

Distributed Algorithms for Approximating Wireless Network Capacity

Michael Dinitz
Carnegie Mellon University
Pittsburgh, PA, USA
Email: mdinitz@cs.cmu.edu

Abstract—In this paper we consider the problem of maximizing wireless network capacity (a.k.a. one-shot scheduling) in both the protocol and physical models. We give the first distributed algorithms with provable guarantees in the physical model, and show how they can be generalized to more complicated metrics and settings in which the physical assumptions are slightly violated. We also give the first algorithms in the protocol model that do not assume transmitters can coordinate with their neighbors in the interference graph, so every transmitter chooses whether to broadcast based purely on local events. Our techniques draw heavily from algorithmic game theory and machine learning theory, even though our goal is a distributed algorithm. Indeed, our main results allow every transmitter to run *any* algorithm it wants, so long as its algorithm has a learning-theoretic property known as *no-regret* in a game-theoretic setting.

I. INTRODUCTION

In this paper we are concerned with maximizing the transmission capacity of wireless networks. In the basic model we are given a collection of transmitter/receiver pairs in the Euclidean plane, and the goal is to maximize the number of successful instantaneous transmissions. Maximizing transmission capacity has been studied in many contexts, and while many variants have been considered, there are two axes along which much of the work can be partitioned. The first axis is random vs. arbitrary networks. If we consider random networks, then the goal is typically to give bounds on the expected capacity, and study how this changes with the density of the network (or with some other interesting parameter). Another option, which is what we consider in this work, is to study arbitrary or worst-case topologies. In this setting it makes no sense to study the “average” capacity, since that could depend heavily on the actual structure of the network. Instead, the goal is to study the problem of maximizing capacity as an optimization problem, and give hardness results, centralized algorithms, and distributed protocols given an arbitrary network as input.

The second axis is the protocol model vs. the physical model, and is concerned with how we model interference and define a successful transmission. In the *protocol model* there is some interference graph on the desired transmissions, and a transmission is successful if and only if none of the neighbors of the transmission in this graph also chose to transmit. It is obvious from this definition that maximizing network capacity is the same problem as finding a maximum independent set in the interference graph, which is a famous and well-studied problem in its own right. In the context of this problem, further assumptions are usually made about the structure of the

interference graph, since physical constraints make it unlikely that this graph is totally arbitrary. One typical assumption is that it is a *unit disk graph* (UDG), which basically means that transmitters interfere if they are too close to each other. There has also been a considerable line of work on weakening this assumption or on variants of it, including the Tx model of [20] and the growth-bounded model of [18] and [11].

In the *physical model*, on the other hand, we do not assume the existence of an interference graph. Instead we let every transmitter choose a power to broadcast at, give a rule for how that power fades with distance, and say that a transmission has been successful if and only if the received signal divided by the sum of the interference and background noise is at least some threshold. This model is significantly more complicated than the protocol model, for a variety of reasons. In the protocol model the success of a transmission depends only on the OR of its neighbors; if any of its neighbors transmit then it fails, irrespective of whether one or 10 transmitted, and any number of transmitters outside of its neighborhood can transmit without affecting its success. But in the physical model interference accumulates and normally spreads out to infinity, so not only is the decision function more complicated than an OR of neighbors it actually depends on every transmitter in the entire network. While not all of the assumptions in the physical model are absolutely true, it is commonly thought to be a more accurate model of reality than the protocol model. In this paper we will consider both the protocol and the physical models.

Furthermore, there is a difference between *centralized* and *distributed* algorithms. While studying the fundamental computational problem is interesting, in many (perhaps most) real world situations there is no central authority to run the algorithm and tell all of the transmitters what to do. Ideally each transmitter would make its own decisions about whether to broadcast (and in the physical model, how much power to use). In the protocol model, since we have an interference graph we can simply abstract out to the graph and run a normal distributed protocol on this graph, and indeed this problem is usually classified under “distributed maximum independent set”. In the physical model, however, there is no underlying communication or interference graph so coordination is more complicated. And even in the protocol model, using standard models for distributed algorithms are problematic: do transmitters really know their neighbors? Can they send different messages to different transmitters? Can a transmitter receive multiple messages at the same time?

In this paper we try to work in a more general model that assumes less about inter-node communication: we consider arbitrary networks in both the physical and protocol models, where every transmitter knows only what happens to its transmissions. We do not even assume that neighbors in the protocol model can communicate (although our assumption about reception knowledge is equivalent to every node knowing whether or not at least one of its neighbors attempted to transmit), and in fact we design algorithms assuming that they cannot. We show that even in this extremely general model, there are simple algorithms that can guarantee that the average number of successful transmissions is a good approximation to the optimal solution. While distributed approximation algorithms for maximum independent set are well studied (e.g. [18]), this is, to the best of our knowledge, the first result that does not include communication among nodes, just information about the result of a transmission. This is also the first decentralized algorithm with provable approximation guarantees in the physical model, which is perhaps a more interesting result as until recently we did not even know of a good centralized algorithm in this model [7], [6], [1].

Moreover, the inspiration and techniques we use come not from the distributed computing literature, but instead from the algorithmic game theory and learning theory literature. In particular, we study a notion that generalizes the famous price of anarchy: the *price of total anarchy*, originally defined by [3] as a way of weakening the rationality assumption behind the price of anarchy. But by definition there are algorithms that do almost as well as the price of total anarchy, unlike the price of anarchy. So by proving that the price of total anarchy is small we have actually proved that if every transmitter runs a *no-regret* algorithm then the average performance will be good. This is a powerful tool when designing distributed algorithms since it allows the algorithm designer to prove approximation guarantees for distributed algorithms simply by proving no-regret for a centralized algorithm. We hope that this technique for designing distributed algorithms will prove useful for other problems, and believe that maximizing network capacity is simply one of a number of problems in which the price of total anarchy can be bounded.

A. Wireless Models

As discussed, in the protocol model every transmitter can either transmit or not transmit, and a particular transmitter is successful if and only if it chooses to transmit and none of the neighbors of its connection in the interference graph choose to transmit. We will briefly consider general graphs, but spend most of our time on *locally growth-bounded* graphs, which are a generalization of the *growth-bounded* graphs of [11]:

Definition 1: A graph $G = (V, E)$ is *locally growth-bounded* if there is some constant k such that for every node $v \in V$ the size of a maximum independent set in $N(v) \cup \{v\}$ is at most k , where $N(v)$ denotes the set of neighbors of v in G .

We note that, as pointed out by [18], growth-bounded graphs generalize unit disk graphs, quasi-unit disk graphs, unit ball graphs, and other popular generalizations of UDGs, and

therefore locally growth-bounded graphs also generalize these models.

For the physical model we use a combination of the model used by Andrews and Dinitz in [1] and the generalized physical model of Moscibroda, Wattenhofer, and Zollinger [14]. In particular, we consider a set of n connections in the plane, where each connection has a transmitter t_i and a receiver r_i . For two points u and v in the plane, let $d(u, v)$ be the normal Euclidean distance between them. Suppose that u is broadcasting with power p . Following the model from [16] and [1], the theoretical signal strength at v is $p \cdot \min\{(d_0/d(u, v))^\alpha, 1\}$, where α and d_0 are some parameters that we assume are constants. We will also make the standard assumption that $\alpha > 2$ (this assumption was used in [7], [6], [1], among others). Note that this model allows nodes to be arbitrarily close together, and just caps the received power by what happens at distance d_0 . This model generalizes the model from much of the previous work in which $d_0 = 1$ and all distances are at least 1 [13], [14], [7].

However, it has long been observed that the theoretical received power is not perfectly accurate in practice, since it does not take into account obstructions, terrain, etc. To take this into account, Moscibroda, Wattenhofer, and Zollinger introduced the *generalized* physical model. In this generalization, there is some constant parameter $\theta \geq 1$ such that the actual received signal strength is within $1/\theta$ and θ times the predicted theoretical received signal strength. More formally, if a node u transmits at power p , then the received signal strength $P_r(u, v)$ at another node v is in the interval

$$\left[\frac{1}{\theta} p \min \left\{ \left(\frac{d_0}{d(u, v)} \right)^\alpha, 1 \right\}, \theta p \min \left\{ \left(\frac{d_0}{d(u, v)} \right)^\alpha, 1 \right\} \right]$$

We do assume, however, that the distortion from the theoretical signal strength does not depend on the choice of power. So if when u broadcasts at power p the received signal strength is $\gamma p \min\{1, (d_0/d(u, v))^\alpha\}$ then the received signal strength if u broadcasts at power p' is $\gamma p' \min\{1, (d_0/d(u, v))^\alpha\}$.

A transmission from t_u to r_u is successful if the ratio of the received signal strength to the interference is at least some threshold τ ; that is, if

$$\frac{P_r(t_u, r_u)}{\sum_{v \neq u} P_r(t_v, r_u)} \geq \tau$$

We will sometimes call this an *SINR constraint*.

For ease of presentation, we will assume throughout this paper that the maximum power of a transmitter is 1. All of our results hold for an arbitrary maximum power. We will make the simplifying assumption that there is no background noise, i.e. the only causes of interference at a receiver are the signals of other transmitters. Our results still hold with nonzero background noise as long as we are guaranteed that every transmitter-receiver pair is at a distance bounded away from their absolute physical limit assuming signal weakening by θ ; that is, there is some constant $\delta > 0$ such that $d(t_i, r_i) \leq (1 - \delta) \left(\frac{1}{\tau \theta W} \right)^{1/\alpha}$ where W is the background noise. We will let $d_{\max} = \max_i d(t_i, r_i)$ be the maximum distance among any transmitter-receiver pair.

B. Game Theory Basics

Before we can state our results we first need a few basic definitions about games. In this paper the only games we will care about will be games with n players in which every player has exactly two possible actions. Let $\mathcal{A} = \{0, 1\}^n$ be the space of possible strategy profiles for the game, i.e. given a point $A \in \mathcal{A}$, the i th coordinate a_i represents the action used by player i in profile A . Each player i will have a function $\alpha_i : \mathcal{A} \rightarrow \mathbb{R}$ that assigns a utility to each strategy profile. We will sometimes want to consider modifications of strategy profiles: given $A \in \mathcal{A}$, let $A \oplus a'_i$ be the strategy set obtained if player i changed its action from a_i to a'_i . We will use superscripts to denote time, so A^t will be the action set at time t and a_i^t will be the action taken by player i at time t .

The following definition will play a central role in this paper:

Definition 2: The *regret* of player i at time T given strategy profiles A^1, A^2, \dots, A^T is

$$\max_{a_i \in \{0,1\}} \frac{1}{T} \sum_{t=1}^T \alpha_i(A^t \oplus a_i) - \frac{1}{T} \sum_{t=1}^T \alpha_i(A^t)$$

Intuitively, having low regret means that you do almost as well as on average as the best single action would have done. This notion of regret has been studied extensively, especially in two different models: the *experts* model and the *bandit* model. The difference between the two models lies in the knowledge gained by a player after each round: in the bandit model a player only finds out the utility that it gained, while in the experts model players also find out the utility they would have gained if they had played the other action. Since the results that we care about are similar in both models, we will choose the more general one and be in the bandit model. In the wireless setting, this means that if a transmitter chooses to transmit then it will find out whether or not it succeeded, but if it chooses not to transmit then it gains no information.

The *price of total anarchy* was introduced by Blum et al. [3] as a way of generalizing the *price of anarchy*. The price of anarchy of a game is the ratio of the value of the worst Nash equilibrium to the social optimum. It is, as the name suggests, supposed to quantify the “price” that is being paid by allowing each player to be a separate rational agent rather than simply being controlled by a centralized authority. Unfortunately there are various problems with this definition, one of which is that, since finding a Nash equilibrium is PPAD-complete [4], it is not clear that rational agents will actually play a Nash equilibrium. In particular, if they always do then we could find a Nash equilibrium simply by letting rational agents play. Blum et al. [3] proposed weakening this rationality assumption by assuming only that the agents use strategies with regret tending to 0 as time goes to infinity (called *no-regret algorithms*). They chose this assumption because it generalizes the Nash assumption, and is plausible since such algorithms actually do exist so smart players should do at least as well. They call the ratio of the average social welfare obtained by players using no-regret algorithms to the optimum social welfare the price of total anarchy. We will use this not as a tool for weakening rationality assumptions, but rather as a tool for designing distributed algorithms, since by definition we are

guaranteed the existence of algorithms that achieve the price of total anarchy (any no-regret algorithm).

C. Results

In all of our results we let OPT denote the optimal solution, i.e. a set of transmitters forming a maximum independent set (in the protocol model) or for which an appropriate setting of powers maximizes the number of successful connections (in the physical model). Our main results are proofs that the price of total anarchy is small in certain games that we have designed for the protocol and physical models. This immediately implies that if every transmitter uses a no-regret algorithm then the average performance over time is good. In particular, in the protocol model we design a game that gives the following result:

Theorem 3: In the protocol model in locally growth-bounded graphs, if all transmitters have regret at most ϵ then the average number of successful connections is at least $\Omega(|OPT| - \epsilon n)$.

We design a similar game for the physical model; in fact, we use a game introduced by Andrews and Dinitz [1], who showed that its price of anarchy is small. We leave as open the question of whether or not there are algorithms for this game that converge to a Nash equilibrium, but we do show that the price of total anarchy is within a constant of the price of anarchy proved in [1].

Theorem 4: In the physical model, if all transmitters have regret at most ϵ , then the average number of successful transmissions is $\Omega(|OPT|/d_{max}^{2\alpha} - \epsilon n)$, where d_{max} is the largest distance between any transmitter and its own receiver.

Thus if every transmitter has regret less than $1/n$, the average number of successful connections is within a constant of the optimal solution. The algorithm that we will use to achieve this is that of Auer, Cesa-Bianchi, Freund, and Schapire [2], who gave an algorithm with the following guarantee:

Theorem 5 ([2]): There is an algorithm that has regret at most $O\left(\sqrt{\frac{\log(T/\delta)}{T}}\right)$ with probability at least $1 - \delta$ for any $\delta > 0$, for any game with a constant number of possible actions per player.

A similar guarantee was given for the Randomized Weighted Majority Algorithm by Littlestone and Warmuth [12], which is the algorithm that we will use later in simulation.

If $\delta = \frac{1}{n^2}$ and $T \geq \Omega(n^2 \log n)$ then this regret is at most $\frac{1}{n}$ with probability at least $1 - \frac{1}{n^2}$, and thus with high probability every transmitter will have regret at most $\frac{1}{n}$. This immediately yields the following corollaries:

Corollary 6: In the protocol model in locally growth-bounded graphs, if every transmitter uses the algorithm of Theorem 5 then after $O(n^2 \log n)$ rounds the average number of successful connections is $\Omega(|OPT|)$ with high probability.

Corollary 7: In the physical model, if every transmitter uses the algorithm of Theorem 5 then after $O(n^2 \log n)$ rounds the average number of successful connections is $\Omega(|OPT|/d_{max}^{2\alpha})$ with high probability.

While the $O(n^2 \log n)$ bound on the number of rounds is rather large, if the rounds are short (i.e. transmitters update their powers often) then it is still reasonable. Furthermore, in

Section V we show by simulation that in practice the number of rounds required is small.

We believe that this is the first distributed algorithm that provably approximates the maximum capacity in the physical model. In the protocol model the current state of the art for distributed approximations to the maximum independent set is the algorithm of Schneider and Wattenhofer [18] who obtain a $(1 + \epsilon)$ -approximation after $O(\log^* n / \epsilon^{O(1)})$ rounds on bounded-growth graphs. Our approximation is not as good and we take many more rounds, but on the other hand their algorithm assumes that every transmitter knows who interferes with it and can communicate with them, whereas our algorithm requires only that in each round every transmitter knows whether or not its broadcast succeeded (if it tried to broadcast). This is a significantly weaker model of distributed computation since it explicitly disallows the complex communication among nodes that is usually a hallmark of distributed protocols, so it is interesting that we can still obtain good approximation guarantees. However, in light of [18] we feel that the main takeaway from our result in the protocol model is that our general technique (bounding the price of total anarchy) can be applied to more than just one problem in one model.

D. Related Work

The capacity of random networks in both the physical and the protocol models (for UDGs) was examined in the seminal paper of Gupta and Kumar [8]. The capacity of arbitrary networks in the protocol model has traditionally been considered in the special case when the interference graphs are UDGs; in this case, $(1 + \epsilon)$ -approximations were obtained in [5], [15]. There has also been significant work on designing distributed algorithms for finding large independent sets in UDGs and generalizations of UDGs, most recently the work of Schneider and Wattenhofer [18] who gave a $(1 + \epsilon)$ -approximation for growth-bounded graphs.

In the physical model, the problem of maximizing capacity in arbitrary networks has only recently been considered. The first to do so were Goussevskaia, Oswald, and Wattenhofer [7], who gave an $O(\log d_{\max})$ -approximation but assumed that transmitters could only broadcast at full power or not at all. Goussevskaia, Halldórson, Wattenhofer, and Welzl [6] then improved this to give a $O(1)$ -approximation. Andrews and Dinitz [1] considered the general power setting, and proved a $O(\log d_{\max})$ -approximation using techniques similar to [7]. Recently, Halldórson and Wattenhofer [9] showed how to extend the techniques of [6] to the general scheduling problem, instead of the one-shot capacity. However, all of this work is extremely centralized; this paper is, to the best of our knowledge, the first to consider the distributed problem.

There has, however, been a significant amount of work on distributed algorithms for related problems, and in fact much of it has used game-theoretic techniques. For example, Stolyar and Viswanathan [19] studied fractional frequency reuse algorithms for joint channel assignment and power control in cellular OFDM systems, and provided a game theoretic solution that always leads to a stable solution. Saraydar, Mandayam, and Goodman [17] designed a game theoretic algorithm for choosing powers on the uplink of a single cell wireless system.

However, while both of these examples (and much previous work on similar problems) give algorithms that provably converge to a stable solution, neither of them study the quality of that solution compared to the global optimum.

A different approach was taken by Andrews and Dinitz [1], who designed a game (in fact, the game we will be analyzing in this paper) for wireless network capacity in the physical model but did not prove convergence. In fact, they showed a simple example on which natural rational play will never converge. On the other hand, they showed that all Nash equilibria, both pure (in which case the players converged) and mixed, have an expected number of successful transmissions within $O(d_{\max}^{2\alpha})$ of optimal. They left designing a distributed algorithm that actually converged to a pure Nash equilibrium (or even a mixed Nash) as an open problem. While we do not answer this, we show that their game actually has the additional property that low-regret algorithms result in solutions that are almost as good as Nash equilibria from an approximation standpoint. This is not always the case, as many games have the property that all Nash equilibria are close to optimal while low-regret play can be arbitrarily bad (see [3] for some examples). But in this case the game is robust enough that, after overcoming some technical difficulties, the techniques from [1] can be used to prove that low-regret play is close to optimal.

II. BASIC GAME THEORY

Some of the basic game theory underlying our results is the same in both the protocol model and the physical model. In particular, the basic game is the same. Each transmitter is a player, with two possible strategies: broadcast at power 0 (i.e. do not broadcast) or broadcast at power 1 (full power). Note that in the physical model we are competing with the optimum solution that can use any power between 0 and 1, but we will only be using powers 0 and 1. A transmitter has utility 1 if it broadcasts successfully, i.e. meets its SINR requirement in the physical model or has no neighbors broadcasting in the protocol model. It has utility -1 if it broadcasts unsuccessfully, and utility 0 if it does not broadcast at all.

Let T be some time at which all transmitters have regret at most ϵ . Our goal is to prove that the average number of successful connections up to time T has been close to $|OPT|$. For each transmitter u , let q_u be the fraction of times at which u chose to transmit (i.e. played action 1), and let s_u be the fraction of times at which u transmitted successfully. Then $Q = \sum_u q_u$ is the average number of attempted transmissions and $S = \sum_u s_u$ is the average number of successful transmissions, so we are trying to prove that S is close to $|OPT|$. The following lemma shows that S can be bounded by Q , and thus will allow us to only look at attempted broadcasts rather than successful broadcasts:

$$\text{Lemma 8: } S \leq Q \leq 2S + \epsilon n$$

Proof: The first inequality is obvious from the definitions, since the average number of successful transmissions is clearly at most the average number of attempted transmissions. For the second inequality, it is sufficient to show that $s_i \geq \frac{1}{2}(q_i - \epsilon)$ for all senders i . Suppose that $s_i < \frac{1}{2}(q_i - \epsilon)$ for some i . Then i 's average utility is $s_i - (q_i - s_i) = 2s_i - q_i < -\epsilon$. But i could

have had average utility of 0 by never broadcasting, which is a contradiction since i has regret at most ϵ . ■

III. PHYSICAL MODEL

We first consider the physical model. Our main theorem is that the price of total anarchy is small; in particular, we show that if all transmitters have low regret then the average number of successful transmissions is close to optimal:

Theorem 9: Suppose that at time T every sender has regret at most ϵ . Then the average number of successful transmissions is $\Omega(|OPT|/d_{max}^{2\alpha}) - \epsilon n$.

In this section we first prove some useful lemmas, then move on to the proof of Theorem 9, and then discuss some interesting extensions.

A. Density Lemmas

The following lemmas were proved in the pure physical model in [1]. We will use them extensively, so give proofs here in the generalized model (i.e. when the received signal strength might be off from the theoretical signal strength by up to a θ factor in either direction). We note that they are trivial from volume arguments in the physical model used in most previous work (e.g. [13], [7], [6]), in which the minimum distance between any pair of points is assumed to be at least 1.

Lemma 10: Consider a square S with side-length d_0 . In any feasible solution the maximum number of connections with a receiver in square S is $\theta^2 3^\alpha / \tau$.

Proof: Suppose that all nodes in the feasible solution transmit at a power such that the theoretical received signal strength is a constant \bar{p} , i.e. $p_i \min\{1, (d_0/d(t_i, r_i))^\alpha\} = \bar{p}$. Let i and i' be two connections such that both r_i and $r_{i'}$ lie in S .

The interference caused by connection i at receiver $r_{i'}$ is at least $\frac{1}{\theta} p_i \cdot \min\{1, (d_0/d(t_i, r_{i'}))^\alpha\} \geq \frac{1}{\theta} p_i \min\{1, (d_0/(d(r_i, r_{i'}) + d(t_i, r_i)))^\alpha\}$. By the geometry of the square S we know that $d(r_i, r_{i'}) \leq 2d_0$, which implies that $\frac{1}{\theta} p_i \min\{1, (d_0/(d(r_i, r_{i'}) + d(t_i, r_i)))^\alpha\} \geq \frac{1}{\theta 3^\alpha} p_i \min\{1, (d_0/d(t_i, r_i))^\alpha\} \geq \frac{\bar{p}}{\theta 3^\alpha}$. Since the actual signal received signal strength of a connection is at most $\theta \bar{p}$, if there are more than $\theta^2 3^\alpha / \tau$ such connections the interference experienced by all of them would be enough to prevent the SINR constraint being satisfied for *all* connections.

We now remove the condition that the received powers for every connection are the same. However, in this case the SINR value for some connection must be worse than it was when the received signal powers were the same, since we assume the distortion from the theoretical strength does not depend on the chosen power. This implies that if there are more than $3^\alpha / \tau$ connections, then for any set of transmission powers there will be *some* connection whose SINR constraint is not satisfied. ■

Corollary 11: Suppose now that square S has side-length d . In any feasible solution the maximum number of connections with a receiver in square S is $\theta^2 3^\alpha d^2 / \tau d_0^2$.

Proof: Divide square S up into subsquares of size d_0 and then apply Lemma 10. ■

Lemma 12: Now consider a ball B of radius d . In any feasible solution the maximum number of connections with a receiver in ball B is $\theta^2 3^\alpha \cdot 4d^2 / \tau d_0^2$.

Proof: Follows immediately from the fact that any circle with radius d is contained in a square with side-length $2d$. ■

B. Proof of Theorem 9

Before we dive into technical details we first give some intuition. We first find some receiver whose transmitter only rarely broadcasts but has “small” interference, where our notion of small is something that increases as S gets closer to OPT (so if the approximation is extremely bad then the interference is extremely small). Since its transmitter does not broadcast very often, but has low regret compared to the strategy of always broadcasting, this small interference must be enough to kill the transmission much of the time. Thus this interference must actually be quite large, so S must be close to OPT .

Let $L = \{i : q_i \geq 1/2 - \epsilon\}$ be the set of transmitters that broadcast at least $1/2 - \epsilon$ fraction of the time. Consider the following procedure. For each receiver x in $OPT \setminus L$ we will keep track of how much it is “bought” with a variable $b(x)$, initially set to 0. Now we order all transmitters in the instance (or just all transmitters with non-zero q_i) arbitrarily. We examine the transmitters one by one in this order. Say we are on transmitter i . Let $\Phi = \lfloor \frac{|OPT \setminus L| - k}{kQ} \rfloor$ (for some parameter k to be defined later) and let $R(i)$ be the Φ closest receivers in $OPT \setminus L$ to i that are currently bought to less than 1, i.e. have $b(x) < 1$. We now increase their b values by q_i , so $b(x) := b(x) + q_i$.

Since each transmitter i increases the sum of the b values by $q_i \Phi$, at the end of this process we know that

$$\sum_x b(x) = \sum_i q_i \Phi \leq \frac{|OPT \setminus L| - k}{k} < \frac{|OPT \setminus L|}{k}$$

since by definition $Q = \sum_i q_i$. This means that there is some receiver a that is in OPT but whose transmitter is not in L that has $b(a) < 1/k$.

Let M' be the set of transmitters that contributed to $b(a)$ during the above process. Note that since $b(a) \leq 1/k$ we know that $\sum_{x \in M'} q_x \leq 1/k$; we will use this later. Let M be all other transmitters, and for every distance d let $z(d) = \sum_{x \in M: d(a, x) \leq d} q_x$ be the average number of transmissions from transmitters in M located inside $B(a, d)$. Consider some transmitter $x \in M$. Since $a \notin R(x)$ and $b(a) < 1$, any receiver $y \in R(x)$ must have $d(x, y) \leq d(x, a)$, or else a would be in $R(x)$. So by the triangle inequality we know that $d(a, y) \leq 2d(a, x)$, and thus that any transmitter x at distance at most d from a must have its entire $R(x)$ at distance at most $2d$ from a .

We will now bound $z(d)$. Since every transmitter x in $M \cap B(a, d)$ contributes $q_x \Phi$ to the sum of the b values, and each receiver that it contributes to must be in $B(a, 2d)$, the sum of the b values of receivers in $B(a, 2d)$ is at least $z(d) \Phi$. Since a receiver’s b value only increases if it is less than 1, and then only increases by at most 1, we know that the b value of any receiver is at most 2. Thus the number of receivers from OPT

in $B(a, 2d)$ is at least $\frac{z(d)}{2}\Phi$. By Lemma 12, this implies that $cd^2 \geq \frac{z(d)}{2}\Phi$ and thus that

$$z(d) \leq \frac{2cd^2}{\Phi} \quad (1)$$

for some constant c depending only on α, τ , and d_0 .

Now that we have a bound on the average number of transmissions inside a ball around a , we want to bound the average interference at a . To do this, we will first bound the average number of transmissions in an annulus of radius d_0 , i.e. $z(d+d_0) - z(d)$. We first note that the interference at a is at most the interference caused if every ball around a actually meets the bound given by (1) and the received strength at a is θ times the theoretical strength. This is easily proved: let d be the first ball that doesn't meet the bound of (1). If there are no transmitters at distance greater than d from a , then clearly the average interference could be increased by adding more transmitters to every annulus past d so that the bound of (1) is met. If there are transmitters at distance greater than d from a , then clearly the average interference would be increased by moving enough of them into $B(a, d)$ to meet the bound. Now we can just keep repeating this process until there are no more transmitters past the first distance that fails to meet (1), reducing to the first case.

This now implies that we can treat inequality (1) as a lower bound as well as an upper bound, and thus the average number of transmissions coming from senders between distance d and $d+d_0$ is at most

$$\frac{2c}{\Phi}((d+d_0)^2 - d^2) \leq \frac{6cd_0d}{\Phi}$$

when $d \geq d_0$, and is at most $2c/\Phi$ when $d = 0$. Since the interference from a transmitter at distance d from a is at most $\theta \min\{1, (\frac{d_0}{d})^\alpha\}$, this means that the average interference at a caused by transmitters at distance between d and $d+d_0$ from a is at most $(6cd_0^{\alpha+1}/\Phi) \cdot \frac{\theta}{d^{\alpha-1}}$ for $d \geq d_0$. For $d = 0$, since the interference caused by a transmitter is at most θ the average interference from transmitters between distances 0 and d_0 from a is at most $2c\theta/\Phi$. Using linearity of expectations, we can sum over the annuli to get that the expected interference at a is at most

$$\begin{aligned} & \frac{2c\theta}{\Phi} + \frac{6c\theta d_0^{\alpha+1}}{\Phi} \sum_{i=1}^{\infty} \frac{1}{(id_0)^{\alpha-1}} \\ &= \frac{2c\theta(1 + 3d_0^2\zeta(\alpha-1))}{\Phi} \leq \frac{8c\theta\zeta(\alpha-1)}{\Phi} \end{aligned}$$

where $\zeta(\alpha-1)$ is the Riemann zeta function (which will be constant for $\alpha > 2$) and we are assuming $d_0 \leq 1$. If $d_0 > 1$ then we will simply have d_0^2 as a constant to carry through the rest of the calculations, which will not matter since we are not attempting to optimize constants anyway.

This gives us a bound on the average interference at a caused by transmitters in M . What about the transmitters in M' ? Since we know that $\sum_{x \in M'} q_x \leq \frac{1}{k}$, it is obvious that they cause at most $\frac{\theta}{k}$ average interference (which is what would happen if they were all at distance d_0 or less from a). Thus the total

expected interference is at most

$$\frac{8c\theta\zeta(\alpha-1)}{\Phi} + \frac{\theta}{k}$$

So now we have a bound on the average interference. Let $p_{\text{bad}}(a)$ denote the fraction of times in which a 's transmitter could not succeed in transmitting, whether it tried or not. If $p_{\text{bad}}(a) < \frac{1}{4}$, then sending at every time would give average utility greater than $\frac{3}{4} - \frac{1}{4} = \frac{1}{2}$. But when we chose a we made sure that its transmitter (call it t) was from $OPT \setminus L$, so we know that $q_t < \frac{1}{2} - \epsilon$. Thus t only tries to transmit less than $\frac{1}{2} - \epsilon$ the time, and hence it has an average utility of less than $\frac{1}{2} - \epsilon$. This is a contradiction since we are assuming that t has regret at most ϵ , and thus $p_{\text{bad}}(a) \geq \frac{1}{4}$.

So a 's transmitter would fail at least $\frac{1}{4}$ of the time if it tried to send every time, and the average interference is at most $(8c\theta\zeta(\alpha-1)/\Phi) + \frac{\theta}{k}$. By Markov's inequality we know that the fraction of times at which a hears interference at least four times the average interference is at most $\frac{1}{4}$, so the interference at a is at least $(32c\theta\zeta(\alpha-1)/\Phi) + \frac{4\theta}{k}$ at most $\frac{1}{4}$ of the time. So this amount of interference must be enough to make it impossible for a to successfully receive (i.e. the SINR constraint would be violated), or else $p_{\text{bad}}(a)$ would be less than $\frac{1}{4}$. Since a is at distance at most d_{max} from its transmitter, the strength of its signal is at least $\frac{d_0^\alpha}{\theta d_{\text{max}}^\alpha}$. Thus we get that $32c\theta\zeta(\alpha-1)/\Phi \geq \frac{d_0^\alpha}{\tau\theta d_{\text{max}}^\alpha} - \frac{4\theta}{k}$. We will now finally set k , to $8\theta^2\tau d_{\text{max}}^\alpha/d_0^\alpha$, giving us that

$$\frac{32c\zeta(\alpha-1)}{\left[\frac{|OPT \setminus L| - 8\tau d_{\text{max}}^\alpha}{8\tau d_{\text{max}}^\alpha Q}\right]} \geq \frac{d_0^\alpha}{2\theta\tau d_{\text{max}}^\alpha}. \quad (2)$$

Solving for Q in this equation, and assuming constant α, τ , and d_0 , gives us that $Q \geq \Omega(|OPT \setminus L|/d_{\text{max}}^{2\alpha})$. To compare Q to $|OPT|$ instead of $|OPT \setminus L|$, we note that if $|OPT \setminus L| < \frac{1}{2}|OPT|$ then at least half of the transmitters in OPT are broadcasting at least $\frac{1}{2} - \epsilon$ of the time, and thus $Q \geq \frac{1}{2}|OPT|(\frac{1}{2} - \epsilon) = \Omega(|OPT|)$. On the other hand, if $|OPT \setminus L| \geq \frac{1}{2}|OPT|$ then we get that $Q \geq \Omega(|OPT|/d_{\text{max}}^{2\alpha})$. Now we can simply apply Lemma 8 to prove that the average number of successful connections is at least $\Omega(|OPT|/d_{\text{max}}^{2\alpha} - \epsilon n)$, thus proving the theorem.

C. Other Metrics

While the generalized physical model allows for signal strength to vary by θ , it still assumes that the fundamental underlying metric is the Euclidean plane. However, we only used this assumption in one place: the proof of Lemma 12, the main density lemma, which proved that the number of receivers from any feasible set of transmissions contained in a ball of radius d is at most $O(d^2)$. We then used this lemma to bound the average interference, which will work whenever the exponent in the density lemma is strictly less than the path-loss exponent α . So actually our proof will work in any metric in which the number of receivers from a feasible solution contained in a ball of radius d is $o(d^\alpha)$.

One example of this is true three-dimensional space. In this case, it makes sense to assume that $\alpha > 3$, since power is being

dissipated in three dimensions (so $\alpha \geq 3$) and some is probably being lost due to being absorbed by objects (or just the air). On the other hand, a sphere of radius d can clearly be covered with $O((d/d_0)^3)$ spheres of radius d_0 , and thus we can immediately derive the appropriate density lemma from Lemma 10 (which still holds as stated).

But even making only the weakest standard assumption, that $\alpha > 2$, we can still handle extra metrics; for example, any metric with doubling dimension 2. The *doubling dimension* of a metric is defined to be the smallest number k for which for all distances $d \in \mathbb{R}_{>0}$, any ball of radius d can be covered by 2^k balls of radius $d/2$. Suppose we have a metric with doubling dimension k . Then by recursive applications of the definition of doubling dimension we get that any ball of radius d can be covered by $2^{k \log(d/d_0)} = (d/d_0)^k$ balls of radius d_0 , and thus we can once again apply Lemma 10 to get Lemma 12, and as long as $\alpha > k$ the rest of the proof will go through as before.

Finally, we consider one class of metrics for which a good density lemma does *not* hold, but for intuitively unrealistic reasons: the *wireless manifold*, introduced by Kanade and Vempala [10]. Intuitively, they define the class of wireless manifolds as the class of distorted two-dimensional grids. In particular, consider a $k \times k$ grid, with an arbitrary nonnegative length assigned to every grid edge. Now let the distance between two points be the length of the shortest path between them in this weighted graph. In their paper [10], Kanade and Vempala give heuristics for finding the best such manifold given signal strength data, and show that for existing data sets the best wireless manifold is significantly more accurate than the best embedding into the Euclidean plane. Thus it is natural to ask whether our techniques extend to these manifolds. Unfortunately they do not: the following theorem shows that the density lemma we require for our proof to work is not true.

Theorem 13: For any $d > 0$ there is a wireless manifold and a set of $\Omega(d^\alpha)$ feasible transmissions on this manifold such that all receivers are in a ball of radius d .

Proof: We first note that we can embed uniformly weighted complete graphs into a wireless manifold. To see this, suppose that we want to embed a complete graph with k vertices and weight d on every edge. We first create a square S with $\lceil k/4 \rceil$ vertices on each side, and set the length of all edges with both endpoints in S to 0. We can then take k edges with one endpoint in S and the other outside of S and set their lengths to $d/2$. All other edges will have weight M for some $M \gg d$. The resulting metric is clearly the required complete graph, where the non- S endpoints of the k edges we chose are the vertices.

Now that we can embed complete graphs, consider the same complete graph on k vertices with distance d between them. Put a transmitter/receiver pair on each vertex of the complete graph (so t_i and r_i are co-located). Set $\theta = 1$. If all transmitters broadcast at power 1, the signal at r_i is 1 and the interference is $(k-1)\frac{1}{d^\alpha}$. So as long as $k \leq \frac{d^\alpha}{\tau} + 1$ the SINR at receiver r_i is at least τ , and thus the connection is supported. Thus in a ball of radius d we can support $\Omega(d^\alpha)$ connections. ■

D. Byzantine Transmitters

In many cases it is not realistic to assume that every single transmitter will be running a no-regret algorithm, so when designing a distributed algorithm we would like to be robust to some fraction of the transmitters behaving in arbitrary ways. We manage to achieve this, but only if the number of transmitters that are Byzantine is a fraction of $|OPT|$, not if it depends on n .

We generalize the proof of Theorem 9 in a straightforward way. First, let B be the set of Byzantine transmitters, and let OPT be optimal relative to whatever the Byzantine transmitters do (but OPT is still a fixed set that would be feasible at every time point if not for B). Let Q be defined as before, but let \bar{Q} be the part of the sum that comes only from transmitters not in B . Note that $Q \leq \bar{Q} + |B|$. Also note that Lemma 8 still holds, except with \bar{Q} instead of Q and where S is only summed over transmitters not in B .

It is easy to see that the analysis of Theorem 9 holds as stated for Q , since the node a we find will be in OPT and thus non-Byzantine and the packing argument works as before since OPT is feasible on its own. So $Q = \Omega(|OPT|/d_{\max}^{2\alpha})$. If $|B| \leq (1-\delta)Q$ then we know that $\bar{Q} \geq \delta Q \geq \Omega(\delta|OPT|/d_{\max}^{2\alpha})$, and can apply Lemma 8 to finish the proof. In particular, by setting δ to $1/2$ we see that there is some constant k such that if at most $k|OPT|/d_{\max}^{2\alpha}$ transmitters are Byzantine then we still get a $O(d_{\max}^{2\alpha})$ approximation to $|OPT|$.

IV. PROTOCOL MODEL

As discussed, in the *protocol model* each transmission is a node in an interference graph, and a transmission is successful if none of its neighbors are also transmitting. Clearly maximizing capacity is just the same problem as finding a maximum independent set in the interference graph. The classic theoretical model used for these graphs are unit disk graphs, but we will generalize to all locally growth bounded graphs. We show how to use the same basic technique as in the SINR model, i.e. proving that the price of total anarchy is small for a particular game, to give a distributed algorithm that has good average performance. While we will not obtain either as good an approximation or as small a running time as [18], our algorithm is totally distributed in the sense that the only information each node gets is whether or not any of its neighbors tried to join the independent set in the last round.

We first show that in general graphs the average number of successful transmissions when every transmitter uses a no-regret algorithm can be arbitrarily far from the size of even the smallest *maximal* independent set. We then show that for growth bounded graphs, after a sufficient number of rounds, the average number of nodes that broadcasted successfully in a round is within a constant of the size of the the maximum independent set.

A. General Graphs

Consider the following interference graph: there are two special nodes u and v that are adjacent. The node u is also adjacent to $(n-2)/2$ nodes $x_1, x_2, \dots, x_{(n-2)/2}$, none of which are adjacent to v , and v is also adjacent to $(n-2)/2$




Fig. 1. Bad interference graph

nodes $y_1, y_2, \dots, y_{(n-2)/2}$, none of which are adjacent to u (see Figure 1). Clearly the smallest maximal independent set has size $\frac{n-2}{2} + 1 = n/2$, as it consists of either u with all the y 's or v with all the x 's. On the other hand, suppose that u and v each choose to broadcast independently with probability $1/2$ and all of the x_i 's and y_i 's never broadcast. Obviously the expected number of successful transmissions given these strategies is $1/2$, so it is only an $\Omega(n)$ -approximation. We claim that in this case, every transmitter is using a no-regret algorithm.

To see this, first consider some x_i . Since u broadcasts with probability $1/2$, the expected average utility of x_i if it chose to broadcast every time is 0, and obviously never broadcasting has average payoff 0. Thus never broadcasting is a no-regret algorithm. The same argument can be used for any y_i , with v taking the place of u . Now consider u . It too is adjacent to one node v that broadcasts with probability $1/2$, so both of its actions would give average utility 0. Clearly randomizing independently over these two actions also has an average utility of 0, and thus is a no-regret strategy. The same argument obviously works for v as well, finishing the proof.

B. Growth-Bounded Graphs

When we restrict to growth bounded graphs, no-regret algorithms actually do result in good behavior. This is essentially because the bad example we designed for general graphs can't occur, since in growth bounded graphs the number of neighbors of a node that are in an independent set is at most a constant. In other words, the definition of growth bounded graphs get us the equivalent of the density lemmas of Section III-A.

Theorem 14: In growth-bounded graphs in which at most k neighbors of any node are in an independent set, if all transmitters have regret at most ϵ then the average number of successful connections is at least $\Omega(|OPT|) - \epsilon n$.

Proof: Let x_0 be some node, with neighbors x_1, \dots, x_m . We first claim that $\sum_{i=0}^m q_{x_i} \geq \frac{1}{3} - \frac{\epsilon}{2}$. The first case is if $q_{x_0} \geq 1/3$, in which case the claim is trivially true. So suppose that $q_{x_0} < 1/3$. Let p_{bad} be the fraction of times that x_0 would be unsuccessful if it chose to broadcast, i.e. the fraction of times at which at least one of its neighbors broadcasts. Then the average utility from always broadcasting is $1 - p_{\text{bad}} - p_{\text{bad}} = 1 - 2p_{\text{bad}}$. If $p_{\text{bad}} \leq 1/3 - \epsilon/2$, then the average utility is from always broadcasting is at least $1/3 + \epsilon$. But this is a contradiction since x_0 has average utility at most $q_0 < 1/3$ but also has ϵ -regret. Thus we know that $p_{\text{bad}} > 1/3 - \epsilon/2$. But clearly $p_{\text{bad}} \leq \sum_{i=1}^m q_{x_i}$ by definition, since in order for a time to contribute to p_{bad} at least one of the x_i 's need to be broadcasting. Hence we have proved that $\sum_{i=0}^m q_{x_i} \geq \frac{1}{3} - \frac{\epsilon}{2}$.

Now we relate $|OPT|$ to Q . For nodes $x \in OPT$, let $b(x) = \sum_{y \in N(x) \cup \{x\}} q_y$. By the above claim we know that $b(x) \geq \frac{1}{3} - \frac{\epsilon}{2}$, and thus $\sum_{x \in OPT} b(x) \geq |OPT|(\frac{1}{3} - \frac{\epsilon}{2})$. But every node is adjacent to at most k nodes from OPT by the growth-boundedness of the graph, and thus $\sum_{x \in OPT} b(x) \leq k \sum_u q_u = kQ$. Putting these together, we get that $Q \geq \frac{1}{k}(\frac{1}{3} - \frac{\epsilon}{2})|OPT|$. We now apply Lemma 8 to get that

$$S \geq \frac{1}{2}(Q - \epsilon n) \geq \frac{1}{2} \left(\left(\frac{1}{3k} - \frac{\epsilon}{2k} \right) |OPT| - \epsilon n \right) \geq \Omega(|OPT| - \epsilon n)$$

as claimed. \blacksquare

V. SIMULATIONS

While the main contributions of this work are theoretical, we also performed simulations to show how no-regret algorithms do in practice. There are two aspects that we would like to test: the quality of the algorithm (i.e. the average number of successful transmissions) and the speed at which the algorithm converges on that average. Our simulations will be in the vanilla physical model, where transmitters and receiver are points in the Euclidean plane and $\theta = 1$. We will be using random topologies, in which n transmitter-receiver pairs are placed uniformly at random in a square of size of size 100×100 in the Euclidean plane, subject to each receiver being at distance at most d_{max} from its transmitter. Throughout these simulations we will set $\alpha = 2.1$ and $\tau = 0.5$, since it turns out that changing these parameters does not change the trends by very much.

For the quality simulations, we will compare our algorithm to the *Best Response* distributed algorithm proposed by Andrews and Dinitz in [1], as it seems to be the only proposed distributed algorithm with worst case guarantees in the literature. In particular, we will compare what happens when every transmitter uses Best Response to what happens when every transmitter uses the classic no-regret algorithm *Randomized Weighted Majority (WMR)* of Littlestone and Warmuch [12].

Our quality simulation shows the relationship between the number of nodes and the average number of successful transmissions per round after simulating for 100 rounds. We did this on 100 instances for each value of n and averaged the results. Figure 2 shows that as n gets larger our algorithm does better, while Best Response does about the same (or slightly worse). Note that large n is the only interesting regime, since only when n is large is there a lot of interference from other transmitters. This figure also shows that the approximation bound we proved, $O(d_{\text{max}}^2)$, is overly pessimistic, since for all three values of d_{max} that we tested the actual performance is significantly better than n/d_{max}^2 , and clearly n is an upper bound on $|OPT|$. For example, when $d_{\text{max}} = 8$ and $n = 1000$, we observed an average of 138.861 successful transmissions, while $1000/8^{2.1}$ is only 12.69.

For the convergence speed simulations, instead of comparing to the Best Response algorithm we just test the average number of successful transmissions after various iterations. Our analysis requires $\Omega(n^2 \log n)$ iterations before the approximation guarantee can be made, but our simulations show that in practice this number of iterations is not necessary. As shown in Figure 3,




Fig. 2. Random topology, 100 iterations




Fig. 3. Random topology, $d_{max} = 10$

substantially fewer than n iterations are required before the average is basically stable. The time to stability is not constant, as it does seem to grow with n , but it is certainly substantially smaller than the $\Omega(n^2 \log n)$ bound required by the analysis.

VI. CONCLUSIONS

In this paper we have designed distributed algorithms for maximizing wireless network capacity in both the physical and the protocol models. Our main techniques were game theoretic, namely proving that the price of total anarchy of an appropriately defined game is small. We hope that this technique will prove fruitful when considering other distributed problems, especially when only extremely limited feedback is allowed. We also showed by simulation that low-regret algorithms do even better in practice than the theoretical worst case, both in terms of their approximation to optimal and the time it takes to achieve this approximation.

However, a number of open problems remain even in the realm of network capacity. Most obviously, in the physical model we only get an $O(d_{max}^{2\alpha})$ -approximation, and a natural question is whether this dependence on the distance could be removed, resulting in an $O(1)$ -approximation. It would also be interesting to try to decrease the time necessary before the approximation holds from $O(n^2 \log n)$ to something more reasonable, perhaps linear or even polylogarithmic.

ACKNOWLEDGEMENTS

We would like to thank Matthew Andrews for introducing us to network capacity problems and Katrina Ligett and Aaron Roth for helpful discussions about the price of total anarchy.

REFERENCES

- [1] M. Andrews and M. Dinitz, "Maximizing capacity in arbitrary wireless networks in the SINR model: complexity and game theory," in *INFOCOM*, 2009.
- [2] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire, "The non-stochastic multiarmed bandit problem," *SIAM J. Comput.*, vol. 32, no. 1, pp. 48–77, 2003.
- [3] A. Blum, M. Hajiaghayi, K. Ligett, and A. Roth, "Regret minimization and the price of total anarchy," in *STOC '08: Proceedings of the 40th annual ACM Symposium on Theory of Computing*. New York, NY, USA: ACM, 2008, pp. 373–382.
- [4] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou, "The complexity of computing a nash equilibrium," in *STOC '06: Proceedings of the thirty-eighth annual ACM Symposium on Theory of Computing*. New York, NY, USA: ACM, 2006, pp. 71–78.
- [5] T. Erlebach, K. Jansen, and E. Seidel, "Polynomial-time approximation schemes for geometric graphs," in *SODA '01: Proceedings of the twelfth annual ACM-SIAM Symposium on Discrete Algorithms*. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2001, pp. 671–679.
- [6] O. Goussevskaia, M. Halldórsson, R. Wattenhofer, and E. Welzl, "Capacity of arbitrary wireless networks," in *INFOCOM*, 2009.
- [7] O. Goussevskaia, Y. A. Oswald, and R. Wattenhofer, "Complexity in geometric SINR," in *MobiHoc '07: Proceedings of the 8th ACM International Symposium on Mobile Ad Hoc Networking and Computing*. New York, NY, USA: ACM, 2007, pp. 100–109.
- [8] P. Gupta and P. Kumar, "The capacity of wireless networks," *IEEE Transactions on Information Theory*, vol. 46, no. 2, pp. 388–404, Mar 2000.
- [9] M. Halldórsson and R. Wattenhofer, "Wireless Communication is in APX," in *36th International Colloquium on Automata, Languages and Programming (ICALP), Rhodes, Greece, July 2009*.
- [10] V. Kanade and S. Vempala, "Life (and routing) on the wireless manifold," in *Sixth Workshop on Hot Topics in Networks (HotNets-VI)*. New York, NY, USA: ACM, 2007.
- [11] F. Kuhn, T. Moscibroda, T. Nieberg, and R. Wattenhofer, "Fast deterministic distributed maximal independent set computation on growth-bounded graphs," in *DISC*, 2005, pp. 273–287.
- [12] N. Littlestone and M. Warmuth, "The weighted majority algorithm," *Symposium on Foundations of Computer Science*, vol. 0, pp. 256–261, 1989.
- [13] T. Moscibroda and R. Wattenhofer, "The complexity of connectivity in wireless networks," *INFOCOM 2006. 25th IEEE International Conference on Computer Communications. Proceedings*, pp. 1–13, April 2006.
- [14] T. Moscibroda, R. Wattenhofer, and A. Zollinger, "Topology control meets SINR: the scheduling complexity of arbitrary topologies," in *MobiHoc '06: Proceedings of the 7th ACM International Symposium on Mobile Ad Hoc Networking and Computing*. New York, NY, USA: ACM, 2006, pp. 310–321.
- [15] T. Nieberg, J. Hurink, and W. Kern, "Approximation schemes for wireless networks," *ACM Trans. Algorithms*, vol. 4, no. 4, pp. 1–17, 2008.
- [16] T. Rappaport, *Wireless Communications: Principles and Practice*. Upper Saddle River, NJ, USA: Prentice Hall PTR, 2001.
- [17] C. Saraydar, N. Mandayam, and D. Goodman, "Efficient power control via pricing in wireless data networks," *IEEE Transactions on Communications*, vol. 50, no. 2, pp. 291–303, Feb 2002.
- [18] J. Schneider and R. Wattenhofer, "A log-star distributed maximal independent set algorithm for growth-bounded graphs," in *PODC '08: Proceedings of the twenty-seventh ACM Symposium on Principles of Distributed Computing*. New York, NY, USA: ACM, 2008, pp. 35–44.
- [19] A. Stolyar and H. Viswanathan, "Self-organizing dynamic fractional frequency reuse in ofdma systems," *INFOCOM 2008. The 27th IEEE Conference on Computer Communications*, pp. 691–699, April 2008.
- [20] S. Yi, Y. Pei, and S. Kalyanaraman, "On the capacity improvement of ad hoc wireless networks using directional antennas," in *MobiHoc '03: Proceedings of the 4th ACM International Symposium on Mobile Ad Hoc Networking & Computing*. New York, NY, USA: ACM, 2003, pp. 108–116.