum this one transmits, while in the no-regret sequence no one transmits then. Note that time steps, in which the channel is unavailable to all links, are omitted in the figure.

We construct the 0-availability-ordering request sequence by scheduling one link ℓ_0 out of set S_0 to send whenever the channel is available for it. All other links choose not to send at all. As ℓ_0 can transmit successfully throughout the whole sequence it has 0 regret. The dependence of the availabilities implies that the (expected) utility of the best response in hindsight for all links but ℓ_0 becomes not to transmit, because, in the long run, half of the time slots where the channel is available are shared with ℓ_0 . In those slots the link would not be able to transmit successfully. Hence, “transmitting in all time steps” and “not transmitting at all” have the same utilities, and the regret is 0 for all links. This sequence provides a successful transmission every h -th of the considered time steps.

In contrast, in the optimum one link of every set S_i transmits when the channel is avail-

able to it alone and one link of the set S_0 otherwise. This yields a successful transmission in every time slot. Note that we can easily leave out all time steps where the channel is not available to any link, as this does not change the argument.

This proves a factor of h between the number of successful transmissions in the optimal schedule and the 0-(availability)-ordering-regret sequence. $\square$

With a slight adjustment of the one-or-all structure, it is possible to show even slightly stronger lower bounds close to $\frac{3h}{2}$ for the number of groups being h . This is done by having at least two links in group S_0 . In the no-regret sequence both transmit in $\frac{2}{3}$ of the time steps, in which the channel is available to S_0 , such that $\frac{1}{2}$ of their transmissions overlap being unsuccessful.

By setting the number of groups $h = n$, we directly get a lower bound in $\Omega(n)$ for correlated availabilities both for 0-availability-ordering-regret sequences as used in Section 4.4.1 and 0-ordering-regret sequences as given by Kanade et al. and used in Section 4.4.2.

Corollary 4.4. *There exist group-correlated availabilities such that for a network with n links with a sequence yielding 0 availability-ordering regret for every link, the number of successful transmissions is at least a factor of n smaller than in an optimal schedule.*

The one-or-all structure of availabilities used in the proof of Theorem 4.4 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 4.4 do not guarantee 0 ordering regret as defined in Section 4.4.2. For 0 ordering regret the probability for this to happen diminishes over time.

4.5 Summary

Concerning wireless networks with multiple channels, one can see that quite similar approaches to the ones considered in previous chapters and by Diniz [Din10] and Ásgeirsson and Mitra [AM11] can be utilized to maximize the capacity. This can even be done when the channels are defined by a stochastic process rendering certain channels unavailable during some time steps. This enables us to use adaptive no-regret-learning approaches in very customizable primary-secondary-user settings. Here, primary users have the exclusive right to use the spectrum. They can allow secondary users to utilize the channels in certain time steps when not using the channel themselves.

No-availability-ordering-regret algorithms and no-ordering-regret algorithms are guaranteed to converge to a constant-factor approximation under constant-independent conflict graphs and uncorrelated availabilities. For more intricate laminar group-correlated settings, we can still guarantee an approximation factor linear in the height of the laminar setting.

In contrast to this, arbitrary correlation among the availabilities can worsen the approximation factors significantly. We considered lower bounds for arbitrary group-correlation. That is, we showed that there exists a group structure such that no-regret learning can only guarantee an approximation factor in $\Omega(n)$.

We proved our positive results under the assumption that availabilities of links are only group-correlated and some of the results rely heavily on the structure of groups. In contrast, the lower bound in Theorem 4.4 heavily relies on correlation and a non-laminar group-correlation structure.

We have argued that one can utilize no-external-regret learning to also guarantee low availability-ordering regret. As the considered approach in general increases the number of actions exponentially, it is an interesting open question whether one can find better no-availability-ordering-regret algorithms.

It is another interesting open problem to characterize the influence of correlation on the performance of no-regret learning algorithms beyond the considered laminar group-correlation. Correlation among channel availabilities of different links can for example result from a locality structure of primary and secondary users. One might even be able to adjust the considered learning algorithms to tackle stronger correlation among the availabilities of different links efficiently.

CHAPTER 5

Power Control for Capacity Maximization

One of the most important aspects of capacity maximization clearly is to decide which links should transmit as considered in the previous chapters. It is also of great importance to decide which power assignment to use.

Power control allows links to change their transmission power, and thus also the interference affecting other links. Choosing lower powers still allowing successful transmissions reduces the interference other links experience and thus it can possibly allow successful transmissions for additional links. Raising powers, on the other hand, increases interference on other links. It can also turn a previously unsuccessful link into a successful one. This way, utilizing power control can efficiently increase the overall capacity and performance of the network.

In fact, modern wireless devices have the ability to adjust their transmission power and hence are able to do power control. These wireless devices need to coordinate the choice of powers throughout the network. In this chapter, we will again consider distributed protocols, as wireless networks are inherently distributed.

One such distributed algorithm for power control in wireless networks was proposed by Foschini and Miljanich [FM92]. They update powers based on the SINR achieved in previous transmission attempts. We will consider this principle and a second one based on no-regret learning. The second approach can be considered to be more accessible than the first one. It uses only the little feedback also considered in the previous chapters. That is, a link experiences whether previous transmission attempts were successful or not. The approach of Foschini and Miljanich actually needs links to know the reached SINR of transmissions.

5.1 Results

For the *Foschini-Miljanich iteration* (FM iteration) we will give an analytical bound on the time, which is needed for the chosen powers to be close to the optimal choice. That is, after $O(n \log n)$ rounds a power assignment is reached satisfying the SINR constraint by at least a factor of $1 - \delta$. Nevertheless, we can argue that there is no guarantee for any transmission to be successful at all when using the FM iteration.

Using this bound we will analyze the behavior and convergence time of an approach based on no-regret learning, where links only need to know whether previous transmissions were successful. To do so, we consider a different kind of regret than in the previous chapters, namely no-swap regret. In contrast to the FM iteration, we can even give a guarantee that there are successful transmissions.

5.2 Power-Control Model

To consider power control for capacity maximization, we resort to the SINR model instead of the more general conflict-graph model used in Chapters 2, 3 and 4. In contrast to the conflict-graph model, the SINR model provides a notion on how different powers affect transmissions and thus gives a model to directly consider power control. This allows us to consider the powers needed for transmissions to be successful.

Recall that in the SINR model a transmission of link ℓ_v is successful if the SINR constraint is fulfilled, which is

$$g_{v,v} \cdot p_v \geq \beta_v \cdot \left(\sum_{u \neq v} g_{v,u} \cdot p_u + \nu \right).$$

Here the sender of link ℓ_v emits a signal at power p_u , $g_{v,u}$ is the *gain* as defined in Section 1.1.1, and ν denotes the ambient noise. Remember, that we can consider $g_{v,u} = \frac{1}{d(s_u, r_v)^\alpha}$ depending on the distance from the sender s_u to receiver r_v for a path-loss exponent α . Thus, for a successful transmission of link ℓ_v the SINR needs to be above some threshold β_v .

In the power-control problem, the task is to compute a feasible power assignment such that the SINR constraint is fulfilled for each link. Furthermore, each link should use the minimal possible power. This ensures that no link introduces interference, which is superfluous for its successful transmission.

More formally, we consider the gains in terms of a *normalized gain matrix* C being the $n \times n$ matrix defined by $C_{v,v} = 0$ for all $v \in [n] := \{1, \dots, n\}$ and $C_{v,u} = \beta_v \frac{g_{v,u}}{g_{v,v}}$ for $u \neq v$. In the same spirit, the *normalized noise vector* η is defined by $\eta_v = \beta_v \frac{\nu}{g_{v,v}}$. In regard to these notions, a *feasible power assignment* can be defined to be a vector $p \in \mathbb{R}_{\geq 0}$ such that $p \geq C \cdot p + \eta$. Note that throughout this chapter, we use the comparators $\leq$ and $\geq$ to denote the respective component-wise inequality.

The set of all feasible power vectors in this setting is a convex polytope. If it is non-empty, there is a unique vector p^* satisfying $p^* = C \cdot p^* + \eta$. In a centralized setting, where the algorithm has complete knowledge, this fixed point p^* can simply be computed by solving the linear equation system $p^* = C \cdot p^* + \eta$.

Wireless networks inherently consist of independent senders, which are not cooperative. Thus, wireless networks afford to consider distributed protocols and the matrix C to be unknown. In our first scenario analyzing an iterative approach by Foschini and Miljanich [FM92] we assume the senders to be able to make communication attempts at different powers. Senders receive feedback in the form of the achieved SINR. In an advanced scenario with regret learning only the feedback whether the transmission has been successful or not is received.

We consider this scenario as a normal-form game. Each link (or sender) ℓ_v is a *selfish agent* – also denoted as a *user* $i \in [n]$ controlling the transmission behavior – and picks a transmission power as its strategy. We first assume that the achieved SINR becomes known after each

transmission attempt. Furthermore, we assume the user i to have a utility $u_i(p_i, p_{-i})$ having a unique maximum at the power level that sets his SINR to exactly β_v , but can otherwise be completely arbitrary. Here p_{-i} denotes the powers chosen by all other users but user i .

In this context, the FM iteration constitutes a sequence of *best responses* on the actions chosen by other users. This means a player always chooses the action that was best in the previous round.

5.3 Foschini-Miljanic Iteration

As a first distributed algorithm for power control, we consider an iterative approach due to Foschini and Miljanic [FM92]. It requires that the senders get to know their achieved SINR. The basic idea of this *Foschini-Miljanic (FM) iteration* is that any sender chooses exactly the power level that would have been necessary to obtain a successful transmission in the previous time step.

Formally, the FM iteration is $p^{(t+1)} = C \cdot p^{(t)} + \eta$ where t is the time step. In the context of games played iteratively, it implements concurrent best-response dynamics. The achieved and the target SINR is actually needed to run this iteration. Later, in Section 5.4, we will omit this assumption and describe an approach, which does not need the SINR to be known.

Foschini and Miljanic showed that the sequence of vectors $p^{(t)}$ actually converges to p^* as t goes to infinity. Obviously, p^* is a stable choice in the sense that no sender wants to increase or decrease its power. That is, it is a Nash equilibrium and thus no sender wants to unilaterally change its transmission power.

In our analyses, we refer to the maximal absolute value of an eigenvalue of the matrix C as $\lambda_{\max}$. One can show that the existence of $p^* \in \mathbb{R}_{\geq 0}^n$ with $p^* \geq C \cdot p^* + \eta$ implies that the absolute value of all eigenvalues of C is strictly less than 1. Thus, it holds $\lambda_{\max} < 1$.

In our analysis, it will turn out to be helpful to consider the closed-form variant of the FM iteration given by

$$p^{(t)} = C^t \cdot p^{(0)} + \sum_{k=0}^{t-1} C^k \eta . \quad (5.1)$$

Possible runs of this best-response iteration are shown in Figure 5.1 and Figure 5.2 below.

An exemplary run of the FM iteration with two links is depicted in Figure 5.1. Points in the light grey regions depict the power choices where at least one link can transmit successfully. The darker region consists of power vectors rendering both links successful. It shows that the actual run of the FM iteration highly depends on which starting point is chosen. Here, the iteration allows both of the two links being alternately successful. Also, the convergence towards the fixed point is clearly visible.

In contrast to the run depicted in Figure 5.1, we can observe one of the major drawbacks of the FM iteration in Figure 5.2. The runs only differ in using a different starting point of the iteration. The algorithm again converges to the fixed point, but this time both links stay within the white region being unsuccessful throughout the whole sequence.

As seen, the fixed-point approach of Foschini and Miljanic has some drawbacks. For example, in sequences starting from 0 the target SINR is never reached as shown in Figure 5.2, because all links increase their powers in each step and therefore the powers are always too small.




Figure 5.1: FM iteration in an example with two links.

Another drawback of the FM iteration is that in order to adapt the power correctly, the currently achieved SINR has to be known. A last disadvantage to be mentioned is its lack of robustness. We assume a fixed-point pure Nash equilibrium to exist for the power control game. If for some reason this does not hold, the iteration might end up where some powers are 0 or $p^{\max}$, even if the transmission is not successful.

Besides these drawbacks, with the assumption that the achieved SINR is known, the FM iteration is guaranteed to converge to the optimal choice of powers. In the following section, we will consider the time it needs to converge close to the fixed point.

Note that, we extend the basic scenario analyzed by Foschini and Miljanic also by considering different SINR thresholds β_v for each link ℓ_v .

5.3.1 Convergence Time

As discussed above, it is clear that in general the FM iteration never actually reaches the fixed point. As the iteration converges to it, the SINR converges to the threshold β_v . That is, for each $\delta > 0$ there is some round T after which the SINR is never below $(1 - \delta)\beta_v$. Since having all links reach their SINR target is the main goal, we strive to bound the time T until each transmission is “almost” feasible.

To analyze the convergence time of the FM iteration in this sense, it is sufficient that the current vector of transmission powers p satisfies $(1 - \delta)p^* \leq p \leq (1 + \delta)p^*$.

As a first result, we bound the convergence time in terms of n when starting from 0. We will see that the bound proven here is independent of the values of p^* or v . The only parameter related to the instance is the maximum eigenvalue of C denoted by $\lambda_{\max}$. This parameter related to the instance has to occur in the convergence time, as for $\lambda_{\max} = 1$ no fixed point can exist at all. Assuming it to be constant, the theorem below directly implies that after $O(n \log n)$ rounds a power assignment is reached that satisfies the SINR constraint of every link by a factor of at least $1 - \delta$ for a constant δ .




Figure 5.2: Example FM iteration with two users being always unsuccessful.

Theorem 5.1. Starting from $p^{(0)} = 0$ after $t \geq \frac{\log \delta}{\log \lambda_{\max}} \cdot n \cdot \log(3n)$ rounds, for all $p^{(t)}$ we have $(1 - \delta)p^* \leq p^{(t)} \leq p^*$.

Proof. For our considerations we combine multiple time steps by defining a matrix M such that the eigenvalues can be bounded easily. This can then be used to consider the number of time steps.

Let $M = C^m$ be an auxiliary matrix, where

$$m = \left\lceil \frac{\log \frac{1}{3n}}{\log \lambda_{\max}} \right\rceil.$$

Note that the absolute value of all eigenvalues of M is bounded by $\frac{1}{3n}$. Furthermore, we define $\eta' = \sum_{k=0}^{m-1} C^k \eta$ and thus we have $p^{(m \cdot t')} = \sum_{k=0}^{t'-1} M^k \eta'$ for the sequence of power vectors starting with $p^{(0)} = 0$. The choice of M also implies $p^* = \sum_{k=0}^{\infty} M^k \eta'$.

Now we consider the characteristic polynomial of M in expanded form as well as in factorized form:

$$\chi_M(x) = x^n + \sum_{i=0}^{n-1} a_i x^i = \prod_{j=1}^n (x - b_j).$$

The (possibly complex) b_j values correspond to the eigenvalues. Therefore, with eigenvalues of M being at most $\frac{1}{3n}$ we have $|b_j| \leq \frac{1}{3n}$ for all $j \in [n]$. This way, the absolute value of the a_i values can be expressed in terms of the b_i values by

$$|a_i| \leq \sum_{\substack{S \subseteq [n] \\ |S|=n-i}} \prod_{j \in S} |b_j| \leq \binom{n}{n-i} \left(\frac{1}{3n} \right)^{n-i}.$$

This yields the following bound for their sum

$$\sum_{i=0}^{n-1} |a_i| \leq \sum_{k=0}^n \binom{n}{k} \left(\frac{1}{3n} \right)^k - 1 = \left(1 + \frac{1}{3n} \right)^n - 1 \leq \frac{1}{2}.$$

We now use the fact that by definition $\chi_M(M) = 0$. This is, $M^n = -\sum_{i=0}^{n-1} a_i M^i$. Since all $M^k \eta'$ are non-negative, the following inequality holds

$$\begin{aligned} M^n p^* &= \sum_{k=n}^{\infty} M^k \eta' \\ &= \sum_{k=0}^{n-1} M^k \eta' \left(-\sum_{i=0}^k a_i \right) + \sum_{k=n}^{\infty} M^k \eta' \left(-\sum_{i=0}^{n-1} a_i \right) \\ &\leq \left(\sum_{i=0}^{n-1} |a_i| \right) \sum_{k=0}^{\infty} M^k \eta' \leq \frac{1}{2} \sum_{k=0}^{\infty} M^k \eta' \\ &= \frac{1}{2} p^* . \end{aligned}$$

Now consider $t \geq m \cdot n \cdot \log \frac{1}{\delta}$. We have $p^* - p^{(t)} = \sum_{k=0}^{\infty} C^k \eta - \sum_{k=0}^{t-1} C^k \eta = \sum_{k=t}^{\infty} C^k \eta$ using the closed form of the FM iteration given in Equation 5.1. Further calculus results in $C^t \cdot \sum_{k=0}^{\infty} C^k \eta = C^t p^* \leq M^{n \log \frac{1}{\delta}} p^* \leq \delta p^*$ proving the theorem. $\square$

Theorem 5.2 provides an upper bound on the time needed until the FM iteration is close to the fixed point. To find a lower bound, we argue that there are instances where $\Omega(n)$ rounds are needed to be arbitrarily close to the fixed point or to even reach it. For example, let C be defined by $C_{i+1,i} = 1$ for all i and $C_{j,i} = 0$ for all other entries, $\eta = (1, 0, \dots, 0)$. The only eigenvalue of this matrix is 0. However, it takes n rounds until the first component 1 has propagated to the n -th component and the fixed point is reached.

Obviously, such instances require a certain construction and structure in the values of p^* and η . Thus, our second result is a bound independent of n holding for every starting point $p^{(0)}$. This bound instead takes p^* and η into consideration.

Theorem 5.2. *Starting from an arbitrary $p^{(0)}$, we have $(1 - \delta)p^* \leq p^{(t)} \leq (1 + \delta)p^*$ for all $t \geq T$ with*

$$T = \frac{\log \delta - \log \max_{i \in [n]} \left| \frac{p_i^{(0)}}{p_i^*} - 1 \right|}{\log \max_{i \in [n]} \left| 1 - \frac{\eta_i}{p_i^*} \right|} .$$

Proof. We consider the weighted maximum norm using the entries of p^* as weights defined as

$$\|x\| = \max_{i \in [n]} \frac{|x_i|}{p_i^*} ,$$

which has also been used by Huang and Yates [HY98] before. The induced matrix norm of a matrix M is now given by

$$\|M\| = \max_{i \in [n]} \frac{1}{p_i^*} \sum_{j \in [n]} |M_{i,j}| \cdot p_j^* .$$

In particular, as p^* is the fixed point of the FM iteration, we have for matrix C that $Cp^* + \eta = p^*$. This directly shows for one component of p^* that $(Cp^*)_i = p_i^* - \eta_i$. The matrix norm of C can then be found by

$$\|C\| = \max_{i \in [n]} \frac{1}{p_i^*} \sum_{j \in [n]} C_{i,j} p_j^* = \max_{i \in [n]} \frac{1}{p_i^*} (Cp^*)_i = \max_{i \in [n]} \left| 1 - \frac{\eta_i}{p_i^*} \right| .$$

If $t \geq T$, using the closed form of the iteration (Equation 5.1) we get

$$\begin{aligned}\|p^{(t)} - p^*\| &= \|C^t \cdot (p^{(0)} - p^*)\| \\ &\leq \|C\|^t \cdot \|p^{(0)} - p^*\| .\end{aligned}$$

Applying the weighted maximum norm for vectors and for the matrix C yields for $t \geq T$

$$\begin{aligned}\|C\|^t \cdot \|p^{(t)} - p^*\| &\leq \|C\|^t \cdot \max_{i \in [n]} \left| \frac{p_i^{(0)}}{p_i^*} - 1 \right| \\ &\leq \frac{\delta}{\max_{i \in [n]} \left| \frac{p_i^{(0)}}{p_i^*} - 1 \right|} \cdot \max_{i \in [n]} \left| \frac{p_i^{(0)}}{p_i^*} - 1 \right| = \delta .\end{aligned}$$

So, for all $i \in [n]$, we have $|p_i^{(t)} - p_i^*| \leq \delta \cdot p_i^*$ and get $(1 - \delta)p_i^* \leq p_i^{(t)} \leq (1 + \delta)p_i^*$. $\square$

Note that the original FM iteration assumes arbitrarily high powers, while in the following sections we assume that a maximum power exists. Our proofs on the convergence time above considered the original FM iteration. However, our bounds directly transfer to the case where there is some vector of maximum powers $p^{\max}$. In this setting, all powers are projected to the respective interval $[0, p_i^{\max}]$ in each round in the same fashion as done by Yates [Yat95]. One can see that this can only have a positive effect on the convergence time since the resulting sequence is component-wise dominated by the sequence on unlimited powers. Thus, the same bounds still hold under this assumption.

5.4 Utilizing No-Regret Learning

Recall that the fixed-point approach by Foschini and Miljanic analyzed above has some major drawbacks. For example, in many sequences – in particular the ones starting from 0 – the target SINR is never reached. Other disadvantages of the FM iteration already discussed in Section 5.3 are that the achieved SINR has to be known and that the approach is not very robust. In order to overcome these, we design different regret-based learning algorithms. This approach uses the notion of no-swap-regret as defined below. As these algorithms are randomized, each player can transmit successfully with positive probability already in the first time steps. This is in contrast to the FM iteration with possibly no link becoming successful at all.

Besides this advantage, we also use the no-swap-regret learning algorithm to overcome the other drawbacks of the FM iteration. Using this kind of algorithm, there is no need to know the achieved SINR, and it is robust against scenarios where no fixed point exists.

While no-external-regret worked quite well in Chapter 2 and no-ordering-regret in Chapter 4, those notions are not sufficient to guarantee convergence to the optimal power assignment or fixed point respectively as in the FM iteration.

Proposition 5.1 below formally states that no-external-regret is not sufficient to converge to the fixed-point power assignment. The proof even constitutes that zero external regret does not suffice. Although there is a fixed point p^* at which all links are successful, a no-external regret sequence might make only 2 of the n links successful at all.

To apply regret learning, we assume each user to be again aiming at having successful transmissions. Additionally, each user strives to use the least power possible that still achieves successful transmissions. For the decision which power $p_i \in [0, p_i^{\max}]$ to use we define a suitable *utility function* $u_i(p): [0, p_{\max}] \rightarrow [0, 1]$. On the one hand, it is quite natural to have a utility of 0 when the chosen power does not suffice to guarantee a successful transmission, i.e., the SINR is below the threshold. On the other hand, given that the transmission is successful, it is also clear that it is beneficial to choose a power as small as possible. This choice obviously reduces the interference on other links. Each user's maximum utility is at the point where the SINR condition is exactly met. This way, best-response dynamics corresponds to the FM iteration. These properties cannot be modeled by a continuous function.

Formally, we capture this by the following utility function.

$$u_i(p) = \begin{cases} f_i(p_i) & \text{if user } i \text{ is successful with } p_i \text{ against } p_{-i} \\ 0 & \text{otherwise} \end{cases},$$

where $f_i: [0, p_i^{\max}] \rightarrow (0, 1]$ is a continuous and strictly decreasing function for each $i \in [n]$. With p_{-i} we denote the powers chosen by all users but user i .

Proposition 5.1. *For suitable functions f_i , there is an instance with a fixed point and a no-external-regret sequence in which only a $\frac{2}{n}$ fraction of all links is ever feasible.*

Proof. We consider a “nested pairs” instance from [FKRV09] appropriately scaled with suitable chosen SINR parameters to allow a fixed point. In this scenario, there are n bidirectional links placed on a line where one end of link ℓ_i lies at -2^i and the other one at 2^i . Fanghänel et al. argue that the square root power assignment renders at least a constant fraction of links successful. Thus, scaling yields an instance with a fixed point. Here, we replace the innermost link by two smaller links of about half the length such that the instance still has a fixed point and their distance is chosen appropriately. The maximum power is $p^{\max}$ with $p_i^{\max} = p^{\max}$ for all links i .

Now we consider the following sequence of actions. All links except for the two inner ones always choose not to transmit by playing action $p_i = 0$. In odd steps, the two inner links play action $p^{\max}$, in even steps they play $\frac{p^{\max}}{2}$. The regret for each player is at most 0. For the outer links this is clear as they cannot transmit successfully at all.

The inner links have some smallest action p' allowing them to be feasible in all steps. We have $\frac{p^{\max}}{2} \leq p' \leq p^{\max}$. The regret compared to this action after T steps is $f(p')T - \frac{1}{2}(f(p^{\max}) + f(\frac{1}{2}p^{\max}))T$ as the utility for choosing p' would be $f(p')T$. The utility actually gained by the played sequence is $\frac{1}{2}(f(p^{\max}) - f(\frac{1}{2}p^{\max}))$.

Note that the regret can evaluate to 0 (or less) by a suitable choice of f . This proves the proposition. $\square$

As no-external regret cannot be used to successfully solve power control in our scenario, we will utilize a more intricate notion of regret being no-swap regret. We will define swap regret in the remainder of this section.

Similar to comparing the utility of an action sequence to one single best action in hindsight, we could also consider another concept. Here, we compare the utility gained by using a certain action sequence to an action sequence, where any action is replaced by a specific other one.

One can use the notion of Φ -regret to define swap regret by choosing a set Φ of departure measurable functions as already introduced in Section 1.4. The notion of Φ -regret indeed is more general than swap regret and also encompasses no-external-regret and other further notions of regret.

In terms of Φ -regret, the best action in hindsight – used to compare the algorithm’s action sequence to – is defined by the supremum over all measurable functions. That way, different actions are possibly mapped to different other actions. More intuitively, for each action we can consider to exchange this action by a specific other one whenever it is used.

Besides using Φ -regret with measurable functions, we can also formally define swap regret as follows. This definition gives a more intuitive insight into swap regret by defining a kind of “best action in hindsight” similar to the ones used for no-external-regret before.

Definition 5.1. Let $a^{(1)}, \dots, a^{(T)}$ be a sequence of action vectors. The swap regret of this sequence for user i is defined by

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

where $\mathcal{A}$ denotes the set of possible actions. An algorithm has the no-swap-regret property if the swap regret of the computed sequence of actions is in $o(T)$.

For power control the set of actions $\mathcal{A} = [0, p_{\max}]$. In total, we will show that all no-swap-regret sequences converge to the optimal power vector $p^* \in [0, p_{\max}]$. Furthermore, the fraction of successful transmissions converges to 1. This is in contrast to the FM iteration, where starting from $p^{(0)} = 0$ all transmissions stay unsuccessful during the entire iteration.

An excerpt of a no-swap-regret sequence is shown in Figure 5.3. It shows the run of no-swap-regret learning of two links in the same setting as depicted in Figure 5.1 using the FM iteration. The light grey regions again depict the possible power choices where at least one link can transmit successfully. The darker region refers to states where both links are successful. Considering the chosen powers from time step 900 onwards, one can already see that the power vectors converge to the optimal solution. It is visible that the no-swap-regret sequence provides successful transmissions and that the power choices concentrate around the fixed point. Here, even no transmission is unsuccessful for both links simultaneously. Though, this is not guaranteed it shows a clear strength of the adaptive learning approach. Over time the sequence becomes more concentrated around the fixed point.

There exist several algorithms obtaining the no-swap-regret property. While those often assume the action space to be discrete, we will show in Section 5.5 how to apply them to our setting where the space of possible powers is continuous. One such approach is due to Blum and Mansour [BM07], which we will actually apply in Corollary 5.1. Blum and Mansour show how to utilize multiple no-external-regret algorithms to ensure no-swap-regret for the combined algorithm. Other approaches yielding no-swap regret are for example presented in [SL07, Zin03, AHK12].

Before considering how a no-swap-regret algorithm converges to the fixed point in a numerical sense, we can see that for an algorithm having swap-regret zero the sequence must only consist of the fixed point p^* .

Proposition 5.2. Given any sequence $p^{(1)}, \dots, p^{(T)}$ with the swap regret for each user being 0, then $p^{(t)} = p^*$ for all t .




Figure 5.3: Last 100 of 1000 iterations of no-swap-regret learning.

Proof. For each user i let $\hat{p}_i = \max_{t \in T} p_i^{(t)}$, meaning $\hat{p}_i$ is the maximal power level ever chosen by i . We show now that this maximal power level is not greater than p_i^* and that it is not smaller than p_i^* . Note that comparators on power vectors below are meant in terms of a partial ordering based on an element-wise comparison. That is, it holds $p' \geq p''$ if $p'_i \geq p''_i$ for all i .

Assume that $\hat{p} \leq p^*$ does not hold. This means there is some user which can still transmit successfully using less power. That is, a user i exists, for which $\hat{p}'_i := (C \cdot \hat{p} + v)_i < \hat{p}_i$. This user obviously has non-zero swap regret because using $\hat{p}'_i$ instead of $\hat{p}_i$ lowers his power while the transmission still stays successful, which clearly increases the utility. By contradiction we have $\hat{p} \leq p^*$.

Now let $\check{p}_i = \min_{t \in T} p_i^{(t)}$. We can argue analogously as before to get $\check{p}_i \geq p_i^*$. Thus, assuming this does not hold, we again have some user i for which $\check{p}_i < (C \cdot \check{p} + v)_i$. Thus, user i is never successful when using power $\check{p}_i$, but it would always be with p_i^* . This again results in a higher utility and user i would encounter a non-zero swap regret.

In total, we have that both $\hat{p} \leq p^*$ and $\check{p} \geq p^*$, yielding $\hat{p} = \check{p} = p^*$. $\square$

While this result shows that an action sequence with zero swap-regret can only consist of the fixed point, it is unclear how such a sequence can actually be computed. This computation should be done in a distributed fashion. We also argued previously that one or probably the most important reason for using no-regret learning algorithms is that they can be applied in a distributed fashion.

Until now, algorithms for no-swap regret focused on finite strategy spaces or convex utility functions. For power control we here assume an infinite set of strategies $[0, p_{\max}]$ and a non-convex utility function. To apply the concept of no-swap regret, we adjust previous algorithms in the following section before we can analyze how these converge afterwards.

5.5 Computing No-Swap-Regret Sequences

To apply no-swap-regret learning to power control, we have to move from the case of finite action spaces considered mostly in literature to infinite action spaces. For example, for the case of N actions, the algorithm devised by Blum and Mansour [BM07] is able to guarantee that after T rounds the expected regret of a user is at most $O(\sqrt{TN \log N})$. This algorithm is randomized and uses multiplicative weights updates. That is, each user has a probability distribution over all possible actions. After each step, the user updates this distribution based on the previously observed utilities.

Blum and Mansour actually give a black-box reduction from external to swap regret. Their algorithm uses one no-external-regret algorithm for each action. The probability distribution of those algorithms are used to get a combined probability distribution to decide on the action to choose. Feedback in terms of utility or loss is then distributed back to the algorithms. Arbitrary no-external-regret algorithms like RWM proposed in Section 2.2 can be used in their construction.

Unfortunately, this and similar algorithms [AHK12] require a finite number of actions, while in our case the action space contains all real numbers within $[0, p_i^{\max}]$. Standard approaches for infinite action spaces are not applicable either, as they require convex action spaces and continuous utility functions [Zin03, SL07]. In order to capture the SINR threshold appropriately, however, the utility functions have to be modeled as non-continuous.

In this section, we show that no-regret sequences can be computed in a distributed way in this power-control setting nevertheless. Although our utility function is not convex, we use its specific structure to construct a no-swap-regret algorithm. This is done by applying no-swap-regret algorithms for finite action spaces on a suitable finite subset of the powers. This finite subset is constructed by dividing the set of powers into intervals of equal length and using the right borders as the input action set for the algorithm. This discretization, however, is not chosen in a fixed way but is iteratively refined to guarantee that the no-swap-regret property holds.

Theorem 5.3. *Let $\mathcal{A}$ be any no-swap-regret algorithm for arbitrary finite action spaces, whose swap regret after T rounds in case of N actions is at most $O(T^a \cdot N^b)$, where a and b are suitable constants with $0 \leq a < 1$, $b \geq 0$.*

Then $\mathcal{A}$ can be used to construct an algorithm for power control on infinite action spaces achieving swap regret at most $O\left(T^{\frac{a+b}{1+b}}\right)$.

Proof. We exploit the structure of our utility functions to use no-swap-regret algorithms for finite actions spaces in this setting. Consider the utility function $u_i(\cdot, p_{-i})$ of some user i provided that the other strategies are fixed.

We cut the set of strategies $[0, p_i^{\max}]$ into N intervals of equal length. Let the maximal decrease of the utility function be $S_i = \max_{p_i, h} \frac{f(p_i) - f(p_i + h)}{h}$. It is easy to observe that for any interval the utility at the right border is at most $S_i \frac{p_i^{\max}}{N}$ worse than the maximum in the respective interval. Thus, for all $x \in \left[k \frac{p_i^{\max}}{N}, (k+1) \frac{p_i^{\max}}{N}\right]$, we have

$$u_i(x, p_{-i}) \leq u_i\left((k+1) \frac{p_i^{\max}}{N}, p_{-i}\right) + S_i \frac{p_i^{\max}}{N}.$$

If the number of steps T is known in advance, we can construct the following algorithm by choosing the finite action space accordingly. We set

$$N = \left\lceil T^{\frac{1-a}{1+b}} \right\rceil$$

and run $\mathcal{A}$ using only the finite action set $\left\{ \frac{p_i^{\max}}{N}, 2\frac{p_i^{\max}}{N}, \dots, p_i^{\max} \right\}$ of size N . If the optimal strategies were restricted to exactly the same set, the resulting swap regret would be at most $O(T^a N^b)$. Due to the restriction to the finite set, the actual actions considered in hindsight in the finite action space differ by at most $S_i \frac{p_i^{\max}}{N}$ in each step to the ones in the infinite action space. So the resulting regret is at most

$$O(T^a N^b) + TS_i \frac{p_i^{\max}}{N} = O\left(T^{\frac{a+b}{1+b}}\right).$$

If T is not known in advance, the “doubling trick” also works for our algorithm. That is, starting with an estimate $T = 1$, the algorithm is executed for T steps with the respective estimate T , which is doubled afterwards. Thus, up to a factor of 2 the correct T is assumed in $\frac{1}{2}$ of the time steps. This way only a constant factor is lost in comparison to the case where the exact T is known. $\square$

Theorem 5.3 provides a framework to use different existing no-swap-regret learning algorithms. Having any algorithm with swap regret being bounded polynomial in T and N we can construct an algorithm with a suitable swap-regret bound. Depending on the specific algorithm used, it yields different regret bounds.

If, after each time step, each link knows which powers would have made it successful, it can calculate the utility for all possible power levels. That affords links to know the reached SINR. Thus, a full-information algorithm using feedback on all actions and not only the chosen ones can be used. In particular, we can use the $O(\sqrt{TN \log N})$ swap-regret full-information algorithm proposed by Blum and Mansour [BM07] for the following result.

Corollary 5.1. *There is a full-information algorithm achieving swap regret $O\left(T^{\frac{3}{4}}\right)$ in expectation over the random choices of the algorithm.*

Considering the partial-information model instead, getting to know the utility only for chosen actions, it is sufficient for links to only get to know if the transmission at the actually chosen power was successful. The value of the utility function for the chosen power can still be computed. Therefore, in this case we can use a partial-feedback algorithm. Here, we can apply the $O(N\sqrt{T \log N})$ algorithm by Blum and Mansour [BM07] to Theorem 5.3 and build an algorithm yielding the following result.

Corollary 5.2. *There is a partial-information algorithm achieving swap regret $O\left(T^{\frac{4}{5}}\right)$ in expectation over the random choices of the algorithm.*

As a matter of fact, discretization turns out to be helpful in this case to apply no-swap-regret algorithms. Nevertheless, one can see that assuming discretized powers in the first place might imply that both the FM iteration and no-swap-regret learning do not work well.

Consider the following simple example of two players

$$C = \begin{pmatrix} 0 & \frac{1}{2} \\ \frac{1}{2} & 0 \end{pmatrix} \quad \eta = \begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix}.$$

In the case of continuous action spaces, the unique fixed point is $(1, 1)$. However, if only natural numbers can be chosen as powers this is different.

In the natural adaptation of the Foschini-Miljanic iteration each transmission still uses the minimal power necessary to compensate the interference, i.e. $p_i^{(t+1)} = \lceil Cp^{(t)} + \eta \rceil$. In our example we do not only get two fixed points – $(1, 1)$ and $(2, 2)$ – but also oscillating sequences – $(1, 2), (2, 1), (1, 2), \dots$ – as outcomes of the FM iteration. Having users choose the powers given in any fixed point results in a zero-swap-regret sequence. The second fixed point is therefore not optimal, and the sequence does not converge to the optimum.

The simple transformation for algorithms yielding no-swap regret under finite action spaces to our infinite action space shows how to utilize existing algorithms in our setting. These algorithms converge to the optimum in a reasonable way as we will show in the next section.

5.6 Convergence of No-Swap-Regret Sequences

We have seen that algorithms with zero swap regret actually have to follow the optimal strategy in Proposition 5.2. We have also seen how to compute no-swap-regret sequences. Now we will complement these results by a quantitative analysis of how the no-swap-regret sequence converges.

We show that not only convergence to the optimal power vector p^* is guaranteed but also the fraction of rounds in which each link is successful converges to 1. This shows one clear advantage of no-regret learning to the FM iteration. In the FM iteration there are starting vectors such that no transmission is ever successful at all.

Theorem 5.4. *For every sequence $p^{(1)}, \dots, p^{(T)}$ with swap regret at most $\epsilon \cdot T$ for every user i and for every $\delta > 0$ the fraction of steps in which user i sends successfully is at least*

$$Q \cdot \frac{f_i((1 + \delta)p_i^*)}{f_i((1 - \delta)p_i^*)} - \frac{\epsilon}{f_i((1 - \delta)p_i^*)} ,$$

where Q denotes the fraction of rounds in which a power vector p with $(1 - \delta)p^* \leq p \leq (1 + \delta)p^*$ is chosen.

Given a sequence with swap regret at most $\epsilon \cdot T$, Theorem 5.4 gives a lower bound for the number of steps in which a user sends successfully. The bound depends on the utility function and the fraction of rounds in which a power vector between $(1 - \delta)p^*$ and $(1 + \delta)p^*$ is chosen. This fraction is bound in Lemma 5.1 and Lemma 5.2 later on. Together, this then provides a bound on the fraction of successful transmissions of a link converging to 1 as the swap regret converges to 0.

In order to prove this theorem, we use the fact that no-swap-regret sequences correspond to *approximate equilibria*. Similar to a mixed Nash equilibrium, an ϵ -correlated equilibrium is a probability distribution over strategy vectors (in our case power vectors) such that no user can unilaterally increase his expected utility by more than ϵ . This notion will be helpful in our considerations and allows us to rewrite the theorem. Thus, we are able to argue in terms of probability mass instead of discrete time steps and chosen actions.

In contrast to mixed Nash equilibria the choices of the different users do not need to be independent and this correlation mirrors the correlation naturally given by no-swap-regret learning. Formally, an ϵ -correlated equilibrium is defined as follows.

Definition 5.2. An ϵ -correlated equilibrium is a joint probability distribution π over the set of power vectors $P_1 \times \dots \times P_n$, where $P_i = [0, p_i^{\max}]$, such that for any user i and measurable function $\phi_i: P_i \rightarrow P_i$, we have

$$\mathbf{E}_{p \sim \pi}[u_i(\phi_i(p_i), p_{-i})] - \mathbf{E}_{p \sim \pi}[u_i(p_i, p_{-i})] \leq \epsilon .$$

In an ϵ -correlated equilibrium, no user can increase his expected utility by more than ϵ when exchanging any action with a fixed other one. More intuitively speaking, no user would consider operations such as “each time π says I play A , I play B instead” (for any A) to increase his utility. This kind of operations are exactly the ones considered in Definition 5.1 of no-swap-regret sequences. Therefore each sequence $p^{(1)}, \dots, p^{(T)}$ of swap regret at most R corresponds to an $(\frac{R}{T})$ -correlated equilibrium.

Using this notion, we can rewrite Theorem 5.4 to the following proposition stating the same claim in terms of ϵ -correlated equilibria.

Proposition 5.3. For every ϵ -correlated equilibrium π and for every $\delta > 0$ the probability that user i sends successfully is at least

$$Q \cdot \frac{f_i((1+\delta)p_i^*)}{f_i((1-\delta)p_i^*)} - \frac{\epsilon}{f_i((1-\delta)p_i^*)} ,$$

where $Q = \mathbf{Pr}_{p \sim \pi}[(1-\delta)p^* \leq p \leq (1+\delta)p^*]$.

Proof. Having an ϵ -correlated equilibrium π we are guaranteed that no user would gain utility more than ϵ or higher by switching any of his actions. Thus, we consider a switching operation and use the bound of ϵ on the possible utility gain.

Instead of the powers in the interval $[(1-\delta)p_i^*, (1+\delta)p_i^*]$ user i could always choose $(1+\delta)p_i^*$. Since π is an ϵ -correlated equilibrium, this operation can increase the expected utility by at most ϵ . We now bound the change of the expected utility due to this switching operation.

Let $\mathcal{E}$ be the event that a vector p is chosen with $p_i \in [(1-\delta)p_i^*, (1+\delta)p_i^*]$ then the expected utility gain is

$$\mathbf{E}_{p \sim \pi}[u_i((1+\delta)p_i^*, p_{-i}) \mid \mathcal{E}] - \mathbf{E}_{p \sim \pi}[u_i(p) \mid \mathcal{E}] \leq \frac{\epsilon}{\mathbf{Pr}_{p \sim \pi}[\mathcal{E}]} . \quad (5.2)$$

Now let us bound the two expectations in this sum.

Considering $\mathbf{E}_{p \sim \pi}[u_i((1+\delta)p_i^*, p_{-i}) \mid \mathcal{E}]$ we see that user i is always successful using the power $(1+\delta)p_i^*$ if the other users use a power vector $p_{-i} \leq (1+\delta)p_{-i}^*$. So when applying this switch operation, user i gets an expected utility conditioned on the event $\mathcal{E}$ of $\mathbf{E}_{p \sim \pi}[u_i((1+\delta)p_i^*, p_{-i}) \mid \mathcal{E}] \geq f_i((1+\delta)p_i^*) \cdot \mathbf{Pr}_{p \sim \pi}[p_{-i} \leq (1+\delta)p_{-i}^* \mid \mathcal{E}]$, which yields

$$\begin{aligned} & \mathbf{Pr}[\mathcal{E}] \cdot \mathbf{E}_{p \sim \pi}[u_i((1+\delta)p_i^*, p_{-i}) \mid \mathcal{E}] \\ & \geq f_i((1+\delta)p_i^*) \cdot \mathbf{Pr}_{p \sim \pi}[(1-\delta)p^* \leq p \leq (1+\delta)p^*] . \end{aligned} \quad (5.3)$$

On the other hand, due to non-increasing f_i we have

$$\mathbf{E}_{p \sim \pi}[u_i(p) \mid \mathcal{E}] \leq f_i((1-\delta)p_i^*) \cdot \mathbf{Pr}_{p \sim \pi}[\mathcal{S} \mid \mathcal{E}] ,$$

where $\mathcal{S}$ is the event that the transmission is successful. This yields

$$\mathbf{Pr}[\mathcal{E}] \cdot \mathbf{E}_{p \sim \pi}[u_i(p) \mid \mathcal{E}] \leq f_i((1-\delta)p_i^*) \cdot \mathbf{Pr}_{p \sim \pi}[\mathcal{S}] . \quad (5.4)$$

Combining Equations 5.2, 5.3, and 5.4, we can substitute the expectations in 5.2 and get

$$f_i((1 + \delta)p_i^*) \cdot \Pr_{p \sim \pi}[(1 - \delta)p^* \leq p \leq (1 + \delta)p^*] - f_i((1 - \delta)p_i^*) \cdot \Pr_{p \sim \pi}[\mathcal{S}] \leq \epsilon .$$

Rearranging and easy calculus results in

$$\Pr_{p \sim \pi}[\mathcal{S}] \geq \frac{1}{f_i((1 - \delta)p_i^*)} (f_i((1 + \delta)p_i^*) \cdot \Pr_{p \sim \pi}[(1 - \delta)p^* \leq p \leq (1 + \delta)p^*] - \epsilon) .$$

This proves the claim. $\square$

This proposition bounds the success probability of one link. The bound depends on the probability that a certain power vector is chosen. Thus, it remains to bound the probability $\Pr_{p \sim \pi}[(1 - \delta)p^* \leq p \leq (1 + \delta)p^*]$. For this purpose, we bound the probability mass of states p with $p \not\leq (1 + \delta)p^*$ in Lemma 5.1 and of the ones with $p \not\geq (1 - \delta)p^*$ in Lemma 5.2. Those bounds can then be combined to get the desired bound by

$$\begin{aligned} & \Pr_{p \sim \pi}[(1 - \delta)p^* \leq p \leq (1 + \delta)p^*] \\ &= 1 - \Pr_{p \sim \pi}[p \not\geq (1 - \delta)p^*] - \Pr_{p \sim \pi}[p \not\leq (1 + \delta)p^*] . \end{aligned}$$

The general proof ideas work as follows. In order to bound $\Pr_{p \sim \pi}[p \not\leq (1 + \delta)p^*]$, we consider which probability mass can at most be on vectors p such that for some user i , we have $p_i > (C \cdot p^{\max} + (1 + \frac{\delta}{2}) \cdot \nu)_i$. The probability mass is bounded, because user i could instead always use power $(C \cdot p^{\max} + \nu)_i$, as this is the maximum power needed to compensate the interference in the case that $p_{-i} = p_{-i}^{\max}$. We then proceed in a similar way always using the bound obtained before until we reach a point component-wise smaller than $(1 + \delta)p^*$. Thus, we iteratively “cut off” probability mass and, hence, virtually lower $p^{\max}$ while having to consider probability mass cut off before. We use the results concerning the FM iteration from Section 5.3. The bound on $\Pr_{p \sim \pi}[p \not\geq (1 - \delta)p^*]$ works in a similar way.

First, we consider the probability mass on $p \not\leq (1 + \delta)p^*$. Afterwards, we will do the same for $p \not\geq (1 - \delta)p^*$. For the proofs the following general observation on recursively defined sequences turns out to be helpful.

Observation 5.1. Consider a sequence $(a_t)_{t \in \mathbb{N}}$ satisfying the recursive inequality $a_t \leq b \sum_{k=0}^{t-1} a_k + c$. Then we have

$$\sum_{k=0}^t a_k \leq \frac{c}{b+1} (b+2)^{t+1} .$$

Proof. In order to prove the bound, we define the following auxiliary sequence $(z_t)_{t \in \mathbb{N}}$ by $z_t := (b+1) \sum_{k=0}^{t-1} z_k$ for $t > 0$ and $z_0 = c$. Observe now that this sequence dominates the “real” sequence $(a_t)_{t \in \mathbb{N}}$, this is $a_t \leq z_t$ for all $t \in \mathbb{N}$.

Furthermore, we have that $z_t \leq c(b+2)^t$, which can be easily proven by induction. For $t = 0$ this is clear. Now observe

$$z_t = (b+1) \sum_{k=0}^{t-1} z_k \leq c(b+1) \sum_{k=0}^{t-1} (b+2)^k \leq c(b+1) \frac{(b+2)^t}{(b+2) - 1} = c(b+2)^t .$$

We conclude for the sum

$$\sum_{k=0}^t a_k \leq \sum_{k=0}^t z_k \leq \frac{z_{t+1}}{b+1} \leq \frac{c}{b+1} (b+2)^{t+1} .$$

$\square$

Using the observation above, we will now begin to prove a bound on the probability mass lying not below $(1 + \delta)p^*$ in the following lemma. Let $s_i = \min_{p_i} \left| \frac{f_i(p_i) - f_i(p_i + \frac{\delta}{2}\eta_i)}{\frac{\delta}{2}\eta_i} \right|$ be the minimal absolute value of the difference quotient of f_i at any point p_i and $p_i + \frac{\delta}{2}\eta_i$.

Lemma 5.1. *Let π be an ϵ -correlated equilibrium for some $\epsilon \geq 0$. Then for all $\delta > 0$, we can bound the probability that $p \not\leq (1 + \delta)p^*$ is chosen by*

$$\Pr_{p \sim \pi}[p \not\leq (1 + \delta)p^*] \leq \epsilon \left(\frac{n}{\delta} \max_{i \in [n]} \frac{2}{s_i \eta_i} + 2 \right)^{T+1},$$

$$\text{where } T = \frac{\log \frac{\delta}{4} - \log \max_{i \in [n]} \left| \frac{p_i^{\max}}{(1 + \frac{\delta}{2})p_i^*} \right|}{\log \max_{i \in [n]} \left| 1 - \frac{\eta_i}{p_i^*} \right|},$$

and s_i denotes the minimal absolute value of the difference quotient of f_i at any point p_i and $p_i + \frac{\delta}{2}\eta_i$.

Proof. We consider an iteration being a shifted FM iteration and bound the probability based on this iteration. We iteratively bound the probability mass which lies right of $p^{(t)}$, where $p^{(t)}$ is given by the following iteration starting from $p^{(0)} := p^{\max}$. For every user i we define

$$p_i^{(t)} = \min \left\{ p_i^{(t-1)}, \left(C \cdot p^{(t-1)} + \left(1 + \frac{\delta}{2} \right) \cdot \eta \right)_i \right\}.$$

The iteration equals the FM iteration starting from $p^{\max}$ with the noise vector being $(1 + \frac{\delta}{2})\eta$ instead of η . The fixed point of this iteration can directly be related to the original fixed point being $p' := (1 + \frac{\delta}{2})p^*$. Let $p'' := (1 + \delta)p^* \geq p'$.

Now let $q_0 := \Pr_{p \sim \pi}[p \not\leq p^{(0)}]$ and $q_{t+1} := \Pr_{p \sim \pi}[p \not\leq p^{(t+1)}] - \Pr_{p \sim \pi}[p \not\leq p^{(t)}]$. We will bound these probabilities by considering the marginal distributions. For each user i we consider $q_{0,i} := \Pr_{p \sim \pi}[p_i > p_i^{(0)}]$ and $q_{t+1,i} := \Pr_{p \sim \pi}[p_i > p_i^{(t+1)}] - \Pr_{p \sim \pi}[p_i > p_i^{(t)}]$. In this notation $q_t \leq \sum_{i=1}^n q_{t,i}$.

We now consider one fixed user i and bound $q_{0,i}$. User i could always use power $(Cp^{\max} + \eta)_i$ instead of the powers between $p_i^{(1)}$ and $p_i^{(0)} = p^{\max}$. No matter which powers the other users use, the transmission would still always be successful. Since π is an ϵ -equilibrium any switching operation can at most provide an expected utility gain of ϵ implying $h_i \cdot q_{0,i} \leq \epsilon$, where $h_i = \frac{\delta}{2} \cdot s_i \cdot \eta_i$.

We relate this to the next interval between $p_i^{(t+1)}$ and $p_i^{(t)}$ in an inductive fashion. We bound its probability mass given the probability mass previously bound. Thus, we basically “cut off” mass and argue using the remaining probability mass.

For $t > 0$, we can adapt this observation with slight modification. We consider the operation that always uses power $(Cp^{(t)} + \eta)_i$ instead of the powers between $p_i^{(t+1)}$ and $p_i^{(t)}$. Obviously this switching operation does not guarantee that user i stays successful. It can indeed render transmissions unsuccessful. This only happens if a vector p with $p \not\leq p^{(t)}$ has been chosen. That is, user i might lose all utility obtained by power vectors “cut off” before. So we get $h_i \cdot q_{t,i} - \sum_{k=0}^{t-1} q_k \leq \epsilon$, or equivalently

$$q_{t,i} \leq \frac{1}{h_i} \left(\epsilon + \sum_{k=0}^{t-1} q_k \right).$$

Summing up over all users then leads to

$$\begin{aligned} q_t &\leq \sum_{i=1}^n \frac{1}{h_i} \left(\epsilon + \sum_{k=0}^{t-1} q_k \right) = \sum_{i=1}^n \frac{2}{\delta s_i \eta_i} \left(\epsilon + \sum_{k=0}^{t-1} q_k \right) \\ &= \frac{2n}{\delta} \sum_{i=1}^n \max_i \frac{1}{s_i \eta_i} \left(\epsilon + \sum_{k=0}^{t-1} q_k \right) \leq \frac{n}{\min_i h_i} \left(\epsilon + \sum_{k=0}^{t-1} q_k \right). \end{aligned}$$

Using Observation 5.1 here yields for all $t \geq 1$

$$\sum_{k=0}^t q_k \leq \frac{n \frac{\epsilon}{\min_i h_i}}{\frac{n}{\min_i h_i} + 1} \left(\frac{n}{\min_i h_i} + 2 \right)^{t+1} \leq \epsilon \left(\frac{n}{\min_i h_i} + 2 \right)^{t+1}.$$

The iteration given here only deviates from the FM iteration by changing the noise vector accordingly. We can use the bound given in Theorem 5.2 to bound the number of iteration steps T needed to reach a state $p^{(T)} \leq p''$. This is

$$t \geq T = \frac{\log \frac{\delta}{4} - \log \max_{i \in [n]} \left| \frac{p_i^{\max}}{(1 + \frac{\delta}{2}) p_i^*} \right|}{\log \max_{i \in [n]} \left| 1 - \frac{\eta_i}{p_i^*} \right|}.$$

Altogether we get a bound on the probability mass lying above $(1 - \delta)p^*$ based on the time a certain FM iteration would need to reach p'' . $\square$

We will bound the probability that vectors below $(1 - \delta)p^*$ are chosen using similar techniques as done in the proof of Lemma 5.2. We define $r = \min_i r_i$, where r_i is a lower bound on the utility of user i for successful transmissions at $(1 + \delta)p^*$, i.e., $r \leq \min_{i \in [n]} f_i((1 + \delta)p_i^*)$. Lemma 5.2 actually assumes $r \geq \frac{1}{2}$ though most of the proof holds for general r and we only use the assumption to apply a constant bound on r . In general, any constant bound on r would do. This is easily justified by the fact that any successful transmission is strictly better than not being successful and so the utility for a successful transmission is greater than 0.

Lemma 5.2. *Given an ϵ -correlated equilibrium and assuming $f_i((1 + \delta)p_i^*) \geq r = \frac{1}{2}$ for all $i \in [n]$. Then for every $\delta > 0$ the probability $\Pr_{p \sim \pi}[p \not\geq (1 - \delta)p^*]$ is at most*

$$\epsilon \left(2 + \left(\frac{n}{\delta} \max_{i \in [n]} \frac{2}{s_i \eta_i} + 2 \right)^{T+1} \right) (2n)^{T'+1}$$

with $T' = \frac{\log \delta}{\log \max_{i \in [n]} \left| 1 - \frac{\eta_i}{p_i^*} \right|}$ and T defined as in Lemma 5.1.

Proof. We will first of all bound $\Pr_{p \sim \pi}[p \not\geq (1 - \delta)p^*]$ depending on r and on the previously bound $\Pr_{p \sim \pi}[p \not\geq (1 + \delta)p^*]$.

Using the results from Lemma 5.1 we can then prove the claim. First we will show

$$\Pr_{p \sim \pi}[p \not\geq (1 - \delta)p^*] \leq \left(\frac{\epsilon}{1 - r} + \Pr_{p \sim \pi}[p \not\geq (1 + \delta)p^*] \right) \left(\frac{n}{r} \right)^{T'+1}.$$

We consider the same iteration steps as they are used in the FM iteration starting with $p^{(0)} := 0$, $p^{(t+1)} = C \cdot p^{(t)} + \eta$. We now recursively “cut off” the probability mass for vectors $p \not\geq p^{(t)}$ similarly as done in the proof of Lemma 5.1.

In order to bound the probability mass $q_t := \Pr[p \not\leq p^{(t)}] - \Pr[p \not\leq p^{(t-1)}]$ in the t -th step, we recursively bound q_t and $q_{t,i} := \Pr[p_i < p_i^{(t)}] - \Pr[p_i < p_i^{(t-1)}]$.

A user can gain at least $r_i \cdot q_{t,i}$ utility by shifting the probability mass lying on powers lower than $p^{(t)}$ to $(1 + \delta)p^*$. But he might lose utility up to $1 - r_i$ from states $p \not\leq (1 + \delta)p^*$ and from states that where “cut off” before.

Let $\rho = \Pr_{p \sim \pi}[p \not\leq (1 + \delta)p^*]$. We can bound the possible gain in utility by

$$r_i \cdot q_{t,i} - (1 - r_i) \left(\rho + \sum_{k=0}^{t-1} q_k \right) \leq \epsilon ,$$

and therefore

$$q_{t,i} \leq \frac{\epsilon + (1 - r_i)\rho + \sum_{k=0}^{t-1} (1 - r_i)q_k}{r_i} .$$

Summing over all users i gives us

$$\begin{aligned} q_t &\leq \sum_{i=1}^n \frac{\epsilon + (1 - r_i)\rho + \sum_{k=0}^{t-1} (1 - r_i)q_k}{r_i} \\ &\leq \frac{n(\epsilon + (1 - r)\rho)}{r} + \sum_{k=0}^{t-1} n \left(\frac{1}{r} - 1 \right) q_k . \end{aligned}$$

Using Observation 5.1 again to solve the recursion, we can bound the sum of all q_t for $t = 0, \dots, T'$ by

$$\sum_{t=0}^{T'} q_t \leq \frac{1}{1 - r} (\epsilon + \rho(1 - r)) \left(\frac{n}{r} (1 - r) + 2 \right)^{T'+1} \leq \left(\frac{\epsilon}{1 - r} + \rho \right) \left(\frac{n}{r} \right)^{T'+1} .$$

Having already found a bound for $\rho = \Pr_{p \sim \pi}[p \not\leq (1 + \delta)p^*]$ in Lemma 5.1, we can directly conclude

$$\sum_{t=0}^{T'} q_t \leq \left(\frac{\epsilon}{1 - r} + \epsilon \left(\frac{n}{\delta} \max_{i \in [n]} \frac{2}{s_i \eta_i} + 2 \right)^{T'+1} \right) \left(\frac{n}{r} \right)^{T'+1} .$$

Assuming r to be lower bounded by any constant suffices to get a bound depending on that constant. For simplicity we consider $r \geq \frac{1}{2}$ here. This yields

$$\sum_{t=0}^{T'} q_t \leq \epsilon \left(2 + \left(\frac{n}{\delta} \max_{i \in [n]} \frac{2}{s_i \eta_i} + 2 \right)^{T'+1} \right) (2n)^{T'+1} .$$

We can again use Theorem 5.2 to bound T' as needed in our analysis. With $p^{(0)} = 0$ this is

$$T' = \frac{\log \delta}{\log \max_{i \in [n]} \left| 1 - \frac{\eta_i}{p_i^*} \right|} .$$

□

Note that T and T' in Lemma 5.1 and 5.2 cannot be identified with the time necessary until a no-regret algorithm converges. Those two bounds are analytically due to the comparison between the ϵ -correlated equilibrium and an FM iteration.

Combining Lemma 5.1 and 5.2, we get for an ϵ -correlated equilibrium an upper bound on the probability $\Pr_{p \sim \pi}[(1 - \delta)p^* \leq p \leq (1 + \delta)p^*]$. For appropriately chosen δ , this bound together with Proposition 5.3 yields that the success probability converges to 1 as ϵ approaches zero. This also shows that in each no-swap-regret sequence for each user the limit of the fraction of successful steps is 1. Furthermore, the chosen powers also converge to p^* . For no-swap-regret learning algorithms the same holds when the regret converges to 0.

Considering an actual no-swap-regret algorithm, the time until the chosen powers are close to p^* with high probability can be bound by the time until the average swap regret ϵ is low enough to have a probability of $\frac{1}{n}$ here. Assume that the average regret of the considered no-swap-regret algorithm reduces in a polynomial way. In total, this shows that no-swap-regret algorithms converge to the actual optimal power assignment in a time polynomial in the number of links.

CHAPTER 6

Conclusion and Further Work

In this thesis, we have considered the capacity-maximization problem in wireless networks. We analyzed algorithms to decide which links in a wireless network should transmit. Wireless networks were assumed to consist of sender-receiver pairs called links and each sender wants to transmit to the respective receiver. In most parts of this thesis links were assumed to use uniform powers. The successful reception of transmissions was assumed to be subject to interference from concurrent transmissions and ambient noise. This way, not all links can be feasible in the same time step. The algorithmic goal considered throughout this thesis is to find the largest set of links, such that they are feasible transmitting successfully.

The interference from other links is assumed to be represented by an interference model. We have considered a rather general one based on weighted conflict graphs. This conflict-graph model generalizes different other interference models like unit-disk graphs or the more intricate SINR model.

Using this interference model we have identified properties, which were utilized to construct a proof template. This proof was applied to analyse adaptive learning approaches for capacity maximization. This way, it is possible to derive approximation guarantees for no-regret learning algorithms under adversarial jamming. Besides this, we extended these techniques to settings with multiple receivers and settings where the standard assumption that links stay in the network indefinitely is omitted.

We introduced no-regret learning algorithms known from game theory for this setting by defining suitable utility functions. These algorithms fulfill the conditions necessary to apply our proof template. In simulations, this approach shows a favorable behavior against stochastic jammers. It converges close to the optimum.

We also proved approximation guarantees for no-regret learning algorithms in settings with multiple channels, which are subject to availabilities. The availabilities can be determined by primary users opening the licensed spectrum to other links.

In the second part of this thesis, we have considered the problem of power control, which is highly related to capacity maximization. The task is to choose a power assignment such that all links can conduct a successful transmission simultaneously. Adaptive no-regret sequences were shown to converge to an optimal power assignment. In contrast to previously considered approaches like the FM iteration, we were even able to show that no-regret learning algorithms yield guarantees on the number of successful transmissions.

The results presented in this thesis show a favorable behavior of no-regret learning algorithms. They can guarantee convergence to constant-factor approximations in many settings.

Nevertheless, there are still some open problems, which we will consider in the remainder of this chapter.

The two main topics taken into consideration in this thesis are different in their general characteristics. While the first chapters considered the question which links should transmit, Chapter 5 examined the question of which powers to use for transmission. For centralized algorithms, there exist works analyzing combinations of both these tracks. For our distributed setting, it is still unknown how to connect both tasks in one approach. The main problem here is to enable links to decide whether to increase the power to be successful or to opt out of transmission instead. This decision is necessary for capacity maximization combining scheduling and power control. Without additional information exceeding the restricted feedback considered in this thesis, this seems impossible. It is an interesting future research direction to see how much information links actually need and how the no-regret property can still be algorithmically utilized.

Similar to the approaches analyzed in this thesis one can consider no-regret learning algorithms for throughput maximization. Here, one does not strive to maximize the number of successful transmissions but the sum of functions depending on the reached SINR. This is related to Shannon capacity, which is often used to model throughput of a link in practical research [Har28, Tel99]. Analyzing flexible no-regret learning algorithms here is clearly promising.

Although our approaches rely on very little feedback by only assuming links to know whether previous transmissions were successful, our analyses do not incorporate how links actually get this feedback. Other approaches with different assumptions already considered this issue [KV10]. However, this was not considered in terms of no-regret learning until now. It is very promising to consider a similar approach here for bidirectional connections, which allow acknowledgments to be transmitted back to the original sender. Nevertheless, direct application of this needs links to have more information about the network. So this topic is still an open future direction, especially when assuming restricted information about the network size and structure or when considering the usage of no-regret learning algorithms.

Answering these open questions would definitely bring the considered settings closer to practical applications. There are also other considerations relevant to applied settings, which should be introduced into theoretical research in the future. One such topic, which was already discussed in related work [DHK12b], is Rayleigh fading. The known results concerning Rayleigh fading indicate that a lot of our results can be transferred to this more realistic setting with only a small loss. Nevertheless, other even more challenging models are yet to be considered and approximation guarantees can also be improved in already considered ones.

For other more applied settings, it is still open how to incorporate robust no-regret learning or similar approaches. There are several aspects in which our theoretical model differs from real-world networks. For example, in general, links do not have to transmit data in every time step. Thus, assuming the distributed algorithm to decide whether to attempt a transmission in every time step is not always feasible.

Another underlying assumption in our and many other works is that the network does not change over time. Especially, mobile networks do not comply with this assumption. Although, our no-regret learning approach is very flexible, performance bounds assuming certain changing network characteristics are unknown. These extensions need to be included directly into the theoretical model to determine whether no-regret learning still yields reasonable approximation guarantees.

Besides extending the theoretical wireless-network model to additional problem variants, there are still almost no practical results in terms of implementations of regret-learning approaches. Implementations and further simulations would allow research to get more insight into the behavior of wireless networks and it can show whether the results of our simulations comply with reality.

Our simulations in Section 2.5 provide very promising results, which showed the constant factors in our approximation guarantees to be almost negligible. It is not obvious whether this holds in general. The theoretical approximation guarantees can be improved in terms of these constants.

While the open problems stated above are rather general ones, there are also some further research directions specifically targeting the settings considered in this thesis.

Concerning the jammer scenario, more intricate adversary models can be considered. In particular, we want to mention the model considered in [ORS⁺13], where the constraints on the jammer are subject to a power budget. This models jamming behavior somewhat more realistically, as adversarial jammers and concurrent networks might be restricted in terms of power. Assuming very restricted feedback utilizable by algorithms was not considered in the link-centered scenarios with power-restricted jammers yet. It would be nice to see this scenario being incorporated in our proof template. This would then provide the guarantees proven in Chapter 2 and Chapter 3.

Besides this kind of jammer, other adversarial models can also be considered in the future. The dual-graph model, which is frequently considered in graph-based interference models [KLN⁺10], is not directly applicable under SINR interference. Thus, this dual-graph model needs to be adapted to the SINR setting appropriately.

Improving approximation factors assuming links to be allowed to leave the network is also a future research direction. The factor of $\log n$ proven in Section 3.3 seems to be an analytical artifact. However, it is not clear whether this factor can actually be improved further or whether it has to be part of the bound.

The multi-channel setting allows to consider primary and secondary users. Although our approach provides rather general bounds depending on (C, D) -correlation independence, we only prove bounds on the parameter D for laminar correlation. Thus, one could consider more practical assumptions on the correlation among channel availabilities to get better insight into applications of these multi-channel algorithms.

In total, the results presented in this thesis are a profound basis for the analysis of regret-learning approaches for capacity maximization. They also offer a wide range of further topics for future research. Especially, driving theoretical research of adaptive protocols further into the direction of applied settings seems to be fruitful in the future. Combining the scheduling and the power-control task also seem to be a very promising future research direction.

Bibliography

- [ABAJ04] Eitan Altman, Dhiman Barman, Rachid El Azouzi, and Tania Jiménez. A game theoretic approach for delay minimization in slotted ALOHA. In *Proceedings of the IEEE International Conference on Communications (ICC)*, pages 3999–4003. IEEE, 2004.
- [Abr70] Norman Abramson. THE ALOHA SYSTEM: Another alternative for computer communications. In *Proceedings of the November 17-19, 1970, Fall Joint Computer Conference (AFIPS)*, pages 281–285. ACM, 1970.
- [ABSA02] Tansu Alpcan, Tamer Basar, R. Srikant, and Eitan Altman. CDMA uplink power control as a noncooperative game. *Wireless Networks*, 8(6):659–670, 2002.
- [ACBFS02] Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The non-stochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002.
- [AD09] Matthew Andrews and Michael Dinitz. Maximizing capacity in arbitrary wireless networks in the SINR model: Complexity and game theory. In *Proceedings of the 28th IEEE International Conference on Computer Communications (INFOCOM)*, pages 1332–1340. IEEE, 2009.
- [AFB⁺08] Tansu Alpcan, Xingzhe Fan, Tamer Basar, Murat Arcak, and John T. Wen. Power control for multicell CDMA wireless networks: A team optimization approach. *Wireless Networks*, 14(5):647–657, 2008.
- [AHK12] Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.
- [ÁHM12] Eyjolfur Ingi Ásgeirsson, Magnús M. Halldórsson, and Pradipta Mitra. A fully distributed algorithm for throughput performance in wireless networks. In *Proceedings of the 46th Annual Conference on Information Sciences and Systems (CISS)*, pages 1–5. IEEE, 2012.
- [AM11] Eyjolfur Ingi Ásgeirsson and Pradipta Mitra. On a game theoretic approach to capacity maximization in wireless networks. In *Proceedings of the 30th IEEE International Conference on Computer Communications (INFOCOM)*, pages 3029–3037. IEEE, 2011.

- [ARS08] Baruch Awerbuch, Andréa Richa, and Christian Scheideler. A jamming-resistant MAC protocol for single-hop wireless networks. In *Proceedings of the 27th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 45–54. ACM, 2008.
- [Aue03] Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research (JMLR)*, 3:397–422, 2003.
- [BCP95] Nicholas Bambos, Shou C. Chen, and Gregory J. Pottie. Radio link admission algorithms for wireless networks with power control and active link quality protection. In *Proceedings of the 14th IEEE International Conference on Computer Communications (INFOCOM)*, pages 97–104. IEEE, 1995.
- [BH14] Marijke H. L. Bodlaender and Magnús M. Halldórsson. Beyond geometry : Towards fully realistic wireless models. *CoRR*, abs/1402.5003, 2014.
- [BJ10] Alireza Babaei and Bijan Jabbari. Transmission probability control game for co-existing random ALOHA wireless networks in unlicensed bands. In *Proceedings of the 71st IEEE Vehicular Technology Conference*, pages 1–5. IEEE, 2010.
- [Blu97] Avrim Blum. Empirical support for winnow and weighted-majority algorithms: Results on a calendar scheduling domain. *Machine Learning*, 26(1):5–23, 1997.
- [Blu98] Avrim Blum. On-line algorithms in machine learning. In Amos Fiat and Gerhard J. Woeginger, editors, *Online Algorithms: the state of the art*, chapter 14. Springer, 1998.
- [BM07] Avrim Blum and Yishay Mansour. From external to internal regret. *Journal of Machine Learning Research (JMLR)*, 8:1307–1324, 2007.
- [BMJ⁺98] Josh Broch, David A. Maltz, David B. Johnson, Yih-Chun Hu, and Jorjeta Jetcheva. A performance comparison of multi-hop wireless ad hoc network routing protocols. In *Proceedings of the 4th Annual ACM/IEEE International Conference on Mobile Computing and Networking*, pages 85–97. ACM, 1998.
- [Chl01] Bogdan S. Chlebus. Randomized communication in radio networks. In *Handbook on Randomized Computing*, volume 9, pages 401–456. Kluwer Academic Publishers, 2001.
- [DHK11] Johannes Dams, Martin Hoefer, and Thomas Kesselheim. Convergence time of power-control dynamics. In *Proceedings of the 38th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 2, pages 637–649. Springer, 2011.
- [DHK12a] Johannes Dams, Martin Hoefer, and Thomas Kesselheim. Convergence time of power-control dynamics. *IEEE Journal on Selected Areas in Communications (J-SAC)*, 30(11):2231–2237, 2012.
- [DHK12b] Johannes Dams, Martin Hoefer, and Thomas Kesselheim. Scheduling in wireless networks with rayleigh-fading interference. In *Proceedings of the 24th ACM*

- symposium on Parallelism in algorithms and architectures (SPAA)*, pages 327–335. ACM, 2012.
- [DHK13a] Johannes Dams, Martin Hoefer, and Thomas Kesselheim. Jamming-resistant learning in wireless networks. *CoRR*, abs/1307.5290, 2013.
- [DHK13b] Johannes Dams, Martin Hoefer, and Thomas Kesselheim. Sleeping experts in wireless networks. In *Proceedings of the 27th International Symposium on Distributed Computing (DISC)*, pages 344–357. Springer, 2013.
- [DHK14] Johannes Dams, Martin Hoefer, and Thomas Kesselheim. Jamming-resistant learning in wireless networks. In *Proceedings of the 41st International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 447–458. Springer, 2014.
- [Din10] Michael Dinitz. Distributed algorithms for approximating wireless network capacity. In *Proceedings of the 29th IEEE International Conference on Computer Communications (INFOCOM)*, pages 1–9. IEEE, 2010.
- [DKV11] Johannes Dams, Thomas Kesselheim, and Berthold Vöcking. Transmission probability control game with limited energy. In *Proceedings of the IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, pages 420–430. IEEE, 2011.
- [DP98] Devdatt P. Dubhashi and Alessandro Panconesi. *Concentration of Measure for the Analysis of Randomised Algorithms*. Cambridge University Press, 1998.
- [DP13] Michael Dinitz and Merav Parter. Braess’s paradox in wireless networks: The danger of improved technology. In *Proceedings of the 27th International Symposium on Distributed Computing (DISC)*, pages 477–491. Springer, 2013.
- [EE04] Tamer A. ElBatt and Anthony Ephremides. Joint scheduling and power control for wireless ad hoc networks. *IEEE Transactions on Wireless Communications*, 3(1):74–85, 2004.
- [EG10] Thomas Erlebach and Tom Grant. Scheduling multicast transmissions under SINR constraints. In *Proceedings of the 6th International Workshop on Algorithms for Sensor Systems, Wireless Ad Hoc Networks, and Autonomous Mobile Entities (ALGOSENSORS)*, pages 47–61. Springer, 2010.
- [EJS05] Thomas Erlebach, Klaus Jansen, and Eike Seidel. Polynomial-time approximation schemes for geometric intersection graphs. *SIAM Journal on Computing*, 34(6):1302–1323, 2005.
- [FGHV13] Alexander Fanghänel, Sascha Geulen, Martin Hoefer, and Berthold Vöcking. Online capacity maximization in wireless networks. *Journal of Scheduling*, 16(1):81–91, 2013.
- [FKRV09] Alexander Fanghänel, Thomas Kesselheim, Harald Räcke, and Berthold Vöcking. Oblivious interference scheduling. In *Proceedings of the 28th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 220–229. ACM, 2009.

- [FM92] Gerard J. Foschini and Zoran Miljanic. A simple distributed autonomous power control algorithm and its convergence. *IEEE Transactions on Vehicular Technology*, 42(4):641–646, 1992.
- [FMGS02] David Famolari, Narayan Mandayam, David Goodman, and Viral Shah. A new framework for power control in wireless data networks: Games, utility, and pricing. In *Wireless Multimedia Network Technologies*, pages 289–309. Springer, 2002.
- [FNK⁺11] Michael Fitch, Maziar Nekovee, Santosh Kawade, Keith Briggs, and Richard MacKenzie. Wireless service provision in tv white space with cognitive radio technology: A telecom operator’s perspective and experience. *IEEE Communications Magazine*, 49(3):64–73, 2011.
- [FS97] Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139, 1997.
- [FSSW97] Yoav Freund, Robert E. Schapire, Yoram Singer, and Manfred K. Warmuth. Using and combining predictors that specialize. In *Proceedings of the 29th ACM Symposium on the Theory of Computing (STOC)*, pages 334–343. ACM, 1997.
- [GÁB⁺14] Helga Gudmundsdottir, Eyjólfur Ingi Ásgeirsson, Marijke H. L. Bodlaender, Joseph T. Foley, Magnús M. Halldórsson, Geir M. Järvelä, Henning Úlfarsson, and Ymir Vigfusson. Measurement based interference models for wireless scheduling algorithms. *CoRR*, abs/1401.1723, 2014.
- [Gal85] Robert G. Gallager. A perspective on multiaccess channels. *IEEE Transactions on Information Theory*, 31(2):124–142, 1985.
- [GGM08] Geoffrey J. Gordon, Amy Greenwald, and Casey Marks. No-regret learning in convex games. In *Proceedings of the 25th International Conference on Machine Learning*, pages 360–367. ACM, 2008.
- [GHK⁺14] Oliver Göbel, Martin Hoefer, Thomas Kesselheim, Thomas Schleiden, and Berthold Vöcking. Online independent set beyond the worst-case: Secretaries, prophets, and periods. In *Proceedings of the 41st International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 508–519. Springer, 2014.
- [GHW13] Olga Goussevskaia, Magnús M. Halldórsson, and Roger Wattenhofer. Algorithms for wireless capacity. *IEEE/ACM Transactions on Networking (TON)*, 2013.
- [GJM11] Amy Greenwald, Amir Jafari, and Casey Marks. No- Φ -regret: A connection between computational learning theory and game theory. In *Games, Norms and Reasons*, volume 353 of *Synthese Library*, pages 119–140. Springer, 2011.
- [GK00] Piyush Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, 2000.

- [GKP⁺14] Seth Gilbert, Valerie King, Seth Pettie, Ely Porat, Jared Saia, and Maxwell Young. (Near) optimal resource-competitive broadcast with jamming. In *Proceedings of the 26th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*. ACM, 2014. to be published.
- [GLN13] Mohsen Ghaffari, Nancy A. Lynch, and Calvin C. Newport. The cost of radio network broadcast for different models of unreliable links. In *Proceedings of the 32nd ACM Symposium on Principles of Distributed Computing (PODC)*, pages 345–354. ACM, 2013.
- [GMW08] Olga Goussevskaia, Thomas Moscibroda, and Roger Wattenhofer. Local broadcasting in the physical interference model. In *Proceedings of the DIALM-POMC Joint Workshop on Foundations of Mobile Computing (DIALM-POMC)*, pages 35–44. ACM, 2008.
- [GOW07] Olga Goussevskaia, Yvonne Anne Oswald, and Roger Wattenhofer. Complexity in geometric SINR. In *Proceedings of the 8th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC)*, pages 100–109. ACM, 2007.
- [GWHW09] Olga Goussevskaia, Roger Wattenhofer, Magnús M. Halldórsson, and Emo Welzl. Capacity of arbitrary wireless networks. In *Proceedings of the 28th IEEE International Conference on Computer Communications (INFOCOM)*, pages 1872–1880. IEEE, 2009.
- [GZY95] Sudheer A. Grandhi, Jens Zander, and Roy D. Yates. Constrained power control. *Wireless Personal Communications*, 1(4):257–270, 1995.
- [Hal12] Magnús M. Halldórsson. Wireless scheduling with power control. *ACM Transactions on Algorithms*, 9(1):7:1–7:20, 2012.
- [Har28] Ralph V. L. Hartley. Transmission of information. *Bell System Technical Journal*, 7(3):535–563, 1928.
- [HBH06a] Jianwei Huang, Randall A. Berry, and Michael L. Honig. Auction-based spectrum sharing. *Mobile Networks and Applications*, 11(3):405–408, 2006.
- [HBH06b] Jianwei Huang, Randall A. Berry, and Michael L. Honig. Distributed interference compensation for wireless networks. *IEEE Journal on Selected Areas in Communications (J-SAC)*, 24(5):1074–1084, 2006.
- [Hei06] Tiina Heikkinen. A potential game approach to distributed power control and scheduling. *Computer Networks*, 50(13):2295–2311, 2006.
- [HGGB04] Tim Holliday, Andrea Goldsmith, Peter W. Glynn, and Nick Bambos. Distributed power and admission control for time varying wireless networks. In *Proceedings of the 47th Global Telecommunications Conference (GLOBECOM)*, pages 768–774, 2004.

- [HHMW13] Magnús M. Halldórsson, Stephan Holzer, Pradipta Mitra, and Roger Wattenhofer. The power of non-uniform wireless power. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1595–1606. SIAM, 2013.
- [HKV11] Martin Hoefer, Thomas Kesselheim, and Berthold Vöcking. Approximation algorithms for secondary spectrum auctions. In *Proceedings of the 23rd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 177–186. ACM, 2011.
- [HM11] Magnús M. Halldórsson and Pradipta Mitra. Wireless capacity with oblivious power in general metrics. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1538–1548. SIAM, 2011.
- [HM12] Magnús M. Halldórsson and Pradipta Mitra. Wireless capacity and admission control in cognitive radio. In *Proceedings of the 31st IEEE International Conference on Computer Communications (INFOCOM)*, pages 855–863. IEEE, 2012.
- [HP05] Marcus Hutter and Jan Poland. Adaptive online prediction by following the perturbed leader. *Journal of Machine Learning Research (JMLR)*, 6(4):639–660, 2005.
- [HW09] Magnús M Halldórsson and Roger Wattenhofer. Wireless communication is in APX. In *Automata, Languages and Programming*, pages 525–536. Springer, 2009.
- [HY98] Ching-Yao Huang and Roy D. Yates. Rate of convergence for minimum power assignment algorithms in cellular radio systems. *Wireless Networks*, 4(4):223–231, 1998.
- [IRK09] Aravind Iyer, Catherine Rosenberg, and Aditya Karnik. What is the right model for wireless channel interference? *IEEE Transactions on Wireless Communications*, 8(5):2662–2671, 2009.
- [JH98] Hongbin Ji and Ching-Yao Huang. Non-cooperative uplink power control in cellular radio systems. *Wireless Networks*, 4(3):233–240, 1998.
- [JJ09] Christopher A. St. Jean and Bijan Jabbari. On game-theoretic power control under successive interference cancellation. *IEEE Transactions on Wireless Communications*, 8(4):1655–1657, 2009.
- [JPPQ05] Kamal Jain, Jitendra Padhye, Venkata N. Padmanabhan, and Lili Qiu. Impact of interference on multi-hop wireless network performance. *Wireless Networks*, 11(4):471–487, 2005.
- [KAY10] Stepán Kucera, Sonia Aïssa, and Susumu Yoshida. Adaptive channel allocation for enabling target SINR achievability in power-controlled wireless networks. *IEEE Transactions on Wireless Communications*, 9(2):833–843, 2010.
- [Kes11] Thomas Kesselheim. A constant-factor approximation for wireless capacity maximization with power control in the sinr model. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1549–1559. SIAM, 2011.

- [Kes12] Thomas Kesselheim. Approximation algorithms for wireless link scheduling with flexible data rates. In *Proceedings of the 20th Annual European Symposium on Algorithms (ESA)*, pages 659–670. Springer, 2012.
- [KG05] Sarah Koskie and Zoran Gajic. A nash game algorithm for SIR-based power control in 3G wireless CDMA networks. *IEEE/ACM Transactions on Networking*, 13(5):1017–1026, Oct 2005.
- [KKT04] Ulaş C. Kozat, Iordanis Koutsopoulos, and Leandros Tassiulas. A framework for cross-layer design of energy-efficient communication with QoS provisioning in multi-hop wireless networks. In *Proceedings of the 23rd IEEE International Conference on Computer Communications (INFOCOM)*, volume 2, pages 1446–1456. IEEE, 2004.
- [KLN⁺10] Fabian Kuhn, Nancy A. Lynch, Calvin C. Newport, Rotem Oshman, and Andréa W. Richa. Broadcasting in unreliable radio networks. In *Proceedings of the 29th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 336–345. ACM, 2010.
- [KMB09] Varun Kanade, H. Brendan McMahan, and Brent Bryan. Sleeping experts and bandits with stochastic action availability and adversarial rewards. In *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 272–279, 2009.
- [KNMS10] Robert Kleinberg, Alexandru Niculescu-Mizil, and Yogeshwer Sharma. Regret bounds for sleeping experts and bandits. *Machine learning*, 80(2-3):245–272, 2010.
- [KS12] Varun Kanade and Thomas Steinke. Learning hurdles for sleeping experts. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (ITCS)*, pages 11–18. ACM, 2012.
- [KV05] Adam Kalai and Santosh Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.
- [KV10] Thomas Kesselheim and Berthold Vöcking. Distributed contention resolution in wireless networks. In *Proceedings of the 24th International Symposium on Distributed Computing (DISC)*, pages 163–178. Springer, 2010.
- [KYR06] M. Kemal Karakayali, Roy D. Yates, and Leonid V. Razoumov. Downlink throughput maximization in cdma wireless networks. *IEEE Transactions on Wireless Communications*, 5(12):3492–3500, 2006.
- [LPPP11] Zvi Lotker, Merav Parter, David Peleg, and Yvonne Anne Pignolet. Distributed power control in the SINR model. In *Proceedings of the 30th IEEE International Conference on Computer Communications (INFOCOM)*, pages 2525–2533. IEEE, 2011.
- [LW94] Nick Littlestone and Manfred K. Warmuth. The weighted majority algorithm. *Information and Computation*, 108(2):212–261, 1994.

- [MB04] H. Brendan McMahan and Avrim Blum. Online geometric optimization in the bandit setting against an adaptive adversary. In *Learning theory*, pages 109–123. Springer, 2004.
- [MCPS06] Farhad Meshkati, Mung Chiang, H. Vincent Poor, and Stuart C. Schwartz. A game-theoretic approach to energy-efficient power control in multicarrier cdma systems. *IEEE Journal on Selected Areas in Communications (J-SAC)*, 24(6):1115–1129, 2006.
- [MW01a] Allen B. MacKenzie and Stephen B. Wicker. Game theory and the design of self-configuring, adaptive wireless networks. *IEEE Communications Magazine*, 39(11):126–131, 2001.
- [MW01b] Allen B. MacKenzie and Stephen B. Wicker. Game theory in communications: motivation, explanation, and application to power control. In *Proceedings of the Global Telecommunications Conference (GLOBECOM)*, pages 821–826. IEEE, 2001.
- [MW05] Thomas Moscibroda and Rogert Wattenhofer. Coloring unstructured radio networks. In *Proceedings of the 7th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 39–48. ACM, 2005.
- [MW06] Thomas Moscibroda and Roger Wattenhofer. The complexity of connectivity in wireless networks. In *Proceedings of the 25th IEEE International Conference on Computer Communications (INFOCOM)*. IEEE, 2006.
- [NC06] Nie Nie and Cristina Comaniciu. Adaptive channel allocation spectrum etiquette for cognitive radio networks. *Mobile Networks and Applications*, 11(6):779–797, 2006.
- [Nek09] Maziar Nekovee. A survey of cognitive radio access to tv white spaces. In *Proceedings of the 1st International Conference on Ultra Modern Telecommunications (ICUMT’09)*, pages 1–8. IEEE, 2009.
- [NHK08] Tim Nieberg, Johann Hurink, and Walter Kern. Approximation schemes for wireless networks. *ACM Transactions on Algorithms*, 4(4), 2008.
- [NRTV07] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. *Algorithmic game theory*. Cambridge University Press, 2007.
- [ORS⁺13] Adrian Ogierman, Andréa W. Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. Competitive mac under adversarial sinr. *CoRR*, abs/1307.7231, 2013.
- [Rap96] Theodore S. Rappaport. *Wireless communications - principles and practice*. Prentice Hall, 1996.
- [RL93] Subramanian Ramanathan and Errol L. Lloyd. Scheduling algorithms for multihop radio networks. *IEEE/ACM Transactions on Networking (TON)*, 1(2):166–177, April 1993.
- [RM92] Theodore S. Rappaport and Laurence B. Milstein. Effects of radio propagation path loss on DS-CDMA cellular frequency reuse efficiency for the reverse channel. *IEEE Transactions on Vehicular Technology*, 41(3):231–242, Aug 1992.

- [RSSZ10] Andréa W. Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. A jamming-resistant mac protocol for multi-hop wireless networks. In *Proceedings of the 24th International Symposium on Distributed Computing (DISC)*, pages 179–193. Springer, 2010.
- [RSSZ11] Andréa W. Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. Competitive and fair medium access despite reactive jamming. In *Proceedings of the 31st International Conference on Distributed Computing Systems (ICDCS)*, pages 507–516. IEEE, 2011.
- [RSSZ13] Andréa W. Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. Competitive throughput in multi-hop wireless networks despite adaptive jamming. *Distributed Computing*, 26(3):159–171, 2013.
- [SBH14] Changxin Shi, Randall A. Berry, and Michael L. Honig. Bi-directional training for adaptive beamforming and power control in interference networks. *IEEE Transactions on Signal Processing*, 62(3):607–618, Feb 2014.
- [SK83] John A. Silvester and Leonard Kleinrock. On the capacity of single-hop slotted aloha networks for various traffic matrices and transmission strategies. *IEEE Transactions on Communications*, 31(8):983–991, 1983.
- [SK11] V. P. Singh and Krishan Kumar. Literature survey on power control algorithms for mobile ad-hoc network. *Wireless Personal Communications*, 60(4):679–685, 2011.
- [SL07] Gilles Stoltz and Gábor Lugosi. Learning correlated equilibria in games with compact sets of strategies. *Games and Economic Behavior*, 59(1):187–208, 2007.
- [SS03] Harkirat Singh and Suresh Singh. DOA-ALOHA: slotted ALOHA for ad hoc networking using smart antennas. In *Proceedings of the 58th IEEE Vehicular Technology Conference*, volume 5, pages 2804–2808 Vol.5. IEEE, Oct 2003.
- [SW09] Johannes Schneider and Roger Wattenhofer. Coloring unstructured wireless multi-hop networks. In *Proceedings of the 28th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 210–219. ACM, 2009.
- [SW10] Johannes Schneider and Roger Wattenhofer. An optimal maximal independent set algorithm for bounded-independence graphs. *Distributed Computing*, 22(5-6):349–361, 2010.
- [Tel99] Emre Telatar. Capacity of multi-antenna gaussian channels. *European Transactions on Telecommunications*, 10(6):585–595, 1999.
- [WB01] Cynara Wu and Dimitri P. Bertsekas. Distributed power control algorithms for wireless networks. *IEEE Transactions on Vehicular Technology*, 50(2):504–514, 2001.
- [Yat95] Roy D. Yates. A framework for uplink power control in cellular radio systems. *IEEE Journal on Selected Areas in Communications*, 13(7):1341–1347, 1995.

- [YH95] Roy D. Yates and Ching-Yao Huang. Integrated power control and base station assignment. *IEEE Transactions on Vehicular Technology*, 44(3):638–644, 1995.
- [Zin03] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning (ICML)*, pages 928–936. IMLS, 2003.