Convergence Time of Power-Control Dynamics

Johannes Dams, Martin Hoefer, and Thomas Kesselheim

***Abstract*—We study convergence of distributed protocols for power control in a non-cooperative wireless transmission scenario. There are n wireless communication requests or *links* that experience interference and noise. To be successful a link must satisfy an SINR constraint. Each link is a rational selfish agent that strives to be successful with the least power that is required. A classic approach to this problem is the fixed-point iteration due to Foschini and Miljanic [1], for which we prove the first bounds on worst-case convergence times – after roughly $O(n \log n)$ rounds all SINR constraints are nearly satisfied. When agents try to satisfy each constraint exactly, however, links might not be successful at all. 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.**

***Index Terms*—**

I. 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 interfering signals of simultaneous transmissions and ambient noise. This condition is frequently expressed by the *signal-to-interference-plus-noise ratio (SINR)*. An important aspect that can be considered in these models is *power control*, i.e., the ability of modern wireless devices to allow their transmission powers to be set by software. Power control has several main advantages. On the one hand, battery life can be increased by using 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. So power control is also in the interest of wireless devices as it leads to lower energy consumption and increased battery life.

In this paper, we study spectrum access with power control in a non-cooperative network of wireless devices. We consider a network consisting of n *links*, i.e., sender/receiver pairs. Each sender is a selfish rational agent and 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 protocols for power control executed by all senders in parallel that allow to find transmission powers in order to make a link successful as quickly as 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 power-control protocols with approaches that solve the additional scheduling problem. However, for algorithms proposed for these problems either do not provide provable performance guarantees [2] or require a strong central authority managing the access of all devices to the spectrum [3] – thereby neglecting the distributed and non-cooperative nature of many wireless transmission scenarios.

A simple and beautiful protocol for the power-control problem is the fixed-point iteration method leading to a pure Nash equilibrium by Foschini and Miljanic [1]. 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 is natural that this is the most desirable solution for the sender and thus his utility function has a unique maximum at this power level. Hence, in game-theoretic terms, this protocol implements concurrent best-response dynamics. It can be shown that this protocol converges to a feasible assignment, even if best responses by the senders are not chosen simultaneously [4]. The obtained power assignment is a fixed point in the sense that it is a Nash equilibrium and component-wise smaller than all other feasible assignments. In this way, the fixed point is the most desirable assignment of the system. It is known that the Foschini-Miljanic (FM) iteration converges at a geometric rate [5] in a numerical sense. However, to the best of our knowledge, no results in the sense of quantitative worst-case convergence times have been shown, neither for this nor for other distributed protocols.

In this paper, we investigate two classes of distributed protocols for non-cooperative 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 FM fixed-point iteration in Section II 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 to a state where all SINR constraints are nearly satisfied. 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

Manuscript received 15 December 2011; revised 31 May 2012. This work has been supported by DFG through UMIC Research Centre, RWTH Aachen University, and grant Ho 3831/3-1. An extended abstract of this work appeared in Proc. of 38th International Colloquium on Automata, Languages and Programming (ICALP 2011).

The authors are with the Department of Computer Science, RWTH Aachen University, Germany (e-mail: {dams,mhoefer,kesselheim}@cs.rwth-aachen.de).

Digital Object Identifier 10.1109/JSAC.2012.1212xx.

fixed point if we start with all powers set to 0, as all links will raise their powers to fulfill the SINR constraint. So if one starts from all powers set to 0, and insists on all links satisfying the SINR constraints exactly via $\delta = 0$, we get an infinite convergence time during which all links remain unsuccessful. To overcome this problem, in Section III we introduce a novel technique to compute power assignments employing distributed regret-learning algorithms. This has the desirable property of being compatible with selfish behavior as no-regret sequences converge to correlated equilibria of the respective game. Furthermore, in our case for algorithms that guarantee no swap regret [6], we can show convergence to the fixed point, which in this scenario coincides with the unique mixed and pure Nash equilibrium and the correlated equilibrium. 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 link 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. In these cases, we can fall back on the respective results for capacity maximization [7], [8].

A. 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. Here, ν denotes the ambient noise. 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 $i \in [n] := \{1, \ldots, n\}$ and $C_{i,j} = \beta g_{i,j}/g_{i,i}$ for $i \neq j$. The *normalized noise vector* η is defined by $\eta_i = \beta\nu/g_{i,i}$. A feasible assignment is 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, this fixed point p^* can simply be computed by solving the linear equation system $p^* = C \cdot p^* + \eta$. However, a wireless network consists of independent non-cooperative 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.

We consider this scenario as a normal-form game as follows. Each sender i is a selfish agent and picks a transmission power as strategy. We first assume that the achieved SINR becomes known after each transmission attempt, in which case sender i has a utility function u_i depending on his chosen power p_i and the achieved SINR. Furthermore, we assume u_i has a unique maximum at the power level that sets his SINR to exactly β, but can otherwise be completely arbitrary. The FM iteration is $p^{(t+1)} = C \cdot p^{(t)} + \eta$ where t is the time step. In this context it implements a concurrent best-response dynamic. Note that 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. Obviously, p^* is a Nash equilibrium in which no sender wants to unilaterally change his transmission power. One can show that the existence of $p^* \geq 0$ with $p^* \geq C \cdot p^* + \eta$ 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 chooses his power out of an interval $[0, p_i^{\max}]$. Here $p_i^{\max}$ is the maximal power level user i can use. The utility functions u_i are of a natural form defined more precisely in Section III below. Let Φ be a set of departure measurable functions such that each $\phi \in \Phi$ is a deviating 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 incurs is

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

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 producing a no-Φ-regret sequence is a no-Φ-regret algorithm.

We will see that for our utility functions, there are distributed no-Φ-regret algorithms as will be explained in Section IV. To evaluate the utility, it suffices for each sender to only know after each transmission attempt if it has been successful.

B. Related Work

Foschini and Miljanic [1] were the first to solve the power-control problem using the iterative distributed protocol outlined above. They showed that their iteration converges from each starting point to a fixed point if it exists. Extending this, Yates [4] 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 [5] proved that all these algorithms converge geometrically, i.e., 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.

Apart from our work the only bound based on the network parameters is derived by Lotker et al. [9] for a special case. 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 [10].

In order to reflect the individual rationality of each agent, the power control problem has also been considered in game-theoretic scenarios. For example, the results by Yates can be extended to games involving selfish agents with certain utility functions [11], [12].

The approach of Foschini and Miljanic can be transferred to a game-theoretic utility function by considering the distance to a target SINR. This can be done by penalizing the quadratic error to the target [13] or using a non-continuous function [14], which is zero when the target SINR is not reached and decreasing otherwise. In either case, the unique Nash equilbrium corresponds to the fixed point of the FM iteration.

A different approach is to consider the link capacity (depending on the SINR) as the utility function. To model the trade off between capacity and energy consumption, the used power is subtracted from this utility [15], [16], [17] or considered in some other way [18], [19]. This, however, neglects the necessity of a minimal SINR for certain applications. More sophisticated utility functions can be derived from considering the bit-error rate [20], [21], [22]. Here, the utility is proportional to the "information received per Joule", i.e., the ratio of the probability of a successful transmission and the used power. This results in a function that abruptly inclines left of a maximum and slowly decreases right of it. Apart from having no discontinuity left of the maximum, these functions resemble the ones considered in this paper.

II. Convergence Time of the FM Iteration

In this section, we analyze the convergence time of the FM 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. 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 modulus of the a_i values can be expressed in terms of the b_i values by

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

This yields the following bound for their sum

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

We now use the fact that $\chi_M(M) = 0$. This is, $M^n = -\sum_{i=0}^{n-1} a_i M^i$. Since all $M^k \eta'$ are non-negative, the following inequality holds

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

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

One can see that this bound is almost tight as there are instances where $\Omega(n)$ rounds are needed. For example, 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 first component 1 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*

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

Proof: We consider the weighted maximum norm, which has been used by Huang and Yates [5] before. We use the entries of p^* as weights by defining $\|x\| = \max_{i \in [n]} \frac{|x_i|}{p_i^*}$. The induced matrix norm of a matrix M is now given by $\|M\| = \max_{i \in [n]} \frac{1}{p_i^*} \sum_{j \in [n]} |M_{i,j}| \cdot p_j^*$.

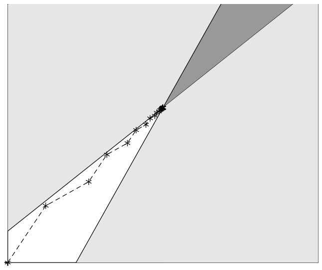

Fig. 1. FM iteration in an example.

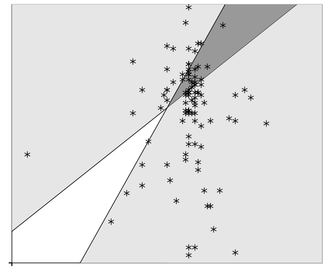

Fig. 2. Last 100 of 1000 iterations of no-swap regret learning.

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

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

If $t \geq T$, using Equation 1 we get $\|p^{(t)} - p^*\| = \|C^t(p^{(0)} - p^*)\| \leq \|C\|^t \cdot \|p^{(0)} - p^*\| = \|C\|^t \cdot \max_{i\in[n]} \left| p_i^{(0)}/p_i^* - 1 \right| \leq \delta$. So, for all $i \in [n]$, we have $|p_i^{(t)} - p_i^*| \leq \delta p_i^*$. ■

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 [4]. 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.

III. 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. An examplary run with two links is shown in Figure 1. The light grey regions depict the possible power choices where at least one link can transmit successfully. The darker region refers to states where both links are successful. All points choosen by the FM iteration lie in the white region.

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

In order to overcome these drawbacks, we design a different approach based on regret learning. As these algorithms are randomized, each player can transmit successfully already in the first time steps. This is in contrast to the FM iteration never becoming successful at all. An example run of such a no-swap-regret algorithm is shown in Figure 2. It shows that it converges towards the region where both links are successful.

Besides the advantage of having successful transmissions, we also use the no-swap-regret algorithm to overcome the other drawbacks of the FM iteration. There is no need to know the achieved SINR and it is robust against scenarios where no fixed point exists.

Again, we assume each user to be aiming at having a successful transmission but using the least power possible. The decision which power $p_i \in [0, p_i^{\max}]$ to use is based on the following kind of utility functions. We assume that each user gets zero utility if the SINR is below the threshold and a positive one otherwise. 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 $f_i\colon [0, p_i^{\max}] \to [0,1]$ is a continuous and strictly decreasing function for each $i \in [n]$. With p_{-i} we denote the powers chosen by all users but user i.

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. This way a best response dynamic corresponds to the FM iteration. 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 $p^{(1)}, \ldots, p^{(T)}$ with the swap regret for each user being 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 $\check{p}_i = \min_{t\in[T]} p_i^{(t)}$. We can argue analogously as before to get $\check{p}_i \geq p^*$. 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 suitable functions f_i, 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. [23]) 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 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 as they cannot get through at all. The inner links have 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. ∎

IV. Computing No-Swap-Regret Sequences

The reason to study no-regret sequences is that they can be computed in a distributed way. Mostly, the case of a finite action space has been studied. For example, for the case of N actions, the algorithm devised by Blum and Mansour [6] is able to guarantee that after T round the expected regret of a user is at most $O(\sqrt{TN \log N})$. This algorithm is randomized and uses multiplicative weights updates. That is, each user has a probability distribution over all possible actions. After each step, he updates this distribution based on the previously observed utilities.

Unforunately, this and similar algorithms [24] require a finite number of actions, while in our case the action space contains all real numbers within $[0, p_i^{\max}]$. Standard approaches for infinite action spaces are not applicable either as they require convex action spaces and continuous utility functions [25], [26]. In order to capture the SINR threshold appropriately, however, the utility functions have to be modeled as non-continuous.

In this section, we show that no-regret sequences can be computed in a distributed way nevertheless. This is achieved by applying no-swap-regret algorithms for finite action spaces on a suitable finite subset of the powers. This finite subset is constructed by dividing the set of powers into intervals of equal length and using the right borders as the input action set for the algorithm. This discretization, however, is not chosen in a fixed way but iteratively refined to guarantee that the no-swap-regret property holds.

Theorem 5. *Let $\mathcal{A}$ be any no-swap-regret algorithm for arbitrary finite action spaces, whose swap regret after T rounds in case of N actions is at most $O(T^a \cdot N^b)$, where a and b are suitable constants with $0 \leq a < 1$, $b \geq 0$.*

Then $\mathcal{A}$ can be used to construct an algorithm for power control on infinite action spaces achieving swap regret at most $O(T^{\frac{a+b}{1+b}})$.

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 an estimate $T = 1$, the algorithm is executed for T steps with the respective estimate T, which is doubled afterwards. This way only a constant factor is lost in comparison to the case where the exact T is known. ∎

Theorem 5 provides a framework to use different suitable existing no-swap-regret algorithms. Depending on the specific algorithm used, it yields different regret bounds. 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 [6] 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 [6] 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.*

V. 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, in the FM iteration there are starting vectors such that no transmission is ever successful at all.

Theorem 8. *For every sequence $p^{(1)}, \cdots, p^{(T)}$ 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 $(1-\delta)p^* \le p \le (1+\delta)p^*$ 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. We bound this fraction in Lemma 10 and Lemma 11 later on. Altogether Theorem 8, Lemma 10, and Lemma 11 yield a bound converging to 1 as the swap regret per step approaches 0.

In order to prove this theorem, we will use the fact that no-swap-regret sequence correspond to approximate equilibria. Similar to a mixed Nash equilibrium, an ϵ-correlated equilibrium is a probability distribution over strategy vectors (in our case power vectors) such that no user can unilaterally increase his expected utility by more than ϵ. 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 $P_1 \times \cdots \times P_n$, 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*

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

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". This kind of operations are exactly the ones considered in the definition of no-swap-regret sequences. Therefore each sequence $p^{(1)}, \ldots, p^{(T)}$ 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*

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

where $Q = \mathbf{Pr}_{p\sim\pi}\left[(1-\delta)p^* \le p \le (1+\delta)p^*\right]$.

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 bound the change of the expected utility due to this switching operation.

Let $\mathcal{E}$ be the event that a vector p is chosen with $p_i \in [(1-\delta)p_i^*, (1+\delta)p_i^*]$ then the expected utility gain is

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

Now let us bound the two expectations in this sum.

When using the power $(1+\delta)p_i^*$, user i will always be successful if the other users use a power vector $p_{-i} \le (1+\delta)p_{-i}^*$. So when applying this switch operation, user i gets an expected utility conditioned on the event $\mathcal{E}$ of $\mathbf{E}_{p\sim\pi}\left[u_i((1+\delta)p_i^*, p_{-i}) \mid \mathcal{E}\right] \ge f_i((1+\delta)p_i^*) \cdot \mathbf{Pr}_{p\sim\pi}\left[p_{-i} \le (1+\delta)p_{-i}^* \mid \mathcal{E}\right]$, which yields

$$\begin{aligned}&\mathbf{Pr}\left[\mathcal{E}\right] \cdot \mathbf{E}_{p\sim\pi}\left[u_i((1+\delta)p_i^*, p_{-i}) \mid \mathcal{E}\right] \\ &\ge f_i((1+\delta)p_i^*) \cdot \mathbf{Pr}_{p\sim\pi}\left[(1-\delta)p^* \le p \le (1+\delta)p^*\right] \ . \qquad (3)\end{aligned}$$

On the other hand, we have $\mathbf{E}_{p\sim\pi}\left[u_i(p) \mid \mathcal{E}\right] \le f_i((1-\delta)p_i^*) \cdot \mathbf{Pr}_{p\sim\pi}\left[\mathcal{S} \mid \mathcal{E}\right]$, where $\mathcal{S}$ is the event that the transmission is successful. This yields

$$\mathbf{Pr}\left[\mathcal{E}\right] \mathbf{E}_{p\sim\pi}\left[u_i(p) \mid \mathcal{E}\right] \le f_i((1-\delta)p_i^*) \cdot \mathbf{Pr}_{p\sim\pi}\left[\mathcal{S}\right] \ . \qquad (4)$$

Combining Equations 2, 3, and 4, we get $f_i((1+\delta)p_i^*) \cdot \mathbf{Pr}_{p\sim\pi}\left[(1-\delta)p^* \le p \le (1+\delta)p^*\right] - f_i((1-\delta)p_i^*) \cdot \mathbf{Pr}_{p\sim\pi}\left[\mathcal{S}\right] \le \epsilon$. This yields the claim. ■

It remains to bound $\mathbf{Pr}_{p\sim\pi}\left[(1-\delta)p^* \le p \le (1+\delta)p^*\right]$. For this purpose, we bound the probability mass of states p with $p \not\le (1+\delta)p^*$ in Lemma 10 and of the ones with $p \not\ge (1-\delta)p^*$ in Lemma 11.

The general proof ideas work as follows. In order to bound $\mathbf{Pr}_{p\sim\pi}\left[p \not\le (1+\delta)p^*\right]$, 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 $\mathbf{Pr}_{p\sim\pi}\left[p \not\ge (1-\delta)p^*\right]$ works in a similar way.

First, we consider how much probability mass can at most lie on states $p \not\le (1+\delta)p^*$. Afterwards, we will do the same for $p \not\ge (1-\delta)p^*$. In the following only the key lemmas are presented. The proofs are omitted and can be found in the full version.

Lemma 10. *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*

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

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

and s_i denotes the minimal absolute value of the difference quotient of f_i at any point p_i and $p_i + \frac{\delta}{2}\eta_i$.

The probability that vectors below $(1-\delta)\,p^*$ are chosen can be bounded in similiar ways yielding Lemma 11.

Lemma 11. *Given an ϵ-correlated equilibrium and assuming $u_i(p^{max}) \ge r = \frac{1}{2}$ for all $i \in [n]$. Then for every $\delta > 0$ the probability $\mathbf{Pr}_{p\sim\pi}\left[p \not\ge (1-\delta)p^*\right]$ is at most*

$$\epsilon \left(2 + \left(\frac{n}{\delta} \max_{i\in[n]} \frac{2}{s_i \eta_i} + 2\right)^{T+1}\right)(2n)^{T'+1}$$

with $T' = \frac{\log \delta}{\log \max_{i \in [n]} \left| 1 - \frac{\eta_i}{p_i^*} \right|}$ *and* T *defined as in Lemma 10.*

Combining Lemma 10 and 11, 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 converge to p^*.

VI. Discussion and Open Problems

In this paper, we studied two distributed power control protocols compatible with individual rational behavior. We obtained the first quantitative bounds on how long it takes for best repsonse dynamics respectively 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 is still applicable in a partial-information model, where the achieved SINR is not known. If all users follow no-swap-regret algorithms, the convergence 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. By adapting the regret-learning approach presented in this paper and tailoring a protocol specifically to power control one could achieve faster and better algorithms.

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] G. Foschini and Z. Miljanic, "A simple distributed autonomous power control algorithm and its convergence," *IEEE Trans. Veh. Technol.*, vol. 42, no. 4, pp. 641–646, 1993.

[2] T. A. ElBatt and A. Ephremides, "Joint scheduling and power control for wireless ad hoc networks," *IEEE Trans. Wireless Commun.*, vol. 3, no. 1, pp. 74–85, 2004.

[3] T. Kesselheim, "A constant-factor approximation for wireless capacity maximization with power control in the SINR model," in *Proc. 22nd SODA*, 2011, pp. 1549–1559.

[4] R. Yates, "A framework for uplink power control in cellular radio systems," *IEEE J. Sel. Area Comm.*, vol. 13, no. 7, pp. 1341–1347, 1995.

[5] C.-Y. Huang and R. Yates, "Rate of convergence for minimum power assignment algorithms in cellular radio systems," *Wireless Networks*, vol. 4, no. 4, pp. 223–231, 1998.

[6] A. Blum and Y. Mansour, "From external to internal regret," *J. Machine Learning Res.*, vol. 8, pp. 1307–1324, 2007.

[7] M. Dinitz, "Distributed algorithms for approximating wireless network capacity," in *Proc. 29th INFOCOM*, 2010, pp. 1397–1405.

[8] E. I. Asgeirsson and P. Mitra, "On a game theoretic approach to capacity maximization in wireless networks," in *Proc. 30th INFOCOM*, 2011, pp. 3029–3037.

[9] Z. Lotker, M. Parter, D. Peleg, and Y. A. Pignolet, "Distributed power control in the sinr model," in *Proc. 30th INFOCOM*, 2011, pp. 2525–2533.

[10] V. Singh and K. Kumar, "Literature survey on power control algorithms for mobile ad-hoc network," *Wireless Personal Communications*, pp. 1–7, 2010.

[11] E. Altman and Z. Altman, "S-modular games and power control in wireless networks," *IEEE Trans. Automat. Contr.*, vol. 48, pp. 839–842, 2003.

[12] T. Heikkinen, "A potential game approach to distributed power control and scheduling," *Comput. Netw.*, vol. 50, pp. 2295–2311, 2006.

[13] S. Koskie and Z. Gajic, "A nash game algorithm for sir-based power control in 3g wireless cdma networks," *IEEE/ACM Trans. Netw.*, vol. 13, no. 5, pp. 1017 – 1026, 2005.

[14] C. Long, Q. Zhang, B. Li, H. Yang, and X. Guan, "Non-cooperative power control for wireless ad hoc networks with repeated games," *IEEE J. Sel. Area Comm.*, vol. 25, no. 6, pp. 1101–1112, 2007.

[15] T. Alpcan and T. Basar, "A hybrid systems model for power control in multicell wireless data networks," *Perform. Eval.*, vol. 57, no. 4, pp. 477–495, 2004.

[16] T. Alpcan, T. Basar, R. Srikant, and E. Altman, "Cdma uplink power control as a noncooperative game," *Wireless Networks*, vol. 8, no. 6, pp. 659–670, 2002.

[17] C. U. Saraydar, N. B. Mandayam, and D. J. Goodman, "Pricing and power control in a multicell wireless data network," *IEEE J. Sel. Area Comm.*, vol. 19, no. 10, pp. 1883–1892, 2001.

[18] J. Huang, R. A. Berry, and M. L. Honig, "Distributed interference compensation for wireless networks," *IEEE J. Sel. Area Comm.*, vol. 24, no. 5, pp. 1074–1084, 2006.

[19] H. Ji and C.-Y. Huang, "Non-cooperative uplink power control in cellular radio systems," *Wireless Networks*, vol. 4, pp. 233–240, 1998.

[20] D. Famolari, N. Mandayam, D. Goodman, and V. Shah, "A new framework for power control in wireless data networks: Games, utility, and pricing," in *Wireless Multimedia Network Technologies*, 2002, vol. 524, pp. 289–309.

[21] C. Saraydar, N. Mandayam, and D. Goodman, "Efficient power control via pricing in wireless data networks," *IEEE Trans. Commun.*, vol. 50, no. 2, pp. 291 –303, 2002.

[22] F. Meshkati, M. Chiang, H. Poor, and S. Schwartz, "A game-theoretic approach to energy-efficient power control in multicarrier cdma systems," *IEEE J. Sel. Area Comm.*, vol. 24, no. 6, pp. 1115 –1129, june 2006.

[23] A. Fanghänel, T. Kesselheim, H. Räcke, and B. Vöcking, "Oblivious interference scheduling," in *Proc. 28th PODC*, 2009, pp. 220–229.

[24] S. Arora, E. Hazan, and S. Kale, "The multiplicative weights update method: a meta algorithm and applications," 2005, manuscript.

[25] M. Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," in *Proc. 20th ICML*, 2003, pp. 928–936.

[26] G. Stoltz and G. Lugosi, "Learning correlated equilibria in games with compact sets of strategies," *Games and Economic Behavior*, vol. 59, no. 1, pp. 187–208, 2007.