Local Broadcasting in the Physical Interference Model

Olga Goussevskaia
Computer Engineering and Networks Laboratory
ETH Zurich
8092 Zurich, Switzerland
golga@tik.ee.ethz.ch

Thomas Moscibroda
Systems and Networking Research Group
Microsoft Research
Redmond, WA 98052, USA
moscitho@microsoft.com

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

ABSTRACT

In this work we analyze the complexity of local broadcasting in the physical interference model. We present two distributed randomized algorithms: one that assumes that each node knows how many nodes there are in its geographical proximity, and another, which makes no assumptions about topology knowledge. We show that, if the transmission probability of each node meets certain characteristics, the analysis can be decoupled from the global nature of the physical interference model, and each node performs a successful local broadcast in time proportional to the number of neighbors in its physical proximity. We also provide worst-case optimality guarantees for both algorithms and demonstrate their behavior in average scenarios through simulations.

Categories and Subject Descriptors

F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems—*Geometrical Problems and Computations, Sequencing and Scheduling*

General Terms

Algorithms, Theory.

1. INTRODUCTION

Achieving efficient spatial reuse is a fundamental issue in wireless networks. Spatial reuse can be investigated from different angles: communication theorists, for example, study the capacity of a wireless network in fading channel models, such as the signal-to-interference-plus-noise ratio (SINR) model; protocol engineers, on the other hand, design media access protocols with high spatial reuse. In this paper we study a fundamental problem of both theoretical and practical interest: the *local broadcasting problem*. Local broadcasting is an operation used as a building block for many higher-layer protocols (such as routing, synchronization, or coordination protocols) in wireless ad-hoc and sensor networks. As a consequence, the time required to successfully transmit a message to all neighbors in the physical proximity of a node frequently lower bounds and often dominates the overall performance of such critical higher-layer protocols.

Local broadcasting is, in essence, the problem of scheduling wireless requests. When analyzing the performance of a scheduling algorithm, a key aspect is the modeling of interference. While a plethora of different interference models have been used in the literature, arguably the most widely adopted interference models have been the *protocol* model and the *physical* model [16] (and their numerous variations).

The *protocol* model is essentially a graph-based representation of a wireless network, which defines a set of Şinterference-edgesŤ, containing pairs of nodes within a certain distance to each other, thus modeling interference as a binary and purely local measure. Scheduling algorithms designed for the protocol model typically employ graph-based techniques (usually an implicit or explicit coloring strategy), which neglect the aggregated interference of faraway nodes.

In the *physical* model, a signal is received successfully if the Signal-to-Interference-plus-Noise-Ratio (SINR)—the ratio of the received signal strength and the sum of the interference caused by nodes sending simultaneously, plus noise—is above a hardware-defined threshold β. In the simplest instantiation of the model, the signal fades with the distance to the power of the path-loss exponent α. Analyzing the performance of a scheduling protocol in the *physical* model is much more intricate than in the *protocol* model, since the notion of a conflict between transmissions cannot be modeled as a binary relation between edges. Whether a certain set of scheduled transmissions is feasible (i.e., can be successfully scheduled simultaneously) can be determined only by considering the SINR at all the receivers, and not by analyzing edge-to-edge mutual interference as was the case with the *protocol* interference model.

In this work we analyze the complexity of local broadcasting in the *physical* interference model. We present two distributed randomized algorithms. To begin with, we study a very simple Aloha-like algorithm that is based on the assumption that each node knows the number of its neighbors, i.e., the number of nodes in geographical proximity. Our second algorithm and its analysis is significantly more involved and constitutes the main contribution of this paper. This algorithm makes no assumptions about topology knowledge, and provably achieves close to optimal performance.

Our analysis reveals some important insight into the structural relationship between the protocol and physical model. In particular, we prove that if the transmission probability of each node meets certain characteristics, the performance

DialM-POMC'08, August 22, 2008, Toronto, Ontario, Canada.

of our algorithms can be decoupled from the global nature of the physical interference model, and each node is capable of performing a successful local broadcast in time proportional to the number of neighbors in its physical proximity. This holds regardless of the density distribution of the nodes in the network. Our analysis also establishes approximation guarantees for both algorithms, showing that their performance is close to optimal even in worst-case situations. We demonstrate their average behavior through simulations.

In addition to the SINR-model, our analysis is based on a particularly harsh model of communication. Specifically, one main characteristics of the model is *asynchronous wake-up*. That is, there is no pre-defined global start time of the algorithm; instead, nodes can wake up (be switched on, etc) at arbitrary times and new nodes can join the network while other nodes have already started executing the algorithm. Our model also assumes that nodes can be placed arbitrarily, possibly even in a worst-case fashion.

The rest of the paper is organized as follows. Sections 2 and 3 describe related work and our communication model, respectively. We begin the analysis by studying a simple algorithm, referred to as *Multi-Hop Aloha* in Section 4. In Section 5 we present our main algorithm, referred to as *SSMA* (*Slow-Start Media Access*). In Section 6, we argue about the approximation ratio of the proposed algorithms. In Section 7 we simulate the algorithms' performance in average networks. Section 8 concludes the paper.

2. RELATED WORK

As already mentioned in the introduction, two important models used to address interference in wireless networks are the (graph-based) *protocol* model and the *physical* model [16].

Graph-based scheduling algorithms usually employ some sort of matching [17, 30] or coloring strategy [22, 23, 26, 32], which neglects the aggregated interference of nodes located farther away. A variety of centralized [17, 26, 30] and decentralized [22, 23, 27] algorithms have been proposed for such models, many presenting approximation guarantees for special-case geometric graphs, such as unit disk graphs or planar graphs [22, 23, 26, 30, 32]. The complexity of many of these problems has been established in [9]. Although these algorithms present extensive theoretical analysis, they are constrained to the limitations of a model that ultimately abstracts away the accumulated interference of a (possibly) large number of distant nodes.

The inefficiency of graph-based scheduling protocols in the *physical interference* (or SINR) model is well documented and has been shown theoretically [20, 24], through simulations [3, 13, 14], and experimentally [25]. Consequently, the question of how to design good protocols specifically designed for the physical communication model has been studied, leading to a large and rich body of literature. For instance, in [4, 7, 8, 34] optimization models and heuristics for the problem of joint scheduling and power control in the SINR model are proposed. A profound discussion of these and other works can be found in [13].

More recently, algorithmic questions in SINR-models have attracted a tremendous amount of attention. In [24], an efficient power-assignment algorithm, which schedules a strongly connected set of links in poly-logarithmic time is presented. In [5], a greedy centralized scheduling algorithm is presented for links with fixed power levels, and an approximation bound is given for the case where nodes are distributed uniformly at random in the plane. There is also very recent work on routing [6], topology control [11], and dominating sets [28] in SINR-based models. To the best of our knowledge, no existing work has studied the time-complexity of local broadcasting in SINR-based models.

Aloha-based MAC schemes have also been analyzed in the SINR model [1, 10, 33]. In contrast to our work, the analysis presented in these papers is primarily based on the assumptions of homogenous and uniformly random node distributions that do not provide any strong worst-case bounds.

Communication models similar to the one used in this work have been studied in the context of multi-hop broadcasting [2], single-hop wake-up problem [12], and clustering [21] in radio networks. As opposed to our work, in [2] all nodes have access to a global clock and start the algorithm simultaneously, whereas in [12] nodes can wake up asynchronously, but only upon successful reception of a message. Moreover, all these models are graph-based.

3. MODEL

The problem can be formulated as follows. Given a set of nodes V, such that each one wishes to locally broadcast a message to all its neighbors within a certain *broadcasting range*, the objective is to schedule all these requests in as few time-slots as possible.

We adopt the signal-to-interference-plus-noise ratio (SINR)-based *physical* model [16]. In this model, a node y successfully receives a message from a sender x if and only if the following condition holds:

$$\frac{\frac{P_x}{d_{xy}^{\alpha}}}{N + I_y} \geq \beta, \tag{1}$$

where P_x is the power level of the transmission, d_{xy} is the distance between nodes x and y, $2 < \alpha \leq 6$ is the path-loss exponent, which depends on external conditions of the medium, $\beta \geq 1$ denotes the minimum signal to interference ratio required for a message to be successfully received, N is the ambient noise, and I_y is the total amount of interference experienced by receiver y. This interference, which is caused by all simultaneously transmitting nodes in the network can be expressed as follows:

$$I_y = \sum_{v \in V \setminus \{x\}} \frac{P_v}{d_{vy}^{\alpha}}. \tag{2}$$

In this work we assume that all nodes transmit with the same power level. This assumption is also referred to as *uniform power assignment scheme* [15]. This kind of power assignment has been widely adopted in practical systems and in the literature [31].

An important aspect of our model the placement of nodes. We assume that nodes can be placed arbitrarily in the plane, possibly in a worst-case fashion (as opposed to uniform random distribution assumption). In practice, networks with heterogenous, non-uniform topologies are quite typical, and protocols should be designed such that they are capable of coping well with such heterogeneous topologies.

In order to reason about our algorithms, we now introduce several new definitions and notation. We define terms *broadcasting range*, *proximity range*, and *transmission range* of a node, all of which are important in the context of our work.

DEFINITION 3.1. *The local broadcasting range R_B of a node x is the distance up to which x intends to broadcast its messages. We refer to the region within this range as broadcasting region B_x and to the number of nodes in it as Δ_x^B. A local broadcast is complete if every node x in the network has transmitted a message to every node in B_x.*

DEFINITION 3.2. *The transmission range R_T of a node x is the maximum distance from which it can receive a clear transmission ($SINR \geq \beta$), assuming no other transmission occurs simultaneously in the network. We refer to the region within this range as T_x and to the number of nodes in it as Δ_x^T. Given fixed power level P and ambient noise N, and assuming zero interference in Equation 1, the transmission range R_T is $R_T \leq \left(\frac{P}{\beta \cdot N}\right)^{1/\alpha}$.*

In addition to these two definitions, we will make use of the novel notion of a *proximity range* R_A, which is a range between the broadcast and transmission range of a node. Intuitively, it describes the distance within which nodes responsible for the most significant part of interference experienced by x are located. The exact definition of the proximity range is determined by parameters α and β of the SINR model, and changes for each of the algorithms (see Equations (4) and (6) for the precise definitions), but in all cases, it is at least twice as big as the broadcasting range ($R_A \geq 2R_B$). We call the region covered by this radius *proximity region* A_x and refer to the number of nodes in it as Δ_x^A.

Finally, we define a *successful local broadcast*.

DEFINITION 3.3. *Consider a transmitter x and a power level P. We define a successful local broadcast to be a transmission of a message, such that it is successfully received by all receivers y located in the local broadcasting region B_x. The successful reception condition is defined in (1).*

The ideas behind the proximity and transmission ranges are reminiscent to those in the *protocol* interference model, where an interference (or carrier sensing) range (maximum distance up to which a node sensing the channel detects an ongoing transmission) and a transmission range (maximum distance up to which a packet can be received) are defined. The proximity range R_A can be viewed as a separator of the deployment area into a "close-in" region (from where the most significant share of interference comes from) and a "far away" region (from where the incoming interference is still significant, but can often be treated as a constant).

In the analysis we show that when the proximity range R_A is carefully chosen, a node can perform a successful local broadcast with high probability whenever it is the only transmitting node in its proximity range. Therefore, in spite of the global nature of the SINR interference model, concurrent local broadcasts are possible when enough spatial separation exists, i.e., the local broadcasting range R_B is sufficiently smaller than the proximity range R_A.

We analyze two topology awareness scenarios:

- *Known competition*: The nodes know the number Δ_x^A of nodes in their proximity range A_x.
- *Unknown competition*: In this more realistic scenario, nodes are clueless about the current number of nodes in close proximity with which they have to compete for the shared medium. However, we assume that all nodes have the same estimate on the total number of nodes in the network $\hat{n} = |V|$.[1] In other words, each node may have between 0 and n nodes in its proximity range, but it does not know how many.

We assume that nodes wake up asynchronously at any time, and new nodes can join at any time during the execution of the protocol. For the sake of the analysis, we assume that time is divided into time-slots. Note that our algorithm does not rely on synchronized time-slots in any way. This would be too unrealistic an assumption, given that nodes do not have access to a global clock and synchronizing time-slots is an expensive task. Assuming a slotted channel in the analysis is justified due to the standard trick which has been introduced in the analysis of slotted vs. unslotted Aloha [29], where it is shown that the two scenarios differ only by a factor of 2.

We conclude the section with some useful facts.

FACT 3.1. *Given a set of probabilities $p_1 ... p_n$ with $\forall i : p_i \in [0, \frac{1}{2}]$, the following inequalities hold:*

$$(1/4)^{\sum_{k=1}^{n} p_k} \leq \prod_{k=1}^{n}(1-p_k) \leq (1/e)^{\sum_{k=1}^{n} p_k}.$$

FACT 3.2. *For all n, t, such that $n \geq 1$ and $|t| \leq n$,*

$$e^t\left(1-\frac{t^2}{n}\right) \leq \left(1+\frac{t}{n}\right)^n \leq e^t.$$

FACT 3.3. *Consider two disks D_1 and D_2 of radii R_1 and R_2, $R_1 > R_2$, we define χ^{R_1,R_2} to be the smallest number of disks D_2 needed to cover the larger disk D_1. Because the limit of the ratio of the area of D_1 to the area of smaller disks D_2 is $2\pi/3\sqrt{3}$ [19], and because all small disks D_2 intersecting D_1 are completely inside the area of radius $R' = R_1 + 2R_2$, it holds that*

$$\chi^{R_1,R_2} \leq \frac{2\pi}{3\sqrt{3}} \cdot \frac{(R_1+2R_2)^2}{R_2^2}.$$

REMARK 3.4. *We assume that the ambient noise level N is upper bounded by a fraction of the maximum tolerable interference level for a successful broadcast ($(I_y + N) \ll P/\beta(R_B)^\alpha$), such that spatial reuse is achievable by concurrent local broadcasts:*

$$N \leq \frac{P}{2\beta(2R_B)^\alpha}. \quad (3)$$

Note that the exact value of the maximum ambient noise level does not influence our analysis in any significant way, the upper bound in (3) is set for the sake of simplicity.

4. KNOWN COMPETITION

4.1 Algorithm

We start the technical part of the paper by analyzing the performance of a simple algorithm, which we call *Multi-Hop Aloha*. *Multi-Hop Aloha* assumes that each node knows the number of nodes Δ_x^A in its proximity range R_A. Then, after

[1]Notice that without this minimal assumption and in absence of a global counter, every algorithm requires at least time $\Omega(n/\log n)$ until a single successful broadcast is achieved, even in a single hop network [18].

waking up, each node x simply transmits with probability $p := 1/\Delta_x^A$ and remains silent with probability $1-p$. Node x repeats this random choice for $\eta\Delta_x^A \log n$ time-slots, where η is a constant to be defined subsequently.

Our goal is to show that, although the SINR model is intrinsically global and the interferences of distant nodes can accumulate and cause collisions, it is possible to guarantee efficient medium access (in particular, local broadcasts) using this simple and completely distributed algorithm. Specifically, in the analysis we prove that with high probability, every node x performs at least one successful local broadcast after $O(\Delta_x^A \log n)$ time-slots.

4.2 Analysis

For the purpose of our analysis, we introduce the concept of *probabilistic interference*, which is the expected value of total interference experienced by a node.

DEFINITION 4.1. *Consider a node $x \in V$. The probabilistic interference at x, ψ_x, is defined as the expected value of interference experienced by x in a certain time-slot.*

$$\psi_x = P \sum_{v \in V \setminus \{x\}} \frac{p_v}{d_{vx}^\alpha},$$

where P is the transmission power, p_v is the sending probability of node v in time-slot t, and d_{vx} is the distance between x and the interfering node v.

In the following lemma we show that, given an upper bound on the sum of sending probabilities inside each broadcasting region $B_v, v \in V$, the probabilistic interference caused by nodes located outside the proximity region A_x of a node x can be bounded by a constant. Given an upper bound on the expected interference coming outside the region A_x, it becomes possible, in a way, to abstract away this interference and to reason mainly about the interference caused by nodes within the proximity range R_A. The analysis in the *physical* interference model then becomes similar to the analysis used in the *protocol* interference model.

LEMMA 4.1. *Consider a node x and its proximity region A_x, of radius R_A. If in a time-slot t, the sum of transmission probabilities inside all broadcasting regions can be bounded by a constant, i.e., if $\sum_{w \in B_v} p_w \leq c, \forall v \in V$, then the probabilistic interference experienced by x, caused by nodes outside region A_x, can be bounded by*

$$\begin{aligned}
\psi_x^{v \notin A_x} &= P \sum_{v \notin A_x} \frac{p_v}{d_{vx}^\alpha} \\
&\leq c \cdot P \left(\frac{\alpha-1}{\alpha-2}\right) 3^3 2^{(\alpha-2)} R_A^{(2-\alpha)} R_B^{-2}.
\end{aligned}$$

PROOF. Consider rings $Ring^l$ of width R_A around x, containing all nodes v, for which $lR_A \leq d_{vx} \leq (l+1)R_A$. The first such layer $Ring^0$ is the proximity region A_x. Consider all nodes $v \in Ring^l$ for some integer $l > 0$. All corresponding broadcasting regions B_v must be located entirely in an extended ring $Ring_+^l$ of area

$$\begin{aligned}
A(Ring_+^l) &= \left[((l+1)R_A + R_B)^2 - (lR_A - R_B)^2\right]\pi \\
&= (2l+1)(R_A^2 + 2R_A R_B)\pi \\
&< (2l+1)(R_A^2 + 2R_A R_B + R_B^2)\pi \\
&= (2l+1)(R_A + R_B)^2\pi \\
&\leq (2l+1)(3/2R_A)^2\pi.
\end{aligned}$$

Each transmitter v in $Ring^l, l \geq 1$ has distance at least lR_A from x, each transmitter $w \in B_v$ has distance $d(w,x) \geq (lR_A - R_B)$ from x. Since $R_B \leq 1/2R_A$ and $l \geq 1$, $d(w,x) \geq lR_A/2$. By applying a standard geometric area argument, we can bound the probabilistic interference $\psi_x^{Ring^l}$ incurred by nodes located in ring $Ring^l, l \geq 1$ as

$$\begin{aligned}
\psi_x^{Ring^l} &= \sum_{v \in Ring^l} \psi_x^v \\
&\leq \frac{A(Ring_+^l)}{A(B_v)} \cdot P \sum_{\substack{w \in B_v, \\ v \in Ring^l}} \frac{p_w}{(lR_A/2)^\alpha} \\
&\leq \frac{(2l+1)}{l^\alpha} \cdot P \cdot c \cdot 3^2 2^{\alpha-2} R_A^{(2-\alpha)} R_B^{-2} \\
&\leq \frac{1}{l^{(\alpha-1)}} \cdot P \cdot c \cdot 3^3 2^{\alpha-2} R_A^{(2-\alpha)} R_B^{-2}.
\end{aligned}$$

Summing up the interferences over all rings yields

$$\begin{aligned}
\psi_x^{v \notin A_x} &< \sum_{l=1}^{\infty} \psi_x^{Ring^l} \\
&\leq c \cdot P \cdot \sum_{l=1}^{\infty} \frac{1}{l^{\alpha-1}} \cdot 3^3 2^{\alpha-2} R_A^{(2-\alpha)} R_B^{-2} \\
&< c \cdot P \cdot \frac{\alpha-1}{\alpha-2} 3^3 2^{\alpha-2} R_A^{(2-\alpha)} R_B^{-2},
\end{aligned}$$

which concludes the proof of the lemma. □

In the following theorem we prove that the algorithm is correct and efficient.

THEOREM 4.2. *After $O(\Delta_x^A \log n)$ time-slots, each node x performs a local broadcast successfully, with probability at least $1 - 1/n^2$. The claim also holds for all nodes with probability at least $1 - 1/n$.*

PROOF. Given the user-defined broadcasting range R_B, we define the proximity range R_A of a node x to be a function of R_B, α and β:

$$R_A = R_B \left(3^3 2^\alpha \beta \cdot \left(\frac{\alpha-1}{\alpha-2}\right)\right)^{\frac{1}{(\alpha-2)}}. \qquad (4)$$

Note that $R_A > 2R_B$, since $\beta \geq 1$ and $2 < \alpha \leq 6$. It follows that if a node y is located inside the broadcasting region of x, then

$$\begin{aligned}
B_x \subset A_{y \in B_x} \quad &\Rightarrow \quad \Delta_x^B \leq \Delta_y^A \Rightarrow \\
p_y = \frac{1}{\Delta_y^A} \leq \frac{1}{\Delta_x^B} \quad &\Rightarrow \quad \sum_{y \in B_x} p_y \leq 1. \qquad (5)
\end{aligned}$$

The main goal is to bound the expected $SINR_{y \in B_x}$ of the intended receiver of x. Consider the proximity region A_y of the receiver y. Using (4), (5) and Lemma 4.1 ($c = 1$), we can bound the probabilistic interference experienced by y caused by nodes located *outside* A_y:

$$\begin{aligned}
\psi_y^{v \notin A_y} &< 1 \cdot P \cdot \frac{\alpha-1}{\alpha-2} 3^3 2^{\alpha-2} R_A^{(2-\alpha)} R_B^{-2} \\
&= \frac{P}{4\beta R_B^\alpha}.
\end{aligned}$$

Given the expected value of interference at the intended receiver y, caused by transmissions outside A_y, we can use

Markov inequality to claim that the probability that the interference at y caused by transmissions outside its proximity region exceeds $2 \cdot \psi_y^{v \notin A_y}$ is less than $1/2$. Consequently, provided that x is the only node transmitting in A_y, with probability $P_{SINR \geq \beta} \geq 1/2$, the $SINR$ at the intended receiver $y \in B_x$ can be lower bounded by

$$SINR_{y \in B_x} \quad \geq \quad \frac{\frac{P}{d_{x,y}^{\alpha}}}{2 \cdot \psi_y^{v \notin A_y} + N} > \beta,$$

which holds since $d_{x,y} \leq R_B$ and ambient noise N is upper bounded by (3).

The probability $P_{none}^{A_y}$ that no node attempts to transmit in the proximity region A_y of y is at least

$$\begin{aligned} P_{none}^{A_y} \quad &\geq \quad \prod_{w \in A_y} (1 - p_w) \\ &\underset{Fact\ 3.1}{\geq} \left(\frac{1}{4}\right)^{\sum_{w \in A_y} p_w} \geq \left(\frac{1}{4}\right)^{\sum_{v \in A_y} \sum_{w \in B_v} p_w} \\ &\underset{Eq.(5)}{\geq} \left(\frac{1}{4}\right)^{\sum_{v \in A_y}} \geq \left(\frac{1}{4}\right)^{\chi^{R_A, R_B}} . \end{aligned}$$

Putting everything together, we define the probability that node x performs a local broadcast successfully at a time-slot as

$$\begin{aligned} P_{\substack{send \\ success}} \quad &\geq \quad P_{SINR \geq \beta} \cdot P_{none}^{A_y} \\ &\geq \quad \left(\frac{1}{2}\right)\left(\frac{1}{4}\right)^{\chi^{R_A, R_B}} . \end{aligned}$$

Since at each time slot each node locally broadcasts successfully with constant probability, the probability P_{fail} that a node does not transmit successfully after $\lambda \lceil \log n \rceil$ time-slots, where $\lambda = 4\Delta_x^A \cdot 4^{\chi^{R_A, R_B}}$, is

$$P_{\substack{fail \\ send}} \quad \leq \quad \left(1 - \frac{1}{2\Delta_x^A}\left(\frac{1}{4}\right)^{\chi^{R_A, R_B}}\right)^{\lambda \lceil \log n \rceil} < \frac{1}{n^2}.$$

Because there are n nodes to be scheduled, the probability that the claim holds for all nodes is at least $P_{all} \geq \left(1 - \frac{1}{n^2}\right)^n \geq \left(1 - \frac{1}{n}\right)$. $\square$

Note that Theorem 4.2 proves that *Multi-Hop Aloha* is not only efficient and provides fast media access, but is also fair, given that each node's schedule depends only on the local parameter Δ_x^A, allowing fast scheduling in low-density areas, regardless of the existence of highly dense regions somewhere else in the network.

5. UNKNOWN COMPETITION

The simple protocol in the previous section crucially depends on nodes knowing the number of neighbors in their proximity. If nodes do not have this information, designing an efficient algorithm becomes substantially more difficult, because nodes do not know at what probability they should transmit. In this section we describe and analyze the *SSMA* (*Slow-Start Media Access*) protocol. Since nodes do not know with how many nodes they have to compete for the medium, we use a technique that allows each node to start with a very low sending probability and exponentially increase it until they make an attempt to transmit or hear a successful broadcast on the channel. The idea is to eliminate conflicts through randomization, but still guarantee fast medium access for all nodes. The only assumption here is that each node has a rough estimate $\hat{n}$ of the total number of nodes in the network. From now on, we will refer to the estimate $\hat{n}$ as n.[2]

The *SSMA* protocol (Algorithm 1) works in rounds, each of which contains $\delta \lceil \log n \rceil$ time-slots. In every time-slot, a node sends with probability p. Starting from a very small value, this sending probability p is doubled in the beginning of every round. For the algorithm to work properly, we must prevent the noise floor (i.e., the sum of sending probabilities) from reaching too high values. Otherwise, too many collisions will occur. Hence, upon making an attempt to send or upon receiving a message (i.e., when $SINR \geq \beta$), a node resets the value of p and starts the incrementing process again. Once a node makes an attempt to broadcast (without knowing whether it was successful or not), it increments a counter. After a node has made $\lambda \lceil \log n \rceil$ attempts, it stops executing the algorithm.

Consider the broadcasting region B_x of a node x. Let t be a time-slot in which a message is sent by a node $y \in B_x$ and received (without collision) by all other nodes $z \in B_x, z \neq y$. We say that a *Drastic Interference Reduction* (DIR) occurs in the broadcasting region B_x in time-slot t, since all nodes decide to reset their sending probability.

Algorithm 1 *SSMA: Slow-Start Media Access*

```
count := 0;
λ := 4 · 4^{(3/2)χ^{R_A,R_B}};
δ := 12 · 4^{(3/2)(1+χ^{R_A,R_B})};
loop
  p := 1/(4n);
  for i := 0 to ⌈log n⌉ do
    p := 2p;
    for j := 0 to δ⌈log n⌉ do
      if (SINR ≥ β) then
        goto line 5; (reset)
      end if
      s := { 1 with probability p
           { 0 with probability (1 − p)
      if (s = 1) then
        transmit();
        count = count + 1;
        goto line 5; (reset)
      end if
    end for
  end for
  if (count > λ⌈log n⌉) then halt; fi
end loop
```

The parameters δ and λ are chosen as to optimize the results and guarantee that all claims hold with high probability. Parameter δ is chosen large enough to ensure that, with high probability, there is a round in which a DIR occurs. Parameter λ is chosen large enough to ensure that each node performs a local broadcast successfully in at least one round with high probability.

[2]Notice that the algorithm's running time depends only poly-logarithmically on the estimate of n. Hence, it degrades only marginally even if the estimate is very inaccurate.

5.1 Analysis

We begin the analysis by defining the proximity range R_A:

$$R_A = R_B \left(3^4 2^{(2\alpha-1)} \beta \left(\frac{\alpha-1}{\alpha-2}\right)\right)^{\frac{1}{(\alpha-2)}} \quad (6)$$

We proceed in the following manner. In Lemma 5.1, we prove that, during the entire execution of the algorithm, the sum of sending probabilities in every broadcasting region B_x is bounded by a constant. In Lemma 5.2, we show that every node x makes $(\lambda \log n)$ attempts to transmit and stops executing the algorithm after $O(\Delta_x^T \log^3 n)$ time-slots, where Δ_x^T is the number of nodes in its transmission region (Def. 3.2). Finally, in Theorem 5.3, we prove that Algorithm 1 is correct and efficient, i. e., after $O(\Delta_x^T \log^3 n)$ time-slots, every node is scheduled successfully, i.e, every node performs a successful local broadcast. All claims hold with high probability.

LEMMA 5.1. *Consider the execution of Algorithm 1. The sum of sending probabilities of nodes in any broadcasting region $B_x, x \in V$ at any time-slot t is upper bounded by*

$$\sum_{y \in B_x} p_y \leq \frac{3}{2}, \quad (7)$$

with probability at least $(1 - 1/n)$.

PROOF. The claim holds in the beginning of execution, since all nodes start with sending probability $1/4n$. Consider a time slot t_1, in which for the first time the sum of sending probabilities exceeds $1/2$ in one of the broadcasting regions, say B_x. We now consider the time interval $\tau = [t_1 \cdots t_1 + \delta\lceil \log n \rceil]$. We first claim that the sum of sending probabilities in the considered interval is at most $3/2$. The claim holds since (1) by choice of t_1, at the beginning of the interval the sum of sending probabilities is at most $1/2$; (2) by definition of Algorithm 1, during the specified interval each node can at most double its sending probability; and (3) there can be only less than n newly awaken nodes, which in $\delta\lceil \log n \rceil$ time slots can achieve sending probability at most $1/2n$ each, yielding

$$\sum_{v \in B_x} p_v^t \leq 2 \cdot \frac{1}{2} + n \cdot \frac{1}{2n} \leq \frac{3}{2}, \quad \forall t \in \tau.$$

Therefore, the following bounds hold for the entire time interval τ:

$$\frac{1}{2} \leq \sum_{v \in B_x} p_v^t \leq \frac{3}{2} \quad \forall t \in \tau \quad (8)$$

$$0 \leq \sum_{\substack{v \in B_y \\ y \neq x}} p_v^t \leq \frac{3}{2} \quad \forall t \in \tau. \quad (9)$$

The second inequality holds because t_1 is the very first time slot in which the sum of sending probabilities exceeds $1/2$. Hence, in each $B_y, y \neq x$, the sum of sending probabilities is at most $3/2$ in the considered time interval. (Otherwise, one of B_y would have reached $1/2$ before B_x and t_1 would not be the first time slot considered).

The proof proceeds by showing that, before the claimed bound is surpassed, the sum of sending probabilities in B_x falls back to less than $1/2$, since, with high probability, a DIR occurs in B_x in the considered interval. Record that a *Drastic Interference Reduction* (DIR) occurs in the broadcasting region B_x in time-slot t when all nodes in B_x decide to reset their sending probability, which happens if every node in B_x either makes an attempt to transmit or receives a clear message ($SINR \geq \beta$).

We proceed by bounding the *probabilistic interference* experienced by a node $z \in B_x$, caused by nodes located outside its proximity region A_z, in interval τ. Using (6), (9), and Lemma 4.1 ($c = 3/2$), we have

$$\begin{aligned}\psi_z^{w \notin A_z} &< \frac{3}{2} \cdot P \left(\frac{\alpha-1}{\alpha-2}\right) 3^3 2^{(\alpha-2)} R_A^{(2-\alpha)} R_B^{-2} \\ &= \frac{P}{4\beta(2R_B)^\alpha}.\end{aligned}$$

By *Markov inequality*, the probability that the interference at $z \in B_x$, caused by transmissions outside its proximity range, exceeds $2 \cdot \psi_z^{w \notin A_z}$ is less than $1/2$. Therefore, with probability $P_{SINR \geq \beta} \geq 1/2$, the signal received by z from transmitter $v \in B_x$ can be lower bounded by

$$SINR_{z \in B_x} > \frac{\frac{P}{(d_{v,z})^\alpha}}{2 \cdot \psi_z^{w \notin A_z} + N} > \beta,$$

which holds since $d_{v,z} \leq 2R_B$ and ambient noise N is upper bounded by (3).

We proceed by calculating the probability that exactly one transmission $(v, z) \in B_x$ occurs:

$$\begin{aligned}P_{one}^{B_x} &\geq \sum_{v \in B_x} \left(p_v \cdot \prod_{\substack{w \in B_x \\ w \neq v}} (1 - p_w)\right) \\ &\geq \sum_{v \in B_x} p_v \cdot \prod_{w \in B_x} (1 - p_w) \\ &\underset{Fact\ 3.1}{\geq} \sum_{v \in B_x} p_v \cdot \left(\frac{1}{4}\right)^{\sum_{w \in B_x} p_w} \\ &\underset{Eq.(8)}{\geq} \frac{1}{2} \left(\frac{1}{4}\right)^{\frac{3}{2}}.\end{aligned}$$

Furthermore, we define the probability that no other node transmits in A_z:

$$\begin{aligned}P_{none}^{A_z} &\geq \prod_{\substack{w \in A_z \\ w \neq z}} \prod_{k \in B_w} (1 - p_k) \underset{Fact\ 3.1}{\geq} \prod_{\substack{w \in A_z \\ w \neq z}} \left(\frac{1}{4}\right)^{\sum_{k \in B_w} p_k} \\ &\underset{Eq.(9)}{\geq} \prod_{\substack{w \in B_x \\ w \neq v}} \left(\frac{1}{4}\right)^{\frac{3}{2}} \underset{Fact\ 3.3}{\geq} \left(\frac{1}{4}\right)^{\frac{3}{2}\chi^{R_A, R_B}}.\end{aligned} \quad (10)$$

Hence, the probability that a DIR occurs in one time slot is

$$\begin{aligned}P_{DIR} &\geq P_{one}^{B_x} \cdot P_{none}^{A_z} \cdot P_{SINR \geq \beta} \\ &\geq \frac{1}{2} \cdot \frac{1}{2} \left(\frac{1}{4}\right)^{\frac{3}{2}(1 + \chi^{R_A, R_B})}.\end{aligned}$$

The probability that a DIR does not occur in the whole interval τ is

$$\overline{P_{DIR}} \leq \left(1 - \frac{1}{4}\left(\frac{1}{4}\right)^{\frac{3}{2}(1+\chi^{R_A, R_B})}\right)^{\delta \log n} < \frac{1}{n^3},$$

where $\delta = 12 \cdot 4^{3/2(1+\chi^{R_A,R_B})}$.

The argument that a DIR occurs with probability $(1 - n^{-3})$ in the critical interval τ is not sufficient, since the number of such intervals could be infinitely large. However, we can bound the total number of intervals using the fact that each node maintains a counter and makes at most $(\lambda \log n)$ attempts to transmit, stopping the execution of the algorithm afterwards. Since there are n nodes, there can be at most $(n \cdot \lambda \log n)$ critical intervals τ during the entire execution of the algorithm. The probability that a DIR occurs in all such intervals is therefore

$$P_{DIR}(\text{all } \tau\text{'s}) \;\geq\; \left(1 - \frac{1}{n^3}\right)^{n\lambda \log n} \;\geq\; \left(1 - \frac{1}{n}\right).$$

□

In the following lemma we prove that the sending probability, although bounded from above as shown in Lemma 5.1, grows quickly enough, allowing each node x to make $\lambda\lceil \log n\rceil$ transmission attempts in time $O(\Delta_x^T \log^3 n)$.

LEMMA 5.2. *Given the number of nodes Δ_x^T in the transmission region T_x of a node x, every node x makes $\lambda\lceil \log n\rceil$ attempts to transmit and stops executing Algorithm 1 after $O(\Delta_x^T \log^3 n)$ time-slots.*

PROOF. The first observation is that, since a node x can only reset its sending probability upon reception of a clear transmission ($SINR \geq \beta$), the reset can only be caused by nodes within its *transmission range* R_T. Given that there are at most $(\Delta_x^T - 1)$ nodes in the transmission region T_x and that each of these nodes makes at most $\lambda\lceil \log n\rceil$ attempts to transmit, node x can reset its sending probability at most $(\Delta_x^T - 1)\lambda\lceil \log n\rceil$ times.

On the other hand, according to the definition of Algorithm 1, every node starts with sending probability $p_0 = 1/(4n)$ and doubles its sending probability after $\delta\lceil \log n\rceil$ consecutive time-slots without resets. Assuming that x does not reset its sending probability, after $\delta\lceil \log n\rceil(\lceil \log n\rceil + 2)$ time slots, x transmits with probability $p = 1$.

Putting everything together, after at most $(\Delta_x^T - 1)\lambda\delta\lceil \log n\rceil^2(\lceil \log n\rceil + 1) + \delta\lceil \log n\rceil(\lceil \log n\rceil + 2) = O(\Delta_x^T \log^3 n)$ time slots, every node makes $\lambda\lceil \log n\rceil$ attempts to transmit and halts the execution of the algorithm. □

Using Lemmas 5.1 and 5.2, we can now prove that Algorithm 1 is correct and efficient.

THEOREM 5.3. *Given the number of nodes Δ_x^T in the transmission region T_x of a node x, every node x performs a local broadcast successfully after $O(\Delta_x^T \log^3 n)$ time-slots with probability at least $1 - 1/n^2$. The bound holds for all nodes with probability at least $1 - 1/n$.*

PROOF. The high probability result is based on the fact that each attempt to transmit has a constant probability of success, i.e., once a node x attempts to transmit, all intended receivers $y \in B_x$ in its broadcasting region will receive the message successfully ($SINR_y \geq \beta$) with constant probability. Since each node makes $\lambda\lceil \log n\rceil$ attempts to transmit, setting λ to high enough a value gives the high probability result.

In Lemma 5.1 we proved that the sum of sending probabilities in every broadcasting region B_x is bounded by 3/2 during the entire execution of Algorithm 1 w.h.p. Using this fact we can apply Lemma 4.1 to bound the *probabilistic interference* experienced by a receiver $y \in B_x$, caused by nodes located outside its proximity range by

$$\psi_y^{w \notin A_y} < \frac{P}{4\beta(2R_B)^\alpha}.$$

As argued earlier, with probability $P_{SINR \geq \beta} \geq 1/2$, the $SINR$ at the intended receiver $y \in B_x$ can be lower bounded by

$$SINR_{y \in B_x} \;\geq\; \frac{\frac{P}{(d_{x,y})^\alpha}}{2 \cdot \psi_y^{w \notin A_y} + N} > 2^\alpha \beta > \beta.$$

Using the result of Lemma 5.1, the probability that the transmission (x, y) is the only one in the proximity range of y can be calculated in the same way as in (10).

Putting everything together, the probability that transmission attempt is successful can be lower bounded by

$$P_{\substack{send\\success}} \;\geq\; P_{SINR \geq \beta} \cdot P_{none}^{A_y} \;\geq\; \frac{1}{2}\left(\frac{1}{4}\right)^{\frac{3}{2}\chi^{R_A,R_B}}.$$

Applying Lemma 5.2, which states that after time $O(\Delta_x^T \log^3 n)$ node x makes $\lambda \log n$ attempts to transmit and the fact that each attempt has constant probability of success, the probability that node x does not broadcast successfully during the entire execution of Algorithm 1 is

$$P_{\substack{fail\\send}} \leq \left(1 - \frac{1}{2}\left(\frac{1}{4}\right)^{\frac{3}{2}\chi^{R_A,R_B}}\right)^{\lambda\lceil \log n\rceil} < \frac{1}{n^2},$$

where $\lambda = 4 \cdot 4^{(3/2)\chi^{R_A,R_B}}$. Because there are at most n nodes, the probability that the claim holds for all nodes is at least $P_{\substack{all\\success}} \geq \left(1 - \frac{1}{n^2}\right)^n \geq \left(1 - \frac{1}{n}\right)$. □

The upper bound on the execution time of the algorithm is proportional to the number of nodes Δ_x^T in the transmission range R_T of each node, which depends on the transmission power level P. Note that, since nodes aim to broadcast messages only to those receivers located within their broadcasting region B_x, and since high power levels require higher energy spending, the power level P should be chosen somehow proportional to the maximum sender-receiver distance, which is R_B. Therefore, R_T/R_B is typically bounded by a constant, and Δ_x^T remains a local property.

6. LOWER BOUND

The algorithms presented in the previous sections achieve local broadcasts in time $O(\Delta_{max}^A \log n)$ and $O(\Delta_{max}^T \log^3 n)$, respectively. We now show that this is close to optimal.

THEOREM 6.1. *Both algorithms schedule all local broadcasts in time at most a poly-logarithmic factor away from the optimum.*

PROOF. Consider a broadcasting region B_x and the number of nodes in it Δ_x^B. A successful broadcast corresponds to a local broadcast within radius R_B around a sender x. Since the receivers inside this area can decode the signal of only one sender at a time, the transmission can succeed only if no other node sends within this area simultaneously. This means that disks of radius R_B do not overlap in the optimum. Therefore, the optimum can schedule only one node in each broadcasting region at a time and, therefore, needs at least Δ^B_{max} time-slots to schedule all nodes, $T_{OPT} \geq \Delta^B_{max}$. *Multi-Hop Aloha* and *SSMA*, on the other hand, need at most $O(\Delta^A_{max} \log n)$ and $O(\Delta^T_{max} \log^3 n)$ time-slots to schedule all broadcasts successfully with high probability. Given that $\Delta^A_{max} \leq \Delta^B_{max} \cdot \chi^{R_A,R_B}$ and $\Delta^M_{max} \leq \Delta^B_{max} \cdot \chi^{R_T,R_B}$, where χ^{R_A,R_B} and χ^{R_T,R_B} are constants defined in Fact 3.3, we have

$$T_{Aloha} \leq T_{OPT} \cdot \chi^{R_A,R_B} \cdot O(\log n), \quad \text{and}$$
$$T_{SSMA} \leq T_{OPT} \cdot \chi^{R_T,R_B} \cdot O(\log^3 n),$$

i.e., our algorithms are only a poly-logarithmic factor away from the optimum. $\square$

7. SIMULATION RESULTS

Our analytical studies show that both algorithms for local broadcasting perform provably well in worst-case scenarios. In this section we use simulations to investigate the performance in the average case, when nodes are distributed uniformly at random in the plane. Our simulations are coded in the Sinalgo [3] simulation framework, which is a packet-level wireless network simulator. The Sinalgo framework can be tuned to model a wide variety of wireless communication models, including the *physical* and the *protocol* models. For our purposes, we used a communication model that accurately captures SINR-based signal propagation in a wireless communication environment, modeling the reception of packets as in Equations 1 and 2. The simulations were set up on a square of area 1000×1000; the number of simulations was chosen in order to reduce the confidence interval to a meaningful value. Due to lack of space, we can only present a small set of simulations in this section.

In Figures 1(a) and 1(b), we evaluate the average and maximum time needed for all nodes to perform a successful local broadcast. The broadcasting range was set to $R_B = 25$, and the total number of nodes was varied from $n = 1000$ to $n = 5000$. The average number of neighbors in a broadcasting region B_x ranged from $\Delta^B_x = 2$ (for $n = 1000$) to $\Delta^B_x = 10$ (for $n = 5000$). The SINR parameters used in the simulations were $\alpha = 6$ and $\beta = 1$, but as we show in Figures 1(c) and 1(d), SSMA is robust to changes in these parameters. In Figure 1(a), it can be seen that the number of time slots needed for a successful broadcast increases with increasing density. In Figure 1(b), we compare the average execution time to the asymptotic bounds presented in the analysis sections. Recall that *Multi-hop Aloha* and *SSMA* have time complexity $O(\Delta^A \log n)$ and $O(\Delta^T \log^3 n)$, respectively. The plotted lines show the hidden constants in the asymptotic bounds, i.e., the ratio of the maximum execution time T_{max} and $\Delta^A \cdot \log n$ (in the simulation of *SSMA*, the transmission range is equal to the proximity range ($\Delta^T = \Delta^A$)). The simulations suggest that, when nodes are distributed uniformly on the plane, the hidden constants are actually very small. Moreover, *SSMA* has similar performance to *Multi-hop Aloha*, even though it uses no information about network topology. Interestingly, the performance of SSMA approaches that of the simple Multi-hop Aloha more closely as the number of nodes in the system (and hence the density) increases.

In Figures 1(c) and 1(d), we analyze the influence of SINR parameters α and β on the average broadcasting time. In Figure 1(c), we use $\beta = 1$ and $\alpha \in \{3, 4, 5, 6\}$. In Figure 1(d), we use $\alpha = 6$ and $\beta \in \{1, 1.5, 2, 2.5, 3\}$. The simulations were performed on $n = 1000$ nodes, and the broadcasting range was set to $R_B = 25$. In Figure 1(c), it can be seen that the performance of *Multi-hop Aloha* strongly depends on the path-loss exponent α. This is due to the fact that the transmission probability is inversely proportional to the number of nodes within proximity range R_A, which decreases with higher path-loss (see Eq. 4). *SSMA*, on the other hand, is less sensitive to the path-loss, given that its transmission probability is not dependent on the topology of the network. In Figure 1(d), it can be seen that, due to the dependency of *Aloha*'s sending probability on β (see Eq. 4), the execution time slightly increases with increasing β. Once again, the influence of β on the performance of *SSMA* is less explicit. Overall, on average, the performance of the *SSMA* protocol was comparable to the performance of *Multi-hop Aloha*, even though the former operates without having topology knowledge.

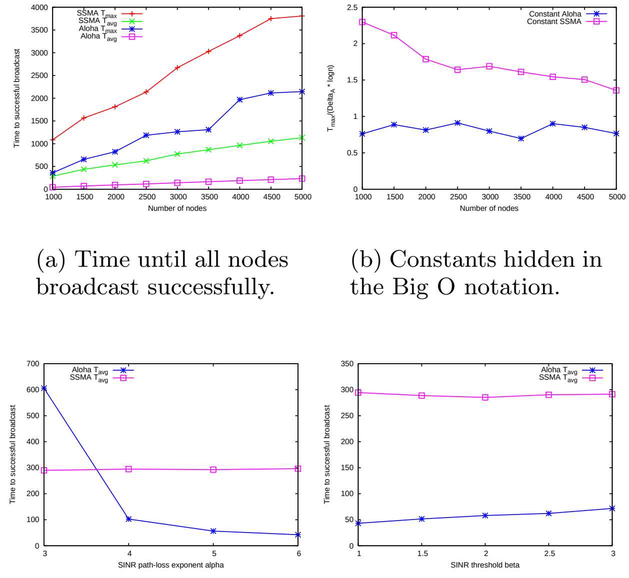

(a) Time until all nodes broadcast successfully.

(b) Constants hidden in the Big O notation.

(c) Influence of the path-loss exponent α.

(d) Influence of the SINR threshold β.

Figure 1: Simulation Results.

8. CONCLUSION

In this work we aim to shed new insight into the complexity of a wireless communication primitive such as local broadcasting in the physical interference model. We analyze the performance of two distributed randomized algorithms and prove that, even when limited knowledge about the topology is provided, close to optimum performance can be achieved. Our analysis greatly relies on the observation that, if the transmission probabilities of nodes are carefully set, the global nature of interference in the physical interference model can be separated into "close-in" and "far-away" regions, which allows the analysis to proceed similarly to analysis in graph-based models, such as the protocol model.

[3] `http://dcg.ethz.ch/projects/sinalgo`

We would like to point out that the analysis presented in this work determines asymptotic bounds. However, an accurate modeling of far-away interference is more involved and will significantly impact practical performance of any MAC protocol.

9. REFERENCES

[1] F. Baccelli, B. Blaszczyszyn, and P. Muhlethaler. An Aloha protocol for multihop mobile wireless networks. *Inf. Theory, IEEE Tran. on*, 52(2):421–436, 2006.

[2] R. Bar-Yehuda, O. Goldreich, and A. Itai. On the time-complexity of broadcast in radio networks: an exponential gap between determinism randomization. In *Proc. 6^{th} Symposium on Principles of Distributed Computing (PODC)*, pages 98–108, 1987.

[3] A. Behzad and I. Rubin. On the Performance of Graph-based Scheduling Algorithms for Packet Radio Networks. In *Globecom*, pages 3432–3436, 2003.

[4] S. A. Borbash and A. Ephremides. Wireless Link Scheduling With Power Control and SINR Constraints. *IEEE Tran. on Inf. Theory*, 52(11):5106–5111, 2006.

[5] G. Brar, D. M. Blough, and P. Santi. Computationally efficient scheduling with the physical interference model for throughput improvement in wireless mesh networks. In *Mobicom*, pages 2–13, 2006.

[6] D. Chafekar, V. Kumar, M. Marathe, S. Parthasarathy, and A. Srinivasan. Cross-layer Latency Minimization for Wireless Networks using SINR Constraints. In *Mobihoc*, 2007.

[7] R. L. Cruz and A. Santhanam. Optimal Routing, Link Scheduling , and Power Control in Multi-hop Wireless Networks. In *Infocom*, 2003.

[8] T. A. ElBatt and A. Ephremides. Joint Scheduling and Power Control for Wireless Ad-hoc Networks. In *Infocom*, 2002.

[9] A. Ephremides and T. V. Truong. Scheduling broadcasts in multihop radio networks. *IEEE Trans. Communications*, 38(4):456–460, Apr. 1990.

[10] G. Ferrari and O. K. Tonguz. Performance of ad hoc wireless networks with Aloha and PR-CSMA MAC protocol. In *Globecom*, pages 2824–2829, 2003.

[11] Y. Gao, J. C. Hou, and H. Nguyen. Topology control for maintaining network connectivity and maximizing network capacity under the physical model. In *Infocom*, 2008.

[12] L. Gasieniec, A. Pelc, and D. Peleg. The Wakeup Problem in Synchronous Broadcast Systems. *SIAM J. Discret. Math.*, 14(2):207–222, 2001.

[13] J. Gronkvist. *Interference-Based Scheduling in Spatial Reuse TDMA*. PhD thesis, Royal Institute of Technology, Stockholm, Sweden, 2005.

[14] J. Gronkvist and A. Hansson. Comparison between graph-based and interference-based STDMA scheduling. In *Mobihoc*, pages 255–258, 2001.

[15] P. Gupta and P. R. Kumar. Critical Power for Asymptotic Connectivity in Wireless Networks. In *Stochastic Analysis, Control, Optimization and Applications*, pages 547–566. Birkhauser, Boston, 1998.

[16] P. Gupta and P. R. Kumar. The Capacity of Wireless Networks. *IEEE Trans. on Inf. Theory*, IT-46(2):388–404, 2000.

[17] B. Hajek and G. Sasaki. Link scheduling in polynomial time. *IEEE Trans. on Inf. Theory*, 34(5):910–917, 1988.

[18] T. Jurdziński and G. Stachowiak. Probabilistic Algorithms for the Wakeup Problem in Single-Hop Radio Networks. *LNCS*, 2518:535–549, 2002.

[19] R. Kershner. The number of circles covering a set. *American Journal of Mathematics*, 62, 1939.

[20] T.-S. Kim, H. Lim, and J. C. Hou. Improving Spatial Reuse through Tuning Transmit Power, Carrier Sense Threshold, and Data Rate in Multihop Wireless Networks. In *Mobicom*, 2006.

[21] F. Kuhn, T. Moscibroda, and R. Wattenhofer. Initializing Newly Deployed Ad Hoc and Sensor Networks. In *Mobicom*, 2004.

[22] V. S. A. Kumar, M. V. Marathe, S. Parthasarathy, and A. Srinivasan. End-to-end packet-scheduling in wireless ad-hoc networks. In *SODA*, pages 1021–1030, 2004.

[23] T. Moscibroda and R. Wattenhofer. Coloring Unstructured Radio Networks. In *SPAA*, July 2005.

[24] T. Moscibroda and R. Wattenhofer. The Complexity of Connectivity in Wireless Networks. In *Infocom*, 2006.

[25] T. Moscibroda, R. Wattenhofer, and Y. Weber. Protocol Design Beyond Graph-Based Models. In *Hotnets*, November 2006.

[26] S. Ramanathan and E. L. Lloyd. Scheduling algorithms for multihop radio networks. *IEEE/ACM Trans. Netw.*, 1(2):166–177, 1993.

[27] I. Rhee, A. Warrier, J. Min, and L. Xu. DRAND: distributed randomized TDMA scheduling for wireless ad-hoc networks. In *Mobihoc*, pages 190–201. ACM Press, 2006.

[28] A. Richa, P. Santi, and C. Scheideler. An O(log n) dominating set protocol for wireless ad-hoc networks under the physical interference model. In *Mobihoc*, 2008.

[29] L. G. Roberts. ALOHA Packet System with and without Slots and Capture. *Computer Communications Review*, 5(2), 1975.

[30] G. Sharma, R. R. Mazumdar, and N. B. Shroff. On the complexity of scheduling in wireless networks. In *Mobicom*, pages 227–238, 2006.

[31] S. Singh and C. Raghavendra. PAMAS: Power Aware Multi-Access Protocol with Signalling for Ad Hoc Networks. *SIGCOMM Comp. Comm. Review*, 28(3):5–26, July 1998.

[32] W. Wang, X.-Y. Li, O. Frieder, Y. Wang, and W.-Z. Song. Efficient interference-aware TDMA link scheduling for static wireless networks. In *Mobicom*, pages 262–273, 2006.

[33] S. Weber, X. Yang, J. Andrews, and G. de Veciana. Transmission capacity of wireless ad hoc networks with outage constraints. *Inf. Theory, IEEE Tran. on*, 51(12):4091–4102, 2005.

[34] J. Zander. Distributed cochannel interference control in cellular radio systems. *IEEE Trans. Veh. Technol.*, vol. 41, Aug. 1992.