Topology Control Meets SINR: The Scheduling Complexity of Arbitrary Topologies

Thomas Moscibroda
Computer Engineering and Networks Laboratory
ETH Zurich
8092 Zurich, Switzerland
moscitho@tik.ee.ethz.ch

Roger Wattenhofer
Computer Engineering and Networks Laboratory
ETH Zurich
8092 Zurich, Switzerland
wattenhofer@tik.ee.ethz.ch

Aaron Zollinger
Department of Electrical Engineering and Computer Science
University of California
Berkeley, CA 94720, USA
zolling@eecs.berkeley.edu

ABSTRACT

To date, topology control in wireless ad hoc and sensor networks—the study of how to compute from the given communication network a subgraph with certain beneficial properties—has been considered as a static problem only; the time required to actually schedule the links of a computed topology without message collision was generally ignored. In this paper we analyze topology control in the context of the physical Signal-to-Interference-plus-Noise-Ratio (SINR) model, focusing on the question of how and how fast the links of a resulting topology can actually be realized over time.

For this purpose, we define and study a generalized version of the SINR model and obtain theoretical upper bounds on the *scheduling complexity of arbitrary topologies* in wireless networks. Specifically, we prove that even in worst-case networks, if the signals are transmitted with correctly assigned transmission power levels, the number of time slots required to successfully schedule all links of an arbitrary topology is proportional to the squared logarithm of the number of network nodes times a previously defined *static interference measure*. Interestingly, although originally considered without explicit accounting for signal collision in the SINR model, this static interference measure plays an important role in the analysis of link scheduling with physical link interference. Our result thus bridges the gap between static graph-based interference models and the physical SINR model. Based on these results, we also show that when it comes to scheduling, requiring the communication links to be symmetric may imply significantly higher costs as opposed to topologies allowing unidirectional links.

Categories and Subject Descriptors

F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems—*geometrical problems and computations, sequencing and scheduling*;
G.2.1 [**Discrete Mathematics**]: Combinatorics—*combinatorial algorithms*;
G.2.2 [**Discrete Mathematics**]: Graph Theory—*network problems*

General Terms

Algorithms, Theory

Keywords

Algorithmic analysis, interference, scheduling complexity, topology control, wireless ad hoc networks

1. INTRODUCTION

A common complaint within the networking community in general, and the ad hoc and sensor networking community in particular, is that there exists a wide chasm between theory and practice. What tends to be overlooked, however, is the fact that there are seemingly insurmountable gaps even within the theory community itself! A particularly distinct divide reveals itself with the question of the underlying communication models. On the one hand, the *graph-theoretically* oriented theory community has tried to come up with efficient network protocols for problems such as routing, clustering, data gathering, or topology control. Typically, these analytical results are derived in a type of static *graph model* that captures certain aspects of wireless communication while abstracting away many others. On the other hand, there has been a large body of theoretical work taking a more *communication* or *information theoretic* approach. Often based on the physical Signal-to-Interference-plus-Noise-Ratio (SINR) model of communication, these papers derive fundamental scaling laws that describe the theoretically achievable capacity in different modalities of communication and model assumptions. It is the goal of this paper to analytically identify previously unknown ties between these two approaches.

Specifically, we study *topology control*—until now an inherently graph-theoretic notion—and analyze its impact in the information-theoretic SINR model. In a very general sense, topology control in wireless ad hoc and sensor networks can be considered the task of—given a network connectivity graph—computing a subgraph with specific desired properties, such as connectivity, short stretches, sparsity, low interference, or low node degree. Accordingly, there has been considerable research effort towards achieving and combining more and more of these properties [4, 18, 19, 20, 27, 30, 33].

All these approaches have in common, however, that they model wireless networks as *static graphs*, hence neglecting one of the most crucial aspects of wireless communication: Eventually messages—or, even more exactly, radio signals—will have to be sent over these static links in the topologies selected by a topology control algorithm, that is, the static graph of communication links must be *scheduled* on the physical layer. In general, not all of these messages can be sent simultaneously, as mutual interference may prevent proper message reception. Hence, what has been inherently lacking in the study of topology control so far is the notion of *time* required to actually *realize* the selected links, that is, to successfully transmit messages over them.

In particular, it has not been clear—even from a theoretical point of view—whether the graph-theoretic measures of topology control really bear any significance when it comes to actually scheduling messages in an SINR environment. In this paper, we demonstrate and prove the existence of fundamental theoretical ties between topology control and the theoretically achievable efficiency of scheduling protocols.

Consider an arbitrary topology computed by a topology control algorithm, or more generally a set of communication requests. In one time slot, only a subset of all the desired communication links can be scheduled in parallel; and in every subsequent time slot, a subset of the remaining unscheduled links may be scheduled, until finally all links are scheduled. Against this background, it is clearly advantageous to find a short schedule, ideally one with minimal length. This allows for higher bandwidth and, ultimately, higher throughput. We call this measure describing the minimal amount of time required to *physically* establish a set of communication requests or a desired network topology the *scheduling complexity* of the topology.

In this paper, we show that the initially mentioned gap in the field of topology control is neither inherent nor inevitable. Specifically, we study two problems: First, what is the scheduling complexity of arbitrary topologies, that is, what is the time required to physically schedule an arbitrary set of requests? And second, which fundamental graph-theoretic measures determine the scheduling complexity? In other words, what properties should a topology have such that provably efficient scheduling becomes possible? Strongly coupled with the first problem is the question of how we can actually find a schedule (and corresponding transmission power assignments) which enable a fast realization of a topology. Interestingly, studying these questions also reveals the intricacy of the interplay between scheduling complexity and transmission power assignment.

More specifically, we make the following contributions:

- As the main technical result of this paper we present a scheduling algorithm which assigns transmission power levels to the network nodes and schedules all links of an arbitrary network topology. The subsequent analysis reveals that this algorithm computes a schedule of length $O(I_{in} \cdot \log^2 n)$, where n is the number of nodes and I_{in} is a previously defined *static interference measure* that reflects a static property of wireless network graphs. Intriguingly, in spite of its static nature, the measure I_{in} appears to play an important role concerning the scheduling complexity of arbitrary network topologies. Our result proves that topology control algorithms choosing good static topologies—topologies with small I_{in}—allow for faster scheduling on the physical SINR communication level. Note that our results also yield the first general scaling laws on the scheduling complexity of an arbitrary set of communication requests.
- We show that there exists an inherent gap with respect to link symmetry: Network topologies preserving connectivity of the given communication network using *unidirectional* links have significantly lower I_{in} values and can therefore be scheduled much faster than connectivity-preserving topologies using exclusively *symmetric* links. This result sheds new light on the question of practicality of directed as opposed to symmetric links in wireless ad hoc and sensor networks: It shows that demanding communication links to be symmetric theoretically incurs a high overhead when it comes to scheduling.
- In this paper we explicitly analyze topology control with special emphasis on the physical definition of interference, or more specifically the Signal-to-Interference-plus-Noise-Ratio (SINR). Simply expressed, this characterization of interference reflects the fact that a radio signal can be correctly decoded by the intended receiver only if the ratio between the sensed power of the actual signal to be received and the sum of all power levels experienced due to other signals concurrently transmitted (plus an ambient noise power level) is above a certain hardware-dependent threshold. In reality, obstacles to signal propagation can shadow, reflect, scatter, and diffract radio signals. In order to capture these effects, we extend the SINR approach by a *generalized signal propagation model*; in particular, we allow the received signal strength to deviate from the theoretically computed value in the absence of any obstacles by a constant factor which is reflected in the subsequent analysis.
- Finally, combining our novel scheduling algorithm with a known low-interference topology control algorithm for strong connectivity with $I_{in} \in O(\log n)$, we show that in the physical model of communication, the scheduling complexity of connectivity is $O(\log^3 n)$, thus improving the bound given in [25] by a logarithmic factor.

As a general remark, we would like to point out that we are aware of the fact that—primarily due to its missing locality, as shown later in the paper,—our scheduling algorithm does not lend itself to direct implementation as a network protocol. Instead, we believe that the algorithm is of great theoretical interest, as it proves a first upper bound on the *scheduling complexity of arbitrary topologies* in wireless networks. Specifically, it explicitly shows that fast scheduling of arbitrary topologies is theoretically possible even in worst-case wireless networks when assigning proper power levels to the nodes.

The paper is organized as follows: After giving an overview of related work in the following section, we formally introduce the considered communication model in Section 3 and define the concepts of scheduling complexity and I_{in} interference in Section 4. In Sections 5 and 6 we present the scheduling algorithm for arbitrary network topologies and its analysis, respectively. Section 7 finally demonstrates the existence of a gap with respect to scheduling complexity between topologies allowing unidirectional links and topologies containing exclusively symmetric connections. Section 8 concludes the paper.

2. RELATED WORK

Early work that can be considered precursors of *topology control* focused on the question of the required node density to achieve connectivity when randomly placing nodes [14, 31]. A first generation of topology control in its modern sense [15, 27, 29] adopted structures from the field of computational geometry and mainly aimed at preserving energy-efficient paths or computing planar subgraphs for geographic routing [5, 17]. In a second wave of research, constructions were proposed which are based on local information and simultaneously reconcile several properties, such as planarity, the spanner property, or constant-bounded node degree [18, 19, 20, 33]. Other approaches tried to build on minimal assumptions about the capabilities of nodes and signal propagation characteristics [34].

All these contributions have in common that they do not consider interference or—if interference is mentioned as an issue—at most implicitly. In contrast, recent work [6, 22, 24, 32] has modeled interference *explicitly* as a graph property, similar to other topology control properties. The main deficiency of this approach is however that—also staying in the tradition of topology control—interference is modeled statically based on graphs, without specifically considering the physical consequences of simultaneously transmitted signals.

An approach that does consider the time aspect of signals and in particular message scheduling is *coloring* in wireless networks. According to the definition of the problem of correct message reception—for instance such that a signal is successfully received if only one of the receiver's neighbors transmits at every instant of time—reduces to coloring problems of various types [21, 23, 28]. Modeling the interference issue as coloring problems is however—compared to the physical SINR—overly pessimistic on the one hand in that it does not reflect the fact that even nearby communication is tolerable if it takes place at a sufficiently low power level; on the other hand such an approach is too optimistic since also weak signals transmitted simultaneously by many remote sources can build up considerable interference [1, 12].

The computation of efficient schedules in the SINR model has been studied in previous work in various flavors. The work of [3] proposes a mathematical programming formulation for deriving optimal schedules. However, the resulting formulations are infeasible from a computational point of view. The authors then propose a heuristic based on a so-called column-generation approach, which they show to produce fast schedules in practical scenarios. Finally, it is shown in [3] that the problem of deriving optimal schedules is NP-hard, even in a much more restricted model. The work of [2, 16] also derives mathematical programming formulations and investigates the impact of power assignments to nodes on the achievable throughput capacity. As these solutions are based on complex non-linear program formulations that can only be solved in exponential time, they provide little insight into the structural property of scheduling in SINR models.

In [9, 11], various protocols for scheduling in SINR-based models were proposed and evaluated under different traffic and random node distribution models. These protocols being evaluated by means of *simulation*, none of them provides theoretical bounds on the efficiency in a worst-case sense. The algorithms in [7, 8] study the problem of finding schedule and power control policies that minimize the total average transmission power in the wireless multi-hop network. The algorithm in [8], for instance, is based on guaranteeing a certain "spatial reuse" distance between all pairs of simultaneously transmitting nodes. As shown in [25], such an approach inherently cannot yield competitive results in worst-case networks. That is, all these protocols may produce schedules such that in certain networks and for certain request sequences, these algorithms may compute a schedule that is significantly worse than the optimal solution. Moreover, none of these contributions provide bounds on the fundamental scheduling complexity in wireless networks.

Recently, an efficient scheduling algorithm in the SINR model has been presented in [25]. This algorithm finds a schedule of length $O(\log^4 n)$ for efficiently scheduling a set of links that combines for a strongly connected topology. In this paper, we improve and generalize the approach of [25] and provide provable bounds on the scheduling complexity for all sets of requests and topologies.

3. GENERALIZED PHYSICAL MODEL

For our analysis we model a wireless network as a set of nodes $X = \{x_1, \ldots, x_n\}$ that are arbitrarily located in the Euclidean plane. The Euclidean distance between two nodes $x_i, x_j \in X$, is denoted by $d(x_i, x_j)$. For a directed communication link (x_i, x_j), $d_{ij} = d(x_i, x_j)$ denotes the distance between its endpoints. The *ball* $B(x_i, r)$ of radius r around node x_i contains all nodes $x_j \in X$ for which $d(x_i, x_j) \leq r$. For simplicity and without loss of generalization, we assume that the minimal distance between any two nodes in X is 1. Furthermore we define the Euclidean diameter Δ to be the largest distance between two nodes.

The core aspect of the communication model underlying our analysis is the description of the circumstances under which a message is correctly received by its intended recipient. In the *Signal-to-Interference-plus-Noise-Ratio* (SINR) model (also called *physical model* in [13]), the successful reception of a transmission depends on the received signal strength, the interference caused by simultaneously transmitting nodes, and the ambient noise level. Let P_r be the received power of a signal sent to a node x_r, and denote by I_r the interference power generated by other nodes in the network. Finally, let N be the ambient noise power level. Then, a node x_r receives a transmission if and only if $\frac{P_r}{N+I_r} \geq \beta$, where β is the minimum signal-to-interference-ratio that is required for a message to be successfully received at x_r.[1]

In wireless networks, the value of the received signal power P_r of a signal is a decreasing function of the distance $d(x_s, x_r)$ between the transmitter node x_s and the receiver node x_r. Theoretically, the received signal power P_r can be modeled as decaying with distance $d(x_s, x_r)$ as

$$P_r = \frac{P_s}{d(x_s, x_r)^\alpha},$$

where P_s is the sending power of the transmitting node. The so-called path-loss exponent α is a constant between 2 and 6 and depends on external conditions of the medium, as well as the exact sender-receiver distance. As common, we assume that $\alpha > 2$ [13].

In practice, the received signal power may however deviate from the above theoretical bound for various reasons.

[1]Our analysis can be generalized such that every node x_i can define its own β_i.

On the one hand, the signal-emitting characteristics of antennas may not be perfectly omni-directional. Moreover, shadowing, reflection, scattering, and diffraction caused by the presence of obstacles to wireless signal propagation may have an impact on the signal power actually sensed at the receiver.

In order to better account for some of these aspects of wireless communication, we define and study the following slight generalization of the physical model, which we call the *generalized physical model*. In this generalized physical model with parameter θ, the received signal power (as well as the interference caused by simultaneously transmitting nodes) can deviate arbitrarily from the theoretically received power by a factor of θ. Formally, if $P_r(x_s)$ is defined to be the actual received power of a signal transmitted by node x_s as sensed by the receiving node x_r, the generalized physical model states that $P_r(x_s)$ is in the range

$$\frac{1}{\theta} \cdot \frac{P_s}{d(x_s, x_r)^\alpha} \;\leq\; P_r(x_s) \;\leq\; \theta \cdot \frac{P_s}{d(x_s, x_r)^\alpha}.$$

Note that the model leaves open the exact received signal power, and hence algorithms working in the generalized physical model must be robust enough to cope with arbitrary (even worst-case) deviations within the stated bounds. Clearly, for $\theta = 1$, the generalized physical model is equivalent to the standard physical model.

As for the notation in this paper, we occasionally use the formulation $I_r(x_i) = P_r(x_i)$ in order to emphasize that the signal power transmitted by a node x_i other than the intending sender is perceived at x_r as interference. In summary, if $P_r(x_s)$ and $I_r(x_i)$ are the received power levels sensed by node x_r in a specific time slot, a signal transmitted by a node $x_s \in X$ is successfully received by x_r if

$$\frac{P_r(x_s)}{N + \sum_{x_i \in X \setminus \{x_s\}} I_r(x_i)} \;\geq\; \beta. \tag{1}$$

Finally, the *total interference* I_r experienced by a receiver x_r is the sum of the interference power values created by all nodes in the network (except the intending sender x_s), that is, $I_r = \sum_{x_i \in X \setminus \{x_s\}} I_r(x_i)$.

4. SCHEDULING COMPLEXITY AND INTERFERENCE

In this section, we formally introduce the concepts of the *scheduling complexity of arbitrary topologies* and of the (static) interference I_{in} of a set of communication requests.

Following the example of [13], we assume without loss of generality that transmissions are slotted into synchronized time slots of equal length. In each time slot t, a node x can either transmit or not transmit a message. If it transmits, it chooses a power level $P_x > 0$ that must be sufficiently large in order to reach the receiver. A *power assignment* determines the power level chosen by each node in a certain time slot. Formally, a power assignment ϕ_t is a function $\phi_t : X \mapsto \mathbb{R}^+$ which maps every node in the network to a power level. We denote by $\phi_t(x_i)$ the power level of node x_i in time slot t. If a node is not scheduled to transmit in this time slot, then $\phi_t(x_i) = 0$. Whenever the considered time slot t is clear from the context, we also use the notational shortcut $P_i = \phi_t(x_i)$. A *schedule* $\mathcal{S} = (\phi_1, \ldots, \phi_{T(\mathcal{S})})$ is a sequence of $T(\mathcal{S})$ power assignments, where ϕ_i denotes the power assignment in time slot i. Finally, we call $T(\mathcal{S})$ the *length* of schedule $\mathcal{S}$. That is, a schedule $\mathcal{S}$ of length $T(\mathcal{S})$ determines the power level P_i for every node $x_i \in X$ for $T(\mathcal{S})$ consecutive time slots.

A construction resulting from the operation of a topology control algorithm is defined as a graph containing links between the network nodes. Every such link can be considered a communication request to send a message over the corresponding link. Consequently, the task of the algorithm we present in this paper is to schedule a set of given communication requests such that the corresponding messages are successfully received. Each such *request* γ_{ij} denotes a directed link (x_i, x_j) and indicates that node x_i is to successfully transmit a message to node x_j. The set of all communication requests is called Γ.

To capture the minimal amount of time required to schedule all requests $\gamma_{ij} \in \Gamma$, we use the *scheduling complexity* defined in the sequel.

DEFINITION 4.1. *Consider a time slot t and a power assignment ϕ_t. We say that a directed link (x_i, x_j) is* successfully scheduled *in time slot t if x_j receives a message from x_i according to the SINR Inequality* (1).

Let L_t be the set of all successfully scheduled links in time slot t. Our goal is that after as few time slots as possible the union of all sets L_t forms the given network topology. We therefore define the scheduling problem for a given network topology as follows:

DEFINITION 4.2. *Let Γ be the set of communication requests of a given topology. The scheduling problem for Γ is to find a schedule $\mathcal{S}$ of minimal length $T(\mathcal{S})$ such that the union of all successfully transmitted links $\bigcup_{t=1}^{T(\mathcal{S})} L_t$ equals Γ.*

Finally, we define the scheduling complexity of an arbitrary topology, that is of an arbitrary set of communication requests Γ.

DEFINITION 4.3. *The* scheduling complexity $T(\Gamma)$ *of a set of communication requests Γ is the minimal number of time slots T such that there exists a valid schedule $\mathcal{S}$ of length $T = T(\mathcal{S})$.*

Given an arbitrary topology or a set of communication requests Γ, the scheduling complexity of Γ reflects how fast all requests in Γ can theoretically be satisfied (that is, when scheduled by an optimal MAC-layer protocol). The scheduling complexity of arbitrary topologies being a fundamental measure in wireless networks, we are interested in obtaining scaling laws that describe the asymptotic behavior of the scheduling complexity as the network grows. Also, we seek algorithms that achieve good performance with regard to the scheduling complexity since such algorithms would come close to an "optimal" MAC-layer protocol.

In general, this *scheduling complexity of arbitrary topologies* may not allow for a concise expression better than the trivial bound of n, which is achieved if nodes are scheduled one after the other. In fact, it is possible to construct examples—as the one depicted in Figure 1—with n requests in which the scheduling complexity grows linearly, $\Omega(n)$, even if all sender and receiver pairs are different.

In view of these trivial tight bounds, it is more interesting to express the scheduling complexity of arbitrary topologies in terms of additional properties besides n. In this regard, it is particularly intriguing to derive the scheduling complexity of arbitrary requests in dependence on a *graph theoretic*

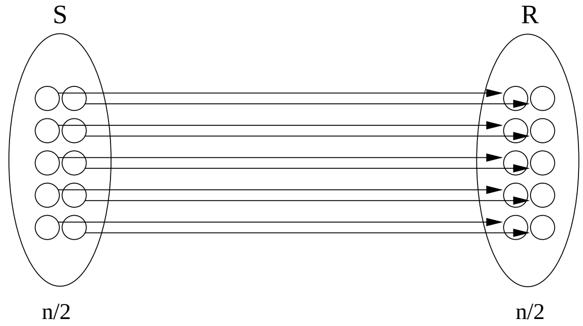

Figure 1: The scheduling complexity of the requests in this example is $n/2$. In each time slot, at most one link can be scheduled successfully.

measure that captures the inherent complexity of scheduling in wireless multi-hop networks. If the scheduling complexity in wireless networks could be expressed by means of a simple, intuitive graph theoretical measure, this would serve as a legitimation for studying graphs when reasoning about scheduling in wireless networks. In a sense, such a correspondence would thus help bridging the gap between communication and information theoretical SINR models and algorithmic graph models.

As a second basic concept we therefore introduce the static interference measure I_{in} of a set of communication requests (that is a network topology), originally introduced in a graph-theoretic context in [32]. This interference measure is based on the question of how many other nodes can potentially disturb a node in the network.[2] Note that the definition of I_{in} is independent of the SINR model and argues using *circular transmission ranges*. Also, we do not assume that actual signal propagation behaves according to this interference measure; this measure is introduced as a static graph property whose significance will become clearer in the analysis of our scheduling algorithm.

With the assumption that the network nodes use perfectly omnidirectional antennas, the maximal disk $\mathcal{D}_i$ of a node x_i represents the transmission range such that all intended receivers of x_i are reached, or, in other words, the disk covering all nodes that are potentially affected by message transmission of x_i to one of its intended receivers. Then the interference of a node x_j is defined as the number of other nodes that potentially affect message reception at node x_j:

Definition 4.4. *Given a set of communication requests Γ, the in-interference of a node $x_j \in X$ is defined as*

$$I_{in}(x_j) = |\{x_i | x_i \in X \setminus \{x_j\}, x_j \in \mathcal{D}_i)\}|.$$

In other words, the interference I_{in} of a node x_j represents the number of nodes covering x_j with their disks induced by their transmission ranges set to a value as to reach all their intended receivers. Note that the in-degree of a node in a given topology Γ does not correspond to its interference; the in-degree merely forms a lower bound for its interference since it can be "covered" by disks of non-neighboring nodes. The node-level interference defined so far is now extended to an interference measure for Γ:

Definition 4.5. *Given a set of nodes X and a set of communication requests Γ, the in-interference of Γ is $I_{in}(\Gamma) = \max_{x_i \in X} I_{in}(x_i)$.*

We end this section with some helpful notation used in the remainder of the paper. If Γ is the set of all communication requests to be scheduled, Γ_i denotes the set of requests for which node x_i is the sender, formally $\Gamma_i = \{\gamma_{ij} \in \Gamma\}$. The set of intended receivers to which a node x_i is to successfully transmit is $R_i = \{x_j \mid \gamma_{ij} \in \Gamma\}$. The Euclidean distance of a node x_i to its most distant intended receiver is the *radius* of the node x_i: $r_i = \max_{x_j \in R_i} d(x_i, x_j)$. In case there is no communication request for which node x_i is the sender, that is, if x_i has no intended receiver at all, we define $r_i = 0$. Finally, the *maximal disk* of a node x_i, formally defined as $\mathcal{D}_i = B(x_i, r_i)$, is the smallest disk centered at x_i that contains all intended receivers, that is, $x_j \in \mathcal{D}_i$ for all $x_j \in \Gamma_i$ (also see the above definition of I_{in}).

5. SCHEDULING ALGORITHM

In this section, we present the paper's main technical contribution, an algorithm that achieves the first known bound on the scheduling complexity of an arbitrary set of requests Γ.

Theorem 5.1. *Given an arbitrary network and a set of communication requests Γ with in-interference I_{in}, all requests $\gamma_{ij} \in \Gamma$ can be successfully scheduled in time $O(I_{in} \cdot \log(n\theta^2) \cdot (\theta^{\frac{4}{\alpha}} + \log n))$. That is, the scheduling complexity $T(\Gamma)$ of a topology with in-interference I_{in} is*

$$T(\Gamma) \in O(I_{in} \cdot \log(n\theta^2) \cdot (\theta^{\frac{4}{\alpha}} + \log n)).$$

We prove the theorem in Section 6 by showing that, given a set of requests Γ, Algorithm 1 computes a schedule of the respective length in which all requests are successfully scheduled. In this section, we give an intuitive overview of Algorithm 1.

The task faced by any MAC-layer or scheduling protocol is twofold. The protocol not only decides which node transmits in which time slot, it also assigns proper power levels. In [25], it was shown that particularly the second task—assigning transmission powers—is a non-trivial task, and MAC-layer protocols that employ uniform or linear power assignments fail to achieve reasonable performance. Therefore, in order to achieve the scheduling complexity stated in Theorem 5.1, Algorithm 1 employs a power-assignment policy that favors nodes with small radii over nodes with large radii. Furthermore, the schedule computed by the algorithm has the characteristic that every node x_i establishes links to all its intended receivers R_i in a single time slot by successfully broadcasting within its radius r_i. However, the transmission power assigned to node x_i is not proportional to r_i^{α}, which would inherently lead to a very slow schedule.

Consider for example the instance depicted in Figure 2. For the following considerations we momentarily neglect the presence of ambient noise N for simplicity. If all of the four

[2]The notation I_{in} reflects the fact that this measure is based on how many other nodes can disturb a given network node. Similar interference measures have been defined assuming in a sense an antipodal perspective by asking how many other nodes a given node can disturb. It can be shown that the results of our analysis in the remainder of the paper asymptotically also hold for such I_{out}-interference measures.

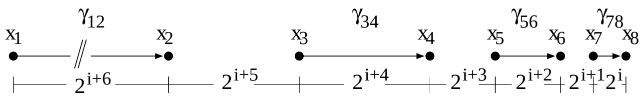

Figure 2: An example in which scheduling many requests in parallel requires a non-trivial power assignment policy.

links are to be scheduled individually, setting the transmission power of the respective sender x_i to $\beta \cdot r_i^\alpha$ is sufficient for the signal to be correctly received. Should however more than one link be scheduled simultaneously, the situation becomes more intricate. If link γ_{12} is scheduled successfully, the signal power received by x_2 is at least β, and hence the intended receivers of all other links face an interference of at least $\beta/2^\alpha$. It follows that if we want some of the smaller links to be scheduled simultaneously, every sender x_i of any one of these links must transmit at a power that is at least by a factor $\beta/2^\alpha$ greater than r_i^α. The problem is the interference created by long link cascades, that is, if all four links are scheduled in the same time slot, the third and fourth senders (x_5 and x_7) must transmit with a power of at least $\beta^3/2^{2\alpha} \cdot r_5^\alpha$ and $\beta^4/2^{3\alpha} \cdot r_7^\alpha$, respectively, in order to guarantee successful simultaneous reception.[3] This observation weighs particularly heavily for the following reason: If we want fast, say, polylogarithmic schedules, there must exist time slots in which at least $n/\log^c n$ nodes are scheduled simultaneously for some constant c. The dependence of the chosen transmission power on other simultaneously scheduled links—together with the requirement to schedule relatively many links at the same time—shows that every provably efficient scheduling protocol must inevitably employ a complex and sophisticated power assignment strategy. In Algorithm 1, the transmission power of a node x_i is scaled by a factor of $(3n\beta\theta^2)^{\tau(x_i)}$, where $\tau(x_i)$ is a value that reflects the relative position of r_i in an ordering of all radii.

Unfortunately, scaling up the transmission powers of nodes with small radii in turn entails new problems. Specifically, since a node x_i with small radius r_i now transmits at a transmission power that is high relative to the length r_i of its longest communication link, x_i may cause significant interference at a receiver x_j even if the distance $d(x_i, x_j)$ is exponentially larger than r_i.

It is important to observe that this *non-locality of interference* is in stark contrast to all generally studied graph-based interference models. In fact, the above intuition shows that *fast scheduling in the SINR model* is inherently a *non-local* task and simple local approaches that typically work in graph-based models fail to produce reasonable solutions.

In the sequel, we describe the algorithm on a more technical level. At the outset of Algorithm 1, the algorithm partitions the set of nodes X into at most $\lfloor \log \Delta \rfloor + 1$ possibly empty disjoint sets $\mathcal{S} = S_0, \ldots, S_{\lfloor \log \Delta \rfloor}$. Each such set S_i contains every node $x_j \in X$ with radius $2^i \le r_j < 2^{i+1}$. If no such node exists, the set S_i remains empty. In the next step, the algorithm removes all these empty sets and

[3] For our illustration we assume that $\beta > 2^\alpha$; otherwise the node distances can be adapted to produce a similar situation in which the nodes' transmission powers are disproportionate compared to their radii.

Algorithm 1 Scheduling Algorithm

Input: - An arbitrarily located set of nodes X
- A set of communication requests Γ

Output: A schedule $\mathcal{S}_{ALG}$ in which all requests $\gamma_{ij} \in \Gamma$ are successfully scheduled

Define two constants ν and μ such that $\nu := 4N$ and $\mu := 1 + 2^{\frac{8}{\alpha}+2} \sqrt[\alpha]{\frac{\beta(\alpha-1)}{\alpha-2}}$; $t := 1$;
Partition X into sets $\mathcal{S} = S_0, \ldots, S_{\lfloor \log \Delta \rfloor}$ such that S_i contains all nodes x_j with $2^i \le r_j < 2^{i+1}$;
Delete all empty sets $S_i \in \mathcal{S}$ and rename $\mathcal{S}$ such that S_i is the i^{th} non-empty set in *decreasing* order of the radii of the contained nodes;
for $k = 1$ **to** $\lceil \log(3n\beta\theta^2) \rceil$ **do**
  Let $\mathcal{F}_k$ be the union of all sets $S_{m\lceil \log(3n\beta\theta^2)\rceil + k} \in \mathcal{S}$ for $m \in \mathbb{N}_0$;
  for each $x_i \in \mathcal{F}_k$ **do**
    $\tau(x_i) := \chi$, where $x_i \in S_\ell$ and S_ℓ is the χ^{th} set in $\mathcal{F}_k$ (in decreasing order of radii);
  end for
  while not all links with intending sender in $\mathcal{F}_k$ have been scheduled **do**
    $E_t := \emptyset$;
    Consider all nodes $x_i \in \mathcal{F}_k$ in decreasing order of r_i:
      if *allowed*(x_i, E_t) **then** $E_t := E_t \cup \{x_i\}$;
    Schedule all $x_i \in E_t$ in time slot t and assign x_i a transmission power of $P_i = \nu(3n\beta\theta^2)^{\tau(x_i)} \cdot r_i^\alpha$;
    Remove all scheduled senders ($\mathcal{F}_k := \mathcal{F}_k \setminus E_t$);
    $t := t + 1$;
  end while
end for

allowed$(\mathbf{x_i}, \mathbf{E_t})$
for each $x_j \in E_t$ **do**
  $\delta_{ij} := \tau(x_i) - \tau(x_j)$;
  if $\tau(x_i) = \tau(x_j)$ and $\mu\theta^{\frac{2}{\alpha}} \cdot r_i > d(x_i, x_j)$
    return false
  else if $r_i \cdot (3n\beta\theta^2)^{\frac{\delta_{ij}+1}{\alpha}} + r_j > d(x_i, x_j)$
    return false
end for
return true

renames the remaining non-empty sets such that S_i is the i^{th} non-empty set in decreasing order of the radii of the contained nodes, for $i = 1, 2, \ldots$ (see Figure 3). In the resulting partition, the radius of all nodes in the same set are still within a factor of 2, whereas the radius of two nodes in sets S_i and S_{i+1} may differ by an arbitrarily large factor if many empty sets were deleted between S_i and S_{i+1}.

The task of each of the $\lceil \log(3n\beta\theta^2) \rceil$ iterations of the subsequent for-loop is to schedule a subset of all the links. In particular, in the k^{th} iteration of the loop, nodes in the sets $S_{m\lceil \log(3n\beta\theta^2)\rceil + k}$ are scheduled for all integers m. All these nodes form the set of nodes $\mathcal{F}_k$ that is to be scheduled in the k^{th} iteration. The reason for partitioning the entire set of requests into $\lceil \log(3n\beta\theta^2) \rceil$ subsets is to guarantee (cf. Lemma 6.6) that two nodes scheduled in the same time slot either have almost the same radius (when they are in the same set of the partition $\mathcal{S}$) or their radii differ significantly. We will use this property in the key Lemmas 6.7 and 6.8.

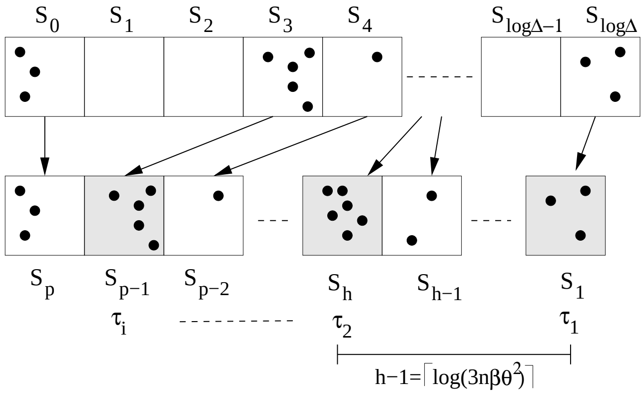

Figure 3: Naming of the sets in partition $\mathcal{S}$. The uppermost row represents the names given in Line 2 of Algorithm 1, the second row reflects the situation after the renaming in Line 3, and the τ_i at the bottom stand for the power scaling factors assigned in Line 7.

It may not be possible to schedule all links in $\mathcal{F}_k$ in a single time slot. Even scheduling the links of $\mathcal{F}_k$ alone turns out to be a challenging task, as we show in the following. In Lines 6 and 7 of Algorithm 1, each node $x_i \in \mathcal{F}_k$ designated to be scheduled in this iteration of the for-loop determines its $\tau(x_i)$ value. If node x_i is in the set with the largest radii of the sets selected in $\mathcal{F}_k$, $\tau(x_i)$ is set to 1. Or generally speaking, if a node x_i is in set S_ℓ and S_ℓ is the χ^{th} set (still in decreasing order of radii) of all sets forming $\mathcal{F}_k$, then $\tau(x_i) := \chi$ (cf. Figure 3). Intuitively, nodes with small radii have a large $\tau(x_i)$, while nodes with large radii have a small $\tau(x_i)$. In other words, $\tau(x_i)$ is a power scaling factor reflecting the fact that—as illustrated in the example of Figure 2—nodes with small radii may have to send with disproportionately high transmission powers compared to nodes with larger radii.

At the heart of Algorithm 1 is the while-loop which schedules all nodes in $\mathcal{F}_k$ using essentially at most $O(I_{in} \cdot \log n)$ time slots, as shown later in Lemmas 6.5 and 6.8. The set of nodes scheduled in parallel in time slot t is denoted by E_t; at the end of each iteration, all nodes that were scheduled are removed from $\mathcal{F}_k$ (Line 14). The selection of nodes for E_t proceeds as follows. The nodes in $\mathcal{F}_k$ are considered one by one in decreasing order of r_i. When considering such a node x_i, the algorithm checks whether scheduling x_i conflicts with previously selected nodes in E_t (that is nodes with larger radius) in the sub-procedure *allowed*(x_i, E_t). This procedure returns true if and only if

1. All (outgoing) links of x_i can be successfully scheduled in spite of the interference created by the nodes already in E_t.

2. All senders in E_t can still successfully transmit in spite of the additional interference caused by x_i.

As shown in Section 6, these two properties can be guaranteed by requiring that for all $x_j \in E_t$ it holds that $d(x_i, x_j) < r_i \cdot \mu\theta^{2/\alpha}$ if $\tau(x_i) = \tau(x_j)$ or $d(x_i, x_j) < r_i \cdot (3n\beta\theta^2)^{(\delta_{ij}+1)/\alpha} + r_j$ otherwise. The constant μ is set to a large enough value as to ensure low interference.

At this point, we would like to emphasize that Algorithm 1 is not primarily intended for being employed as a practical network protocol. Besides being rather complex, the necessity that each node knows its value $\tau(x_i)$ for the purpose of determining its own transmission power renders the algorithm non-trivial to implement in a distributed or local way. Specifically, $\tau(x_i)$ depends on the relative position of the set S_ℓ in $\mathcal{S}$ to which x_i belongs. However, the algorithm proves that, theoretically, even complex topologies can be scheduled efficiently also in large-scale worst-case networks provided that I_{in} is small.

6. ANALYSIS

The analysis of Algorithm 1 consists of two parts. First, we need to guarantee that the obtained schedule $\mathcal{S}_{ALG}$ is valid, that is, all links are successfully received by the intended receivers. Second, we will prove that the number of time slots required in the worst case does not exceed $O(I_{in} \cdot \log(n\theta^2) \cdot (\theta^{\frac{4}{\alpha}} + \log n))$. We start with correctness.

In order to guarantee that each node x_s scheduled by the algorithm in a time slot t is capable of successfully sending to all its intended receivers, we bound the total interference accrued at each of these receivers. For this purpose, we first bound the interference created by simultaneously transmitting nodes y_i with significantly larger radii.

LEMMA 6.1. *Consider a time slot t_s in which the algorithm schedules a node x_s for transmission. It holds for all intended receivers $x_r \in R_s$ and for any simultaneously transmitting node $y_i \in X \setminus \{x_s\}$ with $\tau(y_i) < \tau(x_s)$ that*

$$I_r(y_i) \;\leq\; \nu\theta(3n\beta\theta^2)^{\tau(x_s)-1}.$$

PROOF. Because the radius of y_i must be significantly larger than the radius of x_s, y_i was already in E_t at the time *allowed*(x_s, E_t) evaluated to true and the algorithm selected node x_s for scheduling. Consequently, the distance $d(x_s, y_i)$ must have been at least

$$d(x_s, y_i) \;\geq\; r_s \cdot (3n\beta\theta^2)^{\frac{\delta_{si}+1}{\alpha}} + r_i \;>\; r_s + r_i,$$

where r_i is y_i's radius, and therefore $d(x_r, y_i) > r_i$. The interference caused by y_i at x_r is consequently at most

$$\begin{aligned} I_r(y_i) \;&\leq\; \theta \cdot \frac{P_i}{d(x_r, y_i)^\alpha} \;\leq\; \frac{\theta\nu(3n\beta\theta^2)^{\tau(y_i)} r_i^\alpha}{r_i^\alpha} \\ &=\; \nu\theta(3n\beta\theta^2)^{\tau(y_i)} \;\leq\; \nu\theta(3n\beta\theta^2)^{\tau(x_s)-1}, \end{aligned}$$

which concludes the proof. □

In the following lemma we show that also the interference caused by concurrently transmitting nodes with significantly *smaller* radii than r_s is bounded.

LEMMA 6.2. *Consider a time slot t_s in which the algorithm schedules a node x_s for transmission. It holds for all intended receivers $x_r \in R_s$ and for any simultaneously transmitting node $y_i \in X \setminus \{x_s\}$ with $\tau(y_i) > \tau(x_s)$ that*

$$I_r(y_i) \;\leq\; \nu\theta(3n\beta\theta^2)^{\tau(x_s)-1}.$$

PROOF. The interference $I_r(y_i)$ incurred by a node y_i at x_r is at most

$$I_r(y_i) \;\leq\; \theta \cdot \frac{\nu(3n\beta\theta^2)^{\tau(y_i)} \cdot r_i^\alpha}{d(y_i, x_r)^\alpha}.$$

Assume for contradiction that there exists a node y_i with $\tau(y_i) > \tau(x_s)$ and $I_r(y_i) > \nu\theta(3n\beta\theta^2)^{\tau(x_s)-1}$. Then

$$\theta \cdot \frac{\nu(3n\beta\theta^2)^{\tau(y_i)} \cdot r_i^\alpha}{d(y_i, x_r)^\alpha} \;>\; \nu\theta(3n\beta\theta^2)^{\tau(x_s)-1}$$

holds and hence $d(y_i, x_r) < \sqrt[\alpha]{(3n\beta\theta^2)^{\tau(y_i)-\tau(x_s)+1}} \cdot r_i$. Because it holds by the triangle inequality that $d(y_i, x_s) \leq d(y_i, x_r) + r_s$,

$$\begin{aligned} d(y_i, x_s) &< \sqrt[\alpha]{(3n\beta\theta^2)^{\tau(y_i)-\tau(x_s)+1}} \cdot r_i + r_s \\ &= r_i \cdot (3n\beta\theta^2)^{\frac{\delta_{is}+1}{\alpha}} + r_s \end{aligned}$$

follows. However, this contradicts the fact that y_i and x_s are selected for scheduling in the same time slot. Particularly, at the time *allowed*(y_i, E_t) is invoked for y_i, x_s is already in E_t, and the procedure would evaluate to *false*. Hence, y_i and x_s cannot be scheduled in the same time slot. □

THEOREM 6.3. *For every request $\gamma_{sr} \in \Gamma$, there exists a unique time slot t_s in which x_r successfully receives a message from x_s.*

PROOF. We begin by showing that every x_s with radius $r_s > 0$ is scheduled for sending exactly once during the execution of Algorithm 1. Every node x_s belongs to a single set S_h and each such set is considered in exactly one iteration of the outermost for-loop (more precisely, set S_h is scheduled in iteration k in which $h = m\log(3\beta n) + k$ for some integer $m \geq 0$). Consider this iteration. As long as x_s is not scheduled, it remains in $\mathcal{F}_k$ and the while-loop (Lines 9–16) does not terminate. Termination of this while-loop is guaranteed, however, by the fact that in every iteration at least the node in $\mathcal{F}_k$ with maximal radius is selected for scheduling in E_t and consequently removed from $\mathcal{F}_k$. Because after at most n iterations the set $\mathcal{F}_k$ is empty and the loop terminates, there must be a time slot in which x_s transmits. Further, note that since every $x_s \in E_t$ is removed from $\mathcal{F}_k$, every node has a unique time slot t_s in which it transmits.

Hence, we now need to prove that in this time slot t_s, the message is received successfully by all its intended receivers $x_r \in R_s$. For this purpose, we bound the total interference $I_r = \sum_{y_i \in X\setminus\{x_s\}} I_r(y_i)$ received at any such receiver.

By Lemmas 6.1 and 6.2 we know that for all y_i with $\tau(y_i) < \tau(x_s)$ and $\tau(y_i) > \tau(x_s)$ the interference $I_r(y_i)$ is bounded by $\nu\theta(3n\beta\theta^2)^{\tau(x_s)-1}$. Hence, because there are at most n nodes in these sets, it holds that

$$\sum_{y_i : \tau(x_s) \neq \tau(y_i)} I_r(y_i) \;\leq\; n \cdot \nu\theta(3n\beta\theta^2)^{\tau(x_s)-1}. \qquad (2)$$

What remains to be bounded is the interference created by concurrently transmitting nodes y_i for which $\tau(y_i) = \tau(x_s)$, that is by nodes that are in the same set of the partition $\mathcal{S}$.

Let $\mathcal{T}$ be the set of simultaneously transmitting nodes y_i with $\tau(y_i) = \tau(x_s)$. By the definition of *allowed*(y_j, E_t) a node $y_i \in \mathcal{T}$ prevents all nodes $y_j \in \mathcal{T}$ for which $\mu\theta^{\frac{2}{\alpha}} \cdot r_j > d(y_i, y_j)$ from being added to $\mathcal{T}$. Because the radii of all nodes in $\mathcal{T}$, including x_s, differ at most by a factor of 2, it follows that around each $y_i \in \mathcal{T}$ there can be no other scheduled sender $y_j \in \mathcal{T}$ within distance less than $\frac{1}{2}\mu\theta^{\frac{2}{\alpha}} \cdot r_s$. Hence, disks D_i of radius $\frac{1}{4}\mu\theta^{\frac{2}{\alpha}} r_s$ centered at every node $y_i \in \mathcal{T}$ do not overlap.

Consider rings R_λ of width $\mu\theta^{\frac{2}{\alpha}} r_s$ around x_s, that is, R_λ contains all nodes $y_i \in \mathcal{T}$ for which $(\lambda - \frac{1}{2})\mu\theta^{\frac{2}{\alpha}} r_s < d(x_s, y_i) \leq (\lambda + \frac{1}{2})\mu\theta^{\frac{2}{\alpha}} r_s$. Consider all transmitters $y_i \in \mathcal{T} \cap R_\lambda$ for some integer $\lambda > 0$. All corresponding disks D_i must be located entirely in an extended ring of area

$$\begin{aligned} A(R_\lambda^+) &= \left[\left(\left(\lambda + \frac{3}{4}\right)\mu\theta^{\frac{2}{\alpha}} r_s\right)^2 - \left(\left(\lambda - \frac{3}{4}\right)\mu\theta^{\frac{2}{\alpha}} r_s\right)^2\right]\pi \\ &= 3\lambda\mu^2\theta^{\frac{4}{\alpha}} r_s^2 \pi. \end{aligned}$$

The distance of a transmitter in R_λ to x_r is at least $((\lambda - \frac{1}{2})\mu\theta^{\frac{2}{\alpha}} - 1)r_s$, and each such node transmits with a power of at most $\nu(3n\beta\theta^2)^{\tau(x_s)} \cdot (2r_s)^\alpha$. By applying a standard geometric area argument, we can bound the total interference $I_\lambda = \sum_{y_i \in \mathcal{T} \cap R_\lambda} I_r(y_i)$ incurred by nodes $y_i \in \mathcal{T} \cap R_\lambda$ as

$$\begin{aligned} I_\lambda &\leq \frac{A(R_\lambda^+)}{A(D_i)} \cdot \frac{\theta\nu(3n\beta\theta^2)^{\tau(x_s)} \cdot (2r_s)^\alpha}{(((\lambda - \frac{1}{2})\mu\theta^{\frac{2}{\alpha}} - 1)r_s)^\alpha} \\ &< \frac{3\lambda\mu^2\theta^{\frac{4}{\alpha}} r_s^2\pi}{(\frac{1}{4}\mu\theta^{\frac{2}{\alpha}} r_s)^2\pi} \cdot \frac{\theta\nu(3n\beta\theta^2)^{\tau(x_s)} \cdot (2r_s)^\alpha}{(\frac{1}{2}\lambda\theta^{\frac{2}{\alpha}}(\mu - 1)r_s)^\alpha} \\ &\leq \frac{48\nu(3n\beta\theta^2)^{\tau(x_s)} 2^{2\alpha}}{\theta\lambda^{\alpha-1}(\mu - 1)^\alpha}. \end{aligned}$$

Summing up the interference over all rings R_λ, we obtain

$$\begin{aligned} \sum_{\lambda=1}^{\infty} I_\lambda &\leq \frac{48\nu(3n\beta\theta^2)^{\tau(x_s)} 2^{2\alpha}}{\theta(\mu - 1)^\alpha} \sum_{\lambda=1}^{\infty} \frac{1}{\lambda^{\alpha-1}} \\ &< \frac{48\nu(3n\beta\theta^2)^{\tau(x_s)} 2^{2\alpha}}{\theta(\mu - 1)^\alpha} \cdot \frac{\alpha - 1}{\alpha - 2} \\ &< \frac{\nu}{\theta} \cdot (3\beta)^{\tau(x_s)-1} \cdot (n\theta^2)^{\tau(x_s)}, \end{aligned}$$

where we exploit $\alpha > 2$ to obtain the second inequality and μ's definition for the last inequality.

Adding up the total interference created by nodes for which $\tau(y_i) = \tau(x_s)$ and the total interference by all other nodes as bounded in Inequality (2), we obtain

$$\begin{aligned} I_r &\leq \frac{\nu}{\theta}(3\beta)^{\tau(x_s)-1}(n\theta^2)^{\tau(x_s)} + \nu\theta(3\beta\theta^2)^{\tau(x_s)-1} n^{\tau(x_s)} \\ &= \frac{2\nu\theta}{3} \cdot (\beta\theta^2)^{\tau(x_s)-1} \cdot (3n)^{\tau(x_s)}. \end{aligned}$$

Finally, the SINR experienced at any intended receiver $x_r \in R_s$ is therefore at least

$$\begin{aligned} SINR_r &\geq \frac{\frac{\nu}{\theta}(3n\beta\theta^2)^{\tau(x_s)}}{N + \frac{2}{3}\nu\theta(\beta\theta^2)^{\tau(x_s)-1} \cdot (3n)^{\tau(x_s)}} \\ &= \frac{\frac{4}{\theta}(3n\beta\theta^2)^{\tau(x_s)}}{1 + \frac{8}{3}\theta(\beta\theta^2)^{\tau(x_s)-1} \cdot (3n)^{\tau(x_s)}} \\ &\geq \frac{\frac{4}{\theta}(3n\beta\theta^2)^{\tau(x_s)}}{\frac{11}{3}\theta(\beta\theta^2)^{\tau(x_s)-1} \cdot (3n)^{\tau(x_s)}} \;>\; \beta, \end{aligned}$$

where the second inequality follows from the definition of ν and the third inequality from the fact that $n, \beta \geq 1$. Hence, all scheduled messages are received correctly. □

We now turn our attention to the second aspect of the analysis. Particularly, we prove that the number of time slots required by Algorithm 1 is small, and hence the scheduling complexity is low. We start with a simple helper lemma.

LEMMA 6.4. *In any disk D of diameter d, there can be at most $I_{in}+1$ nodes x_i with $r_i \geq d$.*

PROOF. Assume for contradiction that there are $Q > I_{in}+1$ such nodes in the disk. Since $r_i \geq d$ for all x_i, each node's maximal disk covers the entire disk D. Hence, the interference experienced by each node in D is at least $Q-1$, which contradicts the definition of I_{in} if $Q > I_{in}+1$. □

The idea behind the proof about the schedule lengths produced by Algorithm 1 is to upper-bound the number of nodes which can prevent a node $x_s \in X$ from being selected for scheduling. Since in every iteration of the while-loop at least one node is scheduled, the number of such preventing nodes is an upper bound on the time slots used before x_s is finally scheduled. We say that a node y_i *blocks* node x_s if the presence of y_i in E_t is the reason why $allowed(x_s, E_t)$ evaluates to false. In other words, a node y_i blocks x_s if Algorithm 1 does not schedule x_s simultaneously with y_i. In the sequel, we bound the number of such blocking nodes for each potential sender x_s. We begin with a lemma that captures the number of blocking nodes for which $\tau(y_i) = \tau(x_s)$.

LEMMA 6.5. *Let B_0 be the set of nodes $y_i \in X$ that block x_s with $\tau(y_i) = \tau(x_s)$. For all x_s, $|B_0| \leq \eta\mu^2\theta^{\frac{4}{\alpha}}(I_{in}+1)$ holds for some constant $\eta < 18$.*

PROOF. The proof is based on an area argument. Since all nodes $y_i \in B_0$ are in the same set of the partition $\mathcal{S}$ as x_s, we know that $r_s \leq r_i$ (every y_i is considered before x_s in the algorithm) and $r_s \geq \frac{1}{2}r_i$. By the definition of the algorithm, a node y_i is in B_0 if and only if $\mu\theta^{\frac{2}{\alpha}} \cdot r_s > d(x_s, y_i)$. It follows that all blocking nodes for x_s must be located in a disk D_s of radius $\mu\theta^{\frac{2}{\alpha}} r_s$ around x_s.

Since $r_i \geq r_s$ for all y_i, we know by Lemma 6.4 that there can be at most $I_{in}+1$ blocking nodes in any disk D of diameter r_s. Hence, the number of such disks D required to cover the entire disk D_s times $(I_{in}+1)$ constitutes an upper bound on the number of blocking nodes in B_0. All disks D intersecting D_s are completely inside the disk D'_s, where D'_s has radius $(\mu\theta^{\frac{2}{\alpha}} + \frac{1}{2})r_s$. Furthermore, the disks D can be tesselated in a grid such that the whole area of D'_s is covered while no point in D'_s is covered by more than two disks D. Hence, defining ρ to be the number of disks D required to cover D'_s, we can write

$$\rho \cdot \left(\frac{1}{2}r_s\right)^2 \pi \;\leq\; 2 \cdot \left(\left(\mu\theta^{\frac{2}{\alpha}} + \frac{1}{2}\right) r_s\right)^2 \pi,$$

and by solving for ρ we obtain $\rho \leq 8 \cdot (\mu\theta^{\frac{2}{\alpha}} + 1/2)^2 \leq \eta\mu^2\theta^{\frac{4}{\alpha}}$, which concludes the proof. □

The more intricate part of the analysis is to bound the number of blocking nodes in other sets of the partition. In particular, note that it can be relatively easily shown that there are only a constant number of blocking nodes for x_s in each set of the partition $\mathcal{S}$. However, this bound is not sufficient, as there can be as many as $\Omega(n/\log n)$ different sets that are considered in the same outer for-loop iteration. Hence, we need a much stronger bound in order to guarantee the scheduling complexity as claimed in Theorem 5.1. We start with a helper lemma that characterizes the ratio between the radii of two nodes.

LEMMA 6.6. *Let x_i and x_j be two nodes that are considered in the same iteration of the for-loop, and let $\tau(x_i) \leq \tau(x_j)$. Then, for $\delta_{ij} = \tau(x_i) - \tau(x_j)$, it holds that $r_i \geq \frac{1}{2}(3n\beta\theta^2)^{\delta_{ij}} \cdot r_j$.*

PROOF. In the same iteration of the for-loop, only links in sets $S_k, S_{\lceil\log(3n\beta\theta^2)\rceil+k}, S_{2\lceil\log(3n\beta\theta^2)\rceil+k}, \ldots$ are considered. The value of δ_{ij} denotes the number of these sets that separate the sets containing x_i and x_j. Each such separating set entails at least a doubling of the respective radii. All radii differing at most by a factor of 2 within one set, it follows that $r_i \geq r_j \cdot 2^{\delta_{ij}\lceil\log(3n\beta\theta^2)\rceil-1} \geq r_j \cdot \frac{1}{2}(3n\beta\theta^2)^{\delta_{ij}}$. □

For the next couple of lemmas, we need to introduce some additional notation. Specifically, we define the *reduced distance* ζ_s^i of y_i from x_s to be $\zeta_s^i = d(x_s, y_i) - r_i$. In words, the reduced distance is a lower bound on the minimum possible distance between x_s and an intended receiver of y_i. Note that in procedure $allowed(x_i, E_t)$ node x_s is blocked by a node $y_i \in E_t$, $\tau(x_s) > \tau(y_i)$, if and only if $r_s(3n\beta\theta^2)^{\frac{\delta_{si}+1}{\alpha}} > \zeta_s^i$.

LEMMA 6.7. *For any $\varphi \geq 2$ and $\sigma = 72$, there can be at most $\sigma(I_{in}+1)$ blocking nodes y_i for a node x_s with reduced distance*

$$(3n\beta\theta^2)^{\frac{\varphi}{\alpha}} \cdot r_s \;<\; \zeta_s^i \;\leq\; (3n\beta\theta^2)^{\varphi} \cdot r_s.$$

PROOF. Assume for contradiction that there exists a set of $\sigma(I_{in}+1)+1$ or more nodes y_i that are blocking x_s with each ζ_s^i in the range specified in the lemma. Denote this set of nodes by $\mathcal{B} \subseteq X$. First note that if a node $y_i \in \mathcal{B}$ blocks x_s and the reduced distance $\zeta_s^i > (3n\beta\theta^2)^{\frac{\varphi}{\alpha}} \cdot r_s$, then $\delta_{is} > \varphi - 1$ must hold and consequently $\delta_{is} \geq \varphi$. This is true because if $\delta_{is} < \varphi$, it holds that $(3n\beta\theta^2)^{\frac{\delta_{si}+1}{\alpha}} \cdot r_s < \zeta_s^i$ and consequently, by the definition of the algorithm, y_i does not block x_s. Hence, in combination with Lemma 6.6, we know that all blocking nodes $y_i \in \mathcal{B}$ have radii of at least

$$r_i \;\geq\; \frac{1}{2}(3n\beta\theta^2)^{\varphi} \cdot r_s. \tag{3}$$

We now show that if $|\mathcal{B}| \geq \sigma(I_{in}+1)+1$, then there must be a node that has in-interference at least $I_{in}+1$, which leads to a contradiction. For this purpose we first present a transformation $\mathcal{B}'$ of the node set $\mathcal{B}$ which does not increase I_{in}. We then show that in this transformed instance $\mathcal{B}'$, in-interference is too high.

Consider the following transformation of the node set $\mathcal{B}$ into a node set $\mathcal{B}'$: We replace each node $y_i \in \mathcal{B}$ with radius r_i by a node $y'_i \in \mathcal{B}'$ with radius $r'_i = \frac{1}{2}(3n\beta\theta^2)^{\varphi} \cdot r_s$. Specifically, node y'_i is located on the straight line connecting x_s and y_i at distance $\zeta_s^i + \frac{1}{2}(3n\beta\theta^2)^{\varphi} \cdot r_s$ from x_s, as shown in Figure 4. Note that the disk with radius r'_i centered at y'_i is entirely contained in the disk with radius r_i around y_i (cf. Inequality (3)); this transformation cannot increase the in-interference of any node in the network. Moreover, because the distance $d(x_s, y_i)$ and r_i are reduced by the same amount, this transformation does not change the value ζ_s^i for any transformed node, that is, $\zeta_s^i = \zeta_s^{i'}$ for any $y_i \in \mathcal{B}$ and its transformation $y'_i \in \mathcal{B}'$. Finally, note that transforming $\mathcal{B}$ to $\mathcal{B}'$ is always possible.

We now look at the in-interference of node set $\mathcal{B}'$. According to the second inequality of the Lemma to be proven, the reduced radius ζ_s^i of each node $y'_i \in \mathcal{B}'$ is at most $(3n\beta\theta^2)^{\varphi} \cdot r_s$. All radii being $r'_i = \frac{1}{2}(3n\beta\theta^2)^{\varphi} \cdot r_s$, it follows that all nodes $y'_i \in \mathcal{B}'$ are located in a disk D_s^{φ} of radius $\frac{3}{2}(3n\beta\theta^2)^{\varphi} \cdot r_s$ centered at x_s.

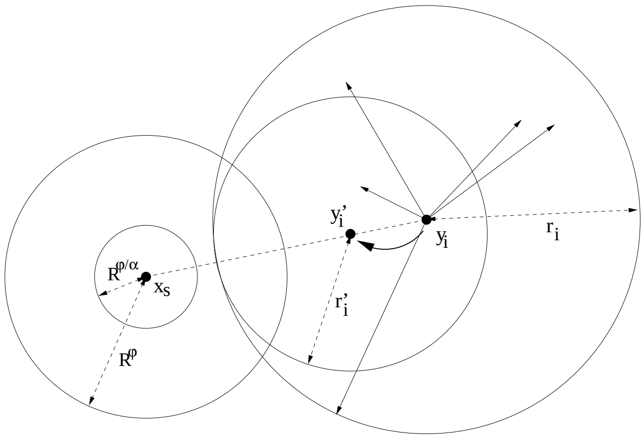

Figure 4: Example for the transformation used in the proof of Lemma 6.7. with $R^{\varphi/\alpha} = (3n\beta\theta^2)^{\frac{\varphi}{\alpha}} \cdot r_s$ and $R^{\varphi} = (3n\beta\theta^2)^{\varphi} \cdot r_s$, respectively. The disk with center y_i and radius r_i is replaced by the smaller disk with center y_i' and radius r_i', such that $\zeta_s^i = \zeta_s^{i'}$. The transformation does not increase I_{in}.

Consider disks D of radius $\frac{1}{4}(3n\beta\theta^2)^{\varphi} \cdot r_s$. By the standard area argument also applied in the proof of Lemma 6.5, the number of disks ρ required to cover the entire disk D_s^{φ} is at most

$$\rho \;\leq\; 2 \cdot \frac{(\frac{3}{2}(3n\beta\theta^2)^{\varphi} r_s)^2 \pi}{(\frac{1}{4}(3n\beta\theta^2)^{\varphi} r_s)^2 \pi} \;=\; \sigma.$$

Since by assumption there are at least $\sigma(I_{in}+1)+1$ nodes in $\mathcal{B}'$, there must exist a disk D that contains at least $I_{in}+2$ of these nodes. This, however, establishes a contradiction because by Lemma 6.4 there can be at most $I_{in}+1$ nodes $y_i' \in \mathcal{B}'$ in any disk D of diameter $d = \frac{1}{2}(3n\beta\theta^2)^{\varphi} r_s$. $\square$

Using Lemmas 6.6 and 6.7, we can now establish the following key lemma.

LEMMA 6.8. *Let B_+ be the set of nodes $y_i \in X$ that block x_s, with $\tau(y_i) < \tau(x_s)$. It holds that for all x_s, $|B_+| \leq (\log_\alpha \frac{n}{\alpha} + 1)\sigma(I_{in}+1)$, for σ as defined in Lemma 6.7.*

PROOF. By Lemma 6.7, we know that there can be at most $\sigma(I_{in}+1)$ blocking nodes y_i for x_s with reduced distance $(3n\beta\theta^2)^{\frac{\varphi}{\alpha}} \cdot r_s \;<\; \zeta_s^i \leq (3n\beta\theta^2)^{\varphi} \cdot r_s$. In particular, this means that there are at most $\sigma(I_{in}+1)$ blocking nodes with reduced distance $(3n\beta\theta^2)^{\frac{1}{\alpha}} \cdot r_s \;<\; \zeta_s^i \leq (3n\beta\theta^2) \cdot r_s$, at most $2\sigma(I_{in}+1)$ such nodes with $(3n\beta\theta^2)^{\frac{1}{\alpha}} \cdot r_s < \zeta_s^i \leq (3n\beta\theta^2)^{\alpha} \cdot r_s$, and so forth. More generally, there are at most $\kappa\sigma(I_{in}+1)$ blocking nodes with

$$(3n\beta\theta^2)^{\frac{1}{\alpha}} \cdot r_s \;<\; \zeta_s^i \;\leq\; (3n\beta\theta^2)^{\alpha^{\kappa-1}} \cdot r_s.$$

Because the partition $\mathcal{S}$ consists of at most n non-empty sets S_i, it holds that $\tau(x_s) - \tau(y_i) \leq n$ for all y_i. Therefore, the reduced distance of any blocking node cannot exceed $r_s(3n\beta\theta^2)^{\frac{n}{\alpha}}$. For $\kappa = \log_\alpha \frac{n}{\alpha} + 1$, on the other hand, the reduced distance is at most $r_s(3n\beta\theta^2)^{\alpha^{\log_\alpha \frac{n}{\alpha}}} = r_s(3n\beta\theta^2)^{\frac{n}{\alpha}}$. Hence, there can be at most $(\log_\alpha \frac{n}{\alpha}+1)\sigma(I_{in}+1)$ blocking nodes for x_s, from which the lemma follows. $\square$

Finally, we can put everything together in the following theorem, which also implies Theorem 5.1.

THEOREM 6.9. *The number of time slots required by Algorithm 1 to successfully schedule all links $\gamma_{ij} \in \Gamma$ is at most $O\left(I_{in} \cdot \log(n\theta^2) \cdot (\theta^{\frac{4}{\alpha}} + \log n)\right)$.*

PROOF. By Lemmas 6.5 and 6.8, there are at most

$$B_0 + B_+ \;\leq\; \eta\mu^2\theta^{\frac{4}{\alpha}}(I_{in}+1) + (\log_\alpha \frac{n}{\alpha} + 1)\sigma(I_{in}+1)$$

blocking nodes for each node x_s. Hence, after at most $\eta\mu^2\theta^{\frac{4}{\alpha}}(I_{in}+1) + (\log_\alpha \frac{n}{\alpha}+1)\sigma(I_{in}+1) + 1$ iterations of the while-loop, all nodes that are considered in the same iteration of the outer for-loop are scheduled for transmission. The number of for-loop iterations being $\lceil \log(3n\beta\theta^2) \rceil$, it follows that for constant β the scheduling complexity $T(S)$ of Algorithm 1 (that is the number of time slots required) is

$$T(S) \in O\left(I_{in} \cdot \log(n\theta^2) \cdot (\theta^{\frac{4}{\alpha}} + \log n)\right).$$

$\square$

In the standard physical model ($\theta = 1$) the scheduling complexity of Algorithm 1 reduces to $O(I_{in} \cdot \log^2 n)$.

7. LOW-INTERFERENCE TOPOLOGIES

Having displayed the significance of static in-interference in the previous section, we specify the value I_{in} in this section more precisely for the most basic network property, that is connectivity. In particular, we distinguish between strong connectivity with directed links and connectivity with undirected links, as it turns out that requesting links to be symmetric significantly complicates the task of quickly scheduling communication requests.

In the following, we first take a look at connected topologies with asymmetric (also called directed or unidirectional) links. In particular, we consider strongly connected topologies, meaning that there exists from every node x_i in the network to every other node x_j a path containing only links oriented from x_i to x_j.

THEOREM 7.1. *Given an arbitrary set of nodes X, there exists a strongly connected topology with asymmetric links having in-interference $I_{in} \in O(\log n)$. On the other hand, there exist node sets for which every strongly connected topology with asymmetric links has interference $I_{in} \in \Omega(\log n)$.*

PROOF. In [10], a sink tree structure is presented which contains directed paths from all nodes to a predefined sink node; it is shown that this structure has interference $I_{in} \in O(\log n)$. If we add directed links from the sink node to all other nodes in the network, the resulting structure is a strongly connected topology with I_{in} increased by at most 1. Furthermore, it was shown that there exist node sets for which all strongly connected topologies with directed links have interference at least $I_{in} \in \Omega(\log n)$. $\square$

It is often argued that communication over asymmetric wireless links is costly [26] or in general unacceptably cumbersome; not even simple acknowledgement of a transmitted packet is easily possible over an asymmetric link. In the following we however show that—from a scheduling perspective—demanding communication links to be symmetric does not come for free. Specifically, it has been shown in [32] that the in-interference experienced at a node if links are required to be symmetric can be as high as $\Omega(\sqrt{n})$. Notice that this is significantly higher than the $O(\log n)$ interference bound that holds in the case where links can be asymmetric.

THEOREM 7.2 ([32]). *There exist node sets for which every connected topology with symmetric links exhibits in-interference at least* $I_{in} \in \Omega(\sqrt{n})$.

A comparative interpretation of Theorems 7.1 and 7.2 therefore leads to the conclusion that the scheduling of a connected topology with exclusively symmetric links is by its nature significantly more costly than the scheduling of a connected topology using asymmetric links. More precisely, combining the two results with the scheduling algorithm proposed in Section 5 shows that essentially a strongly connected topology with asymmetric links can theoretically be scheduled in time $O(\log^3 n)$. On the other hand, any protocol that combines a low-interference algorithm with the algorithm of Section 5 requires as much time as $O(\sqrt{n} \cdot \log^2 n)$ in order to schedule a connected topology consisting of symmetric links. In a sense, this result forms an antithesis to the often made assumption that the use of symmetric links is mandatory for practical reasons. At least, this observation provides a new aspect to the discussion whether asymmetric edges are valid to be considered for network forming or if they ought to be disregarded altogether.

As mentioned above, combining the $O(\log n)$ in-interference topology control algorithm of Theorem 7.1 with the result on the scheduling complexity in Section 5 proves that the *scheduling complexity of connectivity*, as studied in [25] with directed links, is in $O(\log^3 n)$. This bound improves the result presented in [25] by a logarithmic factor:

COROLLARY 7.3. *In the physical model, the scheduling complexity* $T(\mathcal{S})$ *of strong connectivity in arbitrary worst-case networks is at most* $T(\mathcal{S}) \in O(\log^3 n)$.

This result implies that the scheduling of connected topologies can, at least in theory, be performed highly efficiently even in worst-case networks, with schedule lengths that scale well even for large networks.

8. CONCLUSIONS

In this paper we introduce the concept of *scheduling complexity of arbitrary topologies* in the physical SINR model, focusing on the time required to actually schedule the communication requests of a topology over time. Presenting and analyzing a scheduling algorithm, we prove a sharp upper bound on this complexity. Intriguingly, this bound depends on a previously defined static interference measure. In this sense, we put the concept of static interference into the context of physical message realization. With respect to static interference—from the perspective of scheduling complexity—, we also demonstrate the presence of a breach between connected topologies consisting exclusively of symmetric communication links and connected topologies allowed to contain asymmetric links. This analysis not only sheds new light on the question of practicality of directed links, but also improves an existing upper bound on the scheduling complexity of connectivity. Finally, one of the main conclusions of our paper is that—from a theoretical point of view—arbitrary topologies can be realized in short schedules that scale well even for large networks.

It would be interesting to simplify our scheduling algorithm and come up with a practical protocol implementation, since a network protocol with guaranteed low scheduling complexity for arbitrary topologies forms an ideal MAC layer protocol. While we have clearly not yet reached this point, we believe that this paper constitutes a step towards that ultimate goal by laying a theoretical foundation that provides the first known scaling laws for the achievable scheduling complexity of arbitrary topologies.

9. REFERENCES

[1] A. Behzad and I. Rubin. On the Performance of Graph-based Scheduling Algorithms for Packet Radio Networks. In *Proceedings of the IEEE Global Telecommunications Conference (GLOBECOM)*, pages 3432–3436, 2003.

[2] A. Behzad and I. Rubin. Impact of Power Control on the Performance of Ad Hoc Wireless Networks. In *Proceedings of the 24th Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2005.

[3] P. Björklund, P. Värbrand, and D. Yuan. A Column Generation Method for Spatial TDMA Scheduling in Ad Hoc Networks. *Ad Hoc Networks*, 2(4):405–418, 2004.

[4] D. M. Blough, M. Leoncini, G. Resta, and P. Santi. The K-Neigh Protocol for Symmetric Topology Control in Ad Hoc Networks. In *Proceedings of the 4th ACM International Symposium on Mobile Ad Hoc Networking & Computing (MOBIHOC)*, pages 141–152, 2003.

[5] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia. Routing with Guaranteed Delivery in Ad Hoc Wireless Networks. In *Proc. of the 3rd International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIALM)*, pages 48–55, 1999.

[6] M. Burkhart, P. von Rickenbach, R. Wattenhofer, and A. Zollinger. Does Topology Control Reduce Interference? In *Proceedings of the 5th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC)*, pages 9–19, 2004.

[7] R. Cruz and A. Santhanam. Optimal Routing, Link Scheduling, and Power Control in Multi-hop Wireless Networks. In *Proc. of the 22nd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2003.

[8] T. ElBatt and A. Ephremides. Joint Scheduling and Power Control for Wireless Ad-hoc Networks. In *Proc. of the 21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2002.

[9] A. Ephremides and T. Truong. Scheduling Broadcasts in Multihop Radio Networks. *IEEE Transactions on Communications*, 38:456–460, 1990.

[10] M. Fussen, R. Wattenhofer, and A. Zollinger. Interference Arises at the Receiver. In *Proc. of the International Conference on Wireless Networks, Communications, and Mobile Computing (WirelessCom)*, June 2005.

[11] J. Grönkvist. *Interference-Based Scheduling in Spatial*

Reuse TDMA. PhD thesis, Royal Institute of Technology, Stockholm, Sweden, 2005.

[12] J. Grönkvist and A. Hansson. Comparison Between Graph-Based and Interference-Based STDMA Scheduling. In *Proceedings of the 2nd ACM International Symposium on Mobile Ad Hoc Networking & Computing (MOBIHOC)*, pages 255–258, 2001.

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

[14] T. Hou and V. Li. Transmission Range Control in Multihop Packet Radio Networks. *IEEE Transactions on Communications*, 34(1):38–44, 1986.

[15] L. Hu. Topology Control for Multihop Packet Radio Networks. *IEEE Trans. on Communications*, 41(10), 1993.

[16] K. Jain, J. Padhye, V. Padmanabhan, and L. Qiu. Impact of Interference on Multi-hop Wireless Network Performance. In *Proc. of the 9th Annual International Conference on Mobile Computing and Networking (MOBICOM)*, 2003.

[17] F. Kuhn, R. Wattenhofer, Y. Zhang, and A. Zollinger. Geometric Routing: Of Theory and Practice. In *Proc. of the 22nd ACM Symposium on the Principles of Distributed Computing (PODC)*, 2003.

[18] N. Li, C.-J. Hou, and L. Sha. Design and Analysis of an MST-Based Topology Control Algorithm. In *Proc. of the 22nd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2003.

[19] X.-Y. Li, G. Calinescu, and P.-J. Wan. Distributed Construction of Planar Spanner and Routing for Ad Hoc Wireless Networks. In *Proc. of the 21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2002.

[20] X.-Y. Li, W.-Z. Song, and W. Wang. A Unified Energy Efficient Topology for Unicast and Broadcast. In *Proceedings of the 11th International Conference on Mobile Computing and Networking (MOBICOM)*, 2005.

[21] F. Meyer auf der Heide, C. Schindelhauer, K. Volbert, and M. Grünewald. Energy, Congestion and Dilation in Radio Networks. In *Proceedings of the 14th ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 230–237, 2002.

[22] K. Moaveni-Nejad and X.-Y. Li. Low-Interference Topology Control for Wireless Ad Hoc Networks. In *Proc. of the 2nd IEEE Communications Society Conference on Sensor and Ad Hoc Communications and Networks (SECON)*, Santa Clara, California, USA, 2005.

[23] T. Moscibroda and R. Wattenhofer. Coloring Unstructured Radio Networks. In *Proceedings of the 17th Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 39–48, 2005.

[24] T. Moscibroda and R. Wattenhofer. Minimizing Interference in Ad Hoc and Sensor Networks. In *Proc. of the International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIALM)*, pages 24–33, 2005.

[25] T. Moscibroda and R. Wattenhofer. The Complexity of Connectivity in Wireless Networks. In *Proc. of the 25th Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2006.

[26] R. Prakash. Unidirectional Links Prove Costly in Wireless Ad-Hoc Networks. In *Proc. of the 3rd International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIAL-M)*, 1999.

[27] R. Ramanathan and R. Rosales-Hain. Topology Control of Multihop Wireless Networks Using Transmit Power Adjustment. In *Proc. of the 19th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2000.

[28] S. Ramanathan and E. L. Lloyd. Scheduling Algorithms for Multi-Hop Radio Networks. In *Proceedings on Communications Architectures and Protocols (SIGCOMM)*, pages 211–222, 1992.

[29] V. Rodoplu and T. H. Meng. Minimum Energy Mobile Wireless Networks. *IEEE J. Selected Areas in Communications*, 17(8), 1999.

[30] P. Santi. *Topology Control in Wireless Ad Hoc and Sensor Networks*. Wiley, 2005.

[31] H. Takagi and L. Kleinrock. Optimal Transmission Ranges for Randomly Distributed Packet Radio Terminals. *IEEE Transactions on Communications*, 32(3):246–257, 1984.

[32] P. von Rickenbach, S. Schmid, R. Wattenhofer, and A. Zollinger. A Robust Interference Model for Wireless Ad-Hoc Networks. In *Proc. of the 5th Int. Workshop on Algorithms for Wireless, Mobile, Ad Hoc and Sensor Networks (WMAN)*, Denver, Colorado, USA, April 2005.

[33] R. Wattenhofer, L. Li, P. Bahl, and Y.-M. Wang. Distributed Topology Control for Power Efficient Operation in Multihop Wireless Ad Hoc Networks. In *Proceedings of the 20th Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2001.

[34] R. Wattenhofer and A. Zollinger. XTC: A Practical Topology Control Algorithm for Ad-Hoc Networks. In *Proc. of the 4th Int. Workshop on Algorithms for Wireless, Mobile, Ad Hoc and Sensor Networks (WMAN)*, Santa Fe, New Mexico, USA, April 2004.