Broadcast in the Ad Hoc SINR Model

Sebastian Daum
Dept. of Comp. Science
U. of Freiburg, Germany
sdaum@cs.uni-freiburg.de

Seth Gilbert
Dept. of Comp. Science
National University of Singapore
seth.gilbert@comp.nus.edu.sg

Fabian Kuhn
Dept. of Comp. Science
U. of Freiburg, Germany
kuhn@cs.uni-freiburg.de

Calvin Newport
Dept. of Comp. Science
Georgetown University, USA
cnewport@cs.georgetown.edu

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 [5, 7–10, 14, 16, 18, 20]. 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.

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 [6]. 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., α_{min}, α_{max}, β_{min}, β_{max}, 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 at all, while others use significantly different granularities. To produce algorithms

[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 [9]. 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.

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*) *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 high probability,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, 11], 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

[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.

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 [12] 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.

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., [5, 7, 10, 14, 16]. 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. [18] show how to solve strong connectivity broadcast in $O(D + \log n)$ rounds, while Yu et al. [20] present a $O(D + \log^2 n)$ round solution. In the deterministic setting, Jurdzinski et al. [8] describe a $O(\Delta \operatorname{polylog}(n) + D)$ solution, which they recently improved to $O(D \log^2 n)$ (under different assumptions) [9]. 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 [18] 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, [18] makes use of tunable collision detection, [20] allows the algorithm to specify the transmission power level as a function of the network parameters, [8] adds an additional model restriction that forbids communication over weak links,[3] and [9] 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 $O(\Delta \log n)$ [6, 21]. Of these results, the algorithm in [6] by Halldorsson et al. 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

[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.

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 $\Theta\big((D + \log n) \log \frac{n}{D}\big)$ rounds, if nodes start asynchronously (like in this paper) [1, 2, 4, 11, 13, 17]. 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]

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

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

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.[5]

[4] The result of [15] 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.

[5] Notice, for $\delta = 2$ the model strictly generalizes the Euclidean plane. We prefer this general notion of distance over

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 come 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 $SINR(u, v, I)$ is therefore undefined, we determine receive behavior with the following rule: v receives 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 & 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 $r_{\mathrm{w}} := \left(\frac{P}{\beta N}\right)^{1/\alpha}$ be the maximum distance at which any two nodes can communicate. Let $r_{\mathrm{s}} := \frac{r_{\mathrm{w}}}{1+\rho}$, 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* $G_w = (V, E[r_{\mathrm{w}}])$ and the *strong connectivity graph* $G_s = (V, E[r_{\mathrm{s}}])$. The values $R_w = \max_{\{u,v\},\{x,y\} \in E[r_{\mathrm{w}}]} \left\{\frac{d(u,v)}{d(x,y)}\right\}$ and $R_s = \max_{\{u,v\},\{x,y\} \in E[r_{\mathrm{s}}]} \left\{\frac{d(u,v)}{d(x,y)}\right\}$ 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., $\forall v \in V\colon\ \exists s \in S\colon\ v \in N(s) \cup \{s\}$, where $N(s)$ is the set of neighbors of 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 *with high probability* (*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 that a given network is

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).

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 is round 0 for the source).

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 (he *hits* T). 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 (we refer to Appendix A for the proof):

Theorem 3.1. *Let $\mathcal{P}$ be a player that solves the (x, y)-hitting game in $f(x, y)$ rounds, in expectation. It follows that $f(x, y) = \Omega\left(\frac{y}{x}\right)$.*

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 4.1. *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, i.e., $R_s = \Omega(\log n)$,[6] and in the proof of Theorem 4.1 we show how to solve the problem for slowly growing or even constant R_s.

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*

[6] In fact, it is sufficient to assume $\log^{\alpha_{\max}} R_s = \Omega(\log^\star n)$.

while inactive nodes remain silent. 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.

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:
(1) u receives a message from v with probability at least μ and
(2) 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:

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

An ε-close approximation $\tilde{H}_p^\mu[S]$ of $H_p^\mu[S]$ can be computed in time $O\big(\frac{\log n}{\varepsilon^2\mu}\big)$ 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 $\big(1-\frac{\varepsilon}{2}\big)\mu T$ 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.

Lemma 4.2. *W.h.p., the SINR-Induced Graph Computation protocol runs in $O\big(\frac{\log n}{\varepsilon^2\mu}\big)$ rounds and returns a graph $\tilde{H}_p^\mu[S]$ that is an ε-close approximation of $H_p^\mu[S]$.*

Proof. Let $T = \frac{c}{\varepsilon^2\mu}\log n$ for some yet to be defined c. Let X_i be 1, if node u receives a message in round i from node v, and thus $X = \sum_{i=1}^{T} X_i$ is the random variable that counts the number of messages u receives in total from v. Let the probability of a successful transmission be μ'.

First let $\mu' \geq \mu$. Using Chernoff we find that

$$\mathbb{P}\left(X \leq \left(1-\frac{\varepsilon}{2}\right)\mu T\right) \leq \mathbb{P}\left(X \leq \left(1-\frac{\varepsilon}{2}\right)\mu' T\right) \leq e^{-\frac{\varepsilon^2}{12}\mu' T} \leq n^{-\frac{c}{12}}.$$

For $\mu' \leq (1-\varepsilon)\mu$ a similar argumentation gives us $\mathbb{P}\big(X \geq \big(1-\frac{\varepsilon}{2}\big)\mu T\big) \leq n^{-\frac{c}{12}}$.

Choosing c a large enough constant and using a union bound over all nodes, w.h.p., all nodes v that transmit with probability at least μ, are connected to u in $\tilde{H}_p^\mu[S]$, while no nodes with transmission probability less than $(1-\varepsilon)\mu$ are, proving $E\big[H_p^\mu[S]\big] \subseteq E\big[\tilde{H}_p^\mu[S]\big] \subseteq E\big[H_p^{(1-\varepsilon)\mu}[S]\big]$. In the second part of the protocol, for establishing a bidirectional link between two nodes, they must only successfully transmit a single message to each other, which is clearly more likely than transmitting $\big(1-\frac{\varepsilon}{2}\big)\mu T$ messages, thus, it can be included in the union bound above. □

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 $O\big(\frac{\log n}{\mu}\big) = O(\log n)$. 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 4.3. $\forall p \in \big(0, \frac{1}{2}\big]$, $\exists \mu \in (0,p)$ *such that: Let $d_{\min} \leq r_{\mathrm{s}}$ 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 $d(u,v) \leq \min\{2d_{\min}, r_{\mathrm{s}}\}$.*

Proof. We restrict our attention to the case $d_{\min} \leq r_{\mathrm{s}}/2$. If the minimum distance is between $r_{\mathrm{s}}/2$ 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 $p(1-p)^{O(k_0^\delta)} = \Omega(p)$, node v is the only node transmitting among all the nodes within distance $k_0 d_{\min}$ from node u. Further, assuming that all nodes at distance greater than $k_0 d_{\min}$ transmit, the interference $I(u)$ at u can be bounded from above by

$$I(u) \leq \sum_{w:\, d(u,w) \geq k_0 d_{\min}} \frac{P}{d(u,w)^\alpha} \leq \sum_{k=k_0}^{\infty} \sum_{w:\, 1 \leq \frac{d(u,w)}{k d_{\min}} < 1+\frac{1}{k}} \frac{P}{d(u,w)^\alpha} \overset{(*)}{\leq} \sum_{k=k_0}^{\infty} \frac{P}{k^\alpha d_{\min}^\alpha} O\big(\delta k^{\delta-1}\big)$$
$$= \frac{P}{d_{\min}^\alpha} O\left(\delta \int_{k_0}^{\infty} k^{-(1+\alpha_{\min}-\delta)} \mathrm{d}k\right) = \frac{P}{d_{\min}^\alpha} O\left(\delta \frac{k_0^{\delta-\alpha_{\min}}}{\alpha_{\min}-\delta}\right)$$

Step $(*)$ stems from bounding $|\{w : k d_{\min} \leq d(u,w) < (k+1) d_{\min}\}|$, the maximum number of nodes within a ring of diameter $d_{\min}$ at distance $k d_{\min}$. If we define the function κ so as to replace the O-term with $\kappa(k_0) = \kappa(k_0, \alpha_{\min}, \delta) > 0$, which decreases polynomially in k_0, then we get for $SINR(u, v, I)$, where I is the set of all nodes with distance greater than $k_0 d_{\min}$:

$$\frac{\frac{P}{d(u,v)^\alpha}}{N + \kappa(k_o)\frac{P}{d_{\min}^\alpha}} \geq \frac{\frac{P}{(2d_{\min})^\alpha}}{N + \kappa(k_0)\frac{P}{d_{\min}^\alpha}} \geq \frac{\frac{P}{r_{\mathrm{s}}^\alpha}}{\frac{P}{\beta r_{\mathrm{w}}^\alpha} + \kappa(k_0)\frac{2^\alpha P}{r_{\mathrm{s}}^\alpha}} = \frac{\beta}{\frac{1}{(1+\rho)^\alpha} + \kappa(k_0)\beta 2^\alpha} \geq \beta$$

The second inequality follows from $N = \frac{P}{\beta r_{\mathrm{w}}^\alpha}$ and from $d_{\min} \leq r_{\mathrm{s}}/2$. 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 $k_0 d_{\min}$ transmits, then node v can transmit to u with probability μ. □

In the above proof, μ depends on the unknown parameter β, so we use $\beta_{\max}$ 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 $\phi = 1, 2, \ldots, \Phi$, 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-inducted graph computation process

Algorithm 1 High-level pseudo-code for one epoch of STRONGCAST

Input: n, R_s, $\alpha_{\max}$, $\beta_{\max}$, ε, 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 φ = 1 to Φ do
    Compute SINR-induced graph H̃_p^μ[S_φ] within O(log n/(ε²μ)) rounds          ▷ Block 1
    for O(Q log n) rounds do                                                       ▷ Block 2
        Each round transmit M with probability p/Q
    Compute MIS S_{φ+1} on H̃_p^μ[S_φ] within O(log n/(ε²μ) log* n) rounds        ▷ Block 3
```

described in Section 4.1. We choose $\mu > 0$ appropriately as described in Lemma 4.3, while $\varepsilon, p \in (0, 1/2)$ can be chosen freely.[7]

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 $O(\log^\star n)$ MIS algorithm from [19] on this simulated network. (This algorithm requires a growth-bounded 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 transmit 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,w.h.p., the neighborhood dissemination problem for S. 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 4.4. *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 4.3, for a sufficiently small constant μ, w.h.p., $\tilde{H}_p^\mu[S_\phi]$ 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 $\tilde{H}_p^\mu[S_\phi]$,

[7] By Lemma 4.3, μ depends on p; thus p could be chosen to maximize μ.

nodes in $S_{\phi+1}$ are at distance more than $2d_\phi$ and therefore using the induction hypothesis, we get $d_{\phi+1} > 2d_\phi \geq 2^\phi d_{\min}$. □

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 4.5. $\forall p \in \left(0, \frac{1}{2}\right],\ \exists \hat{Q}, \gamma = \Theta(1)$*, such that for all* $Q \geq \hat{Q}$ *the following holds. Consider round* r *in phase* ϕ *where each node in* S_ϕ *transmits with probability* $\frac{p}{Q}$ *the broadcast message* M*. 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* $d(u,v) \leq \left(1+\frac{\rho}{2}\right)r_{\mathrm{s}}$ *and* $d_u \geq \gamma Q^{-1/\alpha} \cdot d(u,v)$*, then node* v *receives* M *in round* r *with probability* $\Theta\left(\frac{1}{Q}\right)$*.*

Proof. 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 $\Theta(1/Q)$ in round r of block 2, in which it transmits with probability p/Q. We start with some notations 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 superscript ^{p} indicates the broadcasting probability of nodes in round r'). 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. For round r, in which nodes use the broadcasting probability p/Q, we use the superscript $^{p/Q}$. Finally, for a set $A \subseteq S_\phi$, we define $\bar{A} := S_\phi \setminus A$.

For any $w \in S_\phi$, the triangle inequality implies that $d(u,w) \leq d(u,v) + d(v,w) \leq 2d(v,w)$. 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 $d(u,u') = d_u$ and therefore both u and u' are in A. 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:
(a) $\frac{P}{d_u^\alpha\left(N+X_{\bar{A}}^p(u)\right)} \geq \frac{P}{d_u^\alpha(N+X^p(u))} \geq \beta$ and
(b) u' transmits and u listens (event $R^{u,u'}$).
Thus we have:

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

Using Lemma B.1, we can therefore bound $X_{\bar{A}}^p(u)$ as

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

For the sake of contradiction, assume that $I_{\bar{A}}(u) > c \cdot \frac{P}{p\beta d_u^\alpha}$ for $c = \max\left\{2, \frac{16\beta}{2^\alpha} \cdot \ln \frac{p(1-p)}{(1-\varepsilon)\mu}\right\}$. Combining (2) and (3), we obtain

$$\frac{(1-\varepsilon)\mu}{p(1-p)} \overset{(2)}{\leq} \mathbb{P}\left(X_{\bar{A}}^p(u) \leq \frac{P}{\beta d_u^\alpha}\right) \leq \mathbb{P}\left(X_{\bar{A}}^p(u) \leq \frac{cP}{2\beta d_u^\alpha}\right) < \mathbb{P}\left(X_{\bar{A}}^p(u) \leq \frac{pI_{\bar{A}}(u)}{2}\right) \leq e^{-\frac{2^\alpha c}{16\beta}},$$

which is a contradiction to the definition of c. We therefore have $I_{\bar{A}}(u) \leq c \cdot \frac{P}{p\beta d_u^\alpha}$.

We now have all tools to show that v receives a message from u in round r, with broadcasting probabilities p/Q. From the fact that the link $\{u,u'\} \in E[\tilde{H}_p^\mu[S_\phi]]$ is reliable, we have seen that with probability at least $(1-\varepsilon)\mu$ fewer than $\frac{2^\alpha}{\beta}$ nodes in A send in round r'. But then in round r with the same probability no more than $\frac{2^\alpha}{\beta Q}$ send within A. Markov's inequality shows that $\mathbb{P}\left(X_{\bar{A}}^{p/Q}(v) < 2\frac{p}{Q} I_{\bar{A}}(v)\right) \geq 1/2$. Finally, u sends with probability p/Q. All those events are independent, thus all of them happen with probability at least $\frac{(1-\varepsilon)\mu p}{2Q} = \Theta(1/Q)$. Let us assume that this is the case. To see that v indeed gets u's message under those conditions, we check whether $SINR(u,v,I) = \frac{Pd(u,v)^{-\alpha}}{N+X_{\bar{A}}^{p/Q}(v)+X_A^{p/Q}} \geq \beta$:

$$
\begin{aligned}
&\beta d(u,v)^\alpha (N + X_{\bar{A}}^{p/Q}(v) + X_A^{p/Q}) \\
&\overset{(*)}{\leq} \quad \beta d(u,v)^\alpha N + 2^{\alpha+1} c_\beta P \frac{d(u,v)^\alpha}{d_u^\alpha} + \beta \sum_{w \in A, w \text{ sends}} P \frac{d(u,v)^\alpha}{Q d(w,v)^\alpha} \\
&\overset{d(u,v)^\alpha \leq Q \frac{d_u^\alpha}{\gamma^\alpha}}{\leq} \quad \left(1 + \frac{\rho}{2}\right)^\alpha r_{\mathrm{s}}^\alpha N \beta + \frac{2^{\alpha+1} c_\beta}{\gamma^\alpha} P + \frac{2^\alpha}{Q} P \\
&\overset{(1+\rho)^\alpha \geq \left(1+\frac{\rho}{2}\right)^\alpha + \alpha \frac{\rho}{2}}{\leq} \quad \left(1 - \frac{\alpha\rho}{2(1+\rho)^\alpha}\right)(1+\rho)^\alpha r_{\mathrm{s}}^\alpha N \beta + \frac{2^{\alpha+1} c_\beta}{\gamma^\alpha} P + \frac{2^\alpha}{\hat{Q}} P \\
&\overset{P = N\beta(1+\rho)^\alpha r_{\mathrm{s}}^\alpha}{\leq} \quad P + P\left(\frac{2^{\alpha+1} c_\beta}{\gamma^\alpha} + \frac{2^\alpha}{\hat{Q}} - \frac{\alpha\rho}{2(1+\rho)^\alpha}\right) \overset{(**)}{\leq} P
\end{aligned}
$$

Inequality $(*)$ holds due to the assumption that $X_{\bar{A}}^{p/Q}(v) < 2I_{\bar{A}}(v)p/Q$ and (2). Inequality $(**)$ holds for properly chosen $\gamma = \Theta(1)$ and $\hat{Q} = \Theta(2^\alpha) = O(\log^{\alpha_{\max}} R_s)$. □

4.3 Proof of Theorem 4.1

Proof. We show that neighborhood dissemination needs $O\left((\log^{\alpha_{\max}+1} R_s)(\log n)\right)$ rounds and that, w.h.p., it correctly passes the message by one hop with each epoch. With the diameter of the strong connectivity graph being D, it suffices to repeat neighborhood dissemination D times in order to broadcast the message to all nodes in the network.

Let $S_1 = S$ be the set of nodes that have the message. We let p and ε be arbitrarily chosen constants in $(0, 1/2)$ and retrieve value μ from Lemma 4.3 and values γ and $\hat{Q}$ from Lemma 4.5. We will fix $Q \geq \hat{Q}$ at the end of the proof.

Lemma 4.2 ensures that block 1 of each phase ϕ is done within $O\left(\frac{\log n}{\varepsilon^2 \mu}\right)$ rounds and that, w.h.p., it returns an ε-close approximation $\tilde{H}_p^\mu[S_\phi]$ of $H_p^\mu[S_\phi]$. The running time of block 2 is fixed to $T_{\text{phase}} = \Theta(Q \log n)$. Because the maximum degree in $\tilde{H}_p^\mu[S_\phi]$ is upper bounded by a constant, in block 3 of each phase, we can simulate the MIS algorithm from [19], which runs in time $O(\log^\star n)$ in a network with reliable links. To simulate this algorithm, we use $O(\log n)$ rounds to guarantee, w.h.p., a transmission over a $(1-\varepsilon)\mu$-reliable link in $\tilde{H}_p^\mu[S_\phi]$. It therefore takes $O(\log n \log^\star n)$ rounds to calculate the active nodes of phase $\phi + 1$. In total, the running time of block 2 dominates the running times of block 1 and 3.[8]

Due to Lemma 4.4, w.h.p., there can be no more than $\log R_s$ phases. Since the length of each phase is $O(Q \log n) = O(\log^{\alpha_{\max}} R_s)(\log n))$, this yields the specified running time. It remains

[8]Recall that we assume R_s to be at least logarithmic in n. Being a bit more careful, it would also be possible to bring the time complexity of computing an MIS to $O(\log n)$. In order to run a $O(\log^\star n)$-round distributed algorithm, it is sufficient if each node u collects the initial states of all nodes in u's $O(\log^\star n)$-neighborhood. As we have a bounded degree graph and we can send a message over each edge with constant probability in each round, this is indeed possible in $O(\log n)$ rounds.

to be shown that every neighbor of S receives the message by the end of the neighborhood dissemination protocol.

Let v be a node in $N(S) \setminus S$. We inductively construct a series of nodes $u_\phi \in S_\phi$, where u_1 is the closest node to v in $S_1 = S$. We denote by d_{u_ϕ} the distance between u_ϕ and its farthest neighbor in S_ϕ. We prove by induction the following claim:

Claim 4.6. *W.h.p., either u_ϕ reaches v in phase ϕ, or $d(u_{\phi+1}, v) \leq r_{\mathrm{s}} \left(1 + \phi \frac{\rho}{2 \log R_s}\right)$.*

Clearly, $d(u_1, v) \leq r_{\mathrm{s}}$.

Let ϕ be any phase. If $d_{u_\phi} \geq \gamma Q^{-1/\alpha} d(u_\phi, v)$, then we can apply Lemma 4.5 and we are done, because u_ϕ sends for $T_{\text{phase}} = \Theta(Q \log n)$ rounds in block 2, i.e., if we choose the constant in T_{phase} large enough, then w.h.p. it reaches v during the execution of block 2. So let this not be the case and let $u_{\phi+1}$ be the closest neighbor to v in $S_{\phi+1}$. Due to the MIS construction we have $d(u_{\phi+1}, v) \leq d(u_\phi, v) + d_{u_\phi}$, and therefore

$$d(u_{\phi+1}, v) \leq \left(1 + \frac{\gamma}{Q^{1/\alpha}}\right) d(u_\phi, v) \leq r_{\mathrm{s}} \left(1 + \phi \frac{\rho}{2 \log R_s} + \frac{2\gamma}{Q^{1/\alpha}}\right) \leq r_{\mathrm{s}} \left(1 + \frac{(\phi+1)\rho}{2 \log R_s}\right)$$

The last inequality holds for properly chosen $Q = \Theta(\log^{\alpha_{\max}} R_s)$, $Q \geq \hat{Q}$, proving the claim.

□

5 Lower Bounds for Strong Connectivity Broadcast

In this section, we present lower bounds for strong connectivity broadcast.

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 at the same distance. Here we prove this assumption makes efficient broadcast impossible.

Theorem 5.1. *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. Assume $\mathcal{A}$ guarantees to solve broadcast in $f(n)$ rounds in any $O(1)$-broadcastable network of size n. We show that this same algorithm solves $(x, n-2)$-hitting in $O(f(n))$ rounds, for some $x = O(1)$. Our theorem statement follows directly from this observation and the result of Theorem 3.1.

In more detail, we construct a player $\mathcal{P}_{\mathcal{A}}$ for the $(\lceil \rho + 1 \rceil, n-2)$-hitting game that operates by simulating an execution of $\mathcal{A}$ in an $\lceil 1 + \rho \rceil$-broadcastable network. We associate each $u_i \in$

$V = \{u_1, \ldots, u_n\}$ with value i in the guessing game and $u_s = u_{n-1}$ with the source node. In the simulation, we use the IDs of broadcasting nodes to help generate guesses.

Fix some arbitrary small $\epsilon \in (0, r_{\mathrm{s}})$ and let $k := \lceil 1 + \rho \rceil$, $l := r_{\mathrm{w}} + \varepsilon$ and $d := l/(k+1)$. In our simulation, we arrange the n nodes in our lower bound network onto one of $k+2$ positions, $p_0, p_1, \ldots, p_{k+1}$, arranged on a line of length l with uniform spacing d. Note that this network is $k+1$-broadcastable with respect to strong connectivity, because these positions have distance $d = l/(k+1) \leq (r_{\mathrm{w}} + r_{\mathrm{s}})(2+\rho) \leq r_{\mathrm{s}}$, and hence there exists a $k+1$-round schedule to broadcast the message along strong links.

The player simulates $\mathcal{A}$ in a network where the broadcast source $u_s := u_{n-1}$ is at position p_0, u_n is at position p_{k+1}, and the k nodes corresponding to the k targets (i.e., the set $\{u_i \mid i \in T\}$) are arranged on positions p_1 to p_k, and the rest of the nodes are arranged at position p_0 as well. See Figure 1. Of course, the player does not know the targets in advance, so, in its simulation, it does not know the positions of all the nodes. We show, however, that its simulation remains consistent until it wins the hitting game.

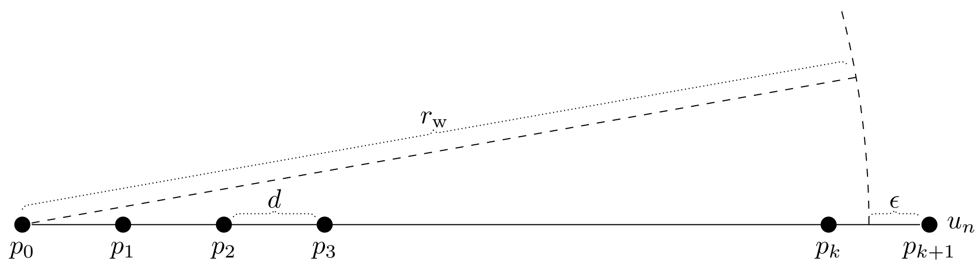

Figure 1: Lower bound for strong connectivity broadcast in the compact ad hoc SINR model

To solve the hitting game, $\mathcal{P}_{\mathcal{A}}$ does the following. Without loss of generality we can assume that u_s broadcasts in round 0, informing all other nodes except u_n, which is too far to receive the message. I.e., in round 1, all nodes but u_n have the message. For a round r of algorithm $\mathcal{A}$ let X_r be the set of broadcasting nodes. $\mathcal{P}_{\mathcal{A}}$'s proposal sets for the hitting game and the following simulation of receive behavior among the nodes depends on X_r as follows – we provide explanations afterward.

(1) If X_r is empty, it creates no proposals and simulates all nodes receiving nothing, and moves on to round $r+1$.

(2) If $X_r = \{u_i\}$, it proposes i. If it does not win, it simulates all nodes except u_n and u_i itself as receiving the message.

(3) If $1 < |X_r| \leq (k+1)^\alpha$, then it uses $|X_r|$ rounds of the hitting game to propose the value set $\{i \mid u_i \in X_r\}$, one by one, before continuing with round $r+1$ of the simulation of $\mathcal{A}$. If no guess wins, then it simulates no node receiving a message in round r.

(4) If $|X_r| > (k+1)^\alpha$, then $\mathcal{P}_\mathcal{A}$ again makes no proposals and simulates no node receiving a message.

We must establish that this simulation is consistent.
(1) & (2) are trivial.
(3): Note that $\mathcal{P}_\mathcal{A}$only continues the simulation if all proposals fail. Which is only the case if all broadcasters are in p_0. Because there are at least 2, there is no node in the network at which the signal can be stronger than the noise. Thus, the simulation of receive behavior is valid.
(4): Note that any set larger than $(k+1)^\alpha$ includes at least two nodes in p_0. Therefore, no node in p_0 can receive a message from any other node, due to the model definition of compact SINR networks. Any incoming signal for any node outside p_0 has signal strength at most $Pd^{-\alpha}$. The total interference for any node, however, is at least

$$|X_r| \cdot P\,l^{-\alpha} \geq P(k+1)^\alpha (d(k+1))^{-\alpha} = Pd^{-\alpha},$$

and therefore, no node can decode a message this round and $\mathcal{P}_\mathcal{A}$'s receive behavior is consistent.

We established, therefore, that until the player wins the hitting game, its simulation is consistent. To solve broadcast, it is a minimal requirement that a node in p_1 to p_k broadcasts in a round where in total no more than $(k+1)^\alpha$ nodes broadcast (otherwise, by our analysis above of case (4), u_n could not receive the message). In such a round, $\mathcal{P}_\mathcal{A}$applies the rules of case 3 and wins the hitting game. Therefore, if $\mathcal{A}$ solves broadcast in $f(n)$ rounds in our simulated network, then $\mathcal{P}_\mathcal{A}$ solves the hitting game in no more than $f(n)(k+1)^\alpha = O(f(n))$ rounds. Since $\rho = O(1)$, Theorem 3.1 states that the $(\lceil \rho + 1 \rceil, n-2)$-hitting game needs $\Omega(n)$ rounds to be solved, which concludes our proof. □

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 5.2. *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. We adopt the same approach as in Theorem 5.1. That is, we show how to implement $\mathcal{P}_\mathcal{A}$, a solution to the $(x, n-2)$-hitting game, for $x = O(1)$, that simulates $\mathcal{A}$ in a special line network to generate its guesses, solving the hitting game with a time complexity within a constant factor of the guarantees of $\mathcal{A}$ for solving broadcast. The linear bound on hitting from Theorem 3.1 then provides the needed linear bound on $\mathcal{A}$.

We begin with the same simulated network setup as in Theorem 5.1. Because we no longer assume the compact assumption we must modify this setup to prevent nodes from occupying the same location. In more detail, we take the nodes located at p_0 and spread them uniformly on a second line l_0 that passes through p_0 perpendicular to the line (p_0, p_{k+1}). We make the spacing between nodes on l_0 small enough such that all nodes on the line remain within distance r_w of p_k and within distance $d := (r_\mathrm{w} + \varepsilon)/(k+1) \leq r_\mathrm{s}$ of p_0. As before, we make $u_s = u_{n-1}$ the source and place it on p_0. See Figure 2.

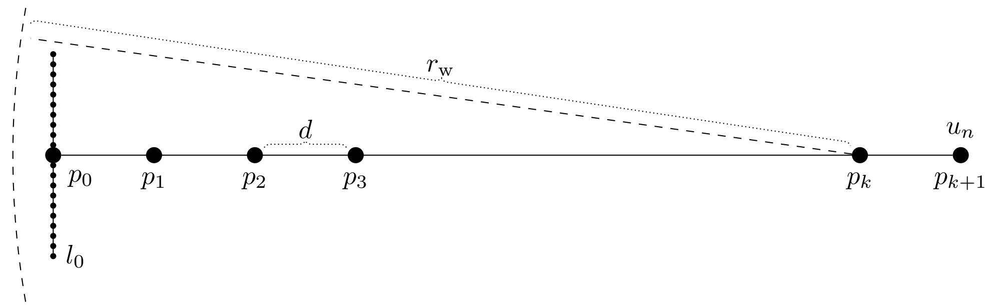

Figure 2: Lower bound for strong connectivity broadcast with back-off style algorithms

Notice, with the nodes near p_0 now no longer occupying the same location, correctly simulating receive behavior has become more complex (it is possible, for example, for multiple nodes assigned to the p_0-line to communicate concurrently). Fortunately, the assumption that $\mathcal{A}$ is back-off style frees us of the need for complex receive behavior simulation. Without loss of generality, u_s broadcasts alone during round 0 (no node can broadcast until u_s broadcasts). The player can simulate all nodes except u_n receiving this message, as all are within r_w distance of u_s. Moving forward, there is no need to simulate their receive behavior *at all* as $\mathcal{A}$ is back-off style and they will ignore all subsequent received messages.

We are left only to specify the guess behavior for the hitting game. In the proof of Theorem 5.1, we noted that if more than $(k+1)^\alpha$ nodes broadcast, then u_n cannot receive the message (as the weakest interference is at least $(k+1)^{-\alpha}$ times the strongest signal). A similar result holds in our modified version of this network. We simply need to increase the threshold slightly to $(k+2)^\alpha$, as all nodes on line l_0 are within distance d of p_0. With this adjusted threshold in mind, we use the same proposal rule for the hitting game: if no more than $(k+2)^\alpha$ nodes broadcast, guess each one by one, then continue the simulation only if no guess wins the game. The rest of the proof follows as in Theorem 5.1. □

6 Weak Connectivity Broadcast

Weak connectivity broadcast seems 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.

In this section, we formalize this difficulty by proving the existence of 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. We reanalyze an algorithm we originally presented in [12], 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 bound in this model, but that it solves broadcast in $O(n \log^2 n)$ rounds essentially in *every reasonable wireless model* (a notion we formalize below).

6.1 Lower Bound

Our lower bound leverages 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 3.1. 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 broadcasts alone. We prove that identifying this bridge node is at least as hard as solving the $(1, n-1)$-hitting game, which requires $\Omega(n)$ rounds.

Theorem 6.1. *Let $\mathcal{A}$ be a 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. Assume $\mathcal{A}$ guarantees to solve broadcast in $f(n)$ rounds in any 2-broadcastable network of size n with at least high probability. We will show that this same algorithm solves $(1, n-1)$-hitting in $O(f(n))$ rounds. Our theorem statement will follow directly from this observation and the result of Theorem 3.1.

In more detail, we construct a player $\mathcal{P}_{\mathcal{A}}$ for the $(1, n-1)$-hitting game that operates by simulating $\mathcal{A}$ on the rotating lollipop network (see Figure 3), which we define as follows: this network consists of nodes $\{u_1, \ldots, u_{n-1}\}$ arranged uniformly on a circle with diameter $\leq r_{\mathrm{w}}$, and node u_n placed at distance r_{w} from u_x on the line defined by the center of the circle and the location of u_x, where x is the target in the current instance of the hitting game (i.e., $T = \{x\}$). Of course, $\mathcal{P}_{\mathcal{A}}$ does not know the identity of u_x, but we will show its simulation remains valid until it wins the hitting game.

In particular, $\mathcal{P}_{\mathcal{A}}$ simulates $\mathcal{A}$ in the rotating lollipop network with u_1 as the broadcast source. In each simulated round r, if exactly one node u_i broadcasts, the player guesses i in the hitting game. If i is not the target, it simulates all nodes in the circle receiving the message and u_n receiving nothing. Note that after the first round, if 1 is not the target, all nodes but u_n have the message. If more than one node broadcasts, then $\mathcal{P}_{\mathcal{A}}$ makes no guess and simulates the receive behavior in the network normally, with u_n not being able to receive the message. This simulation remains valid until $\mathcal{P}_{\mathcal{A}}$ wins the hitting game, because the receive behavior among the nodes in the

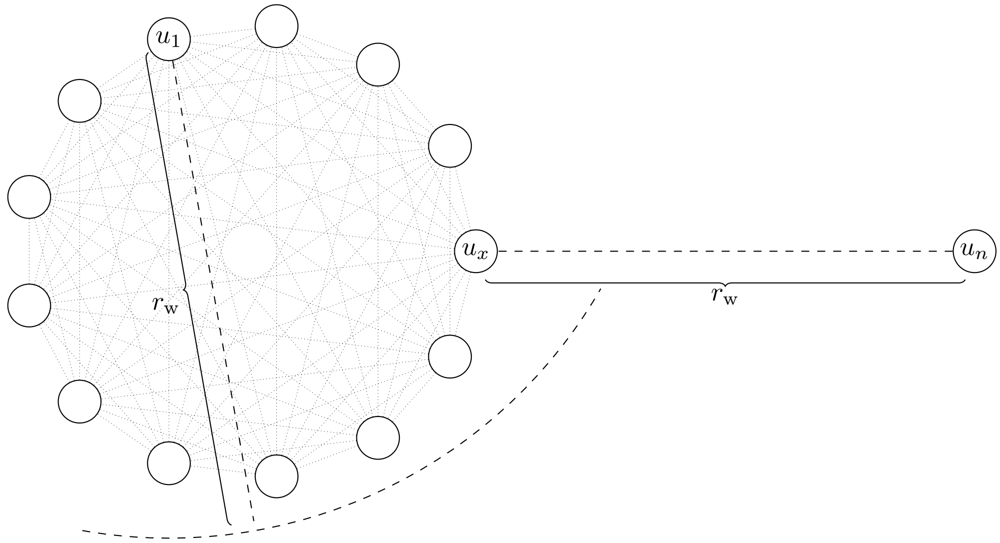

Figure 3: Rotating lollipop network

circle is the same, regardless of which node u_x is the bridge, and u_n can only receive a message if only the bridge u_x broadcasts, in which case the player wins the hitting game before needing to simulate the corresponding receive behavior.

To complete the bound, we note that if $f(n) = o(n)$ then $\mathcal{P}_{\mathcal{A}}$ solves $(1, n-1)$-hitting in $o(n)$ rounds, which contradicts Theorem 3.1. □

6.2 Upper Bound

In [12], we described a simple anti-social broadcast algorithm that solved 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 weakest definition of broadcast in $O(n \log^2 n)$ rounds in every "reasonable" wireless network model. The fact that it does so in the ad hoc SINR model is an immediate corollary.

A Universal Definition of Broadcast. We begin by formalizing our notion of a wireless network model. In particular, we can define a *wireless network* $\mathcal{N} = (V, M, R)$ as a 3-tuple, where V is the non-empty set of nodes in the network, M is the non-empty set of possible messages (that does not include the special *no message* indicator, $\perp$), and the receive function $R : (M \cup \{\perp\})^{|V|} \rightarrow (M \cup \{\perp\})^{|V|}$, that maps a transmission pattern to a receive pattern (where a *pattern* is an assignment of messages—or $\perp$ to indicate no message—to the nodes in the network). A wireless network, in other words, is the formal combination of a particular set of nodes and network topology, with a particular network model that defines the communication rules.

For a given wireless network $\mathcal{N} = (V, M, R)$, let $G_I(\mathcal{N}) = (V, E)$ be the directed *isolation graph* where $(u, v) \in E$ if and only if v would receive message $m \in M$ if u broadcast this message alone in the network. We can now define the broadcast problem with respect to an arbitrary wireless network as follows: We say algorithm $\mathcal{A}$ solves broadcast in wireless network $\mathcal{N}$ if it guarantees, with high probability (defined w.r.t. $|V|$), to propagate the broadcast message to all nodes in the connected component in $G_I(\mathcal{N})$ that contains the broadcast source.

Harmonic Broadcast. We next describe the broadcast algorithm HARMONICCAST, first presented in [12]. 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 6.2. *Let $\mathcal{N}$ be a wireless network. The* HARMONICCAST *algorithm solves broadcast in $\mathcal{N}$ in $O(n \log^2 n)$ rounds.*

Proof. The about results follows immediately from the proof in [12], which assumes pessimistically (due to the difficulties of the dual graph model) that the message only makes progress in the network when it is broadcasting alone in the entire network. □

Since the isolation graph for a wireless network defined with respect to the SINR equation is equivalent to $G(V, E[r_{\mathrm{w}}])$, an immediate corollary of the above is that HARMONICCAST algorithm solves weak connectivity broadcast in the ad hoc SINR model.

Corollary 6.3. *The harmonic broadcast algorithm solves weak connectivity broadcast in the ad hoc SINR model in $O(n \log^2 n)$ rounds.*

References

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

[2] R. Bar-Yehuda, O. Goldreich, and A. Itai. 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] K. Censor-Hillel, S. Gilbert, F. Kuhn, N. Lynch, and C. Newport. Structuring Unreliable Radio Networks. In *Proc. ACM Symp. on Principles of Distributed Computing (PODC)*, pages 79–88, 2011.

[4] A. Czumaj and W. Rytter. Broadcasting algorithms in radio networks with unknown topology. *Journal of Algorithms*, 60:115–143, 2006.

[5] O. Goussevskaia, R. Wattenhofer, M. M. Halldorsson, and E. Welzl. Capacity of Arbitrary Wireless Networks. In *Proc. IEEE Int. Conf. on Computer Communications*, 2009.

[6] M. M. Halldorsson and P. Mitra. Towards Tight Bounds for Local Broadcasting. In *Proc.Int. Workshop on the Foundations of Mobile Computing*. ACM, 2012.

[7] M. M. Halldorsson and P. Mitra. Wireless Connectivity and Capacity. In *Proc. ACM-SIAM Symp. on Discrete Algorithms (SODA)*, 2012.

[8] T. Jurdzinski and D. R. Kowalski. Distributed Backbone Structure for Algorithms in the SINR Model of Wireless Networks. In *Proc. Int. Symp. on Distributed Computing (DISC)*, 2012.

[9] T. Jurdzinski, D. R. Kowalski, and G. Stachowiak. Distributed Deterministic Broadcasting in Uniform-Power Ad Hoc Wireless Networks. *Proc. Int. Symp. on Fundamentals of Computation Theory (FCT)*, 2013.

[10] T. Kesselheim. 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.

[11] D. R. Kowalski and A. Pelc. Broadcasting in Undirected Ad Hoc Radio Networks. *Distributed Computing*, 18(1):43–57, 2005.

[12] F. Kuhn, N. Lynch, C. Newport, R. Oshman, and A. Richa. Broadcasting in Radio Networks with Unreliable Communication. In *Proc. ACM Symp. on Principles of Distributed Computing (PODC)*, 2010.

[13] E. Kushilevitz and Y. Mansour. An $\Omega(D \log(N/D))$ Lower Bound for Broadcast in Radio Networks. *SIAM Journal on Computing*, 27(3):702–712, 1998.

[14] T. Moscibroda. The Worst-Case Capacity of Wireless Sensor Networks. In *Proc. ACM/IEEE Int. Conf. on Information Processing in Sensor Networks (IPSN)*, 2007.

[15] T. Moscibroda and R. Wattenhofer. Maximal Independent Sets In Radio Networks. In *Proc. ACM Symp. on Principles of Distributed Computing (PODC)*, pages 148–157, 2005.

[16] T. Moscibroda and R. Wattenhofer. The Complexity of Connectivity in Wireless Networks. In *Proc. IEEE Int. Conf. on Computer Communications*, 2006.

[17] C. Newport. 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.

[18] C. Scheideler, A. Richa, and P. Santi. 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.

[19] J. Schneider and R. Wattenhofer. A Log-Star Distributed Maximal Independent Set Algorithm for Growth-Bounded Graphs. In *Proc. ACM Symp. on Principles of Distributed Computing (PODC)*, pages 35–44, 2008.

[20] D. Yu, Q. Hua, Y. Wang, H. Tan, and F. Lau. Distributed Multiple-Message Broadcast in Wireless Ad-Hoc Networks Under the SINR Model. In *Proc. Int. Coll. on Structural Information and Communication Complexity*, 2012.

[21] D. Yu, Q.-S. Hua, Y. Wang, and F. C. Lau. 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)*, pages 132–139, 2012.

A Proof of Theorem 3.1

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

Proof. We let the referee choose the set T uniformly at random. In round $i+1$, the player can only make use of the knowledge that its first i guesses were not correct. Thus the probability for it to hit T in round $i+1$ – conditioned that it did not hit T before – is at most $x/(y-i)$. Let Z be the random variable that counts the number of rounds until the player wins the game and let $k \in \{1, \ldots, y/(x+1)\}$.[9] Then we get

$$\mathbb{P}(Z \geq k) \geq \left(1 - \frac{x}{y}\right) \cdot \ldots \cdot \left(1 - \frac{x}{y-(k-2)}\right) = \prod_{i=0}^{k-2} \left(1 - \frac{x}{y-i}\right) \overset{i \leq \frac{y}{x+1}}{\geq} \left(1 - \frac{x+1}{y}\right)^{k-1}.$$

We can now bound $f(x,y)$ as follows:

$$f(x,y) \geq \mathbb{E}[Z] \geq \sum_{k=1}^{\frac{y}{x+1}} \mathbb{P}(Z \geq k) \geq \frac{1 - \left(1 - \frac{x+1}{y}\right)^{\frac{y}{x+1}}}{1 - \left(1 - \frac{x+1}{y}\right)} \geq \frac{1 - e^{-1}}{\frac{x+1}{y}} = \Omega\left(\frac{y}{x}\right).$$

□

B Chernoff Bound Used in Section 4

For completeness, we provide a proof for the following Chernoff bound.

Lemma B.1. *Let $X_1, \ldots, X_k$ be independent random variables such that for each $i \in [k]$, $\mathbb{P}(X_i = a_i) = p$ and $\mathbb{P}(X_i = 0) = 1 - p$, where $a_i > 0$ and p is a given probability. Further, let $A = \sum_{i=1}^k a_i$, $\hat{a} = \max_{i \in [k]} a_i$, and $\mu = \mathbb{E}[X] = pA$, where the random variable X is defined as $X = \sum_{i=1}^k X_i$. For any $\delta > 0$, we then have*

$$\mathbb{P}(X \leq (1-\delta)\mu) \leq e^{-\frac{\delta^2}{2} \cdot \frac{\mu}{\hat{a}}} = e^{-\frac{\delta^2}{2} \cdot \frac{pA}{\hat{a}}}.$$

Proof. Using standard arguments, for any $\eta > 0$ and any $t \geq 0$, we obtain

$$\begin{aligned} \mathbb{P}(X \leq t) &\leq \mathbb{P}(e^{\eta X} \leq e^{\eta t}) \leq \frac{\mathbb{E}\left[e^{\eta X}\right]}{e^{\eta t}} = \frac{\prod_{i=1}^k \mathbb{E}[e^{\eta X_i}]}{e^{\eta t}} \\ &= \frac{\prod_{i=1}^k \left(p(e^{\eta a_i} - 1) - 1\right)}{e^{\eta t}} \leq \frac{e^{\sum_{i=1}^k p(e^{\eta a_i} - 1)}}{e^{\eta t}} \leq e^{\frac{pA}{\hat{a}}(e^{\eta \hat{a}} - 1) - \eta t}. \end{aligned} \tag{4}$$

[9]We assume for simplicity that $y/(x+1)$ is a natural number.

The last inequality follows because for all $0 \leq x \leq y$ and all $0 \leq \varepsilon \leq x$, it holds that

$$(e^x - 1) + (e^y - 1) = \sum_{i=1}^{\infty} \frac{x^i + y^i}{i!} \leq \sum_{i=1}^{\infty} \frac{(x-\varepsilon)^i + (y+\varepsilon)^i}{i!} = (e^{x-\varepsilon} - 1) + (e^{y+\varepsilon} - 1).$$

For $t = (1-\delta)\mu$, the expression in (4) is minimized for $\eta = \ln(1-\delta)/\hat{a}$ and we get

$$\mathbb{P}(X \leq (1-\delta)\mu) \leq e^{\frac{\mu}{\hat{a}} \cdot (-\delta - (1-\delta)\ln(1-\delta))} = \left(\frac{e^{-\delta}}{(1-\delta)^{1-\delta}} \right)^{\frac{\mu}{\hat{a}}} \leq e^{-\frac{\delta^2}{2} \cdot \frac{\mu}{\hat{a}}}. \quad \square$$

$\square$