Scheduling in Wireless Networks with Rayleigh-Fading Interference

Johannes Dams[*] Martin Hoefer[*] Thomas Kesselheim[†]

RWTH Aachen University
Department of Computer Science
52056 Aachen, Germany
{dams, mhoefer, kesselheim}@cs.rwth-aachen.de

ABSTRACT

We study algorithms for wireless spectrum access of n communication requests when interference conditions are given by the Rayleigh-fading model. This model extends the recently popular deterministic interference model based on the signal-to-interference-plus-noise ratio (SINR) using stochastic propagation to address fading effects observed in reality. We consider worst-case approximation guarantees for the two standard problems of capacity maximization (maximize the expected number of successful transmissions in a single slot) and latency minimization (minimize the expected number of slots until all transmissions were successful). Our main result is a generic reduction of Rayleigh fading to the deterministic SINR model. It allows to apply existing algorithms for the non-fading model in the Rayleigh-fading scenario while losing only a factor of $O(\log^* n)$ in the approximation guarantee. This way, we obtain the first approximation guarantees for Rayleigh fading and, more fundamentally, show that non-trivial stochastic fading effects can be successfully handled using existing and future techniques for the non-fading model. Using a more detailed argument, a similar result applies even for distributed and game-theoretic capacity maximization approaches. For example, it allows to show that regret learning yields an $O(\log^* n)$-approximation with uniform power assignments. Our analytical treatment is supported by simulations illustrating the performance of regret learning and, more generally, the relationship between both models.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Wireless Communication, Distributed Networks*; F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems

General Terms

Algorithms, Theory

Keywords

wireless network, transmission scheduling, SINR, Rayleigh fading

[*]Supported by DFG grant Ho 3831/3-1.

[†]Supported by DFG through UMIC Research Centre at RWTH Aachen University.

1. INTRODUCTION

Effective communication in wireless networks depends on successful reception in the presence of interference and noise, which have to be modeled realistically. Since the seminal work of Moscibroda and Wattenhofer [18], attention shifted from simple graph-based interference models to a more realistic model using SINR. This resulted in a variety of nontrivial insights into the algorithmic challenges and limitations of request scheduling [5, 11].

While the SINR model represents a significant improvement over previous approaches, it still uses a limited view of signal propagation. The main assumption is that any signal transmitted at power level p is always received after distance d with strength p/d^α, for some $\alpha > 0$. In contrast, in reality signal propagation is by no means deterministic. For instance, the SINR model does not account for short-term fluctuations such as fading. There exist advanced models using stochastic approaches that take fading effects into account. Most prominently, in the Rayleigh-fading model, signal strength is modeled by an exponentially distributed random variable with mean p/d^α. Stochastic propagation represents a major technical complication in the definition of interference models, and this may be a reason that – up to our knowledge – there are no general algorithmic results for request scheduling in this model or even for a direct comparision between the non-fading and Rayleigh-fading model.

In this paper, we examine the relationship between the non-fading SINR model and the Rayleigh-fading model. Our first main result is a fundamental relation between the models for instances of the same topology. It is based on a detailed analysis of the success probability in the Rayleigh-fading model, and it turns out to allow a surprisingly simple handling of the complicated stochastic propagation. This allows us to transfer existing algorithms and their performance bounds in the SINR model to the Rayleigh-fading model.

Our second main result uses a more detailed reduction to show that a similar result applies even for distributed capacity maximization via distributed regret-learning techniques. As the considered sequences generalize Nash equilibria, this result transfers the respective game-theoretic studies [1]. Our analytical results are supported by simulations illustrating the performance of regret learning and, more generally, the relationship between both models.

On a more fundamental level, our results highlight the inherent robustness of the techniques and bounds derived for the non-fading SINR model. The rather direct adaptation of existing algorithms to Rayleigh fading raises the hope that algorithms and their analyses can also be applied accordingly to interference models capturing further realistic properties.

1.1 Our Contribution

For a set of n communication requests in the non-fading model, we consider the two prominent problems of *capacity maximization* (maximizing the number of simultaneously successful requests in a single slot) and *latency minimization* (minimizing the number of slots such that every request has been successful at least once). In the Rayleigh-fading model, interference becomes stochastic, and thus capacity maximization becomes maximizing the expected number of successful requests in a single slot. Similarly, in latency minimization we strive to minimize the expected number of slots until every request has been successful at least once. In this sense, we adapt a similar perspective as in worst-case analysis of randomized algorithms – we strive to bound the expected performance of the algorithms in an arbitrary (worst-case) topology.

Our first main result characterizes the success probability of a request in the Rayleigh model. This probability is never 0, and thus requests can still be successful if in the non-fading model this is completely impossible (e.g., due to extremely large noise). For a meaningful comparison in terms of approximation factors, we thus focus on interference-dominated scenarios with reasonable noise conditions (for a formal definition see below). Under these conditions, we show in Section 3 that for every set of successful requests with respect to the non-fading model, in the Rayleigh model in expectation a constant fraction of these requests remains successful. Hence, we can use algorithms for capacity optimization in the non-fading model and lose only a constant factor when translating the output to Rayleigh fading. To bound approximation factors, however, we have to relate this to the Rayleigh-fading optimum, i.e., the maximum expected number of successful requests for any subset of transmitting requests. Here we show in Section 5 that this expected number can only be a factor of $O(\log^* n)$ larger than the maximum number of successful requests in the non-fading model. This allows to use existing algorithms and their bounds to derive approximation factors in the Rayleigh-fading model. For capacity maximization, we show, e.g., an $O(\log^* n)$-approximation with power control based on [14] and with distance-based power assignments based on [13]. For latency minimization similar arguments can be applied for algorithms that use repeated single-slot success maximization [7] or ALOHA-style protocols [15] in the non-fading model. For instance, we obtain an $O(\log^* n \cdot \log n)$-approximation for uniform power assignments based on [7]. The algorithms for latency minimization allow to directly apply multi-hop scheduling techniques as in [6, 15, 14]. The transformation does not modify transmission powers or depend on metrical properties of the distances. Thus, the respective properties of the algorithms and also the lower bounds, e.g., on power control [5, 11] are preserved.

In Section 6 we consider distributed approaches for capacity maximization, namely regret-learning algorithms [4]. Here we are not able to plug in the results for the non-fading model in a similar black-box fashion. Instead, we have to argue in a more detailed way to show that for uniform power assignments the expected number of successful requests is only a constant factor smaller than the size of the non-fading optimum. The bound is again completed using previous arguments, and we obtain a $O(\log^* n)$-factor with respect to the Rayleigh-fading optimum. Note that $\log^* n$ is essentially "almost constant", however, deriving a (provable) constant bound remains open.

Finally, we conduct a number of experiments that highlight the relation between the two models and the performance of regret learning. In particular, we observe that with probabilistic spectrum access, the curve for success probability in the Rayleigh-fading model can be seen as a smoothed variant of the success curve in the non-fading model. We observe that the non-fading model predicts more success if total interference is small, while Rayleigh fading allows more requests to become successful if interference is large. Regret-learning algorithms show fast convergence and good performance in both models, and the number of successful requests predicted by the non-fading model is somewhat larger.

1.2 Related Work

In a seminal paper, Gupta and Kumar [9] study the capacity of a wireless peer-to-peer network with a random topology based on non-fading SINR constraints. This brought about a lot of further work in randomly distributed networks [8, 21, 3]. Similar studies have been carried out for the case of regular topologies [20, 22]. Partly, this kind of research has also been generalized to networks with fading channels. For example, Liu and Haenggi [17] consider the capacity of square, triangle, hexagon, and random networks under Rayleigh-fading interference. More often Rayleigh fading is only used to model effects of noise, and interference inside the network itself is neglected [19, 10]. This represents an orthogonal approach because we concentrate on the particular problem of coordinating simultaneous transmissions. To the best of our knowledge, a direct comparision between the non-fading and Rayleigh-fading model like it is done in this paper has not been discussed in literature yet.

Real-world networks are typically neither random nor regular. This motivates the study of arbitrary topologies, as first done by Moscibroda and Wattenhofer [18]. Following this work, approximation algorithms in the non-fading SINR world were treated quite intensively, especially for the pure scheduling problems. Important milestones for capacity maximization are constant-factor approximations for uniform transmission powers [7]. A more sophisticated approach is selecting powers based on the distance between the sender and the respective receiver [13]. For uniform power assignments also a distributed algorithm has been developed [4] that uses regret learning. For latency minimization a distributed, ALOHA-like protocol has been analyzed [15, 12]. It yields an approximation factor of $O(\log n)$ with high probability.

The probably most natural extensions are the combined problems of scheduling and power control. This is, power levels are not fixed but have to be selected by the algorithm. This offers an additional freedom to the optimal solution as well. Using uniform transmission powers yields an $O(\log \Delta)$-approximation factor [1]. Here, Δ denotes the ratio of the maximal and the minimal distance between a sender and the respective receiver. One gets $O(\log\log\Delta + \log n)$-approximations when using square-root power assignments [11], i.e. a link of length d is assigned a transmission power proportional to $\sqrt{d^\alpha}$. The given approximation factors have been shown to be asymptotically almost optimal when restricting to these power assignments. However, for non-oblivious power assignments even a constant-factor approximation exists [14].

2. FORMAL MODEL DEFINITION

We assume that our network consists of n possible communication links $(s_1, r_1), \ldots, (s_n, r_n)$, each consisting of a sender and the respective receiver. In general, we do not make any assumptions on the geometry or distribution of the network nodes.

For the propagation, we consider Rayleigh-fading channels. This is, if a signal is transmitted by sender s_j, it is received by receiver r_i at a strength of $S_{j,i}$, which is a random variable. $S_{j,i}$ is exponentially distributed with mean $\bar{S_{j,i}}$. As usual, we assume this stochastic process to be independent for different (j,i) and from timeslot to timeslot.

The receiver r_i can successfully decode the signal transmitted by its sender s_i, if the SINR is above a certain theshold β, this is

$$\frac{S_{i,i}}{\sum_{j\neq i} S_{j,i} + \nu} \geq \beta \enspace .$$

Here, $\nu \geq 0$ is a constant denoting ambient noise.

We compare this channel model to the standard non-fading propagation model. Here, the received signal strength is always (deterministically) $\bar{S_{j,i}}$. As one can easily see, this comparison might not be fair. In the case of very low transmission powers, the mean received signal strength is already exceeded by the noise. Therefore, the transmission cannot be successful at all in the non-fading model, even in the absence of interference. In the Rayleigh-fading model in contrast, a small success probability remains. Therefore, as we focus on the impact of interference, we assume that $\bar{S_{i,i}}$ is always a constant factor higher than $\beta\nu$. To simplify notation, this factor is assumed to be 2, i.e. $\bar{S_{i,i}} \geq 2\beta\nu$.

When considering approximation algorithms in the non-fading model, it is usually very important that the signal strengths $\bar{S_{j,i}}$ are not arbitrary but determined by certain model parameters. That is for example $\bar{S_{j,i}} = p_j / d(s_j, r_i)^\alpha$ where p_j is the transmission power and $d(s_j, r_i)$ the distance between s_j and r_i. In contrast, our connection between Rayleigh-fading and non-fading models shown below applies in a more general scenario, without any assumptions on the values of the (expected) signal strength $\bar{S_{j,i}}$ – except non-negativity and the relation to noise as detailed in the paragraph above. In particular, this implies that our reduction between the models holds for arbitrary power assignments, path-loss exponents, requests located in metric spaces, etc. For proving bounded approximation factors, however, algorithms for the non-fading model usually rely heavily on $\bar{S_{j,i}}$ being characterized by these parameters. Consequently, our "black-box" translation of these algorithms and their approximation factors also applies only to instances of the Rayleigh model that have expected values $\bar{S_{j,i}}$ with the same characteristics.

3. SUCCESS PROBABILITY

In this section, we consider the following situation under Rayleigh-fading constraints. Assuming each sender s_i transmits with probability q_i, we bound the probability of a successful reception that we refer to as $Q_i(q_1, \ldots, q_n)$. Fortunately, in contrast to the non-fading model, the success probability can be given in a closed-form expression.

THEOREM 1. *The probability that receiver r_i can successfully receive the signal from s_i is*

$$Q_i(q_1, \ldots, q_n) = q_i \cdot \exp\left(-\frac{\beta\nu}{\bar{S_{i,i}}}\right) \prod_{j\neq i} \left(1 - \frac{\beta q_j}{\beta + \bar{S_{i,i}}/\bar{S_{j,i}}}\right) \enspace .$$

The proof of this expression is mainly due to Liu and Haenggi [17]; it can be found the appendix. The expression has the advantage of being an exact probability. However, in order to compare the probability to the one in the non-fading channel model, we need upper and lower bounds.

LEMMA 2. *The success probability for link i is at least*

$$Q_i(q_1, \ldots, q_n) \geq q_i \cdot \exp\left(-\frac{\beta}{\bar{S_{i,i}}}\left(\nu + \sum_{j\neq i} \bar{S_{j,i}} q_j\right)\right) \enspace .$$

The success probability for link i is at most

$$Q_i(q_1, \ldots, q_n) \leq q_i \cdot \exp\left(-\frac{\beta}{2\bar{S_{i,i}}}\left(\nu + \sum_{j\neq i} \bar{S_{j,i}} q_j\right)\right) \enspace .$$

PROOF. The proof of this lemma is based on the following observation concerning the exponential function.

OBSERVATION 3. *For all $x \in (0,1]$, $q \in [0,1]$, we have*

$$\exp(-xq) \leq 1 - \frac{q}{\frac{1}{x} + 1} \leq \exp\left(-\frac{1}{2}xq\right) \enspace .$$

PROOF. We show the first inequality using the fact that $\exp(y) \geq 1 + y$ for all $y \in \mathbb{R}$. Setting $y = xq$ yields

$$\exp(-xq) = \frac{1}{\exp(xq)} \leq \frac{1}{1 + xq} = 1 - \frac{q}{\frac{1}{x} + q} \leq 1 - \frac{q}{\frac{1}{x} + 1} \enspace .$$

Setting $y = -\frac{q}{\frac{1}{x}+1}$, we get

$$1 - \frac{q}{\frac{1}{x} + 1} \leq \exp\left(-\frac{q}{\frac{1}{x} + 1}\right) = \exp\left(-\frac{xq}{1 + x}\right) \enspace .$$

Furthermore, we have for all $x \in (0,1]$ that $\frac{x}{x+1} \geq \frac{1}{2}x$. This yields the second bound. □

Setting now $q = q_j$ and $x = \beta\bar{S_{j,i}}/\bar{S_{i,i}}$ in this observation, we get

$$\exp\left(-\beta\frac{\bar{S_{j,i}}}{\bar{S_{i,i}}} q_j\right) \leq 1 - \frac{\beta q_j}{\beta + \bar{S_{i,i}}/\bar{S_{j,i}}} \leq \exp\left(-\frac{1}{2}\beta\frac{\bar{S_{j,i}}}{\bar{S_{i,i}}} q_j\right) \enspace .$$

Theorem 1 now yields the claim. □

As a first result, this gives us the following relation between the success probability in the Rayleigh-fading channel compared to the one in the non-fading channel.

Corollary 4. *If a set $S \subseteq [n]$ is feasible in the non-fading channel model, setting $q_i = 1$ for all $i \in S$ and $q_i = 0$ for all $i \notin S$, we have $Q_i(q_1, \ldots, q_n) \geq 1/e$ for all $i \in S$.*

If $q_i \in \{0, 1\}$ for all $i \in [n]$ and the Rayleigh success probability is at least $1/\sqrt{e}$ for each link, the set of all links with $q_i = 1$ is feasible in the non-fading channel model.

4. TRANSFORMING SCHEDULING ALGORITHMS

The bounds given in the previous section immediately allow us to estimate the performance of algorithms for the non-fading model in a Rayleigh-fading environment after some minor modifications.

In particular, we can take an arbitrary approximation algorithm for capacity maximization. This might be one of the constant-factor approximations for the setting with uniform transmission powers [7] or monotone transmission powers [13], or even for the case in which the algorithm has to choose the transmission power itself [14]. In any case, a set of links is returned that is feasible in the non-fading model. Making exactly these links transmit with probability 1 (without changes of the transmission powers), Corollary 4 yields that each of them will be successful with probability at least $1/e$. In terms of our objective function "capacity" this means that we are at most a $1/e$-factor worse in expectation. In combination, this means that the resulting algorithm will compute transmission probabilities yielding an expected capacity that is at most a constant factor worse than the optimally achievable capacity in the non-fading model. However, it remains to show that the theoretical optimum in the Rayleigh-fading model cannot be much better than the one in the non-fading model. This will be carried out in Section 5.

Existing approximation algorithms to minimize latency can in general be divided into two classes. On the one hand, there are many algorithms actually attempting to maximize the utilization of the first time slot and then apply this procedure recursively on the remaining links. For these kinds of algorithms and analyses exactly the same argumentation as for capacity maximization can be applied. On the other hand, ALOHA-style protocols have been proposed. Here, in each time slot, each link is assigned a (small) transmission probability, which we assume to be smaller than $1/2$. If it is successful, the sender stops transmitting, otherwise it continues running the algorithm. In order to transform such algorithms to the Rayleigh-fading model, we let each (randomized) step be executed 4 times. This yields a success probability that is at least as large as in the non-fading model. If p is the success probability in the non-fading model, Corollary 4 yields the Rayleigh-fading success probability is at least $p \cdot 1/e$. In 4 independent repeats the probability of at least one success is therefore at least $1 - (1 - p/e)^4$. This is at least p if the transmission probability (and therefore the success probability) is at most $1/2$.

For multi-hop scheduling algorithms [15, 14], the single-hop transformations mentioned above can directly be generalized. Here, in fact, the resulting multi-hop schedule can also be considered as a concatination of single-hop schedules. Transforming each of them in the described way, we still only lose constant factors.

5. TRANSFORMING THE RAYLEIGH-FADING OPTIMUM

The performance of all algorithms constructed in Section 4 were measured in terms of the value of the optimal solution in the non-fading model. However, in order to derive approximation guarantees for the Rayleigh-fading model, the value of the computed solution has to be compared within the Rayleigh-fading model. Here, the optimal solution could potentially be much better than the non-fading one. In this section, we give a possibly surprising result that this indeed cannot happen in an interference-dominated environment. To be more precise, we take an arbitrary assignment of transmission probabilities. In the Rayleigh-fading model this yields a particular success probability for each link. We then simulate this single transmission with $O(\log^* n)$ independent steps in the non-fading model. In the end, for each link the success probability is at least as large as in the single Rayleigh-fading step.

This yields that for both considered scheduling problems, the Rayleigh-fading optimum can be at most an $O(\log^* n)$-factor better than the non-fading optimum. For capacity maximization this holds because we find $O(\log^* n)$ sets that are all feasible in the non-fading sense. In expectation, their summed value is at least as large as the one of the Rayleigh-fading optimum. This means that the best one can be at most an $O(\log^* n)$ factor worse.

When considering latency minimization under Rayleigh-fading conditions, the optimum should rather be considered as an algorithm itself that assigns transmission probabilities in each step. This assignment may arbitrarily depend on previous successes and may be computed using arbitrary computation power. However, our theorem shows for this case that even the perfect algorithm computes schedules that are at most an $O(\log^* n)$ factor shorter than the non-fading optimum, because we could replace each timeslot by the described simulation, increasing the schedule length by a factor of at most $O(\log^* n)$.

Theorem 5. *For each assignment of transmission probabilities $q_1, \ldots, q_n$ there is a simulation using $O(\log^* n)$-steps such that the non-fading success probability in these steps is at least $Q_i(q_1, \ldots, q_n)$ for each link i.*

Proof. We define $(b_k)_{k \in \mathbb{N}}$ recursively by setting $b_0 = 1/4$, $b_{k+1} = \exp(b_k/2)$. The simulation works as follows. For each $k \geq 0$ with $b_k < n$, we let each sender transmit with probability $q_i^{(k)} := q_i/4b_k$ for 19 times independently at random.

Algorithm 1: Formal description of the simulation.

$k := 0, b_0 := 1/4$;
while $b_k < n$ **do**
  for 19 *times* **do**
    transmit with probability $q_i^{(k)} := q_i/4b_k$;
  $b_{k+1} := \exp(b_k/2), k := k + 1$;

Consider an arbitrary $i \in [n]$. We claim: The probability of success during the $O(\log^* n)$ repeats in the non-fading model is at least $Q_i(q_1, \ldots, q_n)$.

We set $A = \sum_{j \neq i} \min\left\{1, \beta_i \bar{S_{j,i}}/\bar{S_{i,i}}\right\} \cdot q_j$. Observe that $0 \leq A \leq n$. In order to bound the success probability, we only take the kth iteration of the *while* loop into account where $b_k \leq \exp(A/2) \leq \exp(b_k/2)$.We will show that in this iteration, the probability of a successful transmission in the non-fading model is at least as large as the original one in the Rayleigh-fading model. Using Lemma 2, we observe the probability of success in the Rayleigh-fading model is at most $\frac{q_i}{e^{A/2}} \leq \frac{q_i}{b_k}$.

Let us first consider a single one of the 19 independent iterations. Let X_j be a 0/1 random variable indicating if sender s_j transmits in this iteration. By definition $\mathbf{E}[X_j] = q_j^{(k)}$. Furthermore, set $Z = \sum_{j \neq i} \min\left\{1, \beta_i \bar{S_{j,i}}/\bar{S_{i,i}}\right\} \cdot X_j$.

To make the transmission successful in the non-fading model, we have to have $X_i = 1$ and $\bar{S_{i,i}} \geq \beta_i(\sum_{j \neq i} \bar{S_{j,i}} X_j + \nu)$. To bound the probability of the latter event, we use the assumption that $\bar{S_{i,i}} \geq 2\beta_i \nu$. Therefore it suffices to have $Z < 1/2$, allowing to estimate the probability of this event by Markov's inequality using

$$\mathbf{Pr}\left[Z \geq \frac{1}{2}\right] \leq 2\mathbf{E}[Z] = 2\sum_{j \neq i} \min\left\{1, \beta_i \frac{\bar{S_{j,i}}}{\bar{S_{i,i}}}\right\} \mathbf{E}[X_j]$$

$$= 2\sum_{j \neq i} \min\left\{1, \beta_i \frac{\bar{S_{j,i}}}{\bar{S_{i,i}}}\right\} \cdot \frac{q_j}{4b_k} \leq 2\frac{A}{4b_k} \quad .$$

For the remaining considerations, we distinguish between the two cases $k = 0$ and $k \geq 1$.

In the case $k \geq 1$, we use the fact that $A \leq b_k$ to get that the success probability in the non-fading model in a single iteration is at least

$$q_i^{(k)} \cdot \left(1 - 2\frac{A}{4b_k}\right) \geq \frac{q_i^{(k)}}{2} = \frac{q_i}{8b_k} \quad .$$

We use now the facts that $k \geq 1$ and therefore $b_k \geq \exp(1/8)$ and furthermore that for all $0 \leq x \leq \exp(-1/8)$ we have $1 - (1 - x/8)^{19} \geq x$. This yields that in 19 independent repeats, we get a total success probability of at least

$$1 - \left(1 - \frac{q_i}{8b_k}\right)^{19} \geq \frac{q_i}{b_k} \geq q_i \exp\left(-\frac{A}{2}\right) \quad .$$

As we have already seen, the success probability in the Rayleigh-fading model is at most $q_i \exp(-A/2)$.

For the case $k = 0$, we use that the probability that the transmission is not successful within a single iteration of the inner loop is at most $q_i(1 - 2A)$. This is, the probability that at least one of the 19 independent repeats is successful is at least $1 - (1 - q_i(1 - 2A))^{19} \geq q_i \exp(-A/2)$ for all $0 \leq q_i \leq 1$ because $A \leq 1/4$. □

Taking this theorem into account, we see that we lose at most an $O(\log^* n)$ factor in all approximation guarantees of non-fading algorithms. In particular, the constant-factor capacity-maximization algorithms of the non-fading case provide without any further modification $O(\log^* n)$ approximations in the Rayleigh-fading case.

6. REGRET LEARNING FOR CAPACITY MAXIMIZATION

Another very useful approach to solve capacity maximization was presented by Dinitz [4]. This approach provides a distributed way to solve the problem based on regret-learning techniques. The idea behind regret-learning algorithms is that the algorithm gets feedback in terms of utility depending on the chosen actions of all users and chooses its next action according to this feedback. In the model introduced by Dinitz each user i has in each step the option to attempt a transmission or not. This is, his actions q_i are to send ($q_i = 1$) or not to send ($q_i = 0$). When sending a user gets a utility of 1 for being successful and -1 for not being successful. Not sending at all yields a utility of 0.

For this model, Ásgeirsson and Mitra [2] showed that in the non-fading model the average number of successful transmissions converges to the optimum up to a constant factor. Unfortunately, due to the sequential computation, our transformation cannot be applied here. However, we are able to prove a similar result showing that the expected number of successful transmissions converges to the non-fading optimum up to a constant factor.

Generally, in regret learning, a sequence of action vectors is computed in a decentralized way. In each step, every user i decides which action a_i to take. Depending on his own choice and the one of the other users, he gets a utility $u_i(a_1, \ldots, a_n)$. The choice which action to choose then depends on the history of utilities experienced before. The *external regret* is defined as the difference between the utility of the best single action in hindsight and the summed utility experienced by the algorithm.

Definition 1. *The external regret of user i at time T given a sequence of action vectors $a^{(1)}, \ldots, a^{(T)}$ is*

$$\max_{a_i' \in A_i} \sum_{t=1}^{T} u_i(a_1^{(t)}, \ldots, a_i', \ldots, a_n^{(t)}) - \sum_{t=1}^{T} u_i(a^{(t)}) \quad ,$$

where A_i is the set of possible actions of user i.

So the user regrets what he might have won by switching to one single action for all time steps instead of using the algorithm. An algorithm has the no-regret property if the average regret per time step converges to 0 for the number of time steps T going to ∞. One famous such algorithm is Randomized Weighted Majority due to Littlestone and Warmuth [16]. For this class of algorithms we prove the following theorem.

Theorem 6. *Consider a sequence $q^{(1)}, \ldots, q^{(t)}$ of action vectors such that each user has regret at most $\epsilon \cdot T$. Then the average number of successful transmissions is in $\Omega(\mathrm{OPT} - \epsilon n)$, for* OPT *being the size of the largest feasible set in the non-fading model under uniform transmission powers.*

Theorem 6 directly follows from Lemma 7 and Lemma 8, which we will prove in the remaining part of this section. Note that this theorem together with Theorem 5 yields a factor of $O(\log^* n)$ in comparison to the Rayleigh-fading optimum. Our analysis extends the one for the non-fading case by Ásgeirsson and Mitra [2] which in some parts relies on Dinitz [4]. The results from [2] also show the $F = \Omega(\mathrm{OPT})$ bound for regret learning in the non-fading channel. This highlights the close relationship between the models.

In the Rayleigh-fading model, the utility function itself is stochastic and therefore hard to deal with. In addition, no-regret algorithms must use internal randomization, thus we consider expected utilities and the expected regret here.

We adapt the utility function from Dinitz for an analysis in expectation. It depends on the success probability $Q_i(q_1, \ldots, q_n)$ of link i. Formally, we define the utility of user i to be

$$u_i(q_1, \ldots, q_n) = \begin{cases} 0 & \text{if } q_i = 0, \\ 2 \cdot Q_i\left(q_1, \ldots, q_n\right) - 1 & \text{if } q_i = 1. \end{cases}$$

In the following, we consider a sequence $q^{(1)}, \ldots, q^{(T)}$ that exhibits external regret $\epsilon \cdot T$ for each user $i = 1, \ldots, n$. We define $f_i = \frac{1}{T} \sum_t q_i^{(t)}$ as the fraction of time steps the user chooses $q_i = 1$. Let $F = \sum_i f_i$. We define x_i to be the average success probability per time step with $x_i = \frac{1}{T} \sum_t Q_i^{(t)} \left(q_1^{(t)}, \ldots, q_n^{(t)}\right)$, and we set $X = \sum_i x_i$.

We examine such sequences and at first bound the average number of successful transmissions. It turns out that for ϵ approaching 0 half of the transmissions are successful in the long run. Besides this result, we will show that the average number of transmitting nodes F is in $\Omega(\text{OPT})$. This together shows that the average number of successful transmissions X is in $\Omega(\text{OPT})$.

Lemma 7. $X \leq F \leq 2X + \epsilon n$

Proof. The first inequality follows by definition. For the second inequality, we use the fact that for each user i the regret is at most ϵ. Therefore, always using action $q_i = 0$ can increase the average utility per step by at most ϵ. Formally this means $2 \cdot x_i - f_i \geq -\epsilon$. Taking the sum over all i, we get $2X - F \geq -\epsilon n$. This yields $F \leq 2X + \epsilon n$. □

We have shown a bound for the average number of successful transmissions that depends on the average number of transmitting nodes F. This allows us next to see that F is in $\Omega(\text{OPT})$.

Lemma 8. *Let* OPT *denote the size of the largest feasible set in the non-fading model under uniform transmission powers, then* $F = \Omega(\text{OPT})$.

Proof. In the following $a(j,i)$ denotes the affectance of link j on link i for uniform powers with

$$a(j,i) = \min\left\{1, \frac{\beta \cdot \frac{d(s_i,r_i)^\alpha}{d(s_j,r_i)^\alpha}}{1 - \beta \cdot \nu \cdot d\left(s_i, r_i\right)^\alpha}\right\} \quad .$$

We will denote the summed affectance from other links on link i by

$$a^{(t)}(i) = \sum_{\substack{j \in [n] \\ q_j^{(t)} = 1}} a(j,i) \quad .$$

Let p_i be the fraction of steps in which $a^{(t)}(i) \leq \frac{1}{2}$ and let $\hat{a}(i) = \frac{1}{T} \sum_t a^{(t)}(i)$.

We define the sets $\text{OPT}' = \left\{i \in \text{OPT} \colon f_i < \frac{1}{2} - \epsilon\right\}$ and $\text{OPT}'' = \left\{i \in \text{OPT}' \colon \sum_{j \in \text{OPT}'} a(i,j) \leq 2\right\}$. So all links in OPT'' attempt to transmit in less than a $\frac{1}{2} - \epsilon$ fraction of the time and affect others doing so by at most 2.

If $|\text{OPT} \setminus \text{OPT}'| > |\text{OPT}|/2$, then F would be at least $\left(\frac{1}{2} - \epsilon\right) \cdot |\text{OPT} \setminus \text{OPT}'|$ and therefore in $\Omega(\text{OPT})$.

So we consider $|\text{OPT}'| \geq |\text{OPT}|/2$ for the rest of the proof. Using [2, Lemma 8], we see $|\text{OPT}''| \geq |\text{OPT}|/4$. Therefore, it is sufficient to show $F = \Omega(|\text{OPT}''|)$ and so we just need to consider links $i \in \text{OPT}''$.

We consider the utility gain for link i by switching to action $q_i = 1$ throughout every step. In an f_i fraction of the steps nothing changes. In at least a $p_i - f_i$ fraction of the steps, link i could have been successful but did not transmit in the original sequence. As the affectance is at most $1/2$, we conclude from Lemma 2 that the success probability in these steps is at least $\exp(-1/2)$. For the remaining steps, we estimate the probability simply by 0. Therefore, the utility gain is at least $(p_i - f_i) \cdot 2\exp(-1/2) - (1 - f_i) \leq \epsilon$.

This yields for all $i \in \text{OPT}''$ and $\epsilon \leq 0.04$ that

$$\begin{aligned} p_i &\leq f_i + \frac{\epsilon + 1 - f_i}{2\exp(-1/2)} \\ &\leq \frac{1}{2}\left(1 + \frac{\exp(1/2)}{2}\right) + \frac{\epsilon \cdot \exp(1/2)}{2} \leq \frac{19}{20} \quad , \end{aligned}$$

because $f_i \leq 1/2$. For $\hat{a}(i)$, we now get by definition of q_i

$$\hat{a}(i) \geq q_i \cdot 0 + (1 - q_i) \cdot \frac{1}{2} \geq \frac{1}{20} \cdot \frac{1}{2} = \frac{1}{40} \quad .$$

Hence, we have

$$\hat{a}(i) = \sum_{j \in [n]} f_j a(j,i) \geq \frac{1}{40} \text{ for all } i \in \text{OPT}''.$$

Taking the sum of all resulting inequalities, we get

$$\sum_{i \in \text{OPT}''} \sum_{j \in [n]} f_j a(j,i) \geq \frac{|\text{OPT}''|}{40}$$

or, equivalently,

$$\sum_{j \in [n]} f_j \left(\sum_{i \in \text{OPT}''} a(j,i)\right) \geq \frac{|\text{OPT}''|}{40} \quad .$$

With [2, Lemma 11] we have that $\sum_{i \in \text{OPT}''} a(j,i) = O(1)$ for all $j \in [n]$ and hence

$$\sum_{j \in [n]} f_j = \Omega(|\text{OPT}''|) \quad . \qquad \square$$

Due to Lemma 7 and 8, for any no-regret algorithm the number of successful transmissions needs to converge to a constant fraction of the non-fading optimum. This proves Theorem 6.

7. SIMULATION RESULTS

In the sections before, we showed a close relation between the Rayleigh-fading and the non-fading models in theory. While bounds are given asymptotically for worst-case instances, our theoretical results can also be verified in simulations.

In particular, we consider the performance of an ALOHA-like protocol and the no-regret capacity-maximization algorithm. Simulations are carried out on random networks constructed by randomly placing receivers on a 1000×1000 plane. Each corresponding sender is placed by choosing the angle and the distance to the receiver uniformly at random from a fixed interval. This way, a minimal and a maximal distance between sender and receiver can be specified.

Comparing the Rayleigh-fading and the non-fading model the simulations show that the number of successful transmissions under uniform powers behave similarly when the

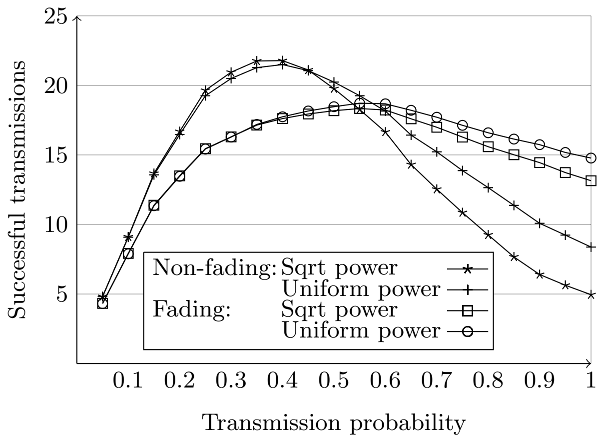

Figure 1: Number of successful transmissions for different transmission probabilities under square-root and uniform power assignment and under the Rayleigh-fading and non-fading SINR model.

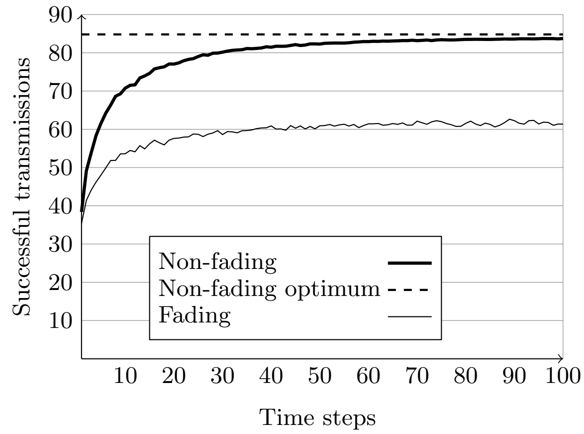

Figure 2: Number of successful transmissions under the Rayleigh-fading and non-fading model when applying a no-regret algorithm.

sending probabilities are chosen uniformly, see Figure 1. The simulation was done over 40 different networks with 100 links each and 25 different seeds for the randomizer to determine which links transmit. The SINR parameters were set to $\beta = 2.5$, $\alpha = 2.2$, and $\nu = 4 \cdot 10^{-7}$. The power for the uniform power assignment and the power from which the square-root power assignment scales was set to 2. For the Rayleigh-fading channel we additionally used 10 different seeds to determine whether a transmission is successful. The distance between a sender and the corresponding receiver was chosen between 20 and 40.

Figure 1 shows the number of successful transmissions averaged over all those runs. Neither the Rayleigh-fading model nor the non-fading model always predicts more success than the other one. The Rayleigh probability distribution leads to a smoothed curve compared to the non-fading model. This is due to the fact that even when the SINR constraint is not fulfilled in the non-fading model, the success probability in the Rayleigh-fading model still remains positive. On the other hand, when a transmission is definitely successful in the non-fading SINR model there is some probability for being not successful in the Rayleigh-fading model. The general characteristics of the curves are the same and show that the Rayleigh-fading and the non-fading model behave alike.

Choosing the optimal set of sending links under uniform powers, we reach on average 49.75 successful transmissions in those networks.

The similarity can also be seen when taking a look at no-regret algorithms. Here we analyzed a version of the Randomized Weighted Majority Algorithm of Littlestone and Warmuth [16]. The weights are initialized with 1 and multiplied by $(1-\eta)^{l_i}$ in every time step, where l_i is the loss of not sending ($i = 0$) or sending ($i = 1$). The loss for sending and not being received is 1 and the loss of not sending at all is 0.5. In all other cases the loss is 0. These losses correspond to the utility function used in Section 6. The factor η starts with $\sqrt{0.5}$ and is multiplied by $\sqrt{0.5}$ every time the number of time steps reaches the next power of 2.

For the simulation shown in Figure 2 we used different networks with 200 links, distances between 0 and 100, $\beta = 0.5$, $\alpha = 2.1$, and $\nu = 0$. The other settings remained as before.

The results behave in the same way as observed by Ásgeirsson and Mitra [2] in their simulations. The Rayleigh-fading model shows more fluctuations due to its stochastic nature. We can also see that the no-regret algorithm converges quite quickly near the optimum of the non-fading model. The number of time steps needed for convergence depends on the specific instance, but a good performance can already be seen after 30 to 40 timesteps.

8. DISCUSSION AND OPEN PROBLEMS

In this paper we showed that from an algorithmic point of view, the non-fading and the Rayleigh-fading model behave similarly in theory as well as in simulations. We regard this as a promising result because it indicates that existing results on approximation algorithms within non-fading models seem to apply more generally. Turning to a different, more realistic scenario does not create a fundamentally new situation as was the case when shifting from graph-based interference models to SINR-based ones.

Future research could take two different directions from this point. On the one hand, it could focus on the similarities, e.g., by improving the obtained bounds. Considering a particular situation, the $O(\log^* n)$-factor in Theorem 5 might be reduced to a constant, which we were not able to prove in general. Futhermore, the similarities could be exploited to take the best of the two worlds, in order to derive more sophisticated, hopefully distributed algorithms. On the other hand, also the differences could be taken into account. For example, the regret-learning simulation in the Rayleigh-fading model reaches a smaller capacity. It would be interesting to see if this is a general effect of the stochastic model or under which conditions this behavior can be observed.

9. REFERENCES

[1] Matthew Andrews and Michael Dinitz. Maximizing capacity in arbitrary wireless networks in the SINR

model: Complexity and game theory. In *Proc. 28th IEEE Conf. Computer Communications (INFOCOM)*, pages 1332–1340, 2009.

[2] Eyjolfur Ingi Asgeirsson and Pradipta Mitra. On a game theoretic approach to capacity maximization in wireless networks. In *Proc. 30th IEEE Conf. Computer Communications (INFOCOM)*, pages 3029–3037, 2011.

[3] Swades De, Dimitrios Pados, Chunming Qiao, and Mainak Chatterjee. Topological and MAI constraints on the performance of wireless CDMA sensor networks. In *Proc. 23rd IEEE Conf. Computer Communications (INFOCOM)*, volume 1, pages 256–265, 2004.

[4] Michael Dinitz. Distributed algorithms for approximating wireless network capacity. In *Proc. 29th IEEE Conf. Computer Communications (INFOCOM)*, pages 1397–1405, 2010.

[5] Alexander Fanghänel, Thomas Kesselheim, Harald Räcke, and Berthold Vöcking. Oblivious interference scheduling. In *Proc. 28th Symp. Principles of Distributed Computing (PODC)*, pages 220–229, 2009.

[6] Alexander Fanghänel, Thomas Kesselheim, and Berthold Vöcking. Improved algorithms for latency minimization in wireless networks. *Theoret. Comput. Sci.*, 412(24):2657–2667, 2011.

[7] Olga Goussevskaia, Magnús Halldórsson, Roger Wattenhofer, and Emo Welzl. Capacity of arbitrary wireless networks. In *Proc. 28th IEEE Conf. Computer Communications (INFOCOM)*, pages 1872–1880, 2009.

[8] Matthias Grossglauser and David Tse. Mobility increases the capacity of ad hoc wireless networks. *IEEE/ACM Trans. Netw.*, 10:477–486, 2002.

[9] Piyush Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Trans. Inf. Theory*, 46(2):388–404, 2000.

[10] Martin Haenggi. On routing in random rayleigh fading networks. *IEEE Trans. Wireless Commun.*, 4(4):1553–1562, 2005.

[11] Magnús Halldórsson. Wireless scheduling with power control. In *Proc. 17th European Symposium on Algorithms (ESA)*, pages 361–372, 2009.

[12] Magnús Halldórsson and Pradipta Mitra. Nearly optimal bounds for distributed wireless scheduling in the SINR model. In *Proc. 38th Intl. Coll. Automata, Languages and Programming (ICALP)*, volume 2, pages 625–636, 2011.

[13] Magnús Halldórsson and Pradipta Mitra. Wireless capacity with oblivious power in general metrics. In *Proc. 22nd Symp. Discrete Algorithms (SODA)*, pages 1538–1548, 2011.

[14] Thomas Kesselheim. A constant-factor approximation for wireless capacity maximization with power control in the SINR model. In *Proc. 22nd Symp. Discrete Algorithms (SODA)*, pages 1549–1559, 2011.

[15] Thomas Kesselheim and Berthold Vöcking. Distributed contention resolution in wireless networks. In *Proc. 24th Intl. Symp. Distributed Computing (DISC)*, pages 163–178, 2010.

[16] Nick Littlestone and Manfred Warmuth. The weighted majority algorithm. *Inform. Comput.*, 108:212–261, 1994.

[17] Xiaowen Liu and Martin Haenggi. Throughput analysis of fading sensor networks with regular and random topologies. *EURASIP J. Wireless Comm. Netw.*, 5(4):554–564, 2005.

[18] Thomas Moscibroda and Roger Wattenhofer. The complexity of connectivity in wireless networks. In *Proc. 25th IEEE Conf. Computer Communications (INFOCOM)*, pages 1–13, 2006.

[19] James Roberts and Timothy Healy. Packet radio performance over slow rayleigh fading channels. *IEEE Trans. Commun.*, 28(2):279–286, 1980.

[20] John Silvester and Leonard Kleinrock. On the capacity of multihop slotted aloha networks with regular structure. *IEEE Trans. Commun.*, 31(8):974–982, 1983.

[21] Stavros Toumpis and Andrea Goldsmith. Capacity regions for wireless ad hoc networks. *IEEE Trans. Wireless Commun.*, 2(4):736–748, 2003.

[22] Liang-Liang Xie and P. R. Kumar. A network information theory for wireless communication: Scaling laws and optimal operation. *IEEE Trans. Inf. Theory*, 50(5):748–767, 2004.

APPENDIX

A. PROOF OF THEOREM 1

PROOF. Two events have to occur for successful transmission. On the one hand, sender s_i has to decide to transmit. By definition this probability is q_i. On the other hand, the SINR for the transmission must be large enough. For the latter event, Liu and Haenggi [17] derived a formula, which can be generalized to our model as follows.

The cumulated interference our transmission is exposed to is given by $I_i = \sum_{j \neq i} S_{j,i} \cdot X_j$, where X_j denotes the 0/1 random variable whether sender s_j makes a transmission attempt. The transmission is successful if $S_{i,i} \geq \beta(I_i + \nu)$. Fixing I_i, we can use the fact that $S_{i,i}$ is expontially distributed to get

$$\mathbf{Pr}\left[S_{i,i} \geq \beta(I_i + \nu) \mid I_i = x\right] = \exp\left(-\frac{\beta(x+\nu)}{\bar{S_{i,i}}}\right) \quad .$$

Taking the expectation over I_i, we get

$$\begin{aligned}
&\mathbf{Pr}\left[S_{i,i} \geq \beta(I_i + \nu)\right] \\
&= \mathbf{E}\left[\exp\left(-\frac{\beta(I_i+\nu)}{\bar{S_{i,i}}}\right)\right] \\
&= \mathbf{E}\left[\exp\left(-\frac{\beta(\sum_{j \neq i} S_{j,i} \cdot X_j + \nu)}{\bar{S_{i,i}}}\right)\right] \\
&= \exp\left(-\frac{\beta \cdot \nu}{\bar{S_{i,i}}}\right) \mathbf{E}\left[\prod_{j \neq i} \exp\left(-\frac{\beta \cdot S_{j,i} \cdot X_j}{\bar{S_{i,i}}}\right)\right] \\
&= \exp\left(-\frac{\beta \cdot \nu}{\bar{S_{i,i}}}\right) \prod_{j \neq i} \mathbf{E}\left[\exp\left(-\frac{\beta \cdot S_{j,i} \cdot X_j}{\bar{S_{i,i}}}\right)\right]
\end{aligned}$$

Since $S_{j,i}$ and X_j are independent, we have

$$\mathbf{E}\left[\exp\left(-\frac{\beta \cdot S_{j,i} \cdot X_j}{\bar{S}_{i,i}}\right)\right]$$
$$= q_j \cdot \mathbf{E}\left[\exp\left(-\frac{\beta \cdot S_{j,i}}{\bar{S}_{i,i}}\right)\right] + (1 - q_j) \ .$$

Using now the fact that $S_{j,i}$ is exponentially distributed, we get

$$\mathbf{E}\left[\exp\left(-\frac{\beta \cdot S_{j,i}}{\bar{S}_{i,i}}\right)\right]$$
$$= \int_0^\infty \frac{1}{\bar{S}_{j,i}} \exp\left(-\frac{x}{\bar{S}_{j,i}}\right) \cdot \exp\left(-\frac{\beta \cdot x}{\bar{S}_{i,i}}\right) dx$$
$$= \frac{1}{1 + \beta \frac{\bar{S}_{j,i}}{\bar{S}_{i,i}}} \ .$$

This yields that $\mathbf{Pr}\left[S_{i,i} \geq \beta(I_i + \nu)\right]$ is

$$\exp\left(-\frac{\beta \cdot \nu}{\bar{S}_{i,i}}\right) \prod_{j \neq i} \left(q_j \cdot \frac{1}{1 + \beta \frac{\bar{S}_{j,i}}{\bar{S}_{i,i}}} + (1 - q_j)\right) \ ,$$

yielding the claim. $\square$