Link Scheduling under Correlated Shadowing

Magnús M. Halldórsson
ICE-TCS, Department of Computer Science
Reykjavik University, Iceland
mmh@ru.is

Tigran Tonoyan
ICE-TCS, Department of Computer Science
Reykjavik University, Iceland
ttonoyan@gmail.com

***Abstract*—We study the effects of stochastic shadowing on scheduling in wireless networks. Previous work has generally assumed that shadowing affects different signals independently. Concentrating on "compact" networks, where very little or no spatial reuse is possible, such as in indoor environments, we form a model of correlation between the shadowing components for different signals. We analyze two fundamental measures: the maximum number of simultaneous transmitting links, and the fewest time/frequency slots or channels needed to schedule all the links. Based on the correlation model, we characterize (up to constant factors) how these measures scale with correlation strength and the number of links. We also give nearly optimal algorithms to compute such schedules, as well as to optimize the maximum weighted sum of simultaneously transmitting links. The latter can as well be extended to arbitrary sets of similar length links, under a suitable model of correlations, by partitioning them into compact subsets.**

***Index Terms*—link capacity, scheduling, lognormal shadowing, correlated shadowing**

Much effort has been spent on understanding the capacity of wireless networks and how to best utilize them. A major challenge has been how to model the wireless environments realistically while maintaining analytical tractability. The model of choice for algorithmic study of general ad hoc networks has been the *physical model* with geometric pathloss. By the geometric nature of the model, the main tool for achieving efficient channel utilization is *spatial reuse*, namely packing as many concurrent transmissions in space as possible, then using time or spectrum division multiplexing. What happens when there is little possibility of spatial reuse, such as in indoor communications? The geometric pathloss model offers in this case only the trivial time/spectrum division. However, the physical model is only the averaged view of what happens in actual networks. In particular, the received signal is almost never symmetric with respect to different directions around the transmitter, due to channel and antenna irregularities, or *shadowing*.

Shadowing is most frequently modeled stochastically. To each pair of points in space, we associate a random variable drawn from a distribution. There is a general agreement that *Lognormal shadowing* (LNS), which is Gaussian on the dBm scale, is the most faithful approximation known of true shadowing [1]–[4]. Though probabilistic models are known to be far from perfect, they are generally understood to be highly useful for providing insight into wireless systems, and certainly more so than using pathloss alone.

The aim of this paper is to go beyond spatial reuse and analyze possible capacity gains due to shadowing effects, using the physical model as a baseline. We tackle two representative fundamental problems, *weighted single-shot scheduling* and *minimum length scheduling*, where in the former, we seek a maximum weight *feasible* subset of a weighted set of links that can successfully transmit simultaneously, while in the latter, the set of links must be partitioned into the minimum number of feasible subsets.

It has been frequently observed in the literature (through simulations or analytically) that throughput and connectivity improve with shadowing. However, such effects have also been attributed to the assumed independence of shadowing across different links. The amount of correlation depends on the network structure, with the highest correlation between links in nearly the same situation.

As a case study, we consider *compact* networks under *correlated* Lognormal shadowing, where the links have roughly similar lengths, and are located in an area comparable to their lengths, such as indoor networks. Clearly, such networks offer very limited spatial reuse, and can be used to best demonstrate possible capacity gains due to shadowing. Also, the problems we consider are already quite challenging even in this simple case. There are many interesting questions that arise when faced with this situation:

- Can shadowing cause non-trivial capacity/scheduling gains?
- If so, how does it scale with the correlation and the size of the network?
- Are there simple and/or natural expressions for the capacity of compact networks?
- Which is the bigger factor in capacity gains: variations in signal strengths of links, or variations in interferences between links?
- Are there efficient algorithms for computing near-optimal throughput schedules?

In this paper, we address all these questions for compact networks. Such an analytic study faces two major challenges. On one hand, the Lognormal distribution is heavy-tailed, meaning that large values are not that unlikely. This precludes the direct use of concentration bounds, in a major departure from exponential distributions, such as the case of Rayleigh fading. The other challenge is how to capture and handle non-independence of stochastic events.

a) Our results: Below, we present a high level overview of our results and the techniques used.

Algorithm for Weighted Single-Shot Scheduling. Our first result (Sec. II) is an algorithmic characterization, up to constant factors, of the expected maximum weight of feasible subsets in a compact network. More concretely, given a compact set with shadowing sampled from correlated Lognormal distribution, we select a feasible subset of links whose weight is only a constant factor away from the expected optimum. In previous work [5], we have shown that under *independent* shadowing, nearly maximal size *unweighted* feasible sets of links can be computed by focusing on *strong* sets of links: sets of t links whose signal strength is about t times larger than the expectation. The heavy-tail property of LNS will give us strong sets of significant size, while the interference within such sets, which is composed of many i.i.d. terms (pairwise interferences), will be concentrated near its mean. Moreover, non-strong sets of links are much less likely to be feasible, leading to the near-optimality of strong sets.

A somewhat different approach is needed in the presence of correlated shadowing. There, roughly speaking, there is a *principal component* of shadowing shared by all signal and interference distributions, which varies simultaneously for all of them. In particular, the interferences are not as well concentrated near their expectation as before, since their values depend on that single component. Therefore, it may not suffice to select the largest strong set of links. If the largest strong set is very small, the reason for that could be the small value of the principal component, which in turn implies smaller value for all interferences, so one could potentially obtain a better solution by selecting sets of weaker links that also happen to have small mutual interference. Thus, our approach is to seek large subsets of links that become strong when both signal strengths and interferences are *normalized* by the value of the principal component, trying all such values.

The main challenge remaining is then to show that those are the only sets we need to consider, in order to get a (expected) constant factor approximation. To that end, we show that achieving feasibility with only (normalized) sets of strong links is much more likely than using even a single non-strong link, given the strong concentration bounds that hold for the aggregate interference received by a link. This fails only when all optimal sets are small, which can occur if the weights are highly unbalanced; for that case, we can apply exhaustive search among the relatively small number of high weight links.

The constant factor approximation can also be extended to the more general setting of sets of nearly equal length links, arbitrarily located on the plane. This is achieved by showing that such a set can be partitioned into a constant number of subsets consisting of well-separated compact subsets. For this, we assume that shadowing effects within well-separated compact subsets are uncorrelated. We leave the consideration of more detailed or accurate models for correlation between disjoint compact subsets to future work.

Scaling Law for Max Size Feasible Set. In our second result (Sec. III), we use the ideas developed above to formulate a *scaling law* for the expected maximum size of a feasible set in a compact network of k links, which shows the dependence of the expected size of a maximum feasible set on the number of links in the network and the correlation coefficient. This requires to estimate the maximum expected size of a normalized strong subset of links. In particular, we show that the expected size is $\exp(c\sqrt{\ln k} \pm \tau)$, where c is a constant (depending on the correlation), while τ is a lower order term.

Min Length Scheduling. In our third result, we characterize the expected minimum number of slots required to schedule a compact set of k links and show, perhaps surprisingly, that considerable scheduling gains can be achieved, even under correlated shadowing (Sec. IV). The difference from one-shot scheduling is that while the main improvement in one-slot schedules is due to very strong links, the min. length scheduling gain is achieved by showing that moderately strong links can be grouped into low-interference subsets, i.e., that there are *many* feasible subsets of significant size, giving us the throughput gain. We show that this number is a constant factor of $k(\ln\ln k)^2/\ln k$, which gives us a throughput gain by a factor of $\ln k/(\ln\ln k)^2$, compared with the bound k in the case without shadowing.

The argument behind this result is based on partitioning the set of links into normalized strong subsets. However, the normalization factor here is not only the principal component, as using just the latter for normalization would yields few feasible subsets of relatively big size, which is not enough to prove our claim. Instead, we look for smaller but still significant feasible subsets that are strong, when normalized by a larger factor. The heavy-tail shadowing distribution shows that there will be many such subsets with significantly lower mutual interference than expected. In order to turn these ideas into the actual result we aim for, we reduce the problem to that of coloring Erdős-Renyi random graphs.

To our knowledge, this is the first work to study the weighted single-shot scheduling and min. length scheduling problems in this setting.

Due to space constraints, some technical details are deferred to the full version of the paper.

b) Related Work: Gupta and Kumar [6] introduced the physical model, in order to study scaling laws regarding throughput capacity in networks. This simplified model gained wide attention in analytic and algorithmic studies of wireless networks. Experimental studies support the physical model over disk-graph-based models, which over-simplify decay behavior of signals [2], [7]–[10]. This model has been accurate and tractable enough to spawn a considerable body of work studying algorithmic problems like (weighted) single-shot and min. length scheduling, among others [11]–[20]. First algorithms with performance guarantees in the physical model were given by Moscibroda and Wattenhofer [21]. These scheduling problems have been fundamental to various MAC layer problems, most notably for TDMA scheduling and maximum thoughput scheduling (e.g. [22], [23]).

Validity of the Lognormal model for shadowing has been confirmed through extensive experiments [1], [2], [24], [25]. There have been numerical and analytical studies showing that log-normal shadowing results in better connectivity [26]–[29] and throughput capacity (under particular scheduling strategies) [30]. However it has been demonstrated that the connectivity gains are mostly due to the independent shadowing assumption, and they decrease dramatically when correlations are introduced [25]. The *unweighted* single-shot scheduling under *independent* shadowing has been considered in a previous work [5]. However, we are not aware of any analytic study quantifying the dependence of the gains on correlation in terms of explicit expressions.

Even though the equal correlation assumption is natural (and has been observed experimentally [1]) for the case of compact networks, it is not the only model. In particular, it may not be applicable for "sparser" networks spread over a large area. Many different correlation models have been studied in the literature. For example, it has been suggested to model correlation as exponentially decreasing with distance [25]. See e.g. [31] for a list of models.

I. Models and Formulations

a) Communication Model and Basic Problems: The main object of our consideration is a set L of *communication links*, numbered from 1 to $k = |L|$. Each link $i \in L$ represents a unit-demand communication request between a sender node s_i and a receiver node r_i, both point-size wireless nodes located on the plane. We assume the links all operate in the same channel, and all (sender) nodes use the same transmission power level P. We refer to a set of links that can successfully communicate in a single time slot as *feasible*. Before formally defining the feasibility model, let us define the main problems we are interested in.

In the *weighted single-shot scheduling* problem, given a set L of links with positive weights, the goal is to select the maximum weight feasible subset of links, where the weight of a subset is the sum of the individual link weights. We refer to the optimum weight by $opt_W(L)$. Of particular interest is the special case when all links have equal weights, i.e., the goal is to find a *maximum cardinality feasible subset*. The optimum cardinality is denoted by $opt_C(L)$. In the *minimum length scheduling* problem, the goal is to partition L into the minimum number of feasible subsets. We call this number the *scheduling number*, denoted $opt_S(L)$.

When S is the subset of links transmitting simultaneously, a given link $i \in S$ succeeds if its signal strength (the power of the transmission of s_i when measured at r_i) is greater than β times the total (sum) interference from other transmissions, where $\beta > 0$ is a threshold parameter, and the interference of link j on link i is the power of transmission of s_j when measured at r_i. We consider interference-limited networks, where the effect of ambient noise is negligible. Formally, link i transmits successfully if

$$SIR(S,i) = \frac{\mathcal{S}_i}{\sum_{j \in S \setminus i} \mathcal{I}_{ji}} > \beta \ , \qquad (1)$$

where $\mathcal{S}_i$ is the received signal strength/power of link i and $\mathcal{I}_{ji}$ is the interference of link j on link i. A link i is *feasible in a set* S if $SIR(S,i) > \beta$. A set S is *feasible* if every link $i \in S$ is feasible in S.

We assume for simplicity that $\beta = 1$. The latter is justified by the following result of [32]: we can state our results for general $\beta \geq 1$ by paying at most a factor of 2β in performance.

Lemma 1. *If a set S of links and number $\beta' > 0$ are such that $SIR(S,i) \geq \beta'$ for each link $i \in S$, then S can be partitioned into at most $\lceil 2\beta/\beta' \rceil$ subsets, each feasible with threshold β.*

b) Geometric Path-Loss: The *Geometric Path-Loss model* (GPL) defines the received signal strength between nodes u and v as $P/f(d(u,v))$, where P is the power used by the sender u, d denotes the Euclidean distance, and f is a deterministic function (e.g., $f(x) = x^\alpha$, in the case of log-distance pathloss). In particular, the signal strength of a link i and the interference of a link j on link i are, respectively,

$$\bar{\mathcal{S}}_i = \frac{P}{f(l_i)} \text{ and } \bar{\mathcal{I}}_{ji} = \frac{P}{f(d(s_j, r_i))},$$

where $l_i = d(s_i, r_i)$ denotes the *length* of link i and $d(s_j, r_i)$ is the distance from the sender node of link j to the receiver node of link i. If the links in a set S transmit simultaneously, the formula determining the success of the transmission on link i is similar to (1).

c) Shadowing: One of the effects that GPL ignores (or models only by an appropriate change of the path loss exponent α), is signal obstruction by objects, or *shadowing*. In generic networks shadowing is often modeled by a *stochastic shadowing model*. Here we adopt the *Lognormal Shadowing model*, or LNS for short. In this case, the signal strength $\mathcal{S}_i$ of a link i at r_i is assumed to have been sampled from a Lognormal distribution with mean $\mathbb{E}[\mathcal{S}_i] = \bar{\mathcal{S}}_i$, and similarly, for any two links i, j, the interference $\mathcal{I}_{ij}$ is sampled from a Lognormal distribution with mean $\mathbb{E}[\mathcal{I}_{ij}] = \bar{\mathcal{I}}_{ij}$. More concretely, there are Normal random variables (r.v.s) $Z_i \sim \mathcal{N}(\mu_i, \sigma)$ and $Z_{ij} \sim \mathcal{N}(\mu_{ij}, \sigma)$ with a fixed $\sigma > 0$, such that $\mathcal{S}_i = e^{Z_i}$ and $\mathcal{I}_{ij} = e^{Z_{ij}}$, for all links i, j. We assume that shadowing does not change during the time period under consideration. In this model too, signal reception is characterized by the signal to interference ratio.

d) Network structure and shadowing correlation: We will primarily be concerned with *compact sets of links*, where the links have lengths in the range $[\ell, 2\ell]$ and are contained in a box of side 4ℓ, for some $\ell > 0$. The main motivation is that in the geometric path-loss model without shadowing, only a few links (a constant number) can be selected from a compact set to transmit in the same time slot, and similarly, in order to schedule a compact set in the geometric path-loss model, one has to use a number of slots that is linear in the size of the set. Thus, the main question is whether one can hope for better under shadowing.

In order to study compact sets, we can partition them into a constant number of *clusters*, where the distance between the sender and receiver of any two links in a cluster is at least $c\ell$, for some constant $0 < c < 1$. Thus, we eliminate the pairs of links which have too high expected interference on each other. In order to keep things simple, we will focus on "ideal" clusters, where we assume that the senders are all located at one point and the receivers are all located at another point, and hence all links have the same length ℓ. This simplification will only affect the constant factors in our results, due to Lemma 1, but will significantly simplify the exposition. Thus, all clusters considered henceforth are assumed to be such.

Within a cluster, since links have approximately similar position, we assume that they are pairwise equally correlated. It is natural to assume that the correlation is such that the shadowing affects links "similarly", i.e. that the correlation is non-negative. More concretely, if L is a cluster, and if for all links $i, j \in L$, $\mathcal{S}_i = e^{Z_i}$ and $\mathcal{I}_{ji} = e^{Z_{ji}}$, where $Z_i \sim \mathcal{N}(\mu_i, \sigma)$, $Z_{ji} \sim \mathcal{N}(\mu_{ji}, \sigma)$ are Normal r.v.s, then we assume that all r.v.s Z_i, Z_{ij} are *jointly Normal* with a covariance matrix that has σ^2 on its diagonal and $\rho\sigma^2$ elsewhere; namely, every pair of (logarithms of) signals has correlation $0 \le \rho < 1$. Note that for jointly Normal variables, $\rho = 0$ implies independence. We emphasize that ideal clusters are only a technically simpler abstraction of actual instances where links do not share the same place in space, hence we assume general correlation ρ and not $\rho = 1$, which is suitable for ideal instances.

e) Technical Preliminaries: We use the following facts about Lognormal random variable $X = e^Z$ with $Z \sim \mathcal{N}(\mu, \sigma^2)$: The expected value is given by $\mathbb{E}[X] = e^{\mu + \frac{\sigma^2}{2}}$, and for any $x > 0$, $\Pr[X > x] = Q\left(\frac{\ln x - \mu}{\sigma}\right)$, where $Q(x)$ is the tail distribution function of standard Normal distribution, given by $Q(x) = \int_x^\infty \phi(t)dt$, and $\phi(t) = \frac{1}{\sqrt{2\pi}} e^{-t^2/2}$. We will mainly be interested in the following probability:

$$\Pr[X > t \cdot \mathbb{E}[X]] = Q\left(\frac{\ln t}{\sigma} + \frac{\sigma}{2}\right) . \quad (2)$$

II. Algorithm for Single-Shot Scheduling

We refer the reader to the introduction for an informal overview on the ideas behind the algorithm and analysis. Formally, the input to our algorithm is a cluster L of $k = |L|$ links i with positive weight w_i with signal strengths $\mathcal{S}_i$ and interferences $\mathcal{I}_{ji}$ sampled from correlated Lognormal shadowing distribution. Namely, we assume that for each pair i, j of links, $\mathcal{S}_i = e^{Z_i}$, $\mathcal{I}_{ji} = e^{Z_{ji}}$, where $\{Z_i, Z_{ij} : i, j \in L\}$ are jointly Normal with correlation matrix Σ that has σ^2 on its diagonal and $\rho\sigma^2$ off the diagonal, and $0 \le \rho < 1$.

The goal is to find a feasible subset $T \subseteq L$ whose total weight is close to the expected optimum $\mathbb{E}[opt_W(L)]$, with expectation taken over the shadowing distribution.

Let $\bar{\mathcal{S}} = \mathbb{E}[\mathcal{S}_i] = \mathbb{E}[\mathcal{I}_{ij}] = e^{\mu + \sigma^2/2}$ denote the expected signal strength of a link (or interference of one link on another). A link i is *t-strong* if $\mathcal{S}_i > t\bar{\mathcal{S}}$. Our approach is divided into two phases. The first one, given in Alg. 1, tries to find the heaviest set consisting of links of strength above a threshold, for various thresholds. The analysis below is made by conditioning on the component (referred to as *principal component* in the introduction) that is common to all signal strengths and interferences, due to correlation. The outer loop in Alg. 1 essentially guesses this component b, then we focus on sets F that consist of (roughly) $b|F|$-strong links. We show that if b is the right value, then such a set F is likely to contain a large feasible subset. Having this, the main challenge is to show that such sets of strong links are likely to contain a near optimal solution, except for the case when the optimal weight is achieved by only sets of few (bounded by a constant) links. To handle the latter case, we complete our algorithm by doing exhaustive search to find the heaviest feasible subset of at most h links.

Algorithm 1 Find heavy strong subsets

```
γ ← ½ · exp(−(1 − ρ)σ²/2)
for b = 1, 2^{±1}, 4^{±1}, ..., 2^{±⌈log W⌉} do
    for t = 1, 2, ... k do
        F_t^b ← (at most) t heaviest of γbt-strong links in L
        E_t^b ← {i ∈ F_t^b : Σ_{j∈F_t^b} I_ji < 4S_i/γ}
        Partition E_t^b into at most 8/γ feasible subsets, using
        Lemma 1, and let G_t^b be the heaviest one
    end for
end for
return Heaviest of all G_t^b
```

Algorithm 2 Main algorithm

```
1: h ← O(1 + max(ln σ, ln³(1/(σ√(1−ρ))) / (σ²(1−ρ))))
2: H ← the heaviest feasible subset of at most h links
3: return the better among H and G, the output of Alg. 1
```

Remark. A naïve implementation of the algorithm has runtime $O(k^3 \log W + k^h)$. When links are very highly correlated, i.e., when $(\sigma\sqrt{1-\rho})^{-1}$ is very large, the second term (due to step 2 in Alg. 2) may be prohibitive. In that case, one may resort to computing the heaviest among smaller subsets. In the most trivial form, one could simply take H to be the heaviest link in L; this would increase the approximation ratio by a factor of h. Note also that with a more detailed calculation, the values of constants can be tuned/traded to give better bounds.

Recall that $opt_W(L)$ is the optimum weight of a feasible subset. We use $opt'_W(L)$ to denote the maximum weight of a feasible subset of size at least h. The rest of this section is devoted to the proof of the following theorem, which shows that the algorithm attains a performance ratio independent of the network size or link weights. For a weighted set S of links, we denote by w_S the total weight of S.

Theorem 1. *Let A be the output of Alg. 2 on a set L of links, and γ be as in Alg. 1. Then there is a constant $c > 0$, s.t. $\mathbb{E}[opt_W(L)] \le (c/\gamma) \cdot \mathbb{E}[w_A]$, where the expectation is with respect to the shadowing distribution.*

Since we handle the subsets of size at most h exactly (by subset H), it suffices to prove the approximation for larger subsets. Therefore, the subsequent analysis focuses on the approximation of $opt'_W(L)$ by the set G produced by Alg. 1.

In order to realize our plan of separating the common part between correlated signals and interferences, we need to formalize this quantity. We achieve that using Lemma 2 below (proved in the full version), which shows that a cluster with correlated shadowing behaves essentially like a cluster with independent shadowing, but with shifted means and scaled (decreased) variances. It implies that the common part is a Normal r.v., which is perturbed by *independent* Normal r.v.s to give individual signals and interferences.

Lemma 2. *Let $X, Y_1, \ldots, Y_t$ be independent $\mathcal{N}(0, \sigma^2)$ r.v.s. Let $0 \le \rho < 1$ be a parameter. Consider the random variables $Z_i = \sqrt{\rho}X + \sqrt{1-\rho}Y_i$ for $i = 1, 2, \ldots, t$. Then Z_i are jointly Normal with correlation matrix Σ, which has σ^2 on its diagonal and $\sigma^2\rho$ elsewhere.*

Let Z_i, Z_{ij}, for $i, j = 1, 2, \ldots, k$, be Normal r.v.s such that $\mathcal{S}_i = e^{Z_i}$ and $\mathcal{I}_{ij} = e^{Z_{ij}}$. Recall that $\mu = \mathbb{E}[Z_i] = \mathbb{E}[Z_{ij}]$. By the assumption on the correlation between the signals and interferences, Lemma 2 implies that there are independent $\mathcal{N}(0, \sigma^2)$ r.v.s X, Y_i and Y_{ij}, $i, j = 1, 2, \ldots, k$, such that $Z_i = \sqrt{\rho}X + \sqrt{1-\rho}Y_i + \mu$ and $Z_{ij} = \sqrt{\rho}X + \sqrt{1-\rho}Y_{ij} + \mu$ (obvious modifications apply for non-ideal instances, e.g., with different $\mathbb{E}[Z_i] = \mu_i$).

Our strategy will be to condition on the value of X, and show that identical bounds can be obtained for all values with significant probability mass, then argue the bound in expectation. Assume that $X = a$ is fixed, with $a \in [-\log W, \log W]$; the remaining values of X have total probability mass $O(1/W)$ and have negligible effect on the expected approximation ratio. Then Z_i and Z_{ij} become *independent* Normal r.v.s with shifted mean $\mu' = \sqrt{\rho}a + \mu$ and variance $\sigma'^2 = (1-\rho)\sigma^2$. Let $\bar{\mathcal{S}}_a = \mathbb{E}[\mathcal{S}_i \mid X = a] = \mathbb{E}[\mathcal{I}_{ji} \mid X = a]$ denote the expected signal strength, conditioned on $X = a$. Note that $\bar{\mathcal{S}}_a = \exp(\mu' + \sigma'^2/2) = \exp(a\sqrt{\rho} - \rho\sigma^2/2) \cdot \bar{\mathcal{S}}$.

For a given value $X = a$, a subset $T \subseteq L$ is *strong*, if for each link $i \in T$, $\mathcal{S}_i > \gamma|T|\bar{\mathcal{S}}_a$ (i.e., if each link is $\gamma|T|$-strong w.r.t. $\bar{\mathcal{S}}_a$). Let $F^a \subseteq L$ denote the maximum weight strong subset of size at least h.

Thm. 1 follows from Lemmas 3 and 4, where the former shows that Alg. 1 essentially captures the optimum weight achieved by strong subsets, while the latter, which is the most challenging part of the argument, shows that the optimum weight in general is expected to be achieved by a strong subset. All probabilities below are conditioned on $X = a$, which we omit from the notation for clarity.

Lemma 3. *$\mathbb{E}[w_G] \ge (\gamma/16) \cdot \mathbb{E}[w_{F^a}]$, where $G \subseteq L$ is the output of Alg. 1.*

Proof. By definition of F^a, for each link $i \in F^a$,

$$\mathcal{S}_i > \gamma t \bar{\mathcal{S}}_a = \gamma t e^{\sqrt{\rho}a} e^{-\rho\sigma^2/2} \cdot \bar{\mathcal{S}} \ , \tag{3}$$

where $t = |F^a|$. Let b be the largest power of two below $e^{\sqrt{\rho}a}e^{-\rho\sigma^2/2}$. Consider the set F_t^b as defined in Alg. 1 (this value of b is considered in the algorithm, since we assume that $a \in [-\log W, \log W]$). Recall that for each link $i \in F_t^b$,

$$\mathcal{S}_i > \gamma t b \cdot \bar{\mathcal{S}} \ge \gamma t \cdot \bar{\mathcal{S}}_a/2 \ , \tag{4}$$

and F_t^b contains the t heaviest of such links. Hence, $w_{F_t^b} \ge w_{F^a}$. To complete the proof, it suffices to show that the corresponding set G_t^b contains a constant fraction of the weight of F_t^b, in expectation. Taking expectations with respect to interferences and using (4), we have, for each link $i \in F_t^b$,

$$\mathbb{E}\left[\sum_{j \in F_t^b} \mathcal{I}_{ji}\right] = (|F_t^b| - 1) \cdot \bar{\mathcal{S}}_a < \frac{2\mathcal{S}_i}{\gamma} \ .$$

By Markov's inequality, $\Pr\left[\sum_{j \in F_t^b} \mathcal{I}_{ji} < 4\mathcal{S}_i/\gamma\right] > \frac{1}{2}$. The latter implies that $\mathbb{E}[w_{E_t^b}] > \frac{1}{2} \cdot w_{F_t^b}$, where E_t^b is, by definition, the subset of links $i \in F_t^b$ with $\sum_{j \in F_t^b} \mathcal{I}_{ji} < 4\mathcal{S}_i/\gamma$. On the other hand, by Lemma 1, E_t^b can be partitioned into at most $\frac{8}{\gamma}$ feasible subsets, the heaviest of which, namely G_t^b, has weight at least $(\gamma/8) \cdot w_{E_t^b}$, which implies the lemma. ☐

Lemma 4. *For any fixed a, conditioned on $X = a$,*

$$\mathbb{E}[opt'_W(L)] \le 2 \cdot \mathbb{E}[w_{F^a}].$$

Proof. The main ingredient is the following "stochastic dominance" property. For a subset T, let $\mathcal{E}_T^s$ and $\mathcal{E}_T^{nsf}$ denote the events that T *is strong*, and T *is not strong but feasible*, respectively.

Claim 1. *There is a constant $c > 0$, such that for any subset $T \subseteq L$ of size at least $h = c \cdot \left(1 + \max\left(\ln \sigma', \frac{\ln^3(1/\sigma')}{\sigma'^2}\right)\right)$, $\Pr[\mathcal{E}_T^s] \ge \Pr[\mathcal{E}_T^{nsf}]$.*

Proof. Fix a subset T and denote $s = |T|$,

$$\begin{aligned} p_s &= \Pr[\mathcal{S}_i > \gamma s \bar{\mathcal{S}}_a], \\ p_{sf} &= \Pr[\mathcal{S}_i > \gamma s \bar{\mathcal{S}}_a \wedge \sum_{j \in T} \mathcal{I}_{ji} < \mathcal{S}_i], \\ p_{wf} &= \Pr[\mathcal{S}_i \le \gamma s \bar{\mathcal{S}}_a \wedge \sum_{j \in T} \mathcal{I}_{ji} < \mathcal{S}_i]. \end{aligned}$$

By the symmetry of the problem, these probabilities are the same for all links. Then using independence, we have:

$$\Pr[\mathcal{E}_T^s] = \prod_{i \in T} \Pr[\mathcal{S}_i > \gamma s \bar{\mathcal{S}}_a] = p_s^s.$$

On the other hand, the Law of Total Probability and independence imply:

$$\Pr[\mathcal{E}_T^{nsf}] = \sum_{t=1}^{s} \binom{s}{t} p_{wf}^t p_{sf}^{s-t} \le \sum_{t=1}^{s} (s p_{wf})^t p_{sf}^{s-t},$$

where we used the simple bound $\binom{s}{t} \le s^t$. Our aim now is to show that for $s \ge h$,

$$p_s \ge s^2 p_{wf}.$$

Using this inequality and the simple fact that $p_{sf} \le p_s$ in the two bounds above, the claim follows easily.

First, let us bound from above the probability p_{wf} of a link not being strong but being feasible in T. Fix a link $i \in T$. Let I_j denote the indicator r.v. that is 1 iff $\mathcal{I}_{ji} \ge \eta\bar{\mathcal{S}}_a$, for a parameter $\eta > 0$. Then I_j are i.i.d. Bernoulli with parameter

$q = Q\left(\frac{\ln\eta}{\sigma'} + \frac{\sigma'}{2}\right)$, and $\mathbb{E}\left[\sum_{j\in T} I_j\right] = qs$, where q is given by (2). Further note that if $\sum_{j\in T} \mathcal{I}_{ji} < \gamma s\bar{\mathcal{S}}_a$, then at most $(\gamma/\eta)\cdot s$ of I_j can be 1, which implies that

$$p_{wf} \le \Pr\left[\sum_{j\in T} \mathcal{I}_{ji} < \gamma s\bar{\mathcal{S}}_a\right] \le \Pr\left[\sum_{j\in T} I_j < \frac{\gamma}{\eta}\cdot s\right]$$
$$= \Pr\left[\sum_{j\in T} I_j < \frac{\gamma}{\eta q}\cdot qs\right] \le \exp\left(-\frac{sq}{2}\cdot\left(1-\frac{\gamma}{\eta q}\right)^2\right),$$

where the first inequality follows from the definition of p_{wf}, the second from the observation above, while in the last one we used a standard Chernoff bound. We set $\eta = e^{-\sigma'^2/2}$. Then $q = Q(0) = 1/2$, and $\gamma = \frac{1}{2}\cdot e^{-\sigma'^2/2} = \frac{\eta q}{2}$. The bound thus simplifies to: $p_{wf} \le e^{-s/8}$.

The probability p_s of a link being strong is, by (2),

$$p_s = \Pr[\mathcal{S}_i > \gamma s] = Q\left(\frac{\ln(\gamma s)}{\sigma'} + \frac{\sigma'}{2}\right) = Q\left(\frac{\ln(s/2)}{\sigma'}\right) \ ,$$

using $\gamma = \frac{1}{2}\cdot e^{-\sigma'^2/2}$. Now, known bounds on the Q function can be used (see the full version) to show that $p_s \ge s^2 p_{wf}$ is satisfied by taking $s \ge c' + c''\cdot\max\left(\ln\sigma', \frac{\ln^3\sigma'^{-1}}{\sigma'^2}\right)$, for absolute constants $c', c'' > 0$. □

For any T as in the Claim, $\Pr[T \text{ feasible}] \le \Pr[\mathcal{E}_T^s] + \Pr[\mathcal{E}_T^{nsf}]$, so by the Claim, $\Pr[\mathcal{E}_T^s] \ge \frac{1}{2}\Pr[T \text{ feasible}]$. This implies that $\mathbb{E}[w_{F^a}] \ge \frac{1}{2}\mathbb{E}[opt'_W(L)]$. □

Putting the pieces together, we have that when restricted to large sets, $\mathbb{E}[w_G] \ge \frac{\gamma}{16}\cdot\mathbb{E}[w_{F^a}]$ and $\mathbb{E}[w_{F^a}] \ge \frac{1}{2}\mathbb{E}[opt'_W(L)]$ hold for each a. The claimed approximation ratio then follows from the inequality $\mathbb{E}[opt_W(L)] \le \max(w_H, \mathbb{E}[opt'_W(L)])$.

A. The Case of Unweighted Links

It is easy to see that in the special case of maximum *cardinality* feasible set, our algorithm gives an *additive* h approximation, even without the exhaustive search step. Moreover, Lemmas 3 and 4 imply the following relationship between strong and feasible subsets, which we will use for deriving a scaling law for feasible sets.

Corollary 1. *For a cluster L, let $F \subseteq L$ be the maximum size strong subset of L, for any value of the component X. Then there are constants $c, c' > 0$, such that*

$$c\gamma\cdot\mathbb{E}[|F|] \le \mathbb{E}[opt_C(L)] \le c'\cdot\mathbb{E}[|F|] + h \ .$$

B. Extension to General Sets of Links

The algorithm above can be extended to the more general case of nearly equal length links arbitrarily placed on the plane. This can be achieved by essentially a direct application of the corresponding result for independent shadowing [5]. To this end, [5, Prop. 4.2] shows that every set of nearly equal length links can be partitioned into a constant number of subsets, each consisting of well-separated clusters, where the distance between two clusters is greater than the length of a link (and can be made bigger by more refined partitioning). Now, if we assume that the *shadowing within each cluster is independent of the shadowing in another one that is well-separated from it*, then we can apply the same reasoning as in [5, Thm. 4.3], to extend the our algorithm for clusters to the more general setting. While this assumption seems reasonable, there are many other ways to model correlation between links at a certain distance, e.g., the correlation could be an exponentially decreasing function of distance. We leave more comprehensive treatment of this issue to the future work.

III. Scaling Law for Maximum Feasible Set

In this section, we examine how the expected maximum size of a feasible set in a cluster scales with the number of links and the shadowing correlation. Given Corollary 1, it suffices to estimate the maximum size of a strong subset to bound the optimal size of a feasible set.

Recall that given the fixed value $X = a$, we say a set T is strong if each link $i \in T$ is $\gamma|T|$-strong, i.e., $\mathcal{S}_i > \gamma|T|\cdot\bar{\mathcal{S}}_a$, where $\bar{\mathcal{S}}_a$ is the expected signal strength conditioned on $X = a$, and $\gamma = \frac{1}{2}\cdot\exp(-(1-\rho)\sigma^2/2)$.

For a Lognormal r.v. $Y = e^Z$ with $Z \sim \mathcal{N}(\mu, \sigma^2)$, let $g(k,\sigma)$ denote the value $g > 0$, such that $\Pr[Y > \gamma g\cdot\mathbb{E}Y] = \frac{g}{k}$. We will use this quantity in our bound. We can estimate $g(k,\sigma)$ using (2) and known bounds on the Q function (details in the full version):

$$g(k,\sigma) = \exp(\sigma\sqrt{2\ln k - O(\ln\ln k)} - \sigma^2).$$

Lemma 5. *Let L be a cluster of $k = |L|$ links under correlated LNS. Let R denote the maximum size of a strong subset of L. Then $\mathbb{E}[R] = \Theta(1)\cdot g(k, \sigma\sqrt{1-\rho})$.*

Proof. Let Z_i, Z_{ij}, for $i, j = 1, 2, \ldots, k$, be the Normally distributed logarithms of signals and interferences, such that $\mathcal{S}_i = e^{Z_i}$ and $\mathcal{I}_{ij} = e^{Z_{ij}}$. Recall that $\mu = \mathbb{E}[Z_i] = \mathbb{E}[Z_{ij}]$. By Lemma 2, there are independent $\mathcal{N}(0,\sigma^2)$ r.v.s X, Y_i and Y_{ij}, $i, j = 1, 2, \ldots, k$, such that $Z_i = \sqrt{\rho}X + \sqrt{1-\rho}Y_i + \mu$ and $Z_{ij} = \sqrt{\rho}X + \sqrt{1-\rho}Y_{ij} + \mu$. We condition on the value $X = a$. Then $\{Z_i, Z_{ij} : i, j \in L\}$ become independent $\mathcal{N}(\mu', \sigma'^2)$, where $\mu' = \mu + a\sqrt{\rho}$ and $\sigma' = \sigma\sqrt{1-\rho}$. All probabilities below are conditioned on $X = a$, so we omit the conditioning notation for simplicity of formulas. Since the bounds below do not depend on a, the lemma follows simply by the law of total expectation.

Let $g = g(k, \sigma')$. Denote by I_i the indicator variable that is 1 iff link i is γg-strong. By the definition of g, $p = \Pr[I_i] = \frac{g}{k}$. Thus, $\mathbb{E}\left[\sum_i I_i\right] = kp = g$, and since I_i are independent, the Chernoff bound applies, giving $\Pr\left[\sum_i I_i < g/3\right] < e^{-2g/9}$, which is at most $1/e$ if $g \ge 9/2$, implying that $\mathbb{E}[R] \ge \frac{e-1}{3e}\cdot g$. If $g < 9/2$, then a crude approximation gives $\Pr[R = 0] = (1-p)^k \le e^{-g}$, and

$$\mathbb{E}[R] \ge \Pr[R \ne 0] \ge 1 - e^{-g} \ge \frac{g}{g+1} > \frac{g}{6} \ ,$$

where we used the inequality $e^g \ge 1 + g$.

In order to show the other direction, let $\mathcal{E}(t)$ denote the event that the number of γg-strong links is at least t, i.e., $\sum_i I_i \ge t$. It follows from the definition of strong sets that for $t \ge g$,

$\Pr[\mathcal{E}(t)] \geq \Pr[R \geq t]$. Also note that the expected number of γg-strong links in L can be expressed as $\mathbb{E}\left[\sum_i I_i\right] = \sum_{t=0}^{k} \Pr[\mathcal{E}(t)]$, and similarly, $\mathbb{E}[R] = \sum_{t=0}^{k} \Pr[R \geq t]$. Thus, we have (assume, for simplicity, that g is an integer):

$$\mathbb{E}[R] = \sum_{t=0}^{k} \Pr[R \geq t] = \sum_{t=0}^{g} \Pr[R \geq t] + \sum_{t=g+1}^{k} \Pr[R \geq t] \ .$$

The first term on the right side is at most $g+1$, while, as observed above, the second term is upper bounded by $\sum_{t=0}^{k} \Pr[\mathcal{E}(t)] = \mathbb{E}\left[\sum_i I_i\right] = g$. Together, these give us the bound $\mathbb{E}[R] \leq 2g+1$. □

The lemma and Cor. 1 together imply our scaling law:

Corollary 2. *Let L be a cluster of $k = |L|$ links under correlated LNS, and let γ and h be as in Algs. 1 and 2. Then there are constants $c, c' > 0$, such that*

$$c\gamma g(k, \sigma\sqrt{1-\rho}) \leq \mathbb{E}[opt_C(L)] \leq c' g(k, \sigma\sqrt{1-\rho}) + h \ .$$

IV. Minimum Length Scheduling

We consider here the min. length scheduling problem for a cluster. In the geometric model, almost all links in a cluster must be scheduled separately, namely the scheduling number is linear in $|L|$. Hence, the question is whether we should expect better schedules due to shadowing.

We showed in previous sections that there will be a significant number of strong links due to shadowing, which can be used to form large feasible sets. For min. length scheduling, it does not suffice to focus on the strong links, as we need to schedule *all* links, and not all of them are strong. Instead, we must leverage the variability in the interference strength and group together links that have very low mutual interference. We model this as a coloring problem in Erdős-Renyi random graphs: the nodes correspond to links and edges correspond to interference above a certain threshold. This turns out to characterize the length of optimal schedules, modulo constant factors: on average, nearly a logarithmic number of links (more precisely, $\Theta((1-\rho)\ln k/(\ln\ln k)^2)$ links) can be simultaneously scheduled in a slot, and with high probability, no feasible subset of links is significantly larger.

Theorem 2. *Let L be a cluster of $k = |L|$ links under correlated LNS. Then $\mathbb{E}[opt_S(L)] = \Theta(1) \cdot f(k, \sigma, \rho)$, where $f(k, \sigma, \rho) = \frac{k(\ln\ln k)^2}{\sigma^2(1-\rho)\ln k}$. The result holds even if power control is available.*

Proof. As before, we can express the signal strengths and interferences in terms of independent Normal r.v.s. Namely, $\mathcal{S}_i = e^{Z_i}$, $\mathcal{I}_{ij} = e^{Z_{ij}}$, where $Z_i = \sqrt{\rho}X + \sqrt{1-\rho}Y_i + \mu$ and $Z_{ij} = \sqrt{\rho}X + \sqrt{1-\rho}Y_{ij} + \mu$ for all i, j, and X, Y_i and Y_{ij} are independent $\mathcal{N}(0, \sigma^2)$. Again, by conditioning on $X = a$ for any a, Z_i and Z_{ij} become independent Normal r.v.s with mean $\mu' = \mu + \sqrt{\rho}a$ and variance $\sigma'^2 = (1-\rho)\sigma^2$.

Now, assume $X = a$ is fixed. All probabilities below are conditioned on this event, unless explicitly stated otherwise.

We begin by proving that $\mathbb{E}[opt_S(L)] = O(1) \cdot f(k, \sigma, \rho)$. Let $t > 1$ be a parameter, which we will specify later. We overload the notation to denote $\bar{\mathcal{S}}$ the expected value of $\mathcal{S}_i$, given $X = a$. A link is *weak* if $\mathcal{S}_i < \frac{\bar{\mathcal{S}}}{t}$. Denote

$$q_1 = \Pr\left[\mathcal{S}_i < \frac{\bar{\mathcal{S}}}{t}\right], \text{ and } q_2 = \Pr\left[\mathcal{S}_i < \frac{\bar{\mathcal{S}}}{t\log k}\right] \ .$$

We start by eliminating the weak links. By the assumption, the expected number of weak links is $q_1 \cdot k$. Since the signals are now independent, we can apply Chernoff bound to see that there are at most $2q_1k$ weak links, with probability at least $1 - e^{-q_1k/3}$. We can schedule all weak links in a separate slot for each, and in $2q_1k$ slots in total.

Now, let us focus on the set L' of non-weak links. Let $k' = |L'|$. Consider any pair of links $i, j \in L'$. Let B_{ij} denote the indicator r.v. that is 1 if and only if $\max\{\mathcal{I}_{ij}, \mathcal{I}_{ji}\} \geq \frac{\bar{\mathcal{S}}}{t\log k'}$. Let $p = \Pr[B_{ij} = 1]$. Note that by independence, $p = 1 - \Pr[\max\{\mathcal{I}_{ij}, \mathcal{I}_{ji}\} \leq \frac{\bar{\mathcal{S}}}{t\log k}] = 1 - q_2^2$. Consider the graph G over the set $1, 2, \ldots, n$ where there is an edge between vertices i, j iff $B_{ij} = 1$. Clearly, this is an instance of the Erdős-Rényi random graph $\mathcal{G}_{k',p}$, which is obtained by taking k' vertices and connecting each pair with an undirected edge with probability p, independently of other pairs. It is not hard to see that small independent sets in G give feasible sets in L'. Namely, given any independent set S of size at most $\log k$ in G, it induces a feasible set S' in L', since no link in L' is weak and by the definition of B_{ij}, the total interference on each link in S' by other links in S' is at most $\frac{\bar{\mathcal{S}}}{t\log k} \cdot |S'| < \frac{\bar{\mathcal{S}}}{t}$. Hence, it is sufficient to show that $\mathcal{G}_{k',p}$ can be colored with not too many colors, using only independent sets of small size, i.e. at most $\log k$. A classic result in random graph theory [33] (see also the full paper) shows, that if $1/(1-p) \geq 2$ then there is a greedy algorithm that colors $\mathcal{G}_{k',p}$ with $O\left(\frac{k'}{\log_{1/(1-p)} k'}\right) = O\left(\frac{k'\log q_2^{-1}}{\log k'}\right)$ colors using only independent sets of size $\log_{1/(1-p)} k'$.

Overall, with probability $1 - e^{-q_1k/3}$, we obtain a schedule of length $O(k) \cdot \left(q_1 + \frac{\ln q_2^{-1}}{\ln k}\right)$. Now, the goal is to choose the parameter t so as to minimize the last sum. By evaluating the expressions for q_1 and q_2, we can see that the best choice for t is $t = (\ln k)^{c''}$ for a constant $c'' > 0$, hence the schedule length is $O\left(\frac{k(\ln\ln k)^2}{\sigma^2(1-\rho)\ln k}\right)$, with probability at least $1 - e^{-\Omega(k/\ln k)}$, and in expectation, conditioned on $X = a$.

Next, let us prove that $\mathbb{E}[opt_S(L)] = \Omega(1) \cdot f(k, \sigma, \rho)$. Let L'' denote the set of links i with $\mathcal{S}_i \leq 2\bar{\mathcal{S}}$. Note that by Markov's inequality, each link is in L'' with probability at least $1/2$. The Chernoff bound then implies that $|L''| \geq k/4$, with probability at least $1 - e^{-k/16}$. Consider any subset $S \subseteq L''$ of size $|S| = t$, where $t < k/4$ is a parameter to be specified below. For each pair $i, j \in S$, let B_{ij} be the indicator r.v. that is 1 if and only if $\min\{\mathcal{I}_{ij}, \mathcal{I}_{ji}\} \geq \frac{18}{t} \cdot \bar{\mathcal{S}}$. The proof of the following claim uses the result of [34] on achieveable SIR for a given set of links (see the full paper).

Claim 2. *If $\sum_{ij\in S} B_{ij} > \frac{t(t-1)}{4}$ then S is not feasible, even with power control.*

Thus, in order for S to be feasible, we must have $\sum_{ij\in S} B_{ij} < T = t(t-1)/4$. Namely, for at least half of the pairs i, j, it must be that $B_{ij} = 0$. Denote $q = \Pr[B_{ij} = 0]$. Let R be a subset of T link pairs. By independence, the probability that $B_{ij} = 0$ for all $(i, j) \in R$ is at most q^T. The number of different subsets R of size T is at most $\binom{2T}{T} < 4^T$. Thus, using the union bound, we see that the probability that there are T pairs of links with $B_{ij} = 0$ is at most $4^T \cdot q^T = (4q)^T = (4q)^{t(t-1)/4}$.

By applying union bound over all subsets of L'' of size t, we see that the probability that L'' contains a feasible subset of size t is at most

$$\binom{k/4}{t} \cdot (4q)^{t(t-1)/4} < \left(\frac{ek}{4t}\right)^t \cdot \left((4q)^{(t-1)/4}\right)^t \leq \frac{1}{k},$$

where the last inequality holds if we take $(4q)^{(t-1)/4} \leq 2tk^{-1-\frac{1}{t}}$, i.e. if $t \geq 4$ and $t > \frac{5 \ln k}{\ln (4q)^{-1}}$. Thus, with high probability, $opt_S(L) > \frac{k}{4t}$, given that $t > \frac{5 \ln k}{\ln(4q)^{-1}}$. In order to see which values of t are admissible, we bound q using the union bound,

$$q \leq 2 \cdot \Pr\left[\mathcal{I}_{ji} < \frac{18}{t}\right] = 2Q\left(\frac{s}{\sigma'} - \frac{\sigma'}{2}\right) ,$$

and use known bounds on the Q function to find that there is a valid choice of t with $t = O\left(\frac{\sigma'^2 \ln k}{(\ln \ln k)^2}\right)$, which gives us $\mathbb{E}[opt_S(L)] = \Theta(1) \cdot \frac{k (\ln \ln k)^2}{\sigma^2 (1-\rho) \ln k}$ w.h.p. and in expectation (details in the full version).

The upper and lower bounds for $\mathbb{E}[opt_S(L)]$ above were obtained by conditioning on $X = a$. The theorem then follows by applying the law of total expectation. □

V. Future Work

While the possibility of power control can only affect the constant factors in our results concerning unweighted single-shot scheduling and min. length scheduling in a compact network, it is an interesting question whether algorithms for weighted single-shot scheduling may profit from power control to obtain better solutions. In any case, power control and link adaptation are crucial for utilizing shadowing gains in practice.

Another interesting question arising from our study is: In what extent is the disagreement of predictions between deterministic SINR and stochastic shadowing models reflected in practice?

References

[1] H. Hashemi, "Impulse response modeling of indoor radio propagation channels," *IEEE J. Sel. Areas Commun.*, vol. 11, no. 7, pp. 967–978, 1993.

[2] M. Z. Zamalloa and B. Krishnamachari, "An analysis of unreliability and asymmetry in low-power wireless links," *ACM Trans. Sensor Networks (TOSN)*, vol. 3, no. 2, p. 7, 2007.

[3] Y. Chen and A. Terzis, "On the implications of the log-normal path loss model: an efficient method to deploy and move sensor motes," in *SenSys*. ACM, 2011, pp. 26–39.

[4] B. Dezfouli, M. Radi, S. A. Razak, T. Hwee-Pink, and K. A. Bakar, "Modeling low-power wireless communications," *Journal of Network and Computer Applications*, vol. 51, pp. 102–126, 2015.

[5] M. M. Halldórsson and T. Tonoyan, "Wireless link capacity under shadowing and fading," in *MobiHoc*, 2017, pp. 27:1–27:10.

[6] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. Inf. Theory*, vol. 46, no. 2, pp. 388–404, 2000.

[7] T. Moscibroda, R. Wattenhofer, and Y. Weber, "Protocol design beyond graph-based models," in *HotNets*, 2006.

[8] D. Son, B. Krishnamachari, and J. Heidemann, "Experimental study of concurrent transmission in wireless sensor networks," in *SenSys*. ACM, 2006, pp. 237–250.

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

[10] Y. Chen and A. Terzis, "On the mechanisms and effects of calibrating RSSI measurements for 802.15.4 radios," in *EWSN*, 2010, pp. 256–271.

[11] O. Goussevskaia, Y.-A. Oswald, and R. Wattenhofer, "Complexity in geometric SINR," in *MobiHoc*, 2007, pp. 100–109.

[12] O. Goussevskaia, M. M. Halldórsson, and R. Wattenhofer, "Algorithms for wireless capacity," *IEEE/ACM Trans. Netw.*, vol. 22, no. 3, pp. 745–755, 2014.

[13] T. Kesselheim, "A constant-factor approximation for wireless capacity maximization with power control in the SINR model," in *SODA*, 2011, pp. 1549–1559.

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

[15] M. Hoefer and T. Kesselheim, "Secondary spectrum auctions for symmetric and submodular bidders," *ACM TEAC*, vol. 3, no. 2, p. 9, 2015.

[16] M. Bodlaender and M. M. Halldórsson, "Beyond geometry: Towards fully realistic wireless models," in *PODC*, 2014.

[17] P.-J. Wan, "Joint selection and transmission scheduling of point-to-point communication requests in multi-channel wireless networks," in *MobiHoc*, 2016, pp. 231–240.

[18] P.-J. Wan, B. Xu, O. Frieder, S. Ji, B. Wang, and X. Xu, "Capacity maximization in wireless mimo networks with receiver-side interference suppression," in *MobiHoc*, 2014, pp. 145–154.

[19] M. M. Halldórsson and T. Tonoyan, "How well can graphs represent wireless interference?" in *STOC*, 2015.

[20] E. I. Ásgeirsson, M. M. Halldórsson, and T. Tonoyan, "Universal framework for wireless scheduling problems," in *ICALP*, 2017, pp. 129:1–129:15.

[21] T. Moscibroda and R. Wattenhofer, "The complexity of connectivity in wireless networks," in *INFOCOM*, 2006, pp. 1–13.

[22] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. Automat. Contr.*, vol. 37, no. 12, pp. 1936–1948, 1992.

[23] T. ElBatt and A. Ephremides, "Joint Scheduling and Power Control for Wireless Ad-hoc Networks," in *INFOCOM*, 2002.

[24] A. J. Coulson, A. G. Williamson, and R. G. Vaughan, "A statistical basis for lognormal shadowing effects in multipath fading channels," *IEEE Trans. Communications*, vol. 46, no. 4, pp. 494–502, 1998.

[25] P. Agrawal and N. Patwari, "Correlated link shadow fading in multi-hop wireless networks," *IEEE Transactions on Wireless Communications*, vol. 8, no. 8, pp. 4024–4036, 2009.

[26] P. Stuedi, O. Chinellato, and G. Alonso, "Connectivity in the presence of shadowing in 802.11 ad hoc networks," in *WCNC*, vol. 4. IEEE, 2005, pp. 2225–2230.

[27] T. Muetze, P. Stuedi, F. Kuhn, and G. Alonso, "Understanding radio irregularity in wireless networks," in *SECON*. IEEE, 2008, pp. 82–90.

[28] C. Bettstetter and C. Hartmann, "Connectivity of wireless multihop networks in a shadow fading environment," *Wireless Networks*, vol. 11, no. 5, pp. 571–579, 2005.

[29] R. Hekmat and P. Van Mieghem, "Connectivity in wireless ad-hoc networks with a log-normal radio model," *MONET*, vol. 11, no. 3, pp. 351–360, 2006.

[30] P. Stuedi and G. Alonso, "Log-normal shadowing meets sinr: A numerical study of capacity in wireless networks." in *SECON*, 2007, pp. 550–559.

[31] S. S. Szyszkowicz, H. Yanikomeroglu, and J. S. Thompson, "On the feasibility of wireless shadowing correlation models," *IEEE Trans. Veh. Technol.*, vol. 59, no. 9, pp. 4222–4236, 2010.

[32] J. Bang-Jensen and M. M. Halldórsson, "Vertex coloring edge-weighted digraphs," *Inf. Process. Lett.*, vol. 115, no. 10, pp. 791–796, 2015.

[33] G. R. Grimmett and C. J. H. McDiarmid, "On colouring random graphs," *Math. Proc. Camb. Philos. Soc.*, vol. 77, no. 2, pp. 313–324, 1975.

[34] J. Zander, "Performance of optimum transmitter power control in cellular radio systems," *IEEE Trans. Veh. Technol.*, vol. 41, no. 1, pp. 57–62, 1992.