The Power of Non-Uniform Wireless Power

Magnús M. Halldórsson* Stephan Holzer† Pradipta Mitra*

Roger Wattenhofer†

September 23, 2012

Abstract

We study a fundamental measure for wireless interference in the SINR model when power control is available. This measure characterizes the effectiveness of using *oblivious* power — when the power used by a transmitter only depends on the distance to the receiver — as a mechanism for improving wireless capacity.

We prove optimal bounds for this measure, implying a number of algorithmic applications. An algorithm is provided that achieves — due to existing lower bounds — capacity that is asymptotically best possible using oblivious power assignments. Improved approximation algorithms are provided for a number of problems for oblivious power and for power control, including distributed scheduling, secondary spectrum auctions, wireless connectivity, and dynamic packet scheduling.

*ICE-TCS, School of Computer Science, Reykjavik University, Reykjavik, Iceland. `mmh@ru.is, ppmitra@gmail.com`

†Distributed Computing Group, ETH Zurich, Switzerland. `{stholzer, wattenhofer}@ethz.ch`

1 Introduction

One of the strongest weapons to improve the capacity of a wireless network is power control. Higher power increases the bandwidth of a single transmission link, while causing more interference to other links that may be transmitting simultaneously. Given this tension, intelligent power control is crucial in increasing the spatial reuse of the available bandwidth. Thus it is not surprising that most contemporary wireless protocols use some form of power control. More recently, this phenomenon has also been studied theoretically; it was shown in a series of works that power control may improve the capacity of a wireless network in an exponential [41, 21] or even unbounded [11] way.

However, unrestricted power control is a double-edge sword. In order to achieve the theoretically best results, one must solve complex optimization problems, where the transmission power of one node often depends on the transmission powers of all other nodes [32]. In real wireless networks where communication demands change over time this may not be an option. In practical protocols, the transmission power should be independent of other concurrent transmissions. In the best case, the transmission power only depends on the distance between sender and receiver. This is known as *oblivious* power control.

Many questions rise immediately in the wake of the previous assertion: What is the price of restricting power control to oblivious powers? Which of the infinitely many oblivious power schemes are good choices? Once an oblivious power scheme is chosen, what algorithmic results can be achieved for it?

In this work, we look at these questions in the context of the physical or SINR model of interference, which is a realistic model of wireless interference gaining accelerating attention (Sec. 1.2 contains historical background and motivation and Sec. 2 contains precise definitions). In this setting, our work answers a number of these questions optimally, completing a significant line of work in the algorithmic study of the SINR model.

The specific problem at the center of our work is the problem of *capacity* maximization: Given a set of transmission links (each a sender-receiver pair), the goal is to find the largest subset of links that can transmit simultaneously.

Before the present work, the state-of-the-art was as follows. The mean power assignment, where a link of length ℓ is assigned power $\ell^{\alpha/2}$ (α being a small physical constant) had emerged as the "star" among oblivious power assignments. It was shown that using mean power, one can approximate capacity maximization for arbitrary power control up to a factor of $O(\log n \cdot \log\log\Delta)$ [21] and $O(\log n + \log\log\Delta)$ [23], where Δ is the ratio between the maximum and minimum transmission distance and n is the number of links in the system. This showed that the somewhat earlier lower bound of $\Omega(n)$ [11] applied only when Δ was doubly exponential. In terms of Δ, it was shown that one *must* pay a $\Omega(\log\log\Delta)$ factor [21]. The best upper bounds were, as mentioned, either dependent on the size of the input [21, 23] and as such unbounded (in relation to Δ), or are exponentially worse ($\log\Delta$) [16, 2].

1.1 Our Contributions

In this paper, we study all power assignments of the form $\ell^{p\cdot\alpha}$ for all fixed $0 < p < 1$ (setting $p = \frac{1}{2}$ gives us mean power). Our first result shows that the lower bound of $\Omega(\log\log\Delta)$ is tight. That is, we give a simple algorithm that uses any oblivious power scheme from the above class, but whose quality is only $O(\log\log\Delta)$-factor worse than the optimum with unrestricted power control. For small to moderate values of Δ, i.e., when Δ is at most polynomial in n (which presumably includes most real-world settings), our bound is an exponential improvement on the previous bounds, including the $O(\log\Delta)$-bound of [2] (see also [16]).

This result, as well as solving the problem of the relation between oblivious and arbitrary power, extends the "star status" from mean power to a large class of assignments. This class has been studied before implicitly in a wide array of work [35, 23, 28, 22] on "length-monotone, sub-linear power" assignments – but their relation to arbitrary power was not understood.

Now that we know what the "right" power assignments are, our second contribution is to improve a large body of algorithmic work on these power assignments. We improve by a logarithmic factor the approximation of a variety of problems for these power assignments — distributed scheduling [35], secondary spectrum auctions [28], wireless connectivity [41, 25, 24], dynamic packet scheduling [34, 3] etc. Using the fundamental capacity relation between oblivious and arbitrary power (our first result), we improve algorithms for these problem in the power control setting as well.

Though we have presented our work above in terms of algorithmic implications; what we actually prove are two *structural* results, out of which this wide class of algorithmic applications fall out essentially immediately. These results are important in their own right, e.g., implying tight bounds on certain efficiently computable measures of interference.

Here is a brief attempt to provide intuitive understanding of these two results in a unified framework. Consider a set of links that can transmit simultaneously (a *feasible* set). What we study is an interference measure between another link (not necessarily in the set) and such a feasible set. Assuming the set is feasible using *some* power assignment (i.e. arbitrary power), we find that the relevant measure can be bound by $O(\log\log\Delta)$ (Thm. 3.4), implying our first capacity result (and its applications). Technically, this is done by carefully extending the analysis of [21]. Assuming the set is feasible under an oblivious power from the class mentioned before, the measure can be bound by $O(1)$ (Thm. 3.8), implying the second set of algorithmic results. We use a potentially novel contradiction technique (at least in the context of SINR analysis) for this result.

Our results apply for general metric spaces and all constant $\alpha > 0$. Apart from the specific applications pinpointed here, we expect any number of future algorithmic questions in the SINR model to directly benefit from these powerful bounds.

1.2 Related Work

Gupta and Kumar [20] were among the first to give analytical results for wireless scheduling in the physical (SINR) model. Those early results analyzed special settings using e.g. certain node distributions, traffic patterns, transport layers etc. In reality, however, networks often differ from these specialized models and no algorithms were provided to optimize the capacity. On the other hand, graph-based models yielded algorithms like [36, 44] but such models do not capture the nature of wireless communication well, as demonstrated in [19, 39, 42]. Six years ago Moscibroda et al. [41] started combining the best of both worlds, studying algorithms for scheduling in arbitrary worst-case networks. Since then, the problems studied in this setting has reflected the diversity of the application areas underlying it – topology control [13, 43, 30], sensor networks [40] , combined scheduling and routing [7], ultra-wideband [29], analog network coding [18].

In spite of this diversity, certain canonical problem have emerged, the study of which have resulted in improvements for other problems as well. The capacity problem is such a problem. After it was quickly shown to be NP-complete [16], a constant factor approximation algorithm for uniform power was achieved in [14, 26], eventually extended to essentially all interesting oblivious power schemes [23]. In [32, 33], a $O(1)$-approximation to the capacity problem for arbitrary powers was obtained. As we have already mentioned, the relation between capacity in oblivious vs. arbitrary power was first studied in [21].

Among fixed power assignments, linear power has turned out to be easiest, being the only one with constant factor approximation for scheduling [12, 47] and constant-bounded inteference measure [12]. Whereas there are instances for which linear and uniform power are arbitrarily bad in comparison with mean power [41], a maximum feasible subset under mean power always within a constant factor of subsets feasible under linear or uniform power [46].

Capacity maximization has recently been studied with respect to several limitations. E.g. [4] investigated the capacity with uniform and non-uniform power-assignments when the network resources are restricted, [10] studies an online-version of capacity maximization with respect to SINR-constraints. In [31] a tradeoff between energy minimization and maximizing the capacity in the SINR. Recently it was shown in [8] that algorithms for capacity-maximization in the SINR model can be transferred to a model that even takes Rayleigh-fading into account by losing only a $O(\log^* n)$ factor in the approximation ratio. This overview is far from being complete, surveys can be found in e.g. [17, 38].

Technically, the idea of looking of the interaction between a feasible set and a link is known. The work of Halldorsson [21] and Kesselheim and Vöcking [35] are specially relevant – the first in the context of oblivious-arbitrary comparison, and the second in the context of oblivious power. Our results improve the bounds in those two papers to the best possible.

1.3 Outline of the Paper

Section 2 lays down the basic setting, including a formal description of the SINR model. In Section 3, we introduce and outline the proof of our two structural results. This is used in Section 4 to give tight approximation algorithm for the capacity problem and in Section 5 to improve various results on topics including distributed scheduling, auctions, and connectivity. Full proofs are given in the appendix; also included there are additional applications of our structural theorems and constructions showing that the assumptions in those theorems are necessary.

2 Model and Definitions

Given is a set $L = \{l_1, l_2, \ldots, l_n\}$ of links, where each link l_v represents a unit-size communication request from a sender s_v to a receiver r_v, both of which are points in an arbitrary metric space. The distance between two points x and y is denoted $d(x, y)$. We write $d_{vw} = d(s_v, r_w)$ for short, and denote by ℓ_v the length of link l_v. Let $\Delta = \Delta(L)$ denote the ratio between the maximum and minimum length of a link in L. d_{vw} ℓ_v

In the *physical model* (or *SINR model*) of interference, a transmission on link l_v is successful if and only if Δ

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

where N is a universal constant denoting the ambient noise, β denotes the minimum SINR (signal-to-interference-noise-ratio) required for a message to be successfully received, $\alpha > 0$ is the so-called path-loss constant, and $S \subseteq L$ is the set of links scheduled concurrently with l_v. β α

Let P_v denote the power assigned to link l_v, or, in other words, s_v transmits with power P_v. We focus on power assignments $\mathcal{P}_p$, where $P_v = \ell_v^{p\cdot\alpha}$. This includes all the specific assignments of major interest: uniform ($\mathcal{P}_0$), mean ($\mathcal{P}_{1/2}$), and linear power ($\mathcal{P}_1$). $\mathcal{P}_p$

We say that S is $\mathcal{P}$*-feasible*, if (1) is satisfied for each link in S when using power assignment $\mathcal{P}$. We say that S is power control feasible (*PC-feasible* for short) if there exists a power assignment $\mathcal{P}$ for which S is $\mathcal{P}$-feasible. We frequently write simply *feasible* when we refer to PC-feasible.

Let PC-Capacity denote the problem of finding a maximum cardinality subset of these links is PC-feasible. Let $OPT^{\mathcal{P}}(L)$ denote the optimal capacity (i.e., size of the largest $\mathcal{P}$-feasible subset) of a linkset L under power assignment $\mathcal{P}$, and $\overline{OPT}(L)$ denote the optimal capacity under any power assignment.

Affectance. We will use the notion of *affectance*, introduced in [14] and refined in [26] and [35]. The affectance $a_w^{\mathcal{P}}(v)$ of link l_v caused by 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 $a_w^{\mathcal{P}}(v)$

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

where the factor $c_v = \beta/(1 - \beta N \ell_v^\alpha / P_v)$ depends only properties of the link l_v and on universal constants. We let $a_v^p(w)$ denote $a_v^{\mathcal{P}_p}(w)$. We shall frequently drop the power assignment reference $\mathcal{P}$, which means then that we assume $\mathcal{P}_p$. Let $a_v(v) = 0$. For sets S and T of links and a link l_v, let $a_v(S) = \sum_{w \in S} a_v(w)$, $a_S(v) = \sum_{w \in S} a_w(v)$, and $a_S(T) = \sum_{w \in S} a_w(T)$. Using this notation, Eqn. 1 can be rewritten as $a_S^{\mathcal{P}}(v) \leq 1$ (assuming S contains more than two links).

We define two more affectance notations. Let $b_v(w) = b_w(v) = a_v(w) + a_w(v)$ be the *symmetric* version of affectance. Let $\hat{a}_v(w)$ ($\hat{b}_v(w)$) be the *length-ordered* version, defined to be $a_v(w)$ ($b_v(w)$) if $\ell_v \leq \ell_w$ and 0 otherwise, respectively. These are extended in similar ways to affectances to and from sets as defined for $a_v(w)$. Notice that $a_S(S) = \hat{b}_S(S) = b_S(S)/2$. $b_v(w)$ $\hat{a}_v(w)$ $\hat{b}_v(w)$

(Non)-weak links A link is said to be *non-weak* if $c_v \leq 2\beta$. This is equivalent to $\frac{P_v}{\ell_v^\alpha} \geq 2\beta N$. Intuitively, this means that the link uses power slightly more than the absolute minimum needed to overcome ambient noise (the constant 2 can be replaced with any fixed constant larger than 1). Our theorems will often assume links to be non-weak. This is reasonable and often-used assumption [35, 2, 9, 15] and can be achieved, if necessary, by scaling the powers.

Length classes A *length class* is any set R of links for which $\Delta(R) \leq 2$ (i.e., link lengths vary by a factor no more than 2). Clearly, any link set L can be partitioned into $\log \Delta(L)$ length classes. We also refer to this as nearly-equilength class.

Independence We refer to links l_v and l_w as *q-independent* if they satisfy $d_{vw} \cdot d_{wv} \geq q^2 \cdot \ell_w \ell_v$. q-independent
A set of mutually q-independent links is said to be *q-independent*.

Independence is a pairwise property, and thus weaker than feasibility. The condition is equivalent to $a_v^{\mathcal{P}}(u) \cdot a_v^{\mathcal{P}}(u) \leq \frac{c_v c_w}{q^{2\alpha}}$, independent of the power assignment $\mathcal{P}$. A feasible set is necessarily $\beta^{1/\alpha}$-independent [21], but there is no good relationship in the other direction.

We give here an independence-strengthening result with better tradeoffs than the so-called "signal-strengthening" result of [26]. The proof is in Appendix A.

Lemma 2.1 *Any feasible set of links can be partitioned into $2q^\alpha/\beta + 1$ or less different q-independent sets.*

3 Structural Properties

We begin with an interference measure.

Definition 3.1 *Let L be a set of links and $\mathcal{P}, \mathcal{Q}$ be two power assignments. Then*

$$I_{\mathcal{Q}}^{\mathcal{P}}(L) \equiv \max_{S \in F_{\mathcal{Q}}(L)} \max_{l_v \in L} \hat{b}_v^{\mathcal{P}}(S),$$

where $F_{\mathcal{Q}}(L)$ is the collection of subsets of L that are $\mathcal{Q}$-feasible. $I_{\mathcal{Q}}^{\mathcal{P}}(L)$

When $\mathcal{P}_p$ is used as one (or both) of the assignments we will use p instead of $\mathcal{P}_p$ in the sub(super)-scripts – thus $I_p^p(L)$ instead of $I_{\mathcal{P}_p}^{\mathcal{P}_p}(L)$.

To get an intuitive handle on this measure, it instructive to look at inductive independence number of a weighted graph, a graph parameter [1] that has recently started to receive increased attention (e.g.[48]). We define this parameter in the context of the SINR model below (a very similar definition for general weighted graphs can be found, for example in [28]).

Definition 3.2 *A set L is d-*inductively independent *for power assignment $\mathcal{P}_p$ if for any link l_v and any $\mathcal{P}_p$-*feasible *subset $S \subseteq L$, $\hat{b}_v^p(S) \leq d$.* d-inductively independent

The following equivalence is easy to see:

Observation 3.3 *A set L of links is d-inductively independent for $\mathcal{P}_p$ iff $I_p^p(L) \leq d$.*

When using different power assignments $I_{\mathcal{Q}}^{\mathcal{P}}(L)$ gives us a handle on how power assignments compare to each other. We will primarily use it in the setting where $\mathcal{P} = \mathcal{P}_p$, for some $p \in (0,1]$ and where $\mathcal{Q}$ is (an) optimal arbitrary power assignment, thus allowing to associate oblivious power to arbitrary power.

Here we give two structural results that characterize the utility of oblivious power assignments. Both of these are best possible and answer long standing open questions. The first characterizes the *price of oblivious power*, i.e., the quality of solutions using oblivious power assignment relative to those achievable by unrestricted power assignments.

Theorem 3.4 *For any set L of non-weak links, any $0 < p < 1$, and any power assignment $\mathcal{Q}$, $I_{\mathcal{Q}}^p(L) = O(\log\log\Delta)$.*

To argue this theorem (Thm. 3.4), we need two lemmas to bound affectances of a link to and from a set of links. Their proofs are given in Appendix B. The first handles the long links in S with relatively high affectance. It originates in [21] (Lemma 4.4), but is generalized here in two ways: to any power assignment $\mathcal{P}_p$, and to sets with the weaker property of 2-independence.

Denote $\hat{p} = \frac{1}{\min(1-p,p)}$ for the rest of this section. $\hat{p}$

Lemma 3.5 *Let p be a constant, $0 < p < 1$, τ be a parameter, $\tau \geq 1$, and $\Lambda = (4(2\beta\tau)^{1/\alpha})^{\hat{p}}$. Let l_v be a link and let Q be a 2-independent set of non-weak links in an arbitrary metric space, where each link $l_w \in Q$ satisfies $\max(a_v^{\mathcal{P}}(w), a_w^{\mathcal{P}}(v)) \geq 1/\tau$ and $\ell_w \geq \Lambda \cdot \ell_v$. Then, $|Q| = O(\log\log\Delta)$.*

For affectances below the threshold of Lemma 3.5, we bound their contributions for each length group separately.

Lemma 3.6 *Let q be a positive real value and l_v be a link. Let S be a 2-independent and feasible set of non-weak links belonging to a single length-class of minimum length at least $q^{\hat{p}/\alpha} \cdot \ell_v$. Then, $b_v^p(S) \leq (\max_{l_w \in S} b_v^p(w)) + O(1/q)$.*

The following corollary will be useful.

Corollary 3.7 *Let S be a feasible set of non-weak nearly-equilength links. Then, S is $O(1)$-inductively independent.*

We are now ready to prove the core result, Thm. 3.4.

Proof: [of Thm. 3.4] Choose any $l_v \in L$ and any feasible subset $S \subseteq L$. We will show that $\hat{b}_v^p(S) = O(\log\log\Delta)$. By definition of $\hat{b}$, we can assume that all links in S are larger than l_v, since $\hat{b}$ is defined in such a way that all shorter links do not contribute to its value. With this assumption, $\hat{b}^p(S) = b^p(S)$. Use the independence-strengthening lemma (Lemma 2.1) to partition S into at most $\frac{2^{\alpha+1}}{\beta} + 1$ different 2-independent feasible sets. Let S' be one such set.

Let $D = \log\Delta(L)$. We say that a link l_w in S is *short* if $\ell_v \leq \ell_w < D^{\hat{p}/\alpha} \cdot \ell_v$ and *long* if $\ell_w \geq D^{\hat{p}/\alpha} \cdot \ell_v$. We partition S' into three sets:

S_1**:** Long links l_w with $b_v(w) \geq 1/D$,

S_2**:** Long links l_w with $b_v(w) < 1/D$, and

S_3**:** Short links.

We bound the affectance $b_v(S_i)$ of each set S_i separately. By Lemma 3.5, $|S_1| = O(\log\log \Delta(S_1))$ and thus $b_v(S_1) \leq 2|S_1| = O(\log\log \Delta(S)) = O(\log\log \Delta(L))$. Due to the choice of D, the set S_2 can be partitioned into D or less length classes. Each such class X satisfies the hypothesis of Lemma 3.6 with $q := D$ (recall that S_1 is a 2-independent subset of S'). This implies that $b_v(X) = O(1/D)$ and $b_v(S_2) = O(1)$. The set S_3 can be partitioned into $\log D \leq \frac{\hat{p}}{\alpha} \log\log \Delta(L)$ length groups. For each group X, we apply Lemma 3.6 with $q = 1$, giving that $b_v(X) = O(1)$, for a total of $b_v(S_3) = O(\log\log \Delta)$. Thus, $b_v(S') = b_v(S_1) + b_v(S_2) + b_v(S_3) = O(\log\log \Delta)$, and $b_v(S) \leq (\frac{2\alpha}{\beta} + 1) b_v(S') = O(\log\log \Delta)$. □

The second main result gives an optimal bound on the inductive independence number for $\mathcal{P}_p$. This improves the previous bound of $O(\log n)$ [35].

Theorem 3.8 *Fix a power assignment $\mathcal{P}_p$ for any $0 < p \leq 1$. Then any set L of non-weak links is $O(1)$-inductively independent under $\mathcal{P}_p$, i.e., $I_p^p(L) = O(1)$.*

The following lemma is the crucial element, after which basic computations arguments lead to Thm. 3.8 which we provide in Appendix C.

Lemma 3.9 *Let L be a $\mathcal{P}_p$-feasible set of non-weak links and l_v be a link (not necessarily in L). Then, $\hat{a}_v(L) = O(1)$.*

Proof: Let $\mathcal{L}(n)$ be the set of all $\mathcal{P}_p$-feasible sets of non-weak links of size n. Define $g(n)$ (a function of n) to be the "optimum upper bound" on $\hat{a}$, that is, $g(n) := \sup_{L \in \mathcal{L}(n)} \sup_{l_v} \hat{a}_v(L)$. Such a function exists, since $\hat{a}_v(L) \leq n$ for any set L of size n and any l_v. We claim that $g(n)$ is indeed $O(1)$, which implies the lemma. For contradiction, assume $g(n) = \omega(1)$.

Since $g(n) = \omega(1)$, we can choose a large enough n_0 such that both of the following hold:

1. There exists $L \in \mathcal{L}(n_0)$ and l_v such that:
$$\hat{a}_v(L) \geq \frac{1}{2} g(n_0) \ . \tag{2}$$
2. Define $f(n) = \frac{1}{2} 2^{\frac{1}{4c_3} g(n)}$. Then,
$$f(n_0) \geq (16 \cdot 3^\alpha \beta)^{1/(p\alpha)} \ . \tag{3}$$
Here c_3 is a fixed constant to be specified later.

We will prove our lemma by deriving a contradiction to Eqn. 2. To prove this, we partition the link set L into L_1 and L_2 where $L_1 := \{l_w : \ell_w \leq f(n_0) \cdot \ell_v\}$ and set $L_2 := L \setminus L_1$.

Claim 3.10 $\hat{a}_v(L_1) < \frac{1}{4} g(n_0)$.

Proof: By definition of $\hat{a}$, we can ignore links in L_1 smaller that l_v. Since the maximum length in L_1 is $\leq f(n_0) \cdot \ell_v$, the remaining links in L_1 can be divided into $\log f(n_0)$ length classes. For each length class C it holds that $a_v(C) \leq c_3$, by Corollary 3.7. Thus

$$a_v(L_1) \leq c_3 \log f(n_0) \stackrel{1}{=} c_3 \left(\frac{1}{4c_3} g(n_0) - 1 \right) < \frac{1}{4} g(n_0) \ ,$$

where we have used the definition of $f(n)$ in Equality 1. □

Claim 3.11 $\hat{a}_v(L_2) \le \frac{1}{4}g(n_0)$,

Proof: Consider $l_w \in L_2$ such that $d(s_v, s_w)$ is minimized. Denote this quantity by D. Let L_3 be the set of links in L_2 with receivers in $B(s_v, D/2)$ (the ball of radius $D/2$ around s_v), and set $L_4 := L_2 \setminus L_3$.

Let us first handle affectances to L_3 using the following (proof in Appendix C):

Proposition 3.12 $|L_3| \le 2 \cdot 4^\alpha + 1$.

Now using this proposition,

$$a_v(L_3 \cup \{l_w\}) \le |L_3| + 1 \le 2 \cdot (4^\alpha + 1) \le \frac{1}{8}g(n_0) \ .$$

The last inequality holds if n_0 is large enough (if not we can choose an larger n_0, the previous bounds will not be affected b this) since $g(n) = \omega(1)$.

We now consider $L_4 \setminus \{l_w\}$. Consider any $l_u \in L_4 \setminus \{l_w\}$. Using that r_u is at least $D/2$ away from s_v (due to being in L_4) and the fact that we chose $D := d(s_v.s_w)$, the triangle inequality yields $d(s_v, r_u) \ge \frac{1}{3}d(s_w, r_u)$. Thus,

$$a_v(L_4 \setminus \{\ell_w\}) \le \sum_{\ell_u \in L_4 \setminus \{\ell_w\}} c_u \cdot \frac{P_v}{d(s_v, r_u)^\alpha} \frac{\ell_u^\alpha}{P_u} \le 3^\alpha 2\beta \sum_u \frac{P_v}{P_w} \frac{P_w}{d(s_w, r_u)^\alpha} \frac{\ell_u^\alpha}{P_u} = 3^\alpha 2\beta \frac{P_v}{P_w} a_w(L_4) \ .$$

The first equality holds because l_w and l_u belong to the same feasible set, thus $a_w(u) = c_u \frac{P_w}{d(s_w,r_u)^\alpha} \frac{\ell_u^\alpha}{P_u}$. Next we use that $c_u \le 2\beta$ as we consider non-weak links.

Since the power function $\mathcal{P}_p$ is non-decreasing and $\ell_w \ge f(n_0) \cdot \ell_v$ due to the choice of $L_2 \supseteq L_4$, $P_w \ge \mathcal{P}_p(f(n_0) \cdot \ell_v) = f(n_0)^{p\alpha} P_v$. Thus, $\frac{P_v}{P_w} \le \frac{1}{f(n_0)^{p\alpha}} \le \frac{1}{16 \cdot 3^\alpha \beta}$ using Eqn. 3. By the definition of $g(n)$, $a_w(L_4) \le g(n_0)$. Thus,

$$a_v(L_4 \setminus \{l_w\}) \le 3^\alpha 2\beta \frac{1}{16 \cdot 3^\alpha \beta} g(n_0) \le \frac{1}{8} g(n_0) \ ,$$

This completes the proof of Claim 3.11. □

Combining Claims 3.10 and 3.11, we get that $a_v(L) \le \frac{1}{2}g(n_0)$, contradicting Eqn. 2. This completes the proof of Lemma 3.9. □

We remark that the bounds in both theorems do not hold when there are weak links. Specifically, we give a $\Omega(\log n)$-lower bound on inductive independence for weak links in Appendix E. Also, the assumption that $p > 0$ is also necessary for both theorems; a similar construction shows that inductive independence is $\Omega(\log n)$ for $\mathcal{P}_p$, where $p = 1 + o(1)$.

4 Capacity Approximation

Using the characterization described above, it is possible to derive a simple single-pass algorithm for maximizing capacity. This is, in fact, the same algorithm as used in [23] to maximize fixed power capacity within a constant factor. It is a type of a greedy algorithm that falls under the notion of "fixed priority", as defined by Borodin et al [6]. We prove that this simple algorithm even delivers the optimal oblivious-power approximation using $\mathcal{P}_p$ power assignments.

Theorem 4.1 *For any $\mathcal{P}_p$ for which L is non-weak,* ***Gr*** *chooses a set X that is $\mathcal{P}_p$-feasible such that $|X| \ge \frac{|R|}{2(2I_{\mathcal{Q}}^p(L)+1)}$ for any power assignment $\mathcal{Q}$ and any set $R \subseteq F_{\mathcal{Q}}(L)$.*

Algorithm 1 Gr(Set $L = \{l_1, l_2, \ldots, l_n\}$ of links in increasing order of length)

```
1: S_0 ← ∅
2: for i = 1 to n do
3:    if b̂^p_{S_{i-1}}(l_i) ≤ 1/2 then
4:       S_i ← S_{i-1} ∪ {l_i}
5:    end if
6: end for
7: X = {l_v ∈ S_n : a^p_{S_n}(v) ≤ 1}
```

Proof: The structure of the proof is inspired by that of, e.g., [32]. First we show that the size of R is not much larger than the size of S, second we relate the size of X to S and conclude the statement. Let $S = S_n$ and X be the sets computed by Algorithm **Gr** on input L. Consider any $\mathcal{Q}$ and R as specified by the statement of the theorem. Let R' be $R' := R \setminus S$.

By definition of $I^p_{\mathcal{Q}}(L)$, $\hat{b}^p_v(R) \le I^p_{\mathcal{Q}}(L)$, for each $l_i \in S$. Thus,

$$\hat{b}^p_S(R) \le I^p_{\mathcal{Q}}(L) \cdot |S| \ , \tag{4}$$

Due to lines 3 and 4 Algorithm **Gr** chose none of the links in R'. Using this and the definition of $\hat{b}^p$ yields that $\hat{b}^p_S(j) \ge \hat{b}^p_{S_{j-1}}(j) \ge 1/2$, for each $l_j \in R'$, implying that

$$\hat{b}^p_S(R') \ge |R'|/2 \ . \tag{5}$$

Combining (4) and (5),

$$|R'| \le 2 \cdot \hat{b}^p_S(R') \le 2 \cdot \hat{b}^p_S(R) \le 2I^p_{\mathcal{Q}}(L) \cdot |S| \ .$$

Thus,

$$|R| \le |R'| + |S| \le (2I^p_{\mathcal{Q}}(L) + 1)|S| \ . \tag{6}$$

Also, the definition of **Gr** ensures that the average affectance of links in S is small (at most half). To see this, observe that,

$$\begin{aligned}\sum_{l_v \in S} a_S(v) = \sum_{l_i \in S}\sum_{l_j \in S} a_j(i) &\overset{1}{=} \sum_{l_i \in S}\sum_{l_j \in S : j<i} (a_j(i) + a_i(j)) \\ &\overset{2}{=} \sum_{l_i \in S}\sum_{l_j \in S : j<i} \hat{b}_j(i) \overset{3}{=} \sum_{l_i \in S} \hat{b}_{S_{i-1}}(i) \le \frac{1}{2}|S| \ ,\end{aligned}$$

which implies that the average affectance $A^p(S)$ is $\frac{1}{|S|}a_S(S) \le \frac{1}{2}$. Explanation of numbered (in)equalities in the above computation are as follows:

1. By rearrangement. Here $j < i$ refers to the indices of the links as sorted by Algorithm **Gr**. We also use the fact that by the definition of affectance it is $\sum_{l_i \in S} a_i(i) = 0$.

2. By the way Algorithm **Gr** iterates over the links, $j < i$ implies that $\ell_j \le \ell_i$. Thus $\hat{b}_j(i) = a_j(i) + a_i(j)$, by definition of $\hat{b}$.

3. Since $S_{i-1} = \{l_j : l_j \in S, j < i\}$ as specified by Algorithm **Gr**.

At least half the links will have at most double the average affectance, or

$$|X| = |\{l_v \in S | a_S(v) \le 1\}| \ge \frac{1}{2}|S| \ . \tag{7}$$

Combining (6) and (7) yields the claim. □

Theorem 4.2 *There is a $O(\log\log\Delta)$-approximation algorithm for* PC-Capacity *that uses $\mathcal{P}_p$.*

Proof: We consider $\overline{OPT}$, a maximum capacity solution with arbitrary power, and a power assignment $\mathcal{Q}$ that makes $\overline{OPT}$ feasible. We can apply Thm. 4.1 to note that **Gr** produces a $O(1 + I^p_{\mathcal{Q}}(L))$ solution. This observation proves the statement. □

When there is a maximum power level and most links are weak, we can still attain the same approximation ratio, as done in [23], by solving the problem separately for the weak links using maximum power.

5 Applications

Apart from the optimal algorithm for capacity approximation in the previous section, both of our structural results have numerous applications, improving the approximation ratio for many fundamental and important problems in wireless algorithms. All our improvements come from noticing that many existing approximation algorithms have bounds that are implicitly based on $I^p_{\mathcal{Q}}(L)$ or $I^p_{\mathcal{P}}(L)$ (or even the combination of both). Plugging in our improved bounds for these thus gives the (poly)-logarithmic improvements for a wide variety of applications. We will often omit proofs of our claims, as they are all of the same flavor. We indicate some of these implications below.

Connectivity

Wireless connectivity — the problem of *efficiently* connecting a set of wireless nodes in an interference aware manner — is one of the most important problems in wireless network research [25]. Such a structure may underlie a multi-hop wireless network, or provide the underlying backbone for synchronized operation of an adhoc network. In a wireless sensor network, the structure can function as an information aggregation mechanism.

Recent results have shown that any set of wireless nodes can be strongly connected in $O(\log n \cdot (\log n + \log\log\Delta))$ slots using mean power, using both centralized [25] and distributed [24] algorithms. These results are directly improved by Thm. 4.2:

Theorem 5.1 *Any set of links can be strongly connected in $O(\log n \cdot \log\log\Delta)$ slots using power assignment $\mathcal{P}_p$. This can be computed by either a poly-time centralized algorithm or a $O(poly(\log n)\log\Delta)$-time distributed algorithm.*

Results for variations of connectivity such as *minimum-latency aggregation scheduling* and applications of connectivity such as maximizing the aggregation rate in a sensor network benefit from similar improvements. We refer the reader to [25] for a discussion of these problems and their numerous applications.

Distributed Scheduling

A fundamental problem in wireless algorithms is to schedule a given set of links in a minimum number of slots. $O(\log n)$ centralized algorithms for $\mathcal{P}_p$-Scheduling, the version with given power assignment $\mathcal{P}_p$. This is obtained via repeated application of an algorithm that schedules a $O(1)$ factor of the links using a subroutine that solves the capacity problem (e.g. [23]). In [35], a first distributed algorithm was given, with a $O(\log^2 n)$ approximation ratio. Since this is a distributed algorithm, the algorithm included an acknowledgement mechanism (via packets sent from receivers to senders) to enable links to know when they have succeeded (and subsequently stop running the algorithm). Assuming "free" acknowledgements, [22] improved the bound to

$O(\log n)$ (using the same algorithm), but [35] remained the best result when acknowledgements have to be implemented explicitly.

To examine this in more detail, we must introduce another complexity measure.

Definition 5.2 *[35] The* maximum average affectance $A^p(L)$ *of a link set* L *is* $A^p(L) := \max_{R \subseteq L} \frac{a_R^p(R)}{|R|}$.

$A^p(L)$

It is easily verified that $A^p(L) = O(I_Q^p(L) \cdot \overline{\chi(L)})$, where $\overline{\chi(L)}$ denotes the minimum number of slots in a feasible schedule of L (using arbitrary power). Similarly $A^p(L) = O(I_p^p(L) \cdot \chi^p(L))$ where $\chi^p(L)$ denotes the minimum number of slots in a $\mathcal{P}_p$-feasible schedule of L.

Corollary 5.3 *For any linkset* L, $A^p(L) = O(\log\log\Delta \cdot \overline{\chi(L)})$ *and* $A^p(L) = O(\chi^p(L))$.

Specifically, it was shown in [35] that the distributed scheduling algorithm completes in $O(A^p(L)\log n)$ rounds, and furthermore that $A^p(L) = O(\chi^p(L)\log n)$ (which we now improve for $0 < p \le 1$). The present work implies a $O(\log n)$ approximation (with acknowledgements). Using Corollary 5.3), the original analysis of [35] can be applied to achieve a $O(\log n)$ approximation factor using power assignment $\mathcal{P}_p$. Along with [22], which covers the case for uniform power, this solves the distributed scheduling problem optimally for all relevant polynomial powers (mean power, in particular). This matches the best bound known for centralized algorithms.

Corollary 5.4 *There is a randomized distributed* $O(\log n)$*-approximate algorithm for* $\mathcal{P}_p$*-Scheduling, for any* $0 < p < 1$.

For comparison with arbitrary power, we can similarly use Corollary 5.3 to achieve a $O(\log n \cdot \log\log\Delta)$ approximation including acknowledgements, improving on the $O(\log n \cdot (\log n + \log\log\Delta))$-factor implied by [35] and [23]. Let PC-Scheduling be the power-control version of the problem. Refer to the power-control

Corollary 5.5 *There is a randomized distributed algorithm for PC-Scheduling that is* $O(\log\log\Delta \cdot \log n)$*-approximate with respect to arbitrary power control optima. It can use any* $\mathcal{P}_p$ *power assignment,* $0 < p < 1$.

Spectrum sharing auctions

In light of recent regulatory changes by the Federal Communications Commission (FCC) opening up the possibility of dynamic white space networks (see, for example, [5]), the problem of dynamic allocation of channels to bidders (these are the wireless devices) via an auction has become highly important [49, 50].

The combinatorial auction problem in the SINR model is as follows: Given k identical channels and n users (links), with each user having a valuation for each of the 2^k possible subset of channels, find an allocation of the users to channels so that each channel is assigned a feasible set and the social welfare is maximized.

For the SINR model, recent work [28, 27] has established a number of results depending on different valuation functions. Since these results are based on the inductive independence number, Thm. 3.8 improves virtually all of them by a $\log n$ factor. For instance, an algorithm was given in [28] for general valuations that achieves a $O(\sqrt{k}\log n \cdot I_p^p(L)) = O(\sqrt{k}\log^2 n)$-approximation. We achieve an improved result by simply plugging in Thm. 3.8.

Corollary 5.6 *Consider the combinatorial auction problem in the SINR setting, for any fixed power assignment* $\mathcal{P}_p$ *with* $0 < p \le 1$. *There exist algorithms that achieve* $O(\sqrt{k}\log n)$*-factor for general valuations [28],* $O(\log n + \log k)$ *approximation for symmetric valuations and* $O(\log n)$ *approximation for Rank-matroid valuations [27].*

Further applications on dynamic packet scheduling are given in Appendix D.

References

[1] K. Akcoglu, J. Aspnes, B. DasGupta, and M.-Y. Kao. Opportunity cost algorithms for combinatorial auctions. *CoRR*, cs.CE/0010031, 2000.

[2] M. Andrews and M. Dinitz. Maximizing capacity in arbitrary wireless networks in the SINR model: Complexity and game theory. In *INFOCOM*, pages 1332–1340. IEEE, 2009.

[3] E. I. Ásgeirsson, M. M. Halldórsson, and P. Mitra. Wireless network stability in the SINR model. In *SIROCCO*, 2012.

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

[5] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh. White space networking with Wi-Fi like connectivity. In *SIGCOMM*, pages 27–38, 2009.

[6] A. Borodin, M. N. Nielsen, and C. Rackoff. (Incremental) priority algorithms. *Algorithmica*, 37(4):295–326, 2003.

[7] D. Chafekar, V. Kumar, M. Marathe, S. Parthasarathy, and A. Srinivasan. Cross-layer latency minimization for wireless networks using SINR constraints. In *Mobihoc*, 2007.

[8] J. Dams, M. Hoefer, and T. Kesselheim. Scheduling in wireless networks with rayleigh-fading interference. In *Proceedings of the 24th ACM symposium on Parallelism in algorithms and architectures (SPAA)*, 2012.

[9] M. Dinitz. Distributed algorithms for approximating wireless network capacity. In *Infocom*, 2010.

[10] A. Fanghänel, S. Geulen, M. Hoefer, and B. Vöcking. Online capacity maximization in wireless networks. In *SPAA*, pages 92–99. ACM, 2010.

[11] A. Fanghänel, T. Kesselheim, H. Räcke, and B. Vöcking. Oblivious interference scheduling. In *PODC*, pages 220–229, August 2009.

[12] A. Fanghänel, T. Kesselheim, and B. Vöcking. Improved algorithms for latency minimization in wireless networks. In *ICALP*, pages 447–458, July 2009.

[13] Y. Gao, J. C. Hou, and H. Nguyen. Topology control for maintaining network connectivity and maximizing network capacity under the physical model. In *Infocom*, 2008.

[14] O. Goussevskaia, M. M. Halldórsson, R. Wattenhofer, and E. Welzl. Capacity of Arbitrary Wireless Networks. In *INFOCOM*, pages 1872–1880, April 2009.

[15] O. Goussevskaia, T. Moscibroda, and R. Wattenhofer. Local Broadcasting in the Physical Interference Model. In *DialM-POMC*, August 2008.

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

[17] O. Goussevskaia, Y. A. Pignolet, and R. Wattenhofer. Efficiency of wireless networks: Approximation algorithms for the physical interference model. *Foundations and Trends in Networking*, 4(3):313–420, 2010.

[18] O. Goussevskaia and R. Wattenhofer. Complexity of scheduling with analog network coding. In *FOWANC*, May 2008.

[19] J. Grönkvist and A. Hansson. Comparison Between Graph-Based and Interference-Based STDMA Scheduling. In *Mobihoc*, pages 255–258, 2001.

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

[21] M. M. Halldórsson. Wireless scheduling with power control. arXiv:1010.3427v2. To appear in *ACM Transactions on Algorithms*. Earlier version appears in *ESA '09*.

[22] M. M. Halldórsson and P. Mitra. Nearly optimal bounds for distributed wireless scheduling in the SINR model. In *ICALP*, 2011.

[23] M. M. Halldórsson and P. Mitra. Wireless Capacity with Oblivious Power in General Metrics. In *SODA*, 2011.

[24] M. M. Halldórsson and P. Mitra. Distributed connectivity of wireless networks. In *PODC*, 2012.

[25] M. M. Halldórsson and P. Mitra. Wireless connectivity and capacity. In *SODA*, 2012.

[26] M. M. Halldórsson and R. Wattenhofer. Wireless Communication is in APX. In *ICALP*, pages 525–536, July 2009.

[27] M. Hoefer and T. Kesselheim. Secondary spectrum auctions for symmetric and submodular bidders. In *ACM Conference on Electronic Commerce*, pages 657–671, 2012.

[28] M. Hoefer, T. Kesselheim, and B. Vöcking. Approximation algorithms for secondary spectrum auctions. In *SPAA*, 2011.

[29] Q.-S. Hua and F. C. M. Lau. The scheduling and energy complexity of strong connectivity in ultra-wideband networks. In *MSWiM*, pages 282–290, 2006.

[30] E. Kantor, Z. Lotker, M. Parter, and D. Peleg. The Topology of Wireless Communication. In *STOC*, 2011.

[31] B. Katz, M. Volker, and D. Wagner. Energy efficient scheduling with power control for wireless networks. In *WiOpt*, pages 160–169. IEEE, 2010.

[32] T. Kesselheim. A Constant-Factor Approximation for Wireless Capacity Maximization with Power Control in the SINR Model. In *SODA*, 2011.

[33] T. Kesselheim. Approximation algorithms for wireless link scheduling with flexible data rates. In *ESA*, 2012.

[34] T. Kesselheim. Dynamic packet scheduling in wireless networks. In *PODC*, 2012.

[35] T. Kesselheim and B. Vöcking. Distributed contention resolution in wireless networks. In *DISC*, pages 163–178, August 2010.

[36] V. Kumar, M. Marathe, S. Parthasarathy, and A. Srinivasan. Algorithmic Aspects of Capacity in Wireless Networks. In *Sigmetrics*, 2005.

[37] L. B. Le, E. Modiano, C. Joo, and N. B. Shroff. Longest-queue-first scheduling under SINR interference model. In *MOBIHOC*, 2010.

[38] Z. Lotker and D. Peleg. Structure and algorithms in the SINR wireless model. *SIGACT News*, 41(2):74–84, 2010.

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

[40] T. Moscibroda. The worst-case capacity of wireless sensor networks. In *IPSN*, pages 1–10, 2007.

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

[42] T. Moscibroda, R. Wattenhofer, and Y. Weber. Protocol design beyond graph-based models. In *Hotnets*, November 2006.

[43] T. Moscibroda, R. Wattenhofer, and A. Zollinger. Topology control meets SINR: The scheduling complexity of arbitrary topologies. In *MOBIHOC*, pages 310–321, 2006.

[44] 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 *MobiHoc*, pages 91–100, 2008.

[45] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Automat. Contr.*, 37(12):1936–1948, 1992.

[46] T. Tonoyan. On the capacity of oblivious powers. In *ALGOSENSORS*, pages 225–237, 2011.

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

[48] Y. Ye and A. Borodin. Elimination graphs. *ACM Transactions on Algorithms*, 8(2):14, 2012.

[49] X. Zhou, S. Gandhi, S. Suri, and H. Zheng. eBay in the Sky: Strategy-proof wireless spectrum auctions. In *MobiCom*, pages 2–13, New York, NY, USA, 2008. ACM.

[50] X. Zhou and H. Zheng. Trust: A general framework for truthful double spectrum auctions. In *INFOCOM*, pages 999–1007, 2009.

A Missing Proof from Sec. 2: Independence Strengthening

Lemma 2.1 *Any feasible set of links can be partitioned into* $\lfloor \frac{2q^\alpha}{\beta} \rfloor + 1$ *or less different* q*-independent sets.*

Proof: Let S be a feasible set and $\mathcal{P}$ a power assignment such that S is feasible for $\mathcal{P}$. We form a graph G on linkset S, such that two links l_v and l_w are adjacent if $b_v^{\mathcal{P}}(w) \geq \beta/q^\alpha$. Let Z be $Z := \lfloor 2q^\alpha/\beta \rfloor$.

We first show that G is Z-inductive (a.k.a. Z-degenerate, or Szekeres-Wilf number Z), which means that there is an ordering of the vertices so that each vertex has at most Z neighbors that appear later in the ordering.

Since S is feasible, $a_S^{\mathcal{P}}(v) \leq 1$, for any l_v in S. Thus, $b_S^{\mathcal{P}}(S)/2 = a_S^{\mathcal{P}}(S) \leq |S|$, so some link l_u satisfies

$$b_u^{\mathcal{P}}(w) \leq 2 \ .$$

It is then clear that for at most $Z = \lfloor 2q^\alpha/\beta \rfloor$ links l_w does it hold that $b_u^{\mathcal{P}}(w) \geq \beta/q^\alpha$. We then form a Z-inductive ordering of S by placing l_u first, followed by the inductively constructed ordering for $S \setminus \{l_u\}$.

Since G is Z-inductive, it is $Z+1$-colorable. Consider a color class (a stable set) C. It holds by definition for any pair l_v, l_w of links in C that

$$a_w^{\mathcal{P}}(v) \cdot a_v^{\mathcal{P}}(w) \leq \frac{\beta}{q^\alpha} \cdot \frac{\beta}{q^\alpha} \leq \frac{c_v c_w}{q^{2\alpha}} \ ,$$

which implies that l_v and l_w are q-independent. Quantifying over all pairs in C, it follows that C is q-independent. □

B Full Proof of Thm. 3.4

We prove here the following main theorem.

Theorem 3.4 *For any set L of non-weak links, any $0 < p < 1$, and any power assignment $\mathcal{Q}$,*

$$I_{\mathcal{Q}}^p(L) = O(\log\log \Delta) \ .$$

To argue the theorem, we need two lemmas to bound affectances of a link to and from a set of links. The first handles the long links in S with relatively high affectance. It originates in [21] (Lemma 4.4), but is generalized here in two ways: to any power assignment $\mathcal{P}_p$, and to sets with the weaker property of 2-independence.

Definition B.1 *We say that links l_v and l_w are t-close under power assignment $\mathcal{P}$ if,*

$$\max(a_v^{\mathcal{P}}(w), a_w^{\mathcal{P}}(v)) \geq t.$$

For the rest of this section, denote $\hat{p} := \frac{1}{\min(1-p,p)}$.

Lemma 3.5 *Let p be a constant, $0 < p < 1$, τ be a parameter, $\tau \geq 1$, and $\Lambda = (4(2\beta\tau)^{1/\alpha})^{\hat{p}}$. Let l_v be a link and let Q be a 2-independent set of non-weak links in an arbitrary metric space, that are both $\frac{1}{\tau}$-close to l_v under power assignment $\mathcal{P}_p$ and at least a Λ-factor longer than l_v. 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 power assignment $\mathcal{P}_p$, and those that are affected by l_v by that amount. We consider first the links of the former type.

Consider a pair $l_w, l_{w'}$ in Q that affect $l_v$by at least $1/\tau$ under $\mathcal{P}_p$, and suppose without loss of generality that $\ell_w \geq \ell_{w'}$. Let l_1 be the shortest link in Q. The affectance of l_w on l_v implies that

$$c_v \left(\frac{\ell_w^p \ell_v^{1-p}}{d_{wv}} \right)^\alpha \geq \frac{1}{\tau} \ ,$$

which can be transformed to $d_{wv} \leq \ell_w^p \ell_v^{1-p} (c_v\tau)^{1/\alpha}$, and similarly, $d_{w'v} \leq \ell_{w'}^p \ell_v^{1-p} (c_v\tau)^{1/\alpha}$. Recall that since l_v is non-weak, $c_v \leq 2\beta$. By the triangular inequality, we have that

$$\begin{aligned} d_{w'w} &\leq d(s_{w'}, r_v) + d(r_v, s_w) + d(s_w, r_w) \\ &= d_{w'v} + d_{wv} + \ell_w \\ &\leq 2\ell_w^p \ell_v^{1-p} (c_v\tau)^{1/\alpha} + \ell_w \\ &\leq 2\ell_w^p \ell_v^{1-p} (2\beta\tau)^{1/\alpha} + \ell_w \\ &\leq \ell_w^p \ell_1^{1-p} + \ell_w \leq 2\ell_w \ , \end{aligned} \tag{8}$$

using that $\Lambda\ell_v \le \ell_1 \le \ell_w$. Similarly,

$$d_{ww'} \le \ell_{w'} + \frac{1}{2}\ell_w^p \ell_1^{1-p} \ . \tag{9}$$

Applying 2-independence, on one hand, and multiplying (8) and (9), on the other, we obtain that

$$4\ell_w\ell_{w'} \le d_{w'w} \cdot d_{ww'} \le 2\ell_{w'}\ell_w + \ell_w^p \ell_1^{1-p} \cdot \ell_w \ , \tag{10}$$

Cancelling a $2\ell_w$-factor, simplifying and rearranging, we have that

$$\ell_w^p \ge \frac{2\ell_{w'}}{\ell_1^{1-p}} \ . \tag{11}$$

Label the links in Q as $l_1, l_2, \ldots, l_{|Q|}$ in increasing order of length, and define $\lambda_i = l_i/l_1$. By dividing both sides of (11) by ℓ_1^p, we get that

$$\lambda_{i+1}^p \ge 2\lambda_i \ .$$

Then, $\lambda_2 \ge 2^{1/p}$ and by induction $\lambda_t \ge 2^{(1/p)^{t-1}}$. Note that $\Delta(Q) = l_{|Q|}/l_1 = \lambda_{|Q|} \ge 2^{(1/p)^t}$, so $|Q| - 1 \le \log_{1/p} \log_2 \Delta$, and the claim follows.

The other case of links l_w with $a_v(w) \ge 1/\tau$ is symmetric, with the roles of p and $1-p$ switched, leading to a bound of $1 + \log_{1/(1-p)} \lg \Delta$. □

We shall in particular apply the lemma with $\tau = \log \Delta$.

Lemma 3.5 bounds the number of longer links that affect a given link by a significant amount. For affectances below that threshold, we bound their contributions for each length group separately.

We first need the following geometric argument. Intuitively, we want to convert statements involving the link l_v into statements about appropriate links within the 2-independent set S.

Lemma B.2 Let l_v be a link. Let S be a 2-independent set of nearly-equilength links and l_u be the link in S with d_{uv} minimum. Then, $\max(d_{wu}, d_{uw}) \le 6d_{wv}$, for any link l_w in S.

Proof: Let $D = d_{wv}$ and note that by definition $d_{uv} \le D$. By the triangular inequality and the definition of l_u,

$$d_{wu} \le d(s_w, r_v) + d(r_v, s_u) + d(s_u, r_u) = d_{wv} + d_{uv} + \ell_u \le 2D + \ell_u \ . \tag{12}$$

Similarly,

$$d_{uw} \le d_{uv} + d_{wv} + \ell_w \le 2D + \ell_w \ . \tag{13}$$

Applying 2-independence, on one hand, and multiplying (12) and (13), on the other hand, we have that

$$4\ell_u\ell_w \le d_{wu} \cdot d_{uw} < (2D + \ell_u) \cdot (2D + \ell_w) \ .$$

It is then easily verified that $D \ge \min(\ell_u, \ell_w)/2 \ge \max(\ell_u, \ell_w)/4$, using that the links are nearly-equilength. The claim then follows from (12) and (13). □

Lemma 3.6 *Let q be a positive real value and l_v be a link. Let S be a 2-independent and feasible set of non-weak nearly-equilength links of minimum length at least $q^{\hat{p}/\alpha} \cdot \ell_v$. Then, $b_v^p(S) \le (\max_{l_w \in S} b_v^p(w)) + 1/O(q)$.*

Proof: Consider the link l_u in S with d_{uv} minimum. By Lemma B.2, $\max(d_{wu}, d_{uw}) \leq 6d_{wv}$, for any link l_w in S.

Since $\ell_v \leq \ell_u$, it holds that $c_v \leq c_u$. Then, we have that

$$a_w^p(v) = c_v \left(\frac{\ell_v^{1-p}\ell_w^p}{d_{wv}}\right)^\alpha \leq c_u \left(\frac{(\ell_u/q^{\hat{p}/\alpha})^{1-p}\ell_w^p}{d_{wu}/6}\right)^\alpha = \frac{6^\alpha}{q^{\hat{p}\cdot(1-p)}} a_w^p(u) \leq \frac{6^\alpha}{q} a_w^p(u) \ .$$

Also, using that the links in S are non-weak, $c_w \leq 2c_u$, and

$$a_v^p(w) = c_w \left(\frac{\ell_v^p \ell_w^{1-p}}{d_{vw}}\right)^\alpha \leq 2c_u \left(\frac{(\ell_u/q^{\hat{p}/\alpha})^p \ell_w^{1-p}}{d_{wu}/6}\right)^\alpha \leq 2c_u \frac{6^\alpha}{q^{\hat{p}\cdot p}} \left(\frac{2\cdot \ell_w^{1-p}\ell_u^p}{d_{wu}}\right)^\alpha \leq 2\frac{2^\alpha\cdot 6^\alpha}{q} a_w^p(u) \ ,$$

where we use in the second-to-last inequality that the links are nearly-equilength.

Thus,

$$b^p(S,v) - b^p(u,v) = a^p_{S\setminus\{u\}}(v) + a^p_v(S\setminus\{u\}) \leq (1+2^{\alpha+1})\frac{6^\alpha}{q} a_u^p(S) \leq (1+2^{\alpha+1})\frac{6^\alpha}{q} \ ,$$

where the last inequality uses the feasibility of S. □

C Remainder of the Proof of Thm. 3.8

We present here the proof of Thm. 3.8.

Theorem 3.8 *Fix $\mathcal{P}_p$ for any $0 < p \leq 1$. Then any set L of non-weak links is $O(1)$-inductively independent under $\mathcal{P}_p$, i.e., $I_p^p(L) = O(1)$.*

Proof: Consider any $S \in F_p(L)$ and any $l_v \in L$. We will show that $\hat{b}_v^P(S) = O(1)$, proving the theorem. By definition, $\hat{b}_v^P(S) \leq \sum_{l_w \geq l_v} a_v(w) + \sum_{l_w \geq l_v} a_w(v)$. For the first term we obtain that

$$\sum_{l_w \geq l_v} a_v(w) = \hat{a}_v(L) = O(1).$$

In the above transformation, the first equality is due to the definition of $\hat{b}_v^P(S)$ and the second is from Lemma 3.9. The second sum $\sum_{l_w \geq l_v} a_w(v)$ is known to be $O(1)$ (Lemma 7, [35]). The proof is completed. □

Proof of Proposition 3.12

Proof: By Lemma 2.1, L_3 can be divided into $2\cdot 4^\alpha + 1$ sets, each of which is 4-independent. For contradiction, if $|L_3| > 2\cdot 4^\alpha + 1$, then at least one of these sets must be of size at least 2. Thus, there would be two different links l_x and l_y that are members of L_3 and are 4-independent.

However, since $l_x, l_y \in L_3$, we can argue that

$$d(x,y) \overset{1}{\leq} \ell_x + d(r_x, r_y) \overset{2}{\leq} \ell_x + D \overset{3}{\leq} \ell_x + 2\ell_x \leq 3\ell_x \ ,$$

Explanation of numbered inequalities:

1. By triangle inequality.
2. Observing that both r_x and r_y are in $B(s_v, D/2)$ (due to the definition of L_3) and using triangle inequality.
3. Since $\ell_x = d(s_x, r_x) \geq D/2$ as $r_x \in B(r_v, D/2)$ (since $l_x \in L_3$) and $d(s_x, r_v) \geq D$ (by definition of D)

We can similarly show that $d(y,x) \leq 3\ell_y$. Then $d(x,y)\cdot d(y,x) \leq 9\ell_x\ell_y$, contradicting 4-independence. □

D Additional Applications

Dynamic Packet Scheduling

Dynamic packet scheduling to achieve network *stability* is one of the fundamental problems in (wireless) network queueing theory [45]. In spite of its long history, this fundamental problem has been considered only recently in the SINR model (see [37, 34, 3]). The problem calls for an algorithm that can keep queue sizes bounded in a wireless network under stochastic arrivals of packets at senders. A measure called *efficiency* between 0 and 1 is used to capture how well a given algorithm does compared to a hypothetical best algorithm. We refer the reader to the aforementioned papers for exact definitions and motivations related to this problem.

The state-of-the-art results for this problem have been achieved very recently and simultaneously in [3] and [34]. In spite of differences in the algorithm and assumptions made, both are based on the scheduling algorithm of [35] and achieve a similar result. Recall that the maximum average affectance is $A^p(L) = \max_{R\subseteq L} \frac{a^p_R(R)}{|R|}$ and $\chi^p(L)$ is the minimum number of slots in a $\mathcal{P}_p$-feasible schedule of L. Let $\phi(L) = \frac{A^p(L)}{\chi^p(L)}$.

The result in [34, 3] can be succinctly expressed as follows.

Theorem D.1 *[34, 3] There exists a distributed algorithm that achieves* $\Omega\left(\frac{1}{\log n\cdot(1+\phi(L))}\right)$*-efficiency for any link set* L*.*

Since the best bound on $\phi(L)$ known was $O(\log n)$ [35], both papers claimed $\Omega(\frac{1}{\log^2 n})$-efficiency. Results in this paper show that $\phi(L) = O(1)$ (Corollary 5.3), we get the following improved result:

Corollary D.2 *There exists a distributed algorithm that achieves* $\Omega\left(\frac{1}{\log n}\right)$*-efficiency for any power assignment* $\mathcal{P}_p$ *(*$0 < p \leq 1$*).*

Since Corollary 5.3 also shows that $\overline{\phi(L)} = \frac{A^p(L)}{\overline{\chi(L)}} = O(\log n \cdot \log\log\Delta)$, we also get the following improved bound for power control:

Corollary D.3 *There is a distributed algorithm with* $\Omega(\frac{1}{\log n\cdot\log\log\Delta})$*-efficiency, with respect to power control optima.*

Multi-hop Scheduling

The following constitute logarithmic improvements over [35]:

Corollary D.4 *There is a distributed algorithm for multi-hop scheduling that runs in* $O(\chi^p(L)\log n + D\log^2 n)$ *slots, where* D *is the maximum pathlength. Also, there exists a centralized* $O(\log n)$*-approximation of the combined problem of routing and multi-hop scheduling.*

E The Case of Weak Links

Let us recall that a link l_v is a weak link if $\frac{P_v}{l_v^\alpha} < 2\beta$. As we have mentioned before, intuitively, a non-weak link is one that succeeds with some slack in the presence of ambient noise N only. An alternative interpretation of this statement is that ambient noise plays a minor role. This is true for many realistic settings.

In any case, for our first result relating oblivious to arbitrary power (Thm. 3.4), the restriction to non-weak links is not relevant. After all, the problem is about power *control*, thus any algorithm can easily choose a power level high enough to ensure non-weakness of links, simply by scaling up the (oblivious) power used.

For our second result (Thm. 3.8), the assumption *is* relevant. Non-weakness is a reasonable and well-known assumption. Indeed, all of the applications improved by Thm. 3.8 refer to claims in [35], which as stated apply to non-weak links only. It has been understood since then, that the results of [35] apply to weak links as well. In our case however, we will show that both of the assumptions for the positive result in Thm. 3.8 are necessary: that $p > 0$ and that the links are non-weak.

We modify a construction that originally was given in [22] for uniform power assignments and adopt it to our setting. The basic construction is as follows: There are links $Z = \{l_1, \ldots, l_n\}$, with the length ℓ_i of link l_i being $(i+1)^{1/\alpha}$. The distance from l_i to l_j with $i < j$ is $(c(j+1))^{2/\alpha}$, for a constant c to be specified. Namely, $d_{ji}^\alpha = d_{ij}^\alpha = (2(\max(i,j)+1))^2$, for any i, j. It is straightforward to verify that this yields a metric instance. The construction satisfies the property that for $i > j$ under uniform power, $a_i(j) = \theta\left(\frac{j}{i^2}\right)$. Thus suffices to ensure that the out-affectance from link 1 is a harmonic sum, $a_1(Z) = \Omega(\log n)$, and that the in-affectance of any link is at most 1, for appropriately chosen c.

We first show that the inductiveness cannot be bounded by a constant when using power functions that grow slower than a polynomial of positive degree. The key property needed is that affectance from a short link is significantly less than that from a long link at the same location, and that this holds even when link lengths are scaled up uniformly. This scale-free property is not shared by functions with $f(x) = x^{o(1)}$, which allows us to apply known lower bounds for the case of uniform power.

Theorem E.1 *Let $\mathcal{P}$ be a power assignment with $p(\ell) = \ell^{1/h(\ell)}$, for $h = \omega(1)$. Then, there is a set L of n non-weak links, $\mathcal{P}$-feasible, satisfying $\max_{l \in L} \hat{a}_l^{\mathcal{P}}(L) = \Omega(\log n)$. Thus, $I_{\mathcal{P}}^p(L) = \Omega(\log n)$.*

Proof: This result was shown to hold for uniform power in the full version of [22]. We show how to scale up that instance so that affectances differ only by a constant factor from the case with uniform power.

Define the inverse $h^{-1}(y)$ as the smallest x such that $h(x) \geq y$, and observe that it is defined on all positive reals. Given n, we scale the instance Z_n by the factor $Q = h^{-1}(\log n)$. Let $N = \ell_n = n^{1/\alpha}$. Then,

$$\frac{p(l_n)}{p(l_1)} = \frac{(QN)^{1/h(QN)}}{Q^{1/h(Q)}} \leq N^{1/h(QN)} \leq N^{1/\log n} \leq 2^{1/\alpha} \ .$$

Thus, powers and affectances differ from the uniform power case by at most a constant factor. Hence, the lower bound of $\Omega(\log n)$ applies. $\square$

We next show that weak links do not have the constant inductiveness property.

Theorem E.2 *There is a set L of weak links, that is $\mathcal{P}_p$-feasible, satisfying $\max_{l \in L} \hat{a}_l^p(L) = \Omega(\log n)$, for any $0 \leq p \leq 1$. Thus, $I_p^p(L) = \Omega(\log n)$.*

Proof: We simulate the instance Z_n using weak links. The inter-link distances remain identical, modulo an appropriate scaling factor. The variations in link lengths are simulated by lengths very close to maximum possible length.

Specifically, link i is assigned length ℓ_i such that $c_i = (i+1)\beta$. That is, $\ell_i^{(1-p)\alpha} = \frac{i}{(i+1)\beta N}$. Set $X = (\frac{1}{\beta N})^{1/(1-p)}$. Then $d_{ji}^\alpha = d_{ji}^\alpha = X \cdot (c \max i, j+1)^2$. Also,

$$\frac{\ell_i^\alpha}{d_{ji}^\alpha} = \frac{i/(i+1)}{(c(\max(i,j)+1))^2} \ ,$$

and

$$\frac{P_i}{P_j} = \left(\frac{\ell_i}{\ell_j}\right)^{p\alpha} = \frac{[i/(i+1)]^{p/(1-p)}}{[j/(j+1)]^{p/(1-p)}} \leq 2^{p/(1-p)} \ .$$

Then,

$$a_i^p(j) = (j+1)\beta\left(\frac{i}{i+1}\cdot\frac{j+1}{j}\right)^{p/(1-p)} \cdot \frac{j}{(j+1)c(\max(i,j)+1)^2} \ .$$

In particular, sett $c = 2\beta 2^{p/(1-p)}$,

$$a_i^p(j) \geq \frac{(j+1)\beta 2^{p/(1-p)}}{c(\max(i,j)+1)^2} = \frac{j+1}{2(\max(i,j)+1)^2} \ .$$

Thus, we see that Z is feasible: for any j,

$$a_Z(j) \leq \sum_{k<j}\frac{1}{2(j+1)} + \sum_{k>j}\frac{1}{2(k+1)}^2 < \frac{1}{2} + \frac{j+1}{2}\sum_{k=j+1}^{\infty}\frac{1}{k(k+1)} = 1 \ .$$

Also, the out-affectance from link 1 is large:

$$a_1(Z) \geq \sum_{j=2}^{n}\frac{c_2}{j+1} = c_2(H_n - 3/2) = \Omega(\log n) \ ,$$

where $c_2 \geq \beta 2^{-(p/(1-p)+1)}/c$ and $H_n = \sum_{i\geq 1} 1/i$ is the harmonic number. $\square$