

Wireless Communication Is in APX

Magnús M. Halldórsson^{1,*} and Roger Wattenhofer²

¹ School of Computer Science, Reykjavik University, 103 Reykjavik, Iceland
`mmh@ru.is`

² Computer Engineering and Networks Laboratory, ETH Zurich, Switzerland
`wattenhofer@tik.ee.ethz.ch`

Abstract. In this paper we address a common question in wireless communication: How long does it take to satisfy an arbitrary set of wireless communication requests? This problem is known as the wireless scheduling problem. Our main result proves that wireless scheduling is in APX. In addition we present a robustness result, showing that constant parameter and model changes will modify the result only by a constant.

1 Introduction

Despite the omnipresence of wireless networks, surprisingly little is known about their algorithmic complexity and efficiency: Designing and tuning a wireless network is a matter of experience, regardless whether it is a Wireless LAN in an office building, a GSM phone network, or a sensor network on a volcano.

We are interested in the fundamental communication limits of wireless networks. In particular, we would like to know what communication throughput can possibly be achieved. This question essentially boils down to spatial reuse, i.e., which devices can transmit concurrently, without interfering. More precisely, formulated as an optimization problem: Given a set of communication requests, how much time does it take to schedule them?

Evidently the answer to this question depends on the wireless transmission model. In the past, algorithmic research has focused on graph-based models, also known as protocol models. Unfortunately, graph-based models are too simplistic. Consider for instance a case of three wireless transmissions, every two of which can be scheduled concurrently without a conflict. In a graph-based model one will conclude that all three transmissions may be scheduled concurrently as well, while in reality this might not be the case since wireless signals sum up. Instead, it may be that two transmissions *together* generate too much interference, hindering the third receiver from correctly receiving the signal of its sender. This *many-to-many* relationship makes understanding wireless transmissions difficult – a model where interference sums up seems paramount to truly comprehending wireless communication. Similarly, a graph-based model oversimplifies wireless attenuation. In graph-based models the signal is “binary”, as if there was an invisible wall at which the signal immediately drops. Not surprisingly, in reality the signal decreases gracefully with distance.

* Work done while visiting Research Institute for Mathematical Sciences (RIMS) at Kyoto University.

In contrast to the algorithmic (“CS”) community which focuses on graph-based models, researchers in information, communication, or network theory (“EE”) are working with wireless models that sum up interference and respect attenuation. The standard model is the signal-to-interference-plus-noise (SINR) model – we will formally introduce it in Section 3. The SINR model is reflecting the physical reality more precisely, it is therefore often simply called the physical model. On the other hand, “EE researchers” are not really looking for algorithmic results. Instead, they usually propose heuristics that are evaluated by simulation. Analytical work is only done for special cases, e.g. the network has a grid structure, or traffic is random. However, these special cases do neither give insights into the complexity of the problem, nor do they give algorithmic results that may ultimately lead to new protocols. Since the SINR model is somewhere between graph-theory and geometry, we believe that it will be interesting for the algorithms community, as a new set of tools will be necessary.

The specific question we are addressing in this paper is a classic question in wireless communication: How long does it take to satisfy an arbitrary set of wireless communication requests? This problem is known as the wireless scheduling problem. It is at the heart of wireless communication. Our solution is hopefully pleasing to the EE community as it is using their models, and it is hopefully pleasing to the CS community because we make no restrictions on the input. Our main result proves that wireless scheduling is in APX.

2 Related Work

Most work in wireless scheduling in the physical (SINR) model is of heuristic nature, e.g. [3,7]. Only after the work of Gupta and Kumar [14], analytical results became *en vogue*. The analytical results however are restricted to networks with a well-behaving topology and traffic pattern. On the one hand this restriction keeps the math involved tractable, on the other hand, it allows for presenting the results in a concise form, i.e., “the throughput capacity of a wireless network with X and Y is Z ”, where X and Y are some parameters defining the network, and Z is a function of the network size. This area of research has been exceptionally popular, with a multi-dimensional parameter space (e.g. node distribution, traffic pattern, transport layer, mobility), and consequently literally thousands of publications. The intrinsic problem with this line of research is that real networks often do not resemble the models studied here, so one cannot learn much about the capacity of a *real* network. Moreover, one cannot devise protocols since the results are not algorithmic.

In contrast there is a body of algorithmic work, however, mostly on graph-based models. Studying wireless communication in graph-based models commonly implies studying some variants of independent set, matching, or coloring [18,26]. Although these algorithms present extensive theoretical analysis, they are constrained to the limitations of a model that ultimately abstracts away the nature of wireless communication. The inefficiency of graph-based protocols in the SINR model is well documented and has been shown theoretically as well as experimentally [13,19,23].

Algorithmic work in the SINR model is fairly new; to the best of our knowledge it was started just three years ago [22]. In this paper Moscibroda et al. present an algorithm that successfully schedules a set of links (carefully chosen to strongly connect an arbitrary set of nodes) in polylogarithmic time, even in arbitrary worst-case networks. In contrast to our work the links themselves are *not* arbitrary (but do have structure that will simplify the problem). This work has been extended and applied to topology control [8,24], sensor networks [20], combined scheduling and routing [5], or ultra-wideband [15], or analog network coding [12]. Recently a moderately exponential-time algorithm has been proposed [16]. Apart from these papers, algorithmic SINR results also started popping up here and there, for instance in a game theoretic context or a distributed algorithms context, e.g., [1,2,4,10,17,25]

So far there are only a few papers that tackle the general problem of scheduling arbitrary wireless links. Goussevskaia et al. give a simple proof that the problem is NP-complete [11], and [21] test popular heuristics. Both papers also present approximation algorithms, however, in both cases the approximation ratio may grow linearly with the network size.

The work most relevant for this paper is by Goussevskaia et al. [9]. Among other things, [9] presents the first wireless scheduling algorithm with approximation guarantee independent of the topology of the network. The paper accomplishes a constant approximation guarantee for the problem of maximizing the number of links scheduled in one single time-slot. Furthermore, by applying that single-slot subroutine repeatedly the paper realizes a $O(\log n)$ approximation for the problem of minimizing the number of time slots needed to schedule a given set of arbitrary requests.

Our present paper removes the logarithmic approximation overhead of [9]. Hence the problem of wireless scheduling is in APX. Moreover, our algorithm is simpler than [9], and will be easier to build on. In addition we are able to present a quite general robustness result, showing that constant parameter and model changes will modify the result only by a constant.

3 Notation and Model

Given is a set of links $\ell_1, \ell_2, \dots, \ell_n$, where each link ℓ_v represents a communication request from a sender s_v to a receiver r_v . We assume the senders and receivers are points in the Euclidean plane; this can be extended to other metrics. The Euclidean distance between two points p and q is denoted $d(p, q)$. The asymmetric distance from link v to link w is the distance from v 's sender to w 's receiver, denoted $d_{vw} = d(s_v, r_w)$. The length of link ℓ_v is denoted $d_{vv} = d(s_v, r_v)$. We shall assume for simplicity of exposition that all links are of different length; this does not affect the results. We assume that each link has a unit-traffic demand, and model the case of non-unit traffic demands by replicating the links. We also assume that all nodes transmit with the same power level P . We show later how to extend the results to variable power levels, with a slight increase in the performance ratio.

We assume the *path loss radio propagation* model for the reception of signals, where the received signal from w at receiver v is $P_{wv} = P/d_{wv}^\alpha$ and $\alpha > 2$ denotes the path-loss exponent. When $w \neq v$, we write $I_{wv} = P_{wv}$. We adopt the *physical interference model*, in which a node r_v successfully receives a message from a sender s_v if and only if the following condition holds:

$$\frac{P_{vv}}{\sum_{\ell_w \in S \setminus \{\ell_v\}} I_{wv} + N} \geq \beta, \quad (1)$$

where N is the ambient noise, β denotes the minimum SINR (signal-to-interference-plus-noise-ratio) required for a message to be successfully received, and S is the set of concurrently scheduled links in the same channel or *slot*. We say that S is *SINR-feasible* if (1) is satisfied for each link in S .

The problems we treat are the following. In all cases are we given a set of links of arbitrary lengths. In the Scheduling problem, we want to partition the set of input links into minimum number of SINR-feasible sets, each referred to as a *slot*. In the Single-Shot Scheduling (SSS) problem, we seek the maximum cardinality subset of links that is SINR-feasible. And, in the k -Thruput problem, for a positive integer k , we seek a collection of k disjoint SINR-feasible sets with maximum combined cardinality. Let χ denote the minimum number of slots in an SINR-feasible schedule.

We make crucial use of the following new definitions.

Definition 1. The relative interference (*RI*) of link ℓ_w on link ℓ_v is the increase caused by ℓ_w in the inverse of the SINR at ℓ_v , namely $RI_w(v) = I_{wv}/P_{vv}$. For convenience, define $RI_v(v) = 0$. Let $c_v = \frac{\beta}{1-\beta N/P_{vv}} = \frac{1}{\frac{1}{\beta} - \frac{N}{P_{vv}}}$ be a constant that indicates the extent to which the ambient noise approaches the required signal at receiver r_v . The affectance¹ of link ℓ_v , caused by a set S of links, is the sum of the relative interferences of the links in S on ℓ_v , scaled by c_v , or

$$a_S(\ell_v) = c_v \cdot \sum_{\ell_w \in S} RI_w(v).$$

For a single link ℓ_w , we use the shorthand $a_w(\ell_v) = a_{\{\ell_w\}}(\ell_v)$. We define a p -signal set or schedule to be one where the affectance of any link is at most $1/p$.

Observation 1. The affectance function satisfies the following properties for a set S of links:

1. (Range) S is SINR-feasible iff, for all $\ell_v \in S$, $a_S(\ell_v) \leq 1$.
2. (Additivity) $a_S = a_{S_1} + a_{S_2}$, whenever (S_1, S_2) is a partition of S .
3. (Distance bound) $a_w(\ell_v) = c_v \cdot \left(\frac{d_{wv}}{d_{vv}}\right)^\alpha$, for any pair ℓ_w, ℓ_v in S .

¹ Affectance is closely related to *affectedness*, defined in [9], but treats the effect of noise more accurately.

4 Robustness of SINR

We present here properties of schedules in the SINR model, which double as tools for the algorithm designer. The results of this section apply equally to scheduling links of different powers, including involving topology control. In the next subsection, we examine the desirable property of link dispersion, and how any schedule can be dispersed at a limited cost.

We now explore how signal requirements (in the value of β), or equivalently interference tolerance, affects schedule length. It is not *a priori* obvious that minor discrepancies cause only minor changes in schedule length, but by showing that it is so, we can give our algorithms the advantage of being compared with a stricter optimal schedule. This also has implications regarding the robustness of SINR models with respect to perturbations in signal transmissions.

The pure geometric version of SINR given in (1) is an idealization of true physical characteristics. It assumes, e.g., perfectly isotropic radios, no obstructions, and a constant ambient noise level. That begs the question, why move algorithm analysis from analytically amenable graph-based models to a more realistic model if the latter isn't all that realistic? Fortunately, the fact that schedule lengths are relatively invariant to signal requirements shows that these concerns are largely unnecessary.

The following result on signal requirement applies also to throughput optimization.

Theorem 1. *There is a polynomial-time algorithm that takes a p -signal schedule and refines into a p' -signal schedule, for $p' > p$, increasing the number of slots by a factor of at most $\lceil 2p'/p \rceil^2$.*

Proof. Consider a p -signal schedule \mathcal{S} and a slot S in \mathcal{S} . We partition S into a sequence S_1, S_2, \dots of sets. Order the links in S in some order, e.g., decreasing order. For each link ℓ_v , assign ℓ_v to the first set S_j for which $a_{S_j}(\ell_v) \leq 1/2p'$, i.e. the accumulated affectance on ℓ_v among the previous, longer links in S_j is at most $1/2p'$. Since each link ℓ_v originally had affectance at most $1/p$, then by the additivity of affectance, the number of sets used is at most $\lceil \frac{1/p}{1/2p'} \rceil = \lceil \frac{2p'}{p} \rceil$.

We then repeat the same approach on each of the sets S_i , processing the links this time in increasing order. The number of sets is again $\lceil \frac{2p'}{p} \rceil$ for each S_i , or $\lceil \frac{2p'}{p} \rceil^2$ in total. In each final slot (set), the affectance on a link by shorter links in the same slot is at most $1/2p'$. In total, then, the affectance on each link is at most $2 \cdot 1/2p' = 1/p'$.

This result applies in particular to optimal solutions. Let OPT_p be an optimal p -signal schedule and let χ_p be the number of slots in OPT_p . It is not *a priori* clear that a smooth relationship exists between χ_p and χ , for $p > 1$.

Corollary 1. $\chi_p \leq \lceil 2p \rceil^2 \chi$.

This has significant implications. One regards the validity of studying the pure SINR model. As asked in [9], “what if the signal is attenuated by a certain

factor in one direction but by another factor in another direction?” A *generalized physical model* was introduced in [24] to allow for such a deviation.

Theorem 1 implies that scheduling is relatively robust under discrepancies in the SINR model. This validates the analytic study of the pure SINR model, in spite of its simplifying assumptions.

Corollary 2. *If a scheduling algorithm gives a ρ -approximation in the SINR model, it provides a $O(\theta^2\rho)$ -approximation in variations in the SINR model with a discrepancy of up to a factor of θ in signal attenuation or ambient noise levels.*

This result can be contrasted with the strong $n^{1-\epsilon}$ -approximation hardness of scheduling in an abstract (non-geometric) SINR model that allows for arbitrary distances between nodes [11]. Alternatively, Theorem 1 allows us to analyze algorithms under more relaxed situations than the optimal solutions that we compare to.

4.1 Dispersion Properties

One desirable property of schedules is that links in the same slot be spatially well separated. This blurs the difference in position between sender and receiver of a link, since it affects distances only by a small constant. Intuitively, we want to measure nearness as a fraction of the lengths of the respective links. Given the affectance measure, it proves to be useful to define it somewhat less restrictively.

Definition 2. *Link ℓ_w is said to be q -near link ℓ_v , if $d_{wv} < q \cdot c_v^{1/\alpha} \cdot d_{vv}$. A set of links is q -dispersed if no (ordered) pairs of links in the set are q -near. A schedule is q -dispersed if all the slots are formed by q -dispersed sets.*

Observation 1, item 3, states that link w is q -near a link ℓ_v iff $a_w(\ell_v) > q^{-\alpha}$. This immediately gives the following strengthening of Lemma 4.2 in [9].

Lemma 2. *Fewer than q^α senders in an SINR-feasible set S are q -near to any given link $\ell_v \in S$.*

At a cost of a constant factor, any schedule can be made dispersed.

Lemma 3. *There is a polynomial-time algorithm that takes a SINR-feasible schedule and refines it into a q -dispersed schedule, increasing the number of slots by a factor of at most $(q + 2)^\alpha$.*

Proof. Let S be a slot in the schedule. We show how to partition S into sets S_1, S_2, \dots, S_t that are q -dispersed, where $t \leq (q + 2)^\alpha + 1$.

Process the links of S in increasing order of length, assigning each link ℓ_v “first-fit” to the first set S_j in which the receiver r_v is at least $\left(qc_v^{1/\alpha} + 2\right) \cdot d_{vv}$ away from any other link. Let ℓ_w be a link previously in S_j , and note that ℓ_w

is shorter than ℓ_v . By the selection rule, $d_{wv} \geq (qc_v^{1/\alpha} + 2) \cdot d_{vv} \geq qc_v^{1/\alpha} \cdot d_{vv}$. Also,

$$d_{vv} \geq d_{wv} - d_{ww} - d_{vv} \geq (qc_v^{1/\alpha} + 1) d_{vv} - d_{ww} \geq qc_v^{1/\alpha} d_{ww}.$$

Since this holds for every pair in the same set, the schedule is q -dispersed. Suppose S_t is the last set used by the algorithm, and let ℓ_v be a link in it. Then, each S_i , for $i = 1, 2, \dots, t-1$, contains a link whose sender is closer than $(qc_v^{1/\alpha} + 2) \cdot d_{vv} \leq (q+2)c_v^{1/\alpha} d_{vv}$ to r_v , i.e., is $(q+2)$ -near to ℓ_v . By Lemma 2, $t-1 < (q+2)^\alpha$.

Let χ^q denote the minimum number of slots in a q -dispersed schedule.

Corollary 3. $\chi^q \leq (q+2)^\alpha \cdot \chi$.

Intuitively, there is a correlation between low affectance and high dispersion in schedules. The following result makes this connection clearer. The converse is, however, not true, since interference can be caused by far-away links.

Lemma 4. *A p -signal schedule is also $p^{1/\alpha}$ -dispersed.*

Proof. Let ℓ_v and ℓ_w be an ordered pair of links in a slot S in a p -signal schedule. By definition, $a_w(\ell_v) \leq a_S(\ell_v) \leq 1/p$. By Observation 1, item 3, $d_{wv} \geq p^{1/\alpha} c_v^{1/\alpha} \cdot d_{vv}$. Hence, the lemma.

5 Scheduling Approximation

The algorithm we analyze is a slightly simplified version of the algorithm of [9]. It involves repeated application of the following algorithm for the Single-Shot Scheduling problem.

$$\text{Let } c = 1/\tau^\alpha, \text{ where } \tau = 2 + \max\left(2, \left(2^6 3 \beta^{\frac{\alpha-1}{\alpha-2}}\right)^{\frac{1}{\alpha}}\right).$$

A(c)
 sort the links $\ell_1, \ell_2, \dots, \ell_n$ by non-decreasing order of length
 $S \leftarrow \emptyset$
 for $v \leftarrow 1$ to n do
 if $(a_S(\ell_v) \leq c)$
 add ℓ_v to S
 output S

We shortly show that this algorithm also gives a $O(1)$ -approximation to the Single-Shot Scheduling problem. It is rather surprising that a $O(1)$ -approximation algorithm can be obtained in a single sweep. This should help in applying the ideas further, e.g., in distributed implementations. Simulation results in [9] also indicate very good practical performance, in relation to previous algorithms, and the simplification given here is likely to perform at least as well.

Instead of applying algorithm **A** repeatedly, we equivalently implement it as the following algorithm **B**:

B(c)
 sort the links $\ell_1, \ell_2, \dots, \ell_n$ by non-decreasing order of length
 $S_i \leftarrow \emptyset$, for $i = 1, 2, \dots$
 for $v \leftarrow 1$ to n do
 assign ℓ_v to the first set S_i for which $a_{S_i}(\ell_v) \leq c$
 output $S = (S_1, S_2, \dots)$

It is not immediate that algorithm **A** (or, equivalently, **B**) produces a feasible solution.

Lemma 5. *Algorithms **A** and **B** produce a $\tau - 2$ -dispersed solution.*

Proof. Let ℓ_w be a link in the set S output by algorithm **A**. Let N^- (N^+) be the set of links in S that are shorter (longer) than ℓ_w . Consider first a link $\ell_u \in N^-$. Since ℓ_w was added by the algorithm after adding ℓ_u , $a_u(\ell_w) \leq c = 1/\tau^\alpha$, which implies by Observation 1, item 3, that $d_{uw} \geq \tau c_w^{1/\alpha} d_{ww} > (\tau - 2)c_w^{1/\alpha} d_{ww}$. Consider next a link $\ell_v \in N^+$. Since ℓ_v was added after ℓ_w , it holds that $a_w(\ell_v) \leq c = 1/\tau^\alpha$. So by Observation 1, $d_{vw} \geq \tau \cdot c_v^{1/\alpha} d_{vv}$. Note that $c_v \geq c_w$ whenever $d_{vv} \geq d_{ww}$. Then, using the triangular inequality,

$$d_{vw} = d(s_v, r_w) \geq d_{vv} - d_{vw} - d_{ww} \geq (\tau c_v^{1/\alpha} - 2) d_{vv} \geq (\tau - 2)c_w^{1/\alpha} d_{ww}.$$

Since this holds for every ordered pair in S , we have that S is $(\tau - 2)$ -dispersed.

The following appeared as part of Lemma 4.1 in [9], and has also been applied in similar forms directly or indirectly elsewhere (e.g. [6]).

Lemma 6. *Let ℓ_v be a link in an SINR-feasible set S . Let N_z^+ be the set of links in S that are at least as long as ℓ_v and whose senders are of distance greater than $z \cdot d_{vv}$ from r_v . Then,*

$$a_{N_z^+}(\ell_v) < \left(\frac{\alpha - 1}{\alpha - 2} 2^5 3 \right) z^{-\alpha} c_v.$$

Theorem 2. *Algorithms **A** and **B** produce an SINR-feasible solution.*

Proof. Let ℓ_w be a link in the set S output by algorithm **A**. Let N^- (N^+) be the set of links in S that are shorter (longer) than ℓ_w . The links in N^- were processed before ℓ_w , so by the if-condition in the algorithm, $a_{N^-}(\ell_w) \leq c$. By Lemma 5, S is $\tau - 2$ -dispersed, so by Lemma 6 and the definitions of τ and dispersion,

$$a_{N^+}(\ell_w) < \left(\frac{\alpha - 1}{\alpha - 2} 2^5 3 \right) \frac{c_v}{(\tau - 2)^\alpha c_v} c_v \leq \frac{1}{2}.$$

Hence, the affectance of each link in S is at most $c + 1/2 < 1$.

5.1 Performance Analysis

We need an extension of a geometric result from [9]. Let \mathcal{R} and \mathcal{B} be two disjoint sets of points in a metric space, called the *red* and the *blue* points. A blue point g guards a red point w , with respect to a point b , if $d(g, w) \leq d(b, w)$ and the angle $\angle gwb$ is at most 30° . That is, g is contained in the 60° sector emanating from w whose centerline goes through b . See Fig. 1. We say that a point b in \mathcal{B} is *blue-shadowed* if each red point has a private guard in \mathcal{B} with respect to b ; i.e., there is an injective function $f : \mathcal{R} \rightarrow \mathcal{B} \setminus \{b\}$ such that $f(w)$ guards w from b for any $w \in \mathcal{R}$.




Fig. 1. Blue point g guards red point w from blue point s_b . If the blue points are sufficiently dispersed, then the receiver r_b will also be closer to g than to w .

The following result is a variation on Lemma 4.4 (“Blue-dominant centers lemma”) of [9].

Lemma 7 (Blue-shadowed lemma). *Let \mathcal{R} and \mathcal{B} be two disjoint sets of red and blue points in 2-dimensional Euclidean space. If $|\mathcal{B}| > 12 \cdot |\mathcal{R}|$, then there is a blue-shadowed point in \mathcal{B} .*

Proof. Process the points in \mathcal{R} in an arbitrary order, we work with a subset \mathcal{B}' of \mathcal{B} initially set at $\mathcal{B}' = \mathcal{B}$. We shall assign each $r \in \mathcal{R}$ a set $\{g_1^r, g_2^r, \dots, g_{12}^r\}$ of guards.

For each point $r \in \mathcal{R}$ in order, let g_i^r be the blue point closest to r among the points in \mathcal{B}' that are contained in the 30° -sector sec_i at angle in the range $[(i - 1) \cdot 30^\circ, i \cdot 30^\circ)$ emanating from r . If a sector i contains no blue point in \mathcal{B}' , then no point is assigned as g_i^r . We then remove these points g_i^r from \mathcal{B}' and continue with the next point in \mathcal{R} .

After going through all the points in \mathcal{R} , the set \mathcal{B}' is still nonempty by the assumption on the relative sizes of \mathcal{R} and \mathcal{B} . We claim that every point in \mathcal{B}' is now blue-shadowed. Let b be such a point and consider a point $r \in \mathcal{R}$. Consider the 60° -sector emanating from r whose centerline goes through b . This sector properly contains one of the 30° -sectors sec_i , and thus contains one of r 's guards. Since b was not selected as a guard, there was a guard selected for that sector and it is closer to r than b is. Since this holds for any point r , b is blue-shadowed.

The following lemma builds on Lemma 4.5 of [9]. Note that the straightforward modification of that lemma appears insufficient. Instead, we need something like our Theorem 1, allowing us to compare the algorithm’s solution with the stricter optimal solution OPT_c . We also utilize the dispersion property to simplify the proof argument.

Lemma 8. *Let $\rho = 12$. Let S_k be the set of links scheduled by algorithm **B** in slot k , and let X_k be the set of links scheduled in slot k of OPT_c . Further, let $S_k = \cup_{i=1}^k S_i$ and $X_k = \cup_{i=1}^k X_i$. Then, for any positive integer k , $|S_{\rho k}| \geq |X_k|$.*

Proof. Suppose the claim is false for some integer k . Then, $|S_{\rho k}| < |X_k|$ or, equivalently, $|S_{\rho k} \setminus X_k| < |X_k \setminus S_{\rho k}|$. Thus, there are slots $i_0, 1 \leq i_0 \leq \rho k$, and $j_0, 1 \leq j_0 \leq k$, for which $|S_{i_0} \setminus X_k| < |X_{j_0} \setminus S_{\rho k}|/\rho$. Let $S = S_{i_0}$, $S' = S \setminus X_k$, $X = X_{j_0}$, and $X' = X \setminus S_{\rho k}$.

Since OPT_c is a $1/c$ -signal schedule and $1/c = \tau^\alpha$, X' is also a τ^α -signal set. By Lemma 4, X' is then τ -dispersed. In particular, it is 3-dispersed.

Let $\mathcal{B} = \{s_v | \ell_v \in X'\}$ and $\mathcal{R} = \{s_w | \ell_w \in S'\}$ be the sets of senders in X' and S' ; we call them blue and red points, respectively. By Lemma 7, there is a blue-shadowed point (sender) s_b in \mathcal{B} . We shall argue that the link $\ell_b = (s_b, r_b)$ would have been picked up by our algorithm for the slot i_0 .

Consider any red point (sender) $w \in \mathcal{R}$, and let $g = f(w)$ be the guard for w guaranteed by the blue-shadowed lemma. Since g guards w , $d(s_b, w) \geq d(s_b, g)$. By the dispersion property, $d(g, r_b) \geq \tau \cdot d_{bb}$. Thus,

$$d(s_b, w) \geq d(s_b, g) \geq d(r_b, g) - d_{bb} \geq (\tau - 1) \cdot d_{bb} = (\tau - 1)d(s_b, r_b).$$

Then, the angle $\angle r_b w s_b$ is at most $\arcsin 1/(\tau - 1) \leq 30^\circ$, since $\tau - 1 \geq 2$. That implies that r_b is contained in the 60° -sector emanating from w with centerline going through s_b , just like the guard g . See Fig. 1 for the relative positions of the points. Then, r_b is closer to g than to w . Thus, $a_g(\ell_b) > a_w(\ell_b)$. Summing up over all links w in \mathcal{R} and their guards $f(w)$, we get

$$a_{S'}(\ell_b) = \sum_{w \in \mathcal{R}} a_w(\ell_b) < \sum_{w \in \mathcal{R}} a_{f(w)}(\ell_b) \leq \sum_{v \in \mathcal{B}} a_v(\ell_b) = a_{X'}(\ell_b).$$

Thus, since the affectance threshold of X is c ,

$$a_S(\ell_b) = a_{S'}(\ell_b) + a_{S \cap X}(\ell_b) < a_{X'}(\ell_b) + a_{S \cap X}(\ell_b) = a_X(\ell_b) \leq c,$$

which contradicts the fact that ℓ_b was not selected into S .

The following result is largely immediate from Lemma 8.

Theorem 3. *Algorithm **B** outputs a schedule that approximates both the Scheduling and k -Thruput problems, for every $k \geq 1$, within a constant factor.*

Proof. By Lemma 8 and Theorem 1, the number ALG of slots used by algorithm **B** is bounded by

$$ALG \leq \rho \chi_c \leq \rho \left\lceil \frac{2}{c} \right\rceil^2 \chi.$$

Also, by Lemma 8, the number of links scheduled by \mathbf{B} in the first $12k$ slots is at least the number of links in an optimal c -signal k -Thruput solution. Again by Theorem 1, we obtain a constant factor approximation to k -Thruput.

5.2 Handling Different Transmission Powers

We can treat the case when links transmit with different powers in two different ways. Let P_{max} (P_{min}) be the maximum (minimum) power used by a link, respectively. By introducing a factor of P_{min}/P_{max} into the affectance threshold c , the algorithm \mathbf{B} still produces a feasible schedule, that is longer by a factor of at most P_{max}/P_{min} .

Alternatively, we can partition the instance into “power regimes”, where each regime consists of links whose powers are equal up to a factor of 2. We schedule each power regime separately, obtaining an approximation factor of at most $\log P_{max}/P_{min}$, or at most the number of different power values.

6 Conclusions

This paper shows that wireless scheduling is in APX. Having a constant approximation algorithm for wireless scheduling implies that we can derive the single-hop throughput capacity of an arbitrary wireless network, up to a constant factor. As such this paper basically solves the scheduling complexity introduced by Moscibroda et al. [22]. However, various parameter combinations are still open, and deserve more research, e.g. power control, multi-hop traffic, scheduling and routing, analog network coding, models beyond SINR such as log-normal shadowing, to name just a few of the obvious ones.

References

1. Auletta, V., Moscardelli, L., Penna, P., Persiano, G.: Interference games in wireless networks. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 278–285. Springer, Heidelberg (2008)
2. Avin, C., Emek, Y., Kantor, E., Lotker, Z., Peleg, D., Roditty, L.: SNR diagrams: Towards algorithmically usable SINR models of wireless networks. arXiv:0811.3284v2 (2008)
3. Brar, G., Blough, D., Santi, P.: Computationally Efficient Scheduling with the Physical Interference Model for Throughput Improvement in Wireless Mesh Networks. In: Mobicom (2006)
4. Brar, G.S., Blough, D.M., Santi, P.: The scream approach for efficient distributed scheduling with physical interference in wireless mesh networks. In: ICDCS, pp. 214–224 (2008)
5. Chafekar, D., Kumar, V., Marathe, M., Parthasarathy, S., Srinivasan, A.: Cross-layer Latency Minimization for Wireless Networks using SINR Constraints. In: Mobihoc (2007)
6. Chafekar, D., Kumar, V., Marathe, M., Parthasarathy, S., Srinivasan, A.: Approximation Algorithms for Computing Capacity of Wireless Networks with SINR Constraints. In: Infocom (2008)

7. ElBatt, T.A., Ephremides, A.: Joint scheduling and power control for wireless ad hoc networks. *IEEE Transactions on Wireless Communications* 3(1), 74–85 (2004)
8. Gao, Y., Hou, J.C., Nguyen, H.: Topology control for maintaining network connectivity and maximizing network capacity under the physical model. In: *Infocom* (2008)
9. Goussevskaia, O., Halldórsson, M., Wattenhofer, R., Welzl, E.: Capacity of Arbitrary Wireless Networks. In: 28th Annual IEEE Conference on Computer Communications (INFOCOM), Rio de Janeiro, Brazil (April 2009)
10. Goussevskaia, O., Moscibroda, T., Wattenhofer, R.: Local Broadcasting in the Physical Interference Model. In: *ACM SIGACT-SIGOPT International Workshop on Foundations of Mobile Computing (DialM-POMC)*, Toronto, Canada (August 2008)
11. Goussevskaia, O., Oswald, Y.A., Wattenhofer, R.: Complexity in Geometric SINR. In: *Mobihoc* (2007)
12. Goussevskaia, O., Wattenhofer, R.: Complexity of scheduling with analog network coding. In: *FOWANC* (2008)
13. Gronkvist, J., Hansson, A.: Comparison between graph-based and interference-based STDMA scheduling. In: *Mobihoc* (2001)
14. Gupta, P., Kumar, P.R.: The Capacity of Wireless Networks. *IEEE Trans. Information Theory* (2000)
15. Hua, Q.-S., Lau, F.C.M.: The scheduling and energy complexity of strong connectivity in ultra-wideband networks. In: *MSWiM*, pp. 282–290 (2006)
16. Hua, Q.-S., Lau, F.C.M.: Exact and approximate link scheduling algorithms under the physical interference model. In: *DIALM-POMC*, pp. 45–54 (2008)
17. Katz, B., Völker, M., Wagner, D.: Link scheduling in local interference models. In: Fekete, S.P. (ed.) *ALGOSENSORS 2008*. LNCS, vol. 5389, pp. 57–71. Springer, Heidelberg (2008)
18. Kumar, V., Marathe, M., Parthasarathy, S., Srinivasan, A.: Algorithmic Aspects of Capacity in Wireless Networks. In: *Sigmetrics* (2005)
19. Maheshwari, R., Jain, S., Das, S.R.: A measurement study of interference modeling and scheduling in low-power wireless networks. In: *SenSys*, pp. 141–154 (2008)
20. Moscibroda, T.: The worst-case capacity of wireless sensor networks. In: *IPSN* (2007)
21. Moscibroda, T., Oswald, Y.A., Wattenhofer, R.: How optimal are wireless scheduling protocols? In: *INFOCOM*, pp. 1433–1441 (2007)
22. Moscibroda, T., Wattenhofer, R.: The Complexity of Connectivity in Wireless Networks. In: *Infocom* (2006)
23. Moscibroda, T., Wattenhofer, R., Weber, Y.: Protocol Design Beyond Graph-Based Models. In: *Hotnets* (November 2006)
24. Moscibroda, T., Wattenhofer, R., Zollinger, A.: Topology Control meets SINR: The Scheduling Complexity of Arbitrary Topologies. In: *Mobihoc* (2006)
25. Scheideler, C., Richa, A.W., Santi, P.: An $o(\log n)$ dominating set protocol for wireless ad-hoc networks under the physical interference model. In: *MobiHoc*, pp. 91–100 (2008)
26. Zussman, G., Brzezinski, A., Modiano, E.: Multihop Local Pooling for Distributed Throughput Maximization in Wireless Networks. In: *Infocom* (2008)