

 Open access • Proceedings Article • DOI:10.1109/INFCOM.2009.5062048

## Maximizing Capacity in Arbitrary Wireless Networks in the SINR Model: Complexity and Game Theory — [Source link](#)

Matthew Andrews, Michael Dinitz

**Institutions:** Bell Labs, Carnegie Mellon University

**Published on:** 19 Apr 2009 - International Conference on Computer Communications

**Topics:** Price of anarchy, Stochastic game, Nash equilibrium, Approximation algorithm and Game theory

Related papers:

- [Capacity of Arbitrary Wireless Networks](#)
- [The capacity of wireless networks](#)
- [Complexity in geometric SINR](#)
- [The Complexity of Connectivity in Wireless Networks](#)
- [A constant-factor approximation for wireless capacity maximization with power control in the SINR model](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/maximizing-capacity-in-arbitrary-wireless-networks-in-the-548wsph6tv>

# Maximizing Capacity in Arbitrary Wireless Networks in the SINR Model: Complexity and Game Theory

Matthew Andrews

Bell Labs, Murray Hill NJ

Email: andrews@research.bell-labs.com

Michael Dinitz

Carnegie Mellon University, Pittsburgh PA

Email: mdinitz@cs.cmu.edu

**Abstract**—In this paper we consider the problem of maximizing the number of supported connections in arbitrary wireless networks where a transmission is supported if and only if the signal-to-interference-plus-noise ratio at the receiver is greater than some threshold. The aim is to choose transmission powers for each connection so as to maximize the number of connections for which this threshold is met.

We believe that analyzing this problem is important both in its own right and also because it arises as a subproblem in many other areas of wireless networking. We study both the complexity of the problem and also present some game theoretic results regarding capacity that is achieved by completely distributed algorithms. We also feel that this problem is intriguing since it involves both continuous aspects (i.e. choosing the transmission powers) as well as discrete aspects (i.e. which connections should be supported). Our results are:

- We show that maximizing the number of supported connections is NP-hard, even when there is no background noise. This is in contrast to the problem of determining whether or not a given set of connections is feasible since that problem can be solved via linear programming.
- We present a number of approximation algorithms for the problem. All of these approximation algorithms run in polynomial time and have an approximation ratio that is independent of the number of connections.
- We examine a completely distributed algorithm and analyze it as a game in which a connection receives a positive payoff if it is successful and a negative payoff if it is unsuccessful while transmitting with nonzero power. We show that in this game there is not necessarily a pure Nash equilibrium but if such an equilibrium does exist the corresponding price of anarchy is independent of the number of connections. We also show that a mixed Nash equilibrium corresponds to a probabilistic transmission strategy and in this case such an equilibrium always exists and has a price of anarchy that is independent of the number of connections.

This work was supported by NSF contract CCF-0728980 and was performed while the second author was visiting Bell Labs in Summer, 2008.

## I. INTRODUCTION

In this paper we consider the problem of maximizing the number of successful transmissions in the physical SINR model. In our basic model we are given a set of transmitter-receiver pairs located in the plane and each has an associated SINR requirement. The aim is to satisfy as many of the requirements as possible.

Maximizing the transmission capacity in wireless networks has been studied in many contexts. Typically this work can

be partitioned along two axes. On one axis we have the two models that are typically used to model channel conditions. The simplest case is the unit-disk graph (UDG) model in which transmissions interfere if and only if they are within distance 1.<sup>1</sup> A more complex model is the SINR model in which each transmission is given a power and we assume a distance-dependent path loss. A transmission is deemed to be successful if the signal-to-interference-plus-noise-ratio (SINR) is more than some specified threshold.

On the other axis is the structure of the networks that are being considered. One option is to look at random networks under a certain distribution of node placements and transmitter-receiver pairings. In this case the typical goal is to calculate the expected capacity of the system and examine how it changes as the density of the network increases. Another option is to simply look at a worst-case topology. In this case it makes no sense to consider some notion of average capacity since that could depend greatly on what the topology looks like, and so we are more interested in the complexity of calculating the optimum capacity and in determining how close we can come to optimality via efficient algorithms (both centralized algorithms and distributed protocols).

For random networks, the problem of calculating transmission capacity has been examined in both the UDG model and the SINR model by Gupta and Kumar [10]. They showed that in both cases the average source-sink capacity scales as  $\Theta(\frac{1}{\sqrt{n \log n}})$  where  $n$  is the number of nodes in the network.<sup>2</sup> For UDGs, the maximum cardinality set of transmissions corresponds to a maximum independent set in the graph. Maximum independent set in UDGs was shown to be NP-

<sup>1</sup>We remark that there are many variants of the UDG model. For example in the Tx model of [23] a transmission from  $u_2$  to  $v_2$  suffers interference from a transmission from  $u_1$  to  $v_1$  if and only if the distance between  $u_2$  to  $u_1$  is at most the distance from  $u_1$  to  $v_1$ .

<sup>2</sup>A difference between Gupta-Kumar and our work (aside from the fact that they look at random networks whereas we look at worst-case networks) is that they consider the case of multihop transmissions. (This is of course necessary in their setting since they consider random source-destination pairs which may not be supportable with one-hop transmissions.) However, for simplicity we shall focus on the case of single hop transmissions. (This is for example the case that arises when we apply our results to the scheduling problems that we discuss in Section I-E.) We plan to extend our results to multi-hop transmissions in future work. We believe that this could be done using ideas similar to those used by Kumar et al. [13] for solving multicommodity flow problems in wireless networks.

hard in [4]. However, it is known due to arguments about the geometry of disks that the simple greedy algorithm in which we continually pick nodes and delete all their neighbors leads to a 6-approximation for Maximum Independent Set in Unit Disk Graphs [14]. Moreover, for any  $\varepsilon > 0$  more complex algorithms can give a  $(1+\varepsilon)$ -approximation in polynomial time [12], [17], [7].

There has been less work on the complexity of calculating the maximum possible capacity in arbitrary networks under the SINR model. (We discuss some previous work on this problem in Section I-D.) We believe that addressing this question is important for two reasons. First, although analyzing capacity in random networks is important for determining what level of transmissions will be possible in completely unstructured networks, there are many situations where the network will have some sort of structure and the transmission capacity may be very different than what is possible in random networks. In these cases we believe that knowing the complexity of calculating the maximum number of transmissions is important. Second, as is well-known the Unit Disk Graph model does not capture many features of wireless networks. One reason for this is that receivers will hear interference from all other transmitters, even if they are far away. A more important reason is that interference at a receiver is a *cumulative* effect of multiple transmitters whereas in the Unit Disk Graph model interference is simply a local binary property.

Before we can describe our results in detail we must present our model. We begin with a basic model that allows us to demonstrate most of our techniques and show what results can be obtained in this model. We then define a number of extensions to the model and describe how our results change in these cases.

#### A. Basic model

We consider a set of  $n$  connections in the plane. Each connection  $i$  has a transmitter  $t_i$  and a receiver  $r_i$ . We let  $d(u, v)$  be the Euclidean distance between two points  $u$  and  $v$ . We use  $d_i$  to denote  $d(t_i, r_i)$  and refer to it as the *distance* of connection  $i$ . Suppose that a node  $u$  is transmitting with power  $p$ . Following [18] we assume that for some parameters  $d_0$  and  $\alpha$  the received signal at another point  $v$  is given by  $p \cdot \min\{(d_0/d(u, v))^\alpha, 1\}$ . We refer to  $\min\{(d_0/d(u, v))^\alpha, 1\}$  as the path loss between  $u$  and  $v$  and denote it by  $g(u, v)$ . We refer to  $\alpha$  as the path loss exponent. We make the traditional assumption that  $\alpha > 2$  since otherwise the total energy received over the plane would be more than the total energy transmitted.

We assume that for any connection  $i$  the distance  $d(t_i, r_i)$  is either 0 or else lies between  $d_{\min}$  and  $d_{\max}$  for some parameters  $d_{\max}, d_{\min} \geq d_0$ . The running times and the performance guarantees of many of our algorithms will depend on the ratio  $d_{\max}/d_{\min}$ . For simplicity in the remainder of the paper we shall normalize distances so that  $d_{\min} = 1$ . Hence in the following all of our formulas with a dependence on  $d_{\max}$  will in fact depend on  $d_{\max}/d_{\min}$  in the unnormalized case.

Let  $p_i$  be the power used by connection  $i$  (which can be zero). We assume that there is a maximum power  $p_{\max}$  with which any node can transmit. The *signal* received at receiver  $r_i$  is given by  $p_i g(t_i, r_i)$  and the interference heard from the

other connections is  $\sum_{j \neq i} p_j g(t_j, r_i)$ . We also assume that there is some background noise level  $W$  and so the signal-to-interference-plus-noise-ratio (SINR) is  $(p_i g(t_i, r_i))/(W + \sum_{j \neq i} p_j g(t_j, r_i))$ . We assume that each connection is for a single application type such as Voice-over-IP for which there is a fixed signal-to-noise requirement that we denote by  $\tau$ . In other words, connection  $i$  is satisfied if and only if  $(p_i g(t_i, r_i))/(W + \sum_{j \neq i} p_j g(t_j, r_i)) \geq \tau$ .

Our aim in this paper is to maximize the number of satisfied connections, i.e. we wish to choose the transmission power levels  $p_i$  so as to maximize,

$$|\{i : \frac{p_i g(t_i, r_i)}{W + \sum_{j \neq i} p_j g(t_j, r_i)} \geq \tau\}|.$$

We refer to this problem as MAX-CONNECTIONS and we denote the maximum achievable value by OPT.

We remark that if OPT equals  $n$ , then the optimum powers can be found using linear programming. This is because all we need to do is find powers such that  $p_i \geq 0$  and,

$$p_i g(t_i, r_i) \geq \tau(W + \sum_{j \neq i} p_j g(t_j, r_i)). \quad (1)$$

If such powers exist then clearly any linear programming algorithm will find them. Moreover, there are also many distributed algorithms for finding these powers. See for example the work of Yates on uplink power control [22]. However, we are mostly concerned with situations where it is not possible to support all connections. In this case linear programming will not work since we need to make the *discrete* decision about which connections to support before we make the *continuous* decision about what power levels to use for the supported connections.

We can think of the problem as being one of maximizing the number of satisfied linear inequalities of the form (1). In general the problem of maximizing the number of satisfied inequalities in a linear system cannot be approximated to within a factor better than  $n^\delta$  for some  $\delta > 0$ . (See Arora et al. [1].) However, our problem has significant geometric structure. Our aim is to exploit this structure to get better bounds than the bounds of [1].

Throughout the paper we shall assume that  $\alpha$  and  $\tau$  are constants. Hence expressions that utilize  $O(\cdot)$  will sometimes be hiding dependencies on  $\alpha$  and  $\tau$ . In the case when  $W \neq 0$ , we will also make the assumption that  $d_{\max}$  is bounded away from the absolute distance limit with no outside interference, i.e. there is some constant  $\epsilon$  such that  $d_{\max} \leq (1 - \epsilon)d_0(p_{\max}/\tau)^{1/\alpha}$ .

#### B. Results from the basic model

- Our first result is a hardness result. In Section II we show that MAX-CONNECTIONS is NP-hard and so we should not expect to obtain a polynomial-time exact algorithm.
- Given that the problem is NP-hard, in Section III we turn our attention to approximation algorithms for MAX-CONNECTIONS. Our first algorithm runs in polynomial time and gives an  $O(\log d_{\max})$ -approximation. For the case of zero background noise we describe a second algorithm that gives an  $O(1)$ -approximation in time  $n^{O(d_{\max}^2)}$ .

- The approximation algorithms presented in Section III are centralized. Although this might be appropriate in a situation where we are given a network configuration and we wish to analyze the capacity, centralized algorithms are unlikely to be useful if we wish to optimize capacity as a network evolves. Distributed algorithms are much more likely to be useful. In Section IV we consider the extreme case of completely decentralized algorithms that do not exchange any information but instead selfishly maximize their payoffs in a game that we design in which a strategy is a transmit power.

We first show that our game does not always have a pure Nash equilibrium. On the other hand, we show that in any mixed Nash equilibrium (of which there is always at least one) the expected number of connections that are supported is always within a  $O(d_{\max}^{2\alpha})$  factor of OPT (i.e. the price of anarchy is  $O(d_{\max}^{2\alpha})$ ). Thus if a pure Nash does exist it is close to optimal.

### C. Extended model

We now briefly describe some ways to extend our model together with the results that we can obtain when these new features are introduced.

- The first extension is to assume that each connection  $i$  has a weight  $w_i$  and the goal is to maximize the weighted total of supported connections. In this model we can slightly modify the proof from the basic model to obtain a similar  $O(\log d_{\max})$ -approximation algorithm.
- Another extension is to assume that there are multiple carriers in the system that do not interfere. Each connection must be assigned to a separate carrier. These carriers might be different channels in an 802.11 system or they might be different frequency bands in an OFDM system such as 3GPP's Long-Term Evolution (LTE) standard. In this case we have three decisions to make, namely which connections should be supported, which powers should they be assigned, and which channels should they be assigned. We remark that the third problem can be thought of as providing a frequency reuse pattern for the connections. In this model all of the results from the basic model continue to hold other than losing another constant factor independent of the number of carriers.

### D. Related work

As already mentioned, Gupta and Kumar [10] looked at the problem of maximizing the number of satisfied connections in random multihop networks in both the SINR model and the UDG model. For arbitrary networks under the UDG model,  $(1+\varepsilon)$ -approximations were obtained in [12], [17], [7]. In terms of completely distributed algorithms, game theoretic results have been obtained in a number of different contexts. In [11], Huang et al. looked at the problem of maximizing an aggregate network utility in a situation where each connection's utility is a concave function of the connection rate and nodes are allowed to share pricing information. They show that this distributed algorithm will converge to a local optimum. In [19], Saraydar et al. look at a game-theoretic algorithm for choosing powers on the uplink of a single cell wireless system. In [20], Stolyar and

Viswanathan study fractional frequency reuse algorithms for joint channel assignment and power control in cellular OFDM systems and provide a game theoretic algorithm that always leads to a stable solution. In [2], Bahl et al. provide distributed algorithms inspired by game theory for the problem of sizing cells and assigning users to basestations. However, none of this work and to the best of our knowledge no other work studies the problem of comparing the quality of a stable solution with the global optimum.

Previous papers that consider the complexity of capacity maximization in the SINR model include [3], [5], [9], [15], [16]. Of these, Goussevskaia et al. [9] is the most closely related to our work. The authors show NP-hardness and provide  $O(\log d_{\max})$  approximation algorithms for a similar objective to ours. They also consider a related objective of minimizing the number of "rounds" required to serve all connections. However, a key distinction between our work and [9] is that [9] assumes that all transmission powers are fixed. In other words it only addresses the combinatorial aspects of the problem (deciding which connections should be scheduled at a given time) and does not consider that the continuous aspects (which transmission powers should be used). In contrast, our results in Sections II and III assume that selection of transmission powers is part of the problem.

The results of [9] were extended to the multihop case by Chafekar et al. in [5]. The papers [15], [16] showed that we can schedule all connections in a number of rounds that is only an  $O(\log^2 n)$  factor than a lower bound based on an intrinsic measure of interference at the receivers. In [3], Borbash and Ephremides present a linear programming formulation of the problem (but which may not always have polynomial size).

### E. Remarks

We remark that throughout this work we focus on a situation where we are simply trying to maximize the capacity of a set of connections transmitting at a given time. We do not explicitly address scheduling issues such as timesharing between different sets of connections.

However, we note that in many scheduling algorithms, the scheduling decision involves finding a set of feasible transmissions that maximizes some notion of weighted capacity. For example, the scheme of Kumar et al. [13] for realising multicommodity flow solutions in multihop wireless networks involves finding maximum feasible sets of transmissions. In addition the well-known backpressure algorithm for stabilizing queues lengths in wireless networks whenever possible (see [21]) involves finding at each time-step a feasible transmission set of maximum weight where the weight of a transmission is derived from the difference between a queue length at the transmitting node and a corresponding queue length at the receiver. In Section V-B we mention how our techniques may be extended to weighted problems. Hence our analysis provides algorithms for solving these subproblems while at the same time showing that finding exact solutions to the subproblems is NP-hard.

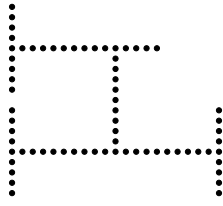


Fig. 1. The graph that forms the basis of our NP-hardness reduction.

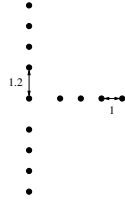


Fig. 2. The gadget.

## II. NP-HARDNESS

In this section we show that the MAX-CONNECTIONS problem in arbitrary networks under the SINR model is NP-hard. Our reduction follows the basic strategy of the NP-hardness reduction for MIS in unit disk graphs. However, the reduction is somewhat more complicated since we have to deal with the fact that interference comes from arbitrary distances. The reduction starts from the NP-hardness of MIS in planar cubic graphs. Specifically it is known (see e.g. [4]) that MIS is NP-hard in graphs where all nodes are on the edges of a grid with squares of size  $M$ , edges are of size 1, all nodes have degree at most 3 and each degree 3 node is incident to linear arrays of size at least  $M/4$  (see Figure 1). Note that any maximum independent set will include at most every other node along an edge of the grid. The proof becomes somewhat complex since we need to show that all power levels will lead to an infeasible solution for any non-independent set.

### A. Gadget

We now describe a gadget that will be used in the eventual hardness proof. The purpose of the gadget is to represent a degree-3 node in our grid. We consider three linear arrays of nodes. (See Figure 2.) Each node serves as both the transmitter and receiver for a single connection. The first linear array is at positions  $(0, 1.2), (0, 2.2), (0, 3.2), \dots$ . The second linear array is at positions  $(1.2, 0), (2.2, 0), (3.2, 0), \dots$ . The third linear array is at positions  $(0, -1.2), (0, -2.2), (0, -3.2), \dots$ . Lastly we have a single node at  $(0, 0)$ . We let the path-loss exponent  $\alpha = 2.05$ , the signal-to-noise ratio threshold  $\tau = 1.00001$  and the maximum power  $p_{\max} = 1$ . We also suppose that each linear array has at least  $\ell$  nodes for some parameter  $\ell$ . The first result about this gadget follows directly from the chosen value of  $\tau$ .

**Lemma 1:** There is no feasible solution that contains adjacent nodes from one of the linear arrays.

*Proof:* Consider two adjacent nodes from a linear array. The distance between them equals 1. Consider the transmission with the smallest power. The SINR for that transmission will be at most 1. Hence the SNR constraint is not satisfied. ■

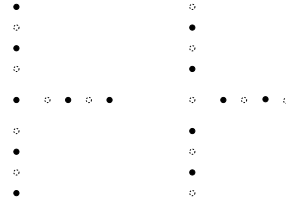


Fig. 3. The feasible configurations.

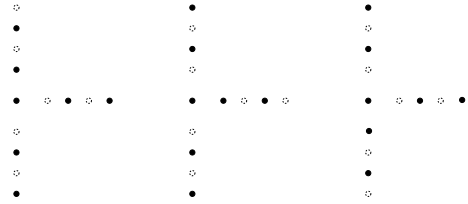


Fig. 4. The infeasible configurations.

Hence it remains to see what configurations are feasible that only use alternating members of a linear array. The following facts can be verified numerically.

**Lemma 2:** The following configurations are feasible for arbitrarily large  $\ell$ , even when there is a background noise level of  $\varepsilon = 0.01$ .

- $(0, 0), (0, 2.2), (0, 4.2), \dots, (2.2, 0), (4.2, 0), \dots, (0, -2.2), (0, -4.2), \dots$  (See Figure 3 (left).)
- $(0, 1.2), (0, 3.2), \dots, (1.2, 0), (3.2, 0), \dots, (0, -1.2), (0, -3.2), \dots$  (See Figure 3 (right).)

For sufficiently large  $\ell$  the following configurations are not feasible, even if there is no background noise level.

- $(0, 0), (0, 1.2), (0, 3.2), \dots, (2.2, 0), (4.2, 0), \dots, (0, -2.2), (0, -4.2), \dots$  (See Figure 4 (left).)
- $(0, 0), (0, 2.2), (0, 4.2), \dots, (1.2, 0), (3.2, 0), \dots, (0, -2.2), (0, -4.2), \dots$  (See Figure 4 (middle).)
- $(0, 0), (0, 2.2), (0, 4.2), \dots, (2.2, 0), (4.2, 0), \dots, (0, -1.2), (0, -3.2), \dots$  (See Figure 4 (right).)

It is easy to see that by making  $M$  sufficiently large we can guarantee that for any node  $a$  the interference caused to  $a$  by nodes at distance at least  $M/4$  from  $a$  is at most  $\varepsilon$ . Note that  $M$  will depend only on  $\varepsilon$ . It is also easy to see from Lemma 2 that in a single linear array it is feasible for every other node to transmit at  $p_{\max} = 1$  even with background noise of 0.01.

We can use the above gadget to show NP-hardness in the following manner. First we can make sure that in the grid example where MIS is hard every node on the corners of the grid have degree 3 and every other node has degree 1 or 2. (See Figure 1.) We then place a copy of the gadget around every degree-3 node so that the linear arrays correspond to degree 1 or 2 nodes.

For the first direction of the reduction we would like to show that for any MIS in the original graph, the corresponding nodes can transmit in our wireless instance. This is easy to see by using Lemma 2, since close to the center of each gadget we know that the interference from outside the gadget is at most  $\varepsilon$ , so it is still feasible. The only non-obvious case is when two gadgets meet at the center of a chain, but this is clearly

still feasible since at the center of the chain everything within distance  $M/4$  is just part of the chain, so is still feasible by broadcasting at power 1 (which is consistent with the feasible gadget solution).

Now we need to show that any maximum feasible solution forms an independent set in the original graph. An important observation is that we can without loss of generality assume that in any maximum feasible solution every other node in a linear array is transmitting. If not, then we could always add to the number of nodes transmitting in the linear array by turning off one of the degree 3 nodes. We can repeat this process until every linear array has half its nodes transmitting. Lemma 2 then implies that we cannot have a degree 3 node transmitting together with one of its neighbors, and Lemma 1 implies that no other adjacent nodes are transmitting. This any maximum feasible set also forms an independent set, completing the reduction.

### III. APPROXIMATION ALGORITHMS

Due to the NP-hardness of our problem we now turn our attention to approximation algorithms. Ideally we would like to adapt one of the polynomial time approximation schemes (that give a  $(1 + \varepsilon)$ -approximation for any  $\varepsilon$ ) to the SINR model. However, we are unable to do that, mainly because the analyses of these algorithms make critical use of the fact that two transmissions only interfere if the transmitters are close to each other. However, in the SINR model interference can occur at arbitrary distances which makes it difficult to directly adapt these analyses. However, in this section we show that if  $d_{\max}$  is constant then we can obtain constant approximation algorithms in polynomial time. More generally, we present an  $O(\log d_{\max})$ -approximation that runs in polynomial time, and for the case in which the background noise  $W = 0$  we give an  $O(1)$ -approximation that runs in time  $O(n^{d_{\max}^2})$ .

Before we present these algorithms we start with a *density lemma* that we shall use both for these results and for our game-theoretic results in Section IV. This lemma states that any feasible solution can only have a limited number of receivers in any fixed area.

**Lemma 3:** 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  $3^\alpha/\tau$ .

*Proof:* Without loss of generality we assume that the background noise is 0. Having a non-zero background noise can only reduce the number of connections that can be supported.

Suppose that all nodes in the feasible solution transmit at a power such that the received signal 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  $p_i \min\{1, (d_0/d(t_i, r_{i'}))^\alpha\} \geq p_i \min\{1, (d_0/(d(r_i, r_{i'}) + d(t_i, r_i)))^\alpha\}$ . Recall that we assume that either  $r_i = t_i$  or  $d(t_i, r_i) \geq d_0$ . In addition, by the geometry of the square  $S$  we know that  $d(r_i, r_{i'}) \leq 2d_0$ . This implies that  $p_i \min\{1, (d_0/(d(r_i, r_{i'}) + d(t_i, r_i)))^\alpha\} \geq \frac{1}{3^\alpha} p_i \min\{1, (d_0/d(t_i, r_i))^\alpha\} \geq \bar{p}/3^\alpha$ . Hence if there are more than  $3^\alpha/\tau$  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. 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 4:** 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  $3^\alpha d^2/\tau(d_0)^2$ .

*Proof:* Divide square  $S$  up into subsquares of size  $d_0$  and then apply Lemma 3. ■

**Lemma 5:** Now consider a ball  $B$  of radius  $d$ . In any feasible solution the maximum number of connections with a receiver in ball  $B$  is  $3^\alpha \cdot 4d^2/\tau(d_0)^2$ .

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

The following extension of Lemma 3 will also be useful.

**Lemma 6:** Consider a square  $S$  with side-length  $d$ . In any feasible solution the maximum number of connections such that  $d(t_i, r_i) \geq d$  and  $r_i$  is in square  $S$  is  $3^\alpha/\tau$ .

*Proof:* The analysis is almost identical to that of Lemma 3 once we note that in this case  $d(r_i, r_{i'}) \leq 2d \leq 2d(t_i, r_i)$  for all  $i, i'$ . ■

In the next two theorems we present our approximation algorithms for the MAX-CONNECTIONS problem in the SINR model.

**Theorem 7:** There exists a polynomial time algorithm that always finds a solution to MAX-CONNECTIONS that is within a factor  $O(\log d_{\max})$  of optimal.

*Proof:* We divide all connections into classes based on distance. Class  $F_j$  contains all connections  $i$  such that  $d_{\max}/2^{j-1} \geq d(t_i, r_i) \geq d_{\max}/2^j$ . (We remark that a similar decomposition was used in [9].) Note that in the optimal solution there must exist a  $j$  such that  $F_j$  contains  $\text{OPT}/\log d_{\max}$  connections. In the following we will consider each  $j$  in turn and obtain a constant approximation for the connections in  $F_j$  only. We focus on a  $j$  for which  $d_{\max}/2^j \geq d_{\min}$ . The connections for which  $d_{t_i, r_i} = 0$  can be handled similarly.

We now divide the problem into squares of side  $d_{\max}/2^j$ . (See Figure 5.) We refer to these squares as *j-squares*. From each *j-square*  $S$ , if there is at least one receiver in  $S$  then we choose one arbitrarily, and restrict ourselves to the problem on these connections. Note that Lemma 6 implies that each *j-square* only contains at most  $3^\alpha/\tau$  receivers from the optimal solution on  $F_j$ , so as long as we can support at least a constant fraction of our chosen connections we are still within a constant of the optimal solution on  $F_j$ .

We now restrict our attention to 1 out of every  $k^2$  *j-squares* in an evenly spaced pattern for some parameter  $k$ , i.e. squares located at the same coordinates mod  $k$ . (See Figure 5). We can partition the plane into  $k^2$  such sets of squares. We show that in each set we can support one connection in each square, so by taking the best set we are only losing another  $k^2$  factor.

Consider some *j-square*  $S$ , and consider the set  $I$  of *j-squares* in the same pattern set that are offset from  $S$  by exactly

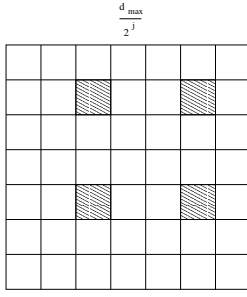


Fig. 5. We form our solution using 1 out of every  $k^2$  squares. Here  $k = 3$ .

$ik$  in one coordinate and at most  $ik$  in the other coordinate (i.e. the set of  $j$ -squares that are on the border of the  $\ell_\infty$  ball of radius  $ik$  around  $S$ ). Since the  $\ell_\infty$  distance is at most the normal  $\ell_2$  distance, it is not hard to verify that the maximum interference caused by the connection in  $I$  to the connection in  $S$  is at most  $p_{\max}(d_0 2^j / (ik - 3)d_{\max})^\alpha \leq p_{\max}(d_0 2^j / (i(k - 3))d_{\max})^\alpha$ . It is also easy to see that there are at most  $8i$  squares in  $I$ . This implies that the total interference suffered by the connection in  $S$  is at most  $\sum_{i=1}^{\infty} 8ip_{\max} \left( \frac{d_0 2^j}{(i(k-3))d_{\max}} \right)^\alpha = 8p_{\max} \left( \frac{d_0 2^j}{(k-3)d_{\max}} \right)^\alpha \zeta(\alpha - 1)$ , where  $\zeta(\alpha - 1)$  is the Riemann zeta function, which is constant for constant  $\alpha > 2$ .

The connection in  $S$  can use power  $p_{\max}$ , so if  $W = 0$  then as long as  $(p_{\max} / (d_{\max} / 2^{j-1})) / (8p_{\max} \left( \frac{d_0 2^j}{(k-3)d_{\max}} \right)^\alpha \zeta(\alpha - 1))$  is at least  $\tau$  the connection in  $S$  can be supported (and thus by the same argument so can all of the rest of the connections in  $I$ ). So it suffices to choose  $k$  such that  $(k - 3)^\alpha \geq 8d_0 2^\alpha \zeta(\alpha - 1)$ , and thus  $k$  is some constant. If  $W \neq 0$  then we have to increase  $k$  by a constant factor depending only on  $\epsilon$  (recall that  $\epsilon$  is a measure of how far  $d_{\max}$  is from the physical limit). The approximation factor that we lose for class  $F_j$  due to all the connections that have been removed is  $\frac{3^\alpha}{\tau} k^2$  which is a constant for fixed  $\alpha$  and  $\tau$ . As already mentioned, our overall approximation ratio is therefore  $O(\log d_{\max})$ . ■

If there is no background noise we can obtain another algorithm whose approximation ratio has a better dependence on  $d_{\max}$  at the expense of a worse dependence in the running time. Due to space constraints we omit the proof, and note only that it is similar to the proof of Theorem 7 except for using squares of size  $d_{\max}$ , scaling so total power from a square is  $p_{\max}$ , and dropping half of the connections in each square.

**Theorem 8:** For the case with no background noise (i.e.  $W = 0$ ) we can find an  $O(1)$  approximate solution in time  $n^{O(d_{\max}^2)}$ .

#### IV. GAME THEORY

As discussed, the approximation algorithms described previously are centralized. We would also like to examine highly distributed algorithms that allow each transmitter to make its own decision based on limited local information. One extreme version of this is the setting in which transmitters are not allowed to exchange any information between themselves, and instead must make a decision on broadcast power based only on knowledge of the signal and noise at their receivers (we

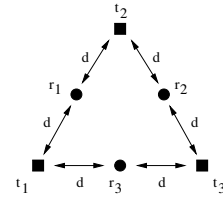


Fig. 6. No pure Nash exists

assume that receivers periodically provide this information to their transmitters). A natural way of viewing this setting is as a game where the transmitters are the players and the pure strategies are power settings. In this section we will define such a game and show that every Nash equilibrium in this game results in an expected number of successful transmissions that is close to optimal if there is no background noise.

For simplicity of notation we will without loss of generality rescale powers and  $W$  so that  $p_{\max} = 1$ . The game that the transmitters will be playing is simple. Each transmitter is a player, whose pure strategies are the reals in  $[0, 1]$ , with a nonzero value representing broadcasting at that power and 0 representing not broadcasting. So a mixed strategy is a probability distribution over  $[0, 1]$ . A transmitter gets payoff 0 if it does not broadcast (i.e. has power 0), payoff 1 if it broadcasts and its receiver has signal to noise ratio at least  $\tau$ , and  $-1$  if it broadcasts but its receiver has SINR less than  $\tau$ . We note that it is easy to see that the same game without the  $-1$  penalty can have bad Nash equilibria (in particular, everyone broadcasting). We first discuss pure Nash equilibria in this game, and then examine the more general mixed Nash case.

##### A. Pure Nash Equilibria

A pure Nash equilibrium is a very natural solution concept, since it would guarantee that everyone broadcasting is doing so successfully while no one not broadcasting could succeed even if they went at maximum power. Unfortunately a simple example shows that pure Nash equilibria do not always exist in our game.

The bad example is as follows, and is given in Figure 6. There are three transmitters  $t_1, t_2, t_3$  in an equilateral triangle with side length  $2d$  for some arbitrary  $d \gg d_0$ . The receiver for  $t_1$  (i.e.  $r_1$ ) is located halfway between  $t_1$  and  $t_2$  (so at distance  $d$  from each). Similarly,  $r_2$  is located halfway between  $t_2$  and  $t_3$  and  $r_3$  is located halfway between  $t_3$  and  $t_1$ . We will set  $\tau = 2$  and  $\alpha = 2.5$ , and will assume no background noise so  $W = 0$ .

We first claim that  $|OPT| = 1$ . To see this, suppose that there are at least two successful broadcasts. Without loss of generality we will assume that connections 1 and 2 are successful, and are broadcasting at powers  $p_1$  and  $p_2$  respectively. Then since connection 1 is successful we know that  $(p_1/d^\alpha)/(p_2/d^\alpha) \geq 2$ , and thus that  $p_1 \geq 2p_2$ . Now we note by simple geometry that the distance from  $t_1$  to  $r_2$  is exactly  $d\sqrt{3}$ , so the SINR of connection 2 at  $r_2$  is  $(p_2/d^\alpha)/(p_1/(d\sqrt{3})^\alpha) = (p_2 3^{1.25})/p_1 \leq (p_2 3^{1.25})/(2p_2) = 3^{1.25}/2 < 2 = \tau$ , which is a contradiction since we assumed that  $b$  was successful.

Now since  $|OPT| = 1$ , any pure Nash equilibrium must have exactly one successful transmission (since obviously it must have more than 0 and at most  $|OPT|$ ). Without loss of generality we will assume that connection 1 is successful with power  $p_1$ , so by the definition of a pure Nash it must be the case that neither connection 2 nor 3 are broadcasting with power greater than 0. But then the interference at  $r_2$  is just  $p_1/(d\sqrt{3})^\alpha$ , so if  $t_2$  broadcasted at power  $p_1$  then the SINR at  $r_2$  would be  $(p_1/d^\alpha)/(p_1/(d\sqrt{3})^\alpha) = (\sqrt{3})^\alpha > 3 > \tau$ . So  $t_2$  would be successful if it transmitted at power at least  $p_1$ , and thus  $a$  broadcasting by itself is not a pure Nash equilibrium.

While pure Nash equilibria do not always exist, when they do exist they have value close to OPT. This is formalized in the next subsection when we prove the same statement about mixed Nash equilibria, but we will provide the intuition for the pure Nash special case. Fix some pure Nash. We will try to find a receiver whose associated transmitter is not broadcasting in the Nash but has “small” interference, where our notion of small is something that increases as the value of the Nash gets closer to OPT. Since this receiver’s transmitter is not broadcasting, the interference must be overcoming any possible signal and thus must actually be quite large, implying that the Nash must actually have value close to OPT.

### B. Mixed Nash Equilibria

While pure Nash equilibria do not always exist, obviously a mixed Nash does. We now show that any mixed Nash (and thus any pure Nash, if one does exist) has value close to OPT. Recall that a mixed strategy is a probability distribution over the possible powers (i.e. over  $[0, 1]$ ), and in a mixed Nash there is no incentive for any transmitter to change its distribution. For our purposes, for a transmitter  $i$  we will only need to consider the probability  $q_i$  that  $i$  broadcasts with non-zero power. We begin with a few useful lemmas. Fix some Nash equilibrium. For each transmitter  $i$ , let  $p_{good}(i)$  be the probability (over the randomness in the strategies of the other transmitters) that  $i$  would be successful if it were to broadcast at power 1. Let  $p_{bad}(i) = 1 - p_{good}(i)$  be the probability that  $i$  would not be successful. Note that if  $i$  has non-zero probability of broadcasting at some power greater than 0 but less than 1 then the probability of it succeeding at that power must be equal to the probability of it succeeding at power 1, since otherwise it could just switch to power 1 and strictly increase its expected payoff. So  $S = \sum_i q_i p_{good}(i)$  is the expected number of successful transmissions (i.e. the value of the equilibrium). Let  $T = \sum_i q_i$  be the expected number of transmissions.

**Lemma 9:** For any Nash equilibrium, for any transmitter  $i$ , if  $q_i < 1$  then  $p_{bad}(i) \geq 1/2$  and if  $q_i > 0$  then  $p_{bad}(i) \leq 1/2$ .

*Proof:* Suppose that  $q_i < 1$  and that  $p_{bad}(i) < 1/2$ , so  $p_{good}(i) > 1/2$ . Then by broadcasting at power 1 with probability  $q_i$ , the expected payoff to  $i$  would be  $q_i(p_{good}(i) - (1 - p_{good}(i))) = q_i(2p_{good}(i) - 1)$ . Since  $2p_{good}(i) - 1 > 0$ , this is maximized by setting  $q_i = 1$ , contradicting our choice of  $i$  and our assumption that this is an equilibrium.

Similarly, suppose that  $q_i > 0$  and that  $p_{bad}(i) > 1/2$ . Then when  $i$  broadcasts it will fail more than  $1/2$  the time, giving negative expected payoff, so  $i$  would just never broadcast (i.e. set  $q_i$  to 0), contradicting our choice of  $i$ . ■

**Lemma 10:** For any Nash equilibrium,  $S \leq T \leq 2S$ .

*Proof:* The first inequality is obvious from the definitions, and the second immediately follows from the second part of Lemma 9, since  $T = \sum_i q_i = 2 \sum_i \frac{1}{2} q_i \leq 2 \sum_i p_{good}(i) q_i = 2S$ . ■

Let  $OPT$  be the set of receivers that achieve their SINR requirement in the optimal solution. We can now prove the main theorem of this section:

**Theorem 11:** Any Nash equilibrium has an expected number of successful transmissions at least  $\Omega(|OPT|/d_{\max}^{2\alpha})$ , where we assume that  $\alpha$  and  $\tau$  are constants.

*Proof:* Fix a Nash equilibrium. Let  $L = \{i : q_i = 1\}$  be the set of transmitters that broadcast at power greater than 0 with probability 1. Consider the following procedure (only for analysis, obviously). For each receiver  $x$  in  $OPT \setminus L$  we will keep track of how much it is “bought” with a variable  $b(x)$ , initially all set to 0. Now we order all transmitters in the instance (or just all transmitters with non-zero  $q_i$  in the Nash) arbitrarily. We examine the transmitters one by one in this order. Say we are on transmitter  $i$ . Let  $R(i)$  be the  $\lfloor \frac{|OPT \setminus L| - k}{kT} \rfloor$  closest receivers in  $OPT \setminus L$  to  $i$  (for some parameter  $k$  to be defined later) 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 increases the sum of the  $b$  values by  $q_i \lfloor \frac{|OPT \setminus L| - k}{kT} \rfloor$ , at the end of this process we know that  $\sum_x b(x) = \sum_i q_i \lfloor \frac{|OPT \setminus L| - k}{kT} \rfloor \leq \frac{|OPT \setminus L| - k}{k} < \frac{|OPT|}{k}$ , since by definition  $T = \sum_i q_i$ . This means that there is some receiver  $a \in OPT \setminus L$  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 probability mass from  $M$  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  $q_x \lfloor \frac{|OPT \setminus L| - k}{kT} \rfloor$  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  $z(d) \lfloor \frac{|OPT \setminus L| - k}{kT} \rfloor$ . 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  $\frac{z(d)}{2} \lfloor \frac{|OPT \setminus L| - k}{kT} \rfloor$ . By Lemma 5, this implies that  $cd^2 \geq \frac{z(d)}{2} \lfloor \frac{|OPT \setminus L| - k}{kT} \rfloor$  and thus that  $z(d) \leq 2cd^2 / \lfloor \frac{|OPT \setminus L| - k}{kT} \rfloor$  for some constant  $c$  depending only on  $\alpha, \tau$ , and  $d_0$ .

Now that we have a bound on the probability mass inside a ball around  $a$ , we want to bound the probability mass in an annulus of thickness 1 around  $a$ . To do this, we note that the interference at  $a$  is maximized if every ball around  $a$  actually meets the above bound. Since in the end we will care about upper bounding the interference, we can say without loss of

generality that every ball meets the above bound, implying that the sum of the probabilities of transmitters between distance  $d$  and  $d + 1$  is at most  $(2c/\lfloor \frac{|OPT \setminus L| - k}{kT} \rfloor)((d + 1)^2 - d^2) \leq 6cd/\lfloor \frac{|OPT \setminus L| - k}{kT} \rfloor$  when  $d \geq 1$ , and is at most  $2c/\lfloor \frac{|OPT \setminus L| - k}{kT} \rfloor$  when  $d = 0$ . Since the expected interference from a transmitter at distance  $d$  from  $a$  is at most its probability of broadcasting times  $1/d^\alpha$ , this means that the expected interference at  $a$  caused by transmitters at distance between  $d$  and  $d + 1$  from  $a$  is at most  $(6c/\lfloor \frac{|OPT \setminus L| - k}{kT} \rfloor) \cdot \frac{1}{d^{\alpha-1}}$  for  $d \geq 1$ . For  $d = 0$ , since the interference caused by a transmitter is at most 1, the expected interference from transmitters between distances 0 and 1 from  $a$  is at most  $2c/\lfloor \frac{|OPT \setminus L| - k}{kT} \rfloor$ . Using linearity of expectations, we can sum over the annuli to get that the expected interference at  $a$  is at most  $2c/\lfloor \frac{|OPT \setminus L| - k}{kT} \rfloor + (6c/\lfloor \frac{|OPT \setminus L| - k}{kT} \rfloor) \sum_{d=1}^{\infty} \frac{1}{d^{\alpha-1}} \leq 8c\zeta(\alpha - 1)/\lfloor \frac{|OPT \setminus L| - k}{kT} \rfloor$ , where  $\zeta(\alpha - 1)$  is the Riemann zeta function (which will be constant for  $\alpha > 2$ ).

This gives us a bound on the expected interference at  $a$  caused by transmitters in  $M$ . What about the transmitters in  $M'$ ? Since we know that  $\sum_{x \in M'} q_x \leq \frac{1}{k}$ , we get that they cause at most  $\frac{1}{k}$  expected interference (which is what would happen if they were all at distance 1 from  $a$ ). Thus the total expected interference is at most  $(8c\zeta(\alpha - 1)/\lfloor \frac{|OPT \setminus L| - k}{kT} \rfloor) + \frac{1}{k}$ .

So now we have a bound on the expected interference. Let us assume (for now) that  $W = 0$ . By using Markov's inequality, we get that the probability that  $a$  hears interference at least twice the expected interference is at most  $1/2$ . But since we know from how we selected  $a$  that the probability that its transmitter tries to transmit is less than 1, Lemma 9 implies that  $p_{bad}(a) \geq 1/2$ . Thus  $(16c\zeta(\alpha - 1)/\lfloor \frac{|OPT \setminus L| - k}{kT} \rfloor) + \frac{2}{k}$  must be enough interference to kill the transmission to  $a$ ; in particular, it must be the case that  $16c\zeta(\alpha - 1)/\lfloor \frac{|OPT \setminus L| - k}{kT} \rfloor \geq \frac{1}{\tau d_{\max}^\alpha} - \frac{2}{k}$ . We will now finally set  $k$ , to  $4\tau d_{\max}^\alpha$ , giving us that  $16c\zeta(\alpha - 1)/\lfloor \frac{|OPT \setminus L| - 4\tau d_{\max}^\alpha}{4\tau d_{\max}^\alpha T} \rfloor \geq \frac{1}{2\tau d_{\max}^\alpha}$ .

Solving for  $T$  in this equation, and assuming constant  $\alpha$  and  $\tau$ , implies that  $T \geq \Omega(|OPT \setminus L|/d_{\max}^{2\alpha})$ , and thus by Lemma 10 we have that  $S \geq \Omega(|OPT \setminus L|/d_{\max}^{2\alpha})$ . If  $|OPT \setminus L| = o(|OPT|)$  then a superconstant fraction of transmitters are broadcasting with probability 1 in the Nash, which by Lemma 9 and Lemma 10 means that the expected number of successful transmissions in the Nash is at least  $\Omega(|OPT|)$ , which would prove the theorem. On the other hand, if  $|OPT \setminus L| = \Omega(|OPT|)$  then the above equation implies that  $S \geq \Omega(|OPT|/d_{\max}^{2\alpha})$ , thus proving the theorem.

If  $W \neq 0$  the theorem is still true but the details are slightly more complicated, so we give only a brief sketch. With background noise, instead of twice the expected interference being enough to kill the signal it must be that twice the expected interference plus the background noise must be enough to kill the signal. But this only causes us to lose another constant, since we assumed from the beginning that the distance from any receiver to its transmitter (and thus from  $a$  to its transmitter) is bounded away from the absolute limit by a constant. ■

## V. EXTENDED MODELS

We now give brief sketches of how our results can be adapted to the extended models discussed in Section I-C.

### A. Multiple Carriers

Recall that in the multiple carrier model there are  $J$  different carriers  $\{1, 2, \dots, J\}$  and transmissions on carrier  $j$  cause interference only to other transmissions using carrier  $j$ . Obviously this version is still NP-hard, as it contains the single-carrier version as a subproblem.

For the approximation algorithms, it is clear that the only thing we need to modify is how we compute the optimal solution in each square, and then the rest of the proof will go through as before. We need to modify the square optimum computation because if  $J$  is large trying all possibilities might take too long. However, if we regard the connections as a ground set and each feasible single carrier solution as a subset of this ground set, it is obvious that finding the best solution is just an instance of the famous max  $k$ -cover problem for  $k = J$ . It is well known (see e.g. Feige [8]) that a simple greedy algorithm is a  $(1 - 1/e)$ -approximation, so we can simply use that algorithm to compute the optimal solution in each square and lose another constant  $(e/(e - 1))$ .

For the game theoretic problem, modulo some technical details it is not difficult to see that the same results will hold. This is because we can repeat the analysis of Theorem 11 for each individual channel, proving that in each channel the Nash must do nearly as well as OPT. We only end up losing an extra constant factor due to technical details related to allowing mixed strategies that randomize over different channels.

### B. Weighted Model

Recall that in the weighted model each connection has a weight, and the goal is to maximize the sum of the weights of the supported connections rather than just the number of such connections. In this model it is easy to see that we can modify the algorithm of Theorem 7 to achieve a similar  $O(\log d_{\max})$ -approximation. In particular, since there are at most  $3^\alpha/\tau$  connections in a square we can lose another  $3^\alpha/\tau$  and simply pick the single connection in each square with the largest weight. Since we guaranteed that a single connection in a square is always feasible, this solution is feasible and is still an  $O(\log d_{\max})$ -approximation. On the other hand, it is trivial to see that in the weighted model a Nash equilibrium can be arbitrarily far from optimal, e.g. if  $\tau$  is greater than 1 and there are two co-located transmitters of which one has weight 1 and the other has arbitrarily large weight.

## VI. SIMULATIONS

In light of our proof that Nash equilibria are reasonably good approximations to OPT, one natural question is whether normal best response dynamics will converge to a Nash equilibrium. It is easy to see that in the worst case the answer is no, but this naturally leads to the question of what the best response dynamics do on random instances.

To test this we ran simulations with various parameters. Our random instances were created as follows: first, each transmitter was given a random location in a 100 by 100 square by drawing two coordinates uniformly from  $[0, 100]$ . For each transmitter we drew a distance from  $[1, d_{\max}]$  uniformly at random, as well as a random direction by drawing two coordinates independently from the normal Gaussian distribution. The receiver was

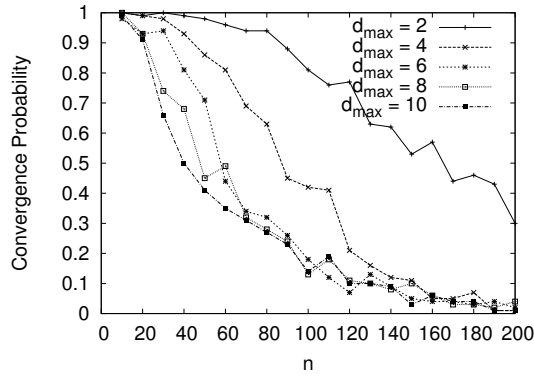


Fig. 7. Convergence probability vs.  $n$  for various values of  $d_{\max}$

then placed in this direction from the transmitter at the random distance. We assumed no background noise and set  $p_{\max} = 1$ . We started out with random powers for each connection, and then iterated the best response either 50 times or until it converged, whichever occurred first. We only allowed powers 0 and 1, and did not allow mixed strategies, so any convergence was to a Nash equilibrium in which some connections were using full power and all others were using none.

For different settings of  $\alpha$ ,  $\tau$ ,  $d_{\max}$ , and  $n$  we ran 100 of the above simulations. Varying  $\alpha$  and  $\tau$  within reasonable ranges did not have a large effect – instead, the dominant trends were the effects of  $d_{\max}$  and  $n$  on the probability of convergence. See Figure 7, which uses  $\alpha = 2.3$  and  $\tau = 0.5$ . This makes intuitive sense, as when  $d_{\max}$  is small and/or  $n$  is small each receiver is likely to be close to its transmitter, so there is less of a chance of a cycle similar to the example from Section IV-A.

## VII. CONCLUSIONS

In this paper we have examined the complexity of maximizing the number of supported connections in a wireless SINR model and we have studied the performance that can be achieved by completely distributed algorithms operating under an appropriate incentive structure. A number of open problems remain. First we would like to know if it possible to remove the dependence on  $d_{\max}$  from the performance guarantees and running times of our algorithms.

Second, throughout this work we have assumed that each connection must satisfy a strict SINR constraint. However, for best-effort types of flows we might be willing to tolerate different SINR values and transmit at a smaller rate. In this case we would have a utility  $u_i$  for each connection  $i$  and the goal is to maximize the sum of the utilities across all connections. If we denote the SINR for connection  $i$  by  $\text{SINR}_i$ , then a natural definition of utility is  $\log \log(1 + \text{SINR}_i)$ . We believe that an extremely interesting open problem is to resolve the complexity of the resulting utility maximization problem. It is known to be non-convex in general and so standard optimization algorithms do not apply. However, as far as we are aware it is also not known to be NP-hard.<sup>3</sup> We would also like to resolve to price

of anarchy due to distributed approaches. Huang et al. [11] showed that if nodes can exchange pricing information then we can obtain distributed algorithms that converge to local optima. However, we are unaware of any work that bounds the difference between local and global optima in this context.

## REFERENCES

- [1] S. Arora, L. Babai, J. Stern, and Z. Sweedyk. The hardness of approximate optima in lattices, codes and systems of linear equations. In *Proceedings of the 34th Annual Symposium on Foundations of Computer Science*, pages 724–733, 1993.
- [2] P. Bahl, M. Hajiaghayi, K. Jain, V. Mirrokni, L. Qiu and A. Saberi. Cell breathing in wireless LANs: Algorithms and evaluation. *IEEE Transactions on Mobile Computing* 6(2):164–178, 2007.
- [3] S. Borbash and A. Ephremides. Wireless link scheduling with power control and SINR constraints. *IEEE Transactions on Information Theory* 52(11): 5106–5111, 2006.
- [4] H. Breu. Algorithmic aspects of constrained unit disk graphs. PhD thesis. University of British Columbia, 1996.
- [5] D. Chafekar, V. Kumar, M. Marathe, S. Parthasarathy and A. Srinivasan. Approximation algorithms for computing capacity of wireless networks with SINR constraints. In *Proceedings of INFOCOM 08*, April 2008.
- [6] M. Chiang, C. Tan, D. Palomar, D. O’Neill, and D. Julian. Power Control by Geometric Programming. *IEEE Trans. on Wireless Communications*, 6(7):2640 – 2651, 2007.
- [7] T. Erlebach, K. Jansen, and E. Seidel. Polynomial-time approximation schemes for geometric graphs. In *Proceedings of the 12th annual ACM-SIAM symposium on Discrete algorithms*, pages 671–679, 2001.
- [8] U. Feige. A threshold of  $\ln n$  for approximating set cover. *Journal of the ACM*, 45(4):634 – 653, 1998.
- [9] O. Goussevskaia, Y. Oswald and R. Wattenhofer. Complexity in geometric SINR. In *Proceedings of MobiHoc 07*, September 2007.
- [10] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388 – 404, 2000.
- [11] J. Huang, R. Berry and M. Honig. Game theoretic analysis of distributed power control for spread spectrum ad hoc networks. In *Proceedings of 2005 IEEE International Symposium on Information Theory*, Adelaide, Australia, September 2005.
- [12] H. Hunt, M. Marathe, V. Radhakrishnan, S. Ravi, D. Rosenkrantz and R. Stearns. NC-approximation schemes for NP- and PSPACE-hard problems for geometric graphs. *Journal of Algorithms*, 26(2):238–274, 1998.
- [13] A. Kumar, M. Marathe, S. Parthasarathy and A. Srinivasan. Algorithmic aspects of capacity in wireless networks. In *SIGMETRICS*, 2005.
- [14] M. Marathe, H. Breu, H. Hunt III, S. Ravi, and D. Rosenkrantz. Simple heuristics for unit disk graphs. *Networks*, 25:59–68, 1995.
- [15] T. Moscibroda, Y. Oswald and R. Wattenhofer. How optimal are wireless scheduling protocols? In *Proceedings of INFOCOM 07*, May 2007.
- [16] T. Moscibroda, R. Wattenhofer and A. Zollinger. Topology control meets SINR: The scheduling complexity of arbitrary topologies. In *Proceedings of MobiHoc 06*, May 2006.
- [17] T. Nieberg, J. Hurink, and W. Kern. A robust PTAS for maximum weight independent sets in unit disk graphs. In *WG 2004*, pages 214–221, 2004.
- [18] T. Rappaport. *Wireless Communications: Principles and Practice*. Wiley, 2001.
- [19] C. Saraydar, N. Mandayam and D. Goodman. Efficient power control via pricing in wireless data networks. *IEEE Transactions on Communications*, 50(2):291–303, 2002.
- [20] A. Stolyar and H. Viswanathan. Self-organizing dynamic fractional frequency reuse in OFDMA systems. In *Infocom*, 2008.
- [21] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, 37(12):1936–1949, 1992.
- [22] R. Yates. A framework for uplink power control in cellular radio systems. *IEEE Journal on Selected Areas in Communications*, 13(7):1341–1347, 1995.
- [23] S. Yi, Y. Pei, and S. Kalyanaraman. On the capacity improvement of ad hoc wireless networks using directional antennas. In *MobiHoc*, pages 108–116, 2003.

<sup>3</sup>However, we remark that if the connection utility is given instead by  $\log(\text{SINR}_i)$  then Chiang et al. [6] showed that the resulting problem is feasible via Geometric Programming.