

THE POWER OF OBLIVIOUS WIRELESS POWER*

MAGNÚS M. HALLDÓRSSON[†], STEPHAN HOLZER[‡], PRADIPTA MITRA[§], AND ROGER
WATTENHOFFER[¶]

Abstract. We study a fundamental measure for wireless interference in the SINR model known as (weighted) inductive independence. 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 inductive independence, implying a number of algorithmic applications. An algorithm is provided that achieves capacity that is — due to existing lower bounds — asymptotically best possible using oblivious power assignments. Improved approximation algorithms are provided for a number of problems involving both oblivious power and arbitrary power control, including connectivity, secondary spectrum auctions, and dynamic packet scheduling. We also show that the price of oblivious power — the relative increase in capacity possible when using unconstrained power control — is only doubly logarithmic in the maximum link length.

Key words. approximation algorithms, wireless communication, SINR model, oblivious power, power control, wireless capacity, inductive independence

AMS subject classifications. 68W40, 68W20, 68W25, 68M10

1. Introduction. Power control is one of the most versatile tools to increase the capacity of a wireless network. Higher power increases the throughput of a transmission link, while causing more interference to other simultaneously transmitting links. Given this tension, intelligent power control is crucial in increasing the spatial reuse of the available bandwidth provided by the shared medium. Thus it is not surprising that most contemporary wireless protocols use some form of power control. It has also been shown theoretically that power control may improve the capacity of a wireless network in an exponential [26, 50] or even unbounded [15] way.

Unrestricted power control is, however, a double-edged sword. In order to achieve the theoretically best results, one must solve complex optimization problems, where transmission power of one node potentially depends on the transmission powers of all other nodes [42]. In real wireless networks, where communication demands change over time, this may not be an option. In practical protocols, the transmission power should preferably be independent of other concurrent transmissions, which leaves it to only depend on the distance between transmitter and receiver. This is known as *oblivious* power control.

Many questions immediately arise in the wake of the preceding observations: What is the price in performance for restricting power control to oblivious powers?

*The preliminary version of this paper appeared in the Proceedings of the 24th ACM-SIAM Symposium on Discrete Algorithms (SODA) 2013 [27].

[†]ICE-TCS, School of Computer Science, Reykjavik University, Reykjavik, Iceland. (mmh@ru.is). Supported by Icelandic Research Fund grant-of-excellence no. 120032011 and grant no. 152679-051.

[‡]Theory of Computation Group, MIT, USA. (holzer@mit.edu). At MIT the author is supported by the following grants: AFOSR Contract Number FA9550-13-1-0042, NSF Award 0939370-CCF, NSF Award CCF-1217506, NSF Award number CCF-AF-0937274. Part of the work was done while at ETH Zurich, Switzerland. Corresponding author.

[§]Google Inc., New York, USA. (ppmitra@gmail.com). Work done at Reykjavik University, Iceland.

[¶]Distributed Computing Group, ETH Zurich, Switzerland. (wattenhofer@ethz.ch).

Which of the infinitely many oblivious power schemes are good choices? Once an oblivious power scheme is chosen, what algorithmic results can be achieved?

In this work, we look at these questions in the context of the physical or SINR model of interference, a realistic model gaining increasing attention (see Section 1.2 for historical background and motivation and Section 2 for precise definitions). In this setting, our work completes a line of investigation by answering a number of these questions optimally.

The specific problem at the center of our work is *capacity* maximization: Given a set of transmission links (each a transmitter-receiver pair), find the largest subset of links that can successfully transmit simultaneously when assigned appropriate transmission powers.

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 (proportional to) $\ell^{\alpha/2}$ (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 with respect to arbitrary power control within a factor of $\mathcal{O}(\log(n) \cdot \log \log \Delta)$ [26] and $\mathcal{O}(\log(n) + \log \log \Delta)$ [29], 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 $\Omega(n)$ [15] applied only when Δ was doubly exponential. In terms of Δ , it was shown that one *must* pay an $\Omega(\log \log \Delta)$ factor [26]. The best upper bounds were, as mentioned, either dependent on the size of the input [26, 29] and as such unbounded (in relation to Δ), or exponentially worse ($\mathcal{O}(\log \Delta)$) [1, 21].

1.1. Our Contributions. In this paper, we study power assignments of the form $\ell^{p\alpha}$ for all fixed $0 < p < 1$ (setting $p = \frac{1}{2}$ results in mean power). Our first result is a simple algorithm using any oblivious power scheme of this form, whose performance matches the known $\Omega(\log \log \Delta)$ lower bound. For small to moderate values of Δ , e.g., when Δ is at most polynomial in n (which presumably includes most real-world settings), our bound is an exponential improvement over previous bounds, including the $\mathcal{O}(\log \Delta)$ -bound of [1] (see also [21]).

This result extends the “star status” from mean power to a large class of assignments. This class has been studied implicitly before in a range of work on “length-monotone, sub-linear” power assignments [28, 29, 37, 45], but its relation to arbitrary power was not well understood.

Our second main contribution is to improve a number of algorithmic results that use these power assignments. We shave a logarithmic (in n) factor off the approximation ratios of a variety of problems, including secondary spectrum auctions [37], wireless connectivity [30, 31, 50], and dynamic packet scheduling [3, 44]. Using the capacity relation between oblivious and arbitrary power (our first result), we strengthen the bounds for these problems 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, from which the algorithmic applications follow essentially immediately. These results are important in their own right, e.g., implying tight bounds on certain efficiently computable measures of interference.

To provide an intuitive understanding of our results in the next paragraph, we recall the graph theoretic notion of *inductive independence* [57].

DEFINITION 1.1 (Inductive independence). *A graph G is d -inductive independent if there is an ordering of the vertices v_1, v_2, \dots, v_n such that each v_i has at most d neighbors in any independent set $I \subseteq \{v_{i+1}, v_{i+2}, \dots, v_n\}$.*




FIG. 1. The graph on the left has inductive independence number 1 (i.e., is chordal), while the graph on the right has inductive independence number 2.

An example is provided in Figure 1.

The inductive independence property is found in many graph classes (e.g., intersection graphs of translates of a convex planar object are 3-inductive independent [57]), and it has powerful algorithmic implications [25, 37, 57]. For example, a simple d -approximation algorithm for the maximum independent set problem in such a graph is as follows: Process the vertices in the prescribed order, adding each vertex to the solution if it has no neighbors already in the solution. By the inductive independence property, the addition of a single vertex disqualifies at most d vertices of the optimal solution from being added in the future, which implies the claimed approximation factor.

In this paper, we deal with an interference measure that is a natural analog of inductive independence, applied to certain weighted graphs that model the SINR interference scenario. In this context, the links to be scheduled are represented by vertices of a graph. The weight of a directed edge is the relative interference that the source link causes on the destination link. The relevant ordering of the vertices is the ascending order of link length. The vertices of an “independent set” in this graph represent a set of links that cause limited interference to each other when transmitting simultaneously. These sets are called *feasible sets* of links, as all links in such a set can transmit successfully at the same time.

When one is interested in feasible sets and allowed to assign *arbitrary* (unrestricted) transmission powers to links, we show that the measure is bounded by $\mathcal{O}(\log \log \Delta)$ (Theorem 3.2), implying our first capacity result (and its applications). This result holds for links on the plane, and in a more general class of metrics that we define here. Technically, this is done by carefully extending the analysis of [26]. When feasibility is with respect to oblivious power from the abovementioned class, the measure is constant-bounded (Theorem 3.3), implying the second set of algorithmic results. This results hold for general metric spaces and all constants $\alpha > 0$.

Apart from the specific applications pinpointed here, we expect future algorithmic questions in the SINR model to directly benefit from these bounds.

1.2. Related Work. Gupta and Kumar [24] were among the first to provide 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 [46, 53], but such models do not capture the nature of wireless communication well, as demonstrated in [23, 48, 51]. In 2006, Moscibroda and Wattenhofer [50] combined the best of both worlds, studying algorithms for scheduling in arbitrary networks. Since then, the problems studied in this setting have reflected the diversity of the application areas underlying it – topology control [17, 41, 52], sensor networks [49], combined scheduling and routing

[9], ultra-wideband [38], and analog network coding [22].

In spite of this diversity, certain canonical problems have emerged, the study of which has resulted in improvements for other problems as well. The capacity problem is one such problem. After it was quickly shown to be NP-complete [21], a constant factor approximation algorithm for uniform power was achieved in [19, 18], and eventually extended to essentially all interesting oblivious power schemes in [29]. In [42, 43], a constant approximation to the capacity problem for arbitrary powers was obtained. The relation between capacity using oblivious power and capacity using arbitrary power was first studied in [26].

An alternative approach to capacity maximization using uniform power is to use regret minimization by distributed algorithms, first proposed in [1], with a constant factor approximation derived in [2], and later extended to handle jamming [12], and changing spectrum availability [11].

Linear power has turned out to be the easiest among fixed power assignments, being the only one with a constant factor approximation for scheduling with respect to the optimum schedule achievable with linear power [16, 56] and a constant-bounded interference measure [16]. Whereas there are instances for which linear and uniform power are arbitrarily bad in comparison with mean power [50], a maximum feasible subset under mean power is known to be always within a constant factor of subsets feasible under linear or uniform power [55]. Recently it was shown in [10] that algorithms for capacity-maximization in the SINR model can be transferred to a model that takes Rayleigh-fading into account, losing only an $\mathcal{O}(\log^* n)$ factor in the approximation ratio. Also, it was argued in [7] that capacity maximization algorithms like ours could be applied to arbitrarily complex environments, with the approximation factor reflecting an innate property of the signal reception matrix (as opposed to the interpolated path loss constant).

Technically, the idea of looking at the interaction between a feasible set and a link was studied before. The works of [26] and [45] are particularly relevant – the first in the context of oblivious/arbitrary-power comparison, and the second in the context of oblivious power. Our results improve the bounds in those papers to the best possible up to a small constant factor.

Since the initial publication of the current work in [27], there has been progress on the related scheduling problem with power control. In particular, a $\mathcal{O}(\log \log \Delta)$ -approximation was achieved in [34] using oblivious power (assuming non-weak links), matching our capacity result. Curiously, that approximation ratio holds only for a more limited set of power assignments (when $p > 2/\alpha$). Also, an improved approximation of $\mathcal{O}(\log^* \Delta)$ was obtained for scheduling with power control in [33], necessarily using non-oblivious power assignments, with evidence provided that this might be best possible. It was shown in [5] that, unlike we show here for capacity, inductive independence is not a strong enough bound to achieve better than logarithmic approximation of scheduling problems.

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 the interference measure and our two structural results. Sections 3.1 and 3.2 are devoted to proofs of the structural results. Section 4 illustrates a key application, with further applications given in Section 5.

2. Model and Definitions. Given is a set L of links, where each link v represents a unit-demand communication request from a transmitter s_v to a receiver r_v , both of which are points in a metric space. The distance between two points x and y is

denoted by $d(x, y)$. We write $d_{vw} = d(s_v, r_w)$ for short, and denote by $d_{vv} = d(s_v, r_v)$ the length of link v .

DEFINITION 2.1 (SINR-formula). *Let Φ_v denote the power assigned to link v , or in other words, s_v transmits with power Φ_v . In the physical model (or SINR model) of interference, a transmission on link v is successful if and only if*

$$(2.1) \quad \frac{\Phi_v/d_{vv}^\alpha}{\sum_{w \in S \setminus \{v\}} \Phi_w/d_{vw}^\alpha + N} \geq \beta ,$$

where N is a universal constant denoting the ambient noise, $\beta \geq 1$ 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 v . Let Φ_{max} denote maximum available transmission power and $\ell_{max} = (\Phi_{max}/\beta N)^{1/\alpha}$ denote the maximum distance for a transmission to be successful in the absence of interference.

We view α as a universal constant that can be hidden in big-oh notation, but not the parameter β .

We focus on oblivious power assignments \mathcal{P}_p , where the power $\Phi_v \sim d_{vv}^{p \cdot \alpha}$ assigned to the sender of link v only depends on the length d_{vv} of the link and the parameter p . More precisely, under \mathcal{P}_p , the power assigned to v is given by

$$\Phi_v = \Phi_{max} \cdot \left(\frac{d_{vv}}{\ell_{max}} \right)^{p\alpha} = \Phi_{max}^{1-p} (\beta N)^p \cdot d_{vv}^{p\alpha} .$$

This includes all the specific assignments of major interest: uniform (\mathcal{P}_0), mean ($\mathcal{P}_{1/2}$), and linear power (\mathcal{P}_1).

DEFINITION 2.2 (Feasibility). *We say that a subset $S \subseteq L$ of links is \mathcal{P} -feasible, if Condition (2.1) is satisfied for each link in S when using power assignment \mathcal{P} . We say that S is feasible if there exists a power assignment \mathcal{P} for which S is \mathcal{P} -feasible.*

Let **Capacity** denote the problem of finding a maximum cardinality feasible subset of links (that is, maximizing the capacity of the wireless channel used).

The notion of relative interference [39], which we refer to as *affectance* following [18, 45], is crucial to our arguments.

DEFINITION 2.3 (Affectance). *The affectance $a_w^{\mathcal{P}}(v)$ on link v caused by another link w , with a given power assignment \mathcal{P} , is the interference of w on v relative to the power received, or*

$$a_w^{\mathcal{P}}(v) := \min \left(1, c_v \frac{\Phi_w/d_{wv}^\alpha}{\Phi_v/d_{vv}^\alpha} \right) = \min \left(1, c_v \frac{\Phi_w}{\Phi_v} \cdot \left(\frac{d_{vv}}{d_{wv}} \right)^\alpha \right) ,$$

where $c_v := \beta/(1 - \beta N d_{vv}^\alpha / \Phi_v)$ depends only on properties of the link v and on universal parameters. Note that $c_v \geq \beta \geq 1$.

Let $a_v^{\mathcal{P}}(w)$ denote $a_w^{\mathcal{P}}(v)$. Conventionally, we define $a_v^{\mathcal{P}}(v) := 0$, since v does not interfere with itself. For sets S and T of links and a link v , let $a_v^{\mathcal{P}}(S) := \sum_{w \in S} a_v^{\mathcal{P}}(w)$, $a_S^{\mathcal{P}}(v) := \sum_{w \in S} a_w^{\mathcal{P}}(v)$, and $a_S^{\mathcal{P}}(T) := \sum_{w \in S} a_w^{\mathcal{P}}(T)$. Given a set of links L and a link $v \in L$, we often refer to $a_L^{\mathcal{P}}(v)$ as the sum of *in-affectance* on link v . Using this notation, Condition (2.1) can be rewritten as $a_S^{\mathcal{P}}(v) \leq 1$ (except for the near-trivial case of S containing only two links).

We introduce two more affectance notations.

DEFINITION 2.4 (Symmetric affectance). *We define*

$$b_v^{\mathcal{P}}(w) := b_w^{\mathcal{P}}(v) := a_v^{\mathcal{P}}(w) + a_w^{\mathcal{P}}(v)$$

to be the symmetric version of affectance.

Without loss of generality, assume that link-lengths form a *total* order \prec , where symmetry is broken arbitrarily; i.e., \prec is an arbitrary linear extension of the partial order given by length comparisons.

DEFINITION 2.5 (Length-ordered affectance). *Denote by $\hat{a}_v^{\mathcal{P}}(w)$ (and $\hat{b}_v^{\mathcal{P}}(w)$) the length-ordered version of affectance, defined as $a_v^{\mathcal{P}}(w)$ ($b_v^{\mathcal{P}}(w)$) if $d_{vv} \prec d_{ww}$ and 0 otherwise, respectively.*

These are extended in similar ways to affectances to and from sets as defined for $a_v^{\mathcal{P}}(w)$. As before, we write $b_v^{\mathcal{P}}(w) := b_v^{\mathcal{P}^p}(w)$, $\hat{b}_v^{\mathcal{P}}(w) := \hat{b}_v^{\mathcal{P}^p}(w)$ and $\hat{a}_v^{\mathcal{P}}(w) := \hat{a}_v^{\mathcal{P}^p}(w)$. These measures are essentially identical when taken over a whole set.

OBSERVATION 2.6. $a_S^{\mathcal{P}}(S) = \hat{b}_S^{\mathcal{P}}(S) = b_S^{\mathcal{P}}(S)/2$.

Proof.

$$b_S^{\mathcal{P}}(S) = \sum_{v,u \in S} b_v^{\mathcal{P}}(w) = \sum_{v,u \in S} (a_v^{\mathcal{P}}(w) + a_w^{\mathcal{P}}(v)) = 2a_S^{\mathcal{P}}(S) .$$

Also,

$$\begin{aligned} \hat{b}_S^{\mathcal{P}}(S) &= \sum_{v,w \in S} \hat{b}_v^{\mathcal{P}}(w) = \sum_{\substack{v,w \in S \\ v \prec w}} b_v^{\mathcal{P}}(w) = \sum_{\substack{v,w \in S \\ v \prec w}} (a_v^{\mathcal{P}}(w) + a_w^{\mathcal{P}}(v)) \\ &= \sum_{\substack{u,x \in S \\ u \prec x}} a_u^{\mathcal{P}}(x) + \sum_{\substack{u,x \in S \\ x \prec u}} a_u^{\mathcal{P}}(x) = \sum_{u,x \in S} a_u^{\mathcal{P}}(x) = a_S^{\mathcal{P}}(S) . \end{aligned}$$

□

Let $\Delta = \Delta(L)$ denote the ratio between the maximum and minimum length of a link in L .

DEFINITION 2.7 (Length-classes). *A set of links is a length-class if the lengths of the links within the set vary by a factor at most 2. We refer to links in the same length-class as nearly-equal length.*

Clearly, every link set L can be partitioned into $\lceil \log \Delta(L) \rceil$ length-classes.

Links that are very close to the longest possible need special treatment.

DEFINITION 2.8 ((Non)-weak links). *A link v is said to be weak if $c_v > 2\beta$ and non-weak if $c_v \leq 2\beta$. The latter is equivalent to the condition $\Phi_v \geq 2\beta N d_{vv}^\alpha$.*

Intuitively, this means that the link uses at least slightly more power than the absolute minimum needed to overcome ambient noise (the constant 2 can be replaced with any fixed constant larger than 1). Some of our theorems assume links to be non-weak, a reasonable and often-used assumption [1, 14, 20, 45].

To handle also weak links, we classify them into groups of roughly equal c_v -values, along similar lines as proposed in [6].

DEFINITION 2.9 (Tolerance-classes). *A set S of links is an (interference) tolerance-class if both the lengths and c_v -values of the links differ by a factor at most 2, i.e., if $\forall v, w \in S, c_v \leq 2c_w$ and $d_{vv} \leq 2d_{ww}$.*

Let $c_{max}(S) = \max_{v \in S} c_v$ ($c_{min}(S) = \min_{v \in S} c_v$) denote the largest (smallest) c_v value among the links in a set S of links, respectively.

DEFINITION 2.10 (Convergent metric). *A metric \mathcal{M} is a pair (V, d) where V is a set and d is a distance function satisfying the triangle inequality. A subset $X \subseteq V$ is a c -packing if $d(x, x') \geq c$, for every $x, x' \in X, x \neq x'$. In the context of a given pathloss parameter α , a metric is convergent if there is a constant $c_{\mathcal{M}}$ such that for any $c > 0$, any c -packing X , and any $v \in V$, it holds that $\sum_{x \in X} d(x, v)^{-\alpha} \leq c_{\mathcal{M}}/c^\alpha$.*

This definition captures the property that is essential for many arguments about wireless algorithms that the interference from a set of properly spaced links that tile the plane converges to a constant, assuming $\alpha > 2$. Convergent metrics generalize all previous definitions of similar type, including fading metrics [26] and bounded-growth metrics [13, 40].

DEFINITION 2.11 (Independence). *Links v and w are q -independent if they satisfy $d_{vw} \cdot d_{wv} \geq q^2 \cdot d_{vw} \cdot d_{vw}$. A set of mutually q -independent links is said to be q -independent.*

An example of q -independence is given in Figure 2.




FIG. 2. Nodes are located on the grid at unit-distances. Links 1 and 2 are 1.92-independent. Set $\{1, 2, 3\}$ is 1.43-independent.

Independence is a pairwise property, and thus weaker than feasibility. A feasible set is necessarily $\beta^{1/\alpha}$ -independent [26], but there is no good relationship in the opposite direction.

In this paper we provide an independence-strengthening result with better trade-offs than the so-called “signal-strengthening” of [35]. This lemma allows us to forget about the specific value of β that the feasible set satisfies and consider some stronger threshold, e.g., $\beta' = 2$, for convenience. Recall $c_{\min}(S) = \min_{v \in S} c_v$.

LEMMA 2.12. *For any given $q > 1$, a feasible set S of links can be partitioned into $\lceil 2q^\alpha / c_{\min}(S) \rceil$ sets, each q -independent.*

Proof. Let $r := c_{\min}(S)/q^\alpha$ and $z := \lceil 2/r \rceil = \lceil 2q^\alpha / c_{\min}(S) \rceil$. Let \mathcal{P} be a power assignment that makes S feasible. We form a graph G on the set S of links, where links v and w are adjacent if and only if $b_v^{\mathcal{P}}(w) > r$.

We first show that G is z -colorable. Suppose otherwise and let $R \subseteq S$ be a minimally $z + 1$ -chromatic subgraph. Since R is \mathcal{P} -feasible, $a_R^{\mathcal{P}}(v) \leq 1$, for each v in R . Thus,

$$b_R^{\mathcal{P}}(R) = 2a_R^{\mathcal{P}}(R) = 2 \sum_{v \in R} a_R^{\mathcal{P}}(v) \leq 2|R|.$$

Then, there is a link u_1 in R with $b_{u_1}^{\mathcal{P}}(R) \leq 2$. Form a z -coloring of $R \setminus \{u_1\}$, which exists by the minimality of R . There are strictly fewer than $z = \lceil 2/r \rceil$ links w in R for which $b_{u_1}^{\mathcal{P}}(w) > r$. Thus, there is a color class that contains no neighbor of u_1 , and assigning u_1 to that class yields a z -coloring of R , which is a contradiction.

Consider now a color class (i.e., a graph-theoretic independent set) $S_c \subseteq S, c \in$

$\{1, \dots, z\}$. It holds by the definition of G that, for any pair v, w of links in S_c ,

$$(2.2) \quad a_w^{\mathcal{P}}(v) \cdot a_v^{\mathcal{P}}(w) \leq b_w^{\mathcal{P}}(v) \cdot b_v^{\mathcal{P}}(w) \leq r \cdot r = \frac{c_{\min}(S)^2}{q^{2\alpha}} \leq \frac{c_v c_w}{q^{2\alpha}}.$$

Since w and v belong to the same feasible set, it holds by the definition of affectance that

$$(2.3) \quad a_v^{\mathcal{P}}(w) \cdot a_w^{\mathcal{P}}(v) = c_w \frac{\Phi_v/d_{vw}^\alpha}{\Phi_w/d_{ww}^\alpha} \cdot c_v \frac{\Phi_w/d_{ww}^\alpha}{\Phi_v/d_{vv}^\alpha} = c_w c_v \left(\frac{d_{ww} d_{vv}}{d_{vw} d_{vw}} \right)^\alpha.$$

Combining the Bounds (2.2) and (2.3) yields that v and w are q -independent. Since this holds for arbitrary links $v \neq w \in S_c$ and arbitrary $c \in \{1, \dots, z\}$, it follows that the coloring of G yields the desired partition of L into q -independent sets. \square

3. Structural Properties. We start by defining the interference measure at the center of this work. As mentioned in the introduction, this definition is a fractional analogue of the inductive independence number of a graph.

DEFINITION 3.1 (Inductive independence). *Let L be a set of links, \mathcal{P}, \mathcal{Q} be two power assignments of L , and $F_{\mathcal{Q}}(L)$ be the collection of \mathcal{Q} -feasible subsets of L . Then,*

$$I_{\mathcal{Q}}^{\mathcal{P}}(L) := \max_{S \in F_{\mathcal{Q}}(L)} \max_{v \in L} \hat{b}_v^{\mathcal{P}}(S).$$

When either \mathcal{P} or \mathcal{Q} denotes \mathcal{P}_p , we replace it by p , e.g., $I_p^p(L) := I_{\mathcal{P}_p}^{\mathcal{P}_p}(L)$.

In our setting, the weighted graph is formed on the links, that is, L is the set of nodes in the graph. The weight of the (undirected) edge between links u and v is $b_u^{\mathcal{P}}(v) = b_v^{\mathcal{P}}(u)$. The ordering is the ascending order of length of the corresponding links. Then, $I_{\mathcal{Q}}^{\mathcal{P}}(L)$ is an upper bound on how much weight/interference (when using power \mathcal{P}) a link can have into a \mathcal{Q} -feasible set containing longer links, just as the inductive independence number in graphs is an upper bound on how many edges a node can have to an independent set consisting of higher-ranked nodes.

When used with different power assignments, $I_{\mathcal{Q}}^{\mathcal{P}}(L)$ gives us a handle on comparing the utility of these power assignments. We primarily use it in the setting where $\mathcal{P} = \mathcal{P}_p$, for some $p \in (0, 1]$, and \mathcal{Q} is (an) optimal arbitrary power assignment (that maximizes Capacity of L), allowing us to relate oblivious power to arbitrary power.

In this section, we give two structural results that characterize the utility of oblivious power assignments. Both of these are best possible and answer important 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. It improves upon the $\mathcal{O}(\log(n) + \log \log \Delta)$ bound stated implicitly in [29] and extends it to a range of power assignments.

THEOREM 3.2. *For every set L of non-weak links in a convergent metric, $0 < p < 1$, and power assignment \mathcal{Q} , it holds that $I_{\mathcal{Q}}^p(L) = \mathcal{O}(\log \log \Delta(L))$.*

The second is a constant upper bound on the function when \mathcal{P} and \mathcal{Q} are the same \mathcal{P}_p assignment, for some $p \in (0, 1]$. This improves upon the $\mathcal{O}(\log n)$ bound in [45].

THEOREM 3.3. *Every set L of links is $\mathcal{O}(1)$ -inductive independent under \mathcal{P}_p , i.e., $I_p^p(L) = \mathcal{O}(1)$, where $0 < p \leq 1$.*

Each theorem is treated in a separate subsection.

3.1. Inductive Independence of Oblivious Powers with Respect to the Optimum Assignment (Proof of Theorem 3.2). We want to bound the symmetric affectance of a link v with respect to a set S of longer links. We do so by partitioning S into three sets and bounding the contributions of each set separately.

For the set S_1 of “long” links with “large” affectance (with respect to v), we extend a result of [26]. That argument (Lemma 3.4) is based on showing that the lengths of the links in S_1 must grow doubly exponentially, and thus there can be at most $\mathcal{O}(\log \log \Delta)$ links in the set.

For the set S_2 of “long” links with “small” affectance with respect to v , we break S_2 into tolerance-classes, show that the affectance of v towards each tolerance-class is small, i.e., $\mathcal{O}(1/\log \Delta)$, and since the number of tolerance-classes is at most $\log \Delta$ (when the links are non-weak), the total affectance is $\mathcal{O}(1)$.

Finally, the set S_3 of “short” links has only $\mathcal{O}(\log \log \Delta)$ tolerance-classes, so it suffices to show that the affectance to each of them is constant. The same Lemma 3.9 is used for the tolerance-class arguments of both S_2 and S_3 .

We first present key lemmata (Lemmata 3.4, 3.7, and 3.9) to bound affectances of a link to and from a set of links. The first lemma handles the set of long links with relatively high affectance. This lemma is key to the $\mathcal{O}(\log \log \Delta)$ bound and essentially shows that if the links in a set are mutually distant, yet all “close” to a given link, then their lengths must grow doubly exponentially. It originates in [26] (Lemma 4.4) and is generalized here to any \mathcal{P}_p power assignment (with $0 < p < 1$) and to the property of independence (which is weaker than affectance).

We introduce a parameter \hat{p} defined as $\hat{p} := \frac{1}{\min(1-p, p)}$ for the rest of this subsection.

LEMMA 3.4. *Let p be a constant, $0 < p < 1$, τ be a parameter, $\tau \geq 1$, and $\Lambda := (2(2\tau)^{1/\alpha})^{\hat{p}}$. Let v be a link and let Q be a $3\beta^{1/\alpha}$ -independent set of non-weak links in an arbitrary metric space, where each link $w \in Q$ satisfies $\max(a_w^p(w), a_w^p(v)) \geq 1/\tau$ and $\Lambda \cdot d_{vv} \leq d_{ww}$. Then, $|Q| = \mathcal{O}(\log \log \Delta(Q))$.*

Proof. Let $\gamma := \beta^{1/\alpha}$. We partition Q into two sets:

1. The set $Q_1 = \{w \in Q : a_w^p(v) \geq 1/\tau\}$ of links that affect v by at least $1/\tau$, and
2. The set $Q_2 = \{w \in Q : a_w^p(w) \geq 1/\tau\}$ that are similarly affected by v .

We prove the statement for each type separately.

Step 1: Consider a pair w, w' of links in Q_1 that each affect v by at least $1/\tau$ under \mathcal{P}_p , and suppose without loss of generality that $d_{w'w'} \leq d_{ww}$. The assumption that $a_w^p(v) \geq 1/\tau$ is equivalent to the relationship $c_v(d_{ww}^p d_{vv}^{1-p})^\alpha \geq d_{ww}^\alpha / \tau$, which implies that

$$(3.1) \quad d_{ww} \leq d_{ww}^p d_{vv}^{1-p} (c_v \tau)^{1/\alpha} .$$

Similarly,

$$(3.2) \quad d_{w'v} \leq d_{w'w'}^p d_{vv}^{1-p} (c_v \tau)^{1/\alpha} \leq d_{ww}^p d_{vv}^{1-p} (c_v \tau)^{1/\alpha} ,$$

additionally using that w' is at most as long as w . Recall that $c_v \leq 2\beta = 2\gamma^\alpha$, since the links are non-weak. By Bounds (3.1) and (3.2) and the definition of Λ ,

$$(3.3) \quad d_{ww} + d_{w'v} \leq 2d_{ww}^p d_{vv}^{1-p} \gamma (2\tau)^{1/\alpha} = d_{ww}^p d_{vv}^{1-p} \gamma \Lambda^{1/\hat{p}} .$$

By the triangle inequality and Bound 3.3,

$$(3.4) \quad d_{w'w} \leq d(s_{w'}, r_v) + d(r_v, s_w) + d(s_w, r_w) \leq d_{w'w'}^p d_{vv}^{1-p} \gamma \Lambda^{1/\hat{p}} + d_{ww}$$

Let the links in Q_1 be ordered $1, 2, \dots, |Q_1|$ in order of non-decreasing length. Observe that by assumption, $d_{vv}\Lambda \leq d_{11} \leq d_{ww}$. Continuing from Bound (3.4), applying the definition of \hat{p} ,

$$(3.5) \quad d_{w'w} \leq d_{ww}^p (d_{vv}\Lambda)^{1-p} \gamma + d_{ww} \leq 2\gamma d_{ww} .$$

We can derive in the same way that

$$(3.6) \quad d_{ww'} \leq d_{wv} + d_{w'v} + d_{w'w'} \leq d_{ww}^p (d_{vv}\Lambda)^{1-p} \gamma + d_{w'w'} \leq d_{ww}^p d_{11}^{1-p} \gamma + d_{w'w'} .$$

Using the definition of independence, on one hand, and Bounds (3.5) and (3.6) on $d_{w'w}$ and $d_{ww'}$, on the other hand, we get that

$$(3.7) \quad 9\gamma^2 d_{ww} d_{w'w'} \leq d_{w'w} \cdot d_{ww'} \leq 2\gamma d_{ww} (d_{w'w'} + \gamma d_{ww}^p d_{11}^{1-p}) ,$$

Canceling a $\gamma^2 d_{ww}$ -factor and simplifying (using that $\gamma \geq 1$), we get that

$$7d_{w'w'} \leq 2d_{ww}^p d_{11}^{1-p} .$$

Rearranging,

$$(3.8) \quad \left(\frac{d_{ww}}{d_{11}} \right)^p \geq \frac{3d_{w'w'}}{d_{11}} .$$

Define $\lambda_i = d_{ii}/d_{11}$ as the ratio of the length between link i and the shortest link 1 in Q_1 . Applying Bound (3.8) with $w' = i$ and $w = i + 1$, we get that, for $i = 1, 2, \dots, |Q_1| - 1$,

$$\lambda_{i+1}^p \geq 3\lambda_i .$$

Then, $\lambda_2 \geq 3^{1/p}$ and by induction $\lambda_t \geq 3^{(1/p)^{t-1}}$. Note that

$$\Delta(Q_1) = d_{|Q_1||Q_1|}/d_{11} = \lambda_{|Q_1|} \geq 3^{(1/p)^{|Q_1|-1}} ,$$

so $|Q_1| - 1 \leq \log_{1/p} \log_3 \Delta(Q_1)$, and the claim follows.

Step 2: The other case of links w with $a_v^p(w) \geq 1/\tau$ is symmetric, with the roles of p and $1 - p$ switched, leading to a bound of $1 + \log_{1/(1-p)} \log_3 \Delta(Q_2)$. \square

Lemma 3.4 bounds the number of longer links that affect (or are affected by) a given link by a significant amount, or at least $1/\tau$. For affectances below that threshold, we bound their contributions for each tolerance-class separately.

We need the following geometric argument to convert statements involving the link v into statements about links within the length-class S . In particular, we bound the affectance involving v by affectances involving its ‘‘guard’’ u , the link in S closest to v . To this end, we first lower bound the distance between the receiver of v to any sender in S by the distance between that sender and the receiver of u .

PROPOSITION 3.5. *Let v be a link. Let S be a 2-independent length-class and let u be the link in S with d_{uv} minimum. Then, $d_{wv} \geq d_{wu}/6$, for any link w in S .*

Proof. The reader may find Figure 3 helpful when reading this proof. Consider a link w in S . Let $D := d_{wv}$ and note that by choice of u , $d_{uv} \leq D$. By the triangle inequality and the choice of u ,

$$(3.9) \quad d_{wu} \leq d_{wv} + d_{uv} + d_{uu} \leq 2D + d_{uu} .$$




FIG. 3. Links u, v and w as used in the proof of Proposition 3.5. The distances d_{wv} and d_{wu} that are related to each other as stated in the proposition’s statement are represented by red dotted lines. The gray dashed lines mark the distances d_{uw} and d_{uv} that are used in the proofs as well.

Similarly,

$$(3.10) \quad d_{uw} \leq d_{uv} + d_{wv} + d_{wv} \leq 2D + d_{wv} .$$

Now we recall the definition of 2-independence, apply it to u and w and bound d_{wu} and d_{wv} by Bounds (3.9) and (3.10) to obtain

$$(3.11) \quad 4d_{uu}d_{wv} \leq d_{wu} \cdot d_{wv} \leq (2D + d_{uu}) \cdot (2D + d_{wv}) .$$

This implies that D must be at least $\min(d_{uu}, d_{wv})/2$, as otherwise Bound (3.11) leads to a contradiction. Thus, since the links are nearly-equilength, $D \geq \max(d_{uu}, d_{wv})/4$. Combined with Bound (3.9), we obtain that $d_{wu} \leq 6D$. \square

The following Proposition serves a similar purpose as Proposition 3.5, treating senders instead of receivers and vice versa.

PROPOSITION 3.6. *Let S be a 2-independent length-class and v be a link not necessarily in S . Let u be the link in S with d_{vu} minimum. Then, $d_{vw} \geq d_{wu}/6$, for any link w in S .*

Proof. The proof is nearly the same as for Proposition 3.5 and is given for completeness in Appendix A. \square

This leads to the second key lemma of this section. It shows that the symmetric affectance of a link relative to a length-class is small when the link is much shorter than the links in the class. We use this lemma for the theorem in the following subsection, but use it also to derive a sibling lemma that is used later in this subsection.

LEMMA 3.7. *Let $q \geq 1$ be a real value and let v be a link. Let S be a 2-independent and \mathcal{P}_p -feasible tolerance-class with links of minimum length at least $q^{\hat{p}/\alpha} \cdot d_{vv}$. Then,*

$$b_v^p(S) = \mathcal{O}(1/q) + \max_{w, w' \in S} b_v^p(\{w, w'\}) .$$

Proof. We prove this in three steps. We show that for a set S' of all but two links in S it holds that (i) $a_{S'}^p(v) = \mathcal{O}(1/q)$ and (ii) $a_v^p(S') = \mathcal{O}(1/q)$. From this we derive the statement of the Lemma in Step (iii).

The intuitive idea is to identify a “guard” for the link v , which is the link in S “closest” to v . We then bound the affectance involving v in terms of that of the guard. By using the triangle inequality (in the form of Propositions 3.5 and 3.6), we argue that the interference on v is not much more than that experienced by the

guard, which by definition is small, since S is feasible. Also, since v is much shorter than the guard – by a factor of at least $q^{\hat{p}/\alpha}$ – the affectance involving v is at least q times smaller than that of the guard. When considering affectance *to* v , the guard is chosen as the link whose sender is closest to v 's receiver, while for affectance *from* v , it is the one with receiver closest to v 's sender.

Step (i): Consider the link u in S with d_{uv} minimum. Since $d_{vv} \leq d_{uu}$, it holds that $c_v \leq c_u$. We bound $a_w^p(v)$, for each $w \in S$, as follows, where the numbered transformations/bounds are explained by:

1. Def. 2.3 of affectance.
2. Choice of u and Proposition 3.5.
3. Def. 2.3 of affectance, and that since S is \mathcal{P}_p -feasible, thresholding does not take place.
4. Choice of $\hat{p} = \frac{1}{\min(1-p, p)}$ and $q \geq 1$.

$$\begin{aligned} a_w^p(v) &\stackrel{1.}{\leq} c_v \left(\frac{d_{vv}^{1-p} d_{vw}^p}{d_{vw}} \right)^\alpha \stackrel{2.}{\leq} c_u \left(\frac{(d_{uu}/q^{\hat{p}/\alpha})^{1-p} d_{uw}^p}{d_{uu}/6} \right)^\alpha = \frac{6^\alpha}{q^{\hat{p} \cdot (1-p)}} \cdot c_u \left(\frac{d_{uu}^{1-p} d_{uw}^p}{d_{uu}} \right)^\alpha \\ &\stackrel{3.}{=} \frac{6^\alpha}{q^{\hat{p} \cdot (1-p)}} a_w^p(u) \stackrel{4.}{\leq} \frac{6^\alpha}{q} a_w^p(u) , \end{aligned}$$

For any subset $S' \subseteq S \setminus \{u\}$, this extends to

$$(3.12) \quad a_{S'}^p(v) \leq \frac{6^\alpha}{q} a_{S'}^p(u) = \mathcal{O}(1/q) ,$$

using that S' is feasible such that $a_{S'}^p(u) \leq 1$.

Step (ii): Consider the link u' in S with $d_{vu'}$ minimum. Then, for each $w \in S$,

$$a_w^p(w) \leq c_w \left(\frac{d_{vv}^p d_{ww}^{1-p}}{d_{vw}} \right)^\alpha \leq c_w \left(\frac{(d_{ww}/q^{\hat{p}/\alpha})^p d_{ww}^{1-p}}{d_{wu'}/6} \right)^\alpha ,$$

using Proposition 3.6 and the bound assumed on the lengths of links in S (and thus of w) relative to v . Since u' and w are in the same tolerance-class, we bound this by

$$a_w^p(w) \leq 2c_{u'} \left(\frac{(d_{ww}/q^{\hat{p}/\alpha})^p (2d_{u'u'})^{1-p}}{d_{wu'}/6} \right)^\alpha = 2^{1+\alpha \cdot (1-p)} \cdot \frac{6^\alpha}{q^{\hat{p} \cdot p}} \cdot a_w^p(u') ,$$

using that $a_w^p(u') < 1$. Using that $q \geq 1$ and that $\hat{p} \cdot p \geq 1$, we get that $a_w^p(w) \leq \frac{1}{2} \cdot 12^\alpha \cdot a_w^p(u')$. Thus, for any subset $S' \subseteq S \setminus \{u'\}$,

$$(3.13) \quad a_v^p(S') \leq \frac{2 \cdot 12^\alpha}{q} a_{S'}^p(u') = \mathcal{O}(1/q) ,$$

since S' is feasible such that $a_{S'}^p(u') \leq 1$.

Step (iii): Combining Bounds (3.12) and (3.13) yields

$$b_v^p(S) - b_v^p(\{u, u'\}) = a_{S \setminus \{u, u'\}}^p(v) + a_v^p(S \setminus \{u, u'\}) = \mathcal{O}(1/q) .$$

Note that we do not require that $u \neq u'$. The theorem follows. \square

For links in a convergent metric, we can replace the assumption of \mathcal{P}_p -feasibility of the preceding lemma with a strengthened independence condition. This is the only place where we use the assumption of a convergent metric.

PROPOSITION 3.8. *Let v be a link and S be a tolerance-class of links in a convergent metric \mathcal{M} . Let $p > 0$, $q \geq 1$ be a real value, and $k \geq (c_{\mathcal{M}} \cdot c_{\max}(S))^{1/\alpha} 2^{p+1} + 4$. If S is k -independent, then it is \mathcal{P}_p -feasible.*

Proof. Let u and w be links in S and let $d(u, w)$ denote the distance between the nearest nodes on the two links: $d(u, w) = \min(d_{uw}, d_{wu}, d(s_u, s_w), d(r_u, r_w))$. Let x be the link in S of largest length, and note that since S is a tolerance-class, $d_{xx} \leq 2 \max(d_{uu}, d_{ww})$. By triangle inequality, $d(u, w) \geq d_{uw} - d_{uu} - d_{ww} \geq d_{uw} - 2d_{xx}$, and similarly $d(u, w) \geq d_{wu} - 2d_{xx}$. Thus by k -independence,

$$(d(u, w) + 2d_{xx})^2 \geq d_{uw}d_{wu} \geq k^2 d_{uu}d_{ww} \geq \left(\frac{k}{2}d_{xx}\right)^2 .$$

Thus,

$$(3.14) \quad d(u, w) \geq (k/2 - 2)d_{xx} .$$

Consider a link a in S . Let X be the set of nodes consisting of the receiver of a and the senders of the other links in S . By the k -independence of S and Bound (3.14), we have that X is a $(k/2 - 2)d_{xx}$ -packing (see Def. 2.10). Thus, $\sum_{u \in S \setminus \{a\}} 1/d_{ua}^\alpha \leq c_{\mathcal{M}} \cdot ((k/2 - 2)d_{xx})^{-\alpha}$, and hence,

$$\sum_{u \in S \setminus \{a\}} \frac{d_{xx}^\alpha}{d_{ua}^\alpha} \leq \frac{c_{\mathcal{M}}}{(k/2 - 2)^\alpha} \leq \frac{1}{c_{\max}(S) \cdot 2^{p\alpha}} .$$

Then, using that $d_{aa} \leq d_{xx}$ and that $\Phi_u \leq 2^{p\alpha}\Phi_a$ under \mathcal{P}_p (since S is a length-class),

$$a_S^p(a) \leq c_a \sum_{u \in S \setminus \{a\}} \frac{\Phi_u d_{aa}^\alpha}{\Phi_a d_{ua}^\alpha} \leq 2^{p\alpha} c_a \sum_{u \in S \setminus \{a\}} \frac{d_{xx}^\alpha}{d_{ua}^\alpha} \leq \frac{c_a}{c_{\max}(S)} \leq 1 .$$

Hence, S is \mathcal{P}_p -feasible. \square

The next main lemma follows immediately from Lemma 3.7 and Proposition 3.8.

LEMMA 3.9. *Let $q \geq 1$ be a real value, v be a link, and S be a tolerance-class with links of length at least $q^{\hat{p}/\alpha} \cdot d_{vv}$ in a convergent metric \mathcal{M} . Let $k \geq (c_{\mathcal{M}} \cdot c_{\max}(S))^{1/\alpha} 2^{p+1} + 4$. If S is k -independent, then,*

$$b_v^p(S) = \mathcal{O}(1/q) + \max_{w, w' \in S} b_v^p(\{w, w'\}) .$$

We are now ready to prove the first core result, Theorem 3.2.

THEOREM 3.2: *For every set L of non-weak links in a convergent metric, $0 < p < 1$, and power assignment \mathcal{Q} , it holds that $I_{\mathcal{Q}}^p(L) = \mathcal{O}(\log \log \Delta(L))$.*

Proof. Consider a link $v \in L$ and a \mathcal{Q} -feasible subset $S \subseteq L \setminus \{v\}$. We first argue that $\hat{b}_v^p(S) = b_v^p(S)$ and then show that $b_v^p(S) = \mathcal{O}(\log \log \Delta(L))$, from which the theorem follows.

By the definition of \hat{b} , we can assume, without loss of generality, that all links in S are of length at least d_{vv} , since \hat{b} is defined in such a way that all shorter links do not contribute to its value. With this assumption, $\hat{b}_v^p(S) = b_v^p(S)$.

We use the independence-strengthening Lemma 2.12 with $q = \max(3\beta^{1/\alpha}, (c_{\mathcal{M}} \cdot c_{\max}(S))^{1/\alpha} 2^{p+1} + 4)$ to partition S into q -independent sets. The number of sets in the partition is $t = \mathcal{O}(\max(\beta/c_{\min}(S), c_{\max}(S)/c_{\min}(S))) = \mathcal{O}(1)$, since S contains only non-weak links. Let S' be one of these sets, and note that it satisfies the independence conditions of both Lemma 3.4 and 3.9.

Let $D := \tau := \log \Delta(L)$ be the number of tolerance-classes of L (recalling that the links are non-weak) and $\Lambda := (2(2\tau)^{1/\alpha})^{\hat{p}}$. We say that a link w in S' is *short* if $d_{vv} \leq d_{ww} < \Lambda \cdot d_{vv}$ and *long* if $d_{ww} \geq \Lambda \cdot d_{vv}$. We partition S' into three sets:

- S_1 : Long links w with $b_v^p(w) \geq 1/\tau$,
- S_2 : Long links w with $b_v^p(w) < 1/\tau$, and
- S_3 : Short links.

We bound the symmetric affectance $b_v^p(S_i)$ of each set S_i separately.

The set S_1 satisfies the hypothesis of Lemma 3.4 (as the set Q) when using the link v and threshold τ . This implies that $|S_1| = \mathcal{O}(\log \log \Delta(S_1))$. Thus, by the Def. 2.4 of symmetric affectance,

$$b_v^p(S_1) \leq 2|S_1| = \mathcal{O}(\log \log \Delta(S_1)) = \mathcal{O}(\log \log \Delta(L)) .$$

Next, consider S_2 and partition it into tolerance-classes X_1, X_2, \dots, X_D . Each such class X_i satisfies the hypothesis of Lemma 3.9 with the choice of $q = 2\tau \geq 1$. Since $b_v^p(w) < 1/\tau$, for each $w \in X_i$, by assumption, Lemma 3.9 yields $b_v^p(X_i) = \mathcal{O}(1/q) + \max_{w, w' \in S_2} b_v^p(\{w, w'\}) = \mathcal{O}(1/\tau)$ for any X_i . Then,

$$b_v^p(S_2) = \sum_{i=1}^D b_v^p(X_i) \leq D \cdot \mathcal{O}(1/\tau) = \mathcal{O}(1) .$$

The set S_3 can be partitioned into $\log \Lambda$ length classes $Y_1, \dots, Y_{\log \Lambda}$. For each such length-class Y_i , we apply Lemma 3.9 with $q = 1$, which yields that $b_v^p(Y_i) = \mathcal{O}(1/q) + \max_{w, w' \in S_2} b_v^p(\{w, w'\}) = \mathcal{O}(1)$. In total, we obtain that $b_v^p(S_3) = \log D \cdot \mathcal{O}(1) = \mathcal{O}(\log \log \Delta(L))$. Thus,

$$b_v^p(S') = b_v^p(S_1) + b_v^p(S_2) + b_v^p(S_3) = \mathcal{O}(\log \log \Delta(L)) ,$$

and, as we can do this for each of the t different q -independent sets S' that make up S , we obtain

$$b_v^p(S) \leq t \cdot b_v^p(S') = \mathcal{O}(\log \log \Delta(L)) .$$

□

Finally, we give trade-offs in terms of the weakness of the links.

COROLLARY 3.10. *For every set L of links in a convergent metric, $0 < p < 1$, and power assignment \mathcal{Q} , $I_{\mathcal{Q}}^p(L) = \mathcal{O}(\log \log \Delta(L) + \log(c_{\max}(L)/\beta))$.*

Proof. Let v be a link and S be a feasible subset of L in a convergent metric \mathcal{M} . Partition S into the non-weak links S_{nw} and the weak links S_w , and further partition S_w into $\lceil \log(c_{\max}/\beta) \rceil$ tolerance classes of weak links. We show below that for each such tolerance class X , $b_v^p(X) = \mathcal{O}(1)$. Hence, $b_v^p(S_w) = \mathcal{O}(1) \cdot \lceil \log(c_{\max}/\beta) \rceil$, and $b_v^p(S) = \mathcal{O}(\log \log \Delta + \log(c_{\max}/\beta))$, by Thm. 3.2. Since this holds for any link and any feasible sets of L , the corollary follows.

It remains to show that $b_v^p(X) = \mathcal{O}(1)$, for a feasible tolerance-class X of weak links. Since it is a tolerance class, $c_{\max}(X) \leq 2c_{\min}(X)$. Let $k \geq (c_{\mathcal{M}} \cdot c_{\max}(X))^{1/\alpha} 2^{p+1} + 4$, and note that $k^\alpha = \mathcal{O}(c_{\max}(X)) = \mathcal{O}(c_{\min}(X))$. By Lemma 2.12, X can be partitioned into $\lceil 2(k/c_{\min}(X))^\alpha \rceil = \mathcal{O}(1)$ sets R_j , $j = 1, 2, \dots$, each k -independent. We apply Lemma 3.9 on each set R_j , obtaining that $b_v^p(R_j) = \mathcal{O}(1)$, and thus $b_v^p(X) = \mathcal{O}(1)$. □

Remark: Metrics The assumption of a convergent metric is necessary. Namely, a set L of equi-length links in a tree metric was constructed in [29] (slightly simplified in [34]) for which $I_{\mathcal{Q}}^p(L) = \Omega(\log n)$ (for some power assignment \mathcal{Q} and any $p \in (0, 1)$). It follows therefore more generally that no bounds in terms of Δ alone can hold in general metrics.

Remark: Power assignments The fact that all \mathcal{P}_p power assignments, $p \in (0, 1)$, seem to result in equivalent approximation factors begs the question whether there is any advantage of using one over another. One characteristic is that the value of \hat{p} grows as p gets closer and closer to either 0 or 1, and the resulting performance guarantees grow about linearly in \hat{p} . On the other hand, both ends of the spectrum have their advantages. Being close to uniform power can be highly useful when the range of power control is limited. Linear power, on the other hand, has the advantage of being energy efficient, requiring only about as much energy as needed to transmit; thus, being close to linear power transfers some of those benefits.

The theorem also shows that capacity under these different power assignments is comparable within doubly logarithmic factors. This may very well be best possible, but that is not known. What is known is that for each oblivious power assignment Φ , there is another oblivious power that allows for $\Omega(\log \log \Delta)$ larger capacity [26].

3.2. Constant-Inductive Independence under Fixed Oblivious Power Assignments (Proof of Theorem 3.3). The following crucial lemma shows that the total affectance of a link v on the longer links in L is constant. Together with a previous result on the affectance *from* the longer links on v , this implies constant-inductive independence.

The proof treats a worst-case instance, i.e., a set S and link v for which the inductive independence $b_v^p(S)$ is maximized (for a given n). We split S into two sets S_1 and S_2 , where S_1 contains “shorter” links, those that are at most a certain factor U longer than v , while S_2 contains the remaining longer links. The set S_1 of shorter links contains few length classes, so the affectance to them can easily be bounded using Lemma 3.7. To handle S_2 , we scale v to a longer link u that is still shorter than all the links in S_2 , and bound u ’s affectance on S_2 by v ’s affectance on S (by the worst-case assumption). The latter (v ’s affectance on S_2) must, however, be considerably less than the former (u ’s affectance on S_2), since shorter links use less power ($p > 0$) and cause correspondingly less affectance.

LEMMA 3.11. *Let $0 < p \leq 1$, S be a \mathcal{P}_p -feasible set of non-weak links, and v be a link (not necessarily in S). Then, $\hat{a}_v^p(S) = \mathcal{O}(1)$.*

Proof. Let \mathbf{L} be the set of all possible non-weak links and let \mathcal{L}_n be the family of all \mathcal{P}_p -feasible subsets of \mathbf{L} with at most n links. Define the function $g : \mathbb{N} \rightarrow \mathbb{R}^+$ to be the “optimum upper bound” on \hat{a}^p , that is, $g(n) := \sup_{S \in \mathcal{L}_n} \sup_{v \in \mathbf{L}} \hat{a}_v^p(S)$. Such a function exists, as it is trivially upper bounded by n .

Let c_1 be the implicit constant in Lemma 3.7 (with $q = 1$) such that for any 2-independent \mathcal{P}_p -feasible tolerance-class X , it holds that $b_v^p(X) \leq c_1$. Define $c_2 = 4 \max(c_1 \lceil \frac{1}{p\alpha} \rceil^{\frac{2\alpha+1}{\beta}}, 3^\alpha)$. We shall show that $g(n) \leq c_2$, for all n , which implies the lemma.

Let n be a number. Let $v \in \mathbf{L}$ and $S \in \mathcal{L}_n$ be a link and a \mathcal{P}_p -feasible set of at most n non-weak links, respectively, for which $\hat{a}_v^p(S) = g(n)$. By Def. 2.5 of \hat{a}^p , we may assume without loss of generality that each link in S is at least as long as v .

Let $U = 2^{1/(p\alpha)}$. Split S into sets S_1 and S_2 of shorter and longer links, where $S_1 := \{w | d_{ww} \leq U \cdot d_{vv}\}$ and $S_2 := S \setminus S_1$. Partition S_1 into $\lceil \lg U \rceil = \lceil \frac{1}{p\alpha} \rceil$ length-classes and further partition each such class into at most $\lceil \frac{2\alpha+1}{\beta} \rceil$ sets that are 2-independent, using Lemma 2.12. For each such 2-independent length-class X , we invoke Lemma 3.7 (with $q = 1$) to obtain that $a_v^p(X) \leq b_v^p(X) \leq c_1$. Thus,

$$(3.15) \quad \hat{a}_v^p(S_1) = a_v^p(S_1) \leq c_1 \left[\frac{1}{p\alpha} \right] \cdot \left[\frac{2\alpha+1}{\beta} \right] \leq \frac{1}{4} c_2 .$$

Consider next set S_2 of links longer than $U \cdot d_{vv}$. Let ℓ_W be the threshold for the length of weak links, i.e., a link is weak if and only if its length is greater than ℓ_W . If $U \cdot d_{vv} > \ell_W$, then S_2 is empty, since S consists of only non-weak links, so we are done. Assume, therefore, in the remainder that $U \cdot d_{vv} \leq \ell_W$. Form a new non-weak link $u = (s_u, r_u)$ with $s_u = s_v$ and r_u chosen arbitrarily so that $d_{uu} = U \cdot d_{vv}$. The power assigned to u satisfies $\Phi_u = U^{p\alpha} \Phi_v = 2 \cdot \Phi_v$. Since S_2 contains only long links, $d_{uu} \leq U \cdot d_{vv} \leq d_{ww}$, for every $w \in S_2$, and thus $a_u^p(S_2) = \hat{a}_u^p(S_2)$.

Partition S_2 into $S'_2 = \{w \in S_2 : a_u^p(w) = 1\}$, the links that u affects heavily, and $S''_2 = S_2 \setminus S'_2$, the rest. For $w \in S'_2$, thresholding of affectance takes place, so

$$1 = a_u^p(w) \leq c_w \frac{\Phi_u \cdot d_{ww}^\alpha}{d_{uw}^\alpha \cdot \Phi_w} = c_w \frac{d_{uu}^{p\alpha} d_{ww}^{(1-p)\alpha}}{d_{uw}^\alpha},$$

that is, $d_{uw} \leq c_w^{1/\alpha} d_{uu}^p d_{ww}^{1-p}$. Thus, using the triangle inequality, we have that for any pair x, y in S'_2 with $d_{xx} \leq d_{yy}$ (and thus $c_x \leq c_y$),

$$d_{xy} \leq d_{ux} + d_{xx} + d_{uy} \leq c_x^{1/\alpha} d_{uu}^p d_{xx}^{1-p} + d_{xx} + c_y^{1/\alpha} d_{uu}^p d_{yy}^{1-p} \leq 3c_y^{1/\alpha} d_{xx}^p d_{yy}^{1-p}.$$

Hence,

$$a_x^p(y) = \min \left(1, c_y \frac{\Phi_x d_{yy}^\alpha}{\Phi_y d_{xy}^\alpha} \right) = \min \left(1, c_y \frac{d_{xx}^{p\alpha} d_{yy}^{(1-p)\alpha}}{d_{xy}^\alpha} \right) \geq \frac{1}{3^\alpha}.$$

Thus, S'_2 contains at most 3^α links, so

$$(3.16) \quad \hat{a}_v^p(S'_2) = a_v^p(S'_2) \leq 3^\alpha \leq c_2/4.$$

For links in S''_2 , affectances scale linearly with transmission power (at least up to Φ_u , since $a_u^p(w) < 1$, for each link $w \in S''_2$). This is the key to proof of the Lemma, and the reason why it does not extend to uniform power. Thus,

$$(3.17) \quad \hat{a}_v^p(S''_2) = a_v^p(S''_2) = \frac{\Phi_v}{\Phi_u} \cdot a_u^p(S''_2) = \frac{1}{2} \hat{a}_u^p(S''_2) \leq \frac{1}{2} g(n).$$

Then, by the definition of S and Bounds (3.15), (3.16) and (3.17),

$$g(n) = \hat{a}_v^p(S) = \hat{a}_v^p(S_1) + \hat{a}_v^p(S'_2) + \hat{a}_v^p(S''_2) \leq \frac{1}{4} c_2 + \frac{1}{4} c_2 + \frac{1}{2} g(n).$$

Thus, $\frac{1}{2} g(n) \leq \frac{1}{2} c_2$, or $g(n) \leq c_2$, as desired. \square

We can now complete the proof of Theorem 3.3.

THEOREM 3.3: *Every set L of links is $\mathcal{O}(1)$ -inductive independent under \mathcal{P}_p , i.e., $I_p^p(L) = \mathcal{O}(1)$, where $0 < p \leq 1$.*

Proof. Let $v \in L$ be a link and $S \subseteq L$ be a \mathcal{P}_p -feasible subset of links. Lemma 7 of [45] states that $\hat{a}_S^p(v) = \mathcal{O}(1)$, while our Lemma 3.11 shows that $\hat{a}_v^p(S) = \mathcal{O}(1)$. Plugging this into the Def. 2.5 of \hat{b}^p , we get that $\hat{b}_v^p(S) = \hat{a}_v^p(S) + \hat{a}_S^p(v) = \mathcal{O}(1)$. Since this holds for an arbitrary combination of a link v and subset S , we conclude that $I_p^p(L) = \mathcal{O}(1)$, by Def. 3.1 of I_p^p . \square

We note that the assumption of L consisting of non-weak links is necessary when there is a fixed upper bound Φ_{max} on maximum power. Indeed, without this assumption, it can be shown that there exists a set L for which $I_p^p(L) = \Omega(\log n)$. Namely,

when dealing with weak links, all the links are operating at near full power, resulting in approximately *uniform* power assignment. A bijective transformation is given in [34] from arbitrary link sets to sets of weak links that preserves feasibility under uniform power. A logarithmic lower bound on inductive independence under uniform power is given in [32].

We can treat instances that include weak links by handling each of the $\lceil \log(c_{max}/\beta) \rceil$ tolerance classes separately using Lemma 3.7 (in the same way as in Corollary 3.10). Combined with Thm. 3.3 for the non-weak links, we obtain the following trade-off, which matches the construction in [32].

COROLLARY 3.12. Every set L of links satisfies $I_p^p(L) = \mathcal{O}(1 + \log(c_{max}(L)/\beta))$, where $0 < p \leq 1$.

4. Application: Capacity Approximation. Using the characterization described above, we can derive a simple single-pass algorithm (Algorithm 1) for maximizing capacity. It is, in fact, identical to Algorithm C of [29] that gives constant approximation for fixed power capacity. We show that it also yields asymptotically best possible approximation when measured against arbitrary power optima.

THEOREM 4.1. *Let $0 < p < 1$. Algorithm 1 is an $\mathcal{O}(\log \log \Delta)$ -approximation algorithm for Capacity in convergent metrics (that uses \mathcal{P}_p), even in the presence of weak links.*

To be more specific, it is a type of a greedy algorithm that falls under the notion of “fixed priority”, as defined by Borodin et al. [8]. Namely, the algorithm uses an initial fixed ordering of the input, and for each item decides irrevocably whether to include that item or not in the solution. Recall the d -approximation algorithm to compute a maximum independent set described in the introduction. It added vertices to the solution set in sequence, where vertices with edges to the solution so far were disqualified. Our algorithm below is its natural weighted variant: each vertex is assigned an affectance-budget of $1/2$, and is disqualified from being in the solution if the weight (affectance) of the edges to it from the solution so far exceeds the budget (Lines 4 and 5). We ensure that the final set of links is indeed \mathcal{P}_p -feasible in Line 8.

Algorithm 1 Input: Set $L = \{1, 2, \dots, n\}$ of links in non-decreasing order of length, parameter $p \in (0, 1)$. Output: \mathcal{P}_p -feasible subset $X \subseteq L$

```

1:  $R_0 \leftarrow \emptyset$ 
2: for  $i = 1$  to  $n$  do
3:    $R_i \leftarrow R_{i-1}$ 
4:   if  $\hat{b}_{R_{i-1}}^p(i) < 1/2$  then
5:      $R_i \leftarrow R_i \cup \{i\}$ 
6:   end if
7: end for
8: return  $X := \{v \in R_n : a_{R_n}^p(v) \leq 1\}$ 

```

THEOREM 4.2. *Given a set L of links, Algorithm 1 finds a \mathcal{P}_p -feasible subset $X \subseteq L$ such that $|X| \geq \frac{|S|}{2 \cdot (2I_{\mathcal{Q}}^p(L)+1)}$ for every power assignment \mathcal{Q} and every \mathcal{Q} -feasible subset $S \in F_{\mathcal{Q}}(L)$.*

The structure of the proof is inspired by that of, e.g., [42].

Proof. Let $R := R_n$ and X be the sets computed by Algorithm 1 on input L . The proof consists of two parts. In Part I we show that S is at most $2I_{\mathcal{Q}}^p(L) + 1$ times larger than R , while in Part II we relate the sizes of X and R .

Part I: Consider a power assignment \mathcal{Q} and \mathcal{Q} -feasible subset $S \subseteq L$. Let $S' := S \setminus R$. By Def. 3.1 of $I_{\mathcal{Q}}^{\mathcal{P}}(L)$, we know that $\hat{b}_i^{\mathcal{P}}(S) \leq I_{\mathcal{Q}}^{\mathcal{P}}(L)$, for each $i \in R$. Thus, since $R \subseteq L$,

$$(4.1) \quad \hat{b}_R^{\mathcal{P}}(S) \leq I_{\mathcal{Q}}^{\mathcal{P}}(L) \cdot |R| .$$

As Algorithm 1 chose none of the links in S' , the acceptance criteria of Line 4 and the definition of $\hat{b}^{\mathcal{P}}$ yield that $\hat{b}_R^{\mathcal{P}}(j) \geq \hat{b}_{R_{j-1}}^{\mathcal{P}}(j) \geq 1/2$, for each $j \in S'$. Summing over S' ,

$$(4.2) \quad \hat{b}_R^{\mathcal{P}}(S') \geq |S'|/2 .$$

Combining Bounds (4.1) and (4.2),

$$|S'| \leq 2 \cdot \hat{b}_R^{\mathcal{P}}(S') \leq 2 \cdot \hat{b}_R^{\mathcal{P}}(S) \leq 2I_{\mathcal{Q}}^{\mathcal{P}}(L) \cdot |R| .$$

Thus,

$$(4.3) \quad |S| = |S'| + |R| \leq (2I_{\mathcal{Q}}^{\mathcal{P}}(L) + 1)|R| .$$

Part II: We next show that the set R found by Algorithm 1 has small mean internal affectance. Observe that the sum of in-affectances is bounded by

$$\begin{aligned} a_R^{\mathcal{P}}(R) &= \sum_{i \in R} \sum_{j \in R} a_j^{\mathcal{P}}(i) \stackrel{1.}{=} \sum_{i \in R} \sum_{j \in R: j < i} (a_j^{\mathcal{P}}(i) + a_i^{\mathcal{P}}(j)) \\ &\stackrel{2.}{=} \sum_{i \in R} \sum_{j \in R: j < i} \hat{b}_j^{\mathcal{P}}(i) \stackrel{3.}{=} \sum_{i \in R} \hat{b}_{R_{i-1}}^{\mathcal{P}}(i) \stackrel{4.}{\leq} \frac{1}{2}|R| , \end{aligned}$$

with the numbered transformation explained as follows:

1. By rearrangement. Here $j < i$ refers to the indices of the links as sorted by Algorithm 1. We also use that by the Def. 2.3 of affectance, $\sum_{i \in R} a_i^{\mathcal{P}}(i) = 0$.
2. From the precondition of Algorithm 1, $j < i$ implies that $d_{jj} \leq d_{ii}$. Thus, $\hat{b}_j^{\mathcal{P}}(i) = a_j^{\mathcal{P}}(i) + a_i^{\mathcal{P}}(j)$, by the Def. 2.5 of \hat{b} .
3. Since $R_{i-1} = \{j : j \in R, j < i\}$ as specified by Algorithm 1.
4. By the acceptance criteria of Line 4 of the algorithm.

This implies that the average in-affectance is $\frac{1}{|R|} a_R^{\mathcal{P}}(R) \leq \frac{1}{2}$.

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

$$(4.4) \quad |X| = |\{v \in R \mid a_R^{\mathcal{P}}(v) \leq 1\}| \geq \frac{1}{2}|R| .$$

Combining Bounds (4.3) and (4.4) yields the statement of the theorem. \square

Now we are ready to prove Theorem 4.1

Proof. [of Theorem 4.1] Let L be the set of non-weak input links and W be the set of weak links. By Theorem 4.2, running Algorithm 1 on L with power \mathcal{P}_p yields a solution with capacity at most $\mathcal{O}(1 + I_{\mathcal{Q}}^{\mathcal{P}}(L))$ -factor smaller than the optimum for Capacity. By Theorem 3.2 this amounts to a $\mathcal{O}(\log \log \Delta)$ factor.

On the set W of weak links, we apply the constant-factor approximation algorithm of [29] for *uniform power* capacity (setting all powers to the maximum power Φ_{max}). We claim that our solution on W is within a constant factor of the size of OPT_W , the optimal solution on W . Observe that links in OPT_W must use power that is

a constant fraction of Φ_{max} , since the links are weak and would otherwise not be feasible. More precisely, one can work out that the power is at least $1/c$ -fraction of Φ_{max} , where $c = 2^{1/(1-p)-1}$. By raising the power on all the links up to Φ_{max} , we increase the affectance by at most a factor of c . By signal strengthening [5], the optimal uniform capacity is then at least $|OPT_W|/(2c)$, and thus the claim.

We finally output the larger of the solutions on L and on W . The theorem follows.

□

5. Further Applications. Both of our structural results have a number of further 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 $I_Q^p(L)$ or $I_p^p(L)$ (or both). Plugging in our improved bounds for these interference measures thus gives the (poly)-logarithmic improvements for a variety of applications.

5.1. Connectivity. Wireless connectivity — the problem of *efficiently* connecting a set of wireless nodes in an interference aware manner — is one of the central problems in wireless network research [31]. Such a structure may underlie a multi-hop wireless network, or provide the underlying backbone for synchronized operation of an ad-hoc 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 $\mathcal{O}(\log(n) \cdot \Upsilon)$ slots using mean power in both centralized [31] and distributed [30] algorithms, where $\Upsilon = \Upsilon(\Delta, n)$ is the maximum ratio, over all instances of n links with length ratio Δ , between the optimal capacity with arbitrary power and the optimal capacity with mean power. More precisely, if T is the minimum spanning tree of a set of points in the Euclidean plane, and T contains no weak links, then it can be scheduled in $\mathcal{O}(\Upsilon \log n)$ slots. Theorem 4.1 implies a tight bound on Υ .

THEOREM 5.1. *Suppose there is no upper bound on transmission power, or that all nodes are within single hop communication distance. Then, any set of nodes in a convergent metric can be strongly connected in $\mathcal{O}(\log(n) \cdot \log \log \Delta)$ slots using power assignment \mathcal{P}_p , where $0 < p < 1$. This can be computed by either a poly-time centralized algorithm or an $\mathcal{O}(\text{poly}(\log n) \cdot \log \Delta)$ -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 [31] for a discussion of these problems and their applications.

5.2. 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, [4]), the problem of dynamic allocation of channels to bidders (these are the wireless devices) via an auction has attracted much attention [58, 59].

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 subsets of channels, find an allocation of the users to channels so that each channel is assigned a feasible set (w.r.t. given restriction of the power control) and the social welfare is maximized.

For the SINR model, recent work [37, 36] has established a number of results depending on different valuation functions. Since these results are based on the in-

ductive independence number, Theorem 3.3 improves virtually all of them by a $\log n$ factor as we argue below. For instance, an algorithm was given in [37] for general valuations that achieves an $\mathcal{O}(\sqrt{k} \log n \cdot I_p^p(L)) = \mathcal{O}(\sqrt{k} \log^2 n)$ -approximation. We achieve an improved result by simply plugging in Theorem 3.3.

COROLLARY 5.2. *Consider the combinatorial auction problem in the SINR setting, for any fixed power assignment \mathcal{P}_p with $0 < p \leq 1$. There exist algorithms that achieve an $\mathcal{O}(\sqrt{k} \log n)$ -factor for general valuations [37], a $\mathcal{O}(\log(n) + \log k)$ -approximation for symmetric valuations and an $\mathcal{O}(\log n)$ -approximation for Rank-matroid valuations [36].*

5.3. Dynamic Packet Scheduling. Dynamic packet scheduling to achieve network *stability* is one of the fundamental problems in (wireless) network queuing theory [54]. In spite of its long history, this fundamental problem has been considered only recently in the SINR model (see [3, 44, 47]). The problem calls for an algorithm that can keep queue sizes bounded in a wireless network under stochastic arrivals of packets at transmitters. A measure called *efficiency* captures the worst-case fraction of the optimal throughput a given algorithm achieves. 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 (for performance guarantees in terms of the parameter n) have been achieved recently and simultaneously in [3] and [44]. In spite of differences in the algorithms and assumptions made, both are based on the scheduling algorithm of [45] and achieve a similar result.

To prove this result, we introduce another complexity measure.

DEFINITION 5.3 ([45], Def. 1). *The maximum average affectance $A^p(L)$ of a link set L is $A^p(L) := \max_{S \subseteq L} \frac{a_S^p(S)}{|S|}$.*

Let $\chi^{\mathcal{Q}}(L)$ be the minimum number of slots in a \mathcal{Q} -feasible schedule of L and let $\bar{\chi}(L)$ denote the minimum number of slots in a \mathcal{P} -feasible schedule of L .

LEMMA 5.4. *Let L be a set of links and \mathcal{P}, \mathcal{Q} be power assignments for L . Then,*

$$A^{\mathcal{P}}(L) \leq I_{\mathcal{Q}}^{\mathcal{P}}(L) \cdot \chi^{\mathcal{Q}}(L) .$$

Proof. Let S be a set that maximizes $A^{\mathcal{P}}(L)$, i.e., $A^{\mathcal{P}}(L) = \frac{a_S^{\mathcal{P}}(S)}{|S|}$. Let v be a link in L and let $S' = S \setminus \{v\}$. Then, by the assumption about S ,

$$A^{\mathcal{P}}(L) \geq \frac{a_{S'}^{\mathcal{P}}(S')}{|S'|} = \frac{a_S^{\mathcal{P}}(S) - b_v^{\mathcal{P}}(S)}{|S| - 1} = \frac{|S|A^{\mathcal{P}}(L) - b_v^{\mathcal{P}}(S)}{|S| - 1} .$$

Rearranging, we obtain that $b_v^{\mathcal{P}}(S) \geq A^{\mathcal{P}}(L)$. In particular, this holds for the *shortest* link w in S , and so by the definition of length-ordered affectance and its additivity,

$$(5.1) \quad A^{\mathcal{P}}(L) = b_w^{\mathcal{P}}(S) \leq \hat{b}_w^{\mathcal{P}}(S) \leq \hat{b}_w^{\mathcal{P}}(L) .$$

Let I_1, I_2, \dots, I_t be a partition of L into \mathcal{Q} -feasible sets, where $t = \chi^{\mathcal{Q}}(L)$. Then, using the additivity of affectance and the definition of the interference measure, we get that

$$\max_{v \in L} \hat{b}_v^{\mathcal{P}}(L) = \max_{v \in L} \sum_{i=1}^t \hat{b}_v^{\mathcal{P}}(I_i) \leq \sum_{i=1}^t \max_{v \in L} \hat{b}_v^{\mathcal{P}}(I_i) \leq \sum_{i=1}^t I_{\mathcal{Q}}^{\mathcal{P}}(L) = t \cdot I_{\mathcal{Q}}^{\mathcal{P}}(L) .$$

Combining this with (5.1) yields the lemma. \square

Our structural theorems imply the following bounds on the maximum average affectance.

COROLLARY 5.5. *For any set of links L and power assignment \mathcal{P}_p with $0 < p < 1$, $A^p(L) = \mathcal{O}(\overline{\chi(L)} \cdot \log \log \Delta)$ (in convergent metrics) and $A^p(L) = \mathcal{O}(\chi^p(L))$ (in general metrics).*

The result in [3, 44] can be succinctly expressed as follows.

THEOREM 5.6. [3, 44] *There exists a distributed algorithm that for any link set L achieves $\Omega\left((1 + \phi(L))^{-1} \cdot \frac{1}{\log n}\right)$ -efficiency.*

Since the best bound on $\phi(L)$ known was $\mathcal{O}(\log n)$ [45], both papers claimed $\Omega\left(\frac{1}{\log^2 n}\right)$ -efficiency. Results in this paper show that $\phi(L) = \mathcal{O}(1)$ for non-weak links (see second part of Corollary 5.5), which gives the following improved result:

COROLLARY 5.7. *There exists a distributed algorithm that achieves $\Omega\left(\frac{1}{\log n}\right)$ -efficiency for any power assignment \mathcal{P}_p ($0 < p \leq 1$) and set L of non-weak links.*

Since Corollary 5.5 also shows that $\phi(L) = \frac{A^p(L)}{\chi(L)} = \mathcal{O}(\log \log \Delta(L))$, we also get the following improved bound for power control:

COROLLARY 5.8. *There is a distributed algorithm for non-weak links in a convergent metric with efficiency $\Omega\left(\frac{1}{\log(n) \cdot \log \log \Delta(L)}\right)$, with respect to power control optima.*

Appendix A. Missing Proof from Section 3.

PROPOSITION 3.6 *Let S be a 2-independent length-class and v be a link not necessarily in S . Let u be the link in S with d_{vu} minimum. Then, $d_{vw} \geq d_{wu}/6$, for any link w in S .*

Proof. The reader might find Figure 4 helpful while reading the proof. Consider a link $w \in S$. Let $D := d_{vw}$ and note that by the choice of u we have that $d_{vu} \leq D$. By the triangle inequality and the choice of u ,

$$(A.1) \quad d_{uw} \leq d_{uu} + d_{vu} + d_{vw} \leq 2D + d_{uu} .$$

Similarly,

$$(A.2) \quad d_{wu} \leq d_{ww} + d_{vw} + d_{vu} \leq 2D + d_{ww} .$$

Now we recall the Def. 1.1 of 2-independence and apply it to u and w . Next we bound d_{wu} and d_{uw} by Bounds (A.2) and (A.1) to obtain that

$$4d_{uu}d_{ww} \leq d_{uw} \cdot d_{wu} \leq (2D + d_{uu}) \cdot (2D + d_{ww}) .$$

We know that $d_{uu}d_{ww} > 0$, therefore $(2D + d_{uu})$ and $(2D + d_{ww})$ must be either both > 0 or both < 0 . This implies that D must be at least $\min(d_{uu}, d_{ww})/2$, which in turn is at least $\max(d_{uu}, d_{ww})/4$, since the links are nearly-equal length. Thus we can bound $d_{ww} \leq 4D$ in Bound (A.2) to obtain that $d_{wu} \leq 6D$. \square

Acknowledgments: We thank Marijke Bodlaender and Tigran Tonoyan for helpful comments.

REFERENCES

- [1] M. ANDREWS AND M. DINITZ, *Maximizing capacity in arbitrary wireless networks in the SINR model: Complexity and game theory*, in Proc. of the 28th International Conference on Computer Communications, INFOCOM, Rio de Janeiro, IEEE, 2009, pp. 1332–1340.




FIG. 4. Links u, v and w as used in the proof of Proposition 3.6. The distances d_{vw} and d_{wu} that are related to each other in the proposition's statement are represented by red dotted lines. The gray dashed lines mark distances d_{uw} and d_{vu} that are used in the proofs as well.

- [2] E. ÁSGEIRSSON AND P. MITRA, *On a game theoretic approach to capacity maximization in wireless networks*, in INFOCOM, 2011.
- [3] E. I. ÁSGEIRSSON, M. M. HALLDÓRSSON, AND P. MITRA, *Wireless network stability in the SINR model*, in Structural Information and Communication Complexity (SIROCCO), G. Even and M. M. Halldórsson, eds., vol. 7355 of Lecture Notes in Comput. Sci., Springer, Heidelberg, 2012, pp. 123–134.
- [4] P. BAHL, R. CHANDRA, T. MOSCIBRODA, R. MURTY, AND M. WELSH, *White space networking with Wi-Fi like connectivity*, in Proc. of the SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, Barcelona, ACM, 2009, pp. 27–38.
- [5] J. BANG-JENSEN AND M. M. HALLDÓRSSON, *Vertex coloring edge-weighted digraphs*, Information Processing Letters, 115 (2015), pp. 791–796.
- [6] D. M. BLOUGH, G. RESTA, AND P. SANTI, *Approximation Algorithms for Wireless Link Scheduling with SINR-based Interference*, IEEE/ACM Trans. Netw., 18 (2010), pp. 1701–1712.
- [7] M. H. L. BODLAENDER AND M. M. HALLDÓRSSON, *Beyond geometry: Towards fully realistic wireless models*, in Proc. of the 33rd ACM Symposium on Principles of Distributed Computing, PODC '14, Paris, 2014, pp. 347–356.
- [8] A. BORODIN, M. N. NIELSEN, AND C. RACKOFF, *(Incremental) priority algorithms*, Algorithmica, 37 (2003), pp. 295–326.
- [9] D. CHAFEKAR, V. S. KUMAR, M. MARATHE, S. PARTHASARATHY, AND A. SRINIVASAN, *Cross-layer latency minimization for wireless networks using SINR constraints*, in Proc. of the 8th International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc, Montreal, ACM, 2007, pp. 110–119.
- [10] J. DAMS, M. HOEFER, AND T. KESSELHEIM, *Scheduling in wireless networks with Rayleigh-fading interference*, in Proc. of the 24th Symposium on Parallelism in Algorithms and Architectures, SPAA, Pittsburgh, PA, ACM, 2012, pp. 327–335.
- [11] J. DAMS, M. HOEFER, AND T. KESSELHEIM, *Sleeping experts in wireless networks*, in Distributed Computing, DISC, Springer, 2013, pp. 344–357.
- [12] ———, *Jamming-resistant learning in wireless networks*, in Automata, Languages, and Programming, ICALP, Springer, 2014, pp. 447–458.
- [13] S. DAUM, S. GILBERT, F. KUHN, AND C. NEWPORT, *Broadcast in the Ad Hoc SINR Model*, in DISC, 2013, pp. 358–372.
- [14] M. DINITZ, *Distributed algorithms for approximating wireless network capacity*, in Proc. of the 29th International Conference on Computer Communications, INFOCOM, San Diego, CA, IEEE, 2010, pp. 1–9.
- [15] A. FANGHÄNEL, T. KESSELHEIM, H. RÄCKE, AND B. VÖCKING, *Oblivious interference scheduling*, in Proc. of the 28th Symposium on Principles of Distributed Computing, PODC, Calgary, ACM, 2009, pp. 220–229.
- [16] A. FANGHÄNEL, T. KESSELHEIM, AND B. VÖCKING, *Improved algorithms for latency minimization in wireless networks*, in Automata, Languages and Programming, ICALP, U. Montanari, J. D. P. Rolim, and E. Welzl, eds., vol. 1853 of Lecture Notes in Comput. Sci., Springer, Heidelberg, 2000, pp. 447–458.
- [17] Y. GAO, J. C. HOU, AND H. NGUYEN, *Topology control for maintaining network connectivity*

- and maximizing network capacity under the physical model, in Proc. of the 27th International Conference on Computer Communications, INFOCOM, Phoenix, AZ, IEEE, 2008, pp. 1013–1021.
- [18] O. GOUSSEVSKAIA, M. M. HALLDÓRSSON, AND R. WATTENHOFER, *Algorithms for wireless capacity*, IEEE/ACM Transactions on Networking (TON), 22 (2014), pp. 745–755.
- [19] O. GOUSSEVSKAIA, M. M. HALLDÓRSSON, R. WATTENHOFER, AND E. WELZL, *Capacity of Arbitrary Wireless Networks*, in Proc. of the 28th International Conference on Computer Communications, INFOCOM, Rio de Janeiro, IEEE, 2009, pp. 1872–1880.
- [20] O. GOUSSEVSKAIA, T. MOSCIBRODA, AND R. WATTENHOFER, *Local Broadcasting in the Physical Interference Model*, in Proc. of the 5th International Workshop on Foundations of Mobile Computing, DialM-POMC, Toronto, ACM, 2008, pp. 35–44.
- [21] O. GOUSSEVSKAIA, Y. A. OSWALD, AND R. WATTENHOFER, *Complexity in geometric SINR*, in Proc. of the 8th International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc, Montreal, ACM, 2007, pp. 100–109.
- [22] O. GOUSSEVSKAIA AND R. WATTENHOFER, *Complexity of scheduling with analog network coding*, in Proc. of the 1st International Workshop on Foundations of Wireless Ad Hoc and Sensor Networking and Computing, FOWANC, Hong Kong, ACM, 2008, pp. 77–84.
- [23] J. GRÖNKVIST AND A. HANSSON, *Comparison Between Graph-Based and Interference-Based TDMA Scheduling*, in Proc. of the 2nd International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc 2001, Long Beach, CA, ACM, 2001, pp. 255–258.
- [24] P. GUPTA AND P. R. KUMAR, *The Capacity of Wireless Networks*, IEEE Trans. Inf. Theory, 46 (2000), pp. 388–404.
- [25] M. M. HALLDÓRSSON, *Approximations of weighted independent set and hereditary subset problems*, J. Graph Algorithms Appl., 4 (2000), pp. 1–16.
- [26] M. M. HALLDÓRSSON, *Wireless scheduling with power control*, ACM Trans. Algorithms, 9 (2012), pp. 7:1–7:20.
- [27] M. M. HALLDÓRSSON, S. HOLZER, P. MITRA, AND R. WATTENHOFER, *The power of non-uniform wireless power*, in Proc. of the 24th Symposium on Discrete Algorithms, SODA, New Orleans, LA, ACM-SIAM, 2013, pp. 1595–1606.
- [28] M. M. HALLDÓRSSON AND P. MITRA, *Nearly optimal bounds for distributed wireless scheduling in the SINR model*, in Automata, Languages and Programming, ICALP, L. Aceto, M. Henzinger, and J. Sgall, eds., vol. 6756 of Lecture Notes in Comput. Sci., Springer, Heidelberg, 2011, pp. 625–636.
- [29] ———, *Wireless Capacity with Oblivious Power in General Metrics*, in Proc. of the 22nd Symposium on Discrete Algorithms, SODA, San Francisco, CA, ACM-SIAM, 2011, pp. 1538–1548.
- [30] ———, *Distributed connectivity of wireless networks*, in Proc. of the 31st Symposium on Principles of Distributed Computing, PODC, Madeira, ACM, 2012, pp. 205–214.
- [31] ———, *Wireless connectivity and capacity*, in Proc. of the 23rd Symposium on Discrete Algorithms, SODA, Kyoto, ACM-SIAM, 2012, pp. 516–526.
- [32] ———, *Nearly optimal bounds for distributed wireless scheduling in the SINR model*, Distributed Computing, (2016). To appear.
- [33] M. M. HALLDÓRSSON AND T. TONoyAN, *How well can graphs represent wireless interference?*, in Proc. of the 47th Annual ACM Symposium on Theory of Computing, STOC, 2015, pp. 635–644.
- [34] M. M. HALLDÓRSSON AND T. TONoyAN, *The Price of Local Power Control in Wireless Scheduling*, in 35th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS, vol. 45 of Leibniz International Proceedings in Informatics (LIPIcs), 2015, pp. 529–542.
- [35] M. M. HALLDÓRSSON AND R. WATTENHOFER, *Wireless Communication is in APX*, in Automata, Languages and Programming, ICALP, vol. 5555 of Lecture Notes in Comput. Sci., Springer, Heidelberg, Germany, 2009, pp. 525–536.
- [36] M. HOEFER AND T. KESSELHEIM, *Secondary spectrum auctions for symmetric and submodular bidders*, ACM Transactions on Economics and Computation, 3 (2015), p. 9.
- [37] M. HOEFER, T. KESSELHEIM, AND B. VÖCKING, *Approximation algorithms for secondary spectrum auctions*, ACM Transactions on Internet Technology (TOIT), 14 (2014), p. 16.
- [38] Q.-S. HUA AND F. C. M. LAU, *The scheduling and energy complexity of strong connectivity in ultra-wideband networks*, in Proc. of the 9th International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems, MSWiM, Terromolinos, ACM, 2006, pp. 282–290.
- [39] K. JAIN, J. PADHYE, V. N. PADMANABHAN, AND L. QIU, *Impact of interference on multi-hop wireless network performance*, Wireless Networks, 11 (2005), pp. 471–487.

- [40] T. JURDZINSKI, D. KOWALSKI, M. ROZANSKI, AND G. STACHOWIAK, *On the impact of geometry on ad hoc communication in wireless networks*, in PODC '14, 2014.
- [41] E. KANTOR, Z. LOTKER, M. PARTER, AND D. PELEG, *The Topology of Wireless Communication*, in Proc. of the 43rd Symposium on Theory of Computing, STOC, San Jose, CA, ACM, 2011, pp. 383–392.
- [42] T. KESSELHEIM, *A Constant-Factor Approximation for Wireless Capacity Maximization with Power Control in the SINR Model*, in Proc. of the 22nd Symposium on Discrete Algorithms, SODA, San Francisco, CA, ACM-SIAM, 2011, pp. 1549–1559.
- [43] ———, *Approximation algorithms for wireless link scheduling with flexible data rates*, in Algorithms - ESA 2012 (proc.), L. Epstein and P. Ferragina, eds., vol. 7501 of Lecture Notes in Comput. Sci., Springer, Heidelberg, 2012, pp. 659–670.
- [44] ———, *Dynamic packet scheduling in wireless networks*, in Proc. of the 31st Symposium on Principles of Distributed Computing, PODC 2012, Funchal, ACM, 2012, pp. 281–290.
- [45] T. KESSELHEIM AND B. VÖCKING, *Distributed contention resolution in wireless networks*, in Distributed Computing, DISC, N. A. Lynch and A. A. Shvartsman, eds., vol. 6343 of Lecture Notes in Comput. Sci., Springer, Heidelberg, 2010, pp. 163–178.
- [46] V. KUMAR, M. MARATHE, S. PARTHASARATHY, AND A. SRINIVASAN, *Algorithmic Aspects of Capacity in Wireless Networks*, ACM SIGMETRICS Perf. Eval. Rev., 33 (2005), pp. 133–144.
- [47] L. B. LE, E. MODIANO, C. JOO, AND N. B. SHROFF, *Longest-queue-first scheduling under SINR interference model*, in Proc. of the 11th International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc, Chicago, IL, ACM, 2010, pp. 41–50.
- [48] R. MAHESHWARI, S. JAIN, AND S. R. DAS, *A measurement study of interference modeling and scheduling in low-power wireless networks*, in Proc. of the 6th Conference on Embedded Network Sensor Systems, SenSys, Raleigh, NC, ACM, 2008, pp. 141–154.
- [49] T. MOSCIBRODA, *The worst-case capacity of wireless sensor networks*, in Proc. of the 6th International Symposium on Information Processing in Sensor Networks, IPSN, Cambridge, MA, ACM, 2007, pp. 1–10.
- [50] T. MOSCIBRODA AND R. WATTENHOFER, *The Complexity of Connectivity in Wireless Networks*, in Proc. of the 25th International Conference on Computer Communications, INFOCOM 2006, Barcelona, IEEE, 2006, pp. 1–13.
- [51] T. MOSCIBRODA, R. WATTENHOFER, AND Y. WEBER, *Protocol design beyond graph-based models*, in Proc. of the 5th Workshop on Hot Topics in Networks, HotNets 2006, Irvine, CA, ACM, 2006, pp. 25–30.
- [52] T. MOSCIBRODA, R. WATTENHOFER, AND A. ZOLLINGER, *Topology control meets SINR: The scheduling complexity of arbitrary topologies*, in Proc. of the 7th International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc, Florence, ACM, 2006, pp. 310–321.
- [53] 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 Proc. of the 9th International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc, Hong Kong, ACM, 2008, pp. 91–100.
- [54] L. TASSIULAS AND A. EPHREIMIDES, *Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks*, IEEE Trans. Autom. Control, 37 (1992), pp. 1936–1948.
- [55] T. TONoyAN, *On the capacity of oblivious powers*, in Algorithms for Sensor Systems, ALGOSENSORS, T. Erlebach, S. E. Nikolettseas, and P. Orponen, eds., vol. 7111 of Lecture Notes in Comput. Sci., Springer, Heidelberg, 2011, pp. 225–237.
- [56] ———, *On some bounds on the optimum schedule length in the SINR model*, in Algorithms for Sensor Systems, 8th International Symposium on Algorithms for Sensor Systems, Wireless Ad Hoc Networks and Autonomous Mobile Entities, ALGOSENSORS 2012, Ljubljana, Slovenia, September 13–14, 2012. Revised Selected Papers, 2012, pp. 120–131.
- [57] Y. YE AND A. BORODIN, *Elimination graphs*, ACM Trans. Algorithms, 8 (2012), pp. 14:1–14:23.
- [58] X. ZHOU, S. GANDHI, S. SURI, AND H. ZHENG, *eBay in the Sky: Strategy-proof wireless spectrum auctions*, in Proc. of the 14th International Conference on Mobile Computing and Networking, MOBICOM, San Francisco, CA, ACM, 2008, pp. 2–13.
- [59] X. ZHOU AND H. ZHENG, *Trust: A general framework for truthful double spectrum auctions*, in Proc. of the 28th International Conference on Computer Communications, INFOCOM, Rio de Janeiro, IEEE, 2009, pp. 999–1007.