

Broadcasting in an Unreliable SINR Model

Fabian Kuhn¹ and Philipp Schneider²

1 University of Freiburg, Germany

kuhn@cs.uni-freiburg.de

2 University of Freiburg, Germany

philipp.schneider@cs.uni-freiburg.de

Abstract

We investigate distributed algorithms for broadcasting in unreliable wireless networks. Our basic setting is the *signal to noise and interference ratio* (SINR) model, which captures the physical key characteristics of wireless communication. We consider a dynamic variant of this model in which an adversary can adaptively control the model parameters for each individual transmission. Moreover, we assume that the network devices have no information about the geometry or the topology of the network and do neither know the exact model parameters nor do they have any control over them.

Our model is intended to capture the inherently unstable and unreliable nature of real wireless transmission, where signal quality and reception depends on many different aspects that are often hard to measure or predict. We show that with moderate adaptations, the broadcast algorithm of Daum et al. [DISC 13] also works in such an adversarial, much more dynamic setting. The algorithm allows to broadcast a single message in a network of size n in time $\mathcal{O}(D \cdot \text{polylog}(n+R))$, where D is the diameter and R describes the granularity of the communication graph.

1998 ACM Subject Classification G.2.2 Network problems, F.2.2 Analysis of Algorithms

Keywords and phrases radio networks, wireless networks, broadcast, SINR model, unreliable communication, dynamic networks

Digital Object Identifier 10.4230/LIPIcs.OPODIS.2017.3

1 Introduction

In the *signal to noise and interference* (SINR) model (a.k.a. the *physical interference model*), a message is received if and only if the ratio between the signal strength at the receiving node and the combined strength of the background noise and any interfering signals is above a given threshold. By now, the SINR model has become the standard communication model to study wireless network algorithms. In the distributed algorithms literature, different variants of the basic SINR model have been studied, based on the properties of the underlying geometric space, how much geometric information the network devices have and how much they know about the model parameters or the network topology.

One of the most general variants of the SINR model has been termed the *ad hoc* SINR model by Daum et al. in [4], where they study the problem of broadcasting a message to all nodes of a wireless network. In [4], it is assumed that the network nodes have no information about the geometry or the topology of the network, prohibiting algorithms that utilize advance knowledge about network topology and layout or the way that signals propagate in space. The distances between the nodes are assumed to form a general *growth-bounded* metric space. Furthermore, the nodes use uniform transmission powers, they have only approximate



© Fabian Kuhn and Philipp Schneider;
licensed under Creative Commons License CC-BY

21st International Conference on Principles of Distributed Systems (OPODIS 2017).

Editors: James Aspnes, Alysson Bessani, Pascal Felber, and João Leitão; Article No. 3; pp. 3:1–3:21

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

knowledge of the hardware and model parameters, and they do not have additional capabilities like carrier sensing. Because the model makes only minimal assumptions, it allows to develop algorithms that can operate in a quite general setting.

However, while the authors of [4] assume that the model parameters are only known approximately, these parameters are still assumed to be uniform over the whole network and static over time. As a result, whether a message is successfully received is determined by a *deterministic* function depending on the system parameters and the geometry of the network. However in real wireless networks, communication often turns out to be rather unreliable and highly volatile, and system parameters tend to not be uniform over space and time [22]. As a result, practical wireless communication behaves in an inherently non-deterministic way. Examples for such unpredictable communication behavior are background noise due to coexisting networks or jamming, various multi-path effects due to changes of the environment, or fluctuations in sending power or signal sensitivity among different wireless devices.

In order to improve practical applicability in the present paper, we consider an unreliable variant of the ad hoc SINR model (which we will term the *unreliable SINR model*) by adding a worst-case adversary, which decides to a given extent whether or not messages are received. The adversary adds a dynamic and non-deterministic component to the model that allows to capture non-uniform and dynamically changing SINR parameters. More concretely, we assume that at each point in time and individually for each node, an adversary can adaptively control the threshold on the signal to interference and noise ratio above which a message can successfully be decoded (for a formal definition of the model, we refer to Section 3).

As our main contribution, we adapt the broadcast algorithm of Daum et al. [4] to work in the unreliable SINR model. We show that with relatively modest adaptations, the algorithm of [4] can also be used to efficiently solve the broadcast problem in the unreliable SINR model. More specifically, let G_C be the communication graph in which two nodes are adjacent if they can reliably communicate with each other in the unreliable SINR model (if the signal to noise ratio for the two nodes is sufficiently above the maximum threshold that the adversary is allowed to choose). We prove that the global broadcast problem can be solved in time $\mathcal{O}(D \cdot \text{polylog}(n+R) \cdot \frac{\beta_{\max}}{\beta_{\min}})$, where D is the diameter of the communication graph G_C , R is the ratio of maximum to minimum distance between any two adjacent nodes in G_C ,¹ and $[\beta_{\min}, \beta_{\max}]$ is the range within which the adversary can adaptively choose the SINR threshold in each time slot and for each node.

The paper is structured as follows. In Section 2 we summarize existing work related to broadcasting in the SINR model and non-deterministic distributed algorithms. In Section 3 we specify our model and give some notions used throughout the following sections. Section 4 gives an overview on the neighborhood dissemination process, which conducts one hop of the global message broadcast. In Section 5 to Section 8 we analyze the neighborhood dissemination protocol. Finally, in Section 9, we show that our abstract model, where the adversary can only control the SINR threshold, in fact allows to also capture a more general adversary, which controls all SINR-parameters.

¹ Theoretically, any functional dependency $R = f(n)$ is possible, e.g., for a series of nodes v_1, \dots, v_{n+1} with decreasing distances $d(v_i, v_{i+1}) = 1/f(i)$. However, we consider $R \in \omega(n)$ unlikely in practice. E.g., placing n nodes uniformly at random inside a square with sidelength $r_e/\sqrt{2}$ (where r_e is the effective communication range), yields an expected granularity $R \in \Theta(n)$.

2 Related Work

Broadcasting in the SINR Model. Broadcasting algorithms in conjunction with the SINR model were first investigated in [10]. They consider local broadcast, where each node needs to broadcast a message to all neighbors in the communication graph. Note that local broadcast can be used as a building block to solve global broadcast. Subsequent publications on broadcasting (local and global) in the SINR model include [13, 14, 25, 26]. However, all the aforementioned publications leverage assumptions that are incompatible with our conception of the SINR model. Incongruities include in particular that nodes have knowledge of their position or of distances to other nodes or that nodes can modify their sending power or have carrier sensing capability.

Broadcasting in the Ad Hoc SINR Model. Most relevant for our work are distributed information dissemination algorithms that work in the ad hoc SINR model as discussed in Section 1. The first such algorithm is the local broadcast algorithm of [12], which can be used to solve global broadcast in the ad hoc SINR model in time $\mathcal{O}(D\Delta_C(\log n)^2)$, where Δ_C is the maximum degree of the communication graph G_C . Note that Δ_C is a potentially large factor (e.g., in single-hop networks). In [4] the global broadcast problem is solved directly in time $\mathcal{O}(D \cdot \text{polylog}(n + R))$. Recall that R is the ratio between the largest and smallest distance between neighbors in the communication graph. The solution of [4] forms the basis of our algorithm and we will point out the differences to the algorithm and analysis alongside.

For large and moderately large values of R , the result of [4] has been improved by Jurdzinski, Kowalski, Rozanski, and Stachowiak [15], where it is shown that one can get rid of the dependency on R and solve the global broadcast problem in time $\mathcal{O}(D \log^2 n)$.² The algorithm of [15] is based on finding an assignment of probabilities such that in each local neighborhood, the sum of probabilities is upper and lower bounded by some constant.

Halldorsson et al. [11] use the concept of *abstract Medium Access Control* (abstract MAC) layers introduced in [16] (and subsequently enhanced in [7]), to implement higher-level procedures such as single and multi-message broadcast and consensus in the ad hoc SINR model. The abstract MAC layer provides basic routines where run-times are bounded by delay-functions. By proving upper bounds on delays in the ad hoc SINR model, they establish an abstract MAC layer that permits single-message broadcast in $\mathcal{O}((D + \log n) \text{polylog } R)$, which is an improvement over [4].

The algorithm of [4] is the simplest and seemingly most robust of the global broadcast algorithms in the ad hoc SINR model and we therefore decided to extend this algorithm to the adversarial, unreliable SINR model considered in the present paper. However, it would certainly be interesting to see whether the approaches of [11] and [15] can also be adapted to work in a more dynamic and unreliable setting.

Adversarial Models in Wireless Networks. To the best of our knowledge, the only previous paper that analyzes non-deterministic, adversarial behavior in combination with the SINR model, is by Ogierman et al. [23]. They assume that the SINR noise parameter is controlled by an adaptive adversary with a restricted energy budget during each interval of some fixed length. The objective of [23] is to design a MAC layer that allows to achieve optimal throughput. The algorithm has a constant competitive ratio and it is based on a protocol

² The algorithm of [4] – which takes $\mathcal{O}(D \log n(\log^* n + \log^{\alpha+1} R))$ time – is faster than the algorithm of [15] if $R \leq 2^{\log^{1/c} n}$ for a sufficiently large constant $c > \alpha$ (α is the path loss exponent, see Section 3).

designed in [24] for a simpler wireless network model. We note that the algorithm of [23] heavily relies on the assumption that the nodes can use carrier-sensing.

Non-deterministic behaviour in the broader context of distributed networks has been studied in various settings. Of particular interest in the context of the present paper is the dual graph model introduced by [17]. The dual graph model considers graphs $G = (V, E)$ and $G' = (V, E')$ with $E \subseteq E'$, where G is reliable and G' contains unreliable edges, which are controlled dynamically by an adaptive worst-case adversary. A message can only be delivered from sender to receiver if no other neighbor (in the graph of currently active edges) of the receiver sends. Broadcast in this framework was studied in [8, 17, 19].

Another line of work [1, 2, 5, 9, 20, 24] studies scenarios, where an adversary with a limited energy budget per time may jam the single shared channel or a subset of multiple shared channels, to the effect that no transmissions via the jammed channels are possible. However, all of the aforementioned non-SINR models share the assumption that interference is a strictly local and binary concept (a message is received if and only if exactly one node or neighbor sends) and they all fail to capture the global and continuous interference concept of the SINR model, which is more faithful to reality (cf. [21]).

3 Model and Preliminaries

Communication Model. Let V be a set of n point-shaped communication devices (we call them nodes) embedded to distinct points of a metric space (X, d) . For $u, v \in V$, let $d(u, v)$ denote the distance between the two points in X to which u and v are embedded. Nodes are reliable and have unique identifiers. Time is divided into synchronous time slots, henceforth also called rounds. In each round, each node can either transmit a message with a fixed transmission power P or it can listen to the channel and thereby receive at most one message. In addition, each node may conduct an arbitrary amount of computations during each round. We do not explicitly restrict the size of messages, however, in addition to the size of the data that needs to be broadcast we require at most $\mathcal{O}(\log n)$ additional bits.

Let $I \subseteq V \setminus \{u, v\}$ be the set of nodes other than $u, v \in V$ that are transmitting during a round. Then v receives a message from u , iff u transmits, v listens, and

$$\text{SINR}(u, v, I) := \frac{P/d(u, v)^\alpha}{N + \sum_{w \in I} P/d(w, v)^\alpha} \geq \beta_v. \quad (1)$$

We call the parameters P , N , α , and β_v in Equation 1 the *SINR parameters*. The value of $P > 0$ denotes the transmission power of all nodes. Further, the parameter $N \geq 0$ describes the background noise and $\alpha \geq 2$ specifies the power loss of the transmitted signals. Finally, for each node $v \in V$, the threshold $\beta_v > 0$ determines how large the signal to interference and noise ratio at node v needs to be, such that v successfully decodes a signal.

In our abstract model, we assume that the SINR parameters P , N , and α are fixed and that only the value of β_v is subject to some variability and uncertainty. More specifically, there are generally known lower and upper bounds $\beta_{\min} \geq 1$ and $\beta_{\max} > \beta_{\min}$ on β_v . In each round and for all nodes $v \in V$, the actual value of β_v is determined by a *strongly adaptive adversary* arbitrarily in the range $[\beta_{\min}, \beta_{\max}]$. That is, the adversary can adjust the value β_v for each node and can thereby control to some extent, whether a given message is received or not. We note that to keep our abstract model as simple as possible, we assume that the adversary can only change β_v while all other SINR parameters are fixed (and globally known). In Section 9, we discuss how our abstract model can be used to capture a more powerful adversary that can vary all SINR parameters within a given range.

Communication Graph. We define the *maximum range* $r_m := (P/N\beta_{\min})^{\frac{1}{\alpha}}$, *maximum safe range* $r_s := (P/N\beta_{\max})^{\frac{1}{\alpha}}$ and the *effective range* $r_e := r_s/(1 + \rho)$ for constant $\rho > 0$. For a given $I \subseteq V$ we call a transmission from u to v *safe* if $\text{SINR}(u, v, I) \geq \beta_{\max}$, otherwise the transmission is called *unsafe*. For $d(u, v) > r_m$ no transmissions among u and v are possible. For $r_s < d(u, v) \leq r_m$ only unsafe transmissions are possible. For $d(u, v) \leq r_s$ and especially $d(u, v) \leq r_e$ safe transmissions are possible (e.g., if $I = \emptyset$). Now we are able to formally define the communication graph G_C as $(V, \{\{u, v\} \mid u, v \in V, u \neq v, d(u, v) \leq r_e\})$.

All nodes are assumed to know polynomial bounds on the number of nodes n and on the ratio $R := \frac{d_{\max}}{d_{\min}}$ of the maximum distance d_{\max} to the minimum distance d_{\min} among neighbors in the communication graph G_C . The values of n and R appear only inside log functions. Since it makes no difference in the *asymptotic* running time of our algorithm, we will assume exact knowledge of n and R for simplicity. Further, we assume that the network is *growth-bounded* in the following sense. There exists a constant $\delta \in [1, \alpha)$ such that for any subset of nodes $S \subseteq V$, for any $v \in S$ and any $x \in \mathbb{R}^+$, the number of nodes from S that are within distance $x \cdot d_{\min}^S$ of v is bounded by $\mathcal{O}(x^\delta)$, where $d_{\min}^S := \min_{u, v \in S, u \neq v} d(u, v)$.³

Global Broadcast and Neighborhood Dissemination. In the present paper, we intend to solve the *global broadcast problem*. There is a distinguished source node s that initially has a broadcast message \mathcal{M} . The global broadcast problem is solved when \mathcal{M} is disseminated to all nodes in V . We assume that a node v cannot participate in a global broadcast algorithm before v learns the broadcast message \mathcal{M} . This assumption is sometimes known as *asynchronous* start. We are interested in *randomized* algorithms that solve the global broadcast problem *with high probability* (w.h.p.) that is, with probability at least $1 - 1/n^c$ for a given, sufficiently large constant $c > 0$. The algorithms' guarantees have to hold for *any* strategy of the adversary, i.e., for every possible way in which the adversary may choose β_v for each round and node. We call an algorithm that fulfills these criteria *robust*.

Our solution to the global broadcast problem is based on a solution to the following *neighborhood dissemination* problem. Assume that a subset $S \subseteq V$ of the nodes knows some message \mathcal{M} . Neighborhood dissemination is solved as soon as all neighbors $N(S)$ ⁴ of S in G_C know \mathcal{M} . Note that if D is the diameter of G_C , the global broadcast problem can clearly be solved by running neighborhood dissemination D times.

Spatial Reuse in Dense Networks. In dense networks we can take advantage of the fact that signals fade polynomially with distance. At the same time, the interference from nodes within a distance around v grows not too fast (when said distance is increased) due to the growth bound. These properties potentially allow large numbers of concurrent transmissions in densely packed areas. This is sometimes known as spatial reuse.

The following lemma gives a formalization of this property and is slightly adapted from Daum et al. [4]. Assume that a set of active nodes all transmit with a fixed probability p and that all other nodes are silent. Then there is a *constant* probability that two nodes can communicate safely, given that they are within safe transmission range of each other and are closer than the smallest distance among active nodes times a constant.

³ Intuitively, the growth bound can be interpreted as follows. Assume we grow the radius x of a sphere around a node v . Then the interference caused by nodes at distance x fades faster (i.e. with $1/x^\alpha$) than the number of nodes within that sphere grows (i.e. with $\mathcal{O}(x^\delta)$). This implies that the total interference from distant nodes decreases (the important feature is $\delta < \alpha$), as we show in Lemma 1. This notion is generalized for arbitrary subsets $S \subseteq V$ and 'normalized for granularity' d_{\min}^S .

⁴ $N(S) := \{v \in V \mid \text{there is } u \in S \text{ such that } (u, v) \text{ is an edge in } G_C\}$

► **Lemma 1** (cf. [4]). *Let $S \subseteq V$ be a set of active nodes sending with fixed probability p . All non-transmitting nodes listen. Consider two nodes $u \in S$ and $v \in V$ with $d(u, v) \leq c_1 \cdot r_e$ and $d(u, v) \leq c_2 \cdot d_{\min}^S$, where $0 < c_1 < 1 + \rho$, $c_2 > 0$ are constants and $d_{\min}^S = \min_{x \neq y \in S} d(x, y)$. If v listens, then a safe transmission from u to v takes place with constant probability $\mu \in (0, p)$.*

Proof. We show $\text{SINR}(u, v, I) \geq \beta_{\max}$. For this purpose we compute a bound on $I_{S'}^{\max}(v)$ defined as the maximum interference at v originating from $S' := \{w \in S \mid d(v, w) \geq k_0 \cdot d_{\min}^S\}$ with $k_0 \in \mathbb{N}$. We partition $S' = \bigcup_{k=k_0}^{\infty} R_k$ along concentric rings $R_k := \{w \in S' \mid k \cdot d_{\min}^S \leq d(v, w) < (k+1) \cdot d_{\min}^S\}$ of thickness d_{\min}^S . Moreover, let $B_k := \{w \in S \mid d(v, w) < (k+1)d_{\min}^S\}$ be the nodes within a ball of radius $(k+1)d_{\min}^S$ centered at v .

First we unwrap the definition of the growth bound $|B_k| \in \mathcal{O}((k+1)^\delta) = \mathcal{O}(k^\delta)$: There are constants $k_0, \zeta > 0$ such that $|B_k| \leq \zeta \cdot k^\delta$ for all $k \geq k_0$. This allows us to bound the interference at v stemming from R_{k_0} . In the worst case (regarding the amount of interference from R_{k_0}) all of the at most $\zeta \cdot k_0^\delta$ nodes in B_{k_0} are located in R_{k_0} , which gives us

$$I_{R_{k_0}}^{\max}(v) = \sum_{w \in R_{k_0}} \frac{P}{d(v, w)^\alpha} \leq \zeta \cdot k_0^{\delta-\alpha} \frac{P}{(d_{\min}^S)^\alpha} \in \mathcal{O}(k_0^{\delta-\alpha}) \frac{P}{(d_{\min}^S)^\alpha}.$$

Now we bound the combined interference from the other rings R_k , with $k > k_0$. Again, we look at the worst case regarding interference at v . The maximum interference without violating the growth bound is obtained when as many nodes as possible are as close to v as possible. Therefore, B_{k-1} contains as many nodes as the growth bound allows, implying that $|R_k| \leq \zeta k^\delta - |B_{k-1}| \leq \zeta(k^\delta - (k-1)^\delta) \leq \zeta \delta k^{\delta-1}$ (the last inequality is given in Appendix C Lemma 17). This leads us to

$$\begin{aligned} \sum_{k=k_0+1}^{\infty} I_{R_k}^{\max}(v) &\leq \zeta \delta \sum_{k=k_0+1}^{\infty} k^{\delta-\alpha-1} \frac{P}{(d_{\min}^S)^\alpha} \\ &\leq \zeta \delta \int_{k=k_0}^{\infty} k^{\delta-\alpha-1} dk \frac{P}{(d_{\min}^S)^\alpha} = \zeta \delta \frac{k_0^{\delta-\alpha} P}{(\alpha - \delta)(d_{\min}^S)^\alpha} \in \mathcal{O}(k_0^{\delta-\alpha}) \frac{P}{(d_{\min}^S)^\alpha}. \end{aligned}$$

Combining these results we observe $I_{S'}^{\max}(v) \leq \kappa(k_0)P/(d_{\min}^S)^\alpha$ for a sequence $\kappa(k_0) \in \mathcal{O}(k_0^{\delta-\alpha})$, which approaches 0 for $k_0 \rightarrow \infty$. Assuming all nodes in $S \setminus S'$ remain silent we obtain

$$\begin{aligned} \text{SINR}(u, v, S') &= \frac{\frac{P}{d(u, v)^\alpha}}{N + I_{S'}^{\max}(v)} \stackrel{(*)}{\geq} \frac{P}{\frac{d(u, v)^\alpha P}{r_s^\alpha \beta_{\max}} + \frac{d(u, v)^\alpha \kappa(k_0) P}{(d_{\min}^S)^\alpha}} \quad (*): N = \frac{P}{r_s^\alpha \beta_{\max}} \text{ (cf. def. of } r_s) \\ &\geq \frac{P}{\frac{c_1^\alpha r_e^\alpha P}{(1+\rho)^\alpha r_e^\alpha \beta_{\max}} + \frac{(c_2 d_{\min}^S)^\alpha \kappa(k_0) P}{(d_{\min}^S)^\alpha}} = \frac{\beta_{\max}}{\underbrace{\frac{c_1^\alpha}{(1+\rho)^\alpha}}_{< 1} + \underbrace{c_2^\alpha \beta_{\max} \kappa(k_0)}_{\rightarrow 0, \text{ for } k_0 \rightarrow \infty}} \geq \beta_{\max} \end{aligned}$$

for sufficiently large $k_0 \in \Theta(1)$.⁵ Thus a safe transmission from u to v takes place with probability $\mu \in p(1-p)^{\mathcal{O}(k_0^\delta)}$ defined as the probability that all nodes in $S \setminus S'$ are silent.⁶ ◀

⁵ Note that k_0 does not depend on the size n , active nodes S , nor on the granularity R of the network. Only on SINR parameters, growth bound exponent δ and c_1, c_2 , which we deem all constant.

⁶ Daum et al. [4] propose to choose p such that $\mu(p, k_0)$ is maximized.

Algorithm 1 ROBUSTDISSEMINATION.	\triangleright <i>high level description</i>
$S_1 \leftarrow S$	\triangleright <i>set of active nodes at start of this procedure</i>
for phase $\phi \leftarrow 1$ to $\lfloor \log R \rfloor + 2$ do	\triangleright <i>number of phases limited due to Lemma 11</i>
for $\Theta(Q \log n)$ rounds do	\triangleright <i>Q determined in Lemmas 13,14</i>
Nodes in S_ϕ send \mathcal{M} with probability p/Q	\triangleright <i>disseminate \mathcal{M} to $N(S)$</i>
Compute DIS $S_{\phi+1}$ of $\tilde{H}^\mu[S_\phi]$ by executing subroutine COMPUTEDIS(μ, p) on all $v \in S$	

4

 Robust Broadcast Algorithm - Overview

The solution of the *neighborhood dissemination problem* lies at the core of the broadcast algorithm by [4]. An algorithm solves neighborhood dissemination for a set of active nodes S that already know a message \mathcal{M} , if after its execution, the neighborhood $N(S)$ of S in the communication graph G_C also knows \mathcal{M} . Starting with a source node v as single active node $S = \{v\}$, global broadcast can be solved by iteratively calling the neighborhood dissemination routine, setting $S = N(S)$ after each iteration. This way \mathcal{M} travels one hop along all shortest paths in G_C emanating from the source, thus solving global broadcast after at most D calls of the neighborhood dissemination routine (D is the diameter of G_C).

Algorithm ROBUSTDISSEMINATION (introduced above) solves the neighborhood dissemination problem in a robust manner. This means that the synchronous execution of ROBUSTDISSEMINATION by all active nodes $v \in S$ disseminates message \mathcal{M} to $N(S)$ w.h.p. and for any strategy of the adversary. It differs from [4] only in its subroutines and its analysis.⁷ The dissemination of \mathcal{M} from S to $N(S)$, relies on a graph structure $H^\mu[S]$ among the set S of active nodes, named *SINR-induced graph* by [4]. The graph $H^\mu[S]$ contains all edges among nodes in S with a probability of transmission success of at least μ .

In [4], active nodes approximate $H^\mu[S]$ by exchanging messages and computing the ratio of successfully received messages in order to determine reliable links. In the unreliable SINR model, the probability of transmission success is subject to adversary influence. In our analysis we account for that by modifying the definition of $H^\mu[S]$ and classify edges into safe and unsafe edges (implying that $H^\mu[S]$ is not unique, since the adversary may 'choose' unsafe edges). The structure $H^\mu[S]$ itself is implicit. We introduce a subroutine TRANSMIT that, when executed simultaneously by all nodes in S , 'probes' connections and passes messages only among those nodes in S which are sufficiently 'reliable', while prohibiting communication among all others, thereby inducing $H^\mu[S]$. The details are covered in Section 5.

Algorithm ROBUSTDISSEMINATION is grouped into phases during which active nodes send. After each phase, active nodes are thinned out to decrease interference and enable dissemination of \mathcal{M} to more distant neighbors. In [4], this is achieved by calculating a maximal independent set (MIS) on $H^\mu[S]$ and deactivating the nodes not in it. Due to unreliable edges we are not able to compute a MIS of $H^\mu[S]$ in our scenario. Instead, we use the well-known coloring algorithm of [18] to construct a structure we call *Dominating Independent Set* (DIS) in $\mathcal{O}(\log n \log^* n)$ rounds and has properties similar to a MIS (this structure was used before in the context of dual graphs by [3]). The algorithm COMPUTEDIS uses TRANSMIT as a subroutine and is given and analyzed in Section 6.

We show that after calculating a DIS on S and deactivating all nodes not in it, the minimal distance among remaining nodes in S at least doubles. Hence, after at most $\mathcal{O}(\log R)$ phases their distance exceeds r_e (remaining nodes are 'sparse') and by then (at the latest) \mathcal{M}

⁷ We adjust the proofs of [4] to this paper, thus knowledge of [4] is recommended but not required. Additionally we mark all lemmas, which have a direct counterpart in [4] with "(cf. [4])".

is received by all nodes in $N(S)$ w.h.p. The properties required for the analysis of Algorithm 1 are given in Section 7 where we adhere closely to the respective proofs provided in [4].⁸

In the following, we present the details and subroutines of ROBUSTDISSEMINATION. We show that key characteristics of the algorithm of [4] are conserved or transformed into similar notions in the unreliable case. This enables us to prove the following theorem in Section 8.

► **Theorem 2.** *Algorithm ROBUSTDISSEMINATION solves neighborhood dissemination in the unreliable SINR model robustly and in $\mathcal{O}(\log n (\log^* n + (\log R)^{\alpha+1} \frac{\beta_{\max}}{\beta_{\min}}))$ rounds.*

We can solve global broadcast by repeating algorithm ROBUSTDISSEMINATION D times, a fact we note in the following corollary.

► **Corollary 3.** *Algorithm ROBUSTDISSEMINATION can be used to solve global broadcast in the unreliable SINR model robustly and in $\mathcal{O}(D \log n (\log^* n + (\log R)^{\alpha+1} \frac{\beta_{\max}}{\beta_{\min}}))$ rounds.*

5 SINR-Induced Graphs in the Unreliable SINR Model

Assume that during each round the active nodes in S send with the same probability $p \in (0, 1)$, which yields a random set $I \subseteq S$ of transmitting nodes. For every pair of nodes u, v let $\sigma_{u,v}$ be the probability that $\text{SINR}(u, v, I) \geq \beta_{\max}$, i.e., the probability that a safe transmission from u to v takes place. Let $\tau_{u,v}$ be the probability that $\text{SINR}(u, v, I) \geq \beta_{\min}$, i.e., an unsafe transmission *may* take place (the adversary might decide). Obviously $\tau_{u,v} \geq \sigma_{u,v}$.

Let $\mu \in (0, p)$ be a given threshold probability. Then we call v a μ -safe neighbor of u and (u, v) a μ -safe edge if both $\sigma_{u,v} \geq \mu$ and $\sigma_{v,u} \geq \mu$. We highlight the fact that (u, v) is μ -safe iff (v, u) is μ -safe as well, with the notation $\{u, v\}$. We call v a μ -unsafe neighbor of u and (u, v) a μ -unsafe edge if $\tau_{u,v} \geq \mu$ and (u, v) is not μ -safe.

Define the (directed) *SINR-induced graph* $H^\mu[S] := (S, E^\mu[S])$ where $(u, v) \in E^\mu[S]$, iff $\tau_{u,v} \geq \mu$. This means $E^\mu[S]$ contains exactly the μ -safe and μ -unsafe edges. Let $E_{\text{safe}}^\mu[S] \cup E_{\text{unsafe}}^\mu[S] = E^\mu[S]$ with $E_{\text{safe}}^\mu[S] \cap E_{\text{unsafe}}^\mu[S] = \emptyset$ be the partition of $E^\mu[S]$ into μ -safe edges and μ -unsafe edges.

We cannot hope to determine $H^\mu[S]$ efficiently and exactly with a distributed algorithm in the unreliable scenario, since the adversary might mask the existence of μ -unsafe edges. Instead we settle for an ε -close approximation $\tilde{H}^\mu[S]$ (for an $\varepsilon \in (0, \frac{1}{2}]$).⁹ A graph $\tilde{H}^\mu[S] := (S, \tilde{E}^\mu[S])$ is an ε -close approximation of $H^\mu[S]$ if and only if

$$E_{\text{safe}}^\mu[S] \subseteq \tilde{E}^\mu[S] \subseteq E^{(1-\varepsilon)\mu}[S].$$

This entails that $\tilde{H}^\mu[S]$ guarantees to include μ -safe edges with $\sigma_{u,v}, \sigma_{v,u} \geq \mu$ and to exclude edges with $\tau_{u,v} < (1-\varepsilon)\mu$. A graph $\tilde{H}^\mu[S]$ may or may not include edges (u, v) with $\tau_{u,v} \geq (1-\varepsilon)\mu$ and $\sigma_{u,v} < \mu$ or $\sigma_{v,u} < \mu$, though. Due to these edges, an approximation $\tilde{H}^\mu[S]$ is not unique and communication in $\tilde{H}^\mu[S]$ along these edges is inherently volatile.

Note that instead of computing $\tilde{H}^\mu[S]$ (as in [4]), we utilize that communication via TRANSMIT(μ, p, M) (Algorithm 2) induces a graph $\tilde{H}^\mu[S]$ by guaranteeing (w.h.p.) communication along edges in $E_{\text{safe}}^\mu[S]$ and inhibiting communication among edges (u, v) with $\tau_{u,v} < (1-\varepsilon)\mu$ w.h.p. (the adversary decides about the remaining edges in $E^{(1-\varepsilon)\mu}[S]$). We

⁸ Algorithm ROBUSTDISSEMINATION uses the following constants: Network param. n, R , SINR param. $P, N, \alpha, \beta_{\min}, \beta_{\max}$, and tuning param. p, ε . Parameter μ is derived from these constants (cf. Lemma 1).

⁹ $\varepsilon = \frac{1}{2}$ optimizes the run-time of TRANSMIT(μ, p, \mathcal{M}_v). The increased in-degree Δ of $\tilde{H}^\mu[S]$ associated with this choice of ε is of comparatively little consequence regarding run-time.

prove this in Lemma 5. The properties of $\tilde{H}^\mu[S]$ are leveraged in the analysis of the DIS (in whose computation we actually employ $\text{TRANSMIT}(\mu, p, M)$, cf. Section 6) and in the analysis of the neighborhood dissemination protocol (Section 7).

A convenient property of $\tilde{H}^\mu[S]$ we immediately observe, is that its maximum in-degree is bounded by $\Delta := \frac{1}{(1-\varepsilon)\mu}$, since a higher in-degree implies a node v that might receive $\sum_{u \in N(v)} \tau_{u,v} \geq \sum_{u \in N(v)} (1-\varepsilon)\mu > 1$ messages in expectation, contradicting the fact that nodes receive at most one message per round. Another property is shown by the subsequent lemma, namely that there is a small enough constant μ , such that $H^\mu[S]$ contains all relatively short edges as μ -safe edges. Since $E_{\text{safe}}^\mu[S] \subseteq \tilde{E}^\mu[S]$, this lemma extends to $\tilde{H}^\mu[S]$.

► **Lemma 4** (cf. [4]). *There exists $\mu \in (0, p)$, s.t. for any $S \subseteq V$ and all $u, v \in S, u \neq v$ with $d(u, v) \leq \min(2d_{\min}^S, r_\varepsilon)$, $d_{\min}^S = \min_{u \neq v \in S} d(u, v)$, it holds that $\{u, v\} \in E_{\text{safe}}^\mu[S]$.*

Proof. Since $d(u, v) \leq 2 \cdot d_{\min}^S$ and $d(u, v) \leq r_\varepsilon$, this lemma is a corollary of Lemma 1. ◀

The following routine guarantees safe and fast message passing among μ -safe neighbors in S w.h.p. Furthermore, it inhibits communication among nodes in S with low transmission probability w.h.p. Edges along which communication takes place form a graph $\tilde{H}^\mu[S]$.

Algorithm 2 $\text{TRANSMIT}(\mu, p, \mathcal{M}_v)$. ▷ to be initiated simultaneously by each $v \in S$

for $T \leftarrow \frac{c \log n}{\varepsilon^2 \mu}$ rounds **do** ▷ c determined in Lemma 5 and $\varepsilon \in (0, \frac{1}{2})$ fixed
 Send pair $(\text{ID}(v), \mathcal{M}_v)$ of message \mathcal{M}_v and own $\text{ID}(v)$ with probability p
for T rounds **do** ▷ list length constant, since nodes have at most $\frac{1}{(1-\varepsilon)\mu}$ neighbors
 With probability p , send list of IDs of which at least $(1 - \frac{\varepsilon}{2})\mu T$ messages were received
for all $\text{ID}(u)$ in own ID-list **do**
 if $\text{ID}(u)$ is in own list **and** $\text{ID}(v)$ is in the ID-list received from u **then**
 Consider message \mathcal{M}_u of u as received ▷ receive messages where $\sigma_{u,v}, \sigma_{v,u} \geq \mu$
 else discard \mathcal{M}_u ▷ neglect messages where $\tau_{u,v} < (1-\varepsilon)\mu$

► **Lemma 5.** *Algorithm $\text{TRANSMIT}(\mu, p, \mathcal{M}_v)$ takes $\mathcal{O}(\frac{\log n}{\varepsilon^2 \mu})$ rounds. If initiated simultaneously by each $v \in S$, then messages are received among μ -safe neighbors u, v ($\sigma_{u,v} \geq \mu$) w.h.p. If $\tau_{u,v} < (1-\varepsilon)\mu$ then w.h.p. no message of u is received by v . Furthermore, edges along which messages are successfully received form an ε -close approximation $\tilde{H}^\mu[S]$ of $H^\mu[S]$ w.h.p.*

Proof. Let $u, v \in S$ and let $X_i := 1$, if $\text{SINR}(u, v, I) \geq \beta_{\max}$ in round i of the first loop, i.e., a safe transmission from u to v takes place, and $X_i := 0$, else. Furthermore, let $X := \sum_{i=1}^T X_i$. Assume (u, v) is μ -safe, i.e., $\sigma_{u,v}, \sigma_{v,u} \geq \mu$. Using a multiplicative Chernoff bound we show that the event $X \leq (1 - \frac{\varepsilon}{2})\mu T$ is very improbable.

$$\mathbb{P}(X \leq (1 - \frac{\varepsilon}{2})\mu T) \leq \exp(-\frac{\varepsilon^2 \sigma_{u,v} T}{8}) = \exp(-\log n \cdot \frac{c \sigma_{u,v}}{8\mu}) \leq n^{-\frac{c}{8}}.$$

Analogously, the same result holds for messages sent from v to u . Let $A_{u,v}$ be the event that v does not receive message \mathcal{M}_u from u via algorithm $\text{TRANSMIT}(\mu, p, \mathcal{M}_v)$ that is $X \leq (1 - \frac{\varepsilon}{2})\mu T$ during the first loop or u does not receive v 's list during the second loop. The second condition is even less likely than the first, hence $\mathbb{P}(A_{u,v}) \leq 2n^{-\frac{c}{8}} \leq n^{1-\frac{c}{8}}$.

This means that (for $c > 8$) event $\bar{A}_{u,v}$ takes place w.h.p., i.e., v receives message \mathcal{M}_u from u . The same is true for event $\bar{A}_{v,u}$. If both $\bar{A}_{u,v}$ and $\bar{A}_{v,u}$ occur, we have $\{u, v\} \in \tilde{E}_{\text{safe}}^\mu[S]$ in the graph $\tilde{H}^\mu[S]$ induced by those edges along which messages are successfully received.

3:10 Broadcasting in an Unreliable SINR Model

Now let $Y_i := 1$, if $\text{SINR}(u, v, I) \geq \beta_{\min}$ in round i of the first loop, i.e., an unsafe transmission from u to v might take place, and $Y_i := 0$, else. Let $Y := \sum_{i=1}^T Y_i$. Assume $\tau_{u,v} < \mu(1 - \varepsilon) =: \mu'$ and let $B_{u,v}$ be the event $Y \geq (1 - \frac{\varepsilon}{2})\mu T$. We obtain

$$\begin{aligned} \mathbb{P}(B_{u,v}) &= \mathbb{P}(Y \geq (1 - \frac{\varepsilon}{2})\mu T) = \mathbb{P}(Y \geq \frac{1 - \frac{\varepsilon}{2}}{1 - \varepsilon} \mu' T) \leq \mathbb{P}(Y \geq (1 + \frac{\varepsilon}{2}) \mu' T) \\ &\leq \exp(-\frac{\varepsilon^2 \mu' T}{12}) = \exp(-\log n \cdot \frac{c \mu'}{12\mu}) = n^{-\frac{c}{12}(1-\varepsilon)} \stackrel{(\varepsilon \leq \frac{1}{2})}{\leq} n^{-\frac{c}{24}}. \end{aligned}$$

Therefore, even if the adversary permits all transmissions with $\text{SINR}(u, v, I) \geq \beta_{\min}$ during the first loop, the total number will be lower than $(1 - \frac{\varepsilon}{2})\mu T$ w.h.p. Thus u is not in v 's list, hence neither receives the others message thus $(u, v), (v, u) \notin \tilde{E}^\mu[S]$ w.h.p. The probability that at least one edge of the clique among active nodes $C(S) := \{(u, v) \mid u, v \in S, u \neq v\}$ is falsely in- or excluded in $\tilde{E}^\mu[S]$ is bounded by

$$\mathbb{P}(\bigcup_{(u,v) \in C(S)} (A_{u,v} \cup B_{u,v})) \leq \sum_{(u,v) \in C(S)} (\mathbb{P}(A_{u,v}) + \mathbb{P}(B_{u,v})) \leq \sum_{(u,v) \in C(S)} (n^{1-c/8} + n^{-c/24}) < n^{4-c/24}.$$

This shows that for $c > 96$ w.h.p. all μ -safe edges are included in $\tilde{E}^\mu[S]$ and all edges which are not even $\mu(1 - \varepsilon)$ -unsafe are excluded w.h.p. \blacktriangleleft

Note that the definition of $\tilde{H}^\mu[S]$ permits directed edges $(u, v) \in \tilde{E}^\mu[S]$ with $(v, u) \notin \tilde{E}^\mu[S]$. In practice, this happens if $\tau_{u,v}, \tau_{v,u} \geq \mu(1 - \varepsilon)$ and the adversary permits all transmissions among u and v in the first loop of Algorithm 2 such that u and v are in each others lists. Subsequently, the adversary allows only the transfer of v 's list to u but blocks all attempts of u sending its list to v . Nevertheless, μ -safe edges are undirected w.h.p., since there will occur sufficiently many safe transmissions between μ -safe neighbors, which the adversary cannot block. In the following we use the notation $\{u, v\}$ to highlight μ -safe edges.

6 Calculating Dominating Independent Sets

After each phase of message dissemination of nodes in S , we rely on the properties of $\tilde{H}^\mu[S]$, to thin out the set of nodes S so that message dissemination becomes feasible in areas of high interference. In the basic algorithm [4] this is accomplished by determining a MIS of $\tilde{H}^\mu[S]$ and deactivating all nodes not in it. In unreliable scenarios we cannot guarantee independence with respect to μ -unsafe edges since the adversary might mask their existence during MIS-calculation. To loosen the requirements, we introduce a more general notion.

► **Definition 6.** Let $G = (V, E, E')$ be a graph with disjoint sets of undirected edges E and directed edges E' . Subset $V' \subseteq V$ is independent w.r.t. E if no two nodes in V' are connected by an edge $\{u, v\} \in E$. Subset $V' \subseteq V$ is dominating w.r.t. $E \cup E'$ if for every node $v \in V$ either $v \in V'$ or there exists a $u \in V'$ such that $(u, v) \in E \cup E'$. A *Dominating Independent Set* (DIS) $D \subseteq V$ of G is independent w.r.t. E and dominating w.r.t. $E \cup E'$.

We exploit the fact that the in-degree of any valid graph $\tilde{H}^\mu[S]$ is bounded by $\Delta \in \mathcal{O}(1)$ to calculate a DIS of $(S, E_{\text{safe}}^\mu[S], E_{\text{unsafe}}^{(1-\varepsilon)\mu}[S])$. First, we determine a $(3\Delta \log \Delta)^2$ -coloring of S with respect to the μ -safe edges $E_{\text{safe}}^\mu[S]$ in $\mathcal{O}(\log n \log^* n)$ rounds by adapting the method of Linial [18], which uses Δ -cover-free families. Second, we calculate a DIS based on this coloring by successively adding nodes of one color while sustaining the DIS condition.

► **Definition 7.** A Δ -cover-free family \mathcal{F} is a family of subsets of $[m]$,¹⁰ such that for any selection of distinct sets $F_0, \dots, F_\Delta \in \mathcal{F}$ it holds that $F_0 \not\subseteq \bigcup_{j=1}^\Delta F_j$.

Erdős et al. show in [6] how such families can be constructed using polynomials over finite fields. The proof of the following lemma provides the according construction procedure and is given in Appendix A.

► **Lemma 8.** For prime power q , and integer $d \in [q-1]$ there is a $\lfloor (q-1)/d \rfloor$ -cover-free family \mathcal{F} of size $|\mathcal{F}| = q^{d+1}$ of subsets of $[q^2]$.

From Lemma 8 we immediately obtain the following Lemma 9, which was suggested by Linial in [18] as constructive alternative to his non-constructive proof of an even smaller Δ -cover-free family. For completeness, the lemma is derived in Appendix A.

► **Lemma 9.** For $k \in \mathbb{N}$, there is a Δ -cover-free family \mathcal{F} of subsets of $[m]$ with $|\mathcal{F}| \geq k$ and $m \leq (3\Delta \log k)^2$.

Algorithm 3 COMPUTEDIS(μ, p). ▷ to be initiated simultaneously by each $v \in S$

$k_0 \leftarrow \text{ID}(v)$ ▷ assign colors, initially unique IDs
 $k \leftarrow n$ ▷ k tracks upper bound of colors currently used
for $\mathcal{O}(\log^* n)$ times **do** ▷ loop computes coloring w.r.t. μ -safe edges
 TRANSMIT(μ, p, k_0) ▷ send own color to neighbors in $\tilde{H}^\mu[S]$
 Let k_1, \dots, k_t be the colors received from v 's neighbors in $\tilde{H}^\mu[S]$. ▷ $t \leq \Delta$, $k_i \leq k$
 Let $\mathcal{F} = \{F_1, \dots, F_k\}$ as in Lemma 9 ▷ construction procedure in proof of Lemma 8
 Choose new color $k_0 \in F_{k_0} \setminus \bigcup_{k_i \neq k_0} F_{k_i}$ ▷ new coloring is valid w.r.t. μ -safe edges
 $k \leftarrow (3\Delta \log k)^2$ ▷ at most $(3\Delta \log k)^2$ colors remaining (Lemma 9)
for $l \leftarrow 1$ **to** k **do** ▷ loop uses $\mathcal{O}(1)$ -coloring to compute DIS
 if $l = k_0$ **then**
 TRANSMIT($\mu, p, \text{'Do not join DIS and terminate!'}$) ▷ inform neighbors
 Join DIS and terminate
 else TRANSMIT(μ, p, \emptyset) ▷ All active nodes must initiate routine (Lemma 5)
 if received message $\text{'Do not join DIS and terminate!'}$ **then** terminate

► **Lemma 10.** Algorithm COMPUTEDIS(μ, p) computes a DIS of $(S, E_{\text{safe}}^\mu[S], E_{\text{unsafe}}^{(1-\epsilon)\mu}[S])$ in $\mathcal{O}(\frac{\log n}{\epsilon^2 \mu} \log^* n)$ rounds in a robust manner, when initiated simultaneously by all nodes in S .

Proof. We prove that during the **first loop** of COMPUTEDIS(μ, p) a valid coloring w.r.t. μ -safe edges is maintained. Initially we have the trivial coloring by IDs. The rest can be shown inductively. Presume that, at the beginning of a given loop cycle, we have a valid coloring with respect to $E_{\text{safe}}^\mu[S]$ of size at most k .

Due to the synchronous execution of TRANSMIT(μ, p, k_v) by each node $v \in S$ (where k_v is the current color of v) every node receives at most $t \leq \Delta$ colors $k_1, \dots, k_t \leq k$ from its neighbors w.h.p., since messages are only transmitted along edges of $\tilde{H}^\mu[S]$ (Lemma 5) whose in-degree is bounded by Δ . In particular, TRANSMIT(μ, p, k_v) guarantees that every node receives the colors from all of its μ -safe neighbors w.h.p. (Lemma 5).

¹⁰ For natural numbers $n \geq 1$ we define $[n] := \{1, \dots, n\}$.

Note that v may also receive colors that conflict with its own, which stem from $\mu(1-\varepsilon)$ -unsafe neighbors which are not μ -safe and whose messages were blocked by the adversary during earlier loop cycles. However, since we only need to ensure a valid coloring with respect to μ -safe edges, this does not concern us.

Now all nodes in S construct the same Δ -cover-free family $\mathcal{F} = \{F_1, \dots, F_k\}$ of subsets of $[m]$ with $m \leq (3\Delta \log k)^2$, in accordance with Lemmas 8, 9.¹¹ Each node $v \in S$ picks a new color $k_0 \in F_{k_0} \setminus \bigcup_{k_i \neq k_0} F_{k_i}$, which exists due to the Δ -cover-free property. Therefore, neighbors that had different colors during previous rounds are (w.h.p.) still differently colored. In particular, this implies that w.h.p. μ -safe neighbors remain differently colored.

In the **second loop** we use the coloring to determine a DIS of S . Let D_l be the subset of nodes in S that have joined the DIS and let T_l be those that have terminated until (and including) loop cycle l . W.h.p., COMPUTEDIS maintains the invariant that D_l is a DIS of T_l , i.e., $D_l \subseteq T_l$ is independent w.r.t. μ -safe edges and D_l dominates T_l w.r.t. μ -safe edges and $\mu(1-\varepsilon)$ -unsafe edges.

For $l = 0$, this is obvious since $T_0 = D_0 = \emptyset$. Presume that D_{l-1} is a DIS w.r.t. T_{l-1} and let $v \in T_l \setminus T_{l-1}$. Since μ -safe neighbors are w.h.p. differently colored, $v \in D_l \setminus D_{l-1}$ cannot have a μ -safe neighbor in $D_l \setminus D_{l-1}$ (w.h.p.). Moreover, since D_{l-1} informed its μ -safe neighbors w.h.p. (Lemma 5), they already terminated during previous cycles. Therefore a node $v \in D_l \setminus D_{l-1}$ that joined the DIS during cycle l and terminated only then, cannot have any μ -safe neighbors in D_{l-1} (w.h.p.).

We already know that D_{l-1} dominates T_{l-1} w.r.t. μ -safe edges and $\mu(1-\varepsilon)$ -unsafe edges (w.h.p.). Every node $v \in T_l \setminus T_{l-1}$ is either in $D_l \setminus D_{l-1}$ or was informed by a neighbor $u \in D_l \setminus D_{l-1}$ (possibly via a $\mu(1-\varepsilon)$ -unsafe edge), which means v is dominated by $u \in D_l \setminus D_{l-1}$. Therefore $D_l \setminus D_{l-1}$ dominates $T_l \setminus T_{l-1}$, thus D_l dominates T_l (w.h.p.).

The first loop has the claimed runtime (cf. Lemma 5). The run-time of the second loop depends on the size of the coloring. After at most $\mathcal{O}(\log^* n)$ cycles of the first loop there are $\mathcal{O}((\Delta \log \Delta)^2) = \mathcal{O}(1)$ colors left. Therefore the run-time of the second loop is dominated by the run-time of the first, proving the claimed overall run-time. ◀

7 Neighborhood Dissemination Analysis

In this section we apply moderate changes to the technical parts given in [4] in order to reuse them in our unreliable SINR model. Let ϕ be a phase of Algorithm 1: ROBUSTDISSEMINATION, let S_ϕ be the set of active nodes during phase ϕ and let $d_\phi := \min_{u \neq v \in S_\phi} d(u, v)$, if $|S_\phi| \geq 2$ and $d_\phi := \infty$, else. The following lemma proves that in a DIS $S_{\phi+1}$ of S_ϕ ¹² the minimum distance $d_{\phi+1}$ among nodes doubles, or is already larger than the effective communication range r_e . The proof is provided in Appendix B.

► **Lemma 11** (cf. [4]). *Let ϕ be a phase of Algorithm 1. There exists a constant $\mu \in (0, p)$, such that $d_\phi \geq 2^{\phi-1} d_{\min}$ or $d_\phi > r_e$, where $d_{\min} = \min_{u \neq v \in V} d(u, v)$.*

The next lemmas show that node $v \in N(S)$ receives the broadcast message \mathcal{M} with a guaranteed probability from its nearest active neighbor $u \in S_\phi$, in case certain conditions are met. These conditions include that u is in safe communication range and that there exists a neighboring node $w \in S_\phi$ of u in $H^\mu[S_\phi]$, which is not too close. With the existence of a

¹¹ According to our model, nodes have unlimited computing power during each round. Nevertheless Δ -cover-free families can be computed within polynomial time using the constructive proof of Lemma 8.

¹² We presume that μ, ε are fixed and abbreviate 'DIS of $(S, E_{\text{safe}}^\mu[S], E_{\text{unsafe}}^{(1-\varepsilon)\mu}[S])$ ' with 'DIS of S '.

relatively long edge $(w, u) \in E^\mu[S_\phi]$ we can bound the interference at u using the knowledge that w was able to transfer messages to u with a certain probability ($\tau_{w,v} \geq (1-\varepsilon)\mu$). Consequently, we are able to limit the interference at u 's (close) neighbor v . The proofs of the following lemmas are given in Appendix B. We adhere closely to the proofs provided in [4], although some adaptations and clarifications are required.

First, we limit the interference stemming from the set of nodes close to v and then from the set of distant nodes. For this purpose, let $V' \subseteq V$ be a subset of nodes and let $X_{V'}^q$ be the random number of nodes in V' that send, when each has individual sending probability q . Moreover, let $I_{V'}^q(y) := \sum_{x \in V', \text{ sends}} P/d(x, y)^\alpha$ be the random interference at $y \in V$ stemming from sending nodes in V' , given that nodes send with probability q . Then $I_{V'}^{\max}(y) := \sum_{x \in V'} P/d(x, y)^\alpha$ defines the maximum interference at $y \in V$ from nodes in V' .

► **Lemma 12** (cf. [4]). *Let ϕ be a phase of Algorithm 1. Let $v \in N(S)$ and let $u \in S_\phi$ be the node closest to v . Presume there exists $w \in S_\phi$ with an edge $(w, u) \in E^\mu[S_\phi]$ and let w be the farthest such neighbor of u . Let $A = \{x \in S_\phi \mid d(u, x) \leq 2d(u, w)\} \setminus \{u, v\}$ and $\bar{A} = S_\phi \setminus (A \cup \{u, v\})$. If active nodes send with probability p/Q , then there exists $\eta \in \Theta(\beta_{\min})$, such that the following events occur simultaneously with probability at least $\frac{(1-\varepsilon)\mu p}{8Q} \in \Theta(1/Q)$:*

$$(i) I_A^{p/Q}(v) \leq \frac{2^{\alpha+1}P}{Q\beta_{\min}d(u,v)^\alpha}, (ii) I_{\bar{A}}^{p/Q}(v) \leq \frac{2^{\alpha+1}\eta P}{Q\beta_{\min}d(u,w)^\alpha}, (iii) u \text{ sends and } v \text{ listens.}$$

Lemma 13 proves that there is a constant probability that $v \in N(S)$ receives \mathcal{M} from u , in case v is in communication range of u and the ratio $d(u, w)/d(u, v)$ is above a threshold.

► **Lemma 13** (cf. [4]). *Let ϕ be a phase of Algorithm 1. Let $v \in N(S)$ and $u, w \in S_\phi$ as in Lemma 12. Then there exists $\hat{Q} \in \mathcal{O}(\frac{\beta_{\max}}{\beta_{\min}} 2^\alpha)$, $\gamma \in \Theta((\frac{\beta_{\max}}{\beta_{\min}})^{1/\alpha})$ such that for all $Q \geq \hat{Q}$, node v receives \mathcal{M} safely with probability $\Theta(1/Q)$, if active nodes send with probability p/Q and $d(u, v) \leq (1 + \frac{\rho}{2})r_e$ and $d(u, w) \geq \gamma Q^{-\alpha}d(u, v)$.*

Lemma 14 guarantees for each phase ϕ except the last that $v \in N(S)$ receives \mathcal{M} w.h.p. in this or previous phases, or v is still in safe communication range of an active neighbor.

► **Lemma 14** (cf. [4]). *Let $\phi \leq \log R + 1$ be a phase of Algorithm 1. Let $v \in N(S)$ and let $u_\phi \in S_\phi$ be the active node closest to v . Then there exists a $Q \in \Theta((\log R)^\alpha)$, $Q \geq \hat{Q}$ such that either v receives \mathcal{M} w.h.p. during phase ϕ or earlier, or $d(u_{\phi+1}, v) \leq r_e(1 + \frac{\rho\phi}{2\log R})$.*

Finally, we show that if S_ϕ with $|S_\phi| \geq 2$ is 'sparse' in the sense that $d_\phi > r_e$ (which is a typical case in at least one phase of Algorithm 1), then $v \in N(S)$ receives \mathcal{M} from $u_\phi \in S_\phi$.

► **Lemma 15.** *Let ϕ be a phase of Algorithm 1. Let $|S_\phi| \geq 2$ and $d_\phi > r_e$. If for $v \in V$ there is a node $u \in S_\phi$ with $d(u, v) < c_1 \cdot r_e$ with $0 < c_1 < 1 + \rho$, then v receives \mathcal{M} in phase ϕ w.h.p.*

8 Proof of Theorem 2

Proof. Using the previous lemmas, we show that algorithm ROBUSTDISSEMINATION is correct and has the claimed run-time. For the correctness, we prove that any node $v \in N(S)$ receives \mathcal{M} w.h.p. during the last phase $\psi := \lfloor \log R \rfloor + 2$, at the latest.

Lemma 11 shows that for the minimum distance $d_\phi = \min_{u \neq v \in S_\phi} d(u, v)$ among nodes in S_ϕ , it holds that $d_\phi \geq 2^{\phi-1}d_{\min}$ or $d_\phi > r_e$. Hence, if $\phi > \log R + 1$, then $2^{\phi-1}d_{\min} > R \cdot d_{\min} = d_{\max}$. Therefore, in the last, ψ -th phase, we have $d_\psi \geq 2^{\psi-1}d_{\min} > d_{\max}$ and consequently $d_\psi > r_e$ (recall that d_{\max} is the maximum distance among any two nodes within range r_e of each other). Thus, during the last phase ψ , either $d_\psi = \infty$ in case $|S_\psi| = 1$, or the remaining nodes in S_ψ are *sparse*, in the sense that their minimum distance is greater than r_e .

First we approach the case $|S_\psi|=1$. For this purpose, let ϕ be the first phase for which only one node remains in S_ϕ , i.e., either $|S_1|=1$ or $|S_{\phi-1}|\geq 2$. In case $S_1=\{u\}$, node $v\in N(S_1)$ receives \mathcal{M} safely and w.h.p., during the inner loop where we disseminate \mathcal{M} for $\Theta(Q\log n)$ rounds with probability p/Q . This is due to the fact that $d(v,u)\leq r_e$ ($v\in N(S)$) and there is no interference from any other node. Additionally, we choose the hidden constant in $\Theta(Q\log n)$ sufficiently large, such that u sends at least once during that loop w.h.p.

Otherwise ($|S_{\phi-1}|\geq 2$), we apply Lemma 14 on phase $(\phi-1)$, and see that either v has received \mathcal{M} already (in that case we are done) or $d(u_\phi, v)\leq r_e(1+\frac{\rho(\phi-1)}{2\log R})<r_e(1+\rho)=r_s$. The latter condition implies that v is in safe communication range of the only active node u_ϕ in phase ϕ , hence v receives \mathcal{M} w.h.p. during round ϕ using the same argument as before.

Now consider $|S_\psi|\geq 2$. Then Lemma 14 applies for phase $(\psi-1)\leq\log R+1$, therefore either v already received \mathcal{M} w.h.p. or $d(u_\psi, v)\leq r_e(1+\frac{\rho(\psi-1)}{2\log R})\leq c_1\cdot r_e$ with $c_1\leq 1+\frac{3}{4}\rho$.¹³ Since in the final phase ψ , S_ψ is also sparse in the sense that $d_\psi>r_e$, the premise of Lemma 15 is fulfilled, thus v receives \mathcal{M} from its closest neighbor u_ψ in S_ψ w.h.p.

During each phase we execute the sub-procedure COMPUTEDIS, which takes $\mathcal{O}(\log n\log^*n)$ rounds (Lemma 10). Sending \mathcal{M} takes $\mathcal{O}(Q\log n)\subseteq\mathcal{O}(\frac{\beta_{\max}}{\beta_{\min}}(\log R)^\alpha\log n)$ rounds each phase (we determined Q in Lemma 14). We have $\lfloor\log R\rfloor+2$ phases, therefore algorithm ROBUST-DISSEMINATION takes at most $\mathcal{O}(\log n(\log^*n+(\log R)^{\alpha+1}\frac{\beta_{\max}}{\beta_{\min}}))$ rounds. \blacktriangleleft

9 Adversary Reduction

Finally, we show that a seemingly stronger adversary, which controls *all* SINR-parameters can be reduced to our abstract model, where it controls only β_v . Assume that nodes know only upper and lower bounds $P^\uparrow>P^\downarrow>0$, $N^\uparrow>N^\downarrow>0$, $\alpha^\uparrow>\alpha^\downarrow\geq 1$, and $\beta^\uparrow>\beta^\downarrow\geq 1$ on the true SINR parameters $P_v, N_v, \alpha(u, v), \beta_v$. In each round and for all pairs of nodes $u, v\in V$, the adversary determines the actual values $P_v, N_v, \alpha(u, v)$, and β_v arbitrarily within the given upper and lower bounds, thereby influencing the outcome of

$$\text{SINR}(u, v, I) := \frac{P_u/d(u, v)^{\alpha(u, v)}}{N_v + \sum_{w\in I} P_w/d(w, v)^{\alpha(w, v)}} \geq \beta_v.$$

The following theorem formalizes the fact that instead of picking values for all SINR parameters $P_v, N_v, \alpha(u, v), \beta_v$ within the given upper and lower bounds, the adversary can equivalently modify only β'_v within some enlarged interval $[\beta_{\min}, \beta_{\max}]$ with $\beta_{\min}\leq\beta^\downarrow<\beta^\uparrow\leq\beta_{\max}$, while the other SINR parameters remain static (and globally known).

► **Theorem 16.** *For sufficiently large β^\downarrow and a fixed set of uniform SINR parameters $P\in[P^\downarrow, P^\uparrow]$, $N\in[N^\downarrow, N^\uparrow]$, $\alpha\in[\alpha^\downarrow, \alpha^\uparrow]$, and $\beta\in[\beta^\downarrow, \beta^\uparrow]$, there are values $\beta_{\max}\geq\beta_{\min}\geq 1$ such that for any choice of $P_v, N_v, \alpha(u, v), \beta_v$ within the bounds, there is a $\beta'_v\in[\beta_{\min}, \beta_{\max}]$ s.t.*

$$\frac{P_u/d(u, v)^{\alpha(u, v)}}{N_v + \sum_{w\in I} P_w/d(w, v)^{\alpha(w, v)}} \geq \beta_v \iff \frac{P/d(u, v)^\alpha}{N + \sum_{w\in I} P/d(w, v)^\alpha} \geq \beta'_v. \quad (2)$$

Proof. For brevity, let $\text{SINR}_{\text{adv}}(u, v, I)$ be the formula with the *adversary* controlled parameters $P_v, N_v, \alpha(u, v)$ (left-hand side of (2)) and analogously let $\text{SINR}_{\text{uni}}(u, v, I)$ be defined with the *uniform* parameters P, N, α (right-hand side of (2)). For a given $I\subseteq V\setminus\{u, v\}$, we interpret $C_{u, v, I} := \text{SINR}_{\text{adv}}(u, v, I) - \beta_v$ as *connection strength* from u to v . Recall that v receives a message from u if and only if $C_{u, v, I}\geq 0$. Similarly we define $C_{u, v, I}^{\text{uni}} := \text{SINR}_{\text{uni}}(u, v, I) - \beta$.

¹³We assume $\log R\geq 2$. Otherwise the neighborhood dissemination can easily be solved with Lemma 1.

Instead of going through the ramifications of the functional dependencies among the SINR parameters occurring in $C_{u,v,I}$, we acknowledge that for the given bounds on the SINR parameters and fixed I , the connection strength $C_{u,v,I}$, treated as function of the SINR parameters, does not diverge. Hence there is a fixed range $[C_{u,v,I}^\downarrow, C_{u,v,I}^\uparrow]$, from which the adversary chooses $C_{u,v,I}$. We define¹⁴ the *largest negative deviation* Δ^- and the *largest positive deviation* Δ^+ of $C_{u,v,I}$ from $C_{u,v,I}^{\text{uni}}$

$$\Delta^- := \min_{u \neq v \in V, I \subseteq V} (C_{u,v,I}^\downarrow - C_{u,v,I}^{\text{uni}}) \leq 0, \quad \Delta^+ := \max_{u \neq v \in V, I \subseteq V} (C_{u,v,I}^\uparrow - C_{u,v,I}^{\text{uni}}) \geq 0.$$

Suppose the adversary influences the deviation of $C_{u,v,I}$ from $C_{u,v,I}^{\text{uni}}$ via a *deviation variable* $\Delta \in [\Delta^-, \Delta^+]$, then

$$\begin{aligned} \text{SINR}_{\text{adv}}(u, v, I) \geq \beta_v &\iff C_{u,v,I} \geq 0 \iff C_{u,v,I}^{\text{uni}} + \Delta \geq 0 \\ &\iff \text{SINR}_{\text{uni}}(u, v, I) \geq \beta - \Delta \iff \text{SINR}_{\text{uni}}(u, v, I) \geq \beta'_v, \end{aligned}$$

for $\beta'_v \in [\beta_{\min}, \beta_{\max}]$ with $\beta_{\min} := \beta - \Delta^+ \leq \beta - \Delta^- =: \beta_{\max}$. Retracing our definitions, we see that $\beta_{\min} = \beta - \Delta^+ \geq 1$ holds if and only if for all $u, v \in V$ and all $I \subseteq V \setminus \{u, v\}$:

$$\beta - (C_{u,v,I}^\uparrow - C_{u,v,I}^{\text{uni}}) \geq 1 \iff \beta^\downarrow \geq \text{SINR}^\uparrow(u, v, I) - \text{SINR}_{\text{uni}}(u, v, I) + 1$$

where $\text{SINR}^\uparrow(u, v, I) := C_{u,v,I}^\uparrow + \beta^\downarrow$ is the maximum value of $\text{SINR}_{\text{adv}}(u, v, I)$ with u, v, I fixed. Hence the requirement $\beta_{\min} \geq 1$ is met iff

$$\beta^\downarrow \geq \max_{u, v \in V, I \subseteq V} (\text{SINR}^\uparrow(u, v, I) - \text{SINR}_{\text{uni}}(u, v, I) + 1),$$

which completes the proof. ◀

Conveniently, our construction allows us to reduce the case where the default SINR parameters in the network are *non-uniform* (e.g. nodes have different sending Power P or background noise N) to the *uniform* case without changes to the model. This can be done by transferring deviation in the connection strengths due to non-uniformity into the adversaries control. Consequently, the range of deviation $[\Delta^-, \Delta^+]$ from which the adversary may choose the deviation Δ of the true $C_{u,v,I}$ from the presumed uniform connection strength $C_{u,v,I}^{\text{uni}}$ may become large depending on the extent of non-uniformity, thereby increasing the ratio $\beta_{\max}/\beta_{\min} = (\beta - \Delta^-)/(\beta - \Delta^+)$ (intuitively speaking: we 'bought' uniformity with additional adversary influence). We saw in the previous sections that $\beta_{\max}/\beta_{\min}$ affects the run-time at most linearly.¹⁵

Obviously, the construction in the proof of Theorem 16 was simplified by accounting for the *largest possible* deviation of the true $C_{u,v,I}$ from the uniform connection strength $C_{u,v,I}^{\text{uni}}$ for any pair $(u, v) \in V^2, u \neq v$ and any interfering subset $I \subseteq V$. This may result in a much bigger ratio $\beta_{\max}/\beta_{\min}$ compared to the average deviation (again some intuition: we exchanged simplicity for adversary influence).

Moreover, we point out that the qualitative construction in the proof (intentionally) omits the intricate analysis of the dependence of $\beta_{\max}/\beta_{\min}$ on the SINR parameters, the network's size, design and its layout. The ratio $\beta_{\max}/\beta_{\min}$ obtained by the above construction does in fact functionally depend on these parameters and might become large. However, we argue

¹⁴In practice, Δ^+ and Δ^- could be determined via measurements.

¹⁵To minimize run-time, the choice of uniform SINR parameters P, N, α, β should minimize $\frac{\beta_{\max}}{\beta_{\min}}$.

that the size of $\beta_{\max}/\beta_{\min}$ is moderate if the local SINR parameters $P_v, N_v, \alpha(u, v), \beta_v$ are relatively homogeneous and have a narrow range of variation due to the adversary.

In order to guarantee $\beta_{\min} \geq 1$, the default sensitivity parameter β of the network devices needs to be designed large enough, so that β^\downarrow does not decrease below the value given in the proof. This shows that our construction has an inherent trade-off, where increased adversary influence due to the aforementioned effects must be compensated with decreased sensitivity to signal reception (and thus also decreased transmission range).

References

- 1 L. Anantharamu, B. S. Chlebus, D. R. Kowalski, and M. A. Rokicki. Medium access control for adversarial channels with jamming. In *Proceedings of the 18th International Conference on Structural Information and Communication Complexity, SIROCCO '11*, 2011.
- 2 Baruch Awerbuch, Andréa W. Richa, and Christian Scheideler. A jamming-resistant MAC protocol for single-hop wireless networks. In Rida A. Bazzi and Boaz Patt-Shamir, editors, *Proceedings of the Twenty-Seventh Annual ACM Symposium on Principles of Distributed Computing, PODC 2008, Toronto, Canada, August 18-21, 2008*, pages 45–54. ACM, 2008. doi:10.1145/1400751.1400759.
- 3 Keren Censor-Hillel, Seth Gilbert, Fabian Kuhn, Nancy A. Lynch, and Calvin C. Newport. Structuring unreliable radio networks. *Distributed Computing*, 27(1):1–19, 2014. doi:10.1007/s00446-013-0198-8.
- 4 Sebastian Daum, Seth Gilbert, Fabian Kuhn, and Calvin C. Newport. Broadcast in the ad hoc SINR model. In Yehuda Afek, editor, *Distributed Computing - 27th International Symposium, DISC 2013, Jerusalem, Israel, October 14-18, 2013. Proceedings*, volume 8205 of *Lecture Notes in Computer Science*, pages 358–372. Springer, 2013. doi:10.1007/978-3-642-41527-2_25.
- 5 Shlomi Dolev, Seth Gilbert, Rachid Guerraoui, and Calvin C. Newport. Secure communication over radio channels. In Rida A. Bazzi and Boaz Patt-Shamir, editors, *Proceedings of the Twenty-Seventh Annual ACM Symposium on Principles of Distributed Computing, PODC 2008, Toronto, Canada, August 18-21, 2008*, pages 105–114. ACM, 2008. doi:10.1145/1400751.1400767.
- 6 P. Erdős, P. Frankl, and Z. Füredi. Families of finite sets in which no set is covered by the union of r others. *Israel Journal of Mathematics*, 51(1), 1985. doi:10.1007/BF02772959.
- 7 Mohsen Ghaffari, Erez Kantor, Nancy A. Lynch, and Calvin C. Newport. Multi-message broadcast with abstract MAC layers and unreliable links. In Magnús M. Halldórsson and Shlomi Dolev, editors, *ACM Symposium on Principles of Distributed Computing, PODC '14, Paris, France, July 15-18, 2014*, pages 56–65. ACM, 2014. doi:10.1145/2611462.2611492.
- 8 Mohsen Ghaffari, Nancy A. Lynch, and Calvin C. Newport. The cost of radio network broadcast for different models of unreliable links. In Panagiota Fatourou and Gadi Taubenfeld, editors, *ACM Symposium on Principles of Distributed Computing, PODC '13, Montreal, QC, Canada, July 22-24, 2013*, pages 345–354. ACM, 2013. doi:10.1145/2484239.2484259.
- 9 Seth Gilbert, Rachid Guerraoui, and Calvin C. Newport. Of malicious motes and suspicious sensors: On the efficiency of malicious interference in wireless networks. In Alexander A. Shvartsman, editor, *Principles of Distributed Systems, 10th International Conference, OPODIS 2006, Bordeaux, France, December 12-15, 2006, Proceedings*, volume 4305 of *Lecture Notes in Computer Science*, pages 215–229. Springer, 2006. doi:10.1007/11945529_16.
- 10 Olga Goussevskaia, Thomas Moscibroda, and Roger Wattenhofer. Local broadcasting in the physical interference model. In Michael Segal and Alexander Kesselman, editors, *Proceed-*

- ings of the DIALM-POMC Joint Workshop on Foundations of Mobile Computing, Toronto, Canada, August 18-21, 2008, pages 35–44. ACM, 2008. doi:10.1145/1400863.1400873.
- 11 Magnús M. Halldórsson, Stephan Holzer, and Nancy A. Lynch. A local broadcast layer for the SINR network model. In Chryssis Georgiou and Paul G. Spirakis, editors, *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing, PODC 2015, Donostia-San Sebastián, Spain, July 21 - 23, 2015*, pages 129–138. ACM, 2015. doi:10.1145/2767386.2767432.
 - 12 Magnús M. Halldórsson and Pradipta Mitra. Towards tight bounds for local broadcasting. In Fabian Kuhn and Calvin C. Newport, editors, *FOMC'12, The Eighth ACM International Workshop on Foundations of Mobile Computing (part of PODC 2012), Funchal, Portugal, July 19, 2012, Proceedings*, page 2. ACM, 2012. doi:10.1145/2335470.2335472.
 - 13 Tomasz Jurdzinski and Dariusz R. Kowalski. Distributed backbone structure for algorithms in the SINR model of wireless networks. In Marcos K. Aguilera, editor, *Distributed Computing - 26th International Symposium, DISC 2012, Salvador, Brazil, October 16-18, 2012. Proceedings*, volume 7611 of *Lecture Notes in Computer Science*, pages 106–120. Springer, 2012. doi:10.1007/978-3-642-33651-5_8.
 - 14 Tomasz Jurdzinski and Dariusz R. Kowalski. Distributed randomized broadcasting in wireless networks under the SINR model. In *Encyclopedia of Algorithms*, pages 577–580. Springer, 2016. doi:10.1007/978-1-4939-2864-4_604.
 - 15 Tomasz Jurdzinski, Dariusz R. Kowalski, Michal Rozanski, and Grzegorz Stachowiak. On the impact of geometry on ad hoc communication in wireless networks. In Magnús M. Halldórsson and Shlomi Dolev, editors, *ACM Symposium on Principles of Distributed Computing, PODC '14, Paris, France, July 15-18, 2014*, pages 357–366. ACM, 2014. doi:10.1145/2611462.2611487.
 - 16 Fabian Kuhn, Nancy A. Lynch, and Calvin C. Newport. The abstract MAC layer. *Distributed Computing*, 24(3-4):187–206, 2011. doi:10.1007/s00446-010-0118-0.
 - 17 Fabian Kuhn, Nancy A. Lynch, Calvin C. Newport, Rotem Oshman, and Andréa W. Richa. Broadcasting in unreliable radio networks. In Andréa W. Richa and Rachid Guerraoui, editors, *Proceedings of the 29th Annual ACM Symposium on Principles of Distributed Computing, PODC 2010, Zurich, Switzerland, July 25-28, 2010*, pages 336–345. ACM, 2010. doi:10.1145/1835698.1835779.
 - 18 Nathan Linial. Locality in distributed graph algorithms. *SIAM J. Comput.*, 21(1):193–201, 1992. doi:10.1137/0221015.
 - 19 Nancy A. Lynch and Calvin Newport. A (truly) local broadcast layer for unreliable radio networks. In Chryssis Georgiou and Paul G. Spirakis, editors, *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing, PODC 2015, Donostia-San Sebastián, Spain, July 21 - 23, 2015*, pages 109–118. ACM, 2015. doi:10.1145/2767386.2767411.
 - 20 Dominic Meier, Yvonne Anne Pignolet, Stefan Schmid, and Roger Wattenhofer. Speed dating despite jammers. In Bhaskar Krishnamachari, Subhash Suri, Wendi Rabiner Heinzelman, and Urbashi Mitra, editors, *Distributed Computing in Sensor Systems, 5th IEEE International Conference, DCOSS 2009, Marina del Rey, CA, USA, June 8-10, 2009. Proceedings*, volume 5516 of *Lecture Notes in Computer Science*, pages 1–14. Springer, 2009. doi:10.1007/978-3-642-02085-8_1.
 - 21 T. Moscibroda, R. Wattenhofer, and Y. Weber. Protocol design beyond graph-based models. In *Proceedings of the ACM Workshop on Hot Topics in Networks (HotNets-V)*, 2006.
 - 22 Calvin C. Newport, David Kotz, Yougu Yuan, Robert S. Gray, Jason Liu, and Chip Elliott. Experimental evaluation of wireless simulation assumptions. *Simulation*, 83(9):643–661, 2007. doi:10.1177/0037549707085632.
 - 23 Adrian Ogierman, Andréa W. Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. Competitive MAC under adversarial SINR. In *2014 IEEE Conference on Computer Com-*

- munications, INFOCOM 2014, Toronto, Canada, April 27 - May 2, 2014*, pages 2751–2759. IEEE, 2014. doi:10.1109/INFOCOM.2014.6848224.
- 24 A. Richa, Ch. Scheideler, S. Schmid, and J. Zhang. A jamming-resistant MAC protocol for multi-hop wireless networks. In *Proceedings of the 24th International Conference on Distributed Computing*, DISC '10. Springer-Verlag, 2010.
 - 25 D. Yu, Q. Hua, Y. Wang, and F. C. M. Lau. An $O(\log n)$ distributed approximation algorithm for local broadcasting in unstructured wireless networks. In *IEEE 8th International Conference on Distributed Computing in Sensor Systems*, ICDCS '12. IEEE, 2012. doi:10.1109/dcoss.2012.39.
 - 26 Dongxiao Yu, Qiang-Sheng Hua, Yuexuan Wang, Haisheng Tan, and Francis C. M. Lau. Distributed multiple-message broadcast in wireless ad hoc networks under the SINR model. *Theor. Comput. Sci.*, 610:182–191, 2016. doi:10.1016/j.tcs.2014.06.043.

A Δ -Cover-Free Families

Proof. (Lemma 8). Let \mathbb{F}_q be the finite field with q elements. Let $\mathcal{P}_d = \{a_0 + a_1x + \dots + a_dx^d \mid a_i \in \mathbb{F}_q\}$ be the set of polynomials with coefficients in \mathbb{F}_q and degree at most d . For a polynomial $f \in \mathcal{P}_d$ let $G_f := \{(x, f(x)) \mid x \in \mathbb{F}_q\} \subseteq \mathbb{F}_q^2$ be the graph of f over \mathbb{F}_q . Consider the family $\mathcal{F} := \{G_f \mid f \in \mathbb{F}_q^d\} \subseteq 2^{\mathbb{F}_q^2}$ of graphs of polynomials from \mathcal{P}_d over \mathbb{F}_q .

Graphs of degree d or less intersect at most d times and since $d < q$, all graphs G_f of polynomials $f \in \mathcal{P}_d$ are distinct, hence $|\mathcal{F}| = |\mathcal{P}_d| = q^{d+1}$. For the same reason we have $G_1 \cap G_2 \leq d$ for graphs $G_1, G_2 \in \mathcal{F}$, hence $\lfloor (q-1)/d \rfloor$ graphs are cover-free. \blacktriangleleft

Proof. (Lemma 9). We choose the smallest prime power $q \geq \Delta \log k + 1$. Considering powers of 2 we obviously have $\Delta \log k + 1 \leq q \leq 3\Delta \log k$. Moreover, we choose $d := \log k$. Then we have $\lfloor (q-1)/d \rfloor \geq \lfloor (\Delta \log k)/\log k \rfloor = \Delta$, and due to Lemma 8, there is a Δ -cover-free family \mathcal{F} of size $|\mathcal{F}| = q^{d+1} \geq (\Delta \log k + 1)^{\log k} \geq k$ of subsets of $[q^2]$. \blacktriangleleft

B Neighborhood Dissemination Analysis

Proof. (Lemma 12) The case $|S_\phi| \leq 1$ is covered, since $d_\phi = \infty \geq 2^{\phi-1}d_{\min}$. In case $|S_\phi| \geq 2$ we prove the lemma by induction on ϕ . Obviously, we have $d_1 \geq d_{\min} = 2^0d_{\min}$. Presume that $d_\phi \geq 2^{\phi-1}d_{\min}$ or $d_\phi > r_e$. If $d_\phi > r_e$ is true, we have $d_{\phi+1} \geq d_\phi > r_e$ and are done. Otherwise, we have $d_\phi \geq 2^{\phi-1}d_{\min}$. In Lemma 10 we showed that $S_{\phi+1}$ is a DIS of S_ϕ , thus independent w.r.t. $E_{\text{safe}}^\mu[S_\phi]$. Therefore no two nodes from $S_{\phi+1}$ are connected via a μ -safe edge in $H^\mu[S_\phi]$.

For a contradiction, assume $d_{\phi+1} < 2d_\phi$ and $d_{\phi+1} \leq r_e$. There are two nodes $u, v \in S_{\phi+1}$ with $d(u, v) = d_{\phi+1}$. Let $\mu \in (0, p)$ be the constant from Lemma 4, such that $E_{\text{safe}}^\mu[S_\phi]$ contains all edges not longer than $\min(2d_\phi, r_e)$. Thus, u, v would be connected via a μ -safe edge in $H^\mu[S_\phi]$, a contradiction to the independence of the DIS $S_{\phi+1}$. Consequently, $d_{\phi+1} \geq 2d_\phi \geq 2^\phi d_{\min}$ (second inequality is due to the induction hypothesis) or $d_{\phi+1} > r_e$. \blacktriangleleft

Proof. (Lemma 11). If $X_A^p > \frac{2^\alpha}{\beta_{\min}}$, then u cannot receive a message from w , because

$$\text{SINR}(w, u, I) \leq \frac{P/d(u, w)^\alpha}{I_A^p(u)} \leq \frac{P/d(u, w)^\alpha}{X_A^p P / 2^\alpha d(u, w)^\alpha} < \beta_{\min}.$$

However, since $(w, u) \in E^\mu[S_\phi]$, we have $\text{SINR}(w, u, I) \geq \beta_{\min}$ with probability at least $(1-\varepsilon)\mu$ and therefore $\mathbb{P}(X_A^p \leq \frac{2^\alpha}{\beta_{\min}}) \geq (1-\varepsilon)\mu$. Now assume nodes send with probability p/Q instead of p . We simulate this with a two step random experiment. First, we randomly

determine a set $C \subseteq A$ of candidate nodes, whereas the probability that $x \in A$ becomes a candidate is p . Second, the probability that a candidate $x \in C$ actually sends is $1/Q$. Using the law of total probability and then Markov's inequality we obtain

$$\begin{aligned} \mathbb{P}\left(X_A^{p/Q} \leq \frac{2^{\alpha+1}}{Q\beta_{\min}}\right) &\geq \mathbb{P}\left(|C| \leq \frac{2^\alpha}{\beta_{\min}}\right) \cdot \mathbb{P}\left(X_C^{1/Q} \leq \frac{2^{\alpha+1}}{Q\beta_{\min}} \mid |C| \leq \frac{2^\alpha}{\beta_{\min}}\right) \\ &\geq \mathbb{P}\left(X_A^p \leq \frac{2^\alpha}{\beta_{\min}}\right) \cdot \mathbb{P}\left(X_C^{1/Q} \leq 2\mathbb{E}(X_C^{1/Q})\right) \geq \frac{(1-\varepsilon)\mu}{2}. \end{aligned}$$

Therefore, with probability $\frac{(1-\varepsilon)\mu}{2}$ the interference $I_A^{p/Q}(v)$ is bounded by

$$I_A^{p/Q}(v) = \sum_{x \in A \text{ sends}} \frac{P}{d(x,v)^\alpha} \leq \frac{X_A^{p/Q} P}{d(u,v)^\alpha} = \frac{2^{\alpha+1} P}{Q\beta_{\min} d(u,v)^\alpha}.$$

Next, we limit the maximum interference $I_A^p(u)$ at u stemming from the set of distant nodes \bar{A} . Again, we utilize that $(w, u) \in E^\mu[S_\phi]$.

$$\begin{aligned} (1-\varepsilon)\mu &\leq \mathbb{P}(SINR(w, u, I) \geq \beta_{\min}) = \mathbb{P}\left(\frac{P/d(u,w)^\alpha}{N + I_V^p(u)} \geq \beta_{\min}\right) \\ &\leq \mathbb{P}\left(\frac{P}{I_{\bar{A}}^p(u)d(u,w)^\alpha} \geq \beta_{\min}\right) = \mathbb{P}\left(I_{\bar{A}}^p(u) \leq \frac{P}{\beta_{\min}d(u,w)^\alpha}\right). \end{aligned} \quad (3)$$

Using the specific Chernoff bound from [4] given in Appendix C Lemma 18, we obtain

$$\mathbb{P}\left(I_{\bar{A}}^p(u) \leq \frac{\mathbb{E}(I_{\bar{A}}^p(u))}{2}\right) \leq \mathbb{P}\left(I_{\bar{A}}^p(u) \leq \frac{pI_{\bar{A}}^{\max}(u)}{2}\right) \leq \exp\left(-\frac{p2^\alpha d(u,w)^\alpha}{8P} \cdot I_{\bar{A}}^{\max}(u)\right). \quad (4)$$

We show $I_{\bar{A}}^{\max}(u) \leq \frac{\eta \cdot P}{p\beta_{\min}d(u,w)^\alpha}$ with $\eta := \max\left\{2, \frac{8\beta_{\min}}{2^\alpha} \cdot \ln \frac{1}{(1-\varepsilon)\mu}\right\} \in \Theta(\beta_{\min})$. Assume that $I_{\bar{A}}^{\max}(u) > \frac{\eta \cdot P}{p\beta_{\min}d(u,w)^\alpha}$. Then we obtain a contradiction to Equation 3 as follows

$$\begin{aligned} \mathbb{P}\left(I_{\bar{A}}^p(u) \leq \frac{P}{\beta_{\min}d(u,w)^\alpha}\right) &\stackrel{\eta \geq 2}{\leq} \mathbb{P}\left(I_{\bar{A}}^p(u) \leq \frac{\eta \cdot P}{2\beta_{\min}d(u,w)^\alpha}\right) \leq \mathbb{P}\left(I_{\bar{A}}^p(u) \leq \frac{pI_{\bar{A}}^{\max}(u)}{2}\right) \\ &\stackrel{(4)}{\leq} \exp\left(-\frac{p2^\alpha d(u,w)^\alpha}{8P} I_{\bar{A}}^{\max}(u)\right) < \exp\left(-\frac{2^\alpha \eta}{8\beta_{\min}}\right) \leq (1-\varepsilon)\mu. \end{aligned}$$

The upper bound of $I_{\bar{A}}^{\max}(u)$ can be used to bound $I_{\bar{A}}^{\max}(v)$. For any node $x \in S_\phi$ we have $d(u, x) \leq d(u, v) + d(v, x) \leq 2d(v, x)$, since u is closest to v among the nodes in S_ϕ . Therefore, we obtain

$$I_{\bar{A}}(v) = \sum_{x \in \bar{A}} \frac{P}{d(v, x)^\alpha} \leq \sum_{x \in \bar{A}} \frac{2^\alpha P}{d(u, x)^\alpha} = 2^\alpha I_{\bar{A}}(u).$$

Markov's inequality yields

$$\mathbb{P}\left(I_{\bar{A}}^{p/Q}(v) \leq 2\mathbb{E}\left(I_{\bar{A}}^{p/Q}(v)\right)\right) \leq \mathbb{P}\left(I_{\bar{A}}^{p/Q}(v) \leq \frac{2p}{Q} I_{\bar{A}}^{\max}(v)\right) \geq \frac{1}{2}.$$

Thus, with probability at least $\frac{1}{2}$ we have $I_{\bar{A}}^{p/Q}(v) \leq \frac{2p}{Q} I_{\bar{A}}^{\max}(v) \leq \frac{2^{\alpha+1} p}{Q} I_{\bar{A}}^{\max}(u) \leq \frac{2^{\alpha+1} \eta P}{Q\beta_{\min}d(u,w)^\alpha}$. Finally, the probability that u sends and v listens is $\frac{p}{Q}(1-\frac{p}{Q}) \leq \frac{p}{2Q}$. Events (i), (ii) and (iii) are independent, hence the probability that all of them to occur simultaneously is $\frac{(1-\varepsilon)\mu p}{8Q}$. ◀

Proof. (Lemma 13). Assume that events (i), (ii) and (iii) from Lemma 12 occur. We use the bounds on the interference to show that a safe transmission from u to v takes place. This is the case if $\text{SINR}(u, v, I) \geq \beta_{\max}$, i.e., $\beta_{\max} d(u, v)^\alpha (N + I_A^{p/Q}(v) + I_{\bar{A}}^{p/Q}(v)) \leq P$.

$$\begin{aligned}
& \beta_{\max} d(u, v)^\alpha \left(N + I_A^{p/Q}(v) + I_{\bar{A}}^{p/Q}(v) \right) \\
& \leq \beta_{\max} \left(d(u, v)^\alpha N + \frac{2^{\alpha+1}P}{Q\beta_{\min}} + \frac{2^{\alpha+1}\eta P d(u, v)^\alpha}{Q\beta_{\min} d(u, w)^\alpha} \right) && \text{Lemma 12 (i), (ii)} \\
& \leq \beta_{\max} \left((1 + \frac{\rho}{2})^\alpha r_e^\alpha N + \frac{2^{\alpha+1}P}{\hat{Q}\beta_{\min}} + \frac{2^{\alpha+1}\eta P}{\gamma^\alpha \beta_{\min}} \right) && d(u, v) \leq (1 + \frac{\rho}{2})r_e, d(u, v)^\alpha \leq \frac{Qd(u, w)^\alpha}{\gamma^\alpha} \\
& \leq \beta_{\max} \left((1 - \frac{\alpha\rho}{2(1+\rho)^\alpha})(1+\rho)^\alpha r_e^\alpha N + \frac{2^{\alpha+1}P}{\hat{Q}\beta_{\min}} + \frac{2^{\alpha+1}\eta P}{\gamma^\alpha \beta_{\min}} \right) && (1 + \frac{\rho}{2})^\alpha \stackrel{\text{Lemma 19}}{\leq} (1+\rho)^\alpha - \frac{\alpha\rho}{2} \\
& \leq P \left(1 - \frac{\alpha\rho}{2(1+\rho)^\alpha} \right) + \beta_{\max} \left(\frac{2^{\alpha+1}P}{\hat{Q}\beta_{\min}} + \frac{2^{\alpha+1}\eta P}{\gamma^\alpha \beta_{\min}} \right) && P = \beta_{\max}(1+\rho)^\alpha r_e^\alpha N \\
& = P + P \left(\frac{\beta_{\max} 2^{\alpha+1}}{\beta_{\min} \hat{Q}} + \frac{\beta_{\max} 2^{\alpha+1} \eta}{\beta_{\min} \gamma^\alpha} - \frac{\alpha\rho}{2(1+\rho)^\alpha} \right) \leq P && \hat{Q} \in \mathcal{O}\left(\frac{\beta_{\max}}{\beta_{\min}} 2^\alpha\right), \gamma \in \mathcal{O}\left(\left(\frac{\beta_{\max}}{\beta_{\min}}\right)^{1/\alpha}\right).
\end{aligned}$$

In the last step, $\hat{Q} \in \mathcal{O}\left(\frac{\beta_{\max}}{\beta_{\min}} 2^\alpha\right)$, $\gamma \in \mathcal{O}\left(\left(\frac{\beta_{\max}}{\beta_{\min}}\right)^{1/\alpha}\right)$ are chosen sufficiently large, such that the term in the bracket becomes negative. Due to event (iii) u sends and v listens, thus v receives \mathcal{M} from u . Events (i), (ii) and (iii) occur with probability $\frac{(1-\epsilon)\mu p}{8Q} \in \Theta(1/Q)$. \blacktriangleleft

Proof. (Lemma 14) Note that we assume $\log R \geq 1$ (otherwise the network is 'sparse' and broadcast is easy, see Lemma 1). We proof the claim inductively. Active node $u_1 \in S_1$ is a neighbor of $v \in N(S_1)$ in the communication graph G_C , thus $d(u_1, v) \leq r_e$ by definition. Presume that during phase ϕ either $d(u_\phi, v) \leq r_e(1 + \frac{\rho(\phi-1)}{2\log R})$ or v has already received \mathcal{M} . We show that the same is true for phase $\phi+1$. If v has already received \mathcal{M} we are done. Therefore, we concentrate on the case that $d(u_\phi, v)$ is bound from above.

If there is no active node $w_\phi \in S_\phi$ for which an edge $(w_\phi, u_\phi) \in E^\mu[S_\phi]$ exists, then $u_{\phi+1}$ is not dominated by a node in S_ϕ and therefore $u_{\phi+1} = u_\phi$. With to the induction hypothesis, we obtain $d(u_{\phi+1}, v) \leq r_e(1 + \frac{\rho(\phi-1)}{2\log R}) \leq r_e(1 + \frac{\rho\phi}{2\log R})$. Otherwise, let w_ϕ to be the farthest node from u_ϕ with an edge $(w_\phi, u_\phi) \in E^\mu[S_\phi]$.

If $d(u_\phi, w_\phi) \geq \gamma Q^{-1/\alpha} d(u_\phi, v)$, then Lemma 13 applies and we have a guaranteed probability of $\Theta(1/Q)$ that v receives \mathcal{M} , in case active nodes send with probability p/Q . Since we do exactly that for $\Theta(Q \log n)$ rounds during phase ϕ of algorithm ROBUSTDISSEMINATION, node v receives \mathcal{M} w.h.p. (we adjust the constant in $\Theta(Q \log n)$ accordingly).

Now consider the case $d(u_\phi, w_\phi) < \gamma Q^{-1/\alpha} d(u_\phi, v)$. Since $u_{\phi+1} \neq u_\phi$, node u_ϕ is dominated by some other node $x \in S_{\phi+1}$, which is at distance at most $d(u_\phi, w_\phi)$ from u_ϕ . Thus, the nearest neighbor $u_{\phi+1} \in S_{\phi+1}$ of v is at distance at most $d(v, u_{\phi+1}) \leq d(v, x) \leq d(v, u_\phi) + d(u_\phi, x) \leq d(v, u_\phi) + d(u_\phi, w_\phi)$. We obtain

$$\begin{aligned}
d(v, u_{\phi+1}) & \leq \left(1 + \frac{\gamma}{Q^{1/\alpha}} \right) d(v, u_\phi) \leq r_e \left(1 + \frac{\gamma}{Q^{1/\alpha}} \right) \left(1 + \frac{\rho(\phi-1)}{2\log R} \right) \\
& \stackrel{(*)}{\leq} r_e \left(1 + \frac{\rho(\phi-1)}{2\log R} + \frac{\gamma(1+\rho/2)}{Q^{1/\alpha}} \right) \leq r_e \left(1 + \frac{\rho\phi}{2\log R} \right) && (*): \phi - 1 \leq \log R
\end{aligned}$$

The last step holds for a sufficiently large $Q \in \Theta\left(\frac{\beta_{\max}}{\beta_{\min}} (\log R)^\alpha\right)$, $Q \geq \hat{Q}$. \blacktriangleleft

Proof. (Lemma 15) Since S_ϕ is sparse the premise of Lemma 1 is fulfilled and v receives \mathcal{M} from u with constant probability $\mu' = \frac{p}{Q}(1 - \frac{p}{Q})^{k_0}$ and $k_0 \in \Theta(1)$ (nodes send with probability p/Q in each phase of Algorithm 1). Therefore, v receives \mathcal{M} from u w.h.p. after sending

with probability p/Q for $\mathcal{O}(\frac{\log n}{\mu'})$ rounds.¹⁶ Define the constant probability $\mu := p(1-p)^{k_0}$. We find

$$\mu' = \frac{p}{Q}(1 - \frac{p}{Q})^{k_0} \geq \frac{p}{Q}(1-p)^{k_0} = \frac{\mu}{Q}.$$

Hence $(\log n)/\mu' \leq (Q \log n)/\mu$. Therefore, v also receives \mathcal{M} w.h.p. if u sends with probability p/Q for $\mathcal{O}(\frac{Q \log n}{\mu})$ rounds (i.e. longer). Since u sends for $\Theta(Q \log n)$ rounds in the inner loop of Algorithm 1 (with an appropriately chosen constant factor), we have proven the claim. ◀

C Inequalities

For the sake of completeness, the following lemmas give the specific bounds which we use in the proofs of Lemma 1, Lemma 12 and Lemma 13. We refer to [4] for the proof of Lemma 18.

► **Lemma 17.** *Let $\delta \geq 1$ and $k \in \mathbb{N}$. Then $k^\delta - (k-1)^\delta \leq \delta k^{\delta-1}$.*

Proof. Let $f(k) := k^\delta$. Then $f'(k) := \frac{df(k)}{dk} = \delta k^{\delta-1}$ is monotonous increasing since $\delta \geq 1$. The mean value theorem states that there is a $\xi \in (k-1, k)$, such that $f(k) - f(k-1) = f'(\xi)(k - (k-1)) = f'(\xi) \leq f'(k)$. ◀

► **Lemma 18** (cf. [4]). *Let X_1, \dots, X_k be independent random variables with $\mathbb{P}(X_i = a_i) = p$ and $\mathbb{P}(X_i = 0) = 1-p$ for $a_i > 0$ and $p \in (0, 1)$. Let $A := \sum_{i=1}^k a_i$ and $\hat{a} := \max_{i \in [k]} a_i$. Further let $X := \sum_{i=1}^k X_i$ and $\lambda := \mathbb{E}(X) = pA$. For any $\delta > 0$ it holds that*

$$\mathbb{P}(X \leq (1 - \delta)\lambda) \leq \exp(-\frac{\delta^2 \lambda}{2\hat{a}}).$$

► **Lemma 19.** *Let $\rho \geq 0$ and $\alpha \geq 2$. Then $(1+\rho)^\alpha \leq (1 + \frac{\rho}{2})^\alpha + \frac{\alpha\rho}{2}$.*

Proof. Let $f(x) := x^\alpha - (\frac{x+1}{2})^\alpha - \alpha \frac{x-1}{2}$. We proof the claim by showing $f(1+\rho) \geq 0$. First we note that $f(1) = 0$ fulfills the claim for $\rho = 0$. The claim is proved for $\rho \geq 0$ if we can show $f'(x) := \frac{df(x)}{dx} \geq 0$ for $x \geq 1$. We have $x \geq \frac{x+1}{2}$ (since $x \geq 1$). This is equivalent to $x^{\alpha-1} \geq (\frac{x+1}{2})^{\alpha-1}$ (since $\alpha \geq 2$). Hence $x^{\alpha-1} - \frac{1}{2}(\frac{x+1}{2})^{\alpha-1} \geq x^{\alpha-1} - \frac{1}{2}x^{\alpha-1} \geq \frac{1}{2}$ (since $x \geq 1, \alpha \geq 2$). Finally, we obtain that $f'(x) = \alpha \underbrace{(x^{\alpha-1} - \frac{1}{2}(\frac{x+1}{2})^{\alpha-1})}_{\geq 1/2} - \frac{\alpha}{2} \geq 0$. ◀

¹⁶The math is similar to what we did it in the proof of Lemma 5.