

Broadcast in the Ad Hoc SINR Model

Sebastian Daum¹, Seth Gilbert^{3,*}, Fabian Kuhn¹, and Calvin Newport^{2,**}

¹ Department of Computer Science, University of Freiburg, Germany
{sdaum,kuhn}@cs.uni-freiburg.de

² Department of Computer Science, Georgetown University, USA
cnewport@cs.georgetown.edu

³ Department of Computer Science, National University of Singapore, Singapore
seth.gilbert@comp.nus.edu.sg

Abstract. An increasing amount of attention is being turned toward the study of distributed algorithms in wireless network models based on calculations of the *signal to noise and interference ratio* (SINR). In this paper we introduce the *ad hoc SINR* model, which, we argue, reduces the gap between theory results and real world deployment. We then use it to study upper and lower bounds for the canonical problem of broadcast on the graph induced by both *strong* and *weak* links. For strong connectivity broadcast, we present a new randomized algorithm that solves the problem in $O(D \log(n) \text{polylog}(R))$ rounds in networks of size n , with link graph diameter D , and a ratio between longest and shortest links bounded by R . We then show that for *back-off* style algorithms (a common type of algorithm where nodes do not explicitly coordinate with each other) and *compact* networks (a practice-motivated model variant that treats the distance from very close nodes as equivalent), there exist networks in which centralized algorithms can solve broadcast in $O(1)$ rounds, but distributed solutions require $\Omega(n)$ rounds. We then turn our attention to weak connectivity broadcast, where we show a similar $\Omega(n)$ lower bound for all types of algorithms, which we (nearly) match with a back-off style $O(n \log^2 n)$ -round upper bound. Our broadcast algorithms are the first known for SINR-style models that do not assume synchronous starts, as well as the first known not to depend on power control, tunable carrier sensing, geographic information and/or exact knowledge of network parameters.

1 Introduction

In this paper, we study distributed broadcast in wireless networks. We model this setting using an *SINR-style* model; i.e., communication behavior is determined by the ratio of signal to noise and interference [6, 8–11, 15, 17, 19, 21]. While we are not the first to study broadcast in an SINR-style model (see *related work* below), we are the first to do so under a specific set of assumptions which we call the *ad hoc SINR* model. It generalizes the SINR-style models previously used to study broadcast by eliminating or reducing assumptions that might conflict with real networks, including, notably, idealized uniform signal propagation and knowledge of exact network parameters or geographic information. In this setting, we produce new efficient broadcast upper bounds as well as new lower bounds that prove key limitations. In the remainder of this section, we detail and motivate our model, then describe our results and compare them to existing work.

* Supported by the Singapore Academic Research Fund MOE2011-T2-2-042.

** Supported in part by the Ford University Research Grant program.

The Ad Hoc SINR Model. In recent years, increasing attention has been turned toward studying distributed wireless algorithms in *SINR-style* models which determine receive behavior with an *SINR formula* (see Section 2) that calculates, for a given sender/receiver pair, the ratio of signal to interference and noise at the receiver. These models differ in the assumptions they make about aspects including the definition of distance, knowledge of network parameters, and power control constraints. In this paper we study an SINR-style model with a collection of assumptions that we collectively call the *ad hoc SINR* model, previously studied (however not named yet) e.g. in [7]. Our goal with this model is to capture the key characteristic of wireless communication while avoiding assumptions that might impede the translation of theoretical results into practical algorithms. The ad hoc SINR model is formally defined in Section 2, but we begin by summarizing and motivating it below.

We start by noting that a key parameter in the SINR formula is the distance between nodes. Distance provides the independent variable in determining signal degradation between a transmitter and receiver. In the ad hoc SINR model, we do not assume that distance is necessarily determined by Euclidean geometry. We instead assume only that the distances form a metric in a “growth-bounded metric space”—describing, in some sense, an *effective distance* between nodes that captures both path loss and attenuation. Crucially, we assume this distance function is *a priori* unknown—preventing algorithms that depend on advance exact knowledge of how signals will propagate.

Another key assumption in the definition of an SINR-style model is the nodes’ knowledge of network parameters. In the ad hoc SINR model, we assume nodes do not know the precise value of the parameters associated with the SINR formula (i.e., α, β, N), but instead know only reasonable upper and lower bounds for the parameters (i.e., $\frac{\alpha_{min}}{\alpha_{max}}, \frac{\beta_{min}}{\beta_{max}}, \frac{N_{min}}{N_{max}}$). This assumption is motivated by practice where ranges for these parameters are well-established, but specific values change from network to network and are non-trivial to measure.¹ We also assume that nodes only know a polynomial upper bound on the relevant deployment parameters—namely, network size and density disparity (ratio between longest and shortest links).

Finally, we assume that all nodes use the same fixed constant power. This assumption is motivated by the reality that power control varies widely from device to device, with some chipsets not allowing it all, while others use significantly different granularities. To produce algorithms that are widely deployable it is easiest to simply assume that nodes are provided some unknown uniform power.

Results. The global broadcast problem provides a *source* with a broadcast message M , which it must propagate to all reachable nodes in the network. We study this problem under the two standard definitions of *reachable* for an SINR-style setting: *weak* and *strong*. In more detail, let d_{max} be the largest possible distance such that two nodes u and v can communicate (i.e., the largest distance such that if u broadcasts alone in the entire network, v receives its message). A link between u and v is considered *weak* if their distance is no more than d_{max} , and *strong* if their distance is no more than $\frac{d_{max}}{1+\rho}$, where $\rho = O(1)$ is a constant parameter of the problem. *Weak* (resp. *strong*)

¹ In addition to keeping the specific values unknown, it might be interesting to allow them to vary over time in the range; e.g., an idea first proposed and investigated in [10]. The difficulty of defining such dynamic models lies in introducing the dynamic behavior without subverting tractability. This is undoubtedly an intriguing direction for future exploration.

connectivity broadcast requires the source to propagate the message to all nodes in its connected component in the graph induced by weak (resp. strong) links.

Existing work on broadcast in SINR-style models focuses on strong connectivity. With this in mind, we begin, in Section 4, with our main result: a new strong connectivity broadcast algorithm that terminates in $O(D \log n \log^{\alpha_{max}+1}(R_s))$ rounds, with probability at least $1 - 1/n^c$, for some $c \geq 1$ (w.h.p.), where D is the diameter of the strong link graph, $\alpha_{max} = \alpha + O(1)$ is an SINR model parameter, and R_s is the maximum ratio between strong link lengths. Notice, in most practical networks, R_s is polynomial in n ,² leading to a result that is in $O(D \text{ polylog}(n))$. This is also, to the best of our knowledge, the first broadcast algorithm for an SINR-style model that does not assume synchronous starts. It instead requires nodes to receive the broadcast message first before transmitting—a practical and common assumption, that prevents nodes from needing advance knowledge of exactly when broadcast messages will enter the system.

We then continue with lower bounds for strong connectivity broadcast. In the graph-based models of wireless networks, the best known broadcast solutions are *back-off style* algorithms [2, 4, 12], in which a node’s decision to broadcast depends only on the current round and the round in which it first received the broadcast message. These algorithms are appealing due to their simplicity and ease of implementation. In this paper, we prove that back-off style algorithms are inherently inefficient for solving strong connectivity broadcast. In more detail, we prove that there exist networks in which a centralized algorithm can solve broadcast in a constant number of rounds, but any back-off style algorithm requires $\Omega(n)$ rounds. This result opens a clear separation between the graph and SINR-style models with respect to this problem.

We also prove an $\Omega(n)$ bound on a *compact* version of our model that allows arbitrarily large groups of nodes to occupy the same position. We introduce this assumption to explore a reality of many real networks: when you pack devices close enough, the differences between received signal strength fall below the detection granularity of the radio hardware, which experiences the signal strength of these nearby devices as if they were all traveling the same distance. This bound emphasizes an intriguing negative reality: efficient broadcast in SINR-style models depends strongly, in some sense, on the theoretical conceit that the ratio between distances is all that matters, regardless of how small the actual magnitude of these distance values is.

We conclude by turning our attention to weak connectivity broadcast. To the best of our knowledge, we are the first to concretely consider this version of broadcast. We formalize the intuitive difficulty of this setting by proving the existence of networks where centralized algorithms can solve broadcast in $O(1)$ rounds, while any distributed algorithm requires $\Omega(n)$ rounds. We then match this bound (within $\log^2 n$ factors) by showing that the back-off style upper bound we first presented in our study of the dual graph model [13] not only solves weak connectivity broadcast in $O(n \log^2 n)$ rounds in the ad hoc SINR model, but also does so in essentially *every reasonable model* of a wireless network.

² There are theoretically possible networks, like the exponential line, in which R_s is exponential in n , but as n grows beyond a small value, those networks become impossible to realize in practice. E.g., to deploy an exponential line consisting of ~ 45 nodes, with a maximum transmission range of $100m$, the network would have to include pairs of devices separated by a distance less than the width of a single atom.

Related Work. The theoretical study of SINR-style models began by focusing on centralized algorithms meant to bound the fundamental capabilities of the setting; e.g., [6, 8, 11, 15, 17]. More recently, attention has turned toward studying distributed algorithms, which we discuss here. In the following, n is the network size, D is the diameter of the strong link graph, and Δ is the maximum degree in the weak link graph. Randomized results are assumed to hold with high probability.

We begin by summarizing existing work on distributed strong connectivity broadcast in SINR-style models. There exist several interesting strategies for efficiently performing strong connectivity broadcast. In more detail, in the randomized setting, Scheideler et al. [19] show how to solve strong connectivity broadcast in $O(D + \log n)$ rounds, while Yu et al. [21] present a $O(D + \log^2 n)$ round solution. In the deterministic setting, Jurdzinski et al. [9] describe a $O(\Delta \text{polylog}(n) + D)$ solution, which they recently improved to $O(D \log^2 n)$ (under different assumptions) [10]. However, all of these above solutions make strong assumptions on the knowledge and capability of devices, which are forbidden by the ad hoc SINR model. In particular, all four results leverage knowledge of the exact network parameters (though in [19] it is noted that estimates are likely sufficient), and assume that all nodes begin during round 1 (allowing them to build an overlay structure on which the message is then propagated). In addition, [19] makes use of tunable collision detection, [21] allows the algorithm to specify the transmission power level as a function of the network parameters, [9] adds an additional model restriction that forbids communication over weak links,³ and [10] heavily leverages the assumption that nodes know their positions in Euclidean space and the exact network parameters, and can therefore place themselves and their neighbors in a precomputed overlay grid with nice properties.

A problem closely related to (global) broadcast is *local* broadcast, which requires a set of senders to deliver a message to all neighbors in the strong link graph. This problem is well-studied in SINR-style models and the best known results are of the form $O(\Delta \log n)$ [7, 22]. Of these results, the algorithm in [7] is the most relevant to our work as it deploys an elegant randomized strategy that can be easily adapted to the ad hoc SINR model. Using this local broadcast algorithm as a building block yields a solution for (global) broadcast that runs in $O(\Delta D \log n)$ time. In our work, we avoid dependency on the degree of the underlying link graph as we only need to propagate a single message.

In the classical graph-based wireless network model, for distributed broadcast there is a tight bound of $\Theta((D + \log n) \log(n/D))$ rounds, if nodes start asynchronously (like in this paper) [1, 2, 4, 12, 14, 18]. For the easier case where all nodes start at the same time, it is currently unknown whether or not better bounds are possible in general graphs, but in unit disk graphs a solution of the form $O(D + \log^2 n)$ is likely possible.⁴

³ In slightly more detail, their model forbids v from receiving a message from u if u is too far away, even if the SINR of the transmission is above β . This restriction makes it easier to build a useful dominating set because it eliminates the chance that you are dominated by a weakly connected neighbor.

⁴ The result of [16] can build a maximal independent set in the UDG graph model in $O(\log^2 n)$ rounds. Once this set is established under these constraints, an additional $O(\log^2 n)$ rounds should be enough to build a constant-degree overlay—e.g., as in [3]—on which broadcast can be solved in an additional $O(D + \log n)$ rounds.

2 Model

We study the ad hoc SINR model, which describes a network consisting of a set of nodes V deployed in a metric space and communicating via radios. We assume time is divided into synchronous rounds and in each round a node can decide to either transmit or listen. We determine the outcome of these communication decisions by the standard *SINR formula*, which dictates that $v \in V$ receives a message transmitted by $u \in V$, in a round where the nodes in $I \subseteq V \setminus \{u, v\}$ also transmit, if and only if v is listening and

$$\text{SINR}(u, v, I) = \frac{\frac{P_u}{d(u,v)^\alpha}}{N + \sum_{w \in I} \frac{P_w}{d(w,v)^\alpha}} \geq \beta,$$

where P_x is the transmission power of node x , d is the distance formula for the underlying metric space, and $\alpha \in [\alpha_{\min}, \alpha_{\max}]$, $\beta \in [1, \beta_{\max}]$, and $N \in [0, N_{\max}]$, where α_{\max} , β_{\max} and N_{\max} are constants.

In this paper, we assume that: (1) Algorithms are distributed. (2) All nodes use the same constant power P . (3) Nodes do not have advance knowledge of their locations, distances to other nodes, or the specific values of the network parameters α , N , and β , though they do know the range of values from which α , N , and β are chosen. In addition, nodes only know a polynomial upper bound on the standard deployment parameters: the network size ($|V| = n$) and the density (ratio of longest to shortest link distance). (4) Nodes are embedded in a general metric space with a distance function d that satisfies the following property: for every $v \in S \subseteq V$ and constant $c \geq 1$, the number of nodes in S within distance $c \cdot d_{\min}(S)$ of v is in $O(c^\delta)$, where $d_{\min}(S) := \min_{u, u' \in V} \{d(u, u')\}$ is the minimum distance between two nodes in S and $\delta < \alpha_{\min}$ is a fixed constant roughly characterizing a dimension of the metric space. Notice, for $\delta = 2$ the model strictly generalizes the Euclidean plane. We prefer this general notion of distance over standard Euclidean distance as it can capture power degradation due to both path loss *and* attenuation (a link-specific loss of power due to the materials through which the signal travels). In this paper, to achieve the strongest possible results, we prove our upper bounds with respect to this general metric, and our lower bounds with respect to the restricted (i.e., easier for algorithms) two-dimensional Euclidean instantiation.

Compact Networks. The SINR equation is undefined if it includes the distance 0. As motivated in the introduction, a natural question is to ask what happens as distances become effectively 0 (e.g., when nodes become too close for the difference in their signal strength to be detectable). To study this case, we define the *compact ad hoc SINR* model, which allows zero-distances and specifies that whenever $\text{SINR}(u, v, I)$ is therefore undefined, we determine receive behavior with the following rule: v receive u 's message if and only if u is the only node in $I \cup \{u\}$ such that $d(u, v) = 0$. We formalize the impact of this assumption in our lower bound in Section 5.1.

3 Problem and Preliminaries

In this section we define the problems we study in this paper and then introduce some preliminary results that will aid our bounds in the sections that follow.

The Broadcast Problem. In the broadcast problem, a designated source must propagate a message M to every reachable node in the network. Let $r_w := \left(\frac{P}{\beta N}\right)^{1/\alpha}$ be the maximum distance at which any two nodes can communicate. Let $r_s := \frac{r_w}{1+\rho}$, for some known constant $\rho > 0$. Fix a set of nodes and a distance metric. We define $E[\ell]$, for some distance $\ell \geq 0$, to be the set of all pairs $\{u, v\} \subseteq V$ such that $d(u, v) \leq \ell$. When defining broadcast, we consider both the *weak connectivity graph* $G_w = (V, E[r_w])$ and the *strong connectivity graph* $G_s = (V, E[r_s])$. The values $R_w = \max_{\{u,v\}, \{x,y\} \in E[r_w]} \left\{ \frac{d(u,v)}{d(x,y)} \right\}$ and $R_s = \max_{\{u,v\}, \{x,y\} \in E[r_s]} \left\{ \frac{d(u,v)}{d(x,y)} \right\}$ capture the diversity of link lengths in the connectivity graphs. For most networks, you can assume this value to be polynomial in n , though there are certain malformed cases, such as an exponential line, where the value can be larger. A subset $S \subseteq V$ of the nodes is called a *maximal independent set (MIS)*, if any two nodes $u, v \in S$ are *independent*, i.e., $\{u, v\} \notin E$, and if all nodes $v \in V$ are *covered* by some node in $s \in S$, i.e., $\forall v \in V: \exists s \in S: v \in N(s)$.

In *weak connectivity broadcast* the source is required to propagate its message to all nodes in its connected component in G_w , while in *strong connectivity broadcast* the source is required only to propagate the message to all nodes in its component in G_s . In this paper, we are interested in randomized solutions to both broadcast problems. In particular, we say algorithm \mathcal{A} *solves* weak or strong connectivity broadcast in a given number of rounds if it solves the problem in this time w.h.p.; i.e., with probability at least $1 - 1/n^c$, for an arbitrary constant $c > 0$.

We assume nodes remain inactive (i.e., they do not transmit) until they receive the broadcast message for the first time, at which point they become active. We say a given network is *T-broadcastable* with respect to strong or weak connectivity, if there exists a T -round schedule of transmissions that solves the relevant broadcast problem. And finally, we say a broadcast algorithm is a *back-off style* algorithm if nodes base their broadcast decisions entirely on the current round and the round in which they first received the broadcast message (which, for the source, we say is round 0).

The (x, y) -Hitting Game. Our lower bound arguments in this paper deploy the high-level strategy of proving that solving the relevant type of broadcast is at least as hard as solving an easily bounded combinatorial game we call *(x, y) -hitting*. This game is defined for two integers, $0 < x \leq y$. The game begins with an adversary choosing some arbitrary target set $T \subseteq [y]$ where $|T| = x$. The game then proceeds in rounds. In each round the player, modeled as a probabilistic automaton \mathcal{P} , guesses a value $w \in [y]$. If $w \in T$ the player wins. Otherwise it moves on to the next round. It is easy to see that for small x the game takes a long time to solve with reasonable probability:

Theorem 1. *Let \mathcal{P} be a player that solves the (x, y) -hitting game in $f(x, y)$ rounds, in expectation. It follows that $f(x, y) = \Omega(y/x)$.*

4 Strong Connectivity Broadcast

In this section, we present STRONGCAST, an algorithm that solves strong connectivity broadcast in the ad hoc SINR model. We prove the following:

Theorem 2. *The STRONGCAST algorithm solves strong connectivity broadcast in the ad hoc SINR model in $O(D(\log^{\alpha_{\max}+1} R_s)(\log n))$ rounds.*

For most practical networks, R_s is polynomial in n , reducing the above result to $O(D \text{polylog}(n))$. In some malformed networks, however, R_s can be as large as exponential in n . Because we assume the ad hoc SINR model, our algorithm leverages no advanced knowledge of the distance metric and uses only the provided constant upper bounds on α and β , and the polynomial upper bounds on n and R_s . To avoid the introduction of extra notation, we use the exact values of n and R_s in our analysis as those terms show up only within \log factors in big- O notation; for simplicity of presenting the protocol, we also assume that R_s grows at least logarithmic in n .⁵ To keep the analysis of the STRONGCAST algorithm concise, in the following we only present proof sketches. Full proofs for all claims of the section appear in [5].

Algorithm Overview. The STRONGCAST algorithm consists of at most D epochs. In each epoch, the broadcast message is propagated at least one hop further along all shortest paths from the source. In more detail, at the beginning of each epoch, we say a node is *active* with respect to that epoch if it has previously received the message and it has not yet terminated. During each epoch, the active nodes for the epoch execute a sub-protocol we call *neighborhood dissemination*. Let S be the set of active nodes for a given epoch. The goal of neighborhood dissemination is to propagate the broadcast message to every node in $N(S)$, where N is the neighbor function over the strong connectivity graph G_s . (Notice that the high-level structure of our algorithm is the same as seen in the classical results from the graph-based setting; e.g., our neighborhood dissemination sub-protocol takes the place of the *decay* sub-protocol in the canonical broadcast algorithm of Bar-Yehuda et al. [2].)

The neighborhood dissemination sub-protocol divides time into phases. As it progresses from phase to phase, the number of nodes still competing to broadcast the message decreases. The key technical difficulty is reducing contention fast enough that heavily contended neighbors of S receive the message efficiently, but not so fast that some neighbors fail to receive the message before all nearby nodes in S have terminated. We achieve this balance with a novel strategy in which nodes in S approximate a subgraph of their “reliable” neighbors, then build an MIS over this subgraph to determine who remains active and who terminates. We will prove that if a node $u \in S$ neighbors a node $v \in N(S)$, and u is covered by an MIS node (and therefore terminates), the MIS node that covered u must be sufficiently close to v to still help the message progress.

In Section 4.1 we detail a process for constructing a reliable subgraph and analyze its properties. Then, in Section 4.2 we detail the neighborhood dissemination sub-protocol (which uses the subgraph process) and analyze the properties it guarantees. We conclude, in Section 4.3, by pulling together these pieces to prove the main theorem from above.

4.1 SINR-Induced Graphs

The neighborhood dissemination sub-protocol requires active nodes to construct, in a distributed manner, a subgraph that maintains certain properties. For clarity, we describe and analyze this process here before continuing in the next section with the description of the full neighborhood dissemination sub-protocol.

⁵ In fact, it is sufficient to assume $\log^{\alpha_{\max}} R_s = \Omega(\log^* n)$.

We start by defining graphs $H_p^\mu[S]$ which are induced by a node set S , a transmission probability p and a reliability parameter $\mu \in (0, p) \cap \Omega(1)$. Given a set of nodes S , assume that each node in S independently transmits with probability p . Further, assume that there is no interference from any node outside the set S . We define $H_p^\mu[S]$ to be the undirected graph with node set S and edge set $E_p^\mu[S]$ such that for any $u, v \in S$, edge $\{u, v\}$ is in $E_p^\mu[S]$ if and only if both: (i) u receives a message from v with probability at least μ and (ii) v receives a message from u with probability at least μ .

Computing SINR-Induced Graphs. It is difficult to compute the graphs $H_p^\mu[S]$ exactly and efficiently with a distributed algorithm. However, for given S , p , and μ , there is a simple protocol to compute a good approximation $\tilde{H}_p^\mu[S]$ for $H_p^\mu[S]$ (assuming that the reception probabilities for nodes in S do not change over time). Formally, we say that an undirected graph $\tilde{H}_p^\mu[S]$ with node set S is an ε -close approximation of $H_p^\mu[S]$ if and only if:

$$E[H_p^\mu[S]] \subseteq E[\tilde{H}_p^\mu[S]] \subseteq E[H_p^{(1-\varepsilon)\mu}[S]].$$

An ε -close approximation $\tilde{H}_p^\mu[S]$ of $H_p^\mu[S]$ can be computed in time $O(\frac{\log n}{\varepsilon^2 \mu})$ as follows. First, all nodes in S independently transmit their IDs with probability p for $T := c \frac{\log n}{\varepsilon^2 \mu}$ rounds (where the constant c is chosen to be sufficiently large). Each node u creates a list of potential neighbors containing all nodes from which u receives a message in at least $(1 - \varepsilon/2)\mu T$ of those T rounds. For a second iteration of T rounds, each node transmits its list of potential neighbors (as before, by independently transmitting with probability p). At the end, node u adds node v as a neighbor in $\tilde{H}_p^\mu[S]$ if and only if v is in u 's list of potential neighbors and u receives a message from v indicating that u is in v 's list of potential neighbors as well.

The following lemma results from a basic Chernoff bound, observing that: (i) if u and v are neighbors in $H_p^\mu[S]$, then u receives at least μT messages from v , in expectation, and (ii) if u and v are not neighbors in $H_p^{(1-\varepsilon)\mu}[S]$ then u receives at most $(1 - \varepsilon)\mu T$ messages from v , in expectation.

Lemma 3. *W.h.p., the SINR-Induced Graph Computation protocol runs in $O(\frac{\log n}{\varepsilon^2 \mu})$ rounds and returns a graph $\tilde{H}_p^\mu[S]$ that is an ε -close approximation of $H_p^\mu[S]$.*

Properties of SINR-Induced Graphs. In addition to the fact that nodes in an SINR-induced graph can communicate reliably with each other, we point out two other properties. First, we remark that the maximum degree of $H_p^\mu[S]$ is bounded by $1/\mu = O(1)$, because in a single time slot, a node u can receive a message from only one other node v . consequently the second iteration requires messages of size $O(\frac{\log n}{\mu}) = O(\log n)$. Further, as shown by the next lemma, for suitable μ , the graph $H_p^\mu[S]$ contains (at least) all the edges that are very short.

Lemma 4. $\forall p \in (0, 1/2], \exists \mu \in (0, p)$ such that: Let $d_{\min} \leq r_s$ be the shortest distance between any two nodes in S . Then the graph $H_p^\mu[S]$ contains all edges between pairs $u, v \in S$ for which $d(u, v) \leq \min\{2d_{\min}, r_s\}$.

Proof Sketch. We restrict our attention to the case $d_{\min} \leq r_s/2$. If the minimum distance is between $r_s/2$ and r_s , the claim can be shown by a similar, simpler argument.

Consider some node $u \in S$. Due to the underlying metric space in our model, there are at most $O(k^\delta)$ nodes in S within distance kd_{\min} of node u . Let v be a node at distance at most $2d_{\min}$ from u . For any constant k_0 , with probability $\Omega(1)$, node v is the only node transmitting among all the nodes within distance k_0d_{\min} from node u . Further, assuming that all nodes at distance greater than k_0d_{\min} transmit, the interference $I(u)$ at u can be bounded from above by $\kappa(k_0) \cdot P/d_{\min}^\alpha$, where $\kappa(k_0) > 0$ goes to 0 polynomially with k_0 . We therefore get

$$\frac{\frac{P}{d(u,v)^\alpha}}{N + \kappa(k_0) \frac{P}{d_{\min}^\alpha}} \geq \frac{\frac{P}{(2d_{\min})^\alpha}}{N + \kappa(k_0) \frac{P}{d_{\min}^\alpha}} \geq \frac{\frac{P}{r_s^\alpha}}{\frac{P}{\beta r_w^\alpha} + \kappa(k_0) \frac{2^\alpha P}{r_s^\alpha}} = \frac{\beta}{\frac{1}{(1+\rho)^\alpha} + \kappa(k_0)\beta 2^\alpha} \geq \beta.$$

The second inequality follows from $N = \frac{P}{\beta r_w^\alpha}$ and from $d_{\min} \leq r_s/2$. The last inequality holds for sufficiently large k_0 . If we choose μ to be the probability that no more than one node in a ball of radius k_0d_{\min} transmits, then node v can transmit to u with probability μ . \square

In the above proof, μ depends on the unknown parameter β , so we use β_{\max} as the base for computing μ . Note also that since $H_p^\mu[S] \subseteq \tilde{H}_p^\mu[S]$, the lemma induces the same properties on $\tilde{H}_p^\mu[S]$ with high probability.

4.2 Neighborhood Dissemination Sub-protocol

We can now describe the full operation of our neighborhood dissemination sub-protocol (depicted in Algorithm 1). We assume the sub-protocol is called by a set $S \subset V$ of nodes that have a message M that they are trying to disseminate to all nodes in $N(S)$, where N is the neighbor function over G_s . Since every node in S has already received the message M , which originated at the source node s , we can assume that all the nodes in S have been synchronized by s and therefore align their epoch boundaries and call the sub-protocol during the same round.

The protocol proceeds in phases $\phi = 1, 2, \dots, \Phi$, with $\Phi = O(\log R_s)$. Each phase ϕ , the protocol computes a set S_ϕ , such that $S_1 = S$ and for all $\phi \geq 2$, $S_\phi \subset S_{\phi-1}$. The nodes in S_ϕ attempt to send M to nodes in $N(S)$, while the remaining “inactive” nodes remain silent. Each phase is divided into three blocks. In block 1 of phase ϕ , the nodes compute an ε -close approximation $\tilde{H}_p^\mu[S_\phi]$ of the graph $H_p^\mu[S_\phi]$ using the SINR-induced graph computation process described in Section 4.1. We choose $\mu > 0$ appropriately as described in Lemma 4, while $\varepsilon, p \in (0, 1/2)$ can be chosen freely.⁶

In block 2, nodes in S attempt to propagate the message to neighbors in $N(S)$. In more detail, during this block, each node in S_ϕ transmits M with probability p/Q for $T_{\text{phase}} = O(Q \log n)$ rounds, where $Q = \Theta(\log^{\alpha_{\max}} R_s)$ has an appropriately large hidden constant.

In block 3, the nodes in S_ϕ compute the set $S_{\phi+1}$ by finding a maximal independent set (MIS) of $\tilde{H}_p^\mu[S_\phi]$. Only the nodes in this set remain in $S_{\phi+1}$. Notice that building this MIS is straightforward. This can be accomplished by simulating the reliable message-passing model on our subgraph and then executing the $O(\log^* n)$ MIS algorithm from [20] on this simulated network. (This algorithm requires a growth-bounded

⁶ By Lemma 4, μ depends on p ; thus p could be chosen to maximize μ .

Algorithm 1. High-level pseudo-code for one epoch of STRONGCASTInput: $n, R_s, \alpha_{\max}, \beta_{\max}, \varepsilon, p$ Initialization: $Q = Q(p, R_s, \alpha_{\max}) = \Theta(\log^{\alpha_{\max}} R_s), \mu = \mu(p, \beta_{\max}) = \Omega(1), \Phi = O(\log R_s), S_1 = S$ **for** $\phi = 1$ **to** Φ **do** Compute SINR-induced graph $\tilde{H}_p^\mu[S_\phi]$ within $O(\frac{\log n}{\varepsilon^2 \mu})$ rounds ▷ Block 1 **for** $O(Q \log n)$ rounds **do** ▷ Block 2 Each round transmit M with probability $\frac{p}{Q}$ Compute MIS $S_{\phi+1}$ on $\tilde{H}_p^\mu[S_\phi]$ within $O(\frac{\log n}{\varepsilon^2 \mu} \log^* n)$ rounds ▷ Block 3

property which is, by definition, satisfied by any sub-graph of G_s .) Turning our attention to the simulation, we note that by the definition of $\tilde{H}_p^\mu[S_\phi]$, a single round of reliable communication on $\tilde{H}_p^\mu[S_\phi]$ can be easily simulated by having each node in S_ϕ transmits with probability p for $O(\log n)$ consecutive $((1 - \varepsilon)\mu$ -reliable) rounds. Therefore, the MIS construction takes $O(\log n \log^* n)$ rounds.

We now turn our attention to analyzing this protocol. The most technically demanding chore we face in this analysis is proving the following: If a node $u \in S_\phi$ has an uninformed neighbor $v \in N(S)$, then either u gets the message to v in block 2, or u remains in $S_{\phi+1}$, or there is some $w \in S_{\phi+1}$ that is sufficiently close to v to take u 's place in attempting to get the message to v .

Neighborhood Dissemination Analysis. In the following, we show that for appropriate parameters μ, Q , and T_{phase} , the described algorithm solves the neighborhood dissemination problem for S , w.h.p. We first analyze how the sets S_ϕ evolve. In the following, let d_ϕ be the minimum distance between any two nodes in S_ϕ .

Lemma 5. *If the constant μ is chosen to be sufficiently small, w.h.p., the minimum distance between any two nodes in S_ϕ is at least $d_\phi \geq 2^{\phi-1} \cdot d_{\min}$.*

Proof. We prove the claim by induction on ϕ . First, by the definition of d_{\min} , we clearly have $d_1 \geq 2^0 d_{\min} = d_{\min}$. Also, by the definition of an ε -close approximation of $H_p^\mu[S_\phi]$ and by Lemma 4, for a sufficiently small constant μ , w.h.p., $H_p^\mu[S_\phi]$ contains edges between all pairs of nodes $u, v \in S_\phi$ at distance $d(u, v) \leq 2d_\phi$. Because $S_{\phi+1}$ is a maximal independent set of $\tilde{H}_p^\mu[S_\phi]$, nodes in $S_{\phi+1}$ are at distance more than $2d_\phi$ and therefore using the induction hypothesis, we get $d_{\phi+1} > 2d_\phi \geq 2^\phi d_{\min}$. \square

Next we consider node v that needs the message, and its closest neighbor u in S_ϕ . We show that if u and v are sufficiently close, and if the farthest neighbor of u in S_ϕ is also ‘‘sufficiently far’’ away, then u can successfully transmit the message to v .

Lemma 6. $\forall p \in (0, 1/2], \exists \hat{Q}, \gamma = \Theta(1)$, such that for all $Q \geq \hat{Q}$ the following holds. Consider a round r in phase ϕ where each node in S_ϕ transmits the broadcast message M with probability p/Q . Let $v \in N(S)$ be some node that needs to receive M , and let $u \in S_\phi$ be the closest node to v in S_ϕ . Further, let d_u be the distance between u and its farthest neighbor in $\tilde{H}_p^\mu[S_\phi]$. If $d(u, v) \leq (1 + \rho/2)r_s$ and $d_u \geq \gamma Q^{-1/\alpha} \cdot d(u, v)$, node v receives M in round r with probability $1/\Theta(Q)$.

Proof Sketch. The lemma states under what conditions in round r of block 2 in phase ϕ a node $v \in N(S) \setminus S$ can receive the message. The roadmap for this proof is to show

that if u is able to communicate with probability $(1 - \varepsilon)\mu$ with its farthest neighbor u' in some round r' of block 1 in phase ϕ , using the broadcast probability p , then u must also be able to reach v with probability $1/\Theta(Q)$ in round r of block 2, in which it transmits with probability p/Q . We start with some definitions and notation, and continue with a connection between the interference at u and at v . We then analyze the interference at u created in a ball of radius $2d_u$ around u , as well as the remaining interference coming from outside that ball. Finally, we transfer all the knowledge we gained for round r' to round r to conclude the proof.

For a node $w \in V$, let $I(w) = \sum_{x \in S_\phi} \frac{P}{d(x,w)^\alpha}$, i.e., the amount of interference at node w if all nodes of S_ϕ transmit. For round r' , the random variable $X_x^p(w)$ denotes the actual interference at node w coming from a node $x \in S$. The total interference at node w is thus $X^p(w) := \sum_{x \in S_\phi} X_x^p(w)$. If we only want to look at the interference stemming from nodes within a subset $A \subseteq S_\phi$, we use $I_A(w)$ and $X_A^p(w)$, respectively. Further for a set $A \subseteq S_\phi$, we define $\bar{A} := S_\phi \setminus A$.

The triangle inequality implies that $d(u, w) \leq d(u, v) + d(v, w) \leq 2d(v, w)$ for any $w \in S_\phi$. By comparing $I_{S'}(u)$ and $I_{S'}(v)$ for an arbitrary set $S' \subseteq S_\phi$ we obtain the following observation:

$$I_{S'}(u) \geq 2^{-\alpha} I_{S'}(v). \tag{1}$$

Let u' be the farthest neighbor of node u in $\tilde{H}_p^\mu[S_\phi]$. Because $\tilde{H}_p^\mu[S_\phi]$ is an ε -close approximation of $H_p^\mu[S_\phi]$, we know that $\tilde{H}_p^\mu[S_\phi]$ is a subgraph of $H_p^{(1-\varepsilon)\mu}[S_\phi]$ and therefore in round r' , u receives a message from u' with probability at least $(1 - \varepsilon)\mu$.

Let $A \subseteq S_\phi$ be the set of nodes at distance at most $2d_u$ from u . Note that both u and u' are in A , because $d(u, u') = d_u$. In round r' , if more than $2^\alpha/\beta = O(1)$ nodes $u'' \in A$ transmit, then node u cannot receive a message from u' . Since node u receives a message from u' with probability at least $(1 - \varepsilon)\mu$ in round r' , we can conclude that fewer than $2^\alpha/\beta$ nodes transmit with at least the same probability.

We now bound the interference from nodes outside of A . Using the fact that node u receives a message from node u' with constant probability at least $(1 - \varepsilon)\mu$ allows us to upper bound $I_{\bar{A}}(u)$ and by (1) also $I_{\bar{A}}(v)$. For node u to be able to receive a message from u' , two things must hold: (I) u' transmits and u listens (event $R^{u',u}$) and (II) $\frac{P}{d_u^\alpha(N+X_A^p(u))} \geq \frac{P}{d_u^\alpha(N+X^p(u))} \geq \beta$. Thus we have

$$(1-\varepsilon)\mu \leq \mathbb{P}(R^{u',u}) \cdot \mathbb{P}\left(X_A^p(u) \leq \frac{P}{\beta d_u^\alpha} - N\right) \leq p(1-p) \cdot \mathbb{P}\left(X_A^p(u) \leq \frac{P}{\beta d_u^\alpha}\right). \tag{2}$$

Using a Chernoff result (see [5]), we can bound $X_A^p(u)$ as

$$\mathbb{P}\left(X_A^p(u) \leq \frac{\mathbb{E}[X_A^p(u)]}{2}\right) = \mathbb{P}\left(X_A^p(u) \leq \frac{pI_{\bar{A}}(u)}{2}\right) \leq e^{-\frac{p2^\alpha d_u^\alpha}{8P} \cdot I_{\bar{A}}(u)}. \tag{3}$$

Together, (2) and (3) imply that $I_{\bar{A}}(u) = O(P/d_u^\alpha)$. Hence if each node transmits with probability p/Q , by (1), with constant probability, the interference from nodes in A at v is bounded by $O(p/Q \cdot P2^\alpha/d_u^\alpha)$. Since in addition, with probability $1/\Theta(Q)$, u is the only node in A transmitting, by choosing $Q = \Omega(2^\alpha)$ sufficiently large, node v receives M with probability $1/\Theta(Q)$.

4.3 Proof Sketch of Theorem 2

Proof Sketch. First note that by construction, every phase of the neighborhood dissemination protocol has a time complexity of $O((\log^* n + Q) \log n) = O(\log^\alpha R_s \log n)$ (recall that we assumed that R_s is at least logarithmic in n). The claim of the theorem immediately follows if we show that assuming that all algorithm parameters are chosen appropriately, (I) the number of phases Φ of the neighborhood protocol is $O(\log R_s)$, and (II) the neighborhood dissemination protocol is correct, i.e., when carried out by a set S of nodes, w.h.p., each node $v \in N(S)$ receives the broadcast message M .

We prove statements (I) and (II) together. Let v be any node in $N(S)$ and let u_ϕ be the closest node to v in S_ϕ . Since $v \in N(S)$, we have $d(u_1, v) \leq r_s$. Recall that in block 2 of a phase ϕ of the neighborhood dissemination protocol, S_ϕ broadcasts M with probability p/Q for sufficiently large interval of $O(Q \log n)$ rounds. Therefore, by choosing $Q = O(\log^\alpha R_s)$ sufficiently large, by Lemma 6, for all $\phi \in \{1, \dots, \Phi\}$, either $d(u_\phi, v) \leq d(u_1, v)(1 + \phi\gamma Q^{-1/\alpha}) \leq (1 + \rho/2)d(u_1, v)$ or v has already received the message at the start of phase ϕ . As we also know by Lemma 5 that the minimum distance between nodes in S_ϕ grows exponentially with ϕ , it follows that for some $\phi \leq \Phi$, the minimum distance between nodes in S_ϕ exceeds r_s at which point a node within distance $(1 + \rho/2)d(u_1, v) \leq (1 + \rho/2)r_s$ of v trivially reaches node v . \square

5 Lower Bounds for Strong Connectivity Broadcast

In this section, we present lower bounds for strong connectivity broadcast. For complete proofs we refer to [5].

5.1 Lower Bound for Compact Networks

In the compact variant of the ad hoc SINR model (defined in Section 2 and motivated in Section 1) nodes can formally occupy the same position (have mutual distance of 0), which informally captures the real world scenario where the difference in strength of signals coming from a group of nodes packed close enough together are too small to detect, making it seem as if they are all traveling the same distance. Here we prove this assumption makes efficient broadcast impossible.

Theorem 7. *Let \mathcal{A} be a strong connectivity broadcast algorithm for the compact ad hoc SINR model. There exists an $O(1 + \rho)$ -broadcastable network in which \mathcal{A} requires $\Omega(n)$ rounds to solve broadcast.*

Proof Sketch. We reduce the $(\lceil(1 + \rho)\rceil, n)$ -hitting game broadcast in a specific difficult compact network. We construct a network with $k + 2 = \lceil \rho + 1 \rceil + 2$ nodes located uniformly along a line of length $r_w + \epsilon$, for some $\epsilon > 0$, and $n - (k + 2)$ additional nodes placed at one end of the line in the same position. Broadcast to the lone node at the opposite end of the line can only succeed when exactly one node in the middle of the line decides to broadcast by itself. Until that happens, interference prevents all nodes from learning anything. Hence solving broadcast requires solving the hitting game (i.e., choosing one of the k internal nodes on the line). \square

5.2 Lower Bound for Back-Off Style Algorithms

In the study of broadcast in *graph-based* models, the best known algorithms are often back-off style algorithms (e.g., the canonical solution of Bar-Yehuda et. al. [2]). We prove below that such algorithms are too simple to solve strong connectivity broadcast efficiently in the SINR setting.

Theorem 8. *Let \mathcal{A} be a back-off style strong connectivity broadcast algorithm for the ad hoc SINR model. There exists an $O(1 + \rho)$ -broadcastable network in which \mathcal{A} requires $\Omega(n)$ rounds to solve broadcast.*

Proof Sketch. The proof is similar to that of Theorem 7 where we reduce an (x, n) -hitting game to broadcast. As before, we begin with $k + 2 = \lceil 1 + \rho \rceil + 2$ nodes distributed along a vertical line of length $r_w + \epsilon$ for some $\epsilon > 0$. Since we are no longer in a compact network, we cannot place the remaining $n - k + 2$ nodes in the same position at one end of the line. Instead, we spread the remaining nodes uniformly on a horizontal line perpendicular to one end of the existing vertical line. The spacing is small enough that the nodes remain within distance r_w of every other node, except for the one lone node at the far end of the line. Since the network is no longer compact, nodes can now succeed in communicating amongst themselves before the hitting game is won. However, since the algorithm is assumed to be back-off style, this additional communication is ignored and cannot affect their behavior. As before, the nodes are reduced to guessing which k nodes among n total with the message are among those able to solve broadcast. \square

6 Weak Connectivity Broadcast

Weak connectivity broadcast is more difficult than strong connectivity broadcast because it might require messages to move across weak links (links at distance near r_w). When communicating over such a long distance, it is possible for most other nodes in the network to be *interferers*—capable of disrupting the message, but not capable of communicating with the receiver themselves—reducing possible concurrency.

In this section we formalize this intuition by proving that there is a 2-broadcastable network in which all algorithms require $\Omega(n)$ rounds to solve weak connectivity broadcast. We then turn our attention to upper bounds by reanalyzing an algorithm we originally presented in [13], in the context of the *dual graph* model, to show that it solves weak connectivity broadcast in the ad hoc SINR model in $O(n \log^2 n)$ rounds. To the best of our knowledge, this is the first known non-trivial weak connectivity broadcast algorithm for an SINR-style model (all previous broadcast algorithms make stronger assumptions on connectivity). To help underscore the surprising universality of this algorithm, we prove that not only does it solve broadcast in this time in *this* model, but that it works in this time essentially in *every standard wireless model* (a notion we formalize below).

6.1 Lower Bound

Theorem 9. *Let \mathcal{A} be weak connectivity broadcast algorithm for the ad hoc SINR model. There exists a 2-broadcastable network in which \mathcal{A} requires $\Omega(n)$ rounds to solve broadcast.*

Proof Sketch. We leverage the same general approach as the lower bounds in Section 5: We reduce (x, y) -hitting to the relevant broadcast problem, and then apply the bound on hitting from Theorem 1. In our reduction, we use a *rotating lollipop* network, consisting of a circle of $n - 1$ nodes with the message and a receiver at distance r_w from some unknown *bridge* node in the circle (and strictly more distant from all others). To get the message from the circle to the receiver requires that this bridge node broadcast alone. We prove that identifying this bridge node is at least as hard as solving the $(1, n - 1)$ -hitting game, which we know requires $\Omega(n)$ rounds. (See [5] for a detailed proof.) \square

6.2 Upper Bound

In [13], we described a simple back-off style algorithm that solves broadcast in the *dual graph* model—a variant of the classical graph-based wireless model that includes unreliable links controlled by an adversary. In this section, we show that this algorithm solves the basic definition of broadcast in $O(n \log^2 n)$ rounds in every “standard” wireless network model. The fact that it does so in the ad hoc SINR model is an immediate corollary.

First, we consider a broadcast algorithm *universal*, if it distributes the message to every node in the *isolation graph*, defined as the directed graph $G = (V, E)$, where $(u, v) \in E$ if and only if v can receive a message M if u broadcasts M alone in the network. (See [5] for a more formal definition.)

We next describe the broadcast algorithm HARMONICAST, first presented in [13], and show that it solves broadcast in most standard wireless network models. The algorithm works as follows: Let t_v be the round in which node v first receives the broadcast message (if v is the source, $t_v = 0$). Let H be the harmonic series on n , then each round $t \in [t_v + 1, t_v + T]$, for $T = n \lceil 24 \ln n \rceil H(n)$, v broadcasts with probability:

$$p_v(t) = \frac{1}{1 + \lceil \frac{t - t_v - 1}{24 \ln n} \rceil}.$$

After these T rounds, the node can terminate. We now establish the (perhaps surprising) universality of this algorithm.

Theorem 10. *Let \mathcal{N} be a wireless network. The HARMONICAST algorithm solves broadcast in \mathcal{N} in $O(n \log^2 n)$ rounds.*

The above results follows immediately from the proof in [13], which assumes pessimistically (due to the difficulties of the dual graph model) that the message only makes progress in the network when it is broadcast alone in the entire network. Since the isolation graph for a wireless network defined with respect to the SINR equation is equivalent to $G(V, E[r_w])$, an immediate corollary of the above is that HARMONICAST algorithm solves weak connectivity broadcast in the ad hoc SINR model.

References

1. Alon, N., Bar-Noy, A., Linial, N., Peleg, D.: A Lower Bound for Radio Broadcast. *Journal of Computer and System Sciences* 43(2), 290–298 (1991)

2. Bar-Yehuda, R., Goldreich, O., Itai, A.: On the Time-Complexity of Broadcast in Multi-Hop Radio Networks: An Exponential Gap between Determinism and Randomization. *Journal of Computer and System Sciences* 45(1), 104–126 (1992)
3. Censor-Hillel, K., Gilbert, S., Kuhn, F., Lynch, N., Newport, C.: Structuring Unreliable Radio Networks. In: *Proc. ACM Symp. on Principles of Distributed Computing (PODC)*, pp. 79–88 (2011)
4. Czumaj, A., Rytter, W.: Broadcasting algorithms in radio networks with unknown topology. *Journal of Algorithms* 60, 115–143 (2006)
5. Daum, S., Gilbert, S., Kuhn, F., Newport, C.: Broadcast in the Ad Hoc SINR Model. Technical Report 274, University of Freiburg, Dept. of Computer Science (2013)
6. Goussevskaia, O., Wattenhofer, R., Halldorsson, M.M., Welzl, E.: Capacity of Arbitrary Wireless Networks. In: *Proc. IEEE Int. Conf. on Computer Communications (2009)*
7. Halldorsson, M.M., Mitra, P.: Towards Tight Bounds for Local Broadcasting. In: *Proc. Int. Workshop on the Foundations of Mobile Computing. ACM* (2012)
8. Halldorsson, M.M., Mitra, P.: Wireless Connectivity and Capacity. In: *Proc. ACM-SIAM Symp. on Discrete Algorithms, SODA* (2012)
9. Jurdzinski, T., Kowalski, D.R.: Distributed backbone structure for algorithms in the SINR model of wireless networks. In: Aguilera, M.K. (ed.) *DISC 2012. LNCS*, vol. 7611, pp. 106–120. Springer, Heidelberg (2012)
10. Jurdzinski, T., Kowalski, D.R., Stachowiak, G.: Distributed deterministic broadcasting in uniform-power ad hoc wireless networks. In: Gąsieniec, L., Wolter, F. (eds.) *FCT 2013. LNCS*, vol. 8070, pp. 195–209. Springer, Heidelberg (2013)
11. Kesselheim, T.: A Constant-Factor Approximation for Wireless Capacity Maximization with Power Control in the SINR Model. In: *Proc. ACM-SIAM Symp. on Discrete Algorithms, SODA* (2011)
12. Kowalski, D.R., Pelc, A.: Broadcasting in Undirected Ad Hoc Radio Networks. *Distributed Computing* 18(1), 43–57 (2005)
13. Kuhn, F., Lynch, N., Newport, C., Oshman, R., Richa, A.: Broadcasting in Radio Networks with Unreliable Communication. In: *Proc. ACM Symp. on Principles of Distributed Computing, PODC* (2010)
14. Kushilevitz, E., Mansour, Y.: An $\Omega(D \log(N/D))$ Lower Bound for Broadcast in Radio Networks. *SIAM Journal on Computing* 27(3), 702–712 (1998)
15. Moscibroda, T.: The Worst-Case Capacity of Wireless Sensor Networks. In: *Proc. ACM/IEEE Int. Conf. on Information Processing in Sensor Networks, IPSN* (2007)
16. Moscibroda, T., Wattenhofer, R.: Maximal Independent Sets In Radio Networks. In: *Proc. ACM Symp. on Principles of Distributed Computing (PODC)*, pp. 148–157 (2005)
17. Moscibroda, T., Wattenhofer, R.: The Complexity of Connectivity in Wireless Networks. In: *Proc. IEEE Int. Conf. on Computer Communications* (2006)
18. Newport, C.: Brief Announcement: A Shorter and Stronger Proof of an $\Omega(D \log(n/D))$ Lower Bound on Broadcast in Radio Networks. In: *Proc. ACM Symp. on Principles of Distributed Computing, PODC* (2013)
19. Scheideleer, C., Richa, A., Santi, P.: An $O(\log n)$ Dominating Set Protocol for Wireless Ad-Hoc Networks under the Physical Interference Model. In: *Proc. ACM Int. Symp. on Mobile Ad Hoc Networking and Computing* (2008)
20. Schneider, J., Wattenhofer, R.: A Log-Star Distributed Maximal Independent Set Algorithm for Growth-Bounded Graphs. In: *Proc. ACM Symp. on Principles of Distributed Computing (PODC)*, pp. 35–44 (2008)
21. Yu, D., Hua, Q.-S., Wang, Y., Tan, H., Lau, F.C.M.: Distributed multiple-message broadcast in wireless ad-hoc networks under the SINR model. In: Even, G., Halldórsson, M.M. (eds.) *SIROCCO 2012. LNCS*, vol. 7355, pp. 111–122. Springer, Heidelberg (2012)
22. Yu, D., Hua, Q.-S., Wang, Y., Lau, F.C.M.: An $O(\log n)$ Distributed Approximation Algorithm for Local Broadcasting in Unstructured Wireless Networks. In: *Proc. IEEE Int. Conf. on Distributed Computing in Sensor Systems (DCOSS)*, pp. 132–139 (2012)