

Link Scheduling in Local Interference Models

Bastian Katz*, Markus Völker, and Dorothea Wagner

Faculty of Informatics, Universität Karlsruhe (TH)
{katz,mvoelker,wagner}@ira.uka.de

Abstract. Choosing an appropriate interference model is crucial for link scheduling problems in sensor networks. While graph-based interference models allow for distributed and purely local coloring approaches which lead to many interesting results, a more realistic and widely agreed on model such as the signal-to-noise-plus-interference ratio (SINR) inherently makes scheduling radio transmission a non-local task, and thus impractical for the development of distributed and scalable scheduling protocols in sensor networks. In this work, we focus on interference models that are *local* in the sense that admissibility of transmissions only depends on local concurrent transmissions, and *correct* with respect to the geometric SINR model.

In our analysis, we show lower bounds on the limitations that these restrictions impose on any such model as well as approximation results for greedy scheduling algorithms in a class of these models.

1 Introduction

Agreeing on good schedules in wireless networks is not only a question of good, i. e., local and distributed scheduling algorithms. The correctness of any scheduling algorithm's output relies on the underlying interference model. Choosing an interference model is thus crucial for any kind of scheduling protocol in sensor networks. Both, interference models and scheduling problems have been studied thoroughly in the "tradition" of sensor networks. Complex interference models incorporating sophisticated signal fading models and antenna characteristics, developed over the years, proved helpful in the simulation and design of sensor networks. In the algorithmic community, however, the need of clear, preferably combinatoric and geometric interference models, led to a focus on graph-based interference models. These graphs all have in common that they are local in the sense that mutual exclusion between transmissions only "connects" nodes that are close to each other. The simple combinatorial character of these models naturally translates scheduling problems to coloring problems in graphs. Moreover, the geometric properties of these graphs allow for tailored coloring protocols.

Despite their simplicity, the downside of these models clearly is that they could neither be proven to be correct nor good in real sensor networks. They cannot

* Partially supported by the German Research Foundation (DFG) within the Research Training Group GRK 1194 "Self-organizing Sensor-Actuator Networks".

model interference from far away nodes summing up and jamming communication, nor can they model that if in reality any pair out of three transmissions can successfully be performed simultaneously, this does not necessarily mean that all three transmissions can be performed simultaneously.

Algorithmic research considering a class of models that renders signal propagation much more realistically did to the best of our knowledge not yet lead to local algorithms. In SINR models, successful or sufficiently probable reception is assumed if at a receiver, the respective sender's signal strength outperforms the sum of all interfering signals plus the background noise by a hardware dependent constant. The geometric SINR models closely cover the main features of sophisticated fading models such as the two-ray-ground model without losing too much of the simplicity needed for algorithmic results.

In this paper, we introduce the concept of locality and correctness of interference models. We prove fundamental limitations of all models that are local in a very straightforward sense and correct with respect to the geometric SINR model. We show under which conditions well known concepts such as graph coloring *can* be used to approximate scheduling problems and a generalization that improves the quality of easy-to-implement scheduling algorithms. We believe that the introduced models open a door to more realistic, yet viable solutions not only for scheduling, but for many protocols that rely on local, dependable communication.

2 Related Work

Interference of concurrent communication, being the most outstanding attribute of wireless networks, has been subject of countless publications. Since in reality, interference is composed of many hard-to-capture phenomena such as multipath fading, algorithmic research developed numerous simplifications. Most of the algorithmic models model interference as a binary relation on transmissions, among them the unit disk graph (UDG) with distance or hop interference or the protocol model. We refer the reader to [1] for a survey. In SINR models, successful reception depends on the ratio between the received signal strength on the one side and the interference from concurrent transmissions plus the background noise on the other side [2, 3]. They differ in whether they assume signal strength decay to be a function of the distance (geometric SINR) or allow an arbitrary gain matrix. In the geometric SINR model, Gupta and Kumar analyzed the capacity of ad-hoc networks and proved an upper bound on the throughput of $\Theta(1/\sqrt{n})$ for networks of n nodes. Until now, the effects of the SINR models to algorithm design raise interesting questions [4].

Scheduling of link transmissions has been addressed in many interference models and, in most cases, proven to be NP-hard. Among others are proofs for scheduling in graph-based models [5, 6], in the abstract SINR model in [7], and, recently, in the geometric SINR model for fixed power assignment by Goussevskaia et al. [8]. The joint problem of scheduling and power assignment is still open in the geometric SINR model [9]. A variety of graph-based scheduling

algorithms has been proposed and analyzed [10, 11, 12]. It is however argued in various works that graph-based scheduling is inferior to scheduling designed for the SINR model [13, 14]. Among the early publications addressing scheduling in geometric SINR models, Moscibroda and Wattenhofer show that uniform or linear power assignments in worst-case scenarios need exponentially longer schedules for a strongly connected set of links [15] than more sophisticated assignments. Moscibroda et al. also propose a scheduling algorithm for arbitrary power assignment in [16] that outperforms previous heuristics by an exponential factor. In [8], Goussevskaia et al. propose an approximation algorithm for link scheduling and the problem of finding a maximum number of links that can transmit concurrently in the geometric SINR model under the fixed power assumption. The latter three works introduced many of the techniques applied in the following under the practically more relevant assumptions that nodes do not feature arbitrarily high transmission powers and cannot rely on a global instance to compute a schedule, but are restricted to a local view. Locality has, to our knowledge, only been looked at in a combinatorial sense [17, 18, 19].

3 Definitions and Models

A *deterministic* interference model \mathcal{M} is a property telling for a fixed set of nodes \mathcal{V} whether a set of transmissions T between nodes in \mathcal{V} can be carried out simultaneously for given transmission powers. More formally, let $\mathcal{T} := \{(u, v, p) \in \mathcal{V}^2 \times \mathbb{R}_+ \mid u \neq v\}$ be the set of all possible transmissions and transmission powers. Then, a model $\mathcal{M} \subseteq \mathcal{P}(\mathcal{T})$ contains all sets of transmissions which are valid. We further assume that less concurrent transmissions cannot cause a transmission to fail, i. e., that for all $T' \subseteq T \subseteq \mathcal{T}$,

$$T \in \mathcal{M} \Rightarrow T' \in \mathcal{M} , \quad (1)$$

which holds for all models which are currently used and most likely for all models which are meaningful. One should note that the restriction to deterministic models alone already is a giant step away from reality and the probabilistic models typically employed by communication theorists. But still, even deterministic models are not understood well. Such models can rely on various kinds of additional input and assumptions of radio propagation, antenna characteristics and so on. In higher layer protocol design, however, there is a need to “model away” the complexity of most of these unrulable phenomena. An aspect that has not received much attention yet is how different approaches to model interference relate to each other, or, in other words: If I choose a simpler model, are my algorithms or schedules still correct with respect to a more realistic one or are they just suboptimal? Do optimal solutions in a simple model approximate optimal solutions in a more complex model? In the following, we will call an interference model \mathcal{M} *conservative* with respect to another model \mathcal{M}' if $\mathcal{M} \subset \mathcal{M}'$.

Most analytical research on scheduling problems has been done in some kind of *graph-based* interference model accordig to the following definition from [4].

Definition 1 (Graph-Based Model). A graph-based model \mathcal{M} can be defined by two directed graphs, one connectivity graph $D_C = (\mathcal{V}, A_C)$ restricting possible transmissions and one interference graph $D_I = (A_C, A_I)$ connecting conflicting transmissions, such that $T \in \mathcal{M}$ if and only if $T \subset A_C$ and $T^2 \cap A_I = \emptyset$.

Usually, a simpler model consisting of two graphs $G_C = (\mathcal{V}, E_C)$ and $G_I = (\mathcal{V}, E_I)$ is used, in which a set of transmissions is valid, if for every sender, the intended receiver is a neighbor in G_C and no receiver of a distinct transmission is connected in G_I . Sometimes, the connectivity graph and interference graph are defined implicitly, i. e., as the result of a geometric setting.

Graph-based models all have in common that they claim that a set of transmissions whose transmissions can *pairwise* be carried out at the same time, collectively may be scheduled into one single time slot. This is unrealistic in general, and the models fail to formulate the assumptions under which they guarantee not to produce schedules that do not comply with more realistic models. On the other hand, in the single-power case, graph-based models reduce scheduling problems to well-known coloring problems.

As opposed to the oversimplification of graph-based interference models, the models capturing the findings of signal propagation and reception best are the *signal-to-noise-plus-interference* (SINR) models. Their main paradigm is that a transmission is (almost) always successful, if the sender's signal strength *at the receiver* is significantly stronger than the sum of all interfering signals, including other sender's signals and (individual) background noise. Thus, in its most general form, an SINR model is defined by a *gain* matrix (G_{uv}) denoting the signal fading between nodes u and v , on the background noise η_v at each of the nodes and the (individual) ratio β_v a node v needs for proper reception. Here, a set of transmissions is valid, i. e. $T \subset \mathcal{M}$, if and only if for all $t = (s, r, p_s) \in T$

$$\frac{p_s G_{sr}}{\eta_r + \sum_{(u,v,p_u) \in T \setminus \{t\}} p_u G_{ur}} \geq \beta_r . \quad (2)$$

Definition 2 (Geometric Model). In a geometric model, \mathcal{M} is defined for $\mathcal{V} = \mathbb{R}^2$ such that \mathcal{M} is invariant under all isometries.

Generally speaking, geometric interference models are incapable of modeling individual characteristics of nodes, but are restricted to those of geometric settings. This does not mean that a geometric model has to be parameter-free, but for geometric SINR models, this definition implies a much simpler structure:

Theorem 1. Every geometric SINR model can also be defined equivalently such that all η_v and all β_v are independent of the respective position v and all G_{uv} can be expressed as $G_{uv} := f(d(u, v))$ for a $f : \mathbb{R} \rightarrow \mathbb{R}$.

Proof. Let \mathcal{M} be a geometric SINR model defined by (G_{ji}) , (η_i) and (β_i) . We get an equivalent model for (G'_{ji}) , (η'_i) , and (β_i) , in which all $\eta'_i = \eta'$ are the same by setting $G'_{ji} = G_{ji} \frac{\eta'_i}{\eta_i}$. Now, for any d , take two pairs u_1, v_1 and u_2, v_2 . We know that $\{(u_i, v_i, p)\} \in \mathcal{M}$ if and only if $p \geq \eta' \beta'_{v_i} / G_{u_i v_i}$. Since there is an isometry

mapping u_1 to u_2 and v_1 to v_2 , and since \mathcal{M} is geometric, transmissions (u_1, v_1, p) are valid, for exactly the same values of p as (u_2, v_2, p) . Thus, $p \geq \eta' \beta'_{v_1} / G'_{u_1 v_1}$ if and only if $p \geq \eta' \beta'_{v_2} / G'_{u_2 v_2}$ and thus $\beta'_{v_1} / G'_{u_1 v_1} = \beta'_{v_2} / G'_{u_2 v_2}$. I. e., all pairs of nodes with distance d have the same ratio of G'_{ji} and β_i and by fixing some β' and setting $G''_{ji} = G'_{ji} \beta' / \beta_i$, we get a representation of the claimed form.

The class of *geometric SINR models* ($SINR_G$) is a quite straightforward application of the above definition. Individual characteristics such as the background noise and the necessary SINR ratio are replaced by common constants η and β and the gain G_{uv} is replaced by a function of the distance, usually $K d_{uv}^{-\alpha}$ for a so-called *path-loss exponent* α and some constant K . Currently, the $SINR_G$ models widely agreed are the best models to reason about in the algorithmics of sensor networks. Thus, we will focus on local models that are conservative with respect to this class of models.

A *scheduling problem* in a wireless network is a set \mathcal{Q} of communication requests, each request (s, r) consisting of a sender s and a receiver r , both from some set \mathcal{V} of nodes. A schedule then is a sequence T_1, T_2, \dots, T_k of sets of transmissions of the form (s, r, p) for some $(s, r) \in \mathcal{Q}$ and some power assignment $p \in \mathbb{R}_+$, such that for every $(s, r) \in \mathcal{Q}$, there is a transmission (s, r, p) in one of the T_i , and every T_i is valid with respect to an interference model. We refer to the problem of finding a schedule of minimum length as SCHEDULE, and to the problem of finding a maximum number of transmissions that can be scheduled to a single slot as ONESHOTSCHEDULE as in [8]. We will also denote the maximum link length occurring in a schedule request \mathcal{Q} by $\ell(\mathcal{Q}) := \max_{(s,r) \in \mathcal{Q}} d_{sr}$. If the scheduling problem is combined with the problem of assigning transmission powers, usually powers must be chosen from some power range $\mathbf{p} = [p_{\min}, p_{\max}]$. In the following, we will focus on the problem of finding schedules for a fixed power p and thus also write (s, r) to denote a transmission (s, r, p) .

4 Local Interference Models

The concept of locality has been introduced for distributed systems and adopted in the context of sensor networks. Usually, a distributed algorithm is said to be k -local, if the outcome for every node only depends on nodes which are inside a k -hop-neighborhood. Unfortunately, this concept is too restrictive to allow for any local scheduling algorithms in a geometric SINR model with nodes that do not feature arbitrarily high transmission powers, but are limited to some maximum power. Even if we define a node's neighborhood as the set of nodes the node can communicate with when no other communication takes place at the same time, in an $SINR_G$ model \mathcal{M}_G , it might be impossible to arrange a schedule at all. If we denote the maximum possible link length of an interference model \mathcal{M} by $\ell(\mathcal{M}) := \limsup_{\{(u,v,p)\} \in \mathcal{M}} d_{uv}$, we get $\ell(\mathcal{M}_G) = \sqrt[k]{Kp/(\beta\eta)}$, since for nodes with higher distance even in the absence of concurrent transmissions sending at maximum power does not result in a received signal strength of $\beta\eta$, which is necessary due to the background noise alone.

In Fig. 1, such a situation is depicted: Out of the two sender/receiver pairs in the transmission request, only the pairs themselves have a distance less than $\ell(\mathcal{M}_G)$, and thus, there is no communication possible between the different pairs, which, however, have to agree not to transmit at the same time since both of them cannot compensate for the interference caused by the other. We will thus in the next section look at a weaker, geometric definition of locality and its consequences for scheduling problems.




Fig. 1. Links with distance $\ell(\mathcal{M}_G) + \epsilon$ may need communication

Definition 3 (Local Model). A ρ -local model is a geometric model \mathcal{M} with the additional constraint that $T \in \mathcal{M}$ if for every $t = (s, r, p) \in T$

$$T(s, \rho) := \{(s', r', p') \in T \mid d(s, s') \leq \rho\} \in \mathcal{M} .$$

In other words, an interference model is *local*, if for a set of transmissions T , it is sufficient that for every sender in T the transmissions in its ρ -neighborhood comply with the model to make T valid. Models of this kind not only allow to tell that a set of transmissions will be successful by only locally looking at the transmissions, but they are also essential for the design of local algorithms. They can be seen as a rule for every node that can only observe nearby nodes, either during a setup phase or, more importantly maintaining a dynamic link transmission schedule. The geometric graph-based models mentioned above quite naturally have this property, but SINR_G models do not, which proved to be one of the main obstacles when tackling scheduling problems in these models. This holds for existing centralized approximation algorithms which try to break the interwoven dependencies into independent subproblems as in [8], and it inherently does so in distributed settings – how could nodes come up with a provably valid schedule with local communication, when the validity of a schedule cannot be judged locally? Local interference models on the other hand seem to be incorrect by design: They are blind for interference that arises from nodes that are far away, and thus cannot factor what these nodes are doing. From this time on, let $\mathcal{M}_G = (K, \eta, \beta, \alpha)$ be a standard SINR_G model. We start with an observation which illustrates the first limitations local reasoning about interference implicates. It is a generalization of the considerations above.

Observation 1. Let \mathcal{M}_L be a \mathcal{M}_G -conservative, ρ -local interference model. The following two inequalities hold:

$$\ell(\mathcal{M}_L) < \frac{\rho}{1 + \sqrt[\alpha]{\beta}} \quad \text{and} \quad \ell(\mathcal{M}_L) < \ell(\mathcal{M}_G) \cdot \left(1 + \frac{3K\rho}{\eta\rho^\alpha}\right)^{-1/\alpha} \quad (3)$$

Proof. Let $T \in \mathcal{M}_L$ be any set of transmissions that is accepted by the local model and $t = (s, r)$ any transmission in T . Let $\ell_t := d_{sr}$. By the subset acceptance property (1) and the fact that \mathcal{M}_L is geometric (and thus invariant under isometries of the plane), \mathcal{M}_L would accept any set of transmissions T in which all senders have pairwise distances of more than ρ and all transmissions from T have length ℓ_t .

Now, consider a set of transmissions T , as depicted in Fig. 2, where senders are placed on a triangular grid with edge length $\bar{\rho} := \rho + \epsilon$, i. e., which complies with the considerations above. First, if we assume that $\ell_t \geq \rho/(1 + \sqrt[3]{\beta})$, the interference of the sender s_2 alone would interfere with the reception of the transmission t to the limit,




Fig. 2. Lower bounding interference in ρ -local models

$$\lim_{\epsilon \rightarrow 0} \frac{Kp\bar{\rho}^{-\alpha}(1 + \sqrt[3]{\beta})^\alpha}{Kp\bar{\rho}^{-\alpha}(1 - 1/(1 + \sqrt[3]{\beta}))^{-\alpha}} = \beta ,$$

and together with the additional interference caused by other senders, reception would become impossible, contradicting with the choice of t .

Second, as we now know that $\ell_t < \rho/2$, we get that the interference of the senders s_1 , s_2 and s_3 at the receiver is at least $3Kp\rho^{-\alpha}$. Thus, since $T \in \mathcal{M}_G$,

$$\frac{Kp\ell_t^{-\alpha}}{3Kp\rho^{-\alpha} + \eta} \geq \beta \Leftrightarrow \ell_t \leq \left(\frac{\beta\eta}{Kp} + \frac{3\beta}{\rho^\alpha} \right)^{-1/\alpha} = \ell(\mathcal{M}_G) \cdot \left(1 + \frac{3Kp}{\eta\rho^\alpha} \right)^{-1/\alpha} ,$$

which concludes the proof.

Note that this bound is by no means tight, but it shows how severe the restrictions are that one can only overcome by globally solving schedule problems: To allow for longer links, especially of lengths close to $\ell(\mathcal{M}_G)$, the radius ρ has to be chosen accordingly. We can derive better bounds by calculating interferences more accurately than above by summing up interference for more senders on the same triangular grid. Fig. 3 shows an exemplary tradeoff between the maximum link length needed, $\ell(\mathcal{M}_L)$ and the resulting analytical and numerical lower bounds on the radius ρ for $\alpha = 4$, $\beta = 4$ (≈ 6 dB) and η, p, K normalized to $\ell(\mathcal{M}_G) = 1$. It shows that in the case that we do not assume that nodes can communicate with nodes outside their transmission radius, e. g., by the assumption that the node density is sufficiently high, no link length longer than 40% of the maximum link length can safely be scheduled in realistic scenarios.




Fig. 3. Lower bounds on ρ ($\alpha = \beta = 4.0$)

The second observation we can make about local interference models regards the case that nodes cannot send with arbitrarily low power:

Observation 2. *Let \mathcal{M}_L be a \mathcal{M}_G -conservative ρ -local interference model for a $\rho < \infty$. Even for requests with $\ell(\mathcal{Q}) \leq \ell(\mathcal{M}_L)$, optimal solutions to SCHEDULE and ONESHOTSCHEDULE in \mathcal{M}_L can be arbitrarily worse than in \mathcal{M}_G .*

For the sake of brevity, we will only give a sketch of the proof here. We look at a request of a ring of n transmissions as depicted in Fig. 4 with sufficiently small transmission lengths $\ell(g)$ plus one transmission t^* of length $\ell(n)$ in the middle. It is easy to see that in \mathcal{M}_G , it is admissible to schedule all transmissions but t^* to the same slot (and t^* to a second). In \mathcal{M}_G , assigning a slot to t^* and to the rest of transmissions is independent. Thus, at no time more than a constant number of the n outer transmissions can be carried out, allowing for concurrent transmission of t^* .




Fig. 4. Ring of transmissions

5 $\Omega(1)$ -Sender Model

In every meaningful local model, acceptance of a (local) set of transmissions must follow this consideration: Given the rules for local acceptance of a set of transmissions – is it guaranteed that if all nodes obey these local rules, no node possibly has to accept more interference from outside the ρ -neighborhood than allowed, given the amount of interference arising from local transmissions. A quite straightforward implementation of this concept is the following: For some function $\mu : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, which serves as an upper bound for interference from far away nodes, a set of transmissions T is licit if for every transmission (s, r) a *local signal-to-noise-plus-interference* condition holds:

$$\frac{Kpd_{sr}^{-\alpha}}{\sum_{(\hat{s}, \hat{r}) \in T(s, \rho)} Kpd_{\hat{s}, r}^{-\alpha} + \eta + \mu(d_{sr})} \geq \beta, \tag{4}$$

and if it is guaranteed that a transmission (s, r) cannot receive more interference than $\mu(d_{sr})$ from senders further away than ρ from s . One way to guarantee the latter is to prohibit that close senders are transmitting concurrently and thus, to limit the density of active senders:

Definition 4 ($\Omega(1)$ -sender model). *In the $\Omega(1)$ -sender model $\mathcal{M} = (\rho, c, \mu)$, a set of transmissions T is valid if and only if for every $(s, r) \in T$ equation (4) holds, and any two senders in T have distance at least c .*

Such a model clearly is ρ -local if $c \leq \rho$, but quite obviously not \mathcal{M}_G -conservative for an arbitrary μ . However, for certain values of ρ, c , and μ , the resulting model (ρ, c, μ) is \mathcal{M}_G -conservative and local:

Lemma 1. *Let $\mathcal{M}_L = (\rho, c, \mu)$ be an $\Omega(1)$ -sender model. \mathcal{M}_L is conservative with respect to \mathcal{M}_G if¹*

$$\mu(\ell) \geq \frac{\sqrt{12}Kp\pi\zeta(\rho^2/c^2 + 2\rho/c)}{(\rho - \ell)^\alpha} =: \mu_1(\ell) \tag{5}$$

¹ For the Riemannian ζ -function and $\zeta := \zeta(\alpha - 1)$, a constant $1 < \zeta < 2$ for $\alpha \geq 3$.

Proof. Let (s, r) be some sender/receiver pair with $d_{sr} = \ell$. We divide the plane into annuli A_k with center s and radii $k\rho$ and $(k+1)\rho$ for $k \in \mathbb{N}$. The maximum number of senders lying within the k th annulus is the maximum number of disks of radius c within an annulus with radii $k\rho - c$ and $(k+1)\rho + c$. Since senders in $\Omega(1)$ -sender models form a Minkowski arrangement, which cannot exceed a density of $\frac{2\pi}{\sqrt{3}} \approx 3.638$ [20], we get that the number of senders in A_k is at most

$$N_k := \left\lfloor \frac{2\pi}{\sqrt{3}} \cdot \frac{\pi \left(((k+1)\rho + c)^2 - (k\rho - c)^2 \right)}{\pi c^2} \right\rfloor \leq k \underbrace{\sqrt{12}\pi (\rho^2/c^2 + 2\rho/c)}_{:=N^*},$$

The interference received from any of the senders in A_k can be bounded by

$$I_k := Kp(k\rho - \ell)^{-\alpha} \leq k^{-\alpha} \underbrace{Kp(\rho - \ell)^{-\alpha}}_{:=I^*},$$

and the total interference received from any sender can then be bounded by $\sum_{k=1}^{\infty} N_k I_k \leq N^* I^* \sum_{k=1}^{\infty} k^{-\alpha+1} = N^* I^* \zeta$.

Let (ρ, c) denote a shortcut for the \mathcal{M}_G -conservative model (ρ, c, μ_1) . Obviously, the bound μ_1 can very straightforward be replaced by a better numerical bound. Fig. 5(a) shows how these bounds compare to each other and to the lower bound from the last section. Note that all bounds are correct, and, given the SINR parameters, easy to calculate. We will use the closed-form result for further analysis and the improved bounds for simulation.

With the approximation above, we still have the choice of first the maximum possible link length and second, the balance of the locality factor ρ and the exclusion radius c .

Corollary 1. *Let $\ell = u \cdot \ell(\mathcal{M}_G)$ for some $0 < u < 1$ be a link length and $a \geq 1$, the \mathcal{M}_G -conservative $\Omega(1)$ -sender model $\mathcal{M}_L = (\rho, \rho/a)$ with*

$$\rho = \left(1 + \sqrt[\alpha]{\frac{\sqrt{12}(a^2 + 2a)\pi\beta\zeta}{1 - u^\alpha}} \right) \cdot \ell$$

has $\ell(\mathcal{M}_L) = \ell$. It is graph-based for $a = 1$, yielding (ρ, ρ) .


Proof. According to Lemma 1, \mathcal{M}_L is \mathcal{M}_G -conservative. It remains to show that $\ell(\mathcal{M}_L) = \ell$. First, by $\ell = u \sqrt[\alpha]{Kp/\beta\eta}$, we observe that

$$\rho - \ell = \sqrt[\alpha]{\frac{\sqrt{12}(a^2 + 2a)\pi\beta\zeta}{1 - u^\alpha}} \cdot \ell = \sqrt[\alpha]{\frac{\sqrt{12}(a^2 + 2a)\pi Kp\zeta}{\eta(u^{-\alpha} - 1)}} \quad (6)$$

and therefore by the definition of u

$$\ell(\mathcal{M}_L) = \sqrt[\alpha]{\frac{Kp}{\beta\eta + \sqrt{12}(a^2 + 2a)\pi Kp\beta\zeta (\rho - \ell)^{-\alpha}}} = \sqrt[\alpha]{\frac{Kp}{\beta\eta u^{-\alpha}}} = \ell$$

Obviously, it is graph-based for $a = 1$.



(a) Minimum locality as function of maximum link length, analytical and numerical bounds compared to lower bound

(b) Tradeoff between the minimum pairwise sender distance and the locality radius for fixed $\ell(\mathcal{M}_L)/\ell(\mathcal{M}_G)$

Fig. 5. Tradeoffs for $\Omega(1)$ -sender model (ρ, ρ)

This corollary in a way justifies the work that has been done on scheduling problems in graph-based models as it provides a very simple graph-based model that is provably correct with respect to the geometric SINR models (and, at the same time shows the price for reducing an SINR_G model to a graph-based model). Fig. 5(b) shows this tradeoff between the maximum schedulable link length and the respective ρ for different ratios. As argued in Section 4, no local model can allow for “good” solutions in the single-power setting in the sense that the scheduling problems can be approximated within a constant factor in any such model if links can be arbitrarily short. We will show that \mathcal{M}_G -conservative $\Omega(1)$ -sender models (ρ, ρ) are only by constant and comparably small factors worse than any local model in two dimensions – first the locality needed to allow for a given maximum link length and second the quality of optimal solutions to the scheduling problems.

Lemma 2. *Let $\mathcal{M}_L = (\rho, \rho)$ be the \mathcal{M}_G -conservative $\Omega(1)$ -sender model according to Corollary 1 for some ℓ . Then for any ρ' -local model \mathcal{M}'_L which is \mathcal{M}_G -conservative,*

$$\frac{\rho}{\rho'} \leq \frac{1}{\sqrt[\alpha]{3\beta}} + 2\sqrt{3}\pi\zeta$$

Proof. From (3) and with $u := \ell/\ell(\mathcal{M}_G)$, we get that

$$\rho' = \sqrt[\alpha]{3\beta} \left(\frac{1}{\ell(\mathcal{M}_L)^\alpha} - \frac{1}{\ell(\mathcal{M}_G)^\alpha} \right)^{-1/\alpha} = \ell \sqrt[\alpha]{3\beta} (1 - u^\alpha)^{-1/\alpha}$$

and, from Corollary 1 and $u < 1$,

$$\rho = \left(1 + \sqrt[\alpha]{\frac{6\sqrt{3}\pi\beta\zeta}{1 - u^\alpha}} \right) \cdot \ell < (1 - u^\alpha)^{-1/\alpha} \cdot \left(1 + \sqrt[\alpha]{6\sqrt{3}\pi\beta\zeta} \right) \cdot \ell,$$

which directly implicates the claimed approximation.

This bound depends only on α and β , and is thus constant for a fixed SINR_G model. For a given set of SINR_G parameters, this approximation ratio can be improved using the non-closed-form lower bounds for local models and upper bounds for the c -distant sender model. E.g., for the exemplary values used throughout this paper, the best bounds guarantee a ratio of less than $5/4$ for arbitrary ℓ .

Lemma 3. *Let $\mathcal{M}_L = (\rho, \rho)$ be the \mathcal{M}_G -conservative $\Omega(1)$ -sender model according to Lemma 1 and Corollary 1. Let \mathcal{Q} be a schedule request with $\ell(\mathcal{Q}) < \ell(\mathcal{M}_L)$. Optimal solutions to SCHEDULING and ONE-SHOT-SCHEDULE in \mathcal{M}_L are only by a constant factor worse than optimal solutions in any other ρ -local \mathcal{M}_G -conservative model.*

Proof. First we show that any ρ -local model \mathcal{M}'_L with $\ell(\mathcal{M}'_L) \geq \ell$ cannot accept any set of transmissions T such that for any transmission (s, r, p) , $T(s, \rho)$ contains more than $h := 4^\alpha 6\sqrt{3}\zeta$ transmissions. To this extent, let T be a set of h transmissions such that $T(s, \rho) = T$ for some $(s, r) \in T$. We add a transmission (s', r') to T with $d_{s, s'} = 2\rho$, pointing towards s . Note that if T is valid in \mathcal{M}'_L , then $T' = T \cup \{(s', r')\}$ must be valid, too. But if all transmissions in T' are carried out simultaneously, the SINR-level at r' is below

$$\frac{Kp\ell^{-\alpha}}{hKp(2\rho)^{-\alpha} + \eta} = \frac{\ell^{-\alpha}}{h4^{-\alpha} \frac{\ell^\alpha(1-u^\alpha)}{6\sqrt{3}\pi\beta\zeta} + \frac{\eta}{Kp}} = \frac{\ell^{-\alpha}}{\frac{\ell^{-\alpha}(1-u^\alpha)}{\beta} + \frac{\eta}{Kp}} = \beta .$$

Now take any schedule request \mathcal{Q} . Let A be the square with side $\rho/\sqrt{2}$ that contains the most senders in \mathcal{Q} . Let m denote this number. Since in every ρ -local model at most h of the senders in A can transmit concurrently, leading to a schedule length of at least $\lceil m/h \rceil$. In \mathcal{M}_L , in turn, we can construct a schedule of length $4m$ by the same construction as in [8]: We extend the square to a grid of grid-length $\rho/\sqrt{2}$, 4-color the grid cells, and cyclically choose a color and pick an unscheduled sender from each cell with that color. This guarantees a $4h$ -approximative schedule compared to an optimal solution in any ρ -local model. Similar arguments lead to a $4h$ -approximation of ONE-SHOT-SCHEDULE: Take an optimal solution T in any ρ -local model and picture a 4-colored grid with grid-length $\rho/\sqrt{2}$. Focus only on grid-squares that contain a sender of T . Each of the squares contains at most h transmissions in T . Now pick the color of the most non-empty squares and pick one transmission of each square. This set of transmissions contains at least $\lceil |T|/4h \rceil$ transmissions that can all be carried out concurrently in \mathcal{M}_L . We get a slightly worse approximation for greedy scheduling, where we will only look at the SCHEDULING problem: In the very same grid as above, if one cell does not contain an active sender in a slot, then for two possible reasons: First, since all senders have been scheduled to earlier slots, and second, because of some active sender in one of the adjacent cells. Thus, greedy scheduling uses at most $9m$ slots, which is $9h$ -approximative.

6 Implementation and Simulation Results

We implemented a very basic scheduling algorithm for $\Omega(1)$ -sender models which greedily assigns slots to senders in a random order: Each sender is assigned to the first allowed slot according to the respective model. This is not only a very simple centralized approach, but also a reasonable distributed scheduling algorithm. Given that the node density is sufficient for nodes to have their ρ -neighborhood some constant number of hops away, nodes can draw random numbers and decide on their slot after all neighbors with lower numbers did so only by local communication. This approach is also suited to schedule online requests. We compare the results to three different global scheduling algorithms. First, we select nodes in random order and add them to the first slot allowed by the plain SINR_G model. Second, we compare to the algorithm given in [8]. Please note that this algorithm is not designed to produce good schedulings, but only as a proof of approximability. It is thus not surprising that it returns comparably poor results. Third, since solving the SCHEDULE problem optimally is hard, and solving the corresponding mixed-integer linear problem only works for a very small number of transmissions, we compare to a heuristic, which produced near-optimal results for small instances of random transmission requests. We fill the slots one after another, at any time adding the transmission which causes the least drop of the minimum signal-to-noise-plus-interference ratio for all transmissions earlier added to that slot. We ran all of the above algorithms on schedule requests with at most 80% of the maximum link length in the SINR_G model. Instances were random sets of 20000 transmissions and random unit disk graphs with 5000 nodes on a 50x50 square unit area, i. e. some 10000 edges leading to some 20000 transmissions to schedule links symmetrically. For the $\Omega(1)$ -sender models, we compare three configurations. First, the graph-based (ρ, ρ) model with minimum ρ , second a $(\rho, \ell(\mathcal{M}_G))$ model with minimum ρ and third, a $(\rho, \ell(\mathcal{M}_G))$ model for a ρ higher than necessary. Additionally, we compare to a variant of the $\Omega(1)$ -sender models, where the locality radius ρ is centered at the receiver, where the interference occurs. We call this receiver-centered locality. This class of models never is

Table 1. Comparison of scheduling in different interference models

Algorithm	locality	excl. rad.	random links		
	$[\ell(\mathcal{M}_G)]$	$[\ell(\mathcal{M}_G)]$	length	max. util.	avg. util.
Greedy scheduling in \mathcal{M}_G	∞	0	71.91	482.45	278.43
Intelligent scheduling in \mathcal{M}_G	∞	0	34.15	2726.00	586.11
Goussevskaia et al. [8]	∞	0	605.03	206.01	33.06
Greedy scheduling in ...					
... (ρ, ρ) , min. ρ (graph-based)	2.82	2.82	117.38	224.31	170.45
... $(\rho, \ell(\mathcal{M}_G))$, min. ρ	4.59	1.00	58.63	626.59	341.31
... $(\rho, \ell(\mathcal{M}_G))$, incr. ρ	6.00	1.00	47.89	629.41	417.86
... rc.-local (ρ, r) , min. ρ	2.55 (rc)	1.75	60.88	487.68	328.79
... rc.-local $(\rho, \ell(\mathcal{M}_G))$, min. ρ	3.82 (rc)	1.00	54.84	675.81	364.91
... rc.-local $(\rho, \ell(\mathcal{M}_G))$, incr. ρ	6.00 (rc)	1.00	46.48	646.45	430.50




Fig. 6. Utilization of slots for different scheduling algorithms, for 20000 random links (left) and random UDGs with ≈ 20000 links, averaged over 250 runs. Plotted are results for the graph-based $\Omega(1)$ -sender models $(2.82\ell(\mathcal{M}_G), 2.82\ell(\mathcal{M}_G))$ and $(4.59\ell(\mathcal{M}_G), \ell(\mathcal{M}_G))$ and the model $(2.55\ell(\mathcal{M}_G), 1.75\ell(\mathcal{M}_G))$ for the receiver-centered locality. Values for ρ and c are minimal in the sense that by decreasing these parameters we cannot prove correctness with respect to \mathcal{M}_G . Results are compared to the following global algorithms : greedy scheduling, intelligent scheduling and the algorithm from [8].

graph-based, but allows for better bounds and schedules. Table 1 shows values for random links, averaged over 250 runs, a selection is also plotted in Fig. 6 together with results from scheduling UDG links. Not surprisingly, the more far-seeing global algorithm performs best among all compared schemes and the global algorithm from [8] by far worst. Among the greedy schedule algorithms, the global view did in general not give an advantage. Greedy scheduling in the SINR_G model only outperformed the graph-based variant (whose big advantage is its simplicity). Increasing locality a little more or switching to the receiver-centered locality, the local models even led to better results since the exclusion radius prevented scheduling of close links and receiver-centered locality reflects the nature of interference better.

7 Conclusion and Future Work

In this work, we introduced the concept of local interference models capturing the natural demand for scalable and distributed scheduling protocols to have a local yet provably correct characterization of successful concurrent transmissions. We prove lower bounds that arise in those models and introduce a very simple class, $\Omega(1)$ -sender models which provably allow for constant approximation compared to any other local interference model that produces correct results with respect to the widely agreed on geometric SINR model. We believe that these models will be helpful to attack other local problems such as the problem to construct good topologies not only with respect to properties of spanning ratios and low degree or other heuristics to minimize interference, but with respect to the local construction of short schedules of these topologies. To this extent, it will also be of interest to generalize the findings of this work to the case of variable power

assignment, which, unfortunately is not well understood even with respect to global schedule algorithms.

References

1. Schmidt, S., Wattenhofer, R.: Algorithmic Models for Sensor Networks. In: 20th IEEE Int. Parallel and Distributed Processing Symposium (IPDPS 2006), pp. 450–459 (2007)
2. Rappaport, T.: Wireless Communications: Principles and Practices. Prentice-Hall, Englewood Cliffs (1996)
3. Gupta, P., Kumar, P.R.: The Capacity of Wireless Networks. *IEEE Transactions on Information Theory* 46(2), 388–404 (2000)
4. Moscibroda, T., Wattenhofer, R., Weber, Y.: Protocol Design Beyond Graph-Based Models. In: Proc. of the 5th Workshop on Hot Topics in Networks (HotNets) (2006)
5. Hunt, H., Marathe, M., Radhakrishnan, V., Ravi, S., Rosenkrantz, D., Stearns, R.: NC-Approximation Schemes for NP- and PSPACE-Hard Problems for Geometric Graphs. *Journal of Algorithms* 26 (1998)
6. Krumke, S.O., Marathe, M., Ravi, S.S.: Models and approximation algorithms for channel assignment in radio networks. *Wireless Networks* 6, 575–584 (2000)
7. Björklund, P., Värbrand, P., Yuan, D.: A Column Generation Method for Spatial TDMA Scheduling in Ad hoc Networks. *Ad Hoc Networks* 2(4), 4005–4418 (2004)
8. Goussevskaia, O., Oswald, Y.A., Wattenhofer, R.: Complexity in Geometric SINR. In: Proceedings of the 8th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC 2007), pp. 100–109. ACM Press, New York (2007)
9. Locher, T., von Rickenbach, P., Wattenhofer, R.: Sensor Networks Continue to Puzzle: Selected Open Problems. In: Rao, S., Chatterjee, M., Jayanti, P., Murthy, C.S.R., Saha, S.K. (eds.) ICDCN 2008. LNCS, vol. 4904, pp. 25–38. Springer, Heidelberg (2008)
10. Hajek, B., Sasaki, G.: Link Scheduling in Polynomial Time. *IEEE Transactions on Information Theory* 34(5), 910–917 (1988)
11. Kumar, V.S.A., Marathe, M.V., Parthasarathy, S., Srinivasan, A.: End-to-end packet scheduling in wireless ad-hoc networks. In: Proc. of the 15th annual ACM-SIAM symposium on Discrete Algorithms (SODA 2004), pp. 1021–1030 (2004)
12. Moscibroda, T., Wattenhofer, R.: Coloring Unstructured Radio Networks. In: Proc. of the 17th Annual ACM Symp. on Parallel Algorithms and Architectures (SPAA 2005) (2005)
13. Behzad, A., Rubin, I.: On the Performance of Graph-based Scheduling Algorithms for Packet Radio Networks. In: Proc. of the IEEE Global Telecommunications Conference (GLOBECOM) (2003)
14. Grönkvist, J., Hansson, A.: Comparison Between Graph-Based and Interference-Based STDMA Scheduling. In: Proc. of the 2nd ACM International Symposium on Mobile Ad Hoc Networking & Computing (MOBIHOC), pp. 255–258 (2001)
15. Moscibroda, T., Wattenhofer, R.: The Complexity of Connectivity in Wireless Networks. In: Proceedings of the 25th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2006) (2006)
16. Moscibroda, T., Oswald, Y.A., Wattenhofer, R.: How Optimal are Wireless Scheduling Protocols? In: Proceedings of the 26th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2007) (2007)

17. Linial, N.: Locality in Distributed Graph Algorithms. *SIAM Journal on Computing* 21, 193–201 (1992)
18. Naor, M., Stockmeyer, L.: What can be computed locally? *SIAM Journal on Computing* 24(6), 1259–1277 (1995)
19. Kuhn, F., Moscibroda, T., Wattenhofer, R.: The Price of Being Near-Sighted. In: *Proc. of the 17th ACM–SIAM Symp. on Discrete Algorithms (SODA 2006)* (2006)
20. Toth, L.F.: Minkowskian distribution of discs. *Proceedings of the AMS* 16(5), 999–1004 (1965)