Wireless Scheduling with Power Control

Magnús M. Halldórsson[*]

July 1, 2009

Abstract

We consider the scheduling of arbitrary wireless links in the physical model of interference to minimize the time for satisfying all requests. We study here the combined problem of scheduling and power control, where we seek both an assignment of power settings and a partition of the links so that each set satisfies the signal-to-interference-plus-noise (SINR) constraints.

We give an algorithm that attains an approximation ratio of $O(\log n \cdot \log\log \Lambda)$, where Λ is the ratio between the longest and the shortest linklength. Under the natural assumption that lengths are represented in binary, this gives the first *polylog*(n)-approximation. The algorithm has the desirable property of using an oblivious power assignment, where the power assigned to a sender depends only on the length of the link. We show this dependence on Λ to be unavoidable, giving a construction for which any oblivious power assignment results in a $\Omega(\log\log \Lambda)$-approximation.

We also give a simple *online* algorithm that yields a $O(\log \Lambda)$-approximation, by a reduction to the coloring of unit-disc graphs. In addition, we obtain improved approximation for a bidirectional variant of the scheduling problem, give partial answers to questions about the utility of graphs for modeling physical interference, and generalize the setting from the standard 2-dimensional Euclidean plane to doubling metrics.

1 Introduction

We are interested in fundamental limits on communication in wireless networks. How much communication throughput is possible? This is an issue of efficient spatial separation, keeping the interference from simultaneously communicating links sufficiently low. The interference scheduling problem is then to schedule an arbitrary set of communication links in the least amount of time while satisfying interference constraints. In this paper, we focus on the power control version, where we also choose the power settings for the links.

The scheduling problem depends strongly on the model of interference. Until recently, previous algorithmic work has revolved around various graph-based models, where interference is modeled as a pairwise constraint. This, however, fails to capture the accumulative property of actual radio signals. In contrast, 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, to be formally introduced in Section 2. The SINR model reflects physical reality more accurately and is therefore often simply called the physical model. On the other hand, "EE researchers" tend to propose heuristics that are evaluated by simulation, which neither give insights into the complexity of the problem nor give algorithmic results that may ultimately lead to new protocols.

---

[*] School of Computer Science, Reykjavik University, & Icelandic Center of Excellence in Theoretical Computer Science (ICE-TCS). Work done in part while visiting Research Institute for Mathematical Sciences (RIMS) at Kyoto University. Supported by Icelandic Research Fund grant 90032021. Email: `mmh@ru.is`

In a seminal work, Moscibroda and Wattenhofer [16] initated the study of scheduling in the SINR model. Formally, given is an arbitrary set of links, each a sender-receiver pair of points in the plane. We seek an assignment of power settings to the senders and a partition of the linkset into minimum number of slots, so that the links in each slot satisfy the SINR-constraints. We refer to this as the PC-Scheduling problem. In the related bidirectional scheduling problem, both nodes in a link may be transmitting, which implies a stronger, symmetric form of interference.

We seek algorithms that result in good schedules. Beyond this rather obvious objective, we seek to map the landscape of the issues surrounding SINR-scheduling. In particular, we shall study in this paper two relevant questions: the utility of "simple" power allocation strategies, and the extent to which graphs can capture interference in the SINR-model.

For reasons of simplicity of use, it is desirable to use power assignments that are precomputable independent of other links. Such *oblivious* assignments depend only on the length of the given link. In fact, oblivious assignments appear unavoidable in the distributed setting. The two most frequently used power assignment strategies are indeed of this type, using either *uniform* (or fixed) power for all the links, or *linear* assignment that ensures that the signals received at the intended receivers are identical.

The other issue of particular interest is the utility of graphs for modeling interference. It is *a priori* given that graphs are imperfect models, given both the non-locality and the additive nature of interference in the SINR model. The perceived difficulty in reasoning analytically about these additional complications has been cited as a factor against SINR model. Still, graphs have proved to be highly versatile tools for analysis and algorithm design, and pairwise constraints are in general much easier to handle than many-to-many constraints. We would therefore like to quantify the cost of doing business using graphs, or the overhead that amenable graph models have over non-graphic models, as well as pinpointing particular situations where graphs work especially well.

1.1 Our contributions

We present a simple scheduling algorithm that works for any oblivious power assignment strategy, resulting in a $O(\log \Lambda)$ approximation ratio, where Λ is the ratio between the maximum and minimum link length. In particular, when all links are of nearly equal length, we obtain the first constant approximation. This matches the constructions in [15] that show that both fixed and linear assignments can be as much as $\Omega(\log \Lambda)$ factor from optimal.

We then examine a new oblivious assignment, the *mean* assignment studied recently in [7], which is the geometric mean of the uniform and linear power assignments. We give a simple scheduling algorithm that uses mean assignment and obtains a $O(\log \log \Lambda \cdot \log n)$-approximation. Under the natural assumption that lengths can be represented in binary, this implies also $O(\log^2 N)$-approximation, where N is the size of the input. In the bidirectional version, the algorithm results in a $O(\log n)$-approximation, improving on the previous $O(\log^c n)$-factor with $c > 6$ [6] using considerably simpler arguments.

We show that the dependence on Λ is unavoidable for oblivious power assignment strategies. Namely, any oblivious assignment forces $\Omega(\log \log \Lambda)$-approximate schedules. Thus, within the framework of power assignments that are oblivious of the link instance, our results are best possible up to logarithmic factor. Our results also pinpoint the issues of essential difficulty in the scheduling of wireless links. When the input instance is "well-behaving" in that no sender is much closer to the receiver of another link than its sender, or when the criteria is changed from unidirectional to bidirectional scheduling, the bound improves to a single logarithmic factor. Thus, we can characterize more precisely the structural property of link arrangements that make scheduling hard.

Our results apply to the standard setting of the two-dimensional Euclidean plane with the path-loss constant $\alpha > 2$ (see Section 2). More generally, they hold for a general class of distance metrics that are *doubling metrics*, for the case when α is greater than the doubling constant of the metric. The requirement on α is to ensure that the cumulative power of a transmission fades away. This is a natural assumption, since preservation or amplification would contradict the second law of thermodynamics. We can also extend all the results to general metrics, with a roughly logarithmic increase in the approximation factor.

The simplicity of our algorithms leaves them suitable for distributed implementation. For links of nearly equal length, we can use uniform power, and we show that the problem reduces, within a constant factor, to the coloring of unit-disc graphs, a very well-studied problem. This also implies an $O(\log \Lambda)$-competitive online scheduling algorithm.

Our work also gives partial answer to a nagging question regarding the utility of graphs in representing physical models of interference. Our results indicate that graphs can still play a useful role. For nearly equal length links, the basic class of unit-disc graphs is in fact sufficient. The $O(\log n \cdot \log\log \Lambda)$-approximation result is also relative to the underlying graph.

1.2 Related Work

Most work in wireless scheduling in the physical (SINR) model has been of heuristic nature, e.g. [5]. Only after the work of Gupta and Kumar [11] did analytical results became *en vogue*, but were largely non-algorithmic and restricted to networks with a well-behaving topology and traffic pattern such as uniform geometric distribution.

In contrast, the body of algorithmic work is mostly on graph-based models that ultimately abstract 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 [10, 14, 17].

Approximation algorithms for the problem of scheduling wireless links in the SINR model were given in [18], [15] and [3]. In all cases the performance ratios obtained consist of the product of structural properties and a function of the number of nodes. The structural properties are different but can all grow linearly with the size of the network.

A number of recent related results have featured a $O(\log \Lambda)$-approximation. Andrews and Dinitz [1] gave a $O(\log \Lambda)$ approximation for the Single-Slot scheduling problem. Fanghänel, Kesselheim and Vöcking [7] gave a randomized algorithm for the scheduling problem using linear power assignment that uses $O(OPT \log \Lambda + \log^2 n)$ slots, matching our results for dense instances. And finally, Avin, Lotker and Pignolet [2] show that assumption of $\alpha > 2$ used by all previous work may not be necessary, in that the ratio between non-oblivious and oblivious Single-Slot schedules is $O(\log \Lambda)$, at least in the 1-dimensional metric.

In [6], Fanghänel et al. give a construction that shows that any schedule based on any oblivious power assignment can be a factor of n from optimal. We show that in terms of Λ, the gap is actually $\Omega(\log\log \Lambda)$, using similar constructions. They also introduce the bidirectional version of the scheduling problem and give a $O(\log^{4.5+\alpha} n)$-approximation factor using the mean power assignment in general metrics. Their proof involves non-trivial embeddings into tree metric spaces.

In contrast, the scheduling complexity of arbitrary links in the case of fixed, uniform power is now fairly well understood. Constant factor approximations were recently obtained for the Single-Slot scheduling problem [8] and the scheduling problem [12]. Both of these problems are known to be NP-complete [9]. The results obtained here for power control build on and extend the techniques and properties derived in the case of uniform power in [8, 12].

2 Notation and Preliminaries

Given is a set $L = \{\ell_1, \ell_2, \ldots, \ell_n\}$ of links, where each link ℓ_v represents a communication request from a sender s_v to a receiver r_v. The distance between two points x and y is denoted $d(x, y)$. 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 simply ℓ_v. We shall assume for simplicity of exposition that all links are of different length; this does not affect the results materially. We assume that each link has a unit-traffic demand, and model the case of non-unit traffic demands by replicating the links.

$d(x, y)$
d_{vw}
ℓ_v

The nodes can transmit with different power. Let P_v denote the power assigned to node v. We assume the *path loss radio propagation* model for the reception of signals, where the signal received from w at receiver v is P_w/d_{wv}^α and $\alpha > 2$ denotes the path-loss exponent. 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:

P_v
α

$$\frac{P_v/\ell_v^\alpha}{\sum_{\ell_w \in S \setminus \{\ell_v\}} P_w/d_{wv}^\alpha + N} \geq \beta, \tag{1}$$

where N is the ambient noise, $\beta \geq 1$ denotes the minimum SINR (signal-to-noise-ratio) required for a message to be successfully received, and S is the set of concurrently scheduled links in the same *slot*. Note that by scaling the power of all the senders, the effect of the noise can be made arbitrarily small, thus we ignore this term. Of course, in real situations, there are upper bounds on maximum power, etc, which we ignore here. We shall also assume that $\beta = 1$; by Prop. 2.2 this does not affect the results materially. We say that S is *SINR-feasible* if (1) is satisfied for each link in S.

The *affectance* of link ℓ_v caused by a set S of links, is the sum of the interferences of the links in S on ℓ_v relative to the power received, or

$a_S(\ell_v)$

$$a_S(\ell_v) = \sum_{\ell_w \in S \setminus \{v\}} \frac{P_w/d_{wv}^\alpha}{P_v/\ell_v^\alpha} = \sum_{\ell_w \in S \setminus \{v\}} \frac{P_w}{P_v} \cdot \left(\frac{\ell_v}{d_{wv}}\right)^\alpha$$

For a single link ℓ_w, we use the shorthand $a_w(v) = a_{\{\ell_w\}}(\ell_v)$. Note that affectance is additive in that for disjoint sets of links S_1, S_2, $a_{S_1 \cup S_2}(\ell_v) = a_{S_1}(\ell_v) + a_{S_2}(\ell_v)$. For convenience, let $a_v(v) = 0$.

$a_w(\ell_v)$

Let $OPT = OPT(L)$ denote a SINR-feasible schedule with minimum number of slots. Let $\Gamma = \Gamma(L)$ denote the number of slots in OPT. A *p-signal* set or a schedule is one where the affectance of any link is at most $1/p$. A set is SINR-feasible iff it is a 1-signal set. Let OPT_p be a p-signal schedule with minimum number of slots, and let Γ_p denote its number of slots. Let Λ denote the ratio between the maximum and minimum length of a link.

OPT
Γ
p-signal
OPT_p
Γ_p
Λ

For a graph G, let $\Delta(G)$ denote the maximum degree of a vertex, and $\chi(G)$ denote the chromatic number.

$\Delta(G)$
$\chi(G)$

We extend the setting from the Euclidean plane to *doubling metrics* (see Clarkson [4]). A metric space is a pair $(\mathcal{U}, d)$, where $\mathcal{U}$ is a set and d is a distance function, satisfying: $d(x, x) = 0$, $d(x, y) = d(x, y)$ (symmetry), and $d(x, y) + d(y, z) \leq d(x, z)$ (triangular inequality), for any points $x, y, z \in \mathcal{U}$. Intuitively, a metric space is doubling if the volume of a ball increases by at most a constant times the radius. Let $B(y, \epsilon) = \{x \in \mathcal{U} | d(x, y) < \epsilon\}$ be the ϵ-ball centered at y. A set $Y \subset \mathcal{U}$ is an *ϵ-packing* if $d(x, y) > 2\epsilon$, for any $x, y \in Y$. That is, the set of balls $\{B(y, \epsilon) | y \in Y\}$ are disjoint. The packing number $\mathcal{P}(\mathcal{U}, \epsilon)$ is the size of the largest ϵ-packing. The *Assouad dimension* $\dim_A(\mathcal{U}, d)$ (also known as uniform metric dimension, doubling dimension) for a space $(\mathcal{U}, d)$ is the value t, if it exists, such that

$\dim_A$

$$\sup_{x \in \mathcal{U}, r > 0} \mathcal{P}(B(x, r), \epsilon r) = C \cdot 1/\epsilon^t,$$

as $\epsilon \rightarrow 0$, where C is an absolute constant. It is known that $\dim_A(\Re^n) = n$ [13]. We require that the path loss exponent α be strictly greater than the doubling dimension $\dim_A(\mathcal{U}, d)$ of the metric. We shall refer to such a combination of distance metric and path loss function as a *fading metric.* C

Preliminaries We shall refer to two links ℓ_v and ℓ_w as *q-independent* if they satisfy the constraint q-independent

$$d_{vw} \cdot d_{wv} \geq q^2 \cdot \ell_w \ell_v \ .$$

A set S of links is a *q-independent set* if the links in S are mutually q-independent.

Define the *link graph* $G_q(L)$ on a link set L, parameterized by a constant q such that a pair of links are adjacent in G_q iff they are not q-independent. The following observation shows that a schedule of a linkset forms a coloring of the corresponding link graph. The converse, however, does not necessarily hold, as we shall see. Thus, the graph representation is more relaxed than required. $G_q(L)$

Lemma 2.1 *Links that belong to the same q^α-signal slot are q-independent.*

Proof: Since the links belong to the same p-signal slot, they satisfy

$$\frac{P_v/\ell_v^\alpha}{P_w/d_{wv}^\alpha} \geq p, \quad \text{and} \quad \frac{P_w/\ell_w^\alpha}{P_v/d_{vw}^\alpha} \geq p \ .$$

By multiplying these inequalities together and rearranging, we get that

$$d_{vw} \cdot d_{wv} \geq p^{2/\alpha} \cdot \ell_w \ell_v = q^2 \cdot \ell_w \ell_v \ .$$

□

One of the main difficulty in scheduling is the asymmetry of the links. In more stringent schedules, the links are kept further apart, which diminishes the problems of asymmetry. For this purpose, the following result from [12] is crucial, which shows that increasing stringency affects only the constant in the approximation ratio.

Proposition 2.2 ([12]) *For any $p \geq 1$ and any linkset L, $\Gamma_p(L) \leq \lceil 2p \rceil^2 \Gamma(L)$.*

3 Uniform power assignment

One of the most widely used power assignment is the uniform one, where senders use the same power setting. This might be viewed as ultra-oblivious, as transmissions are now independent of link length.

We show in Sec. 3.1 that uniform power assignment performs very well when links are of nearly equal lengths. This results in a $O(\log \Lambda)$-approximation of PC-Scheduling, using any oblivious power assignment. Additionally, the global nature of the problem disappears, and local strategies become sufficient. In fact, as we show in Sec. 3.2, it suffices to color a unit-disc graph, with discs of radius proportional to the link lengths, to obtain a constant approximation.

3.1 Nearly equal linklengths

We say that a set of links is *nearly equilength* if lengths of any pair of links in the set differ by a factor of at most 2. We first observe that the scheduling complexity of a linkset is at least proportional to the degree of its link graph.

Lemma 3.1 *Let L be a set of nearly equilength links and q be a constant. Then, $\Gamma(L) = \Omega(\Delta(G_q))$.*

Proof: Let ℓ_v be a link with a set N_v of $\Delta(G_q)$ neighbors in G_q. We shall argue that the links in N_v must belong to distinct slots in any p'-signal schedule, for some $p' = O(q^\alpha)$. The theorem then follows from the signal-strengthening Proposition 2.2.

Let ℓ_u and ℓ_w be links in N_v and let D be the longest length of a link in N_v. By the link relationship in G_q, we have that $d_{uv} \cdot d_{vu} \le q^2 \ell_v \ell_u \le q^2 D^2$ and $d_{wv} \cdot d_{vw} \le q^2 D^2$. For any pair of links ℓ_x, ℓ_y in N_v, we have by the triangular inequality that $d_{xy} \le d_{yx} + 2D$. Thus, we have that $\max(d_{uv}, d_{vu}, d_{wv}, d_{vw}) \le (q+2)D$. Again by the triangular inequality, $d_{uw} \le d(s_u, r_v) + d(r_v, s_w) + d(s_w, r_w) = d_{uv} + d_{wv} + \ell_w \le (2q+5)D$. Similarly, $d_{wu} \le (2q+5)D$. Thus,

$$d_{wu} \cdot d_{uw} \le (2q+5)^2 D^2 \le (4q+10)^2 \ell_w \ell_u \ .$$

Hence, N_v forms a clique in G_{4q+10}. Hence, by Lemma 2.1 $\Gamma_{p'} \ge \Delta(G_q)$, for $p' = (4q+10)^\alpha$. By Proposition 2.2, the theorem now follows. □

The following results extends similar lemmas in previous works (see [8, 12]) from the setting of the Euclidean plane to the more general class of doubling metrics. It yields a converse of Lemma 2.1 for the case of nearly equilength links. This is the only place where we use the fading property of the metric, i.e., that α is strictly greater than the doubling dimension.

Lemma 3.2 *Let S be a z-independent set of nearly equilength links in a fading metric, with uniform power assignment. Then, S is a $\Omega(z^\alpha)$-signal set.*

Proof: Let d be the shortest linklength in S. We first observe that senders of links in S are of mutual distance at least $(z-2)d$. Suppose otherwise that $d(s_u, s_w) \le (z-2)d$, for some pair ℓ_u, ℓ_w. Then, by the triangular inequality, $d_{uw} \le d(s_u, s_w) + \ell_w \le zd$, and similarly $d_{wu} \le zd$. But then ℓ_u and ℓ_w are not z-independent, which is a contradiction.

Let S' be the set of senders of links in S. Let $Z = (z-2)d/2$. The dispersion claim implies that S' is a Z-packing. The definition of a doubling metric implies that

$$\mathcal{P}(B(x, tZ), Z) \le C t^A, \tag{2}$$

where $A = \dim_A(\mathcal{U}, d)$ is the doubling dimension of the metric, C is an absolute constant, and x is any point. Namely, any packing of balls of radius Z inside a ball of radius tZ contains at most Ct^A balls.

Let g be a number. Let x be a sender in S', belonging to link ℓ_x. Let $S_g = \{y \in S' | d(x,y) < gZ\}$. The balls of radius Z centered at points in S_g are all contained within the ball $B(x, (g+1)Z)$ of radius $(g+1)Z$ around x. By z-independence, $S_2 = \emptyset$. The packing bound (2) implies that g

$$|S_g| \le C(g+1)^A \ . \tag{3}$$

Let $T_g = S_g \setminus S_{g-1}$. The senders in T_g are of distance at least $(g-1)Z$ from x, and $\ell_x \le 2d$, so the affectance of each sender y in T_g on ℓ_x is at most

$$a_y(x) = \frac{1/d_{yx}^\alpha}{1/\ell_x^\alpha} \le \left(\frac{2d}{(g-1)Z}\right)^\alpha = \left(\frac{4}{(g-1)(z-2)}\right)^\alpha \ .$$

Observe that

$$\frac{1}{(g-1)^\alpha} - \frac{1}{g^\alpha} = \frac{g^\alpha - (g-1)^\alpha}{g^\alpha (g-1)^\alpha} \le \frac{\alpha g^{\alpha-1}}{g^\alpha (g-1)^\alpha} < \frac{\alpha}{(g-1)^{\alpha+1}} \ .$$

Then,

$$\begin{aligned}
a_S(x) &= \sum_{g>2} a_{T_g}(x) \\
&\le \sum_{g>2} |S_g \setminus S_{g-1}| \cdot \left(\frac{4}{(g-1)(z-2)}\right)^{\alpha} \\
&= \left(\frac{4}{z-2}\right)^{\alpha} \sum_{g>2} |S_g| \left(\frac{1}{(g-1)^{\alpha}} - \frac{1}{g^{\alpha}}\right) \\
&\le \left(\frac{4}{z-2}\right)^{\alpha} \sum_{g>2} |S_g| \frac{\alpha}{(g-1)^{\alpha+1}} \; .
\end{aligned}$$

Observe that for $g \ge 3$, we have using (3) that

$$\frac{|S_g|}{(g-1)^{\alpha+1}} \le \frac{C(g+1)^A}{(g-1)^{\alpha+1}} \le \frac{2^A C}{(g-1)^{\alpha+1-A}} \; .$$

Thus, continuing,

$$a_S(x) \le \left(\frac{4}{z-2}\right)^{\alpha} \cdot \alpha 2^A C \cdot \zeta(\alpha+1-A) = O(1/z^{\alpha}),$$

where $\zeta(x) = \sum_{t\ge1} \frac{1}{t^x}$ is the zeta-function, which is well-defined for any $x > 1$.

Hence, if we choose z so that $(z-2)^{\alpha} \ge 8^{\alpha}\alpha \cdot \zeta(1+\alpha-A)$, or $z \ge 8(\alpha \cdot \zeta(1+\alpha-A))^{1/\alpha} + 2$, we get that $a_S(x) \le 1$ so S is SINR-feasible. □

The two preceding lemmas imply the following result. Lemma 3.2 shows that when q is sufficiently large, any coloring of $G_q(L)$ gives a SINR-feasible schedule, while Lemma 3.1 gives a matching lower bound on the optimal solution.

Theorem 3.3 *Let L be a set of nearly equilength links, and let q be appropriately chosen constant. A coloring of $G_q(L)$ with $O(\Delta(G_q(L)))$ yields a uniform-power schedule that is a constant approximation of PC-Scheduling of L.*

We can handle links of arbitrary lengths by partitioning them into groups, where lengths of links in each group differ by a factor of at most 2. A simple approach is to schedule each group separately using Theorem 3.3. We can choose an arbitrary fixed power to apply to each length class, or modify the powers within each class up to a constant factor. Thus, we can apply any length-consistent power assignment.

Let $g(L)$ denote the *length diversity* of the link set L, or the number of length groups. Note that $g(L) \le \log \Lambda$.

Theorem 3.4 *The PC-Scheduling problem is $O(g(L))$-approximable, using any oblivious power assignment.*

Moscibroda and Wattenhofer [16] showed that uniform and linear power scheduling can be highly suboptimal, and Moscibroda, Oswald and Wattenhofer [15] showed that they can can be as far as n or $\Omega(g(L))$ from optimal.

3.2 Unit Disc Graphs and SINR Scheduling

We can represent the link graph $G_q = G_q(L)$ approximately with a unit-disc graph (UDG). Suppose the links have lengths in the range $[d, 2d)$. Let G'_q be the UDG formed by the points r_v, for $\ell_v \in L$, with radius $\frac{q}{2} \cdot d$. We find that the link graphs and UDGs are closely related, in G'_q that pairs of graphs of one type sandwich graphs of the other type.

Lemma 3.5 *For any $q \geq 1$ and any linkset L, $G'_q \subseteq G_{q+1}$ and $G_q \subseteq G'_{2(q+1)}$.*

Proof: Let v and w be neighbors in G'_q. Then, $d(r_v, r_w) \leq q \cdot d$. Thus, $d_{vw} \leq \ell_v + d(r_v, r_w) \leq (q+1)\ell_v$, and similarly $d_{wv} \leq (q+1)\ell_w$. Hence, $d_{vw} \cdot d_{wv} \leq (q+1)^2 \ell_v \ell_w$, so ℓ_v and ℓ_w are neighbors in G_{q+1}.

On the other hand, suppose we have neighbors ℓ_u and ℓ_w in G_q. Notice that $d(r_u, r_w) \leq d_{uw} + \ell_v \leq d_{uw} + 2d$, and similarly $d(r_u, r_w) \leq d_{wu} + 2d$. Then,

$$(d(r_u, r_w) - 2d)^2 \leq d_{uw} \cdot d_{wu} \leq q^2 \ell_v \ell_w \leq (2qd)^2 \ .$$

Thus, $d(r_u, r_w) \leq 2(q+1)d$. Hence, ℓ_u and ℓ_w are neighbors in $G'_{2(q+1)}$. □

The scheduling problem reduces then, within a constant factor, to the coloring of UDGs. Using Lemma 3.5, we can now simply find an ordinary graph coloring of the UDG graph $G'_{2(q+1)}$, where q is as in Theorem 3.3. Any minimal coloring suffices, leading to an easy online algorithm.

Theorem 3.6 *Applying an algorithm that colors unit disc graphs with $O(\Delta(G))$ colors yields a constant approximation for PC-Scheduling on nearly equilength links. In particular, there is an online algorithm that is constant competitive on nearly equilength links and $O(\log \Lambda)$-competitive in general.*

We can design a distributed algorithm, with UDG coloring as a primitive. The links are divided into length classes, where in class i we have links of length $[d_i, 2d_i)$. The simplest method would be to schedule the length classes independent of each other. Then, the scheduling problem becomes a straightforward reduction to coloring of UDGs.

Proposition 3.7 *A distributed $O(g(L))$-approximation to PC-Scheduling can be obtained in complexity equivalent to the coloring problem on unit-disc graphs.*

4 Oblivious power assignments

We present in this section an scheduling algorithm using a certain oblivious power assignment that achieves a ratio of $O(\log \log \Lambda \cdot \log n)$. In the bidirectional setting, the algorithm obtains an improved $O(\log n)$-ratio, as shown in Sec. 4.1. We also give a construction that shows a $\Omega(\log \log \Lambda)$-separation between the lengths of optimal schedules with or without oblivious power assignments.

We consider the *mean* power assignment (or, square-root assignment [7]) given by $P_v = \ell_v^{\alpha/2}$. The affectance of link ℓ_w on link ℓ_v under mean power assignment is

$$a_w(v) = \frac{P_w / d_{wv}^\alpha}{P_v / \ell_v^\alpha} = \left(\frac{\ell_w}{\ell_v}\right)^{\alpha/2} \left(\frac{\ell_v}{d_{wv}}\right)^\alpha = \left(\frac{\sqrt{\ell_v \ell_w}}{d_{wv}}\right)^\alpha \ .$$

The following observation motivates the consideration of this power assignment.

Observation 4.1 *Suppose $d_{wv} = d_{vw}$, for two links ℓ_v, ℓ_w. Then, $a_w(v) = a_v(w)$ iff we use mean power assignment.*

We say that a set S of links is *well-separated* if any pair of links differ by a factor that is either less than 2 or greater than $8n^{2/\alpha}$. We say that a link ℓ_v and ℓ_w are *τ-close* under mean power assignment if, $\max(a_v(w), a_w(v)) \ge \tau$.

τ-close

The key observation that we make is that each link affects (or is affected by) very few links that are of widely different length. We can then treat those affectance relationships in a graph-theoretic manner.

Lemma 4.2 *Let Q be a well-separated SINR-feasible set of links, and let ℓ_v be a link that is shorter than the links in Q (by a factor of at least $n^{2/\alpha}$). Suppose all the links in Q are $\frac{1}{2n}$-close to ℓ_v under mean power assignment. Then, $|Q| = O(\log\log\Lambda)$.*

Proof: Let Q' be a maximum 3^α-signal subset of Q. By Prop. 2.2, $|Q'| \ge |Q|/9^\alpha$. Q' consists of two types of links: those that affect ℓ_v by at least $\frac{1}{2n}$ under mean power, and those that are affected by ℓ_v by that amount. We shall consider the former type; the argument is nearly identical for the latter type, and will be omitted.

Consider a pair $\ell_w, \ell_{w'}$ in Q' that affect ℓ_v by at least $\frac{1}{2n}$, and suppose without loss of generality that $\ell_w \ge \ell_{w'}$. Thus, $\sqrt{\ell_v\ell_w}^{\alpha} \ge d_{wv}^{\alpha} \cdot \frac{1}{2n}$, which implies that $d_{wv} \le \sqrt{\ell_v\ell_w}(2n)^{1/\alpha}$. Similarly, $d_{w'v} \le \sqrt{\ell_v\ell_{w'}}(2n)^{1/\alpha}$. By the triangular inequality we have that

$$d_{w'w} \le d(s_{w'}, r_v) + d(r_v, s_w) + d(s_w, r_w) = d_{w'v} + d_{wv} + \ell_w \le \ell_w + 2^{1+1/\alpha} n^{1/\alpha} \sqrt{\ell_v\ell_w} \le 3\ell_w,$$

where the last inequality holds because $\ell_w \ge 8n^{2/\alpha}\ell_v$ and $\alpha \ge 1$. Similarly,

$$d_{ww'} \le d_{vw} + d_{vw'} + \ell_{w'} \le \ell_{w'} + 2^{1+1/\alpha} n^{1/\alpha} \sqrt{\ell_v\ell_w}\ .$$

Multiplying together, we obtain that

$$d_{w'w} \cdot d_{ww'} \le 3\ell_{w'}\ell_w + 12n^{1/\alpha}\sqrt{\ell_v\ell_w}\ell_w\ .$$

However, since Q' is a 3^α-signal set, $d_{w'w} \cdot d_{ww'} \ge 9\ell_w\ell_{w'}$. By combining the last two inequalities and cancelling a $6\ell_w$ factor, we have that

$$\ell_{w'} \le 2n^{1/\alpha}\sqrt{\ell_v\ell_w}\ . \tag{4}$$

Note that (4) implies that $\ell_w \ge 2\ell_{w'}$, and thus, by well-separation, that $\ell_w \ge 8n^{2/\alpha}\ell_{w'}$.

Label the links in Q' by $\ell_1, \ell_2, \ldots, \ell_t$ in increasing order of length. Equation (4) implies that

$$\ell_{i+1} \ge \frac{\ell_i^2}{\ell_v n^{2/\alpha}} \ge 2\frac{\ell_i^2}{\ell_1},$$

for any $i = 2, 3, \ldots, t$. Thus, if we let $\lambda_i = \ell_i/\ell_1$, we get that $\lambda_{i+1} \ge 2\lambda_i^2$, and by induction that $\lambda_t \ge 2^{2^{t-1}-1}$. Hence, $|Q'| = t \le \lg\lg\lambda_t + 2 = \lg\lg\Lambda + 2$, and the lemma follows. □

We find that links of widely different lengths can be scheduled together rather easily.

Lemma 4.3 *Let L be a set of links partitioned into length groups $L_1, L_2, \ldots, L_t$ such that links in the same group differ by a factor of at most 2 but links in different groups differ by a factor of at least n^2. Suppose each group L_i has been scheduled with uniform power using Γ_i slots. Then, there is an algorithm that produces a combined schedule of L with mean power assignment using $O(\log\log\Lambda \cdot \max_i \Gamma_i)$ slots.*

Proof: Let $p = O(\log\log\Lambda)$ denote the bound of Lemma 4.2 on the size of the set Q. First, we transform the schedules of the length groups into f-signal schedules, where $f = 2^{\alpha/2+1}$. By Proposition 2.2, this stretches each schedule by a factor of at most $(f+1)^2$. Let S_i be some slot in the resulting schedule for L_i, for $i = 1, 2, \ldots, t$. We show that $S = \cup_i S_i$ can be scheduled in $p + 1$ slots, resulting in a total of $(p+1) \cdot (f+1)^2 \cdot \max_i \Gamma_i$ slots for L. This yields the claimed result.

Process the links in S in decreasing order of length, and consider a link ℓ_v. By Lemma 4.8, there are at most p longer links ℓ_w in S that are $\frac{1}{2n}$-close to ℓ_v. Assign ℓ_v to a slot T_j, $j \in \{1, 2, \ldots, p+1\}$, that does not contain a $\frac{1}{2n}$-close link. This completes the specification of the algorithm.

It remains to argue that this assignment yields a SINR-feasible schedule. Consider a link ℓ_v in slot T_j, that originally came from slot S_k. The affectance $a_{S_k \cap T_j}(\ell_v)$ by links of nearly equal length in T_j is at most $1/f$ by the f-signal property. Changing the power assignment in the length group S_k from uniform to mean power assignment increases affectance by at most a factor of $2^{\alpha/2}$, for a total of $2^{\alpha/2}/f = 1/2$. The affectance of all other links (from different length classes) in T_j on ℓ_v is at most $1/(2n)$ each, by construction, or at most $1/2$ in total. Hence, the total affectance is at most one, resulting in a SINR-feasible schedule. □

We obtain an algorithm that processes the length groups in decreasing order, greedily assigning the links to the first class in which it affects no link by a non-trivial amount.

Proposition 4.4 *Suppose there exists a ρ-approximate algorithm for PC-Scheduling on equi-length links that uses uniform power. Then, there exists a $O(\rho \cdot \log\log\Lambda \cdot \log n)$-approximate algorithm for PC-Scheduling that uses mean power assignment.*

Proof: Given a set L of links, divide L into length groups $S_1, S_2, \ldots$, such that $S_i = \{\ell_v \in S | \ell_v \in [2^{i-1}\ell_{min}, 2^i \ell_{min})\}$, where ℓ_{min} denotes the length of the shortest link in S. Then, partition S into classes $B_i = \cup_j S_{i+j\cdot\frac{2}{\alpha}\log n}$, for $i = 1, 2, \ldots, \frac{2}{\alpha}\log n$. The theorem follows from applying Lemma 4.3 on each class B_i separately. □

Using our result of Thm. 3.6 on equi-length links, we obtain the following result.

Theorem 4.5 *PC-Scheduling is $O(\log\log\Lambda \cdot \log n)$-approximable in fading metrics.*

For general metrics, we obtain approximations with higher logarithmic factors. Fanghänel et al [7] gave an algorithm using linear power assignment that yields a schedule of length $O(\log\Lambda \log n(\Gamma(L) + \log^2 n))$. On nearly equilength links, this implies a schedule length of $O(\log n(\Gamma(L) + \log^2 n))$. They also gave a simpler algorithm with a ratio of $O(\log^2 n)$.

Corollary 4.6 *PC-Scheduling is $O(\log\log\Lambda \cdot \log^3 n)$-approximable in general metrics, and within a factor of $O(\log\log\Lambda \cdot \log^2 n)$ when $\Gamma(L) = \Omega(\log^2 n)$.*

We obtain as corollary, a relationship between schedule length and the chromatic number of the link graph.

Corollary 4.7 *For any link set L, there is a schedule using $O(\log\log\Lambda \cdot \log n)\chi(G_p)$ slots in fading metrics.*

4.1 Bidirectional scheduling

In the bidirectional variant introduced by Fanghänel et al [6], a stronger separation criteria applies, since communication along each link can occur in either direction. The distance between two links is now the shortest distance between any endpoints of the links. Thus, $d_{uv} = d_{vu} = \min(d(r_v, r_u), d(r_v, s_u), d(s_v, s_u), d(s_v, r_u))$. Other definitions are unchanged.

We can obtain a better approximation ratio for this problem, with essentially the same algorithm, via the following stronger version of Lemma 4.2.

Lemma 4.8 *Let S be a set of links inducing an independent set in G_p and let ℓ_v be a link. Then, there is at most one link ℓ_w in S with $\ell_w \geq n^{2/\alpha} \cdot \ell_v$ such that $a_v(w) \geq 1/(2n)$ under mean power assignment and bidirectional measure.*

Proof: Suppose the lemma is false and let $\ell_w, \ell_{w'}$ be two links in S that are longer than $n^{2/\alpha}$ times ℓ_v and affect it by at least $1/(2n)$ each. Suppose without loss of generality that $\ell_w \geq \ell_{w'}$. The assumption of affectance under mean power assignment implies that

$$\left(\frac{\sqrt{\ell_v \ell_u}}{d_{vu}}\right)^{\alpha} \geq 1/(2n),$$

for $u \in \{w, w'\}$. Thus, $d_{vu} \leq (2n)^{1/\alpha}\sqrt{\ell_v \ell_u}$. In the bidirectional case, $d_{vu} = d_{uv}$. Thus, by the triangular inequality, we have that

$$d_{ww'} \leq d_{wv} + d_{vw'} \leq 2(2n)^{1/\alpha}\sqrt{\ell_v \ell_w} \leq 2(2n)^{1/\alpha}\sqrt{(\ell_{w'}/n^{2/\alpha})\ell_w} \leq 2^{1+1/\alpha}\sqrt{\ell_{w'}\ell_w} \ .$$

Then, ℓ_w and ℓ'_w are not independent in G_p, for $p \geq 2$, which contradicts our assumption. □

The rest of the argument is identical to the unidirectional case.

Theorem 4.9 *Suppose there exists a ρ-approximate algorithm for the bidirectional scheduling problem on equi-length links. Then, there exists a $O(\rho \log n)$-approximate algorithm for the general bidirectional problem. In particular, there is a $O(\log n)$-approximation algorithm for bidirectional scheduling in fading metrics.*

We obtain as corollary a connection between length of bidirectional schedules and the chromatic number of the link graph representation.

Corollary 4.10 *For any link set L with link graph G_p, there is a bidirectional schedule using $O(\log n)\chi(G_p)$ slots in fading metrics.*

4.2 Construction

We can show that the bound obtained is best possible for oblivious power functions, up to the logarithmic factor. The following result follows also from the constructions in [7] by analyzing the dependence on Λ.

Theorem 4.11 *For any length-consistent power function ϕ, there is a SINR-feasible instance for which any schedule under ϕ requires $\Omega(\log \log \Lambda)$ slots.*

Proof: We break this into two cases, first considering power functions that grow no faster than the mean assignment. We implicitly assume that power assignment functions either grow no faster or no slower than mean assignment. Let $t = 2$. Let $a_i = t^{2^0} + t^{2^1} + \cdots + t^{2^i} = \sum_{j=0}^{i} t^{2^i}$. Consider the following set of links $L = \{\ell_1, \ell_2, \ldots, \ell_n\}$ located on the real line, where $\ell_i = t^{2^i}$.

Let ℓ_0 denote $t^{2^0} = t$. Position the receiver r_i of ℓ_i at location $+a_{i-1}$ and the sender s_i at location $-(\ell_i - a_{i-1})$. Observe that $\sqrt{\ell_i} = \ell_{i-1}$. Note that $d_{i-1,i} = \ell_{i-1} + \ell_{i-1} = 2\ell_{i-1}$, and more generally $d_{j,i} = (\ell_j - a_{j-1}) + a_{i-1} \leq 2\ell_{i-1}$, and $d_{ji} > \ell_{i-1}$ for $j < i$. Also, $d_{i,j} = \ell_i - a_{i-1} + a_{j-1} > \ell_i/2$.

Under mean power assignment, we have that for $j < i$,

$$a_j(i) = \frac{\sqrt{\ell_i \ell_j}}{d_{ji}} \geq \frac{\ell_{i-1}}{2\ell_{i-1}} = \ell_{j-1} > 1 \ .$$

Hence, in any schedule based on the mean assignment, the n links must be assigned to different slots. Observe that this holds for any power function that grows no faster than mean assignment, and can be extended to functions that grow as ℓ_c, $c < \alpha$.

Consider instead the oblivious power assignment function $P_v = \ell_v^\alpha / \log \ell_v$. Under this power function, for $j < i$,

$$a_j(i) = \frac{P_j/d_{ji}^\alpha}{P_i/\ell_i^\alpha} = \frac{\ell_j^\alpha}{d_{ji}^\alpha \log d_j} \cdot \log \ell_i = \frac{\ell_j^\alpha 2^{i-j}}{d_{ji}^\alpha} \leq \frac{\ell_j^\alpha 2^{i-j}}{\ell_{i-1}^\alpha} = t^{2^j - 2^{i-1}} \cdot 2^{i-j}.$$

Observe that $\sum_{j=2}^{i-2} a_j(i) < 1$. Also, for $i > j$,

$$a_i(j) = \frac{P_i/d_{ij}^\alpha}{P_j/\ell_j^\alpha} = \frac{\ell_j^\alpha}{d_{ij}^\alpha \log d_j} \cdot \log \ell_j = \frac{\ell_i^\alpha 2^{j-i}}{d_{ij}^\alpha} \leq \frac{\ell_i^\alpha 2^{j-i}}{\ell_i^\alpha/2} = 2^{j-i+1}.$$

So, again $\sum_{i=j+2}^{n} a_i(j) < 1$. It follows that the set of links can be partitioned into two SINR-feasible sets: the even and the odd numbered. Hence, for mean power assignment, and any function growing no faster than it, throughput is only $n = O(\log\log\Lambda)$, while there are assignments with throughput $\Omega(1)$.

Consider now the complementary instance, where the direction or the role of senders and receivers, has been reversed. Then, nearly identical computation shows that any function that grows no slower than mean assignment can only schedule a single link in a single slot. On the other hand, using power assignment $f(\ell_v) = \lg \ell_v$, two slots suffice.

Finally, we can combine the two constructions into a single instance that is hard to schedule for all oblivious power functions, by taking disjoint copies that are sufficiently separated in space.

□

5 Conclusions

Another way of evaluating the communication throughput is to bound the capacity: how many links can be simultaneously scheduled. This is also known as the Single-Shot Scheduling problem. It can be verified that all of the results presented here give identical bounds on capacity.

We have encountered a fundamental asymmetry property that makes scheduling hard, at least with oblivious strategies. If extreme levels of asymmetry are disallowed, either in the input or in the resulting schedule, then we obtain a simple logarithmic approximation.

The main open question is whether a polylogarithmic approximation is possible for the (uni-directional) scheduling problem. From a practical perspective, it would be interesting if the logarithmic factor could be removed, giving a $O(\log\log\Lambda)$-approximation. Alternatively, non-oblivious power strategies that could be implemented in a distributed setting would be interesting.

Acknowledgement

The author is grateful to Roger Wattenhofer for posing the questions of scheduling complexity in the physical model, as well as for helpful comments and advice.

References

[1] M. Andrews and M. Dinitz. Maximizing capacity in arbitrary wireless networks in the sinr model: Complexity and game theory. In *INFOCOM*, April 2009.

[2] C. Avin, Z. Lotker, and Y. A. Pignolet. On the power of uniform power: Capacity of wireless networks with bounded resources. In *these proceedings*, 2009.

[3] D. Chafekar, V. Kumar, M. Marathe, S. Parthasarathy, and A. Srinivasan. Cross-layer Latency Minimization for Wireless Networks using SINR Constraints. In *Mobihoc*, 2007.

[4] K. L. Clarkson. Nearest neighbor queries in metric spaces. *Discrete and Computational Geometry*, 22:63–93, 1999.

[5] T. A. ElBatt and A. Ephremides. Joint scheduling and power control for wireless ad hoc networks. *IEEE Transactions on Wireless Communications*, 3(1):74–85, 2004.

[6] A. Fanghänel, T. Keßelheim, H. Räcke, and B. Vöcking. Oblivious interference scheduling. In *PODC*, August 2009.

[7] A. Fanghänel, T. Keßelheim, and B. Vöcking. Improved algorithms for latency minimization in wireless networks. In *ICALP*, July 2009.

[8] O. Goussevskaia, M. M. Halldórsson, R. Wattenhofer, and E. Welzl. Capacity of Arbitrary Wireless Networks. In *INFOCOM*, April 2009.

[9] O. Goussevskaia, Y. A. Oswald, and R. Wattenhofer. Complexity in Geometric SINR. In *Mobihoc*, pages 100–109, 2007.

[10] J. Gronkvist and A. Hansson. Comparison between graph-based and interference-based STDMA scheduling. In *Mobihoc*, pages 255–258, 2001.

[11] P. Gupta and P. R. Kumar. The Capacity of Wireless Networks. *IEEE Trans. Information Theory*, 46(2):388–404, 2000.

[12] M. M. Halldórsson and R. Wattenhofer. Wireless Communication is in APX. In *ICALP*, July 2009.

[13] J. Heinonen. *Lectures on Analysis in Metric Spaces.* Springer, 1999.

[14] R. Maheshwari, S. Jain, and S. R. Das. A measurement study of interference modeling and scheduling in low-power wireless networks. In *SenSys*, pages 141–154, 2008.

[15] T. Moscibroda, Y. A. Oswald, and R. Wattenhofer. How optimal are wireless scheduling protocols? In *INFOCOM*, pages 1433–1441, 2007.

[16] T. Moscibroda and R. Wattenhofer. The Complexity of Connectivity in Wireless Networks. In *INFOCOM*, 2006.

[17] T. Moscibroda, R. Wattenhofer, and Y. Weber. Protocol Design Beyond Graph-Based Models. In *Hotnets*, November 2006.

[18] T. Moscibroda, R. Wattenhofer, and A. Zollinger. Topology Control meets SINR: The Scheduling Complexity of Arbitrary Topologies. In *Mobihoc*, 2006.