THE POWER OF NON-UNIFORM WIRELESS POWER

MAGNÚS M. HALLDÓRSSON*, STEPHAN HOLZER†, PRADIPTA MITRA*,
AND ROGER WATTENHOFER†

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 — 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, connectivity, secondary spectrum auctions, and dynamic packet scheduling.

1. Introduction

One of the strongest weapons for increasing 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 simultaneously transmitting links. 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 [38, 18] or even unbounded [8] 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 [30]. 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, which leaves it to only depend on the distance between transmitter and receiver. This is known as *oblivious* power control.

Many questions immediately rise 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?

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 answers a number of these questions optimally, completing an extensive line of work in the algorithmic study of the SINR model.

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 transmit simultaneously.

*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.

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 $\mathcal{O}(\log n \cdot \log\log \Delta)$ [18] and $\mathcal{O}(\log n + \log\log \Delta)$ [22], 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 $\Omega(n)$ [8] applied only when Δ was doubly exponential. In terms of Δ, it was shown that one *must* pay an $\Omega(\log\log \Delta)$ factor [18]. The best upper bounds were, as mentioned, either dependent on the size of the input [18, 22] and as such unbounded (in relation to Δ), or exponentially worse ($\mathcal{O}(\log \Delta)$) [1, 13].

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 $p = \frac{1}{2}$ results in 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, achieving a solution quality within an $\mathcal{O}(\log\log \Delta)$-factor of the optimum with unrestricted power control. This is an asymptotically optimal solution quality for this class of schemes according to [18]. 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 all previous bounds, including the $\mathcal{O}(\log \Delta)$-bound of [1] (see also [13]).

This result extends the "star status" from mean power to a large class of assignments. This class has been studied implicitly before in a wide array of work [22, 21, 27, 33] on "length-monotone, sub-linear" power assignments, but its relation to arbitrary power was not understood.

Our second main contribution is to improve a number of algorithmic results that use these power assignments. We shave a logarithmic approximation factor off a variety of problems, including distributed scheduling [33], secondary spectrum auctions [27], wireless connectivity [23, 24, 38], and dynamic packet scheduling [2, 32]. 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 this host of 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, it is useful to recall the graph theoretic notion of *inductive independence* [45]. A graph G has inductive independence number d if there is an ordering of the vertices $v_1, v_2, \ldots, v_n$ such that each v_i has at most d edges to any independent set $I \subseteq \{v_{i+1}, v_{i+2}, \ldots, v_n\}$. An example is provided in Figure 1. The inductive independence

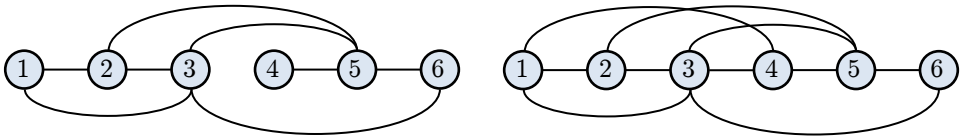

FIGURE 1. The graph on the left has inductive independence number 1, the graph on the right has inductive independence number 2.

property is found in many graph classes (e.g., intersection graphs of convex planar objects are 3-inductive independent [45]), and it has powerful algorithmic implications [19, 27, 45]. 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 edges to nodes 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, links are vertices, and the edge weights represent the extent of interference between links. The relevant ordering of the links is the ascending order by length, and "independent sets" are represented by *feasible sets* of links (links that can transmit simultaneously).

When feasibility is with respect to *arbitrary* power assignment, we show that the measure is bounded by $\mathcal{O}(\log\log\Delta)$ (Theorem 3), implying our first capacity result (and its applications). Technically, this is done by carefully extending the analysis of [18]. When feasibility is with respect to oblivious power from the above mentioned class, the measure can be bounded by a constant (Theorem 4), implying the second set of algorithmic results. This involves a potentially novel contradiction technique (at least in the context of SINR analysis).

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

1.2. **Related Work.** Gupta and Kumar [17] 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 [34, 41] but such models do not capture the nature of wireless communication well, as demonstrated in [16, 36, 39]. Seven years ago, Moscibroda and Wattenhofer [38] started combining 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 [10, 29, 40], sensor networks [37], combined scheduling and routing [5], ultra-wideband [28], and analog network coding [15].

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 [13], a constant factor approximation algorithm for uniform power was achieved in [11, 25], and eventually extended to essentially all interesting oblivious power schemes in [22]. In [30, 31], 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 [18].

Linear power has turned out to be the easiest among fixed power assignments, being the only one with a constant factor approximation for scheduling [9, 44] and a constant-bounded interference measure [9]. Whereas there are instances for which linear and uniform power are arbitrarily bad in comparison with mean power [38], a maximum feasible subset under mean power is known to be always within a constant factor of subsets feasible under linear or uniform power [43]. Recently it was shown in [6] 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. This overview is far from being complete, surveys can be found in e.g. [14].

Technically, the idea of looking at the interaction between a feasible set and a link was studied before. The works of Halldórsson [18] and Kesselheim and Vöcking [33] are particularly 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 papers to the best possible up to a constant factor.

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. We follow this in Section 4 by illustrating two applications of these results,

one for each of the main theorems. Section 5 contains the proofs of the structural results, and Section 6 contains a medley of further applications.

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 transmitter 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.

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

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

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.

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 ($\mathcal{P}_1$).

We say that S is $\mathcal{P}$-*feasible*, if Eqn. 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 the links in L that is PC-feasible (that is we maximize the capacity of the channel used). 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 arbitrary power assignment (i.e., size of the largest PC-feasible subset).

Affectance We use the notion of *affectance*, introduced in [11] and refined in [25] and [33]. 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

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

where the factor $c_v := \beta/(1 - \beta N \ell_v^\alpha/P_v)$ depends only on properties of the link l_v and on universal constants. We let $a_v^p(w)$ denote $a_v^{\mathcal{P}_p}(w)$. We frequently drop the power assignment reference $\mathcal{P}$, which means then that we assume $\mathcal{P}_p$. Conventionally, we define $a_v(v) := 0$, since v does not interfere with itself. 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 $a_S^{\mathcal{P}}(v) \leq 1$ (except for the near-trivial case of S containing only two links).

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

(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 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). Our theorems often assume links to be non-weak. This reasonable and often-used assumption [1, 7, 12, 33] can be achieved, if necessary, by scaling the powers.

Length classes A *length class* is any set R of links with $\Delta(R) \leq 2$ (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 $d_{vw} \cdot d_{wv} \geq q^2 \cdot \ell_w \ell_v$. A set of mutually q-independent links is said to be *q-independent*. An example of 1-independence is given in Figure 2.

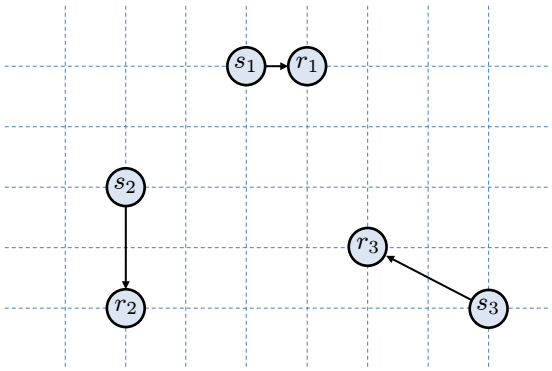

FIGURE 2. Links l_1 and l_2 are 1.92-independent. Set $\{l_1, l_2, l_3\}$ is 1.43-independent.

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 [18], but there is no good relationship in the opposite direction.

In this paper we provide an independence-strengthening result with better tradeoffs than the so-called "signal-strengthening" result of [25]. The proof is in Appendix A.

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

3. Structural Properties

We start by defining the interference measure at the center of this work.

Definition 2. *Let L be a set of links, $\mathcal{P}, \mathcal{Q}$ be two power assignments of L, and let $F_{\mathcal{Q}}(L)$ be the collection of subsets of L that are $\mathcal{Q}$-feasible. Then,*

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

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

As mentioned in the introduction, this definition is analogous to the inductive independence number of a graph. 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 l_u and l_v is $b_u(v)$ $= b_v(u)$ (computed according to power assignment $\mathcal{P}$). The ordering is the ascending order of length. 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 is an upper bound on how many edges a vertex can have to an independent set consisting of higher-ranked vertices.

When using 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 the capacity with respect to L), allowing us to relate 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 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. The second is a constant upper bound on the function when both $\mathcal{P}$ and $\mathcal{Q}$ are the same assignment (specifically, $\mathcal{P}_p$ for some $p \in (0, 1]$).

Theorem 3. *For any set L of non-weak links, any $0 < p < 1$, and any power assignment $\mathcal{Q}$, $I^p_{\mathcal{Q}}(L) = \mathcal{O}(\log\log\Delta)$.*

Theorem 4. *Fix a power assignment $\mathcal{P}_p$ for any $0 < p \leq 1$. Then any set L of non-weak links is $\mathcal{O}(1)$-inductively independent under $\mathcal{P}_p$, i.e., $I^p_p(L) = \mathcal{O}(1)$.*

Both theorems will be proven in Section 5. Theorem 3 improves upon the $\mathcal{O}(\log\log\Delta + \log n)$ bound that is stated implicitly in [22] (and extends it to many more power assignments). Theorem 4 improves upon the $\mathcal{O}(\log n)$ bound proven in [33]. Both of these new theorems are optimal (up to constant factors).

4. Applications

Before embarking upon the rather technical proofs of Theorems 3 and 4, we highlight two applications, one for each theorem. Further implications are provided in Section 6.

4.1. **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 [22] 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. [4]. Recall the d-approximation to the max-independent set problem described in the introduction. We added vertices to the solution set in order, and vertices with edges to the solution set so far were disqualified. Our algorithm below is the natural weighted version of it – each vertex is assigned a budget of 1/2, and is disqualified from being in the solution if the weight 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 Gr(Set $L = \{l_1, l_2, \ldots, l_n\}$ of links in increasing order of length)

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

Theorem 5. *Let L be a set of links. For any $\mathcal{P}_p$ for which L is non-weak, **Gr** chooses a $\mathcal{P}_p$-feasible set X such that $|X| \geq \frac{|S|}{2\cdot(2I^p_{\mathcal{Q}}(L)+1)}$ for any power assignment $\mathcal{Q}$ and any set $S \in F_{\mathcal{Q}}(L)$.*

Proof. The structure of the proof is inspired by that of, e.g., [30]. Let $R := R_n$ and X be the sets computed by Algorithm **Gr** on input L. First, we show that the size of S is not much larger than the size of R, and then relate the size of X to R to conclude the statement.

Consider any power assignment $\mathcal{Q}$ and feasible set S as specified by the statement of the theorem. Let S' be $S' := S \setminus R$. By definition of $I^p_{\mathcal{Q}}(L)$, we know that $\hat{b}^p_i(S) \leq I^p_{\mathcal{Q}}(L)$, for each $l_i \in R$. Thus,

$$\hat{b}^p_R(S) \leq I^p_{\mathcal{Q}}(L) \cdot |R| \ , \tag{2}$$

Now, Algorithm **Gr** chose none of the links in S'. Using the acceptance criteria of Line 4 and the definition of $\hat{b}^p$ yields that $\hat{b}^p_R(l_j) \geq \hat{b}^p_{R_{j-1}}(l_j) \geq 1/2$, for each $l_j \in S'$, implying that

$$\hat{b}^p_R(S') \geq |S'|/2 \ . \tag{3}$$

Combining Eqn. 2 and Eqn. 3,

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

Thus,

$$|S| \leq |S'| + |R| \leq (2I^p_{\mathcal{Q}}(L) + 1)|R| \ . \tag{4}$$

Also, the definition of **Gr** ensures that the average affectance of links in R is small (at most half). To see this, observe that the sum of in-affectances is bounded by

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

with the numbered transformation explained as follows:

(1) By rearrangement. Here $j < i$ refers to the indices of the links as sorted by Algorithm **Gr**. We also use that by the definition of affectance, $\sum_{l_i \in R} a_i(i) = 0$.
(2) From the way **Gr** iterates over the links, $j < i$ implies that $\ell_j \leq \ell_i$. Thus $\hat{b}_j(i) = a_j(i) + a_i(j)$, by definition of $\hat{b}$.
(3) Since $R_{i-1} = \{l_j : l_j \in R, j < i\}$ as specified by **Gr**.
(4) By the acceptance criteria of Line 4 of the algorithm.

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

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

$$|X| = |\{l_v \in R | a_R(v) \leq 1\}| \geq \frac{1}{2}|R| \ . \tag{5}$$

Combining Eqn. 4 and Eqn. 5 yields the statement of the theorem. □

Theorem 6. *For any $\mathcal{P}_p$, there is an $\mathcal{O}(\log\log\Delta)$-approximation algorithm for PC-Capacity that uses $\mathcal{P}_p$.*

Proof. By Thm. 5, **Gr** uses $\mathcal{P}_p$ in producing a solution with capacity at most $\mathcal{O}(1 + I^p_{\mathcal{Q}}(L))$-factor smaller than the optimum for PC-Capacity. By Thm. 3 this amounts to a $\mathcal{O}(\log\log\Delta)$ factor. □

[1]

When there is a maximum power level and most links are weak, we can still attain the same approximation ratio, as done in [22], by solving the problem separately for the weak links using maximum power.

[1]SH: Why $\mathcal{O}(1 + I^p_{\mathcal{Q}}(L))$, not $\mathcal{O}(I^p_{\mathcal{Q}}(L))$? Because the interference measure can be much smaller than 1.

4.2. **Distributed Scheduling.** A fundamental problem in wireless algorithms is to schedule a given set of links in a minimum number of slots. For $\mathcal{P}_p$ $(0 \le p \le 1)$, $\mathcal{O}(\log n)$-approximate centralized algorithms are known [22]. In [33], the first *distributed* algorithm was given, with an $\mathcal{O}(\log^2 n)$-approximation ratio. Since it is distributed, the algorithm includes an acknowledgment mechanism (via packets sent from receivers to transmitters) to enable links to know when they have succeeded (and subsequently stop running the algorithm). Assuming "free" acknowledgments, [21] improved the bound to $\mathcal{O}(\log n)$ (using the same algorithm), but [33] remained the best result when acknowledgments have to be implemented explicitly.

Here we show that,

Theorem 7. *There is a randomized distributed $\mathcal{O}(\log n)$-approximate algorithm for $\mathcal{P}_p$-Scheduling which implements explicit acknowledgments, for any $0 \le p \le 1$.*

The cases of $p = 0$ and $p = 1$ was shown in [21]; thus, we only need to focus on $p \in (0, 1)$. To explain this result, we introduce another complexity measure.

Definition 8. [33] *The* maximum average affectance $A^p(L)$ *of a link set L is* $A^p(L) := \max_{S \subseteq L} \frac{a_S^p(S)}{|S|}$.

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

Corollary 9. *For any set L of links, $A^p(L) = \mathcal{O}(\log\log\Delta \cdot \overline{\chi(L)})$ and $A^p(L) = \mathcal{O}(\chi^p(L))$.*

It was shown in [33] that the distributed scheduling algorithm completes in $\mathcal{O}(A^p(L) \log n)$ rounds. Thus, the second bound in Corollary 9 immediately gives us the $\mathcal{O}(\log n)$-approximation. The approximation bound in [33] was worse because it only showed that $A^p(L) = \mathcal{O}(\chi^p(L) \log n)$.

For comparison with arbitrary power, we can similarly use Corollary 9 to achieve an $\mathcal{O}(\log n \cdot \log\log\Delta)$-approximation including acknowledgments, improving on the $\mathcal{O}(\log n \cdot (\log n + \log\log\Delta))$-factor implied by [33] and [22]. Let PC-Scheduling be the power-control version of the problem.

Corollary 10. *There is a randomized distributed algorithm for PC-Scheduling that is $\mathcal{O}(\log\log\Delta \cdot \log n)$-approximate with respect to arbitrary power control optima. It can use any $\mathcal{P}_p$ power assignment, $0 < p < 1$.*

5. Proofs of the Structural Results

5.1. **Proof of Theorem 3.** We need two lemmas (Lemma 11 and 14) to bound affectances of a link to and from a set of links. Denote $\hat{p} := \frac{1}{\min(1-p,p)}$ for the rest of this section.

The first lemma handles the set of long links that have relatively high affectance. It originates in [18] (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. The proof is given in Appendix B.

Lemma 11. *Let p be a constant, $0 < p < 1$, τ be a parameter, $\tau \ge 1$, 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 $\max(a_v^{\mathcal{P}}(w), a_w^{\mathcal{P}}(v)) \ge 1/\tau$ and $\ell_w \ge \Lambda \cdot \ell_v$. Then, $|Q| = \mathcal{O}(\log\log\Delta)$.*

Lemma 11 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 class 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.

Proposition 12. *Let l_v be a link. Let S be a 2-independent set of nearly-equilength links and let l_u be the link in S with d_{uv} minimum. Then, $d_{wv} \ge d_{wu}/6$, for any l_w in S.*

The reader may find Figure 3 helpful when reading the proof of Proposition 12.

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

$$\begin{aligned} d_{wu} &\leq d_{wv} + d_{uv} + d_{uu} \\ &= d_{wv} + d_{uv} + \ell_u \leq 2D + \ell_u \ . \end{aligned} \tag{6}$$

Similarly,

$$d_{uw} \leq d_{uv} + d_{wv} + \ell_w \leq 2D + \ell_w \ . \tag{7}$$

Applying 2-independence, on one hand, and multiplying Eqn. 6 and Eqn. 7, on the other hand, we have that

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

This implies that D must be at least $\min(\ell_u, \ell_w)/2$ which in turn is at least $\max(\ell_u, \ell_w)/4$, using that the links are nearly-equilength. Thus we can bound $l_u \leq 4D$ in Eqn. 6 and obtain $d_{wu} \leq 6D$. □

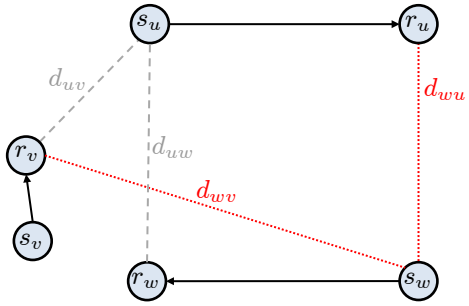

FIGURE 3. Displays links l_u, l_w and l_w as used in the proofs of Propositions 12 and 13. The distances d_{wv} and d_{wu} that are related to each other in the Proposition's statement are represented by red dotted lines. The gray dashed lines mark distances d_{uw} and d_{uv} that are used in the proofs as well.

Proposition 13. *Let l_v be a link. Let S be a 2-independent set of nearly-equilength links and let l_u be the link in S with d_{uv} minimum. Then, $d_{vw} \geq d_{wu}/6$, for any l_w in S.*

Proof. The proof is essentially the same as the proof of Proposition 13 with the roles of senders and receivers switched. For completeness, the proof can be found in Appendix B. □

This leads to the second lemma of this section.

Lemma 14. *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,*

$$b_v^p(S) \leq \max_{\{l_w, l_{w'}\} \subseteq S} b_v^p(\{w, w'\}) + \mathcal{O}(1/q).$$

Proof. Consider the link l_u in S with d_{uv} minimum. Since $\ell_v \leq \ell_u$, it holds that $c_v \leq c_u$. Then, we have that

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

using Proposition 12. Continuing from above, The above equals

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

where the last inequality follows from our choice of $\hat{p} = \frac{1}{\min(1-p,p)}$.

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

$$a^p_{S'}(v) \leq \frac{6^\alpha}{q} a^p_{S'}(u) = \mathcal{O}(1/q), \tag{8}$$

as S' is feasible.

Now consider the link $l_{u'}$ in S with $d_{vu'}$ minimum. Since links in S are non-weak, $c_w \leq 2c_{u'}$. Thus,

$$a^p_v(w) = c_w \left(\frac{\ell_v^p \ell_w^{1-p}}{d_{vw}} \right)^\alpha \leq 2c_{u'} \left(\frac{(\ell_w / q^{\hat{p}/\alpha})^p \ell_w^{1-p}}{d_{wu'}/6} \right)^\alpha ,$$

using Proposition 13 and the assumed bound on link lengths. Since $l_{u'}$ and l_w are nearly-equilength, this is bounded by

$$a^p_v(w) \leq 2c_{u'} \left(\frac{(\ell_w / q^{\hat{p}/\alpha})^p (2\ell_{u'})^{1-p}}{d_{wu'}/6} \right)^\alpha ,$$

Rearranging, we get that

$$a^p_v(w) \leq 2^{1+\alpha} \cdot \frac{6^\alpha}{q^{\hat{p} \cdot p}} \cdot c_{u'} \frac{\ell_w^{p\alpha}}{\ell_{u'}^{p\alpha}} \cdot \left(\frac{\ell_{u'}}{d_{wu'}} \right)^\alpha \leq 2^{1+\alpha} \cdot \frac{6^\alpha}{q} \cdot a^p_w \left(u' \right) \ .$$

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

$$a^p_v \left(S' \right) \leq 2^{1+\alpha} \cdot \frac{6^\alpha}{q} a^p_{S'} \left(u' \right) = \mathcal{O}(1/q), \tag{9}$$

since S' is feasible. Combining (8) and (9) yields

$$b^p_v(S) - b^p_v \left(\left\{ u, u' \right\} \right) = a^p_{S \setminus \{u,u'\}}(v) + a^p_v \left(S \setminus \left\{ u, u' \right\} \right) = \mathcal{O}(1/q),$$

from which we conclude that

$$b^p_v(S) \leq \max_{\{l_w, l_{w'}\} \subseteq S} b^p_v \left(\left\{ w, w' \right\} \right) + \mathcal{O}(1/q).$$

□

We are now ready to prove the core result, Theorem 3.

Proof. [of Theorem 3] Choose any $l_v \in L$ and any $\mathcal{Q}$-feasible subset $S \subseteq L$. We shall show that $\hat{b}^p_v(S) = \mathcal{O}(\log \log \Delta)$. By the 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, $\hat{b}^p(S) = b^p(S)$. We use the independence-strengthening lemma (Lemma 1) to partition S into at most $\frac{2^{1+\alpha}}{\beta} + 1$ different 2-independent feasible sets. Let S' be one of these sets.

Let $D := \log \Delta(L)$ and let $\Lambda = (4(4\beta D)^{1/\alpha})^{\hat{p}}$. We say that a link l_w in S is *short* if $\ell_v \leq \ell_w < \Lambda \cdot \ell_v$ and *long* if $\ell_w \geq \Lambda \cdot \ell_v$. 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.

As $S_1 \subseteq S$ is 2-independent, since its superset S is 2-independent, we can apply Lemma 11 with $\tau = 2D$. This implies that $|S_1| = \mathcal{O}(\log \log \Delta(S_1))$ and thus

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

Next we observe that due to the choice of D, the set S (and thus S_2) can be partitioned into $\mathcal{O}(D)$ length classes $X_1, X_2, \ldots$. Each such class X_i satisfies the hypothesis of Lemma 14 with $q := D \geq 1$. Since $b_v(w) < 1/D$, for each $l_w \in X_i$, by assumption, Lemma 14 gives that

$$b_v(X_i) = \mathcal{O}(1/D) \text{ and } b_v(S_2) = \mathcal{O}(D) \cdot b_v(X_i) = \mathcal{O}(1).$$

The set S_3 can be partitioned into $\log \Lambda = \mathcal{O}(D)$ length classes $Y_1, \ldots, Y_{\log \Lambda}$ as well. For each such length class Y_i, we apply Lemma 14 with $q = 1$, giving that $b_v(Y_i) = \mathcal{O}(1)$, for a total of $b_v(S_3) = \mathcal{O}(D) = \mathcal{O}(\log \log \Delta)$. Thus,

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

and

$$b_v(S) \leq \left(\frac{2^\alpha}{\beta} + 1\right) b_v(S') = \mathcal{O}(\log \log \Delta).$$

□

To provide more intuition behind this proof, consider Figure 4, where Δ is a small constant. ThereforeTODO: MMH: ?? the set S_3 contains only links of similar length to l_v. Since S_3 is feasible for some power assignment, these links are not too close to each other and can thus be scheduled within a few time slots using $\mathcal{P}_p$. Since all long links in this example are of roughly the same length, they can be partitioned in S_1 and S_2 using a disc around l_v. Note that in general, the radius of the disc that is used to decide whether a link is in S_1 or S_2 depends on the link's length relative to l_v. Now we see that l_v hardly interferes with a set of feasible links far away from l_v. There can also not be too much interference with long links that are close to l_v since there can't be too many of them within the disc as they also need to be feasible for some power assignment.

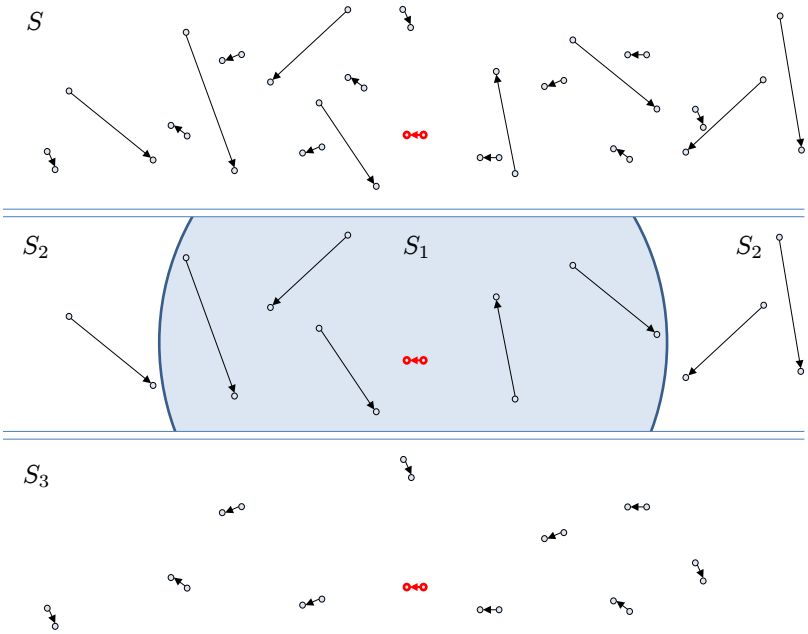

FIGURE 4. The top part of the figure displays a set S of links that is (assumed to be) feasible under some power assignment. Let l_v in the proof of Theorem 3 be the red bold link. The center part displays long links of S partitioned (in this particular example) into S_1 and S_2 using a disc around l_v, the lower part displays S_3 consisting of short links.

5.2. **Proof of Theorem 4.** The following lemma is the crucial element in the proof.

Lemma 15. *Let L be a $\mathcal{P}_p$-feasible set of non-weak links and l_v be a link (not necessarily in L). Then, $\hat{a}_v(L) = \mathcal{O}(1)$.*

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 $\hat{a}_v(L) \le n$ for any set L of size n and any l_v. We claim that $g(n)$ is indeed $\mathcal{O}(1)$, 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 all of the following hold:

(a) There exists $L \in \mathcal{L}(n_0)$ and l_v such that:

$$\hat{a}_v(L) \ge \frac{1}{2} g(n_0) \ . \tag{10}$$

Observe that such an L and l_v always exist independent of n_0, by the definition of g.

(b) Define $f(n) := \frac{1}{2} 2^{\frac{1}{4c_3} g(n)}$, where c_3 is a fixed constant to be specified later. Then,

$$f(n_0) \ge (8 \cdot 3^\alpha)^{1/(p\alpha)} \ . \tag{11}$$

(c) Lastly,

$$g(n_0) \ge 16 \cdot (4^\alpha + 1) \tag{12}$$

We prove our lemma by deriving a contradiction to Eqn. 10. To prove this, we partition the link set L into L_1 and L_2 where $L_1 := \{l_w : \ell_w \le f(n_0) \cdot \ell_v\}$ and $L_2 := L \setminus L_1$.

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

Proof. By definition of $\hat{a}$, we can ignore links in L_1 smaller than l_v. Since the maximum length in L_1 is at most $f(n_0) \cdot \ell_v$, the remaining links in L_1 can be divided into $\log f(n_0)$ length classes. Consider any such length class C. By Lemma 1, C can be partitioned into $\frac{2^{\alpha+1}}{\beta} + 1$ sets that are feasible and 2-independent. For any such set C', we can invoke Lemma 14 to show that $a_v(C') = \mathcal{O}(1)$ and thus $a_v(C) = (2^{1+\alpha}/\beta + 1)\mathcal{O}(1) = \mathcal{O}(1)$. By setting c_3 to be this constant, we get that

$$\hat{a}_v(L_1) \le c_3 \log f(n_0) = c_3 \left(\frac{1}{4c_3} g(n_0) - 1 \right) < \frac{1}{4} g(n_0) \ ,$$

where we used the definition of $f(n)$ in the equality. □

Claim 2. $\hat{a}_v(L_2) \le \frac{1}{4} g(n_0)$,

Proof. Consider $l_w \in L_2$ such that $D := d(s_v, s_w)$ is minimized. Let L_3 be the set of links in L_2 with receivers within the ball $B(s_v, D/2)$ 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 B):

Proposition 16. $|L_3| \le 2 \cdot 4^\alpha + 1$.

Using this proposition, we get that

$$\hat{a}_v(L_3 \cup \{l_w\}) \le |L_3| + 1 \le 2 \cdot (4^\alpha + 1) \le \frac{1}{8} g(n_0) \ ,$$

[2] where the last inequality follows from Eqn. 12.

Now 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,

$$a_v(L_4 \setminus \{\ell_w\}) \le \sum_{\ell_u \in L_4 \setminus \{\ell_w\}} c_u \cdot \frac{P_v}{d(s_v, r_u)^\alpha} \frac{\ell_u^\alpha}{P_u} \le 3^\alpha \sum_u \frac{P_v}{P_w} \frac{P_w}{d(s_w, r_u)^\alpha} \frac{\ell_u^\alpha}{P_u} = 3^\alpha \frac{P_v}{P_w} a_w(L_4) \ . \tag{13}$$

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 $L_2 \supseteq L_4$,

$$P_w \ge \mathcal{P}_p(f(n_0) \cdot \ell_v) = f(n_0)^{p\alpha} P_v.$$

[2] SH: Why can we assume $a_v(w) \le 1$? There is no restriction on l_v nor its location. MMH: Definition of affectance.

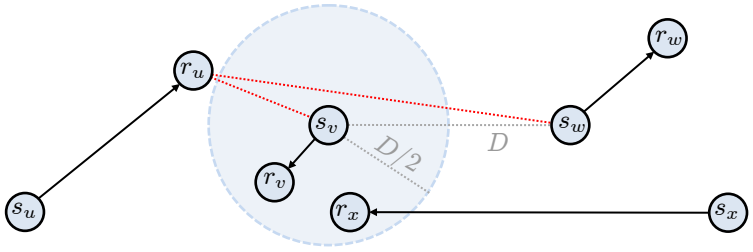

FIGURE 5. Nodes s_v, s_w and r_u play the role described in the proof. Here, $L_2 := \{l_u, l_w, l_x\}, L_3 := \{l_x\}$ and $L_4 := \{l_u, l_w\}$. The red dotted lines indicate the relevant distances in the triangle inequality yielding $d(s_v, r_u) \geq \frac{1}{3} d(s_w, r_u)$.

Thus, $\frac{P_v}{P_w} \leq \frac{1}{f(n_0)^{p\alpha}} \leq \frac{1}{16 \cdot 3^\alpha \beta}$ using Eqn. 11. Combining this insight with Inequality 13 and using that $a_w(L_4) \leq g(n_0)$ due to the definition of $g(n)$, we conclude that

$$\hat{a}_v(L_4 \setminus \{l_w\}) \leq 3^\alpha \frac{1}{8 \cdot 3^\alpha} g(n_0) = \frac{1}{8} g(n_0) \ .$$

[3] This completes the proof of Claim 2. □

Combining Claims 1 and 2, we get that $\hat{a}_v(L) < \frac{1}{2} g(n_0)$, contradicting Eqn. 10. This completes the proof of Lemma 15. □

We can now complete the proof of Theorem 4, which we recall states that $I_p^p(L) = O(1)$, for any $p \in (0, 1)$ and any set L of links. TODO: MMH: Restate the theorem?

Proof. Consider any $S \in F_{\mathcal{P}_p}(L)$ and any $l_v \in L$. Starting with the definition of $\hat{b}$,

$$\hat{b}_v(S) = \hat{a}_v(S) + \hat{a}_S(v) = O(1) + O(1) = O(1) \ ,$$

where we apply Lemma 15 on the first term and Lemma 7 of [33] on the second term. □

We remark that the bound in neither theorem remains true when there are weak links.

6. 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 thus gives the (poly)-logarithmic improvements for a variety of applications. Here we often omit proofs of our claims, as they are all of the same flavor.

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 [24]. 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 $\mathcal{O}(\log n \cdot (\log n + \log\log \Delta))$ slots using mean power in both centralized [24] and distributed [23] algorithms. These results are directly improved by Theorem 6:

Theorem 17. *Any set of links can be strongly connected in $\mathcal{O}(\log n \cdot \log\log \Delta)$ slots using power assignment $\mathcal{P}_p$. This can be computed by either a poly-time centralized algorithm or an $\mathcal{O}(poly(\log n) \log \Delta)$-time distributed algorithm.*

[3] MMH: What was the meaning of this statement here: "Therefore the assumption $g(n) = \omega(1)$ is wrong"?

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 [24] for a discussion of these problems and their numerous applications.

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, [3]), the problem of dynamic allocation of channels to bidders (these are the wireless devices) via an auction has attracted much attention [46, 47].

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 [27, 26] has established a number of results depending on different valuation functions. Since these results are based on the inductive independence number, Theorem 4 improves virtually all of them by a $\log n$ factor. For instance, an algorithm was given in [27] 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 4.

Corollary 18. *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* [27]*, a $\mathcal{O}(\log n + \log k)$-approximation for symmetric valuations and an $\mathcal{O}(\log n)$-approximation for Rank-matroid valuations* [26]*.*

Dynamic Packet Scheduling. Dynamic packet scheduling to achieve network *stability* is one of the fundamental problems in (wireless) network queuing theory [42]. In spite of its long history, this fundamental problem has been considered only recently in the SINR model (see [35, 32, 2]). 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* between 0 and 1 is used to capture how well a given algorithm performs 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 [2] and [32]. In spite of differences in the algorithm and assumptions made, both are based on the scheduling algorithm of [33] and achieve a similar result. Recall that the maximum average affectance is $A^p(L) = \max_{S \subseteq L} \frac{a_S^p(S)}{|S|}$ 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 [32, 2] can be succinctly expressed as follows.

Theorem 19. [32, 2] *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 $\mathcal{O}(\log n)$ [33], both papers claimed $\Omega(\frac{1}{\log^2 n})$-efficiency. Results in this paper show that $\phi(L) = \mathcal{O}(1)$ (see second part of Corollary 9), which gives the following improved result:

Corollary 20. *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 9 also shows that $\overline{\phi(L)} = \frac{A^p(L)}{\overline{\chi(L)}} = \mathcal{O}(\log n \cdot \log\log \Delta)$, we also get the following improved bound for power control:

Corollary 21. *There is a distributed algorithm with $\Omega\left(\frac{1}{\log n \cdot \log\log \Delta}\right)$-efficiency, with respect to power control optima.*

Acknowledgments: We would like to thank Marijke Bodlaender for helpful comments and for pointing out errors in the conference version [20].

References

[1] M. Andrews and M. Dinitz. Maximizing capacity in arbitrary wireless networks in the SINR model: Complexity and game theory. In *Proc. 28th IEEE Conference on Computer Communications (INFOCOM)*, pages 1332–1340. IEEE, 2009.

[2] E. I. Ásgeirsson, M. M. Halldórsson, and P. Mitra. Wireless network stability in the SINR model. In *Proc. 19th International Colloquium in Structural Information and Communication Complexity (SIROCCO)*, pages 123–134, 2012.

[3] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh. White space networking with Wi-Fi like connectivity. In *Proc. of the ACM SIGCOMM 2009 conference on Data communication*, pages 27–38, 2009.

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

[5] D. Chafekar, V. Kumar, M. Marathe, S. Parthasarathy, and A. Srinivasan. Cross-layer latency minimization for wireless networks using SINR constraints. In *Proc. of the 8th ACM international symposium on Mobile ad hoc networking and computing (MobiHoc)*, pages 110–119, 2007.

[6] J. Dams, M. Hoefer, and T. Kesselheim. Scheduling in wireless networks with rayleigh-fading interference. In *Proc. of the 24th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 327–335, 2012.

[7] M. Dinitz. Distributed algorithms for approximating wireless network capacity. In *Proc. 29th IEEE Conference on Computer Communications (INFOCOM)*, pages 1–9, 2010.

[8] A. Fanghänel, T. Kesselheim, H. Räcke, and B. Vöcking. Oblivious interference scheduling. In *Proc. 28th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC)*, pages 220–229, 2009.

[9] A. Fanghänel, T. Kesselheim, and B. Vöcking. Improved algorithms for latency minimization in wireless networks. In *Proceedings of the 36th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 447–458, 2009.

[10] Y. Gao, J. C. Hou, and H. Nguyen. Topology control for maintaining network connectivity and maximizing network capacity under the physical model. In *Proc. 27th IEEE Conference on Computer Communications (INFOCOM)*, pages 1013–1021, 2008.

[11] O. Goussevskaia, M. M. Halldórsson, R. Wattenhofer, and E. Welzl. Capacity of Arbitrary Wireless Networks. In *Proc. 28th IEEE Conference on Computer Communications (INFOCOM)*, pages 1872–1880, April 2009.

[12] 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)*, pages 35–44, 2008.

[13] O. Goussevskaia, Y. A. Oswald, and R. Wattenhofer. Complexity in geometric SINR. In *Proc. of the 8th ACM international symposium on Mobile ad hoc networking and computing (MobiHoc)*, pages 100–109, 2007.

[14] 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.

[15] O. Goussevskaia and R. Wattenhofer. Complexity of scheduling with analog network coding. In *Proc. of the 1st ACM international workshop on Foundations of wireless ad hoc and sensor networking and computing (FOWANC)*, pages 77–84, 2008.

[16] J. Grönkvist and A. Hansson. Comparison Between Graph-Based and Interference-Based STDMA Scheduling. In *Proc. of the 2nd ACM international symposium on Mobile ad hoc networking and computing*, pages 255–258, 2001.

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

[18] M. M. Halldórsson. Wireless scheduling with power control. arXiv:1010.3427v2. To appear in *ACM Transactions on Algorithms*. Earlier version appeared in Proc. of the 17th Annual European Symposium on Algorithms (ESA) 2009.

[19] M. M. Halldórsson. Approximations of weighted independent set and hereditary subset problems. *Journal of Graph Algorithms and Applications*, 4(1):1–16, 2000.

[20] M. M. Halldórsson, S. Holzer, P. Mitra, and R. Wattenhofer. The power of non-uniform wireless power. In *SODA*, 2013.

[21] M. M. Halldórsson and P. Mitra. Nearly optimal bounds for distributed wireless scheduling in the SINR model. In *Proceedings of the 38th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 625–636, 2011.

[22] M. M. Halldórsson and P. Mitra. Wireless Capacity with Oblivious Power in General Metrics. In *Proc. of the 22nd annual ACM-SIAM symposium on Discrete algorithms (SODA)*, pages 1538–1548, 2011.

[23] M. M. Halldórsson and P. Mitra. Distributed connectivity of wireless networks. In *Proc. 31st Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC)*, pages 205–214, 2012.

[24] M. M. Halldórsson and P. Mitra. Wireless connectivity and capacity. In *Proceedings of the 23rd annual ACM-SIAM symposium on Discrete algorithms (SODA)*, pages 516–526, 2012.

[25] M. M. Halldórsson and R. Wattenhofer. Wireless Communication is in APX. In *Proceedings of the 36th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 525–536, 2009.

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

[27] M. Hoefer, T. Kesselheim, and B. Vöcking. Approximation algorithms for secondary spectrum auctions. In *Proc. of the 23rd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 177–186, 2011.

[28] 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 ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, pages 282–290, 2006.

[29] E. Kantor, Z. Lotker, M. Parter, and D. Peleg. The Topology of Wireless Communication. In *Proceedings of the 43rd annual ACM symposium on Theory of computing (STOC)*, pages 383–392, 2011.

[30] T. Kesselheim. A Constant-Factor Approximation for Wireless Capacity Maximization with Power Control in the SINR Model. In *Proc. of the 22nd annual ACM-SIAM symposium on Discrete algorithms (SODA)*, pages 1549–1559, 2011.

[31] T. Kesselheim. Approximation algorithms for wireless link scheduling with flexible data rates. In *Proc. of the 20th Annual European Symposium on Algorithms (ESA)*, pages 659–670, 2012.

[32] T. Kesselheim. Dynamic packet scheduling in wireless networks. In *Proc. 31st Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC)*, pages 281–290, 2012.

[33] T. Kesselheim and B. Vöcking. Distributed contention resolution in wireless networks. In *Proc. 24th International Symposium on Distributed Computing (DISC)*, pages 163–178, 2010.

[34] V. Kumar, M. Marathe, S. Parthasarathy, and A. Srinivasan. Algorithmic Aspects of Capacity in Wireless Networks. *ACM SIGMETRICS Performance Evaluation Review*, 33(1):133–144, 2005.

[35] L. B. Le, E. Modiano, C. Joo, and N. B. Shroff. Longest-queue-first scheduling under SINR interference model. In *Proc. of the 11th ACM international symposium on Mobile ad hoc networking and computing (MobiHoc)*, pages 41–50, 2010.

[36] 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 ACM conference on Embedded network sensor systems (SenSys)*, pages 141–154, 2008.

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

[38] T. Moscibroda and R. Wattenhofer. The Complexity of Connectivity in Wireless Networks. In *Proc. 25th IEEE Conference on Computer Communications (INFOCOM)*, pages 1–13, 2006.

[39] T. Moscibroda, R. Wattenhofer, and Y. Weber. Protocol design beyond graph-based models. In *Proc. of the 5th ACM Workshop on Hot Topics in Networks (HotNets)*, pages 25–30, 2006.

[40] T. Moscibroda, R. Wattenhofer, and A. Zollinger. Topology control meets SINR: The scheduling complexity of arbitrary topologies. In *Proc. of the 7th ACM international symposium on Mobile ad hoc networking and computing (MobiHoc)*, pages 310–321, 2006.

[41] 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 ACM international symposium on Mobile ad hoc networking and computing (MobiHoc)*, pages 91–100, 2008.

[42] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, 37(12):1936–1948, 1992.

[43] T. Tonoyan. On the capacity of oblivious powers. In *Proc. of the 7th International Symposium on Algorithms for Sensor Systems, Wireless Ad Hoc Networks and Autonomous Mobile Entities (ALGOSENSORS)*, pages 225–237, 2011.

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

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

[46] X. Zhou, S. Gandhi, S. Suri, and H. Zheng. eBay in the Sky: Strategy-proof wireless spectrum auctions. In *Proc. of the 14th ACM international conference on Mobile computing and networking (MobiCom)*, pages 2–13, 2008.

[47] X. Zhou and H. Zheng. Trust: A general framework for truthful double spectrum auctions. In *Proc. 28th IEEE Conference on Computer Communications (INFOCOM)*, pages 999–1007, 2009.

Appendix A. Missing Proof from Section 2: Independence Strengthening

Lemma 1 *Any feasible set of links can be partitioned into* $\lfloor \frac{2q^\alpha}{\beta} \rfloor + 1$ *or fewer* 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,

$$b_S^{\mathcal{P}}(S)/2 = a_S^{\mathcal{P}}(S) \leq |S|,$$

implying that 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 it is true 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. □

Appendix B. Missing proofs from Section 5

Proposition 13 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_{vu} minimum. Then $d_{vw} \geq d_{wu}/6$.

The reader might find it useful to use Figure 6 while reading the proof of Proposition 13.

Proof. Let $D = d_{vw}$ and note that by definition $d_{vu} \leq D$. By the triangular inequality and the definition of l_u,

$$\begin{aligned} d_{uw} &\leq d_{uu} + d_{vu} + d_{vw} \\ &= \ell_u + d_{vu} + d_{vw} \leq 2D + \ell_u \ . \end{aligned} \tag{14}$$

Similarly,

$$d_{wu} \leq d_{vu} + d_{vw} + \ell_w \leq 2D + \ell_w \ . \tag{15}$$

Applying 2-independence, on one hand, and multiplying Eqn. 15 and Eqn. 14, on the other hand, we have that

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

This implies that D must be at least $\min(\ell_u, \ell_w)/2$, which in turn is at least $\max(\ell_u, \ell_w)/4$, since the links are nearly-equilength. Thus we can bound $\ell_w \leq 4D$ in Eqn. 15 to obtain that $d_{wu} \leq 6D$. □

4

Definition 22. *We say that links l_v and l_w are* t-close *under power assignment $\mathcal{P}$ if,*

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

[4] MMH: Add an intuition here on the idea behind the lemma below?

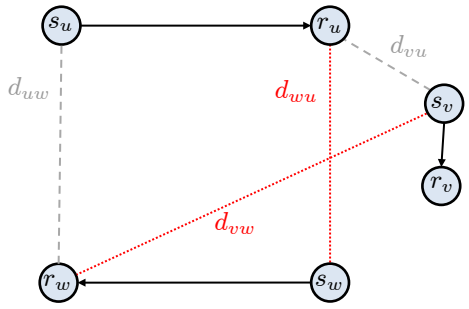

FIGURE 6. Links l_u, l_w and l_w as used in the proofs of Propositions 12 and 13. The distances d_{vw} and d_{wu} that are related to each other in the Proposition's statement are represented by red dotted lines. The gray dashed lines mark distances d_{uw} and d_{vu} that are used in the proofs as well.

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

Lemma 11 *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, $|Q| = \mathcal{O}(\log\log\Delta)$.*

Proof. The set Q consists of links that have at least one of the following properties:

(1) a link can affect l_v by at least $\frac{1}{\tau}$ under power assignment $\mathcal{P}_p$
(2) the link itself can be affected by l_v by that amount

We consider first the links with the first property. 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

$$\begin{aligned} d_{w'w} &\leq d(s_{w'}, r_v) + d(r_v, s_w) + d(s_w, r_w) \\ &= d_{w'v} + d_{wv} + \ell_w \\ &\leq 2\ell_w^p \ell_v^{1-p} (c_v\tau)^{1/\alpha} + \ell_w \\ &\leq 2\ell_w^p \ell_v^{1-p} (2\beta\tau)^{1/\alpha} + \ell_w \\ &\leq \ell_w^p \ell_1^{1-p} + \ell_w \leq 2\ell_w , \end{aligned} \tag{16}$$

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

$$d_{ww'} \leq \ell_{w'} + \frac{1}{2}\ell_w^p \ell_1^{1-p} . \tag{17}$$

Applying 2-independence, on one hand, and multiplying Eqn. 16 and 17, on the other, we obtain that

$$4\ell_w\ell_{w'} \leq d_{w'w} \cdot d_{ww'} \leq 2\ell_{w'}\ell_w + \ell_w^p \ell_1^{1-p} \cdot \ell_w , \tag{18}$$

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

$$\ell_w^p \geq \frac{2\ell_{w'}}{\ell_1^{1-p}} . \tag{19}$$

We label the links in Q as $l_1, l_2, \ldots, l_{|Q|}$ in increasing order of length, and define $\lambda_i = \ell_i/\ell_1$. By dividing both sides of Eqn. 19 by ℓ_1^p, we get that

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

Then, $\lambda_2 \geq 2^{1/p}$ and by induction $\lambda_t \geq 2^{(1/p)^{t-1}}$. Note that

$$\Delta(Q) = \ell_{|Q|}/\ell_1 = \lambda_{|Q|} \geq 2^{(1/p)^{|Q|-1}},$$

so

$$|Q| - 1 \leq \log_{1/p} \log_2 \Delta,$$

and the claim follows.

The other case of links l_w with $a_v(w) \geq 1/\tau$ is symmetric, with the roles of p and $1-p$ switched, leading to a bound of $1 + \log_{1/(1-p)} \log_2 \Delta$.

The case that a link can have both properties does not affect the asymptotic statement of the result. □

Proposition 16 $|L_3| \leq 2 \cdot 4^\alpha + 1$.

Proof. By Lemma 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 the 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. □