Convergence Time of Power-Control Dynamics*

Johannes Dams Martin Hoefer Thomas Kesselheim

April 20, 2011

Abstract

We study two (classes of) distributed algorithms for power control in a general model of wireless networks. There are n wireless communication requests or *links* that experience interference and noise. To be successful a link must satisfy an SINR constraint. The goal is to find a set of powers such that all links are successful simultaneously. A classic algorithm for this problem is the fixed-point iteration due to Foschini and Miljanic [8], for which we prove the first bounds on worst-case running times – after roughly $O(n \log n)$ rounds all SINR constraints are nearly satisfied. When we try to satisfy each constraint exactly, however, convergence time is infinite. For this case, we design a novel framework for power control using regret learning algorithms and iterative discretization. While the exact convergence times must rely on a variety of parameters, we show that roughly a polynomial number of rounds suffices to make every link successful during at least a constant fraction of all previous rounds.

1 Introduction

A key ingredient to the operation of wireless networks is successful transmission in spite of interference and noise. Usually, a transmission is successful if the received signal strength is significantly stronger than the disturbance due to decayed signals of simultaneous transmissions and ambient noise. This condition is frequently expressed by the *signal-to-interference-plus-noise ratio (SINR)*. Over the last decade, a large amount of research work has studied the problem of throughput or *capacity maximization*, i.e., determining the maximum number of wireless transmissions that can be executed successfully in a network in parallel. Very recently, algorithms for capacity maximization are starting to receive interest also from an analytical and theoretical point of view. Most of the algorithms proposed and analyzed so far require a strong central authority managing the access of all devices to the spectrum. In addition, most works neglect *power control*, i.e., the ability of modern wireless devices to allow their transmission powers to be set by software.

*Department of Computer Science, RWTH Aachen University, Germany. {dams,mhoefer,kesselheim}@cs.rwth-aachen.de. This work has been supported by DFG through UMIC Research Centre, RWTH Aachen University, and grant Ho 3831/3-1.

Power control has two main advantages. On the one hand, battery life can be increased by using only minimal powers that are necessary to guarantee reception. On the other hand, reduced transmission power causes less interference, and thereby the throughput of a wireless network can increase significantly when using power control.

In this paper, we study spectrum access with power control in a network of wireless devices. We consider a network consisting of n *links*, i.e., sender/receiver pairs. Each sender attempts a successful transmission to its corresponding receiver using a transmission power. The chosen power has to be large enough to compensate the interference and ambient noise. In contrast, choosing a smaller transmission power is desirable as it results in less energy consumption. We investigate distributed algorithms to find transmission powers in order to make links successful as quickly possible. A standard assumption in the analysis of power-control problems of this kind is the existence of a solution, in which all transmissions are successful. For networks, in which this assumption does not hold, it is possible to combine the algorithms with approaches that solve the additional scheduling problem [13].

A simple and beautiful distributed algorithm for the power-control problem is the fixed-point iteration due to Foschini and Miljanic [8]. In each step every sender sets his power to the minimum power that was required to overcome interference and noise in the last round. It can be shown that powers converge to a feasible assignment, even if updates are not simultaneous [17]. The obtained power assignment is minimal in the sense that it is component-wise smaller than all other feasible assignments. It is known that this algorithm converges at a geometric rate [12] in a numerical sense. However, to the best of our knowledge, no convergence results in the sense of quantitative worst-case running times have been shown, neither for this nor for other distributed algorithms.

In this paper, we investigate two classes of distributed algorithms for power control and analyze the dependencies of running time and solution quality on several parameters of the structure of the instance. For example, our analysis of the Foschini-Miljanic fixed-point iteration in Section 2 uses the largest eigenvalues of the normalized gain matrix and the degree, to which the SINR constraint is fulfilled. Assuming that both these parameters are constant, our first main result (Theorem 1) shows that the FM iteration achieves polynomial convergence time. In particular, starting from all powers set to 0, for any constant $\delta > 0$ we reach in $O(n \log n)$ steps a power assignment that satisfies the SINR constraint of every link by a factor of at least $1 - \delta$.

It is easy to see that the FM iteration might never reach the fixed point if we start with all powers set to 0. Thus, if we insist on links satisfying the SINR constraint exactly, we get an infinite convergence time during which all links remain unsuccessful. To overcome this problem, in Section 3 we introduce a novel technique to compute power assignments employing distributed regret-learning algorithms. For algorithms that guarantee no swap regret [4], we can also guarantee convergence to the fixed point. The convergence properties rely on our analysis of the FM iteration and depend additionally on the position of the fixed point compared to noise vector and maximum allowed power. Assuming these ratios are bounded by a constant, our second main result (Theorem 8) is that for every constant $\epsilon > 0$ after a polynomial number of steps, we can reach a situation in which every request has been successful with respect to the exact SINR constraint during at least a $(1-\epsilon)$ fraction of the previous steps. Our regret learning technique has the advantage of being applicable also to instances, in which not all links can be successful simultaneously.

1.1 Formal Problem Statement

We consider transmissions in general interference models based on SINR. If the sender of link j emits a signal at power p_j, then it is received by the receiver of link i with strength $g_{i,j} \cdot p_j$, where $g_{i,j}$ is called the *gain*. This includes the well-known special case of the *physical model*, where the gain depends polynomially on the distance between sender and receiver. The transmission within link i is successful if the SINR constraint

$$\frac{g_{i,i} \cdot p_i}{\sum_{j \neq i} g_{i,j} \cdot p_j + \nu} \geq \beta$$

is fulfilled, i.e., the SINR is above some threshold β. In the power control problem, our 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. More formally, let the *normalized gain matrix* C be the $n \times n$ matrix defined by $C_{i,i} = 0$ for all $i \in [n]$ and $C_{i,j} = \beta g_{i,j}/g_{i,i}$ for $i \neq j$. The *normalized noise vector* η is defined by $\eta_i = \beta\nu/g_{i,i}$. The task is to find a vector p such that $p \geq C \cdot p + \eta$. Note that throughout this paper, we use $\leq$ and $\geq$ to denote the respective component-wise inequality.

The set of all feasible power assignments is a convex polytope. If it is non-empty, there is a unique vector p^* satisfying $p^* = C \cdot p^* + \eta$. In a full-knowledge, centralized setting, the optimal power vector p^* can simply be computed by solving the linear equation system $p^* = C \cdot p^* + \eta$. However, a wireless network consists of independent devices with distributed control and the matrix C is not known. We assume the devices can only make communication attempts at different powers and they receive feedback in the form of the achieved SINR or (in an advanced scenario) only whether the transmission has been successful or not.

For the scenario in which the achieved SINR is known after each transmission attempt, the FM iteration is $p^{(t+1)} = C \cdot p^{(t)} + \eta$, where the achieved and the target SINR are needed to run this iteration. Foschini and Miljanic showed that the sequence of vectors $p^{(t)}$ converges to p^* as t goes to infinity. One can show that the existence of $p^* \geq 0$ with $p^* \leq C \cdot p^*$ implies that the modulus of all eigenvalues of C must be strictly less than 1. In our analyses, we will refer to the maximal modulus of an eigenvalue as $\lambda_{\max}$.

For the regret-learning technique we assume that each link i uses a no-regret learning algorithm to select from a suitably defined discrete subset power values in an interval $[0, p_i^{\max}]$. So $p_i^{\max}$ is the maximal power level user i might choose. Let Φ be a set of measurable functions such that each $\phi \in \Phi$ is a map $\phi\colon [0, p_i^{\max}] \to [0, p_i^{\max}]$. Given a sequence of power vectors $p^{(1)}, \ldots, p^{(T)}$, the Φ-regret link i encounters is

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

where u_i is a suitable utility function defined below. For our analyses, we consider two cases for the set Φ. For *external regret* each $\phi \in \Phi$ maps every power value to a single power p_ϕ. In contrast, to define *swap regret*, Φ contains all measurable functions. An infinite sequence is called no Φ-regret if $R_i^\Phi(T) = o(T)$. An algorithm that produces a no Φ-regret sequence is a no-Φ-regret algorithm.

We will see that under the utility functions we assume, there are distributed no-Φ-regret algorithms. It suffices for each user to only know after each transmission attempt if it has been successful.

1.2 Related Work

For about two decades, wireless networks with power control have been extensively studied. While at first research focused on engineering aspects, recently the topic has attracted interest among computer scientists. Algorithmic research so far focused on scheduling problems, where for a given network of senders and receivers the goal is to select a maximum feasible subset (the "independent set" problem) or to partition the links into the minimal number of feasible subsets (the "coloring" problem). Allowing a scheduling algorithm to choose powers has a significant impact on size and structure of links scheduled simultaneously in practice, as was shown by [15].

As a consequence, much effort has been put into finding algorithms for scheduling with power control. More recently, theoretical insights on the problem are starting to develop [9, 14, 6]. In particular, the independent set problem with power control has been shown to be NP-hard [1]. In most related work, however, the algorithmic power control problem is neglected by setting powers according to some "oblivious" schemes, e.g., using some polynomial depending on distance between sender and receiver. For independent set and coloring problems in metric spaces, usually the mean or square-root function achieves best results [7, 10, 11]. In addition, there are distributed approaches for the independent set problem using no-regret learning and uniform power [5, 3]. In general, there are strong lower bounds when using oblivious power, as algorithms provide only trivial $\Omega(n)$-approximations in instances with n links of greatly varying length [7]. Power control can significantly improve this condition as exemplified by the recent centralized constant-factor approximation algorithm for the independent set problem by Kesselheim [13]. Being a centralized combination of scheduling and power control, this algorithm is rather of fundamental analytical interest and of minor relevance in heavily distributed systems like wireless networks.

Distributed algorithms exist especially for the power control problem without the scheduling aspect. In this case, a feasible power assignment is assumed to exist that makes all links feasible. The goal is to find feasible power levels with minimal power consumption. This problem can be seen as a second step after scheduling and can be solved in a centralized fashion by solving a system of linear equations as noted above. Foschini and Miljanic [8] solved the problem using a simple iterative distributed algorithm. They showed that their iteration converges from each starting point to a fixed point if it exists.

Extending this, Yates [17] proved convergence for a general class of iterative algorithms including also a variant for limited transmission powers and an iteration, in which users update powers asynchronously. Besides this, Huang and Yates [12] proved that all these algorithms converge geometrically. This means that the norm distance to the fixed point in time step t is given by a^t for some constant $a < 1$. However, this is only a bound on the convergence rate in the numerical sense and does not imply a bound on the time until links actually become successful.

In addition, more complex iterative schemes have been proposed in the literature. For a general survey about these algorithms and the power control problem, see Singh and Kumar [16].

2 Convergence Time of the Foschini-Miljanic iteration

In this section, we analyze the convergence time of the Foschini-Miljanic iteration with $p^{(t)} = C \cdot p^{(t-1)} + \eta$. It will turn out to be helpful to consider the closed form variant

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

The iteration will never actually reach the fixed point, although getting arbitrarily close to it. However, during the iteration the SINR will converge to the threshold β. For each $\delta > 0$, there is some round T from which the SINR will never be below $(1-\delta)\beta$. Since maximizing the SINR is the main target, we strive to bound the time T until each transmission is "almost" feasible. That is, the SINR is above $(1-\delta)\beta$. For this purpose, it is sufficient that the current vector 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 time is independent of the values of p^* or η. The only parameter related to the instance is $\lambda_{\max}$ the maximum eigenvalue of C, which has to occur as for $\lambda_{\max} = 1$ no fixed point can exist at all. Assuming it to be constant, we show that after $O(n \log n)$ rounds we reach a power assignment that satisfies the SINR constraint of every link by a factor of at least $1 - \delta$.

Theorem 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. Define the following auxiliary matrix $M = C^m$, where $m = \left\lceil \log \frac{1}{3n} / \log \lambda_{\max} \right\rceil$. As we can see, the modulus of all eigenvalues of M is bounded by $\frac{1}{3n}$. Furthermore, defining $\eta' = \sum_{k=0}^{m-1} C^k \eta$, we have $p^{(mt')} = \sum_{k=0}^{t'-1} M^k \eta'$. This also implies $p^* = \sum_{k=0}^{\infty} M^k \eta'$.

Now we consider the characteristical polynomial of M in expanded as well as in factored 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, we have $|b_j| \leq \frac{1}{3n}$ for all $j \in [n]$. The a_i values can be computed from the b_i values using

$$a_i = \sum_{\substack{S \subseteq [n] \\ |S| = n-i}} \prod_{j \in S} (-b_j) \quad .$$

For the modulus this gives us

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

This yields the following bound for their sum

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

We now use the fact that $\chi_M(M) = 0$. 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

$$\begin{aligned} M^n p^* = \sum_{k=n}^{\infty} M^k \eta' &= \sum_{k=0}^{n-1} M^k \eta' \left(-\sum_{i=0}^{k} a_i\right) + \sum_{k=n}^{\infty} M^k \eta' \left(-\sum_{i=0}^{n-1} a_i\right) \\ &\leq \left(\sum_{i=0}^{n-1} |a_i|\right) \sum_{k=0}^{\infty} M^k \eta' \leq \frac{1}{2} \sum_{k=0}^{\infty} M^k \eta' = \frac{1}{2} p^* \quad . \end{aligned}$$

Now consider $t \geq m \cdot n \cdot \log \frac{1}{\delta}$. We have $p^* - p^{(t)} = C^t p^* \leq M^{n \log \frac{1}{\delta}} p^* \leq \delta p^*$. This proves the theorem. □

One can see that this bound is almost tight as there are instances where $\Omega(n)$ rounds are needed. A simple example can be given as follows. Let C be defined by $C_{i+1,i} = 1$ for all i and all other entries 0, $\eta = (1, 0, \ldots, 0)$. The only eigenvalue of this matrix is 0. However, it takes n rounds until the 1 of the first component has propagated to the nth component and the fixed point is reached.

These instances require a certain structure in the values of p^* and η. As a second result, we would like to present a bound independent of n and for every possible starting point $p^{(0)}$ that takes p^* and η into consideration.

Theorem 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*

$$T = \frac{\log \delta - \log \max_{i \in [n]} \left| \frac{p_i^{(0)}}{p_i^*} - 1 \right|}{\log \max_{i \in [n]} \left| 1 - \frac{\eta_i}{p_i^*} \right|} \quad .$$

Proof. We consider the weighted maximum norm, which has been used by Huang and Yates [12] before. We use the entries of p^* as weights

$$\|x\| = \max_{i\in[n]} \frac{|x_i|}{p_i^*} \ .$$

The induced matrix norm of a matrix M is now given by

$$\|M\| = \max_{i\in[n]} \frac{1}{p_i^*} \sum_{j\in[n]} |M_{i,j}|\, p_j^* \ .$$

In particular, we have for matrix C that $Cp^* + \eta = p^*$, that is $(Cp^*)_i = p_i^* - \eta_i$. This yields for the matrix norm of C

$$\|C\| = \max_{i\in[n]} \frac{1}{p_i^*} \sum_{j\in[n]} C_{i,j} p_j^* = \max_{i\in[n]} \frac{1}{p_i^*} (Cp^*)_i = \max_{i\in[n]} \left| 1 - \frac{\eta_i}{p_i^*} \right| \ .$$

If $t \geq T$, using Equation 1 we get for the distance of $p^{(t)}$ and p^*

$$\|p^{(t)} - p^*\| = \|C^t(p^{(0)} - p^*)\| \leq \|C\|^t \cdot \|p^{(0)} - p^*\| = \|C\|^t \cdot \max_{i\in[n]} \left| \frac{p_i^{(0)}}{p_i^*} - 1 \right| \leq \delta \ .$$

So, for all $i \in [n]$, we have $|p_i^{(t)} - p_i^*| \leq \delta p_i^*$. This is $(1-\delta)p_i^* \leq p_i^{(t)} \leq (1+\delta)p_i^*$. □

As assumed for the original FM iteration, we also focused on the case that powers can be chosen arbitrarily high so far. 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 [17]. 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.

3 Power Control via Regret Learning

The fixed-point approach analyzed above has some major drawbacks. For example, in many sequences – in particular the ones starting from 0 – the target SINR is never reached, because all powers increase in each step and therefore they are always too small. Another drawback 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 lacking robustness. We assumed the fixed-point to exist. 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.

In order to overcome these drawbacks, we design a different approach based on regret learning. Here each link is a user striving to have a successful transmission but using the least power possible. The user is assumed to decide which power $p_i \in [0, p_i^{\max}]$ to use based on an utility function. In particular, we assume that each user gets zero utility if the SINR is below the threshold and a positive one otherwise. However, this utility increases when using a smaller power. Formally, we assume utility functions of the following form:

$$u_i(p) = \begin{cases} f_i(p_i) & \text{if user } i \text{ is successful with } p_i \text{ against } p_{-i} \\ 0 & \text{otherwise} \end{cases}$$

where $f_i \colon [0, p_i^{\max}] \to [0, p_i^{\max}]$ 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.

The utility functions have to be considered this way in order to capture the SINR constraint appropriately. On the one hand, each user's maximum is at the point where the SINR condition is exactly met. On the other hand and at least as important, for each user having a successful transmission is always better than an unsuccessful one. This property cannot be modeled by a continuous function.

As a consequence, we can ensure that all no-swap-regret sequences converge to the optimal power vector p^*. 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.

As a first result, we can see that the only possibility that all links encounter zero swap regret is the sequence only consisting of p^*.

Proposition 3. *Given any sequence $p^{(1)}, \ldots, p^{(T)}$ such that the swap regret for each user is 0, then $p^{(t)} = p^*$ for all t.*

Proof. For each user i let $\hat{p}_i = \max_{t \in [T]} p_i^{(t)}$. Now assume that $\hat{p} \leq p^*$ does not hold. This means there is some user i for which $\hat{p}'_i := (C \cdot \hat{p} + \eta)_i < \hat{p}_i$. This user encounters non-zero swap regret because he could always use $\hat{p}'_i$ instead of $\hat{p}_i$. The user would still be successful in the same steps as before but get a higher utility each time he chose $\hat{p}_i$. Since this is a contradiction we have $\hat{p} \leq p^*$.

Now let $\check{p}_i = \min_{t \in [T]} p_i^{(t)}$. Assume that $\check{p} \geq p^*$ does not hold. This implies that for some user $\check{p}_i < (C \cdot \check{p} + \eta)_i$. So user i is never successful when using power $\check{p}_i$ but would always be with p_i^* (since $\hat{p} \leq p^*$). This is again a contradiction because 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 $\hat{p} = \check{p} = p^*$. □

In contrast, 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.

Proposition 4. *For $\beta \geq 1$, $p_i^{\max} < \infty$ and $f_i(p_i^{\max}) > \frac{1}{2}$, there is an instance with a fixed point and a no-external-regret sequence in which only a $2/n$ fraction of all links is ever feasible.*

Proof. (Sketch) We consider a "nested pairs" instance (c.f. [7]) appropriately scaled to allow 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 that 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: All links except for the two inner ones always play action $p_i = 0$. In odd steps, the two inner links play action $p^{\max}$, in even steps they play $p^{\max}/2$. We claim that the regret for each player is at most 0. For the outer links this is clear because they cannot get through at all, even not with $p^{\max}$. For the inner links there is some smallest action p' allowing them to be feasible in all steps. We have $p^{\max}/2 \leq p' \leq p^{\max}$. The regret compared to this action after T steps is $f(p')T - \frac{1}{2}(f(p^{\max}) + f(\frac{1}{2}p^{\max}))T$. Note that this can evaluate to 0 by a suitable choice of f. □

4 Computing No-Swap-Regret Sequences

Observe that in our case, the users are given an infinite number of possible choices. Furthermore, in order to capture the SINR threshold appropriately the utility functions have to be modeled as non-continuous. Unfortunately, this yields standard no-swap-regret algorithms cannot be used in this scenario because, to the best of our knowledge, they require a finite number of actions [2] or convex action spaces and continuous and concave utility functions [18].

Luckily, no-regret sequences can be computed in a distributed way nevertheless. In order to achieve swap regret $\epsilon \cdot T$, we execute an arbitrary existing no-swap-regret algorithm on a finite subset of the available powers, which is chosen depending on ϵ. 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. Bounding the loss due to the restriction on the right borders, we can prove the following theorem.

Theorem 5. *Let Algorithm $\mathcal{A}$ be a no-swap-regret algorithm for finite action spaces achieving swap regret at most $O(T^a \cdot N^b)$ after T rounds in case of N actions for suitable constants $0 \leq a < 1$, $b \geq 0$. Then $\mathcal{A}$ can be used to compute a no-swap-regret sequence for power control, achieving swap-regret at most $O(T^{\frac{a+b}{1+b}})$ in T steps.*

Proof. We exploit the structure of our utility functions. 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. Now observe that the utility at the right border of each interval is at most $S_i p_i^{\max}/N$ worse than the maximum in the respective interval, where $S_i = \max_{p_i,h} \frac{f(p_i) - f(p_i+h)}{h}$. This is for all $x \in [kp_i^{\max}/N, (k+1)p_i^{\max}/N]$, we have $u_i(x, p_{-i}) \leq u_i((k+1)p_i^{\max}/N, p_{-i}) + S_i p_i^{\max}/N$.

If the number of steps T is known in advance this allows us to construct the following algorithm. We set

$$N = \left\lceil T^{\frac{1-a}{1+b}} \right\rceil$$

and run $\mathcal{A}$ using only the finite strategy set $\{p_i^{\max}/N, 2p_i^{\max}/N, \ldots, p_i^{\max}\}$ of size N. If optimal strategies were also restricted to this finite set, the resulting swap regret would be at most $O(T^a N^b)$. Due to the restriction to the finite set, we additionally lose at most $S_i p_i^{\max}/N$ in each step. So the resulting regret is at most $O(T^a N^b) + T S_i p_i^{\max}/N = O(T^{\frac{a+b}{1+b}})$.

If T is not known in advance the "doubling trick" also works for our algorithm. Starting with $T = 1$, it is executed for T steps with the respective T value, which is doubled afterwards. This way only a constant factor is lost in comparison to the case where T is known. □

In particular, if each link knows after each step which powers would have made it successful, we can use the $O(\sqrt{TN \log N})$ full-information algorithm proposed by Blum and Mansour [4] for the following result.

Corollary 6. *There is an algorithm achieving swap regret* $O(T^{\frac{3}{4}})$.

If each link only gets to know if the transmission at the actually chosen power suffices, it can nevertheless compute the value of the utility function for the chosen power. Therefore, in this case we are in the partial-feedback model. Here, we can apply the $O(N\sqrt{T \log N})$ algorithm by Blum and Mansour [4] to build the following algorithm.

Corollary 7. *There is an algorithm achieving swap regret* $O(T^{\frac{4}{5}})$ *that only needs to know if the transmissions carried out were successful.*

5 Convergence of No-Swap-Regret Sequences

So far, we have seen how to compute no-swap-regret sequences. In this section, the result is complemented by a quantitative analysis of a no-swap-regret sequence. We see 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. In contrast, starting from certain vectors in the FM iteration no transmission is ever successful at all.

Theorem 8. *For every sequence* $p^{(0)}, \cdots, p^{(T)}$ *with swap regret at most* $\epsilon \cdot T$ *and for every* $\delta > 0$ *the fraction of steps in which user* i *sends successfully is at least*

$$Q \cdot \frac{f_i\left((1+\delta)p_i^*\right)}{f_i\left((1-\delta)p_i^*\right)} - \frac{\epsilon}{f_i\left((1-\delta)p_i^*\right)} \ ,$$

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 8 gives a lower bound for the number of steps in which a user can send 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. For this we give a bound in Lemma 11 and Lemma 13 later on. Altogether Theorem 8, Lemma 11, and Lemma 13 yield a bound converging to 1 as the swap regret per step approaches 0.

In order to prove this theorem, we will switch to a more convenient notation from game theory, namely correlated 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. In contrast to mixed Nash equilibria the choices of the different users do not need to be independent. Formally, an ϵ-correlated equilibrium is defined as follows.

Definition 1. *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 $\phi_i \colon P_i \to P_i$, we have*

$$\mathbf{E}_{s \sim \pi}\left[u_i(\phi_i(p_i), p_{-i})\right] - \mathbf{E}_{s \sim \pi}\left[u_i(p_i, p_{-i})\right] \leq \epsilon \enspace .$$

This is, in an ϵ-correlated equilibrium no user can increase his expected utility by operations such as "each time π says I play A, I play B instead". These kinds of operations are exactly the ones considered in the definition of no-swap-regret sequences. Therefore each sequence $p^{(1)}, \dots, p^{(T)}$ of swap regret at most R corresponds to an R/T-correlated equilibrium.

Using this notion, we can rewrite Theorem 8 to the following proposition.

Proposition 9. *For every ϵ-correlated equilibrium π and for every $\delta > 0$ the probability that user i sends successfully is at least*

$$\mathbf{Pr}_{p \sim \pi}\left[(1-\delta)p^* \leq p \leq (1+\delta)p^*\right] \frac{f_i\left((1+\delta)p_i^*\right)}{f_i\left((1-\delta)p_i^*\right)} - \frac{\epsilon}{f_i\left((1-\delta)p_i^*\right)} \enspace .$$

Proof. Consider the following switching operation. 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 have to 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

$$\mathbf{Pr}_{p \sim \pi}\left[\mathcal{E}\right] \cdot \left(\mathbf{E}_{p \sim \pi}\left[u_i((1+\delta)p_i^*, p_{-i}) \mid \mathcal{E}\right] - \mathbf{E}_{p \sim \pi}\left[u_i(p) \mid \mathcal{E}\right]\right) \leq \epsilon \enspace . \tag{2}$$

Now let us bound the two expectations in this sum.

When using power $(1+\delta)p_i^*$, user i will always be successful if the other users use a power vector $p_{-i} \leq (1+\delta)p_{-i}^*$. So when applying the switch operation, user i gets an expected utility conditioned on $\mathcal{E}$ of

$$\mathbf{E}_{p \sim \pi}\left[u_i((1+\delta)p_i^*, p_{-i}) \mid \mathcal{E}\right] \geq f_i((1+\delta)p_i^*) \cdot \mathbf{Pr}_{p \sim \pi}\left[p_{-i} \leq (1+\delta)p_{-i}^* \mid \mathcal{E}\right] \enspace ,$$

which yields

$$\mathbf{Pr}\left[\mathcal{E}\right]\cdot\mathbf{E}_{p\sim\pi}\left[u_i((1+\delta)p_i^*, p_{-i}) \mid \mathcal{E}\right] \geq f_i((1+\delta)p_i^*)\cdot\mathbf{Pr}_{p\sim\pi}\left[(1-\delta)p^* \leq p \leq (1+\delta)p^*\right] \ . \quad (3)$$

On the other hand, we have

$$\mathbf{E}_{p\sim\pi}\left[u_i(p) \mid \mathcal{E}\right] \leq f_i((1-\delta)p_i^*) \cdot \mathbf{Pr}_{p\sim\pi}\left[\text{transmission } i \text{ is successful} \mid \mathcal{E}\right] \ ,$$

yielding

$$\mathbf{Pr}\left[\mathcal{E}\right]\mathbf{E}_{p\sim\pi}\left[u_i(p) \mid \mathcal{E}\right] \leq f_i((1-\delta)p_i^*) \cdot \mathbf{Pr}_{p\sim\pi}\left[\text{transmission } i \text{ is successful}\right] \ . \qquad (4)$$

Combining Equations 2, 3, and 4, we get

$$\begin{aligned} & f_i((1+\delta)p_i^*) \cdot \mathbf{Pr}_{p\sim\pi}\left[(1-\delta)p^* \leq p \leq (1+\delta)p^*\right] \\ & - f_i((1-\delta)p_i^*) \cdot \mathbf{Pr}_{p\sim\pi}\left[\text{transmission } i \text{ is successful}\right] \ \leq \ \epsilon \ . \end{aligned}$$

This yields the claim. □

It remains to bound the probability $\mathbf{Pr}_{p\sim\pi}\left[(1-\delta)p^* \leq p \leq (1+\delta)p^*\right]$. For this purpose, we bound the probability mass of states p with $p \not\leq (1+\delta)p^*$ in Lemma 11 and of the ones with $p \not\geq (1-\delta)p^*$ in Lemma 13. This way, we get the desired bound by

$$\begin{aligned} & \mathbf{Pr}_{p\sim\pi}\left[(1-\delta)p^* \leq p \leq (1+\delta)p^*\right] \\ = & \ 1 - \mathbf{Pr}_{p\sim\pi}\left[p \not\geq (1-\delta)p^*\right] - \mathbf{Pr}_{p\sim\pi}\left[p \not\leq (1+\delta)p^*\right] \ . \end{aligned}$$

The general proof ideas work as follows. In order to bound $\mathbf{Pr}_{p\sim\pi}\left[p \not\leq (1+\delta)p^*\right]$, we consider which probability mass can at most lie on vectors p such that for some user i, we have $p_i > \left(C \cdot p^{\max} + (1+\delta/2)\cdot\eta\right)_i$. This probability mass is bounded, because user i could instead always use power $\left(C \cdot p^{\max} + \eta\right)_i$, 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 $(1+\delta)p^*$. The bound on $\mathbf{Pr}_{p\sim\pi}\left[p \not\geq (1-\delta)p^*\right]$ works in a similar way.

To see which probability mass lies on states by a factor δ away from the fixed point p^*, we will now consider how much probability mass can at most lie on states $p \not\leq (1+\delta)p^*$. 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 10. *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 $a_t \le z_t$ for all $t \in \mathbb{N}$.

Furthermore, we have that $z_t \le c(b+2)^t$. This can be proven by induction. For $t = 0$ this is clear. Now observe

$$z_t = (b+1)\sum_{k=0}^{t-1} z_k \le c(b+1)\sum_{k=0}^{t-1}(b+2)^k \le c(b+1)\frac{(b+2)^t}{(b+2)-1} = c(b+2)^t \ .$$

This yields for the sum

$$\sum_{k=0}^{t} a_k \le \sum_{k=0}^{t} z_k \le \frac{z_{t+1}}{b+1} \le \frac{c}{b+1}(b+2)^{t+1} \ .$$

□

We now turn to the proofs of the key lemmas.

Lemma 11. *Let π be an ϵ-correlated equilibrium for some $\epsilon \ge 0$. Then for all $\delta > 0$, we can bound the probability that $p \not\le (1+\delta)p^*$ is chosen by*

$$\mathbf{Pr}_{p\sim\pi}\left[p \not\le (1+\delta)p^*\right] \le \epsilon\left(\frac{n}{\delta}\max_{i\in[n]}\frac{2}{s_i\eta_i} + 2\right)^{T+1}$$

$$\textit{where } T = \frac{\log\frac{\delta}{4} - \log\max_{i\in[n]}\left|\frac{p_i^{max}}{(1+\frac{\delta}{2})p_i^*}\right|}{\log\max_{i\in[n]}\left|1 - \frac{\eta_i}{p_i^*}\right|} \ ,$$

where 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. To prove this lemma, we will 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

$$p_i^{(t)} = \min\left\{p_i^{(t-1)}, \left(C\cdot p^{(t-1)} + \left(1+\frac{\delta}{2}\right)\cdot\eta\right)_i\right\} \ .$$

This iteration is the FM iteration shifted by a factor of $1+\frac{\delta}{2}$. It shows the same behaviour as the fixed-point iteration starting from $p^{\max}$ with the noise vector being $\left(1+\frac{\delta}{2}\right)\eta$ instead of η. The fixed point of this iteration then is $p' := \left(1+\frac{\delta}{2}\right)p^*$. Let $p'' := (1+\delta)\,p^*$. Note that $p'' > p'$.

Now let be $q_0 := \mathbf{Pr}_{p\sim\pi}\left[p \not\le p^{(0)}\right]$ and $q_{t+1} := \mathbf{Pr}_{p\sim\pi}\left[p \not\le p^{(t+1)}\right] - \mathbf{Pr}_{p\sim\pi}\left[p \not\le p^{(t)}\right]$. We will bound these probabilities by considering the marginal distributions. This is, for each user i we consider $q_{0,i} := \mathbf{Pr}_{p\sim\pi}\left[p_i > p_i^{(0)}\right]$ and $q_{t+1,i} := \mathbf{Pr}_{p\sim\pi}\left[p_i > p_i^{(t+1)}\right] - \mathbf{Pr}_{p\sim\pi}\left[p_i > p_i^{(t)}\right]$. In this notation $q_t \leq \sum_{i=1}^{n} q_{t,i}$.

Now fix some user i. We can bound $q_{0,i}$ by the following observation. This user could always use power $(Cp^{\max} + \eta)_i$ instead of the powers between $p_i^{(1)}$ and $p_i^{(0)}$. No matter which powers the other users use, the transmission would still always be successful. Since π is an ϵ-equilibrium this switching operation can only yield an expected utility gain of ϵ. Therefore it holds $g_i \cdot q_{0,i} \leq \epsilon$, where $g_i = \frac{\delta}{2} \cdot s_i \cdot \eta_i$.

For $t > 0$, we can adapt this observation. We again consider the operation that always uses power $(Cp^{(t)} + \eta)_i$ instead of the powers between $p_i^{(t+1)}$ and $p_i^{(t)}$. Unfortunately, under this operation the transmission can become unsuccessful, but only if a vector p with $p \not\leq p^{(t)}$ has been chosen. This is user i might lose all utility obtained by power vectors "cut off" before. So we get

$$g_i \cdot q_{t,i} - \sum_{k=0}^{t-1} q_k \leq \epsilon \ ,$$

or equivalently

$$q_{t,i} \leq \frac{1}{g_i}\left(\epsilon + \sum_{k=0}^{t-1} q_k\right) \ .$$

Summing up over all users then leads to

$$\begin{aligned} q_t &\leq \sum_{i=1}^{n} \frac{1}{g_i}\left(\epsilon + \sum_{k=0}^{t-1} q_k\right) = \sum_{i=1}^{n} \frac{2}{\delta \cdot s_i\eta_i}\left(\epsilon + \sum_{k=0}^{t-1} q_k\right) \\ &\leq \frac{2n}{\delta} \cdot \max_i \frac{1}{s_i\eta_i}\left(\epsilon + \sum_{k=0}^{t-1} q_k\right) = \frac{n}{\min_i g_i}\left(\epsilon + \sum_{k=0}^{t-1} q_k\right) \ . \end{aligned}$$

Using Observation 10 here yields for all $t \geq 1$

$$\sum_{k=0}^{t} q_k \leq \frac{n\frac{\epsilon}{\min_i g_i}}{\frac{n}{\min_i g_i} + 1}\left(\frac{n}{\min_i g_i} + 2\right)^{t+1} \leq \epsilon\left(\frac{n}{\min_i g_i} + 2\right)^{t+1} \ .$$

The iteration given here only deviates from the FM iteration by changing the noise vector accordingly. So we can directly deduce the number of iteration steps to reach a state $p^{(T)} \leq p''$ from Theorem 2. 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 it is now possible to bound the sum of all q_t for $t \geq T$ as needed. □

The probability that vectors below $(1-\delta)\,p^*$ are chosen can be bounded in similiar ways. For this, we define $r = \min_i r_i$, and r_i is a lower bound on the utility of user i at $(1+\delta)p^*$, i.e., $r \le \min_{i\in[n]} f_i((1+\delta)p_i^*)$.

Lemma 12. *Let π be an ϵ-correlated equilibrium for some $\epsilon \ge 0$. Then for all $\delta > 0$, we can bound the probability that $p \not\geq (1-\delta)p^*$ is chosen by*

$$\mathbf{Pr}_{p\sim\pi}\left[p \not\geq (1-\delta)p^*\right] \le \left(\frac{\epsilon}{1-r} + \mathbf{Pr}_{p\sim\pi}\left[p \not\leq (1+\delta)p^*\right]\right)\left(\frac{n}{r}\right)^{T'+1}$$

$$\textit{where } T' = \frac{\log\delta}{\log\max_{i\in[n]}\left|1-\frac{\eta_i}{p_i^*}\right|} \quad .$$

Proof. We consider the same iteration steps as they are used in the fixed-point iteration starting with $p^{(0)} := 0$, $p^{(t+1)} = C \cdot p^{(t)} + \eta$. We now recursively "cut off" the probability mass for vectors $p \not\geq p^{(t)}$.

In order to bound the probability mass $q_t := \mathbf{Pr}\left[p \not\geq p^{(t)}\right] - \mathbf{Pr}\left[p \not\geq p^{(t-1)}\right]$ in the tth step, we will recursively bound q_t and $q_{t,i} := \mathbf{Pr}\left[p_i < p_i^{(t)}\right] - \mathbf{Pr}\left[p_i < p_i^{(t-1)}\right]$ similiar to the proof of Lemma 11.

A user can gain $r_i \cdot q_{t,i}$ utility by shifting the probability mass left of $p^{(t)}$ 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.

This is, we have

$$r_i \cdot q_{t,i} - (1-r_i)\left(\rho + \sum_{k=0}^{t-1} q_k\right) \le \epsilon \ \text{, where } \rho = \mathbf{Pr}_{p\sim\pi}\left[p \not\leq (1+\delta)p^*\right],$$

and therefore

$$q_{t,i} \le \frac{\epsilon + (1-r_i)\rho + \sum_{k=0}^{t-1}(1-r_i)q_k}{r_i} \quad .$$

Summing over all users j gives us

$$\begin{aligned} q_t &\le \sum_{i=1}^{n} \frac{\epsilon + (1-r_i)\rho + \sum_{k=0}^{t-1}(1-r_i)q_k}{r_i} \\ &\le \frac{n\left(\epsilon + (1-r)\rho\right)}{r} + \sum_{k=0}^{t-1} n\left(\frac{1}{r} - 1\right) q_k \quad . \end{aligned}$$

Using again Observation 10 to solve the recursion, we can bound the sum of all q_t for $t = 0, \ldots, T$ by

$$\sum_{t=0}^{T} q_t \le \frac{1}{1-r}\left(\epsilon + \rho(1-r)\right)\left(\frac{n}{r}(1-r) + 2\right)^{T+1} \le \left(\frac{\epsilon}{1-r} + \rho\right)\left(\frac{n}{r}\right)^{T+1} \quad .$$

□

Having already found a bound for $\mathbf{Pr}_{p\sim\pi}\left[p \not\leq (1+\delta)p^*\right]$ in Lemma 11, we can directly conclude the following.

Lemma 13. *Given an ϵ-correlated equilibrium and $u_i(p^{max}) \geq r = \frac{1}{2}$ for all $i \in [n]$. Then for every $\delta > 0$ the probability that a vector $p \not\geq (1-\delta)p^*$ is chosen is at most*

$$\mathbf{Pr}_{p\sim\pi}\left[p \not\geq (1-\delta)p^*\right] \leq \epsilon \left(2 + \left(\frac{n}{\delta}\max_{i\in[n]}\frac{2}{s_i\eta_i} + 2\right)^{T+1}\right)(2n)^{T'+1}$$

$$\text{with } T' = \frac{\log\delta}{\log\max_{i\in[n]}\left|1-\frac{\eta_i}{p_i^*}\right|} \text{ and } T = \frac{\log\frac{\delta}{4} - \log\max_{i\in[n]}\left|\frac{p_i^{max}}{(1+\frac{\delta}{2})p_i^*}\right|}{\log\max_{i\in[n]}\left|1-\frac{\eta_i}{p_i^*}\right|} .$$

Combining Lemma 11 and 13, we get an upper bound on $\mathbf{Pr}_{p\sim\pi}\left[(1-\delta)p^* \leq p \leq (1+\delta)p^*\right]$. For appropriately chosen δ, this bound and Proposition 9 yield that the success probability converges to 1 as ϵ approaches 0. This also yields 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 have to converge to p^*.

6 Discussion and Open Problems

In this paper, we studied two distributed power control algorithms. We obtained the first quantitative bounds on how long it takes in the FM iteration until the SINR is close to its target value. Furthermore a novel approach based on regret learning was presented. It overcomes some major drawbacks of the FM iteration. It is robust against users that deviate from the protocol and it still converges in a partial-information model, where the achieved SINR is not known. For no-swap-regret algorithms the convergence of the regret-learning approach is guaranteed.

Considering general no-swap-regret sequences is only a weak assumption and therefore the obtained bounds are not as good as the ones of the FM iteration. This yields a perspective for possible future work. An algorithm particular for power control could be designed based on the regret-learning approach presented in this paper.

Another aspect to be considered in future work could be discretization of the power levels. The standard assumption is that users can choose arbitrary real numbers as powers. In realistic devices this assumption might not be applicable. To the best of our knowledge, the additional challenges arising in this case have not been considered so far.

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] Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta algorithm and applications. Manuscript, 2005.

[3] 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)*, 2011.

[4] Avrim Blum and Yishay Mansour. From external to internal regret. *J. Machine Learning Res.*, 8:1307–1324, 2007.

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

[6] Alexander Fanghänel, Sascha Geulen, Martin Hoefer, and Berthold Vöcking. Online capacity maximization in wireless networks. In *Proc. 22nd Symp. Parallelism in Algorithms and Architectures (SPAA)*, pages 92–99, 2010.

[7] 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.

[8] Gerard Foschini and Zotan Miljanic. A simple distributed autonomous power control algorithm and its convergence. *IEEE Trans. Vehicular Technology*, 42(4):641–646, 1993.

[9] 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.

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

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

[12] Ching-Yao Huang and Roy Yates. Rate of convergence for minimum power assignment algorithms in cellular radio systems. *Wireless Networks*, 4(4):223–231, 1998.

[13] 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.

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

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

[16] V.P. Singh and Krishan Kumar. Literature survey on power control algorithms for mobile ad-hoc network. *Wireless Personal Communications*, pages 1–7, 2010.

[17] Roy Yates. A framework for uplink power control in cellular radio systems. *IEEE J. Sel. Area Comm.*, 13(7):1341–1347, 1995.

[18] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proc. 20th Intl. Conf. Machine Learning (ICML'03)*, pages 928–936, 2003.