

NEARLY OPTIMAL BOUNDS FOR DISTRIBUTED WIRELESS SCHEDULING IN THE SINR MODEL

MAGNÚS M. HALLDÓRSSON AND PRADIPTA MITRA

ABSTRACT. We study the wireless scheduling problem in the SINR model. More specifically, given a set of n links, each a sender-receiver pair, we wish to partition (or *schedule*) the links into the minimum number of slots, each satisfying interference constraints allowing simultaneous transmission. In the basic problem, all senders transmit with the same uniform power.

We give a distributed $O(\log n)$ -approximation algorithm for the scheduling problem, matching the best ratio known for centralized algorithms. It holds in arbitrary metric space and for every length-monotone and sublinear power assignment. It is based on an algorithm of Kesselheim and Vöcking, whose analysis we improve by a logarithmic factor. We show that every distributed algorithm uses $\Omega(\log n)$ slots to schedule certain instances that require only two slots, which implies that the best possible absolute performance guarantee is logarithmic.

1. INTRODUCTION

Given a set of n wireless links, each a sender-receiver pair, what is the minimum number of slots needed to schedule all the links, given interference constraints? This is the canonical problem of *scheduling* wireless communication, which we study here in a distributed setting.

In a wireless network, simultaneous transmissions on the same channel interfere with each other. Algorithmic questions for wireless networks depend crucially on the model of interference considered. In this work, we use the physical, a.k.a. *SINR*, model of interference, precisely defined in Section 2. It is known to capture reality more faithfully than the graph-based models most common in the theory literature, as shown theoretically as well as experimentally [21, 23]. Early work on scheduling in the SINR model focused on heuristics and/or non-algorithmic average-case analysis (e.g. [11]). In seminal work, Moscibroda and Wattenhofer [22] proposed the problem of scheduling an arbitrary set of links. Numerous works on various problems in the SINR setting have appeared since.

The *scheduling* problem has primarily been studied in a centralized setting. In many realistic scenarios, however, it is imperative that a distributed solution be found, since a centralized controller may not exist, and individual nodes in the link may not be aware of the overall topology of the network. For the scheduling problem, the only rigorous result previously known is due to Kesselheim and Vöcking [20], who show that a simple and natural distributed algorithm provides an $O(\log^2 n)$ -approximation.

In this work, we adopt the algorithm of Kesselheim and Vöcking, but provide an improved analysis of an $O(\log n)$ -approximation. This matches the best upper bound known for centralized algorithms. Moreover, we show this to be best possible for distributed algorithms that use no external communication infrastructure.

2. PRELIMINARIES AND CONTRIBUTIONS

Given is a set $L = \{l_1, l_2, \dots, l_n\}$ 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 by $d(x, y)$.

Date: May 21, 2018.

Supported by grants 90032021 and 120032011 from the Icelandic Research Fund. Preliminary version appeared in ICALP 2011.

The asymmetric distance from link l_v to link l_w is the distance from v 's sender to w 's receiver, denoted by $d_{vw} = d(s_v, r_w)$. Let $\ell_v = d(r_v, s_v)$ denote the length of link l_v .

Let P_v denote the power assigned to link l_v , or, in other words, s_v transmits with power P_v . We adopt the SINR model (a.k.a., *physical model*) of interference, in which a node r_v successfully receives a message from a sender s_v if and only if the following condition holds:

$$(1) \quad \frac{P_v/\ell_v^\alpha}{\sum_{l_w \in S \setminus \{l_v\}} P_w/d_{wv}^\alpha + N} \geq \beta,$$

where N is a universal constant denoting the ambient noise, $\alpha > 0$ denotes the path loss exponent, $\beta > 0$ denotes the minimum SINR (signal-to-interference-noise-ratio) required for a message to be successfully received, and S is the set of concurrently scheduled links in the same *slot*. We say that S is *SINR-feasible* (or simply *feasible*) if (1) is satisfied for each link in S .

A power assignment \mathcal{P} is *length-monotone* if $P_v \geq P_w$ whenever $\ell_v \geq \ell_w$ and *sub-linear* if $\frac{P_v}{\ell_v^\alpha} \leq \frac{P_w}{\ell_w^\alpha}$ whenever $\ell_v \geq \ell_w$ [20]. Two widely used power assignments in this class are the *uniform* power assignment, where every link transmits with the same power; and the *linear* power assignment, where P_v is proportional to ℓ_v^α . A third one, *mean power* [6, 12] has also proved to be versatile.

Given a set of links L , the *scheduling problem* is to find a partition of L of minimum size such that each subset in the partition is feasible. The size of the partition equals the minimum number of slots required to schedule all links. We will call this number the *scheduling number* of L , and denote it by $\chi(L)$ (or χ when clear from context).

Distributed algorithms. A communication infrastructure for running distributed algorithms is generally assumed to exist in the traditional distributed setting. The current setting, which abstracts the MAC layer in networks, is different, as the goal actually is to construct such an infrastructure. Thus, our algorithm will work with very little global knowledge and minimal external input.

Communication is only available over the channel. Algorithms operate in synchronous rounds with the senders either transmitting or listening in each round. When transmission is successful, the sender stops transmitting. This necessitates an acknowledgment from the receiver, so that the sender knows when his message has been heard. These acknowledgments are sent over the same channel as the message; thus, there are no side-channels for control messages. We shall assume this model, which we call *ack-only*, in the rest of the paper.

We assume that nodes have a rough estimate of the network size n and (senders of) links are assigned a fixed length-monotone, sublinear power function. The power assignment indirectly requires knowledge of distances and the path loss constant α and the technological parameters β and N . No information of locations is needed.

We note that the assumptions are particularly minimal when using uniform power. The algorithm then needs no knowledge of distances, the path loss constant α , nor the technological parameters β and N . Only the polynomial bound on the number n of nodes is needed.

Affectance. We will use the notion of *affectance*, introduced in [9, 17] and refined in [20] to the thresholded form used here. The affectance $a_w^{\mathcal{P}}(v)$ on link l_v from another link l_w , with a given power assignment \mathcal{P} , is the interference of l_w on l_v relative to the power received, or

$$a_w^{\mathcal{P}}(v) = \min \left\{ 1, c_v \frac{P_w/d_{wv}^\alpha}{P_v/\ell_v^\alpha} \right\},$$

where $c_v = \beta/(1 - \beta N \ell_v^\alpha / P_v)$ depends only on model constants and on the length of l_v . We will drop \mathcal{P} and assume it to be an arbitrary length-monotone sub-linear power strategy, unless otherwise stated. Let $a_v(v) = 0$. For a set S of links and a link l_v , let $a_S(v) = \sum_{l_w \in S} a_w(v)$, referred to as *in-affectance*, and $a_v(S) = \sum_{l_w \in S} a_v(w)$, the *out-affectance* from l_v . For sets S and R , $a_R(S) = \sum_{l_v \in R} \sum_{l_u \in S} a_v(u)$. Using such notation, (1) can be rewritten as

$$(2) \quad a_S(v) \leq 1,$$

whenever $|S| > 2$, and this is the form we will use.

2.1. Related Work. In the centralized setting, scheduling results have closely followed results on the related *capacity* problem, where one wants to find the maximum subset of L that can be transmitted in a single slot). Goussevskaia et al. [10] showed the problem to be NP-hard for the case of uniform power on the plane and gave $O(\log \Delta)$ -approximation result (on the plane), where Δ denotes the ratio between the maximum and minimum length of a link. Same bound was shown by Andrews and Dinitz [1] but in comparison with optimum that is allowed to choose arbitrary power. Constant factor approximation was obtained for uniform power, also on the plane, by Goussevskaia et al. [9], which was generalized to all length-monotone, sublinear power assignments and arbitrary metrics space by Halldórsson and Mitra [14]. Kesselheim [18] gave a constant-factor approximation for the joint problem of selecting links and assigning them feasible power (see also earlier work of Chafekar et al. [4]).

All the results lead to equivalent bounds for the centralized scheduling problem with $O(\log n)$ -factor overhead. In particular, $O(\log n)$ -approximation holds for scheduling with length-monotone, sublinear power [14] and with arbitrary power control [18]. Also, the problem remains NP-hard [10]. For the results in terms of Δ on the plane [10, 1], this overhead can be avoided (see, e.g., [12]). Scheduling with arbitrary power control can also be approximated within a factor of $O(\log n \log \log \Delta)$ when the algorithm uses mean power. For linear power on the plane, an algorithm using $O(\chi + \log^2 n)$ slots for instances with optimal schedule length χ was given by Fanghänel et al. [7]; on the plane, this can be improved to a constant factor [26]. A bi-directional version was studied by Fanghänel et al. [6] and further treated in [12, 14] and the joint multi-hop scheduling and routing was treated by Chafekar et al. [4].

In the distributed setting, the capacity problem was treated with no-regret learning by Dinitz [5] culminating in a $O(1)$ -approximation algorithm for uniform power of Ásgeirsson and Mitra [2]. However, these game-theoretic algorithms take time polynomial in n to converge, and thus can be viewed more appropriately as determining capacity instead of realizing it in “real time”.

For distributed scheduling, the only work that we are aware of is the groundbreaking paper of Kesselheim and Vöcking [20], who give a distributed $O(\log^2 n)$ -approximation algorithm for the scheduling problem with fixed length-monotone and sublinear power assignment. Our results constitute a $\Omega(\log n)$ -factor improvement. Kesselheim and Vöcking also extend their results to multi-hop scheduling, with the same approximation factor, for which our improvements do not apply, and to routing, with an extra logarithmic factor.

A versatile measure introduced in [20] is the maximum average affectance \bar{A} of a link set L , defined as

$$\bar{A}(L) := \max_{R \subseteq L} \text{avg}_{l \in R} a_R(l) = \max_{R \subseteq L} \frac{a_R(R)}{|R|}.$$

They then show two results that combined yield the $O(\log^2 n)$ -approximation factor. On the one hand, they show that $\bar{A}(L) = O(\chi(L) \log n)$. On the other hand, they present a natural algorithm (which we also use in this work) that schedules links in $O(\bar{A}(L) \log n)$ slots. We show that both of these bounds are tight. Thus, it is not possible to obtain improved approximation using the measure \bar{A} .

Following the original publication of this work, the results have been applied to distributed connectivity and aggregation [15, 3]. A different approach for distributed capacity was proposed by Pei and Kumar [24], with complexity that is a function of the link lengths. In a recent follow-up work, Halldórsson et al. [13] have shown that $\bar{A}(L) = O(\chi)$ for all sublinear, length-monotone power assignments other than uniform power.

2.2. Our Contributions. We achieve the following results:

Theorem 1. *There is a $O(\log n)$ -approximate distributed algorithm for the scheduling problem, in arbitrary metric space and for all length-monotone sublinear power assignments.*

Theorem 2. *For every n , there is an instance L_n of links on the real line that can be scheduled in two slots but for which every distributed algorithm uses $\Omega(\log n)$ -slots (w.h.p). Thus, $\Theta(\log n)$ is the best absolute approximation factor for a distributed scheduling algorithm.*

As in [20], our upper bound results hold in arbitrary distance metrics (and do not require the common assumption that $\alpha > 2$). We also show that the results hold independent of the ambient noise term N , extending [20]. The lower bound result necessarily holds independent of power assignment strategy and for all settings of the technological constants α , β and N .

One of our main technical insights is to devise a different measure that involves *median* rather than average affectance. The measure $\Lambda = \Lambda(L)$ is given by

$$\Lambda(L) := \max_{R \subseteq L} \text{median}(A(R)) ,$$

where $A(R) = \{a_R(l) : l \in R\}$ is the multi-set of in-affectance values of links in the subset R , and $\text{median}(X)$ denotes the median of a multi-set X . Since we only insist that half of the given subset R of links have affectance bounded by Λ , the value of Λ may be much smaller than \bar{A} . Indeed, we show that $\Lambda = O(\chi)$ and that the algorithm schedules all links in time $O(\Lambda \log n)$, achieving the claimed approximation factor.

The other main technical contribution of the paper is the introduction of the concept of *anti-feasibility*. A set S of links is anti-feasible¹ if $a_v(S) \leq 2$, for every l_v in S ; i.e., if the out-going affectance from each link is small. A set is *bi-feasible* if it is both feasible and anti-feasible. We observe in this paper that every feasible set contains a large bi-feasible set and that certain analyses are easier on bi-feasible sets. This has proved useful in later works, e.g., in giving simplified analysis of capacity approximation algorithms [19, 16].

In the next section, we give the improved analysis of a $O(\log n)$ -factor for distributed scheduling, via the measure Λ ; the treatment of acknowledgments is given in Section 3.2. We show in Section 4 that this logarithmic factor is best possible, and give a construction in Section 5 that shows that this result cannot be obtained in terms of the measure \bar{A} .

3. $O(\log n)$ -APPROXIMATE DISTRIBUTED SCHEDULING ALGORITHM

The algorithm from [20], listed below as **Distributed**, is a natural backoff scheme, in the tradition of ALOHA [25]. It is run synchronously, but independently, on each sender of a link. The algorithm, and all the results in this section, work for an arbitrary fixed sublinear length-monotone power assignment.

The algorithm is mostly self-descriptive. The constant c_1 is to be chosen to satisfy the high probability bound desired. One point to note is that Line 2 necessitates some sort of acknowledgment mechanism for the distributed algorithm to stop. For simplicity, we will defer the issue of acknowledgments to Section 3.2 and simply assume their existence for now. Thm. 3 below implies our main positive result. Let $\Lambda = \Lambda(L)$.

Theorem 3. *If all links of a set L of n links run **Distributed**, then L is fully scheduled in $O(\Lambda \log n)$ slots, with high probability.*

To prove Thm. 3, we claim the following.

Lemma 3.1. *Consider a subset $R \subseteq L$ of links and a particular time slot t in which each sender of R transmits with probability $q \leq \frac{1}{2\Lambda}$. Then, the expected number of successful transmissions is at least $\frac{q \cdot |R|}{4}$.*

¹For a technical reason we use a different constant here than for feasibility; the signal-strengthening result of [17] implies that this only affects constants in the approximation factors.

Algorithm 1 Distributed

```

1:  $k \leftarrow 0$ 
2: loop
3:    $q = \frac{1}{4 \cdot 2^k}$ 
4:   for  $\frac{4}{q} c_1 \ln n$  slots do
5:     transmit with i.i.d. probability  $q$ 
6:     if successful (and acknowledged) then
7:       halt
8:     end if
9:   end for
10:   $k \leftarrow k + 1$ 
11: end loop

```

Proof. Define $M = M_\Lambda(R) = \{l_u \in R : a_R(u) \leq \Lambda\}$. By the definition of Λ , $|M| \geq |R|/2$. Thus, it suffices then to show that at least $q|M|/2$ transmissions in slot t are successful in expectation.

Intuitively, the success probability of a link is proportional to its in-affectance. The links in M are the ones with low in-affectance, so as long as the transmission probability q is less than $1/(2\Lambda)$, they will succeed with probability $1/2$ if transmitting.

For $l_u \in R$, recall that $T_u = T_u(t)$ is the indicator random variable that link l_u transmits, and let $S_u = S_u(t)$ be the indicator random variable that l_u succeeds.

We shall make use of a few elementary facts about probabilities. For a (Bernoulli) indicator random variable X , $\mathbb{E}(X) = \Pr(X)$. For random variables X_1, X_2, \dots , it holds by the linearity of expectation that $\sum_i \mathbb{E}(X_i) = \mathbb{E}(\sum_i X_i)$. And, for a random variable X that assumes non-negative values, $\mathbb{P}(X > 1) \leq \mathbb{E}(X)$.

Armed with these facts, we can now bound the probability that a transmitting link $l_u \in M$ is unsuccessful:

$$\begin{aligned}
\mathbb{P}(S_u = 0 | T_u = 1) &= \mathbb{P}\left(\sum_{l_v \in R} a_v(u) T_v > 1\right) \\
&\leq \mathbb{E}\left(\sum_{l_v \in R} a_v(u) T_v\right) \\
&= \sum_{l_v \in R} a_v(u) \mathbb{E}(T_v) \\
&= q \sum_{l_v \in R} a_v(u) \leq q \cdot \Lambda,
\end{aligned}$$

where the first equality uses (2), and the last inequality uses the definition of M .

Thus, when $q \leq \frac{1}{2\Lambda}$,

$$\mathbb{P}(S_u = 0 | T_u = 1) \leq 1/2,$$

which allows us to bound the probability of link l_u transmitting in the time slot by

$$\begin{aligned}
\mathbb{E}(S_u) &= \mathbb{P}(S_u = 1) \\
&= \mathbb{P}(T_u = 1) \mathbb{P}(S_u = 1 | T_u = 1) \\
&= q(1 - \mathbb{P}(S_u = 0 | T_u = 1)) \\
&\geq q/2.
\end{aligned}$$

The expected number of successful links in the time slot is then

$$\begin{aligned} \mathbb{E} \left(\sum_{l_u \in R} S_u \right) &= \sum_{l_u \in R} \mathbb{E}(S_u) \\ &\geq \sum_{l_u \in M} \mathbb{E}(S_u) \geq |M| \cdot q/2 \geq |R| \cdot q/4, \end{aligned}$$

implying the lemma. \square

of *Thm. 3*. Given Lemma 3.1, the theorem follows essentially from the arguments in Thms. 2 and 3 of [20]. Let $\hat{q} = 2^{-(1+\lceil \lg \Lambda \rceil)}$, i.e., the unique power of two satisfying $\frac{1}{4\Lambda} \leq \hat{q} \leq \frac{1}{2\Lambda}$.

We first bound the probability that not all links are scheduled during the iteration of the outer loop when q in Line 3 equals \hat{q} .

Let \hat{t} be the first time slot where $q \leq \hat{q}$. Let n_t be the random variable indicating the number of links that did not successfully transmit in the first t time slots.

Lemma 3.1 implies that for any given value s and time slot $t \geq \hat{t}$,

$$\mathbb{E}(n_t | n_{t-1} = s) \leq s - \frac{\hat{q}}{4}s,$$

and thus

$$\mathbb{E}(n_t) \leq \sum_{s=0}^{\infty} \mathbb{P}(n_{t-1} = s) \cdot (1 - \hat{q}/4)s = (1 - \hat{q}/4)\mathbb{E}(n_{t-1}).$$

Noting that $n_0 = n$, this yields that

$$\mathbb{E}(n_t) \leq (1 - \hat{q}/4)^t n.$$

Now, after $\hat{t} + 4c_1 \ln n/\hat{q}$ time slots, the expected number of remaining requests is

$$\begin{aligned} \mathbb{E}(n_{\hat{t}+4c_1 \ln n/\hat{q}}) &\leq (1 - \hat{q}/4)^{4c_1 \ln n/\hat{q}} \mathbb{E}(n_{\hat{t}}) \\ &\leq \left(\frac{1}{e}\right)^{c_1 \ln n} n = n^{1-c_1}. \end{aligned}$$

By Markov's inequality,

$$\begin{aligned} \mathbb{P}(n_{\hat{t}+4c_1 \ln n/\hat{q}} \neq 0) &= \mathbb{P}(\hat{t} + n_{4c_1 \ln n/\hat{q}} \geq 1) \\ &\leq \mathbb{E}(n_{\hat{t}+4c_1 \ln n/\hat{q}}) \leq n^{1-c_1}. \end{aligned}$$

Thus, with high probability all the links are scheduled while $q \geq \hat{q}$.

Finally, to bound the total running time of the algorithm, we sum up the spent for values of q smaller than \hat{q} , bounding t_0 . This is a geometric series given by

$$\begin{aligned} t_0 &= \sum_{i=2}^{\lg(1/\hat{q})} \frac{8c_1 \ln n}{2^{-i}} \\ &= 8c_1 \ln n \sum_{i=2}^{\lg(1/\hat{q})} 2^i \\ &\leq 8c_1 \ln n \cdot 2^{\lg(1/\hat{q})+1} \\ &= 8c_1 \ln n \cdot \frac{2}{\hat{q}} \\ &\leq 64c_1 \Lambda \ln n, \end{aligned}$$

establishing the time complexity. \square

3.1. Bounding the Measure. We need the following lemma to get a handle on affectances. Recall that we assumed that the implicit power assignment is length-monotone and sublinear.

Lemma 3.2 (Lemma 7, [20]). *Let L be a feasible set and $l_u \notin L$ be link with $\ell_u \leq \ell_v$ for all $l_v \in L$. Then, $a_L(u) = O(1)$.*

We now prove the following complementary result. It can be contrasted with Lemma 9 of [20], which without the anti-feasibility condition can only give $a_v(L) = O(\log n)$. The second part of the lemma essentially follows Lemma 11 of [2] (which had the unnecessary assumption that L is feasible).

We first need the following result.

Lemma 3.3 ([12]). *Let l_u, l_v be links with $\min(a_u(v), a_v(u)) \leq 1/q$. Then, $d_{uv} \cdot d_{vu} \geq q^2 \cdot \ell_u \ell_v$.*

Lemma 3.4. *Let L be an anti-feasible set with length-monotone and sublinear power and let $l_v \notin L$ be a link with $\ell_v \leq \ell_u$, for every $l_u \in L$. Then, $a_v(L) = O(1)$.*

Proof. We first use a variation of the signal strengthening technique of [17], given as Thm. 7 in the Appendix. This allows us to decompose the set L into $\lceil 4 \cdot 3^\alpha \rceil^2$ sets, where each set S satisfies $a_w(S) \leq \frac{1}{3^\alpha}$, for all $l_w \in S$. We shall prove the claim for S ; the claim will then hold for L by summing over the $\lceil 4 \cdot 3^\alpha \rceil^2$ sets.

Let $l_u = (s_u, r_u)$ ($l_w = (s_w, r_w)$) be the link in S whose sender (receiver) is closest to s_v , i.e., $d(s_v, s_u) \leq \min_{l_x \in S} d(s_v, s_x)$ ($d(s_v, r_w) \leq \min_{l_x \in S} d(s_v, r_x)$), respectively. Let $h = d(s_v, s_u)$. We claim that for all links l_x in S , $l_x \neq l_w$, it holds that

$$(3) \quad d(s_v, r_x) \geq \frac{1}{2}h .$$

To prove this, assume, for contradiction, that $d(s_v, r_x) < \frac{1}{2}h$. Then, by the definition of l_w , $d(s_v, r_w) < \frac{1}{2}h$, and by the definition of l_u , $d(s_v, s_x) \geq d(s_v, s_u) \geq h$ and $d(s_v, s_w) \geq h$. Thus, $\ell_w \geq d(s_v, s_w) - d(s_v, r_w) > \frac{h}{2}$ and similarly $\ell_x > \frac{h}{2}$. On the other hand, by the triangular inequality and the assumed inequality,

$$d(r_w, r_x) \leq d(r_w, s_v) + d(s_v, r_x) < \frac{h}{2} + \frac{h}{2} < h .$$

Now,

$$\begin{aligned} d_{wx} \cdot d_{xw} &\leq (\ell_w + d(r_w, r_x))(\ell_x + d(r_w, r_x)) \\ &< (\ell_w + h)(\ell_x + h) \\ &< 9\ell_w \ell_x , \end{aligned}$$

contradicting Lemma 3.3. This establishes (3).

Now, by the triangular inequality, the definition of h and (3),

$$d_{ux} = d(s_u, r_x) \leq d(s_u, s_v) + d(s_v, r_x) \leq 3d(s_v, r_x) = 3d_{vx} .$$

We observe that $P_v \leq P_u$ holds by length-monotonicity. Also, note that since the maximum affectance between links in S is $\frac{1}{3^\alpha}$, the thresholding in the affectance definition does not take effect, implying that $a_u(x) = c_x \frac{P_u \ell_x^\alpha}{d_{ux}^\alpha P_x}$. Thus,

$$a_v(x) = c_x \frac{P_v \ell_x^\alpha}{d_{vx}^\alpha P_x} \leq c_x \frac{3^\alpha P_u \ell_x^\alpha}{d_{ux}^\alpha P_x} = 3^\alpha a_u(x) .$$

Finally, summing over all links in S ,

$$\begin{aligned} a_v(L) &= a_v(w) + \sum_{l_x \in S \setminus \{l_w\}} a_v(x) \\ &\leq 1 + 3^\alpha \sum_{l_x \in S \setminus \{l_w\}} a_u(x) \\ &\leq 1 + 3^\alpha \cdot 2 = O(1), \end{aligned}$$

using anti-feasibility in the last inequality. The lemma follows. \square

We can now derive the needed bound on the measure.

Theorem 4. *Let L be a set of links. Then, $\Lambda(L) = O(\chi(L))$.*

Proof. Let $\chi = \chi(L)$ and let R be an arbitrary subset $R \subseteq L$. To prove the theorem, it suffices to show that at least half of the links in R have in-affectance at $O(\chi(L))$.

Consider a partition of R into χ feasible subsets S_1, S_2, \dots, S_χ , and define $S'_i = \{l_v \in S_i : a_v(S_i) \leq 3\}$. We claim that S'_i contains at least two thirds of the links in S_i .

Claim 3.5. *For all i , $|S'_i| \geq \frac{2|S_i|}{3}$.*

Proof. Since S_i is feasible, it follows from (2) that $a_{S_i}(v) \leq 1$, for every link $l_v \in S_i$. Let $\hat{S}_i = S_i \setminus S'_i$. Now,

$$a_{\hat{S}_i}(S_i) \leq \sum_{l_v \in S_i} a_{S_i}(v) \leq \sum_{l_v \in S_i} 1 \leq |S_i|.$$

But, $a_{\hat{S}_i}(S_i) = \sum_{l_v \in \hat{S}_i} a_v(S_i) \geq 3 \cdot |\hat{S}_i|$, by the definition of \hat{S}_i . Thus, $|\hat{S}_i| \leq 2|S_i|/3$, proving the claim. \square

Let $R' = \cup_i S'_i$. By the above claim, $3|R'|/4 \geq |R|/2$.

We next show the following. Let c_2 (c_3) be the constant implicit in the big-oh notation in Lemma 3.2 (Lemma 3.4), respectively.

Claim 3.6. $a_R(R') \leq (c_2 + c_3)|R| \cdot \chi$.

Proof. We first observe that for every i, j ,

$$\begin{aligned} a_{S_j}(S'_i) &= \sum_{l_u \in S'_i} \sum_{l_v \in S_j} a_v(u) \\ &\leq \sum_{l_u \in S'_i} \sum_{\substack{l_v \in S_j \\ \ell_v \geq \ell_u}} a_v(u) + \sum_{l_u \in S'_i} \sum_{\substack{l_v \in S_j \\ \ell_v \leq \ell_u}} a_v(u) \\ &\leq \sum_{l_u \in S'_i} c_2 + \sum_{l_v \in S_j} \sum_{\substack{l_u \in S'_i \\ \ell_u \geq \ell_v}} a_v(u) \\ &\leq c_2|S'_i| + \sum_{l_v \in S_j} c_3 \\ (4) \quad &\leq c_2|S_i| + c_3|S_j|, \end{aligned}$$

using Lemma 3.2 and rearrangement in the second inequality, and Lemma 3.4 in the third inequality. We then obtain that

$$\begin{aligned}
a_R(R') &= \sum_{i=1}^{\chi} \sum_{j=1}^{\chi} a_{S_j}(S'_i) \\
&\leq \sum_{i,j=1}^{\chi} c_2 |S_i| + c_3 |S_j| && \text{(By (4))} \\
&= \sum_{i,j=1}^{\chi} (c_2 + c_3) |S_i| && \text{(By symmetry)} \\
&= (c_2 + c_3) \chi \sum_{i=1}^{\chi} |S_i| \\
&= (c_2 + c_3) \chi |R| && \text{(Defn. of } S_i)
\end{aligned}$$

□

It follows that the average in-affectance $a_R(l')$ over the links $l' \in R'$ is at most

$$\frac{a_R(R')}{|R'|} \leq \frac{(c_3 + c_4) |R| \cdot \chi}{|R'|} \leq \mu := \frac{3(c_2 + c_3)}{2} \chi .$$

Recall that $M_{4\mu}(R) = \{l \in R : a_R(l) \leq 4\mu\}$ is the set of links in R of in-affectance at most four times the average. By Markov's inequality, at least three fourths of the links have in-affectance at most four times the average; namely,

$$|M_{4\mu}(R)| \geq |M_{4\mu}(R')| \geq 3|R'|/4 \geq |R|/2 .$$

That is, at least half the links in R have in-affectance at most 4μ . Hence, the median in-affectance of links in R is bounded above by

$$\text{median}(A(R)) \leq 4\mu = O(\chi) .$$

Since this holds for every given R , the theorem follows. □

3.2. Acknowledgments. In the preceding exposition, we ignored the issue of sending acknowledgments from receivers to senders. We can treat acknowledgments in a fashion similar to Kesselheim and Vöcking [20]. We outline their approach briefly, but direct the reader to their paper for the details.

A special slot for acknowledgments is inserted between the time slots used by Algorithm 1. A node that successfully received a packet will transmit an acknowledgment with probability $p = 1/8$. The power P_v^* used for the acknowledgment on link l_v is chosen to be proportional to $P_v^* = \ell^\alpha / P_v$ (using the right scaling factor). Kesselheim and Vöcking show that at least half of these acknowledgments are successful in expectation. That implies that we can modify Lemma 3.1 to claim that the expected number of successfully acknowledged transmissions is at least $p \cdot q|R|/4 = q|R|/32$, losing only a constant factor. The rest of the arguments are then identical.

The only catch is that they do assume in their analysis that there are no weak links in the instance; a link l_v is said to be *weak* iff $c_v > C\beta$, for an appropriately chosen constant C (whose value affects the choice of p). We show here how to extend the approach to deal with weak links. For simplicity of exposition, we illustrate it for the case of uniform power and assume that weak links satisfy $c_v > 3 \max(\beta, 1)$.

The original transmissions, using Algorithm 1, are unchanged, but we allocate a separate time slot for the acknowledgments of weak links. Each receiver of a successfully transmitting weak link sends an acknowledgment in that time slot with probability p' (to be chosen).

The key observation in the following lemma is that weak links must be spatially well-separated. This implies that differences between the positions of the sender and receiver of a link are minor, allowing us to relate the success probability for an acknowledgment in terms of the observed success of the original transmission.

Lemma 3.7. *Assume the use of uniform power. Let l_v be a weak link that transmits successfully in a given time slot t of Algorithm 1. Then, the transmission is successfully acknowledged with probability $p'/2$ when $p' \leq (1 + 2 \ln 3 \cdot \alpha)^{-\alpha}$.*

Proof. Let l_u be another weak link that successfully transmitted at time t . Since both were successful, (2) is satisfied in both directions, which implies that

$$(5) \quad d_{uv}^\alpha \geq c_v \ell_v^\alpha \geq 3\ell_v^\alpha, \text{ and } d_{vu}^\alpha \geq 3\ell_u^\alpha.$$

By the triangular inequality, $d_{uv} \leq d_{vu} + \ell_u + \ell_v$, which by (5) implies that

$$d_{uv} \left(1 - \frac{1}{3^{1/\alpha}}\right) \leq d_{uv} - \ell_v \leq d_{vu} + \ell_u \leq \left(1 + \frac{1}{3^{1/\alpha}}\right) d_{vu}.$$

Now, observe that

$$\frac{\left(1 + \frac{1}{3^{1/\alpha}}\right)}{\left(1 - \frac{1}{3^{1/\alpha}}\right)} \leq 1 + \frac{2}{e^{(\ln 3)/\alpha} - 1} \leq 1 + \frac{2\alpha}{\ln 3}.$$

Thus,

$$(6) \quad d_{uv}^\alpha \leq \left(1 + \frac{2\alpha}{\ln 3}\right)^\alpha d_{vu}^\alpha.$$

Now, let $l_v^* = (s_v^*, r_v^*) = (r_v, s_v)$ be the *dual* link of l_v , with the roles of sender and receiver reversed. A transmission on l_v is acknowledged on l_v^* . We use (6) to bound the in-affectances of a dual link l_v^* from another dual link l_u^* :

$$\begin{aligned} a_{u^*}(v^*) &= c_v \left(\frac{\ell_v}{d_{vu}}\right)^\alpha \\ &\leq \left(1 + \frac{2\alpha}{\ln 3}\right)^\alpha c_v \left(\frac{\ell_v}{d_{uv}}\right)^\alpha \\ &= \left(1 + \frac{2\alpha}{\ln 3}\right)^\alpha a_u(v). \end{aligned}$$

Let S be the set of weak links that successfully transmitted in slot t and S^* the set of the corresponding dual links. Suppose each link in S transmits an acknowledgment with probability p' .

Then, the expected in-affectance of a link l_v^* that transmits an ack is bounded by

$$\begin{aligned} \mathbb{E} \left(\sum_{l_u^* \in S} p' \cdot a_{u^*}(v^*) \right) &\leq \sum_{l_u^* \in S^*} \mathbb{E}(p' a_{u^*}(v^*)) \\ &= p' \sum_{l_u^* \in S^*} a_{u^*}(v^*) \\ &\leq p' \left(1 + \frac{2\alpha}{\ln 3} \right)^\alpha \sum_{l_u^* \in S^*} a_u(v) \\ &\leq \frac{1}{2} a_S(v) \leq \frac{1}{2}, \end{aligned}$$

using the feasibility of S . Hence, the probability that a link receives less than twice the expected in-affectance is at least $1/2$, i.e., a dual link that does attempt to transmit an acknowledgment has at least 50% chance of success. The probability that a given link both attempts to send an acknowledgment and that the transmission is successful, is then at least $p'/2$. \square

4. $\Omega(\log n)$ -FACTOR LOWER BOUND FOR DISTRIBUTED SCHEDULING

We construct a set of $2n$ unit length links on the line that can be scheduled in two slots while no distributed algorithm can schedule the set in less than $\Omega(\log n)$ slots.

We assume that all senders start at the same time in the same state and use the same (randomized) algorithm. Note that the algorithm presented operates under these assumptions.

For simplicity, we assume the noise $N = 0$, but note that the construction can be modified to hold for different values of N . We allow α and β to be arbitrary positive values. We start with a *gadget* F with two identical links of length 1, in a yin-yang position, i.e., with the sender of one link in the same position as the receiver of the other (it suffices that they be separated by at most $(P_{max}/(\beta P_{min}))^{1/\alpha}$, where P_{max} (P_{min}) is the maximum (minimum) power that can be used, respectively). Let $x = (2\beta n)^{1/\alpha}$. The construction consists of n such gadgets $F_i, i = 1, 2, \dots, n$, placed on the line as follows: The sender of one link and the receiver of the other link in F_i are placed at point $i(x+1)$ and the other two nodes of F_i are placed at $i(x+1)+1$.

The construction ensures that *a link successfully transmits only if the other link in the gadget does not transmit*. This holds independent of the power used on these links. On the other hand, when using uniform power, the affectance from links of other gadgets is negligible. To see this, consider the affectance on a link $l_u \in F_i$ from all links of other gadgets, i.e., from all links $l_v \in \hat{F} := \cup_{j \neq i} F_j$. There are $2n - 2$ links in \hat{F} . The distance $d_{vu} \geq x$. Therefore, $\sum_{\hat{F}} a_v(u) \leq (2n - 2) \frac{\beta}{x^\alpha} < 1$. Thus, behavior of links in other gadgets is immaterial to the success of a link. This also implies that the scheduling number of this set of links is 2. Note that since the construction uses equi-length links, the only possible oblivious power assignment is the uniform one.

To prove the lower bound, we say that gadget F_i is *active* at time t if neither link of F_i has succeeded by time $t - 1$, and denote the event by $A_i(t)$. Let $T_u(t)$ denote the indicator random variable that link l_u transmits at time t .

Lemma 4.1. *Let F_i be a gadget and $t \geq 0$ be a time. The transmission probabilities of the two links in F_i at time t are identical and independent, conditioned on F_i being active at time t .*

Proof. Let l_u and l_v be the links in gadget F_i . Let $T_u = T_u(t)$ and $T_v = T_v(t)$, for short. By symmetry, the distributions of T_u and T_v are identical, thus we need only to prove their independence.

We can model the randomness used by the algorithms as an i.i.d. random choice over a set \mathcal{F} of functions. Each $f \in \mathcal{F}$ is a function that takes a history of past transmissions and receptions

over previous slots, and returns a binary transmission decision. Note that if $A_i(t)$ occurs then the histories of l_u and l_v over the previous $t - 1$ slots are identical. The different histories that can result in $A_i(t)$ occurring are disjoint; thus, it is enough to prove independence for a fixed history H . Let f_u and f_v denote the functions chosen by l_u and l_v , and allow them also to represent the event that they get chosen. Once again, by symmetry, there is some $\mathcal{F}' \subseteq \mathcal{F}$ such that H happens iff $f_u \in \mathcal{F}'$ and $f_v \in \mathcal{F}'$. We will use the Iverson bracket $[X]$ to denote the value 1 if X is true and 0 otherwise.

Then, for fixed Boolean outcomes a and b ,

$$\begin{aligned} & \mathbb{P}(T_u = a, T_v = b | H) \\ &= \sum_{f_u \in \mathcal{F}', f_v \in \mathcal{F}'} \mathbb{P}(f_u f_v) [f_u(H) = a] [f_v(H) = b] \\ &= \sum_{f_u \in \mathcal{F}', f_v \in \mathcal{F}'} \mathbb{P}(f_u) \mathbb{P}(f_v) [f_u(H) = a] [f_v(H) = b] \\ &= \sum_{f_u \in \mathcal{F}'} \mathbb{P}(f_u) [f_u(H) = a] \cdot \sum_{f_v \in \mathcal{F}'} \mathbb{P}(f_v) [f_v(H) = b] \\ &= \mathbb{P}(T_u = a | H) \mathbb{P}(T_v = b | H), \end{aligned}$$

thereby proving independence. We have used that $\mathbb{P}(f_u f_v) = \mathbb{P}(f_u) \mathbb{P}(f_v)$ in the second equality, which follows from the fact that f_u and f_v are chosen *a priori* and independently. \square \square

Let p_t denote the i.i.d. probability that some link in a given gadget F_i transmits at time t . Now $\mathbb{P}(A_i(t+1) | A_i(t)) = p_t^2 + (1 - p_t)^2$, which is minimized for $p_t = \frac{1}{2}$ with value $\frac{1}{2}$. Thus,

$$(7) \quad \mathbb{P}(A_i(t+1) | A_i(t)) \geq \frac{1}{2}.$$

Intuitively, on average, at most half of the active gadgets become inactive in any given round, and thus it takes $\lg n$ rounds for all gadgets to become inactive.

Theorem 5. *Let $z(n)$ be a random variable whose value is the smallest time t at which none of the gadgets are active. Then, $\mathbb{E}(z(n)) = \Omega(\log n)$.*

Proof. Consider gadget F_i . Note that for every $t > 0$, $A_i(1) \cap A_i(2) \cap \dots \cap A_i(k) = A_i(k)$ and $\mathbb{P}(A_i(0)) = 1$. Let $t_0 = \lceil \lg n \rceil$. Then, for every $t' \geq t_0$,

$$\begin{aligned} \mathbb{P}(A_i(t')) &= \mathbb{P}(A_i(0)) \prod_{t=2}^{t'} \mathbb{P}(A_i(t) | \cap_{j < t} A_i(j)) \\ &= 1 \cdot \prod_{t=2}^{t'} \mathbb{P}(A_i(t) | A_i(t-1)) \\ &\geq 2^{-(t'-1)} > \frac{1}{n}, \end{aligned}$$

by (7). Let $Q_{t'} = \overline{\cap_i A_i(t')}$ be the event that none of the n gadgets are active at time t' . Since events of different gadgets are independent, it holds for any $t' \geq t_0$ that

$$\mathbb{P}(Q_{t'}) = \prod_{i=1}^n (1 - \mathbb{P}(A_i(t'))) \leq \left(1 - \frac{1}{n}\right)^n \leq e^{-1}.$$

Then, by definition of expectation,

$$\mathbb{E}(z(n)) = \sum_{t=1}^{\infty} \Pr(\overline{Q_t}) \geq t_0 \cdot \Pr(\overline{Q_{t_0}}) \geq (1 - e^{-1}) \lg n.$$

□

Note that bounding $\mathbb{E}(z(n))$ suffices to lower bound the expected time before all links successfully transmit, since by definition a link cannot succeed as long as the corresponding gadget is active.

5. TIGHT BOUND ON ANALYSIS VIA \bar{A}

We achieved a $O(\log n)$ -approximation by avoiding the measure \bar{A} in our analysis. In contrast, the $O(\log^2 n)$ bound in [20] is achieved by proving two separate bounds involving \bar{A} : first $ALG = O(\bar{A} \log n)$, and second $\bar{A} = O(\chi \log n)$, where ALG is the expected time taken by the algorithm. The tightness of the bound on ALG under any oblivious power assignment follows from Section 4, as it is easy to verify that $\bar{A} = \Theta(1)$ in that construction. We give a construction below for which the second bound is tight. Thus, going through \bar{A} is not sufficient to obtain improved bounds, and different analysis is required.

Our construction uses uniform power. This is necessary, since for other oblivious power assignments $\bar{A} = O(\chi)$, by recent results of [13].

Theorem 6. *For every numbers \hat{n} and every number t , there is a set \hat{L} of \hat{n} links with $\chi(\hat{L}) = \Theta(t)$ and $\bar{A}(\hat{L}) = \Omega(\chi(\hat{L}) \log \hat{n})$ under uniform power.*

This lemma shows, perhaps surprisingly, that there can be a huge difference between the in-affectance and out-affectance of a link in a feasible set, thereby illustrating the need for the bi-feasibility concept.

Lemma 5.1. *For every n , there is a set L of n links on the line and a link $l_0 \in L$, such that under uniform power, L is feasible while $a_0(L) = \Omega(\log n)$.*

Proof. We form the set $L = \{l_0, l_1, \dots, l_{n-1}\}$ as follows. The sender s_i of link l_i is positioned at coordinate $d(s_0, s_i) = c \cdot i^{1/\alpha} \cdot 2^i$, where $c > 1$ is a constant to be determined. The length of the link l_i is $\ell_i = 2^i$ and the receiver r_i is positioned at $r_i = s_i + \ell_i = (c \cdot i^{1/\alpha} + 1)2^i$.

Then,

$$\begin{aligned} a_0(L) &= \sum_{i=1}^{n-1} \left(\frac{\ell_i}{d_{0i}} \right)^\alpha \\ &= \sum_{i=1}^{n-1} \left(\frac{2^i}{(c \cdot i^{1/\alpha} + 1)2^i} \right)^\alpha \\ &< \frac{1}{(2c)^\alpha} \sum_{i=1}^{n-1} \frac{1}{i} \\ &= \Omega(\log n) . \end{aligned}$$

To show feasibility, we first bound distances between links by:

$$\begin{aligned} d_{i-1,i} &= d_{0i} - d(s_0, s_{i-1}) \\ &= (c \cdot i^{1/\alpha} + 1)2^i - c(i-1)^{1/\alpha}2^{i-1} \\ &> c \cdot i^{1/\alpha}2^{i-1} , \end{aligned}$$

and for $m > 0$,

$$\begin{aligned}
d_{i+m,i} &= d(s_0, s_{i+m}) - d(s_0, r_i) \\
&> c(i+m)^{1/\alpha} 2^{i+m} - (c \cdot i^{1/\alpha} + 1) 2^i \\
&> c i^{1/\alpha} 2^{i+m} - (c \cdot i^{1/\alpha} + 1) 2^{i+m-1} \\
&= (2c \cdot i^{1/\alpha} - (c \cdot i^{1/\alpha} + 1)) 2^{i+m-1} \\
&\geq (c-1) 2^{i+m-1}.
\end{aligned}$$

We then bound the in-affectance of each link by

$$\begin{aligned}
a_L(i) &= \sum_{k,k < i} a_k(i) + \sum_{k,k > i} a_k(i) \\
&\leq i \cdot a_{i-1}(i) + \sum_{m=1}^{n-i-1} \left(\frac{\ell_i}{d_{i+m,i}} \right)^\alpha \\
&\leq i \cdot \left(\frac{\ell_i}{d_{i-1,i}} \right)^\alpha + \sum_{m=1}^{n-i} \left(\frac{2^i}{(c-1) \cdot 2^{i+m-1}} \right)^\alpha \\
&\leq i \cdot \left(\frac{2^i}{c \cdot i^{1/\alpha} \cdot 2^{i-1}} \right)^\alpha + \frac{1}{(c-1)^\alpha} \sum_{m=0} \left(\frac{1}{2^\alpha} \right)^m \\
&= \frac{2}{c^\alpha} + \frac{1}{(c-1)^\alpha} \cdot \frac{2^\alpha}{2^\alpha - 1}.
\end{aligned}$$

Thus, when $c \geq 1 + \left(3 \left(1 + \frac{1}{2^\alpha - 1} \right) \right)^{1/\alpha}$, it holds that $a_L(i) \leq 1$ for each link l_i , i.e., L is feasible. □ □

We now turn to proving Thm. 6. We construct the set \hat{L} that satisfies the claim of the theorem. Let L be the set feasible under uniform power and the link $l_0 = (s_0, r_0) \in L$ with $a_L(l_0) = \Omega(\log n)$, promised by Lemma 5.1. Let L_1 denote t isometric copies of L with links in the same position as L .

We next take an arbitrary set S on n links that is feasible under uniform power, and scale its distances so that maximum pairwise distance is $o(1)$. For instance, we can let $S = \{l'_1, l'_2, \dots, l'_n\}$ be \hat{L} scaled by a factor of 4^{-n} so that the length of $l'_i = (s'_i, r'_i)$ is 2^{i-2n} , $s'_i = c i^{1/\alpha} 2^{i-2n}$, and $r'_i = s'_i + 2^{i-2n}$. By the same argument as Lemma 6, S is feasible; observe also that pairwise distances of points within S are $O(n/2^n) = o(1)$. Let L_2 denote t copies of S , with the same coordinates; thus, the nodes of L_2 are all close to the node s_0 in L_1 . Finally, we form the combined instance $\hat{L} = L_1 \cup L_2$ with a total of $\hat{n} = 2tn$ links.

Observe that for every $l_j \in L_2$ and $l_i \in L_1$ ($i > 0$), that $d_{ji} = d_{0i}(1 + o(1))$. Since we use uniform power, it holds for each of the tn links $l_j \in L_2$ that

$$a_j(L_1) = \Theta(a_0(L_1)) = \Theta(t \cdot a_0(L)) = \Omega(t \log n).$$

Thus,

$$a_{\hat{L}}(\hat{L}) \geq a_{L_2}(L_1) = |L_2| \Omega(t \log n),$$

implying that

$$\bar{A} \geq \frac{1}{|\hat{L}|} a_{\hat{L}}(\hat{L}) = \Omega(t \log(\hat{n}/t)).$$

On the other hand, the set \hat{L} clearly has a scheduling number of $2t$, as it is formed by $2t$ feasible sets. Hence, the theorem.

6. CONCLUSIONS

We have given a distributed scheduling algorithm that is $O(\log n)$ -approximate in the scheduling model, and shown this factor cannot be improved in general. Our lower bound construction, however, applies only to instances with small scheduling number.

A similar randomized scheduling algorithm was shown by Fanghänel et al. [7] to yield an asymptotic constant-factor approximation for the case of linear power assignment. One key difference is that in the case of linear power, all links have low affectance ($O(\chi)$), while for general sublinear length-monotone power assignments this only holds on average.

It remains an important and intriguing open question whether a better asymptotic approximation ratio can be obtained.

Acknowledgement. We thank Marijke Bodlaender for helpful discussions leading to the derivation of Lemma 5.1.

REFERENCES

- [1] Andrews, M., Dinitz, M.: Maximizing capacity in arbitrary wireless networks in the SINR model: Complexity and game theory. In: INFOCOM, pp. 1332–1340. IEEE (2009)
- [2] Ásgeirsson, E.I., Mitra, P.: On a game theoretic approach to capacity maximization in wireless networks. In: INFOCOM (2011)
- [3] Bodlaender, M.H., Halldórsson, M.M., Mitra, P.: Connectivity and Aggregation in Multihop Wireless Networks. In: PODC (2013)
- [4] Chafekar, D., Kumar, V., Marathe, M., Parthasarathy, S., Srinivasan, A.: Cross-layer Latency Minimization for Wireless Networks using SINR Constraints. In: Mobihoc (2007)
- [5] Dinitz, M.: Distributed algorithms for approximating wireless network capacity. In: INFOCOM (2010)
- [6] Fanghänel, A., Kesselheim, T., Räcke, H., Vöcking, B.: Oblivious interference scheduling. In: PODC, pp. 220–229 (2009)
- [7] Fanghänel, A., Kesselheim, T., Vöcking, B.: Improved algorithms for latency minimization in wireless networks. In: ICALP, pp. 447–458 (2009)
- [8] Goussevskaia, O., Halldórsson, M.M., Wattenhofer, R.: Algorithms for wireless capacity. IEEE/ACM Transactions on Networking (2013). To appear
- [9] Goussevskaia, O., Halldórsson, M.M., Wattenhofer, R., Welzl, E.: Capacity of Arbitrary Wireless Networks. In: INFOCOM, pp. 1872–1880 (2009)
- [10] Goussevskaia, O., Oswald, Y.A., Wattenhofer, R.: Complexity in Geometric SINR. In: Mobihoc, pp. 100–109 (2007)
- [11] Gupta, P., Kumar, P.R.: The Capacity of Wireless Networks. IEEE Trans. Information Theory **46**(2), 388–404 (2000)
- [12] Halldórsson, M.M.: Wireless scheduling with power control. ACM Transactions on Algorithms **9**(1), 7 (2012)
- [13] Halldórsson, M.M., Holzer, S., Mitra, P., Wattenhofer, R.: The power of non-uniform wireless power. In: SODA (2013)
- [14] Halldórsson, M.M., Mitra, P.: Wireless Capacity with Oblivious Power in General Metrics. In: SODA (2011)
- [15] Halldórsson, M.M., Mitra, P.: Distributed connectivity of wireless networks. In: PODC (2012)
- [16] Halldórsson, M.M., Mitra, P.: Wireless capacity and admission control in cognitive radio. In: INFOCOM, pp. 855 – 863 (2012)
- [17] Halldórsson, M.M., Wattenhofer, R.: Wireless Communication is in APX. In: ICALP, pp. 525–536 (2009)
- [18] Kesselheim, T.: A Constant-Factor Approximation for Wireless Capacity Maximization with Power Control in the SINR Model. In: SODA (2011)
- [19] Kesselheim, T.: Approximation algorithms for spectrum allocation and power control in wireless networks. Ph.D. thesis, RWTH Aachen University, Aachen, Germany (2012)
- [20] Kesselheim, T., Vöcking, B.: Distributed contention resolution in wireless networks. In: DISC, pp. 163–178 (2010)
- [21] Maheshwari, R., Jain, S., Das, S.R.: A measurement study of interference modeling and scheduling in low-power wireless networks. In: SenSys, pp. 141–154 (2008)
- [22] Moscibroda, T., Wattenhofer, R.: The Complexity of Connectivity in Wireless Networks. In: INFOCOM (2006)
- [23] Moscibroda, T., Wattenhofer, R., Weber, Y.: Protocol Design Beyond Graph-Based Models. In: Hotnets (2006)
- [24] Pei, G., Kumar, V.A.: Distributed link scheduling under the physical interference model. In: INFOCOM (2012)
- [25] Tanenbaum, A.: Computer Networks, 4th edn. Prentice Hall Professional Technical Reference (2002)

[26] Tonoyan, T.: On the problem of wireless scheduling with linear power levels. CoRR **abs/1107.4981** (2011)

AFFECTANCE REDUCTION

The following is given (with minor modification) in [8, Theorem 4.1].

Theorem 7. *Let S be an anti-feasible set and $p < 2$ be a value. Then, S can be partitioned into $t = \left(\frac{4}{p}\right)^2$ sets S_1, S_2, \dots, S_t , each satisfying $a_v(S_i) \leq p$, for every $l_v \in S_i$.*

Proof. We first partition S into a sequence T_1, T_2, \dots of sets as follows. Order the links in S in decreasing order. For each link l_v , assign l_v to the first set T_j for which $a_v(T_j) \leq p/2$, i.e. the accumulated affectance of l_v on the previous, longer links in T_j is at most $p/2$. Since each link l_v originally had out-affectance at most 2, then by the additivity of affectance, the number of sets used is at most $\lceil \frac{2}{p/2} \rceil = \lceil \frac{4}{p} \rceil$.

We then repeat the same approach on each of the sets T_i , processing the links this time in increasing order. The number of sets is again $\lceil \frac{4}{p} \rceil$ for each T_i , or $\lceil \frac{4}{p} \rceil^2$ in total. In each final slot (set), the affectance of a link on the shorter links in the same slot is at most $p/2$. In total, then, the out-affectance of each link is at most $2 \cdot p/2 = p$. \square

ICE-TCS, SCHOOL OF COMPUTER SCIENCE, REYKJAVIK UNIVERSITY, ICELAND.
E-mail address: mmh@ru.is, ppmitra@gmail.com