Algorithms for Wireless Capacity

Olga Goussevskaia, Magnús M. Halldórsson, and Roger Wattenhofer

***Abstract*—In this paper, we address two basic questions in wireless communication. First, how long does it take to schedule an arbitrary set of communication requests? Second, given a set of communication requests, how many of them can be scheduled concurrently? Our results are derived in the signal-to-interference-plus-noise ratio (SINR) interference model with geometric path loss and consist of efficient algorithms that find a constant approximation for the second problem and a logarithmic approximation for the first problem. In addition, we show that the interference model is robust to various factors that can influence the signal attenuation. More specifically, we prove that as long as influences on the signal attenuation are constant, they affect the capacity only by a constant factor.**

***Index Terms*—Approximation algorithms, capacity, physical model, scheduling, wireless networks.**

I. Introduction

DESPITE the omnipresence of wireless networks, surprisingly little is known about their algorithmic complexity and efficiency: Designing and tuning a wireless network is a matter of experience, regardless whether it is a WLAN in an office building, a GSM phone network, or a sensor network on a volcano.

We are interested in the fundamental communication limits of wireless networks. In particular, we would like to know what communication throughput can possibly be achieved. This question essentially boils down to spatial reuse, i.e., which devices can transmit concurrently, without interfering.

The answer to the question stated above depends, among other factors, on the topology of the network. One could be interested in networks where nodes are randomly distributed, or are positioned on a regular grid, as examples of best-case scenarios, i.e., where capacity is maximized. The problem of determining the capacity of such networks has been extensively studied, starting with the seminal work of Gupta and Kumar [24]. Another direction is to restrict attention to link sets with special properties. In [43], a power-assignment algorithm that schedules a strongly connected set of links in polylogarithmic time was presented. This is probably the first algorithmic result in the physical model with guaranteed performance in worst-case topologies; it cannot, however, be extended to schedule arbitrary sets of links and relies strongly on the connectivity requirement.

In this paper, we generalize this research to consider the capacity of *any* network: one with arbitrary topology and an arbitrary set of communication requests. The computational aspect is fundamental: We need to be able to compute the capacity efficiently. Since general instances defy simple laws, the algorithm becomes the means to express capacity. Therefore, if one wants to know the capacity of any network, this paper provides the tool to do that, as it computes the capacity of any network up to a logarithmic factor in the number of communication requests.

In the past, computational research has focused on graph-based models, also known as protocol models. Unfortunately, graph-based models, despite being a useful abstraction, are too simplistic. They fail to capture some essential characteristics of wireless communication, such as the many-to-many relationships underlying wireless interference and the gradual signal attenuation with distance.

Fading channel models, such as the *physical model* (formally introduced in Section III), offer a more realistic representation of wireless communication. A signal is received successfully if the signal-to-interference-plus-noise ratio (SINR)—the ratio of the received signal strength to the sum of the interference caused by all other nodes sending simultaneously, plus noise—is above a hardware-defined threshold. This definition of a successful transmission, as opposed to the graph-based definition, accounts also for interference generated by transmitters located far away. Observe that since the SINR depends on combinations of the transmissions scheduled concurrently, interference is no longer a binary relation (or a graph). This makes the analysis of algorithms more challenging than in graph-based models.

The capacity of wireless networks in fading channel models has received a lot of attention from researchers in information, communication, and network theory. In contrast to the results in graph-based models, which are of algorithmic nature and concerned with arbitrary instances, the results in the physical model have been typically based on heuristics evaluated by simulation of average scenarios. Analytical work in this context has been done only for special cases, e.g., when the network has a grid structure or when traffic is random. Therefore, these results give little insight into the computational complexity of the problem and cannot be translated into algorithms that can ultimately lead to new protocols.

In this paper, we focus on a specific part of the problem of determining the throughput capacity of a wireless network. We study the problem of scheduling one-hop communication requests without power control, i.e., we do not consider routing nor power control problems. The specific questions we address

Manuscript received December 07, 2011; revised October 22, 2012; accepted April 01, 2013; approved by IEEE/ACM Transactions on Networking Editor S. Shakkottai. Date of publication May 01, 2013; date of current version June 12, 2014. This work was supported in part by the Icelandic Research Fund under Grant 90032021 and Grant-of-Excellence 120032011, CNPq, and Fapemig.

O. Goussevskaia is with the Department of Computer Science, Federal University of Minas Gerais, Belo Horizonte 31270-010, Brazil (e-mail: olga@dcc.ufmg.br).

M. M. Halldórsson is with the School of Computer Science, Reykjavik University, 101 Reykjavik, Iceland (e-mail: mmh@ru.is).

R. Wattenhofer is with the Computer Engineering and Networks Laboratory, ETH Zurich, 8092 Zurich, Switzerland (e-mail: wattenhofer@ethz.ch).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TNET.2013.2258036

are two classic issues in wireless communication: Given a set of arbitrary communication requests: 1) how many of them can be scheduled concurrently; and 2) how long does it take to schedule all of them?

We can solve the first problem asymptotically optimally. The solution of the first problem then directly leads to an understanding of the second problem. In particular, it gives an approximation that is optimal up to a factor that is logarithmic in the number of requests. Note that we compute any network's capacity up to a small insecurity, whereas the complete understanding is out of bound since the problem is NP-hard [19].

Our third contribution is a proof of robustness of the physical model with geometric path loss. One may argue that, in reality, path loss will not follow a perfect geometric pattern. Instead, various factors can affect the transmission, e.g., antenna gain may be higher in some directions, obstacles may influence attenuation, and noise may be location-dependent. We show that as long as influences are constant, results will only be affected by a constant. As such, the physical model is *robust*. This result holds in a variety of settings, including power-controlled transmissions.

In the remainder of the Introduction, let us quickly address the two main limitations of our work: single-hop and uniform power. Even though a large body of recent research in wireless communication is about multihop communication, in reality, wireless relaying still is a rare exception, as most wireless systems (e.g., GSM, WLAN) are single-hop. Moreover, understanding the single-hop case also helps understanding the multihop case, as multihop research papers often use a single-hop scheduling algorithm as a basic building block. In Section II, we will give a few examples of how our work was extended to more general scenarios. More surprisingly, this is true also for power control, as the best algorithms with power control [34] can be seen as a generalization of the uniform power algorithm presented in this paper.

II. Related and Current Results

Most work in wireless scheduling in the physical (SINR) model is of heuristic nature, e.g., [6] and [10]. Only after the work of Gupta and Kumar [24] did analytical results become *en vogue*, but only restricted to networks with a well-behaving topology and traffic pattern. On the one hand, this restriction keeps the math involved tractable; on the other hand, it allows for presenting the results in a concise form, i.e., "the throughput capacity of a wireless network with X and Y is Z," where X and Y are some parameters defining the network, and Z is a function of the network size. This area of research has been exceptionally popular, with a multidimensional parameter space (e.g., node distribution, traffic pattern, transport layer, mobility, etc. [23], [37], [39]). An intrinsic problem with this line of research is that, in practice, networks often do not resemble the models studied here, so one cannot learn much about the capacity of an arbitrary network. Moreover, it is difficult to deduce protocols since the results are not algorithmic.

Mathematical programming techniques can be used to formulate the capacity problem and various extensions, typically in the form of convex programming (see, e.g., [47]). The NP-hardness of the problem [19] tells us, however, that one can only hope to solve small instances using such formulations.

In contrast, there is a body of algorithmic work, but mostly on graph-based models. Studying wireless communication in graph-based models commonly implies studying some variants of independent set, matching, or coloring, e.g., [38]. Although these algorithms present extensive theoretical analysis, they are constrained to the limitations of a model that ultimately abstracts away the nature of wireless communication. The inefficiency of graph-based protocols in the SINR model is well documented and has been shown theoretically as well as experimentally [22], [40], [44].

Algorithmic work in the SINR model is fairly new; to the best of our knowledge, it was started just a few years ago [43]. In [43], Moscibroda and Wattenhofer present an algorithm that successfully schedules a set of links (carefully chosen to strongly connect an arbitrary set of nodes) in polylogarithmic time, even in arbitrary worst-case networks. In contrast to our work, the links themselves are *not* arbitrary, but have structure that simplifies the problem. In a follow-up paper, Moscibroda *et al.* [45] first define the link scheduling problem, whose single-shot variant is the focus of this paper. These concepts have been extended and applied to topology control [16], [45], sensor networks [41], combined scheduling and routing [8], ultra-wideband [32], and analog network coding [21], just to name a few. Apart from these papers, algorithmic SINR results also started appearing here and there, such as in a game-theoretic or distributed algorithms context, e.g., [4], [5], [7], [18], [33], and [46].

Previous to our work, few papers appeared that tackle the problem of scheduling arbitrary wireless links. Goussevskaia *et al.* [19] showed that the problem is NP-complete, and Moscibroda *et al.* [42] evaluated popular heuristics. Both papers also present approximation algorithms, with approximation ratios that depend on network parameters and can become linear in the network size.

Since the original publication of our work [17], numerous results have appeared on different aspects of scheduling in the SINR model. The scheduling problem with linear power assignment was treated by Fanghänel *et al.* [14], including a nearly constant approximation. Online algorithms for the dynamic scheduling problem, where communication requests arrive dispersed over time, have been examined in [11], [12], and [25]. Game theory was treated in [1], [2], and [9], and auctioning of spectrum in [31]. Distributed algorithms have been proposed in [2], [9], and [36]. The weighted version of the scheduling problem was studied in [20] and [28].

Halldórsson and Mitra [27] have extended the results of this paper in two ways: from the Euclidean plane to general metric spaces, and to more general range of fixed power assignments. For the case of arbitrary power assignments, Kesselheim [34] gave a constant factor approximation algorithm. Also, if one seeks to maximize capacity with a fixed power assignment, but compare to the optimum that uses arbitrary power, this can be obtained a price of a multiplicative factor of $\Theta(\log\log\Delta)$ [26], which is the best possible [13], [25].

The single-hop capacity problem also plays a central role in more complex scenarios and higher-layer functions, including multihop capacity and flow maximization [48], multirate communication [35], spectrum auctions [26], [31], connectivity and aggregation capacity [29], and the stability of networks under stochastic packet injections [3].

A. Our Results

In this paper, we present the first results that provide approximation guarantees independent of the topology of the network. Our main contributions are the following.

- Given an arbitrary set of requests, we present a simple greedy algorithm that chooses a subset of the requests that can be transmitted concurrently without violating the SINR constraints. This subset is guaranteed to be within a constant factor of the optimal subset.
- Furthermore, by applying the single-slot subroutine repeatedly, we realize an $O(\log n)$-approximation (where n is the number of communication links) for the problem of minimizing the number of time slots needed to schedule a given set of arbitrary requests. Simulation results indicate that this approximation algorithm, besides having an exponentially better approximation ratio in theory, is also practical. It is easy to implement and achieves superior performance in various network scenarios.
- We also present a nonapproximability result for the scheduling problem in the nongeometric SINR model. More specifically, we show that in the SINR model where path loss is set arbitrarily (i.e., not determined by the Euclidean coordinates of the nodes), it is NP-hard to approximate the scheduling problem to within $n^{1-\varepsilon}$ factor (where n is the number of communication links), for any constant $\varepsilon > 0$.
- Finally, we present a general robustness result for the physical model, showing that constant parameter changes, such as path loss and minimum signal ratio, will modify the capacity of the network only by a constant factor.
- All our results rely on a new definition to understand physical interference: *affectance*. This definition has been proved to be of general utility for analyzing algorithms in the SINR context, both for scheduling with fixed-but-different power assignments [27], [36] and in power-controlled scheduling [25], [27], [34].

One may argue that media access and scheduling are fundamental problems when it comes to wireless communication. Although power-controlled cases are interesting from a theoretical point of view, practically the most important cases are those with constant power. Although there are many actual wireless networks where nodes can choose different transmission powers, the selection is then either restricted to a small set of possible power levels, or a bounded power range. The analytical results of this paper hold for both extensions. Apart from constants, all our findings are directly transferrable to bounded power set and to bounded ratio of maximum and minimum power. As such, we believe that our results are practically relevant.

The main features of the current paper, including the general style of the algorithm, affectance analysis, and signal strengthening, factor in and influence nearly all recent work.

This paper fixes several minor plus one larger mistake (an erroneous claim on the scheduling complexity in [30]) from the preliminary conference versions [17] and [30].

III. Notation and Model

Given is a set of links $L = \{\ell_1, \ell_2, \ldots, \ell_n\}$, where each link ℓ_v represents a communication request from a sender s_v to a receiver r_v. We assume the senders and receivers are points in the Euclidean plane; this can be extended to other metric spaces. The Euclidean distance between two points p and q is denoted $d(p,q)$. The asymmetric distance from link ℓ_v to link ℓ_w is the distance from ℓ_v's sender to ℓ_w's receiver, denoted $d_{vw} = d(s_v, r_w)$. The length of link ℓ_v is denoted $d_{vv} = d(s_v, r_v)$. We shall assume for simplicity of exposition that all links are of different length; this does not affect the results. We assume that each link has a unit-traffic demand, and model the case of non-unit-traffic demands by replicating the links. We also assume that all nodes transmit with the same power level P. We show later how to extend the results to variable power levels, with a slight increase in the performance ratio.

We assume the *path-loss radio propagation* model for the reception of signals, where the received signal from transmitter s_w at receiver r_v is $P_{wv} = P/d_{wv}^{\alpha}$ and $\alpha > 2$ denotes the path-loss exponent. We adopt the *physical interference model*, in which a node r_v successfully receives a message from a sender s_v if and only if the following condition holds:

$$\frac{P_{vv}}{\sum_{\ell_w \in S \setminus \{\ell_v\}} P_{wv} + N} \geq \beta \tag{1}$$

where N is the ambient noise, β denotes the minimum SINR required for a message to be successfully received, and S is the set of concurrently scheduled links in the same channel or *slot* as ℓ_v. We say that S is *SINR-feasible* if (1) is satisfied for each link in S.

The problems we treat are the following. In all cases, we are given a set of links of arbitrary lengths. In the `Scheduling` problem, we want to partition the set of input links into minimum number of SINR-feasible sets, each referred to as a *slot*. In the `Single-Shot Scheduling` problem, we seek the maximum cardinality subset of links that is SINR-feasible.

We make crucial use of the following new definitions.

Definition 3.1: The *relative interference* (RI) of link ℓ_w on link ℓ_v is the increase caused by ℓ_w in the inverse of the SINR at ℓ_v, namely $RI_w(v) = P_{wv}/P_{vv}$. For convenience, define $RI_v(v) = 0$. Let $c_v = 1/(1 - \beta N/P_{vv})$ be a link-dependent constant that indicates the extent to which the ambient noise approaches the required signal at receiver r_v. The *affectance* of link ℓ_v, caused by a set S of links, is the sum of the relative interferences of the links in S on ℓ_v, scaled by c_v, or

$$a_S(\ell_v) = c_v \cdot \sum_{\ell_w \in S} RI_w(v).$$

For a single link ℓ_w, we use the shorthand $a_w(\ell_v) = a_{\{\ell_w\}}(\ell_v)$. We define a *p-signal* set or schedule to be one where the affectance of any link is at most $1/p$.

The constant c_v is monotone increasing with the length of the link: $d_{vv} \geq d_{ww}$ implies that $c_v \geq c_w$. Note that $c_v \geq 1$, with equality holding only in the absence of noise.

Observation 3.2: The affectance function satisfies the following properties for a set S of links.

1) (*Range*) S is SINR-feasible if and only if, for all $\ell_v \in S$, $a_S(\ell_v) \leq 1/\beta$.
2) (*Additivity*) $a_S = a_{S_1} + a_{S_2}$, whenever (S_1, S_2) is a partition of S.
3) (*Distance bound*) $a_w(\ell_v) = c_v \cdot (d_{vv}/d_{wv})^{\alpha}$, for any pair ℓ_w, ℓ_v in S.

Note that the concepts of affectance and relative interference are equally useful in contexts of nonuniform power assignments. If P_v is the power of link ℓ_v, the affectance of link ℓ_w on ℓ_v is given by $a_w(\ell_v) = c_v \cdot ((P_w/d_{wv}^\alpha)/(P_v/d_{vv}^\alpha))$.

IV. Properties of SINR Schedules

We present in Section IV-A properties of schedules in the SINR model, which double as tools for the algorithm designer. Then, in Section IV-B, we examine the desirable property of link dispersion and how any schedule can be dispersed at a limited cost.

A. Robustness of the SINR Model

We now explore how signal requirements (in the value of β), or equivalently interference tolerance, affects schedule length. It is not *a priori* obvious that minor discrepancies cause only minor changes in schedule length, but by showing that it is so, we can give our algorithms the advantage of being compared to a stricter optimal schedule. This also has implications regarding the robustness of SINR models with respect to perturbations in signal transmissions.

The pure geometric version of SINR given in (1) is an idealization of true physical characteristics. It assumes, e.g., perfectly isotropic radios, no obstructions, and a constant ambient noise level. That begs the question, why move algorithm analysis from analytically amenable graph-based models to a more realistic model if the latter is not all that realistic? Fortunately, the fact that schedule lengths are relatively invariant to signal requirements shows that these concerns are largely unnecessary.

The results of this section apply equally to scheduling links of different powers. It also applies to throughput optimization.

Theorem 4.1: There is a polynomial-time algorithm that takes a p-signal schedule and refines it into a p'-signal schedule, for $p' > p$, increasing the number of slots by a factor of at most $\lceil 2p'/p \rceil^2$.

Proof: Consider a p-signal schedule $\mathcal{S}$ and a slot S in $\mathcal{S}$. We partition S into a sequence $S_1, S_2, \ldots$ of sets. Initially, set each $S_i = \emptyset$. Order the links in S in decreasing order of length. For each link ℓ_v, assign ℓ_v to the first set S_j for which $a_{S_j}(\ell_v) \leq 1/(2p')$, i.e., the accumulated affectance on ℓ_v among the previous, longer links in S_j is at most $1/(2p')$. Since each link ℓ_v originally had affectance at most $1/p$, then by the additivity of affectance, the number of sets used is at most $\lceil (1/p)/(1/(2p')) \rceil = \lceil 2p'/p \rceil$.

We then repeat the same approach on each of the sets S_i, processing the links this time in increasing order. The number of sets is again $\lceil 2p'/p \rceil$ for each S_i, or $\lceil 2p'/p \rceil^2$ in total. In each final slot (set), the affectance on a link by shorter links in the same slot is at most $1/2p'$. In total, then, the affectance on each link is at most $2 \cdot 1/2p' = 1/p'$. ■

This result applies in particular to optimal solutions. Let $\psi(L)$ denote the minimum number of slots in an SINR-feasible schedule of a link set L, and let $\psi_p(L)$ denote the same quantity for an optimal p-signal schedule. It is not *a priori* clear that a smooth relationship exists between ψ_p and $\psi = \psi_\beta$, for $p > 1$.

Corollary 4.2: For any link set L and any $p > 1$, $\psi_p(L) \leq \lceil 2p/\beta \rceil^2 \psi(L)$.

This has significant implications. One regards the validity of studying the pure SINR model. As asked in [17], "what if the signal is attenuated by a certain factor in one direction but by another factor in another direction?" A generalized physical model was introduced in [45] to allow for such a deviation. Theorem 4.1 implies that scheduling is relatively robust under discrepancies in the SINR model. This validates analytic studies of the pure SINR model, in spite of its simplifying assumptions.

Corollary 4.3: If a scheduling algorithm gives a ρ-approximation in the SINR model, it provides an $O(\theta^2\rho)$-approximation in variations in the SINR model with a discrepancy of up to a factor of θ in signal attenuation or ambient noise levels.

This result can be contrasted with the result of Section VII, which shows a strong $n^{1-\epsilon}$-approximation hardness of scheduling in an abstract (nongeometric) SINR model that allows for arbitrary distances between nodes. Alternatively, Theorem 4.1 allows us to analyze algorithms under more relaxed situations than the optimal solutions to which we compare.

It is important to note that these results do not depend on the power assignment and apply equally well in the power-control setting. Also, they actually do not depend on the formula used to compute affectance or relative interference, and apply also in nongeometric and nonmetric settings.

Remark: Note that the converse of Theorem 4.1—that a schedule can be shortened by a constant factor so that the signal decreases only by a constant factor—does not hold. An easy example is found by making t copies of a feasible set S (possibly separating the nodes by a sufficiently small distance) for any number t. Any attempt to use fewer than t slots results in an arbitrarily bad signal.

B. Dispersion Properties

One desirable property of schedules is that links in the same slot be spatially well separated. This blurs the difference in position between sender and receiver of a link since it affects distances only by a small constant. Intuitively, we want to measure nearness as a fraction of the lengths of the respective links. Given the affectance measure, it proves to be useful to define nearness somewhat less restrictively.

Definition 4.4: Link ℓ_w is said to be q-*near* link ℓ_v if $d_{wv} < q \cdot c_v^{1/\alpha} \cdot d_{vv}$. A set of links is q-*dispersed* if no (ordered) pairs of links in the set are q-near.

Observation 3.2, item 3, states that link ℓ_w is q-near a link ℓ_v if and only if $a_w(\ell_v) > q^{-\alpha}$. This immediately gives the following strengthening of [17, Lemma 4.2].

Lemma 4.5: Fewer than q^α/β senders in an SINR-feasible set S are q-near to any given link $\ell_v \in S$.

For q constant, any schedule can be made q-dispersed at a cost of a constant factor.

Lemma 4.6: There is a polynomial-time algorithm that takes an SINR-feasible schedule and refines it into a q-dispersed schedule, increasing the number of slots by a factor of at most $\lceil (q+1)^\alpha \rceil$.

Proof: Let S be a slot in the schedule. We show how to partition S into sets $S_1, S_2, \ldots, S_t$ that are q-dispersed, where $t \leq (q+1)^\alpha + 1$.

Initially, all S_i are empty. Process the links of S in increasing order of length, assigning each link ℓ_v "first-fit" to the first set S_j in which the receiver r_v is at least $(qc_v^{1/\alpha} + 1) \cdot d_{vv}$ away from the receiver of any other link. Let ℓ_w be a link previously in S_j, and note that ℓ_w is shorter than ℓ_v. By the selection rule, $d_{wv} \geq (qc_v^{1/\alpha} + 1) \cdot d_{vv} - d_{ww} > qc_v^{1/\alpha} \cdot d_{vv}$. Also

$$\begin{aligned} d_{vw} &\geq d_{wv} - d_{vv} \\ &\geq \left(qc_v^{\frac{1}{\alpha}} + 1\right) d_{vv} - d_{vv} \\ &\geq qc_v^{1/\alpha} d_{ww} \\ &\geq qc_w^{\frac{1}{\alpha}} d_{ww}. \end{aligned}$$

Since this holds for every pair in the same set, the schedule is q-dispersed. Suppose S_t is the last set used by the algorithm, and let ℓ_v be a link in it. Then, each S_i, for $i = 1, 2, \ldots, t-1$, contains a link whose sender is closer than $(qc_v^{1/\alpha} + 2) \cdot d_{vv} \leq (q+1)c_v^{1/\alpha} d_{vv}$ to r_v, i.e., is $(q+1)$-near to ℓ_v. By Lemma 4.5, $t - 1 < (q+1)^\alpha/\beta$. Hence, $t \leq \lceil (q+1)^\alpha/\beta \rceil$. ■

Intuitively, there is a correlation between low affectance and high dispersion in schedules. The following result makes this connection clearer. The converse, however, is not true, since high interference can be caused by shorter faraway links.

Lemma 4.7: A p-signal schedule is also $p^{1/\alpha}$-dispersed.

Proof: Let ℓ_v and ℓ_w be an ordered pair of links in a slot S in a p-signal schedule. By definition, $a_w(\ell_v) \leq a_S(\ell_v) \leq 1/p$. By Observation 3.2, item 3, $d_{wv} \geq (c_v p)^{1/\alpha} \cdot d_{vv}$. ■

We remark that the results given in this section apply only to uniform power assignments, unlike the Section IV-A.

V. Approximation Algorithms

We now give a constant-factor approximation algorithm for `Single-Shot Scheduling`. We aim for conceptual simplicity, rather than optimizing the constants.

Let $C = 2^3\, 9 = 72, \tau = 2 + \max(2, ((C+1)\beta((\alpha-1)/(\alpha-2)))^{1/\alpha})$, and $c = 1/\tau^\alpha$.

It is rather surprising that an $O(1)$-approximation algorithm can be obtained in a single sweep. This should help in applying the ideas further, e.g., in distributed implementations. Note that recent research shows that such a single sweep is also feasible when using power control [34].

It is not immediate that Algorithm 1 produces a feasible solution.

Algorithm 1: One-Slot Scheduling (Algorithm A)

```
input: Set of links L = {ℓ_1, ..., ℓ_n};
output: Feasible subset S of links;
sort the links ℓ_1, ℓ_2, ..., ℓ_n in nondecreasing order of
    length;
S ← ∅;
for v ← 1 to n do
    if a_S(ℓ_v) ≤ c then
        add ℓ_v to S;
    end if
end for
return S;
```

Lemma 5.1: Algorithm 1 produces a $(\tau - 2)$-dispersed solution.

Proof: Let ℓ_w be a link in the set S output by Algorithm 1. Let N_w^- (N_w^+) be the set of links in S that are shorter (longer) than ℓ_w. Consider first a link $\ell_u \in N_w^-$. Since ℓ_w was added by the algorithm after adding ℓ_u, $a_u(\ell_w) \leq c = 1/\tau^\alpha$, which implies by Observation 3.2, item 3, that $d_{uw} \geq \tau c_w^{1/\alpha} d_{ww} > (\tau - 2)c_w^{1/\alpha} d_{ww}$. Consider next a link $\ell_v \in N_w^+$. Since ℓ_v was added after ℓ_w, it holds that $a_w(\ell_v) \leq c = 1/\tau^\alpha$. Hence, by Observation 3.2, $d_{wv} \geq \tau \cdot c_v^{1/\alpha} d_{vv}$. Recall that $c_v \geq c_w$ whenever $d_{vv} \geq d_{ww}$. Then, using the triangular inequality

$$\begin{aligned} d_{vw} = d(s_v, r_w) &\geq d_{wv} - d_{vv} - d_{ww} \\ &\geq \left(\tau c_v^{\frac{1}{\alpha}} - 2\right) d_{vv} \\ &\geq (\tau - 2) c_w^{\frac{1}{\alpha}} d_{ww}. \end{aligned}$$

Since this holds for every ordered pair in S, we have that S is $(\tau - 2)$-dispersed [see Fig. 1(a)]. ■

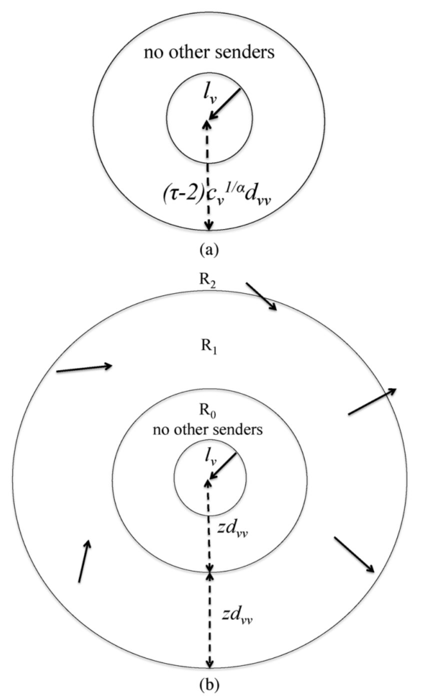

Fig. 1. Illustrations to Lemmas 5.1 and 5.2. (a) $(\tau - 2)$-dispersed link set. (b) Concentric rings around the receiver r_v.

Lemma 5.2: Let S be a Z-dispersed feasible set of links, where $Z \geq 2$. Then, for any link ℓ_v in S, it holds that

$$a_{S_v^+}(\ell_v) < \left(\frac{\alpha - 1}{\alpha - 2} C\right) Z^{-\alpha}$$

where S_v^+ is the set of links in S at least as long ℓ_v.

Proof: Let $z = Zc_v^{1/\alpha}$. Form a disc D_w of radius $r = (z-1)d_{vv}/2$ around each sender s_w in S_v^+. We claim that these discs are disjoint. By the dispersion property, the distance from any sender $s_u \in S$ to any receiver $r_w \in S_v^+$, $w \neq u$, is at least $Zc_w^{1/\alpha} d_{ww} \geq zd_{ww}$, using that $c_w \geq c_v$ since $\ell_w \geq \ell_v$. It follows by the triangular inequality that the separation between two senders s_u, s_w in S is at least $(z-1)d_{ww} \geq (z-1)d_{vv} = 2r$, and thus the discs are disjoint.

We next partition the sender set in S_v^+ into concentric rings R_k of width $z \cdot d_{vv}$ around the receiver r_v [see Fig. 1(b)]. Each ring R_k contains all senders $s_w \in S_v^+$ satisfying $k(z \cdot d_{vv}) \leq d_{wv} \leq (k+1)(z \cdot d_{vv})$. We know that the first ring R_0 contains no sender (since such links would be incompatible with ℓ_v). For each $k > 0$, the senders in R_k are contained in an annulus A_k centered at r_v of width $zd_{vv} + 2r = (2z-1)d_{vv}$ that has r added both to the inside and outside of R_k. The area of A_k is

$$\begin{aligned}\mathrm{Area}(A_k) &= \left[(d_{vv}(k+1)z + r)^2 - (d_{vv}kz - r)\right]\pi \\ &= (2k+1)d_{vv}^2 z(2z-1)\pi.\end{aligned}$$

Since discs D_w of area $\mathrm{Area}(D_w) = r^2\pi$ around senders in S_v^+ do not intersect, and the minimum distance between r_v and $s_w \in R_k$, $k > 0$ is $k(z \cdot d_{vv})$, we can use an area argument to bound the number of senders inside each ring. The total relative interference from senders in R_k, $k \geq 1$ on ℓ_v is bounded by

$$\begin{aligned}RI_{R_k}(\ell_v) &\leq \sum_{s_w \in R_k} RI_{s_w}(\ell_v) \\ &\leq \frac{A(A_k)}{A(D_w)} \cdot \frac{1}{(kz)^\alpha} \\ &\leq \frac{(2k+1)}{k^\alpha} \cdot \frac{4}{z^\alpha}\frac{z(2z-1)}{(z-1)^2} \\ &\leq \frac{1}{k^{(\alpha-1)}} \cdot \frac{2^3 9}{z^\alpha}\end{aligned}$$

where the last inequality holds since $k \geq 1 \Rightarrow 2k+1 \leq 3k$ and $z \geq 2 \Rightarrow z-1 \geq z/2$ and $2z-1 \leq 3(z-1)$. Summing up the interferences over all rings yields

$$RI_{S_v^+}(\ell_v) < \sum_{k=1}^{\infty} RI_{R_k}(\ell_v) \leq \sum_{k=1}^{\infty} \frac{1}{k^{\alpha-1}} \cdot \frac{C}{z^\alpha} < \frac{\alpha-1}{\alpha-2} \cdot \frac{C}{z^\alpha}$$

where the last inequality holds since $\alpha > 2$. This results in affectance of

$$a_{S_v^+}(\ell_v) = c_v RI_{S_v^+}(\ell_v) < \frac{\alpha-1}{\alpha-2} C \cdot \left(\frac{c_v^{1/\alpha}}{z}\right)^\alpha$$

as claimed. ■

Theorem 5.3: Algorithm 1 produces an SINR-feasible solution.

Proof: Let ℓ_w be a link in the set S output by Algorithm 1. Let S_w^- (S_w^+) be the set of links in S that are shorter (longer) than ℓ_w. The links in S_w^- were processed before ℓ_w, so by the if-condition in the algorithm, $a_{S_w^-}(\ell_v) \leq c$. Note that $c \leq (1/(C+1)\beta)$. By Lemma 5.1, S is $\tau - 2$-dispersed, so by Lemma 5.2 and the definitions of τ and dispersion

$$a_{S_w^+}(\ell_w) < \left(\frac{\alpha-1}{\alpha-2}C\right)\frac{1}{(\tau-2)^\alpha} \leq \frac{C}{(C+1)\beta}.$$

Hence, the affectance of each link ℓ_v in S is at most $a_{S_w^-}(\ell_v) + a_{S_w^+}(\ell_v) \leq 1/\beta$. ■

A. Performance Analysis

Definition 5.4: Let $\mathcal{R}$ and $\mathcal{B}$ be disjoint pointsets in a metric space $(\mathcal{V}, d)$, referred to as the *red* and *blue* points, respectively. A point $b \in \mathcal{B}$ is *blue-dominant* if every ball $B_\delta(b) = \{w \in \mathcal{B} | d(w,b) \leq \delta\}$ around b contains more blue points than red points. Formally, $\forall \delta \in \mathbb{R}_0^+ : |B_\delta(b) \cap \mathcal{B}| > |B_\delta(b) \cap \mathcal{R}|$.

For a red point $r \in \mathcal{R}$ and a set $G \subseteq \mathcal{B}$ of blue points, we say that G *guards* r if for all $b \in \mathcal{B} \setminus G$, we have that $B_{d(b,r)}(b) \cap G \neq \emptyset$.

Lemma 5.5: (Blue-Dominant Centers Lemma): Let $\mathcal{R}$ and $\mathcal{B}$ be disjoint sets of red and blue points in a two-dimensional Euclidean space. If $|\mathcal{B}| > 5 \cdot |\mathcal{R}|$, then there exists at least one blue-dominant point in $\mathcal{B}$.

Proof: Process the points in $\mathcal{R}$ in an arbitrary order while maintaining a subset $\mathcal{B}'$ of $\mathcal{B}$ as follows (initially, $\mathcal{B}' = \mathcal{B}$). For each $r \in \mathcal{R}$, we construct a guarding set $G(r) \subseteq \mathcal{B}'$ (guarding r relative to the current $\mathcal{B}'$) and remove $G(r)$ from $\mathcal{B}'$.

We claim that it is possible to construct a guarding set $G(r)$ of size at most 5. The procedure to construct $G(r)$ is as follows. Consider a red point r. Include a closest blue point $b_1 \in \mathcal{B}'$ in $G(r)$. Draw five sectors originating at r in the following manner. The first sector has 120° and is centered at b_1, the remaining four sectors have 60° each and evenly divide the remaining area around r. For each of these four sectors sec_j, include the closest blue point $b_j \in sec_j$ in $G(r)$ (if sec_j has no blue points from $\mathcal{B}'$, skip this sector). Now $G(r)$ has size at most 5, and we claim that it is guarding r. Suppose not. Then, there is a point $b^* \in \mathcal{B}' \setminus G(r)$ with $B_{d(b^*,r)}(r) \cap G(r) = \emptyset$. Suppose b^* is located in sec_j and we selected blue point b_j from sec_j into $G(r)$. This means that $d(b^*, b_j) > d(b^*, r)$, which implies that the sector angle is larger than 60°. (Note that if $G(r)$ contains no point b_j from sector sec_j, then b^* would have been picked to guard r in that sector, also establishing a contradiction.)

After going through all the points in $\mathcal{R}$, the set $\mathcal{B}'$ is still nonempty by the assumption on the relative sizes of $\mathcal{R}$ and $\mathcal{B}$. We claim that every point in $\mathcal{B}'$ is now blue-dominant. This holds since: 1) the guarding sets of points in $\mathcal{R}$ are pairwise disjoint; and 2) every ball $B_\delta(b^*)$, $b^* \in \mathcal{B}'$, that contains a red point r contains also a blue point in $G(r)$. Hence, for every blue node $b^* \in \mathcal{B}'$, every ball $B_\delta(b^*)$ contains more blue points than red points ("more" since the center b^* is also blue). ■

Lemma 5.6: Let $\nu = 2(3\tau/2)^\alpha$ be a constant. Let ALG be the solution output by Algorithm 1 on the given instance and OPT_ν be an optimal ν-signal solution. Then, $|OPT_\nu| \leq 5|ALG|$.

Proof: Let $\mathcal{R} = \{s_w | \ell_w \in ALG \setminus OPT_\nu\}$ and $\mathcal{B} = \{s_v | \ell_v \in OPT_\nu \setminus ALG\}$ be the sets of senders in exactly one of ALG and OPT_ν; we call them red and blue points, respectively. Suppose the claim is false. It follows that $|\mathcal{B}| > 5|\mathcal{R}|$. By Lemma 5.5, there is a blue-dominant s_b in $\mathcal{B}$. We shall argue that the blue link $\ell_b = (s_b, r_b)$ would have been picked by our algorithm, which is a contradiction.

Consider any red point $s_x \in \mathcal{R}$. Let $D = d(s_x, s_b)$. Let s_y denote the guard for s_x w.r.t. s_b, i.e., the blue point that is closer to s_b than s_x is, i.e., within distance D from s_b. Note that by Lemma 4.7, OPT_ν is a s-dispersed set, where $s = \nu^{1/\alpha} \geq 3\tau/2 \geq 6$. Applying Definition 4.4, we know that $d_{xb} \geq s \cdot c_v^{1/\alpha} \cdot d_{bb}$. Using $c_v \geq 1$, we get $d_{xb} \geq 6d_{bb}$. The guarding property and the triangular inequality ensure that

$$d_{yb} \leq d(s_y, s_b) + d_{bb} \leq D + d_{bb} \leq d_{xb} + 2d_{bb} \leq \frac{4}{3}d_{xb}.$$

Thus

$$a_x(b) = c_b\left(\frac{d_{bb}}{d_{xb}}\right)^\alpha \leq c_b\left(\frac{4}{3}\cdot\frac{d_{bb}}{d_{yb}}\right)^\alpha = \left(\frac{4}{3}\right)^\alpha a_y(b).$$

Let t denote $(3/4)^\alpha$. This holds for any $s_x \in \mathcal{R}$, so the total interference that ℓ_b receives from the red senders (those in ALG) is at least t times that from the blue senders. Since ℓ_b is in OPT_ν, it is affected by at most $1/\nu$ by OPT_ν. Using that each node in OPT_ν participates in at most one guardset, we get that

$$\begin{aligned} a_{ALG\setminus OPT_\nu}(\ell_b) &= \sum_{s_x \in \mathcal{R}} a_x(\ell_b) \\ &\leq \sum_{\ell_g \in \mathcal{B}} t \cdot a_g(b) \\ &= t \cdot a_{OPT\setminus ALG}(\ell_b) \\ &\leq t/\nu < c/2. \end{aligned}$$

Furthermore, since OPT_ν is a ν-signal solution, $a_{ALG\cap OPT_\nu}(\ell_b) \leq 1/\nu < c/2$. Thus

$$a_{ALG}(\ell_b) = a_{ALG\setminus OPT_\nu}(\ell_b) + a_{ALG\cap OPT_\nu}(\ell_b) < c$$

which contradicts the assumption that ℓ_b was not selected by the algorithm. ■

The following result is now immediate from Lemma 5.6 in combination with the correctness result in Theorem 5.3 and the signal-strengthening property of Corollary 4.2.

Theorem 5.7: Algorithm 1 approximates the `Single-Shot Scheduling` problem within a constant factor.

B. Scheduling Approximation

Given the constant factor approximation for the `Single-Slot Scheduling` problem, we get an $O(\log n)$-approximation for the `Scheduling` problem by repeatedly executing the `Single-Slot Scheduling` algorithm, and as such always removing a large set of links that can be scheduled concurrently, without interference. See Algorithm 2.

Algorithm 2: MultiSlot Scheduling (ApproxA)

input: Set of links $L = \{\ell_1, \ldots, \ell_n\}$;
output: Schedule $\mathcal{S} = \{\mathcal{S}_1, \cdots, \mathcal{S}_T\}$;
$t := 0$;
repeat
  $\mathcal{S}_t := OneSlotScheduling(L)$; (Algorithm 1)
  $L := L \setminus \mathcal{S}_t$;
  $t := t + 1$;
until $L = \emptyset$
return $\mathcal{S}$;

Theorem 5.8: Repeated application of Algorithm 1 yields an $O(\log n)$-approximation for the `Scheduling` problem.

Proof: Recall that ψ is the minimum number of slots in a feasible solution, and let $\rho = O(1)$ be the performance guarantee of Algorithm 1. Any subset S' of the input instance with N links contains a feasible set of size N/ψ. Thus, Algorithm 1 applied to S' results in a feasible subset of size at least $N/(\rho\psi)$, with the number of remaining unscheduled links becoming at most $N(1 - 1/(\rho\psi))$. Starting with n links, the number of unscheduled links remaining after s iterations is at most $n(1 - 1/(\rho\psi))^s < ne^{-s/(\rho\psi)}$. Thus, when $s \geq \ln n \cdot \rho\psi$, less than one link remains unscheduled, that is, all the links have been scheduled. Hence, $\ln n \cdot \rho\psi$ slots suffice, for an approximation factor of $\rho \ln n$. ■

C. Handling Different Transmission Powers

We can treat the case when links transmit with different powers in two different ways. Let $P_{\max}$ ($P_{\min}$) be the maximum (minimum) power used by a link, respectively. By introducing a factor of $P_{\min}/P_{\max}$ into the affectance threshold c, our algorithm still produces a feasible schedule, that is longer by a factor of at most $P_{\max}/P_{\min}$.

Alternatively, we can partition the instance into "power regimes," where each regime consists of links whose powers are equal up to a factor of 2. We schedule each power regime separately, obtaining an approximation factor of at most $\log P_{\max}/P_{\min}$, or at most the number of different power values.

If $P_{\max}/P_{\min}$ cannot be bounded, and if more generally the number of power levels cannot be bounded, we refer to recent work of [27] and [34].

VI. Simulation Results

In this section, we present some simulation results to better illustrate the practical appeal of the scheduling approximation algorithm (we use the multislot version (Algorithm 2) and refer to it as ApproxA). We compare the performance of ApproxA to the performance of three other scheduling algorithms: ApproxLogN (first proposed by us in [17]), GreedyPhysical (proposed in [6]), and ApproxDiversity (proposed in [19]). All are polynomial-time algorithms, specifically designed for the SINR model. ApproxLogN is very similar in nature to ApproxA. The two algorithms are asymptotically equivalent, but ApproxA yields a cleaner analysis, while ApproxLogN might result in a constant-factor performance gain because it uses an additional distance-based constraint to select links.

We generated two kinds of topologies: *random* and *clustered* [see Fig. 2(a) and (b)]. In the random topology, n receiver nodes were distributed uniformly at random on a plane field of size 1000×1000 units, and n senders were positioned uniformly at random inside discs of radius $l_{\max}$ around each of the receivers. In the clustered topology, n_C cluster center positions were selected uniformly at random on the plane, and n/n_C sender–receiver pairs were positioned uniformly at random inside discs of radius r_C around each of them. The clustered topology aims to simulate a scenario of heterogeneous density distribution.

In all experiments, mean and standard deviation values were plotted based on multiple simulation runs of random instances.

First, we analyze the lengths of the schedules as a function of the number of nodes ($n \in \{100 \cdot 2^0, 100 \cdot 2^1, \cdots, 100 \cdot 2^8\}$). In Fig. 3(a) and (b), the results for the random topology are shown. Since this scenario is not particularly challenging, all four algorithms have good performance, computing schedules of comparable sizes. The performance ratio between ApproxA and ApproxLogN is, as expected, constant with the number of nodes. ApproxLogN presents a slightly better performance on average (25% shorter schedules). In very low-density scenarios [see zoomed-in plots on Fig. 4(a) and (b)], GreedyPhysical presents better performance among all algorithms. As the density increases, however, ApproxA and ApproxLogN present increasingly better relative performance. ApproxDiversity computes schedules that are, on average, twice as long as those computed by ApproxA.

In Fig. 5(a) and (b), the results for the clustered topology are shown. As could be expected, the greedy algorithm is not able to

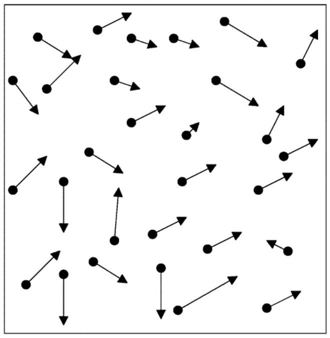

(a)

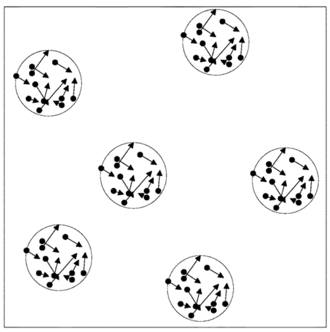

(b)

Fig. 2. Simulated topologies: 1000×1000 field, $\alpha = 3$, $\beta = 1.2$, $N = 0$. (a) Random. (b) Clustered.

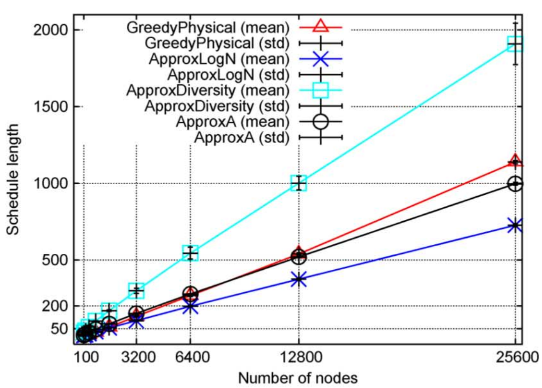

(a)

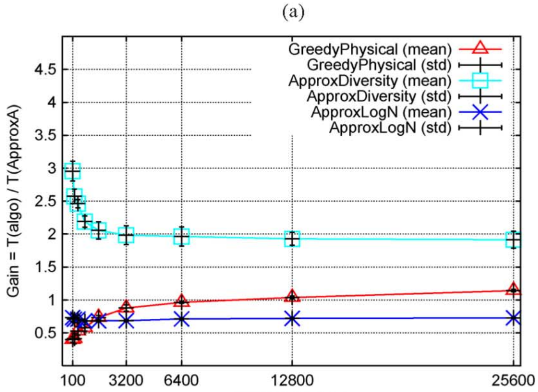

(b)

Fig. 3. Random topology: $l_{\max} = 20$. (a) Schedule length. (b) Gain.

deal with this more difficult scenario as efficiently. Even in very sparse topologies [Fig. 6(a) and (b)], GreedyPhysical computes three times longer schedules than ApproxA. As the density

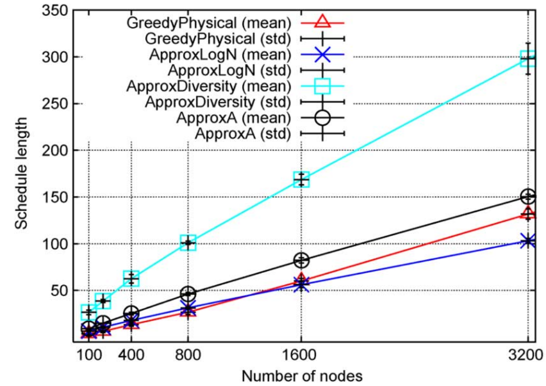

(a)

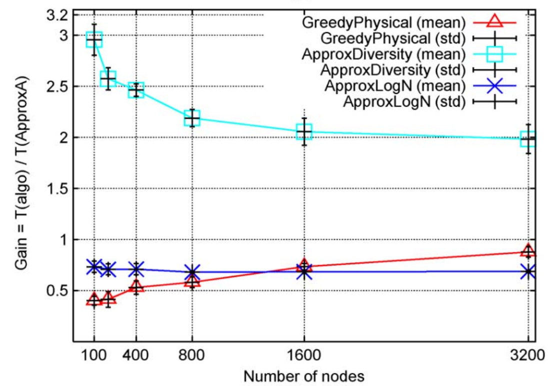

(b)

Fig. 4. Random topology (zooming into small instances): $l_{\max} = 20$. (a) Schedule length. (b) Gain.

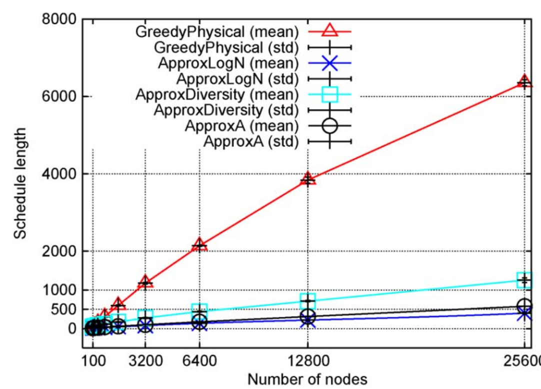

(a)

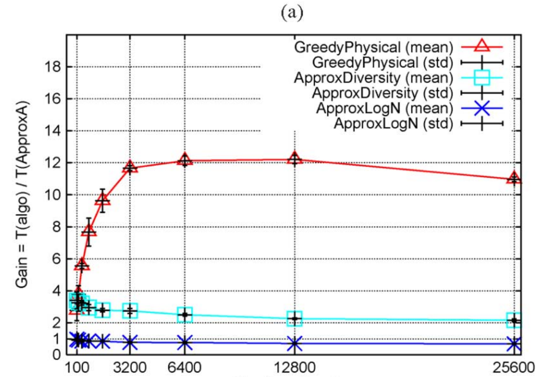

(b)

Fig. 5. Clustered topology: $n_C = n/10$, $r_C = 10$. (a) Schedule length. (b) Gain.

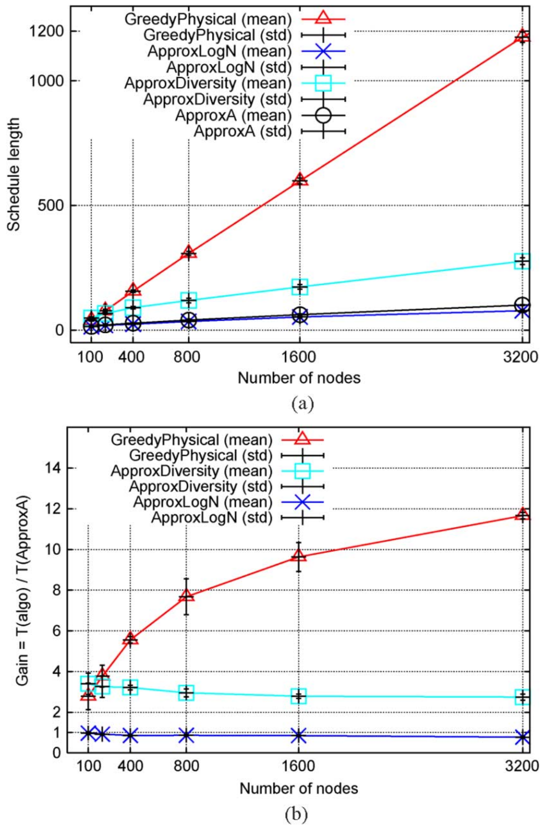

Fig. 6. Clustered topology (zooming into small instances): $n_C = n/10, r_C = 10$. (a) Schedule length. (b) Gain.

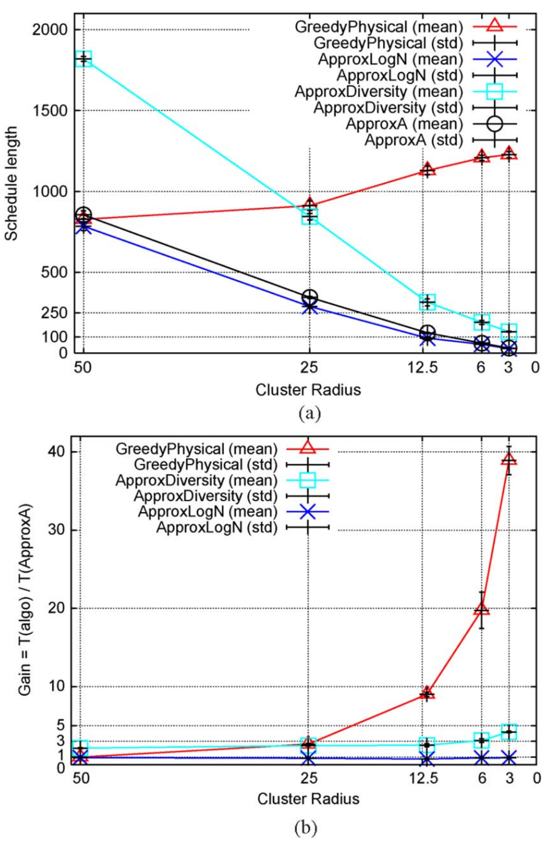

Fig. 7. Clustered topology: $n = 3.2K$, $n_C = n/10$. (a) Schedule length. (b) Gain.

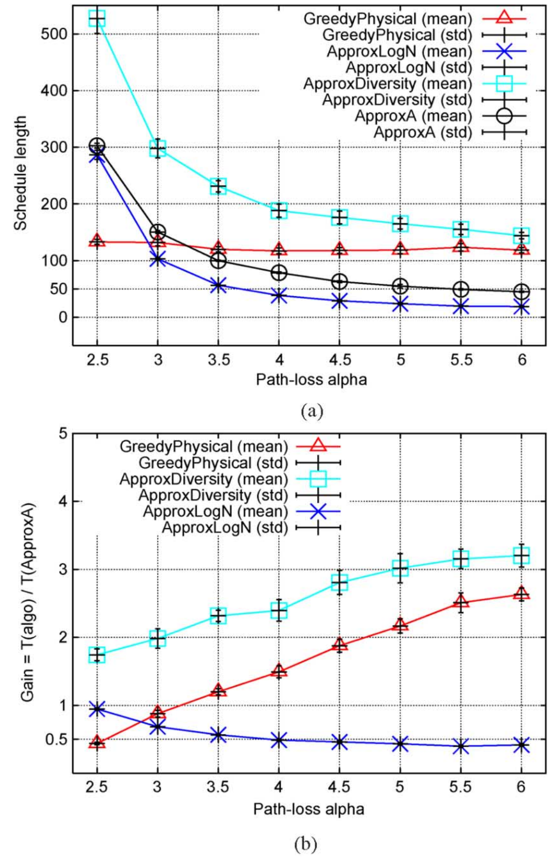

Fig. 8. Random topology: $n = 3.2K$, $l_{\max} = 20$. (a) Schedule length. (b) Gain.

increases, the relative performance of the greedy algorithm deteriorates. ApproxA, ApproxLogN, and ApproxDiversity compute even shorter schedules than in the random case, which indicates that they are able to schedule many clusters in parallel. The performance of ApproxA and ApproxLogN is still superior to that of ApproxDiversity by a factor of 3.

In Fig. 7(a) and (b), we analyze the influence of the cluster radius. In topologies with smaller clusters, i.e., in scenarios with higher density heterogeneity, the difference in performance becomes more accentuated. Whereas GreedyPhysical's performance slightly decreases with decreasing cluster radius, ApproxA and ApproxLogN (and ApproxDiversity) are able to compute ever shorter schedules. Smaller cluster radius means more separate clusters, which makes it easier to schedule clusters in parallel. GreedyPhysical, however, is not able to take advantage of this possibility. Among all three algorithms, ApproxLogN presents the best performance in all cases.

Next, we analyze the influence of the path-loss exponent α in both random [Fig. 8(a) and (b)] and clustered [Fig. 9(a) and (b)] topologies. It can be seen that the performances of ApproxA, ApproxLogN, and ApproxDiversity improve with increasing α, whereas GreedyPhysical is more or less invariant to the path-loss exponent. For $\alpha < 3$, in the random topology, GreedyPhysical presents a better performance than the other three algorithms. In the clustered topology, however, its performance is very poor even for low α and deteriorates relative to the other approaches with increasing α in both kinds of topologies. Among all four algorithms, ApproxLogN presents the

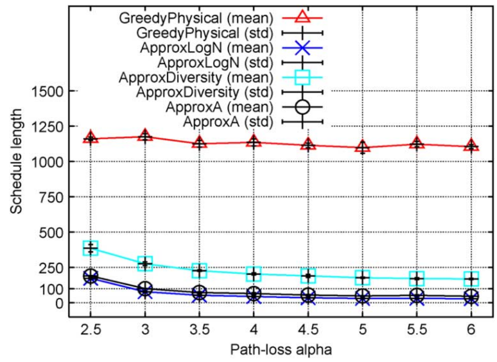

(a)

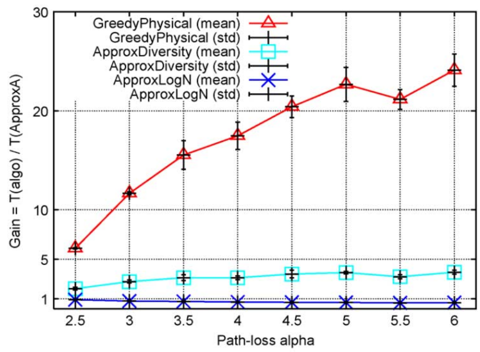

(b)

Fig. 9. Clustered topology: $n = 3.2K$, $n_C = n/10$, $r_C = 10$. (a) Schedule length. (b) Gain.

best performance for all values of α in the clustered topology and for $\alpha \geq 3$ in the random case.

To sum up, the simulations show that ApproxA and ApproxLogN, besides having an exponentially better analytical approximation ratio, present advantages in challenging practical scenarios, such as high-density and heterogeneous-density networks.

VII. Nonapproximability in Abstract SINR

In this section, we show that scheduling is extremely hard if the path-loss function can be nongeometric.

We distinguish "abstract SINR" $(\mathrm{SINR_A})$ from "geometric SINR" $(\mathrm{SINR_G})$ model according to the freedom with which the gain (or path-loss) matrix can be defined. In the $\mathrm{SINR_A}$ model, as opposed to the $\mathrm{SINR_G}$ model, path loss between nodes is not constrained by their Euclidean coordinates, but can be set arbitrarily (i.e., triangular inequality need not be preserved when defining the path-loss matrix). Note that $\mathrm{SINR_A}$ is more general and, therefore, a "harder" model than $\mathrm{SINR_G}$, which we have been using to derive the results in the previous sections. We also remark that these results do not depend on complications due to noise.

Theorem 7.1: The scheduling problem in the $\mathrm{SINR_A}$ model is at least as hard to approximate as the graph coloring problem, and the single-shot scheduling problem is as hard as the maximum independent set problem in graphs. In particular, the scheduling problem is NP-hard to approximate within $n^{1-\varepsilon}$-factor, for any $\varepsilon > 0$.

Proof: Let $G = (V, E)$ be a graph on n vertices. We form an instance I to the scheduling problem, containing a link ℓ_i for each node v_i and a symmetric gain matrix $A = (a_{ij})$. The value of a_{ij} corresponds to the affectance of ℓ_i on ℓ_j (and, by symmetry, the affectance of ℓ_j on ℓ_i). We define

$$a_{ij} = \begin{cases} 2, & \text{if } (v_i, v_j) \in E \\ 1/n, & \text{if } (v_i, v_j) \notin E. \end{cases}$$

Consider an independent set S in G and let S' be the corresponding set of links in I. Observe that for any $\ell_v \in S'$, $a_{S'}(\ell_v) = (|S'| - 1) \cdot (1/n) < 1$, and thus S' is feasible. Similarly, in any feasible set of links there can be no pair that correspond to adjacent vertices in G. It follows that there is one-to-one correspondence between independent sets in G and feasible link sets in I. Hence, approximation algorithms for single-slot scheduling (scheduling) yield equivalent performance guarantees for the maximum independent set (minimum coloring) problem in graphs, respectively.

The last claim follows from the approximation hardness of graph coloring of [15] and [49]. ■

VIII. Conclusion

The main open question is to obtain a constant factor approximation to the scheduling problem (as erroneously claimed in [30]). Additionally, various parameter combinations are still open and deserve more research, e.g., multihop traffic, scheduling and routing, analog network coding, and stochastic fading models such as Rician fading.

References

[1] M. Andrews and M. Dinitz, "Maximizing capacity in arbitrary wireless networks in the SINR model: Complexity and game theory," in *Proc. 28th Annu. IEEE INFOCOM*, Apr. 2009, pp. 1332–1340.
[2] E. Ásgeirsson and P. Mitra, "On a game theoretic approach to capacity maximization in wireless networks," in *Proc. 30th Annu. IEEE INFOCOM*, 2011, pp. 3029–3037.
[3] E. I. Asgeirsson, M. M. Halldórsson, and P. Mitra, "A fully distributed algorithm for throughput performance in wireless networks," in *Proc. 46th Annu. CISS*, 2012, pp. 1–5.
[4] V. Auletta, L. Moscardelli, P. Penna, and G. Persiano, "Interference games in wireless networks," in *Proc. 4th WINE*, 2008, vol. 5385, LNCS, pp. 278–285.
[5] C. Avin, Y. Emek, E. Kantor, Z. Lotker, D. Peleg, and L. Roditty, "SINR diagrams: Towards algorithmically usable SINR models of wireless networks," in *Proc. 28th Annu. ACM SIGACT-SIGOPS PODC*, 2008, pp. 200–209.
[6] G. Brar, D. Blough, and P. Santi, "Computationally efficient scheduling with the physical interference model for throughput improvement in wireless mesh networks," in *Proc. 12th Annu. MobiCom*, 2006, pp. 2–13.
[7] G. S. Brar, D. M. Blough, and P. Santi, "The SCREAM approach for efficient distributed scheduling with physical interference in wireless mesh networks," in *Proc. 28th IEEE ICDCS*, 2008, pp. 214–224.
[8] D. Chafekar, V. Kumar, M. Marathe, S. Parthasarathy, and A. Srinivasan, "Cross-layer latency minimization for wireless networks using SINR constraints," in *Proc. 8th ACM MobiHoc*, 2007, pp. 110–119.
[9] M. Dinitz, "Distributed algorithms for approximating wireless network capacity," in *Proc. 29th Annu. IEEE INFOCOM*, Apr. 2010, pp. 1–9.
[10] T. A. ElBatt and A. Ephremides, "Joint scheduling and power control for wireless ad hoc networks," *IEEE Trans. Wireless Commun.*, vol. 3, no. 1, pp. 74–85, Jan. 2004.
[11] T. Erlebach and T. Grant, "Scheduling multicast requests in the SINR model," in *Proc. 6th ALGOSENSORS*, 2010, pp. 47–61.
[12] A. Fanghänel, S. Geulen, M. Hoefer, and B. Vöcking, "Online capacity maximization in wireless networks," in *Proc. 22nd Annu. ACM SPAA*, 2010, pp. 92–99.
[13] A. Fanghänel, T. Kesselheim, H. Räcke, and B. Vöcking, "Oblivious interference scheduling," in *Proc. 28th Annu. ACM PODC*, Aug. 2009, pp. 220–229.
[14] A. Fanghänel, T. Kesselheim, and B. Vöcking, "Improved algorithms for latency minimization in wireless networks," in *Proc. 37th ICALP*, Jul. 2009, pp. 447–458.

[15] U. Feige and J. Kilian, "Zero knowledge and the chromatic number," *J. Comput. Syst. Sci.*, vol. 57, pp. 187–199, 1998.
[16] 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 Annu. IEEE INFOCOM*, 2008, pp. 1013–1021.
[17] O. Goussevskaia, M. M. Halldórsson, R. Wattenhofer, and E. Welzl, "Capacity of arbitrary wireless networks," in *Proc. 28th Annu. IEEE INFOCOM*, Apr. 2009, pp. 1872–1880.
[18] O. Goussevskaia, T. Moscibroda, and R. Wattenhofer, "Local broadcasting in the physical interference model," in *Proc. 5th ACM SIGACT-SIGOPS Int. Workshop Found. Mobile Comput.*, Aug. 2008, pp. 35–44.
[19] O. Goussevskaia, Y. A. Oswald, and R. Wattenhofer, "Complexity in geometric SINR," in *Proc. 8th ACM MobiHoc*, 2007, pp. 100–109.
[20] O. Goussevskaia, L. Vieira, and M. Vieira, "Wireless multi-rate scheduling: From physical interference to disk graphs," in *Proc. 37th Annu. IEEE LCN*, 2012, pp. 651–658.
[21] O. Goussevskaia and R. Wattenhofer, "Complexity of scheduling with analog network coding," in *Proc. 1st ACM FOWANC*, May 2008, pp. 77–84.
[22] J. Gronkvist and A. Hansson, "Comparison between graph-based and interference-based STDMA scheduling," in *Proc. 2nd ACM MobiHoc*, 2001, pp. 255–258.
[23] M. Grossglauser and D. N. C. Tse, "Mobility increases the capacity of ad hoc wireless networks," *IEEE/ACM Trans. Netw.*, vol. 10, no. 4, pp. 477–486, Aug. 2002.
[24] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. Inf. Theory*, vol. 46, no. 2, pp. 388–404, Mar. 2000.
[25] M. M. Halldórsson, "Wireless scheduling with power control," *Trans. Algor.*, vol. 9, no. 1, p. 7, 2012.
[26] M. M. Halldórsson, S. Holzer, P. Mitra, and R. Wattenhofer, "The power of non-oblivious wireless power," in *Proc. 24th ACM-SIAM SODA*, 2013, to be published.
[27] M. M. Halldórsson and P. Mitra, "Wireless capacity with oblivious power in general metrics," in *Proc. 22nd ACM-SIAM SODA*, Jan. 2011, pp. 1538–1548.
[28] M. M. Halldórsson and P. Mitra, "Wireless capacity and admission control in cognitive radio," in *Proc. 31st Annu. IEEE INFOCOM*, 2012, pp. 855–863.
[29] M. M. Halldórsson and P. Mitra, "Wireless connectivity and capacity," in *Proc. 23rd ACM-SIAM SODA*, 2012, pp. 516–526.
[30] M. M. Halldórsson and R. Wattenhofer, "Wireless communication is in APX," in *Proc. 37th ICALP*, Jul. 2009, pp. 525–536.
[31] M. Hoefer, T. Kesselheim, and B. Vöcking, "Approximation algorithms for secondary spectrum auctions," in *Proc. 23rd Annu. ACM SPAA*, 2011, pp. 177–186.
[32] Q.-S. Hua and F. C. M. Lau, "The scheduling and energy complexity of strong connectivity in ultra-wideband networks," in *Proc. 9th MSWiM*, 2006, pp. 282–290.
[33] B. Katz, M. Völker, and D. Wagner, "Link scheduling in local interference models," in *Proc. 4th ALGOSENSORS*, 2008, pp. 57–71.
[34] T. Kesselheim, "A constant-factor approximation for wireless capacity maximization with power control in the SINR model," in *Proc. 22nd ACM-SIAM SODA*, Jan. 2011, pp. 1549–1559.
[35] T. Kesselheim, "Approximation algorithms for wireless link scheduling with flexible data rates," in *Proc. 20th Annu. ESA*, 2012, pp. 659–670.
[36] T. Kesselheim and B. Vöcking, "Distributed contention resolution in wireless networks," in *Proc. 24th DISC*, 2010, pp. 163–178.
[37] U. Kozat and L. Tassiulas, "Throughput capacity of random ad hoc networks with infrastructure support," in *Proc. MobiCom*, 2003, pp. 55–65.
[38] V. Kumar, M. Marathe, S. Parthasarathy, and A. Srinivasan, "Algorithmic aspects of capacity in wireless networks," in *Proc. SIGMETRICS*, 2005, pp. 133–144.
[39] S. Li, Y. Liu, and X.-Y. Li, "Capacity of large scale wireless networks under Gaussian channel model," in *Proc. 14th ACM MobiCom*, 2008, pp. 140–151.
[40] R. Maheshwari, S. Jain, and S. R. Das, "A measurement study of interference modeling and scheduling in low-power wireless networks," in *Proc. 6th SenSys*, 2008, pp. 141–154.
[41] T. Moscibroda, "The worst-case capacity of wireless sensor networks," in *Proc. 6th IPSN*, 2007, pp. 1–10.
[42] T. Moscibroda, Y. A. Oswald, and R. Wattenhofer, "How optimal are wireless scheduling protocols?," in *Proc. 26th Annu. IEEE INFOCOM*, 2007, pp. 1433–1441.
[43] T. Moscibroda and R. Wattenhofer, "The complexity of connectivity in wireless networks," in *Proc. 25th Annu. IEEE INFOCOM*, 2006, pp. 1–13.
[44] T. Moscibroda, R. Wattenhofer, and Y. Weber, "Protocol design beyond graph-based models," in *Proc. 5th HotNets*, Nov. 2006, pp. 25–30.
[45] T. Moscibroda, R. Wattenhofer, and A. Zollinger, "Topology control meets SINR: The scheduling complexity of arbitrary topologies," in *Proc. 6th ACM MobiHoc*, 2006, pp. 310–321.
[46] 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. 9th ACM MobiHoc*, 2008, pp. 91–100.
[47] Y. Shi, Y. T. Hou, S. Kompella, and H. D. Sherali, "Maximizing capacity in multihop cognitive radio networks under the SINR model," *IEEE Trans. Mobile Comput.*, vol. 10, no. 7, pp. 954–967, Jul. 2011.
[48] P.-J. Wan, O. Frieder, X. Jia, F. Yao, X. Xu, and S. Tang, "Wireless link scheduling under physical interference model," in *Proc. 30th Annu. IEEE INFOCOM*, 2011, pp. 838–845.
[49] D. Zuckerman, "Linear degree extractors and the inapproximability of max clique and chromatic number," in *Proc. 38th ACM STOC*, 2006, pp. 681–690.

Olga Goussevskaia received the doctorate degree in computer science from the Swiss Federal Institute of Technology (ETH Zurich), Zurich, Switzerland, in 2009.

She is an Assistant Professor with the Department of Computer Science of Federal University of Minas Gerais (UFMG), Belo Horizonte, Brazil. Her main research interests include modeling, algorithm design, and analysis of problems related to different kinds of networks, such as wireless communication networks, the Internet, the Web, social networks, and smart grids.

Magnús M. Halldórsson received the doctorate degree in computer science from Rutgers University, New Brunswick, NJ, USA, in 1991.

He then spent four years in Japan, first at the Tokyo Institute of Technology, Tokyo, followed by the Japan Advanced Institute of Science and Technology, Nomi, before returning to his home country of Iceland. He was with the University of Iceland, Reykjavik, Iceland, from 1995 to 2007, and with Reykjavik University, Reykjavik, Iceland, since 2007. He is now a Full Professor with the School of Computer Science, Reykjavik University. His research interests are in the design and analysis of algorithms for computationally hard problems, ranging from combinatorics and graph theory to wireless networking.

Roger Wattenhofer received the doctorate degree in computer science from ETH Zurich, Zurich, Switzerland, in 1998.

He is a Full Professor with the Information Technology and Electrical Engineering Department, ETH Zurich. From 1999 to 2001, he was in the US, first with Brown University, Providence, RI, then with Microsoft Research, Redmond, WA. He then returned to ETH Zurich, originally as an Assistant Professor with the Computer Science Department. His research interests are a variety of algorithmic and systems aspects in computer science and information technology, currently in particular distributed complexity, wireless networks, mobile computing, and social networking.