Wireless Scheduling with Power Control*

Magnús M. Halldórsson†

September 27, 2010

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 $O(\log n \cdot \log\log \Delta)$, where n is the number of links and Δ 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 give evidence that this dependence on Δ is unavoidable, showing that any reasonably-behaving oblivious power assignment results in a $\Omega(\log\log \Delta)$-approximation. We also give a simple *online* algorithm that yields a $O(\log \Delta)$-approximation, by a reduction to the coloring of unit-disc graphs.

These results hold also for the (weighted) capacity problem of finding a maximum (weighted) subset of links that can be scheduled in a single time slot. 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

*A preliminary version of this paper appears in the proceedings of the 17^{th} *Annual European Symposium on Algorithms (ESA'09)*[23].

†School of Computer Science & Icelandic Center of Excellence in Theoretical Computer Science (ICE-TCS), Reykjavik University, 101 Reykjavik, Iceland. 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 Section 2. The SINR model reflects physical reality more accurately and is therefore often simply called the physical model. On the other hand, most research in the SINR model has focused on 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.

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 set of links in each slot satisfies the SINR-constraints. We refer to this as the PC-Scheduling problem. We also consider two closely related throughput maximization problems, both with power control. In the PC-Capacity problem, we seek a maximum cardinality subset of links satisfying the SINR constraints, while in the PC-Weighted-Capacity problem, the links have given weights and we seek to maximize the total weight of a feasible subset. Finally, we also touch on the bidirectional setting, where both nodes in a link may be transmitting, implying a stronger, symmetric form of interference.

For reasons of simplicity of use, it is strongly 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 essential 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 clear 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 simple scheduling algorithms that work for any oblivious power assignment strategy, resulting in a $O(\log \Delta)$ approximation ratio, where Δ is the ratio between the maximum and minimum link length. This holds for PC-Scheduling, PC-Capacity and PC-Weighted-Capacity. In particular, when all links are of nearly equal length, we obtain the first constant approximation. This matches the constructions in [32] that show that both fixed and linear assignments can be as much as $\Omega(\log \Delta)$ factor from optimal. 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 $O(\log \Delta)$-competitive online scheduling algorithm, and randomized $O(\log \Delta)$-competitive online capacity algorithm.

To tackle instances where Δ is very large, we analyze closely conflicts between links of widely different lengths. We obtain algorithms for all three problems that attain a $O(\log\log \Delta \cdot \log n)$-approximation, an exponential improvement in terms of Δ. Under the natural assumption that lengths can be represented in binary, this implies also $O(\log^2 N)$-approximation, where N is the length of the input. This uses a recently introduced oblivious power assignment, *mean* (or *square root* [15]) power. In the bidirectional setting, we obtain a $O(\log n)$-approximation, improving on the previous $O(\log^c n)$-factor for PC-Scheduling with $c > 5$ [14] using considerably simpler arguments.

All the results in this paper are for oblivious power assignments, which are of special practical utility. We show that the dependence on Δ is likely to be unavoidable for oblivious power assignment strategies. Namely, any smoothly growing oblivious power function forces $\Omega(\log\log\Delta)$-approximate schedules.

We have generalized the setting from the plane to the class of *doubling metrics*, the first work to do so. The requirement is that the path-loss constant α be greater than the so-called doubling constant of the metric, which is equivalent in the plane to the standard assumption that $\alpha > 2$ (see Section 2). 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.

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. As mentioned, the basic class of unit-disc graphs suffices for links of nearly equal length. The $O(\log n \cdot \log\log\Delta)$-approximation result is also relative to the underlying graph.

The current paper refines the results and the arguments in the earlier draft [23], and adds to it approximations of PC-Weighted-Capacity. Additionally, the draft [23] contained a faulty lemma (Lemma 4.3), which is corrected here by proving the main results in Section 4 differently. The new treatment involves an extension of a graph-theoretic approximation property, which may be of independent interest.

1.2 Related Work

Most work in wireless scheduling in the physical (SINR) model has been of heuristic nature, e.g. [11]. Only after the work of Gupta and Kumar [22] 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 both theoretically and experimentally [21, 30, 35].

Algorithmic work in the SINR model started just four years ago with the seminal work of Moscibroda and Wattenhofer [33]. In this paper, Moscibroda and Wattenhofer 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 [17, 36], sensor networks [31], and combined scheduling and routing [8]. However, arbitrary networks are beyond the scope of these papers. Apart from these papers, algorithmic SINR results also started showing up here and there, for instance in a game theoretic context or a distributed algorithms context, e.g., [4, 5, 7, 19, 37, 27].

Approximation algorithms for the problem of scheduling wireless links with power control in the SINR model were given in [36], [32] and [8]. 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 $O(\log\Delta)$-like approximation. Fanghänel, Kesselheim and Vöcking [15] gave a randomized algorithm for the scheduling problem using linear power assignment that uses $O(OPT\log\Delta + \log^2 n)$ slots, matching our results for dense instances. Andrews and Dinitz [2] gave a $O(\log\hat{\Delta})$-approximation for the Capacity problem, where $\hat{\Delta} \geq \Delta$ is the ratio between the longest link length and the shortest distance between any pair of points in the instance. Avin, Lotker and Pignolet [6] show that the assumption of $\alpha > 2$ used by all previous work may not be necessary, in that the ratio between optimal non-oblivious and oblivious capacity is $O(\log\Delta)$, at least in the 1-dimensional metric. Finally, the earlier work of Goussievskaia, Oswald and Wattenhofer [20] featured a $O(\log \Delta)$-factor approximation for both the scheduling and the capacity problems, but in comparison with optimal solutions that are constrained to use uniform power assignment.

In [14], Fanghänel et al. gave a construction that shows that any schedule based on any oblivious power assignment can be a factor of n from optimal. They also introduced the bidirectional version of the scheduling problem and give a $O(\log^{3.5+\alpha} n)$-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 better understood. Constant factor approximation for the corresponding capacity problem in the plane was given in [18], yielding a $O(\log n)$-approximation for the scheduling problem. Both of these problems are known to be NP-complete [20]. The results obtained here for power control build on and extend the techniques and properties derived in the case of uniform power in [18, 25].

In developments since the original presentation of this work [23], Erlebach and Grant [12] gave a $O(\log \Delta)$-factor approximation algorithm for the problem of multicast scheduling, where each transmission is to be sent to a collection of receivers. Their work uses in a fundamental way the results of the current paper on nearly-equilength links and unit-disc graphs. Fanghänel et al. [13] studied the online version of PC-Capacity problem, obtaining a tight bound of $\theta(\Delta^{d/2})$ on the competitive ratio of deterministic algorithms in d-dimensional Euclidean space.

In a breakthrough, Kesselheim [28] has very recently obtained a $O(1)$-approximation algorithm for PC-Capacity. It necessarily uses instance-specific power assignment, and the question of optimal schedules using oblivious power assignment remains interesting both from a theoretical and practical viewpoint. Halldórsson and Mitra [24] have generalized our results for PC-Capacity to arbitrary metric spaces. They additionally obtained the improved approximation factors of $O(\log n + \log\log \Delta)$ and $O(1)$ in the uni-directional and bi-directional cases, respectively. For PC-Scheduling with oblivious power, however, our approximation factors are still the best known.

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 ℓ_v to link ℓ_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.

The nodes can transmit with different power. Let P_v denote the power assigned to link ℓ_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 α 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:

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

where N is the ambient noise, β 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 N can be made arbitrarily small, thus we ignore this term. Of course, in real situations, there are upper bounds on maximum power which we ignore here. We shall also assume that $\beta \geq 3^\alpha$; by the signal-strengthening results of [25], this can only affect the constants in the approximation results. We say that S is *SINR-feasible* if (1) is satisfied for each link ℓ_v in S.

This paper deals with *power control*, i.e., determining the power assignment to the links is a part of the problem. In particular, we focus on *oblivious power* assignments, where the power depends only on the length of the link, while we compare it to an optimal solution that is free to use any power assignment. The most basic assignment is *uniform power*, where each link ℓ_v uses the same power $P_v = P$. Another common oblivious assignment is *linear power*, where $P_v = \ell_v^\alpha$. We will focus on uniform power, along with another oblivious assignment, the *mean* ((or, square-root [15]) power $\mathcal{M}$ given by $P_v = \mathcal{M}_v = \ell_v^{\alpha/2}$.

The *affectance* of link ℓ_v caused by a set S of links [18, 25] under a given power assignment P, is the sum of the interferences of the links in S on ℓ_v relative to the power received, or

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

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, $a_{S_1 \cup S_2}(\ell_v) = a_{S_1}(\ell_v) + a_{S_2}(\ell_v)$. Observe that a set S is SINR-feasible iff $a_S(\ell_v) \leq 1/\beta$, for each link $\ell_v \in S$.

A *p-signal* set or a schedule is one where the affectance of any link is at most $1/p$, with respect to the given power assignment. 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. Let Δ denote the ratio between the maximum and minimum length of a link.

For a graph G, let $\chi(G)$ denote its chromatic number, and $\alpha(G)$ its independence number (or the maximum cardinality of a subset of mutually non-adjacent vertices). Define the neighborhood $N(v)$ of a vertex v to be the set consisting of v's neighbors, and the closed neighborhood $N[v]$ to include v as well. For a vertex subset S, let $G[S]$ denote the subgraph induced by S.

We say that a collection of links is *q-independent* if any two of them, ℓ_v and ℓ_w, satisfy the constraint

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

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.

Lemma 2.1 *If S is a q^α-signal set under some power assignment, then S is q-independent.*

Proof: Let P be a power assignment for which S is a q^α-signal set. Since the links belong to the same p-signal set, for $p = q^\alpha$, 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 $d_{vw} \cdot d_{wv} \geq p^{2/\alpha} \cdot \ell_w \ell_v = q^2 \cdot \ell_w \ell_v$. □

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.

Notation	*Meaning*	*Topic*	*Page*
ℓ_v	Link $\ell_v = (s_v, r_v)$; denotes also its length		4
P_v	Power assigned to link ℓ_v		4
$\mathcal{M}_v$	Mean power assignment $\mathcal{M}_v = \ell_v^{\alpha/2}$		5
α	Path loss constant (signal decay exponent).	*SINR*	4
β	SINR requirement (assumed to be at least 3^α).		4
$d(x, y)$	Distance between points x and y.		4
d_{vw}	$= d(s_v, r_w)$		4
$a_S(\ell_v)$	Affectance of linkset S on link ℓ_v		5
$a_w(v)$	$= a_{\{\ell_w\}}(\ell_v)$		5
OPT_p	Optimal p-signal schedule	*Analysis*	5
Δ	Ratio of longest to shortest link length		5
$\zeta(x)$	Riemann zeta-function.		7
q-independent	$d_{vw} \cdot d_{wv} \geq q^2 \cdot \ell_v \ell_w$		5
t-close	$\max(a_v(w), a_w(v)) \geq t$		12
well-separated	Linklengths differ by factor ≤ 2 or $\geq \Lambda$	*Link relationships*	12
τ	$2\beta n$		12
Λ	$2\tau^{1/\alpha}$		12
$\chi(G)$	Chromatic number of graph G		5
$\alpha(G)$	Independence number of graph G		5
$G_q(L)$	q-independence relation on linkset L		5
$U_z(L)$	Unit-disc graph on the senders in L	*Graphs*	8
$G[X]$	Graph induced by vertex subset X		5
$N_G(v)$	Set of neighbors of node v in graph G		5
$N_G[v]$	Closed neighborhood of v, $= N_G(v) \cup \{v\}$		5
$A = \dim_A(\mathcal{U}, d)$	The Assouad (doubling) dimension (2, for $\mathbb{R}^2$)		7
$B(y, \epsilon)$	Ball of radius ϵ centered at y.	*Metrics*	7
C	Constant in doubling dimension definition		7
C'	$= \alpha 2^A(\zeta(\alpha + 1 - A) - 1)$		7
$z_1 = z_1(p)$	$= 2(pC')^{1/\alpha}$. Sufficient sender separation.		7
$z_2 = z_2(p)$	$= p^{1/\alpha} - 1$. Necessary sender separation.		9

Table 1: List of notation

We introduce fading metrics in Sec. 3.1 and show that a well-dispersed set of links in such a metric has good signal properties. In Sec. 3.2, we show that uniform power assignment performs well when links are of nearly equal lengths. The global nature of the problem disappears, and local strategies become sufficient. This results in a $O(\log \Delta)$-approximation algorithms using any oblivious power assignment, which is argued in Sec. 3.3. In fact, our algorithms for PC-Scheduling and PC-Capacity are $O(\log \Delta)$-competitive online algorithms.

3.1 Scheduling in Fading Metrics

We extend the traditional setting from the Euclidean plane to doubling metrics (see Clarkson [9]).

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, i.e., the maximum number of ϵ-balls that can be packed into the body $\mathcal{U}$. The *Assouad dimension* $\dim_A(\mathcal{U}, d)$ [3] (also known as uniform metric dimension or doubling dimension) for a metric space $(\mathcal{U}, d)$ is the value t, if it exists, such that

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

as $\epsilon \to 0$, where C is an absolute constant. It is known that $\dim_A(\Re^k) = k$ for the k-dimensional Euclidean space [26], and in particular for the plane $C = \frac{1}{6}\pi\sqrt{3} \approx 0.907$ [16, 38].

We require that the path loss exponent α be strictly greater than the doubling dimension $A = \dim_A(\mathcal{U}, d)$ of the metric. We shall refer to such a combination of distance metric and path loss constant as a *fading metric*.

The following result extends similar lemmas in previous works (see [18, 25]) from the setting of the Euclidean plane to the more general class of fading 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.

Let $\zeta(x) = \sum_{t \geq 1} \frac{1}{t^x}$ be the Riemann zeta-function, which is well-defined for any $x > 1$. Let $C' = 2^A \alpha(\zeta(\alpha + 1 - A) - 1)$ and let $z_1(p) = 2(pC')^{1/\alpha}$.

Lemma 3.1 *(Far-away lemma) Let p be positive and let S be a set of links whose senders are of mutual distance at least zD, where D is the length of the longest link in S and $z = z_1(p)$. Then, using uniform power assignment, S forms a p-signal set in any fading metric.*

Proof: Let S' be the set of senders of links in S. Let $Z = zD/2$. The separation of the senders implies that S' is a Z-packing. The definition of a doubling metric implies that for any $t > 0$, the packing number of the tZ-ball centered at any point x is bounded by

$$\mathcal{P}(B(x, tZ), Z) \leq Ct^A . \tag{2}$$

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 $S_g = \{y \in S' | d(x, y) < gZ\}$ be the set of senders within distance less than gZ from x, and let $T_g = S_g \setminus S_{g-1}$. By assumption, $S_2 = \emptyset$. The senders in T_g are of distance at least $(g-1)Z$ from x, and $\ell_x \leq D$, so the affectance of each sender y in T_g on ℓ_x is at most

$$a_y(x) = \frac{1/d_{yx}^\alpha}{1/\ell_x^\alpha} \leq \left(\frac{D}{(g-1)Z}\right)^\alpha = \left(\frac{2}{(g-1)z}\right)^\alpha, \quad \forall \ell_y \in T_g.$$

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,

$$\begin{aligned} a_S(x) &= \sum_{g \ge 3} a_{T_g}(x) \\ &\le \sum_{g \ge 3} |S_g \setminus S_{g-1}| \cdot \left(\frac{2}{(g-1)z}\right)^\alpha \\ &= \left(\frac{2}{z}\right)^\alpha \sum_{g \ge 3} |S_g| \left(\frac{1}{(g-1)^\alpha} - \frac{1}{g^\alpha}\right) \\ &\le \left(\frac{2}{z}\right)^\alpha \sum_{g \ge 3} |S_g| \frac{\alpha}{(g-1)^{\alpha+1}} \ . \end{aligned} \tag{3}$$

The balls of radius Z centered at points in S_g are all contained within the ball $B(x, (g+1)Z)$. For $g \ge 3$, the packing bound (2) then implies that $|S_g| \le \mathcal{P}(B(x, (g+1)Z), Z) \le C(g+1)^A$, and thus we have that

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

Continuing from (3),

$$a_S(x) \le \left(\frac{2}{z}\right)^\alpha \alpha \cdot 2^A C \sum_{x \ge 2} \frac{1}{x^{\alpha+1-A}} = \left(\frac{2}{z}\right)^\alpha \alpha \cdot 2^A C(\zeta(\alpha+1-A) - 1) = \left(\frac{2}{z}\right)^\alpha C' \ .$$

Thus, S is an s-signal set, where $s = \frac{1}{C'}(\frac{z}{2})^\alpha = p$. □

Remark 3.2 *Lemma 3.1 does not hold in general for arbitrary distance metrics. In particular, it fails for* $\mathbf{R}^2$ *when* $\alpha \le 2$*. In fact, unit-length links arranged in a grid with separation* q *will be* q^α*-independent, while maximum affectance becomes* $\Omega(\log n)$*.*

3.2 Scheduling Nearly-Equilength Links

We observe here that if the links are of nearly equal length, then we can simplify the many-to-many interference relationships by a pairwise relationship, modulo small constant factors in the approximation. These pairwise relationships correspond to the graphs formed by discs of fixed radius in the plane. With one radius, we capture the necessary distance between any pair of links in a feasible solution, while with another larger radius, we have the sufficient distance so that any set of links of such mutual separation is guaranteed to be SINR-feasible (in any given fading metric). This leads to simple and effective approximation algorithms that can be made online and turned into distributed algorithms.

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 less than 2. The key observation is that we can represent the link graph $G_q = G_q(L)$ of a set L of nearly-equilength links approximately with a unit-disc graph (UDG).

Definition 3.3 *Let* L *be a linkset in a fading metric, and let* d *denote the minimum linklength in* L*. Given a number* z*, the* unit-disc graph $U_z(L)$ *of* L *is the graph with a node for each sender of* L *with two nodes adjacent if the distance between the two senders is less than* $z \cdot d$*.*

That is, $U_z(L)$ is the graph formed by the intersection of balls of radius $zd/2$ that are centered at the senders. We find that the link graphs and UDGs are closely related, in that pairs of graphs of one type sandwich graphs of the other type.

Lemma 3.4 *For any $q \geq 1$ and any nearly-equilength linkset L, $U_{q-1}(L) \subseteq G_q(L)$ and $G_q(L) \subseteq U_{2(q+1)}(L)$.*

Proof: Recall that the links have lengths in the range $[d, 2d)$. Let v and w be neighbors in $U_{q-1}(L)$. Then, $d(s_v, s_w) < (q-1) \cdot d$, by definition. Thus, $d_{vw} \leq d(s_v, s_w) + \ell_w < q\ell_w$, and similarly $d_{wv} < q\ell_v$. Hence, $d_{vw} \cdot d_{wv} < q^2 \ell_v \ell_w$, so ℓ_v and ℓ_w are neighbors in G_q.

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

$$(d(s_u, s_w) - 2d)^2 < d_{uw} \cdot d_{wu} < q^2 \ell_v \ell_w < (2qd)^2 \ .$$

Thus, $d(s_u, s_w) < 2(q+1)d$. Hence, ℓ_u and ℓ_w are neighbors in $U_{2(q+1)}(L)$. □

Read differently, the above lemma implies that sender separation and signal strength of a linkset go hand in hand. Namely, if S is a q-independent set of nearly-equilength links, then the senders in S are of mutual distance at least $(q-1)d$, and thus $U_{q-1}(S)$ is an empty graph (independent set). Conversely, if X is an independent set in unit-disc graph $U_{2(q+1)}(L)$, where L is nearly-equilength linkset, then X is q-independent.

We can now argue our claim that unit-disc graphs capture well nearly-equilength links in fading metrics. Define $z_2 = z_2(p) = p^{1/\alpha} - 1$. We show that a linkset with minimum linklength d and pairwise sender separation of at least $z_1(p) \cdot d$ will be a p-signal set, while any p-signal set must obey a separation of at least $z_2(p) \cdot d$.

Theorem 3.5 *For a set L of nearly-equilength links, any independent set in $U_{z_1}(L)$ is a p-signal linkset, and any p-signal subset of L is an independent set in $U_{z_2}(L)$.*

Proof: By Lemma 3.1, an independent set X in $U_{z_1}(L)$ is a p-signal set. By Lemma 2.1, a p-signal subset S of L is $(p^{1/\alpha})$-independent (i.e., an independent set in $G_{p^{1/\alpha}}$). By Lemma 3.4, it is then an independent set in $U_{p^{1/\alpha}-1}(L)$. □

We note that unit-disc graphs in fading metrics satisfy a *bounded-independence* property as follows. Recall that $\alpha(G)$ is the cardinality of a maximum independent set in G.

Observation 3.6 *Let a and b be given constants, $a \geq b$. Let $U_a = U_a(L)$ and $U_b = U_b(L)$ be unit-disc graphs on the same linkset L but with different radii. Let v be a link in L with closed neighborhood $N_1 = N_{U_a}[v]$ in U_a. Then, $\alpha(U_b[N_1]) \leq C(1+2a/b)^A$.*

Proof: The nodes in an independent set I in U_b form disjoint balls of radius $bd/2$ centered at the senders of the links. All senders of links in N_1 are contained in the ball $B(s_v, ad)$, where d is the minimum link length in L. Thus, all the balls corresponding to I are contained in the larger ball $B(s_v, (a+b/2)d)$. The packing constraint of the metric ensures that a limited number of the smaller disjoint balls fit inside the large ball, implying that $|I| \leq \mathcal{P}(B(s_v, (a+b/2)d), bd/2) \leq C(1+2a/b)^A$. □

Our problems reduces then, within constant factors, to coloring and (weighted) independent sets in UDGs. We say that an independent set in a weighted graph is *greedy* if it is obtained by the iterative process of selecting a vertex whose weight is greater than each of its neighbors', deleting the neighbors, and recursing on the remaining graph.

The following result is immediate from Thm. 3.5 and Obs. 3.6.

Theorem 3.7 *Let L be a nearly-equilength linkset. Then, any maximal independent set of $U_{z_1(\beta)}(L)$ is an $O(1)$-approximation of PC-Capacity and any greedy independent set of $U_{z_1(\beta)}(L)$ is an $O(1)$-approximation of PC-Weighted-Capacity.*

We define a coloring of a graph G to be *minimal* if it uses at most $D(G)+1$ colors, where $D(G)$ is the maximum degree of a vertex in G.

Theorem 3.8 *Let L be a nearly-equilength linkset. Let $\mathcal{S}$ be a minimal coloring of $U_{z_1}(L)$. Then, using uniform power, $\mathcal{S}$ induces a schedule that yields a $O(1)$-approximation to PC-Scheduling.*

Proof: The coloring $\mathcal{S}$ forms an SINR-feasible schedule of L, by Thm. 3.5, and uses at most $D(U_{z_1}(L))+1$ colors, by the minimality of the coloring. Consider the closed neighborhood $N_1 = N_{U_{z_1}}[v]$ of a maximum degree node v in U_{z_1}. By Obs. 3.6, at most $s = (1+2z_1/z_2)^A$ nodes in N_1 can be in any feasible slot. Hence, the optimal solution uses at least $|N_1|/s = (D(U_{z_1}(L))+1)/s$ slots, for a performance ratio of s. □

The performance ratio of our algorithms is bounded by $C(1+\frac{2z_1}{z_2})^A$.

Efficient distributed algorithms are known for coloring unit-disc graphs in the plane [10] and more generally bounded-independence graphs [34]. Thus, our characterization can be translated into distributed constant-factor approximation algorithms of PC-Scheduling and PC-Capacity in nearly-equilength linksets, when given the appropriate communication primitives.

3.3 Scheduling Arbitrary Linksets

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 Thm. 3.8, or to select the largest of the approximately maximum (weighted) capacity subsets from each of the groups. 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 $g(L) \leq \log\Delta$.

Theorem 3.9 *The PC-Scheduling and PC-Weighted-Capacity problems are $O(g(L))$-approximable, using any oblivious power assignment.*

Moscibroda and Wattenhofer [33] showed that uniform and linear power scheduling can be highly suboptimal, and Moscibroda, Oswald and Wattenhofer [32] showed that they can can be as much as a factor of n or $\Omega(g(L))$ from optimal. Hence, the ratio of $O(\log\Delta)$ is best possible for these power regimes.

We can also claim easy *online* algorithms. The algorithm for PC-Scheduling is in fact online.

Corollary 3.10 *There is a deterministic online algorithm for PC-Scheduling that is constant competitive on nearly-equilength links and $O(\log\Delta)$-competitive in general.*

A similar result can be attained by PC-Capacity by a *randomized* online algorithm that randomly picks one of the length groups and then picks greedily from that group. If the value of Δ is not known, then an approach of Lipton and Tomkins [29] can be used.

Corollary 3.11 *There is a randomized $O(\log\Delta)$-competitive algorithm for PC-Capacity, when Δ is known in advance, and a $O(\Delta^{1+\epsilon})$-competitive algorithm otherwise, for any $\epsilon > 0$.*

4 Oblivious Power Assignments

We explore in this section the power of oblivious assignments. The results of the preceding section apply to all oblivious power functions, but are tight only for uniform and linear power assignments. We can greatly surpass these bounds by being selective about the oblivious function used; in particular, we obtain these improvements for the mean power assignment $\mathcal{M}$.

We present in Sec. 4.2 a scheduling algorithm using $\mathcal{M}$ that achieves a ratio of $O(\log\log\Delta \cdot \log n)$. In the bidirectional setting, the algorithm obtains an improved $O(\log n)$-ratio, as shown in Sec. 4.3. The same results hold also for the (weighted) capacity problem. We complement these results with a construction in Sec. 4.4 that suggests a $\Omega(\log\log\Delta)$-separation between the lengths of optimal schedules with or without oblivious power assignments.

We first introduce our approximation technique, which may be of independent interest.

4.1 Approximation Via Inductiveness

A common heuristic for subset problems is to find a "good" item, and then recurse on the set of remaining items that are compatible with the first one. This yields good approximations if we can show that only a small number of the items eliminated in each round can belong to any optimal solution. For instance, if the incompatibilities are in the form of a graph and the set of nodes eliminated can be covered by k cliques, a k-approximation follows. We generalize this well-known concept to fit to our situation.

We have a set of links L, and a set property on L in the form of p-signal sets. We shall partially capture this property with a graph in the following sense. For a set of links to be feasible it is a sufficient but not a necessary condition for it to form an independent set in the graph. E.g., for set L of nearly-equilength links, this property holds for the graph $U_{z_1}(L)$, by Thm. 3.5. This motivates the extensions we put forth below.

A set property π is said to be *hereditary* if, whenever $\pi(S)$ holds for a set S, it also holds for any $S' \subseteq S$. In other words, π represents a monotone Boolean function. A (sub)set satisfying π is said to be a π-(sub)set. Let $\Pi(S)$ be the maximum cardinality of a π-subset of S. We say that a graph $G = (V, E)$ is *compatible* with a property π on V if any independent set in G satisfies π.

Definition 4.1 *Let V be a set of elements, π be a hereditary set property defined on V, and $G = (V, E)$ be a graph on V that is compatible with π. Then, G is* k-π-inductive *if there is an ordering $v_1, v_2, \ldots, v_n$ of the elements, such that for any v_i, $1 \le i \le n$, it holds that $\Pi(N[v_i] \cap \{v_i, v_{i+1}, \ldots, v_n\}) \le k$.*

To be useful in this context, property π needs to be polynomial-time checkable; this holds for the case of SINR feasibility by solving a system of linear constraints. Additionally, there needs to be an oracle to determine the inductive ordering of the vertices.

This property generalizes the property of being *sequentially k-independent*, where π is the property of a vertex set being independent in the graph. This latter property has been around for a while, but was first studied explicitly in [1], followed by [39]. Various optimization problems can be approximated on sequentially k-independent graphs within a factor of k, including Weighted Independent Set [1] and Graph Coloring [39], when an appropriate vertex ordering can be determined.

The *(Weighted) Maximum π-subset* problem is defined as follows for a given hereditary property π: Given a set V of items and a (vertex-weighted) graph $G = (V, E)$ compatible with π, find a maximum (weight) subset $X \subset V$ that satisfies π. In the *Minimum Partition into π-Subsets* problem, we seek a partition of V into fewest number of π-subsets. Note that optimal solutions to these problems do not depend on the graph G; rather, the structure of G specifies the inductiveness characteristic.

Proposition 4.2 *Let π be a polynomially-time verifiable hereditary property with a polynomial-time oracle to find k-π-inductive orderings. Then, there are k-approximation algorithms for the* Weighted Maximum π-Subset *and* Minimum Partition into π-Subsets *problems on k-π-inductive instances.*

Proof: Let $G = (V, E)$ denote an input instance, which by definition is compatible with π. To approximate the unweighted π-subset problem, we process the nodes in the k-π-inductive order. For each vertex we encounter, we add it to our solution if it has no neighbor among the previously added vertices. This results in an independent set in G, which is a feasible π-subset since G is compatible with π. For each node added to the solution, at most k nodes from any feasible solution are eliminated from consideration, by Def. 4.1. Hence, our solution is within a factor of k from optimal.

To approximate the partitioning problem, we process the nodes in the reverse k-π-inductive order and assign each node to the first class to which no neighbor in G has previously been assigned. Again, each set is independent and thus we obtain a proper partition into π-sets. Let v_i be a node assigned the largest numbered class by this algorithm and let $N_i = \{v_j : (j > i) \wedge (v_j, v_i) \in E(G)\}$ be the neighbors of v_i that follow it in the inductive order. Observe that the number of the class that v_i is assigned, and thus the total number of classes used by the algorithm, is at most $|N_i| + 1$. On the other hand, by the definition of k-π-inductiveness, at most k nodes in $N_i \cup \{v_i\}$ belong to any π-set, and thus the optimal partition of V uses at least $(|N_i| + 1)/k$ classes. Hence, the algorithm is k-approximate.

To approximate the weighted π-subset problem, we use the local ratio algorithm of [39]. The algorithm and its proof are given in the appendix for completeness. □

4.2 Unidirectional Scheduling

We shall utilize the *mean* power assignment (or, square-root assignment [15]) given by $\mathcal{M}_v = \ell_v^{\alpha/2}$. The affectance of link ℓ_w on link ℓ_v under $\mathcal{M}$ is

$$a_w(v) = \frac{\mathcal{M}_w/d_{wv}^\alpha}{\mathcal{M}_v/\ell_v^\alpha} = \left(\frac{\ell_w}{\ell_v}\right)^{\alpha/2}\left(\frac{\ell_v}{d_{wv}}\right)^\alpha = \left(\frac{\sqrt{\ell_v \ell_w}}{d_{wv}}\right)^\alpha .$$

The following observation motivates the consideration of this power assignment.

Observation 4.3 *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.*

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

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. This central observation holds independent of metric. Recall that we assumed that $\beta \geq 3^\alpha$, and thus it follows from Lemma 2.1 that any slot in an optimal solution is a 3-independent linkset.

Lemma 4.4 *Let Q be a well-separated 3-independent set of links in an arbitrary metric space, and let ℓ_v be a link that is shorter than the links in Q by a factor of at least Λ. Suppose all the links in Q are $\frac{1}{\tau}$-close to ℓ_v under mean power assignment. Then, $|Q| = O(\log\log\Delta)$.*

Proof: The set Q consists of two types of links: those that affect ℓ_v by at least $\frac{1}{\tau}$ 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 $1/\tau$, and suppose without loss of generality that $\ell_w \geq \ell_{w'}$. The affectance of ℓ_w on ℓ_v implies that $\sqrt{\ell_v \ell_w}^{\alpha} \geq d_{wv}^{\alpha} \cdot 1/\tau$, or

$$d_{wv} \leq \sqrt{\ell_v \ell_w} \tau^{1/\alpha} = \sqrt{\ell_v \ell_w \Lambda/2} \, .$$

Similarly, $d_{w'v} \leq \sqrt{\ell_v \ell_{w'} \Lambda/2}$. By the triangular inequality we have that

$$d_{w'w} \leq d(s_{w'}, r_v) + d(r_v, s_w) + d(s_w, r_w) \leq \ell_w + \sqrt{2\Lambda \ell_v \ell_w} < 3\ell_w,$$

using that $\sqrt{\ell_w} \geq \sqrt{\Lambda}\sqrt{\ell_v}$. Similarly,

$$d_{ww'} \leq d_{wv} + d_{w'v} + \ell_{w'} \leq \ell_{w'} + \sqrt{2\Lambda \ell_v \ell_w} \, .$$

Multiplying together, we obtain that

$$d_{w'w} \cdot d_{ww'} \leq 3\ell_{w'}\ell_w + 3\sqrt{2\Lambda \ell_v \ell_w} \cdot \ell_w \, .$$

By 3-independence, $d_{w'w} \cdot d_{ww'} \geq 9\ell_w \ell_{w'}$. By combining the last two inequalities and cancelling a $6\ell_w$ factor, we have that $\ell_{w'} \leq \sqrt{\Lambda \ell_v \ell_w / 2}$, or

$$\ell_w \geq \frac{2\ell_{w'}^2}{\Lambda \ell_v} \, . \tag{4}$$

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

$$\frac{\ell_{i+1}}{\ell_i} \geq \frac{2\ell_i}{\ell_v \Lambda} \geq 2\frac{\ell_i}{\ell_1}, \tag{5}$$

for any $i = 2, 3, \ldots, t$. Thus, if we let $\lambda_i = \ell_i/\ell_1$, we get from (5) that $\lambda_{i+1} \geq 2\lambda_i^2$, and by induction that $\lambda_t \geq 2^{2^{t-1}-1}$. Hence, $|Q| = t \leq \lg\lg \lambda_t + 2 \leq \lg\lg \Delta + 2$, and the lemma follows. □

We now proceed as follows. We partition a given linkset L into classes $L_1, L_2, \ldots, L_M$, where $M = \lceil \lg 2\Lambda \rceil$, such that $L_i = \{\ell : \exists k, \lg \ell \in (i-1+kM, i+kM]\}$. Namely, each L_i is a well-separated set. We shall solve the problems independently on the classes L_i and combine the subsolutions in the obvious way.

We say that link ℓ_w is *length-separated* from link ℓ_v if $\ell_w > \Lambda \ell_v$. Let S be a well-separated linkset and d be the minimum linklength in S. Let $z = z_1(2\beta)$. Define the graph $H(S)$ on S where two links ℓ_v and ℓ_w are adjacent if they are either: a) nearly-equilength and the distance between their senders is at least zd, or b) length-separated and $1/\tau$-close. We show the scheduling and capacity problems are capture well as coloring and independent set problems on the graph H.

Lemma 4.5 *Let S be a well-separated linkset in a fading metric. Then, any subset of S that is independent in $H(S)$ is SINR-feasible.*

Proof: Let X be a subset of S that is independent in $H(S)$. Consider a link ℓ_v in X. Let S_v be the set of links in S that are nearly-equilength to ℓ_v (including ℓ_v), $X_v = X \cap S_v$ and $\hat{X} = X \setminus X_v$. We bound the affectance on ℓ_v separately for X_v and $\hat{X}$. None of the links in $\hat{X}$ are $1/\tau$-close to ℓ_v, so each affects ℓ_v by at most $1/\tau$, for a total of $a_{\hat{X}}(\ell_v) \leq n \cdot 1/\tau = 1/(2\beta)$. By definition, X_v is independent in $U_z(S_v)$, where $z = z_1(2\beta)$, and so by Thm 3.5 it is a 2β-signal set. Hence, $a_{X_v}(\ell_v) \leq 1/(2\beta)$ and $a_X(\ell_v) = a_{X_v}(\ell_v) + a_{\hat{X}}(\ell_v) \leq 1/\beta$. □

The graph $H(S)$ of a well-separated linkset S has good inductiveness properties. Denote the case of k-π-inductiveness when π refers to SINR feasibility as *k-SINR-inductive*.

Lemma 4.6 *Let S be a well-separated linkset in a fading metric. Then, $H(S)$ is $O(\log\log\Delta)$-SINR-inductive. The inductive ordering is that of non-decreasing linklength.*

Proof: Let ℓ_v be the shortest link in S. Let X be an SINR-feasible subset of $N_H[\ell_v]$, the closed neighborhood of ℓ_v in $H(S)$. We shall show that $|X| = O(\log\log\Delta)$. We can then order ℓ_v first and apply the claim inductively on $S \setminus \{\ell_v\}$ to obtain the remainder of the inductive order, yielding the lemma.

Let S_v be the subset of nearly-equilength links in S of length at most double that of ℓ_v. The nearly-equilength feasible (β-signal) linkset $X_v = X \cap S_v$ is an independent set in $U_{z_2(\beta)}(S_v)$, by Thm. 3.5. Note that X_v is contained in the closed-neighborhood of ℓ_v in $U_{z_1(2\beta)}(S_v)$. Then, by Obs. 3.6,

$$|X_v| \le \alpha(U_{z_2(\beta)}[X]) \le C(1 + 2z_1(2\beta)/z_2(\beta))^A = O(1) \ .$$

The other neighbors of ℓ_v, those in $X \setminus X_v$, are length-separated from ℓ_v. By Lemma 4.4, ℓ_v has at most $O(\log\log\Delta)$ length-separated neighbors in X. Hence, $|X| = O(\log\log\Delta) + O(1) = O(\log\log\Delta)$. □

We now apply Prop. 4.2 on each of the $O(\log n)$ classes L_i separately to obtain our main result.

Theorem 4.7 *PC-Scheduling, PC-Capacity, and PC-Weighted-Capacity are $O(\log\log\Delta\cdot\log n)$-approximable in fading metrics.*

Finally, we obtain as corollary, a relationship between schedule length and the chromatic number of a certain graph on the links. Let $G'(L)$ be the graph on the linkset L formed by the complete union of the graphs $H(L_i)$, for $i = 1, 2, \ldots$. Namely, links in different length classes are adjacent in G', while links in the same length class L_i induce the subgraph $H(L_i)$.

Corollary 4.8 *There is an algorithm that outputs a feasible scheduling using $O(\log\log\Delta\cdot\log n)\cdot\chi(G'(L))$ slots in fading metrics.*

4.3 Bidirectional Scheduling

In the bidirectional variant introduced by Fanghänel et al [14], 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 4.4.

Lemma 4.9 *Let S be a set of 2-independent links in a bidirectional fading metric and let ℓ_v be a link. Then, there is at most one link ℓ_w in S with $\ell_w > \tau^{2/\alpha}\cdot\ell_v$ that is $1/\tau$-close under mean power assignment.*

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

$$\left(\frac{\sqrt{\ell_v\ell_u}}{d_{vu}}\right)^\alpha \ge 1/\tau,$$

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

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

Then, ℓ_w and ℓ'_w are not 2-independent, which contradicts our assumption. □

The rest of the argument is identical to the unidirectional case. Lemma 4.5 still holds, while we get a stronger version of Lemma 4.6.

Lemma 4.10 *Let S be a well-separated linkset in a bidirectional fading metric and let q be as in Lemma 4.5. Then, $H(S)$ is $O(1)$-SINR-inductive. The inductive ordering is that of non-decreasing link length.*

As before, we partition L into $O(\log n)$ well-separated subsets $L_1, L_2 \ldots$, using Prop. 4.2 on each of them. This results in the following approximation results.

Theorem 4.11 *There is an $O(\log n)$-approximation for the bidirectional versions of PC-Scheduling and PC-Capacity in fading metrics.*

Finally, we get a tighter relationship with graphs. Let $G'(L)$ be defined as in the previous subsection.

Corollary 4.12 *There is an algorithm that outputs a feasible scheduling using $O(\log n)\cdot\chi(G'(L))$ slots in fading metrics, in the bidirectional setting.*

4.4 Construction

We now give evidence that the upper bounds obtained are close to the best possible for oblivious power functions. A similar result follows also from the constructions in [15] by analyzing the dependence on Δ.

We say that a function f is *well-behaved* if there is an $\epsilon > 0$, such that either a) for any $x > x' > 0$, it holds that $f(x) = O((x/x')^{\alpha-\epsilon})f(x')$ or b) for any $x > x' > 0$, it holds that $f(x) = \Omega((x/x')^{\epsilon})f(x')$.

Theorem 4.13 *For any well-behaved power function ϕ, there is a SINR-feasible instance for which any schedule under ϕ requires $\Omega(\log\log\Delta)$ slots.*

Proof: Consider first the case when ϕ grows moderately slowly, i.e., there are fixed constants ϵ, c, c_0 such that for any x, x' with $x > c_0 x'$, $\phi(x) \leq c \cdot x^{\alpha-\epsilon}\phi(x')$. We assume for simplicity that $\beta = 1$.

Let $t = \max(\lceil(\alpha + \lg c)/\epsilon\rceil, 4)$ and $c_1 = \lg\lg c_0$. Consider the following set of links $L = \{\ell_1, \ell_2, \ldots, \ell_n\}$ located on the real line, where the length of link ℓ_i is $\ell_i = 2^{t^{i+c_1}}$. Let $a_i = \sum_{j=0}^{i} \ell_j$, where ℓ_0 denotes $2^{t_1^c}$. 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 for any $i > j$, we have that $\ell_i \geq c_0\ell_j$, and thus

$$\frac{\phi(\ell_i)}{\phi(\ell_j)} \leq c\left(\frac{\ell_i}{\ell_j}\right)^{\alpha-\epsilon} < c\ell_i^{\alpha-\epsilon} \ , \tag{6}$$

where the second inequality uses that $\ell_j > 1$. Observe that for $i > j$,

$$d_{ji} = (\ell_j - a_{j-1}) + a_{i-1} \leq \ell_{i-1} - a_{i-2} + a_{i-1} = 2\ell_{i-1},$$

and that

$$\lg \frac{\ell_i^\epsilon}{\ell_{i-1}^\alpha} = t^{i+c_1}\epsilon - t^{i-1+c_1}\alpha \geq t\epsilon - \alpha \geq \lg c + \alpha,$$

which together imply that

$$\ell_i^\epsilon \geq c2^\alpha \ell_{i-1}^\alpha \geq c d_{ji}^\alpha \ . \tag{7}$$

Thus, using Inequalities (6) and (7), respectively, we have that for $i > j$,

$$a_j(i) = \frac{\phi_j}{\phi_i} \cdot \frac{\ell_i^\alpha}{d_{ji}^\alpha} > \frac{1}{c\ell_i^{\alpha-\epsilon}} \cdot \frac{\ell_i^\alpha}{d_{ji}^\alpha} = \frac{\ell_i^\epsilon}{c d_{ji}^\alpha} \geq 1 \ .$$

Hence, in any schedule based on the mean assignment, each of the n links must be assigned to distinct slots.

Consider instead the oblivious power assignment function $\Psi(v) = \ell_v^\alpha / \log \ell_v$. Note that for $i > j$ in the configuration above, $d_{ji} = \ell_j - a_{j-1} + a_{i-1} > \ell_j$. Then, under Ψ, we have that for $i > j$,

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

Note that for $k > i$, it holds that $d_{ki} = a_{i-1} + \ell_k - a_{k-1} \geq \ell_k/2$. Thus, for $k > i$,

$$a_k(i) = \frac{\Psi(\ell_k)/d_{ki}^\alpha}{\Psi(\ell_i)/\ell_i^\alpha} = \frac{\ell_k^\alpha}{d_{ki}^\alpha \log \ell_k} \cdot \log \ell_i = \frac{\ell_k^\alpha t^{i-k}}{d_{ki}^\alpha} \leq 2t^{i-k} \ .$$

It follows that under Ψ, for any link $\ell_i \in L$, it holds that

$$a_L(i) \leq \sum_{j<i} t^{j-i} + \sum_{k>i} 2t^{i-k} < 3\sum_{k=1}^{\infty} t^{-k} = \frac{3}{t-1} \leq 1 \ ,$$

using that $t \geq 4$. It follows that the linkset L is SINR-feasible. We thus obtain a lower bound on the performance ratio of any schedule using ϕ of $n = \Omega(\log\log\Delta)$.

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 $\Omega((x/x')^\epsilon)$ can also only schedule a single link in a single slot. On the other hand, using power assignment $f(\ell_v) = \lg \ell_v$, shows that the construction is SINR-feasible, giving the same $\Omega(\log\log\Delta)$ lower bound.

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

5 Conclusions

From a practical perspective, it would be interesting if the logarithmic factor could be removed, giving a $O(\log\log\Delta)$-approximation. Alternatively, non-oblivious power strategies that could be implemented in a distributed setting would be highly desirable.

Acknowledgement

I would like to extend sincere thanks to Tigran Tonoyan for kindly pointing out an error in an earlier version and suggesting a correction. I also thank Thomas Erlebach and Pradipta Mitra for helpful comments. Finally, the existence of this paper owes much to Roger Wattenhofer for introducing me to these fascinating questions.

References

[1] K. Akcoglu, J. Aspnes, B. DasGupta, and M.-Y. Kao. Opportunity-cost algorithms for combinatorial auctions. In E. J. Kontoghiorghes, B. Rustem, and S. Siokos, editors, *Applied Optimization 74: Computational Methods in Decision-Making, Economics and Finance*, pages 455–479. Kluwer Academic Publishers, 2002.

[2] M. Andrews and M. Dinitz. Maximizing capacity in arbitrary wireless networks in the sinr model: Complexity and game theory. In *Proceedings of the 28th Annual IEEE Conference on Computer Communications (INFOCOM)*, April 2009.

[3] P. Assouad. Plongements lipschitziens dans $\mathbf{R}^n$. *Soc. Math. France*, 111(4):429–448, 1983.

[4] V. Auletta, L. Moscardelli, P. Penna, and G. Persiano. Interference games in wireless networks. In *Proceedings of the 4th International Workshop On Internet And Network Economics (WINE)*, volume 5385 of *LNCS*, pages 278–285, 2008.

[5] C. Avin, Y. Emek, E. Kantor, Z. Lotker, D. Peleg, and L. Roditty. SNR diagrams: Towards algorithmically usable SINR models of wireless networks. In *Proceedings of the Twenty-Eighth Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC)*, 2008.

[6] C. Avin, Z. Lotker, and Y. A. Pignolet. On the power of uniform power: Capacity of wireless networks with bounded resources. In *Proceedings of the 17^{th} Annual European Symposium on Algorithms (ESA)*, 2009.

[7] G. S. Brar, D. M. Blough, and P. Santi. The SCREAM approach for efficient distributed scheduling with physical interference in wireless mesh networks. In *Proceedings of the 28th IEEE International Conference on Distributed Computing Systems (ICDCS 2008)*, pages 214–224, 2008.

[8] D. Chafekar, V. Kumar, M. Marathe, S. Parthasarathy, and A. Srinivasan. Cross-layer Latency Minimization for Wireless Networks using SINR Constraints. In *Proceedings of the 8th ACM Int. Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2007.

[9] K. Clarkson. Nearest-neighbor searching and metric space dimensions. In *Nearest-Neighbor Methods for Learning and Vision: Theory and Practice*. MIT Press, 2005.

[10] M. Couture, M. Barbeau, P. Bose, P. Carmi, and E. Kranakis. Location oblivious distributed unit disk graph coloring. In *Proceedings of the 14th International Conference on Structural Information and Communication Complexity (SIROCCO):*, pages 222–233. Springer-Verlag, 2007.

[11] 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.

[12] T. Erlebach and T. Grant. Scheduling multicast requests in the sinr model. In *Proceedings of the Sixth International Workshop on Algorithmic Aspects of Wireless Sensor Networks (ALGOSENSORS)*, 2010.

[13] A. Fanghänel, S. Geulen, M. Hoefer, and B. Vöcking. Online capacity maximization in wireless networks. In *Proceedings of the 22nd Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 92–99, 2010.

[14] A. Fanghänel, T. Keßelheim, H. Räcke, and B. Vöcking. Oblivious interference scheduling. In *Proceedings of the 28^{th} Annual ACM Symposium on Principles of Distributed Computing (PODC)*, August 2009.

[15] A. Fanghänel, T. Keßelheim, and B. Vöcking. Improved algorithms for latency minimization in wireless networks. In *Proceedings of the 37th International Conference on Algorithms, Languages and Programming (ICALP)*, July 2009.

[16] G. Fejes Tóth. *Lagerungen in der Ebene, auf der Kugel und in Raum.* Springer-Verlag, Berlin, 2nd ed. edition, 1972.

[17] Y. Gao, J. C. Hou, and H. Nguyen. Topology control for maintaining network connectivity and maximizing network capacity under the physical model. In *Proceedings of the 27th Annual IEEE Conference on Computer Communications (INFOCOM)*, 2008.

[18] O. Goussevskaia, M. M. Halldórsson, R. Wattenhofer, and E. Welzl. Capacity of Arbitrary Wireless Networks. In *Proceedings of the 28th Annual IEEE Conference on Computer Communications (INFOCOM)*, April 2009.

[19] O. Goussevskaia, T. Moscibroda, and R. Wattenhofer. Local Broadcasting in the Physical Interference Model. In *Proceedings of the Fifth ACM SIGACT-SIGOPS International Workshop on Foundations of Mobile Computing*, August 2008.

[20] O. Goussevskaia, Y. A. Oswald, and R. Wattenhofer. Complexity in Geometric SINR. In *Proceedings of the Eigth ACM Int. Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 100–109, 2007.

[21] J. Gronkvist and A. Hansson. Comparison between graph-based and interference-based STDMA scheduling. In *Proceedings of the Second ACM Int. Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 255–258, 2001.

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

[23] M. M. Halldórsson. Wireless scheduling with power control. In *Proceedings of the 17^{th} European Symposium on Algorithms (ESA)*, September 2009.

[24] M. M. Halldórsson and P. Mitra. Wireless Capacity with Oblivious Power in General Metrics. In *Proceedings of the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, January 2011.

[25] M. M. Halldórsson and R. Wattenhofer. Wireless Communication is in APX. In *Proceedings of the 37th International Conference on Algorithms, Languages and Programming (ICALP)*, July 2009.

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

[27] B. Katz, M. Völker, and D. Wagner. Link scheduling in local interference models. In *Proceedings of the Fourth International Workshop on Algorithmic Aspects of Wireless Sensor Networks (ALGOSENSORS)*, pages 57–71, 2008.

[28] T. Kesselheim. A Constant-Factor Approximation for Wireless Capacity Maximization with Power Control in the SINR Model. In *Proceedings of the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, January 2011.

[29] R. J. Lipton and A. Tomkins. Online interval scheduling. In *Proceedings of the 5^{th} Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 302–311, 1994.

[30] R. Maheshwari, S. Jain, and S. R. Das. A measurement study of interference modeling and scheduling in low-power wireless networks. In *Proceedings of the 6th International Conference on Embedded Networked Sensor Systems (SenSys)*, pages 141–154, 2008.

[31] T. Moscibroda. The worst-case capacity of wireless sensor networks. In *Proceedings of the 6th International Conference on Information Processing in Sensor Networks (IPSN)*, pages 1–10, 2007.

[32] T. Moscibroda, Y. A. Oswald, and R. Wattenhofer. How optimal are wireless scheduling protocols? In *Proceedings of the 26th Annual IEEE Conference on Computer Communications (INFOCOM)*, pages 1433–1441, 2007.

[33] T. Moscibroda and R. Wattenhofer. The Complexity of Connectivity in Wireless Networks. In *Proceedings of the 25th Annual IEEE Conference on Computer Communications (INFOCOM)*, 2006.

[34] T. Moscibroda and R. Wattenhofer. Coloring Unstructured Radio Networks. *Distributed Computing*, 21:271–284, 2008.

[35] T. Moscibroda, R. Wattenhofer, and Y. Weber. Protocol Design Beyond Graph-Based Models. In *Proceedings of the Fifth Workshop on Hot Topics in Networks (HotNets)*, November 2006.

[36] T. Moscibroda, R. Wattenhofer, and A. Zollinger. Topology Control meets SINR: The Scheduling Complexity of Arbitrary Topologies. In *Proceedings of the 6^{th} ACM Int. Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2006.

[37] C. Scheideler, A. W. Richa, and P. Santi. An $O(\log n)$ dominating set protocol for wireless ad-hoc networks under the physical interference model. In *Proceedings of the 9^{th} ACM Int. Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 91–100, 2008.

[38] E. W. Weisstein. Circle packing. From MathWorld–A Wolfram Web Resource. `http://mathworld.wolfram.com/CirclePacking.html`.

[39] Y. Ye and A. Borodin. Elimination graphs. In *Proceedings of the 37th International Conference on Algorithms, Languages and Programming (ICALP)*, pages 774–785, 2009.

A Approximating Weighted Capacity on k-π-Inductive Graphs

We apply the algorithm of Ye and Borodin [39] for Weighted Independent Set in sequentially k-independent graphs to the Weighted Maximum Π-subgraph problem in k-π-inductive graphs.

Theorem A.1 *[39] Let π be a polynomially-time verifiable property such that any independent set satisfies π. Then, the Weighted Maximum π-Subgraph problem is k-approximable on graphs that are k-π-inductive.*

Proof: Let $G = (V, E)$ be a k-π-inductive graph with weight function $w : V \to \mathbf{R}$. The algorithm maintains a stack S of nodes, first pushing the nodes onto the stack, and then popping them off.

```
Let v_1, v_2, ..., v_n be the nodes in the k-π-inductive order.
Initialize ŵ(v_i) = w(v_i)
for i ← 1 to n do   // Push phase
  if (ŵ(v_i) > 0)
     push v_i on S
     for each neighbor v_j ∈ N(v_i) ∩ {v_{i+1}, ..., v_n} do
        Subtract ŵ(v_i) from ŵ(v_j) A ← ∅
while (S is not empty) do    // Pop phase
  u ←pop(S)
  if (u ∪ A is a π-set)
    add u to A
output A
```

Let A be the output of the algorithm and O be an optimal solution. Let S be the set of vertices in the stack at the end of the push phase and S_i be the contents of the stack when v_i is being considered in the push phase.

We first prove that the stack algorithm achieves at least the total weight of the stack. For a given node $v_i \in A$, let $\bar{v}_i$ be the weight of v_i on the stack. Then,

$$w(v_i) = \bar{v}_i + \sum_{v_j \in S_i \cap N(v_i)} \bar{w}(v_j) \ .$$

If we sum up for all $v_i \in A$, we have

$$\sum_{v_i \in A} w(v_i) = \sum_{v_i \in A} \bar{w}(v_i) + \sum_{v_i \in A} \sum_{v_j \in S_i \cap N(v_i)} \bar{w}(v_j) \geq \sum_{v_t \in S} \bar{w}(v_t),$$

where the second equality holds because for any $v_t \in S$, we either have $v_t \in S_i \cap N(v_i)$ for some $v_i \in A$, or we have $v_t \in A$.

Now we prove that the optimal solution achieves at most K times the weight of the stack. For a given vertex $v_i \in O$, we have

$$w(v_i) \leq \bar{w}(v_i) + \sum_{v_j \in S_i \cap N(v_i)} \bar{w}(v_j) \ .$$

If we sum up for all $v_i \in O$, we have

$$\sum_{v_i \in O} w(v_i) \leq \sum_{v_i \in O} \bar{w}(v_i) + \sum_{v_i \in O} \sum_{v_j \in S_i \cap N(v_i)} \bar{w}(v_j) \leq k \sum_{v_t \in S} \bar{w}(v_t),$$

where the second inequality holds because when we sum up for all $v_i \in O$, each of the terms $\bar{w}(v_t)$ for any vertex $v_t \in S$ can at most appear k times, since the ordering $v_1, v_2, \ldots, v_n$ is a k-π-inductive ordering. Therefore, we have

$$\sum_{v_i \in O} w(v_i) \leq k \sum_{v_i \in A} w(v_i) \ .$$

□