Contention Resolution on a Fading Channel

Jeremy T. Fineman*
Georgetown University
jfineman@cs.georgetown.edu

Seth Gilbert†
Natl. Univ. of Singapore
seth.gilbert@comp.nus.edu.sg

Fabian Kuhn‡
U. of Freiburg, Germany
kuhn@cs.uni-freiburg.de

Calvin Newport§
Georgetown University
cnewport@cs.georgetown.edu

ABSTRACT

In this paper, we study upper and lower bounds for contention resolution on a single hop *fading* channel; i.e., a channel where receive behavior is determined by a signal to interference and noise ratio (SINR) equation. The best known previous solution solves the problem in this setting in $O(\log^2 n/\log\log n)$ rounds, with high probability in the system size n. We describe and analyze an algorithm that solves the problem in $O(\log n + \log R)$ rounds, where R is the ratio between the longest and shortest link, and is a value upper bounded by a polynomial in n for most feasible deployments. We complement this result with an $\Omega(\log n)$ lower bound that proves the bound tight for reasonable R. We note that in the classical *radio network model* (which does not include signal fading), high probability contention resolution requires $\Omega(\log^2 n)$ rounds. Our algorithm, therefore, affirms the conjecture that the spectrum reuse enabled by fading should allow distributed algorithms to achieve a significant improvement on this $\log^2 n$ speed limit. In addition, we argue that the new techniques required to prove our upper and lower bounds are of general use for analyzing other distributed algorithms in this increasingly well-studied fading channel setting.

CCS Concepts

•Networks → Network algorithms; •Theory of computation → Distributed algorithms;

Keywords

contention resolution, leader election, wireless channel, wireless algorithms, SINR model

*Research supported by NSF CCF 1314633
†Research supported by NUS FRC T1 251RES1404
‡Research supported by ERC Grant No. 336495 (ACDC)
§Research supported by NSF CCF 1320279

PODC'16, July 25 - 28, 2016, Chicago, IL, USA

ISBN 978-1-4503-3964-3/16/07…$15.00
DOI: http://dx.doi.org/10.1145/2933057.2933091

1. INTRODUCTION

Contention resolution is one of the oldest and most important problems in distributed computing. In its basic form, an unknown set of nodes are activated and connected to a shared multiple-access channel (MAC). The protocol completes when any active node transmits alone on the channel. This problem (also, in some contexts, called *wake-up* or *leader election*) was introduced in the seminal 1970 paper on the ALOHA radio network [1]. In the decades that followed, researchers have studied it under a wide variety of assumptions concerning the MAC model and algorithm constraints; e.g., [2, 5, 7, 9, 10, 15, 17, 18, 20, 21, 23]. This interest in contention resolution is well-deserved. From a theoretical perspective, the problem reduces to most non-trivial tasks in MAC models, and therefore lower bounds provide fundamental speed limits for distributed computation in these settings. And from a practical perspective, algorithms for contention resolution have been integrated into link-layer implementations for numerous real world systems, including Ethernet, fiber optic, packet-switch radio, and satellite networks.

Our Results. We study contention resolution in a *fading* MAC model—i.e., a model in which radio signal reception is determined by a *signal to interference and noise ratio* (SINR) equation [19] (see Section 2). Signal fading is a characteristic behavior of radio communication. Therefore, by studying contention resolution on a fading MAC, we can probe the possibility of leveraging this property to speed-up distributed computation in the increasingly important radio network setting.

We study randomized algorithms using a fixed transmission power. We develop an algorithm that solves the problem in $O(\log n + \log R)$ rounds, where n is the number of participating nodes, and R describes the ratio between the longest and shortest link in the network. In most feasible deployments, R is upper bounded by a polynomial in n,[1] yielding performance in $O(\log n)$. Our algorithm solves the problem with high probability in n, and does not require any advance knowledge of n. By contrast, the best existing bound requires just under $\log^2 n$ rounds (specifically,

[1]It is, of course, mathematically possible for R to be super-polynomial in n—say, exponential in n—which would cause the $\log^2 n$ bound to dominate. In practice, however, once n grows beyond relatively small numbers, to maintain such a large gap between link distances becomes increasingly infeasible. Even a network size of only 20 nodes, for example, would require that the longest link be on the order of a million times longer than the shortest link.

$O(\log^2 n/\log\log n)$ rounds), and assumes advance knowledge of n [16] (see the related work below).

We then match this upper bound with an $\Omega(\log n)$ lower bound that establishes our result's optimality in most networks. Our lower bound reduces a basic two-player symmetry breaking problem to contention resolution in a carefully constructed large fading network. This technique is potentially applicable to other lower bounds for distributed computation in fading models and is therefore of standalone interest.

Finally, we note that in the standard non-fading *radio network model* [2,3], the lower bound for contention resolution—and therefore for many problems—is $\Omega(\log^2 n)$ rounds. Our result improves this time by a square root, and therefore resolves the long-standing conjecture that the ability to leverage spatial reuse in fading models should allow distributed algorithms to significantly improve on this $\Omega(\log^2 n)$ speed limit.

Our Algorithm. A surprising attribute of our result is the simplicity of the algorithm from which it derives:

> Each participating node starts in an *active* state; at the beginning of each round, each node that is still active broadcasts with a constant probability p (that we fix in our analysis); if an active node receives a message, it becomes inactive.

Arguably, this algorithm represents the simplest possible strategy for reducing contention, and yet it nonetheless implicitly leverages spectrum reuse to provide a time bound that is optimal in a fading model and beats the lower bounds for non-fading models. This outcome was unexpected when we began this investigation, and it is something we find to be pleasingly elegant.

Our Techniques. The simplicity of our algorithm is counterbalanced by the complexity of its analysis. The bulk of this paper is dedicated to showing that it terminates in $O(\log n + \log R)$ rounds. In more detail, a standard approach for analyzing algorithms in a fading model is to divide nodes into *link classes* based on the distance to their nearest neighbors. Existing techniques based on packing arguments prove that nodes in the smallest non-empty link class have limited contention. Our algorithm should then "knock out" a constant fraction of these nodes. These observations yield an $O(\log n \log R)$ bound, i.e., in the worst-case, we spend $\log n$ rounds emptying each of the $\leq \log R$ link classes, in order of smallest to largest.

To break past the $\log n \log R$ threshold, however, we must show that *many* link classes can concurrently experience limited contention. To show this fact requires two analytical innovations that we argue are of general use for the study of distributed algorithms on fading channels. First, we prove that a larger (in terms of link length) link class can have many (though not all) nodes enjoying low contention if there are not *too many* nodes in smaller link classes. We show this in Section 3.2, fully exploiting the geometry of the fading model. Among other tactics, when studying the interference at a receiver from the same (and larger) link classes, we take advantage of the small—but non-trivial—gap between the quadratic growth of interfering nodes as we move away from the receiver, and the super-quadratic fading of signals, as captured by the standard assumption that the fading exponent α is greater than 2. In the space created by this gap, we are able, in some sense, to fit the interference from the bounded number of smaller link class nodes.

Second, we introduce a *fitting strategy* for bounding the fluctuations in link class sizes as the execution proceeds. As described in Section 3.3, we define a sequence of vectors, $q_1, q_2, \ldots$, where each q_t is an upper bound on the link class sizes. These vectors describe, roughly speaking, how these link class sizes would decrease in an ideal execution. We then show that every execution eventually obeys this steady progression. We note that the process of bounding the link class sizes is complicated by two factors in particular: (1) nodes can move between link classes during an execution as their nearest neighbors are deactivated, so progress is not monotonic; and (2) as link class sizes get small (i.e., $o(\log n)$), we can no longer make high probability arguments about their sizes continuing to reduce. We overcome these issues by carefully defining the q vectors, and a careful treatment of probabilities in which we show that the higher error probabilities of small link classes are balanced by the geometrically decreasing error probabilities of the larger classes.

Related Work. The study of contention resolution began with the random access method introduced in the ALOHA paper in 1970 [1], and continued in the 1980's [9, 10, 17, 18, 21, 23] (see [6] for a good survey of early work). Over the next decade, there was an increasing interest in the so-called *radio network model* [2, 3], which assumes concurrent transmissions are lost at all receivers due to collision and—in a break from earlier work—transmitters do not learn the fate of their transmissions. (Once in this model, the contention resolution problem is now also sometimes called *wake-up* [7].) In this model, high probability contention resolution requires $\Theta(\log^2 n)$ rounds—a bound that improves to $\Theta(\log n)$ if you assume receivers can detect collisions [20]. Given an upper bound N on the network size n, the strategy of [2] can be adapted to yield a solution that solves the problem in $O(\log N)$ expected rounds. (Our algorithm, by contrast, works without any information about n, and with high probability.)

Early in the new millennium came a renewed interest in algorithms for *fading models* of radio networks (sometimes called the *SINR* or *physical* model) which claim to better capture the real behavior of radio communication. The bulk of these initial efforts focused on centralized approximation algorithms for core problems such as link scheduling and capacity maximization. Early on, Moscibroda and Wattenhofer [19] proved that algorithms for fading models can achieve better performance than for the radio network model for certain centralized scheduling problems due to spatial reuse of the spectrum (a capability enabled by super-quadratic signal fading). This result implies that a similar advantage should hold for distributed algorithms as well. As detailed below, although others have studied distributed algorithms in the fading model, our paper is one of the first to validate this conjecture by demonstrating a solution to a fundamental distributed problem that significantly outperforms the relevant radio network model lower bound. (A key caveat to this claim is that we restrict our attention to the standard model for the distributed setting where the transmission power is fixed and provided. Under the assumption of power control, it is sometimes possible to do better; e.g., [11]. Similarly, under the assumption of tunable carrier sensing—a generalization of receiver collision detection—it it is also possible to do better than the radio network model without collision detection; e.g., [22].

Specifically, Jurdziński and Kowalski [13] describe how to create a low-contention backbone in $O(\Delta \text{polylog}(n))$ rounds, for local density bound Δ, Daum et al. [4] show how to solve global broadcast in $O(D \log n \cdot \text{polylog}(R))$ rounds, for network diameter D, while Jurdziński et al. [14] provide a $O(D \log^2 n)$ round broadcast solution, which is more efficient for sufficiently large R. Turning to local communication, Goussevskaia et al. [8] show how to solve local broadcast in $O(\Delta \log^3 n)$ nodes, which Halldorsson and Mitra later improved to $O(\Delta \log n + \log^2 n)$ rounds [12]. All of these results can be used to solve contention resolution in a single-hop fading MAC. None, however, perform better than $\Omega(\log^2 n)$ rounds in this context. Moreover, we also note that each of these distributed algorithms at best matches the best known result for their problem in the non-fading radio network model. It follows that none of these previous results demonstrate the potential to *beat* these bounds in a fading setting—a key contribution of our result.

In a recent breakthrough, Jurdziński and Stachowiak [16] confirmed the permeability of the $\log^2 n$ barrier by showing how to solve contention resolution on a fading MAC in $O(\log^2 n / \log\log n)$ rounds. Their algorithm speeds up a standard $O(\log^2 n)$ strategy from the radio network model to now progress a factor of $\log\log n$ times faster. To compensate for this speed-up, they also add a dampening strategy that takes advantage of the spatial reuse possible on a fading channel to slow down the algorithm *just enough* at the right phase to allow it to succeed with the needed probability. Unlike our algorithm, their solution requires advance knowledge of a (polynomial) upper bound on the network size. On the other hand, their solution is insensitive to R, whereas our algorithm slows as R increases. To the best of our knowledge, our algorithm remains the first to achieve a significant improvement on the $\log^2 n$ bound for the contention resolution problem on a fading channel.

2. MODEL AND PROBLEM

We model randomized algorithms in a synchronous single hop radio network with fading signals. Let V be a set of *nodes* (representing wireless devices) deployed in the two-dimensional Euclidean plane. Time is divided into synchronous rounds, labeled $1, 2, 3, \ldots$. In each round, each node can either (i) transmit at a fixed power P, or (ii) listen. The standard signal to interference and noise ratio (SINR) equation governs which messages are received by listening nodes. That is, node $v \in V$ receives a message transmitted by a node $u \in V$, in a round where the nodes in $I \subseteq V \setminus \{u, v\}$ also transmit, if and only if v is listening and:

$$SINR(u, v, I) = \frac{\frac{P}{d(u,v)^\alpha}}{N + \sum_{w \in I} \frac{P}{d(w,v)^\alpha}} \geq \beta, \qquad (1)$$

where P is the constant transmission power, $d(x, y)$ is the distance between x and y, $\alpha > 2$ is the path-loss exponent, and $N \geq 0$ is the noise. We refer to $\frac{P}{d(w,v)^\alpha}$ as the interference caused by w at v.

Let R be the ratio of the longest to shortest link in the network. To simplify, we assume that link lengths are normalized so that the shortest is 1 and the longest is R. We assume that $\lg R$ is a whole number.

We assume a *single hop* network. In the context of fading models, this means that all node pairs are sufficiently close to communicate in the absence of interference. Formally, the power P must be sufficiently large so that, for every pair $u, v \in V$: $P > c \cdot \beta \cdot N \cdot d(u, v)^\alpha$, for some constant $c > 1$ (for the purposes herein, it is sufficient to assume $c \geq 4$). This assumption that node pairs have a signal strength at least a constant factor larger than β is standard in defining "single hop" in the SINR context. (Allowing node pairs to be placed at distances too near the communication threshold trivially eliminates the possibility of spatial reuse.)

Problem. The *contention resolution* problem assumes an unknown subset of nodes in V are activated. The problem is solved in the first round in which a participating node transmits alone among all participating nodes. (We emphasize that these nodes receive no *a priori* information about the number or identity of the other active nodes.) We study probabilistic solutions that work with *high probability*, i.e., at least $1 - \frac{1}{n}$.

3. UPPER BOUND ANALYSIS

We now analyze the performance of the simple contention resolution algorithm described in the introduction. Our goal is to prove that the algorithm solves the problem in $O(\log n + \log R)$ rounds. After defining some useful notation and assumptions in Section 3.1, we split the analysis into two parts. The first part (Section 3.2) leverages the geometric properties of the model to analyze conditions under which nodes experience bounded contention. The second part (Section 3.3) makes use of this result to bound the number of rounds required for our algorithm to resolve contention. All omitted proofs (along with more detailed calculations) can be found in the full version.

3.1 Preliminaries

For this analysis, we assume R is upper bounded by a polynomial or a quasipolynomial in n—as this is the interesting case for which our $O(\log n + \log R)$ bound beats existing $O(\log^2 n)$ strategies. (For the case where R is larger, one can default to existing results. If R is unknown, then our algorithm can be interleaved with an existing algorithm.) Accordingly, we use $O(\log n)$ and $O(\log R)$ interchangeably in the following.

For a given round, we partition the active nodes into at most $\lg R$ ***link classes***, $d_0, d_1, \ldots, d_{\log R - 1}$, where d_i contains all nodes whose nearest neighbor is at a distance in the range $[2^i, 2^{i+1})$. For a given link class d_i, we define V_i to be the set of nodes in link class d_i that are active, and n_i to indicate $|V_i|$. We can replace subscript i with $< i$, $\leq i$, $> i$, and $\geq i$, in defining V and n, as needed. When there is only one node left, it has no closest active neighbor, and therefore it is not in any link class. Once we have emptied all link classes the problem is solved: the broadcast that deactivates the final remaining nodes is a solo broadcast.

3.2 Interference Analysis

Fix a single round of an execution of our contention resolution algorithm. All definitions and claims in this section are defined with respect to this fixed round. In the following, for a fixed distance d, let $B(u, d)$ be the set of active nodes within distance d of u. For two natural numbers i and t, and node $u \in V$, let the ***exponential annulus*** $A_t^i(u)$ be the set of active nodes $B(u, 2^{(t+1)}2^i) \setminus B(u, 2^t 2^i)$.

Good Nodes. We use the definition of an exponential annulus to define the notion of a ***good node***. Intuitively, a good node is one for which each annulus centered on the node does not contain too many other nodes, for a definition of "too many" that grows with distance as a function of α (defined in Equation 1). Formally:

Definition 1. *Fix a node $u \in V_i$ (i.e., it is active and in link class d_i). We say u is* **good** *if for every $t \in \{0, \ldots, \lg R\}$, $|A_t^i(u)| \le 96 \cdot 2^{t(\alpha-\epsilon)}$, for $\epsilon = \alpha/2 - 1$.*

We want to show that a large fraction of the good nodes have a good probability of receiving a message and becoming inactive—an event we refer to as being "knocked out." To do so, we first identify a subset of the good nodes that are (at the risk of overusing this word) good candidates to be knocked out. In more detail, for a given link class V_i and constant $s > 0$ (that is fixed in Lemma 4), we define $S_i \subseteq V_i$ to be the largest subset of good nodes in V_i such that for every $u, v \in S_i$, the distance $d(u, v) > (s+1)2^i$. That is, every pair of nodes in S_i is sufficiently far apart. Applying a standard circle-packing argument, we note that S_i contains at least a constant fraction of the good nodes in V_i:

Lemma 2. *If link class d_i contains k good nodes, then $|S_i| = \Theta(k)$.*

Now we bound the expected interference at the nodes in S_i during our fixed round. (It is in this lemma that the probability of broadcast p is fixed.) We begin by bounding the interference from nodes not in S_i. For a given node $u \in S_i$, let its *partner* be the closest active node v to u (breaking ties arbitrarily). We use T_i to be the set of partner nodes of the nodes in S_i. For the purpose of the below lemma, we define *outside interference* for u to be the sum of the signal strengths of all transmitting nodes that are *not* in $S_i \cup T_i$.

Lemma 3. *For any constant $c > 0$, there exists a choice of constant broadcast probability $p > 0$ and a constant c' (depending only on c and α) such that the following holds: Let d_i be a non-empty link class. Then with probability at least $1 - e^{-c'|S_i|}$, for at least $|S_i|/2$ nodes in S_i, the outside interference at each of these nodes is $\le cP/2^{i\alpha}$.*

Proof. The proof proceeds in several steps. We first observe that the collective interference on nodes in S_i is bounded by $c_{max}|S_i|P/2^{i\alpha}$, for some constant c_{max} that depends only on α; similarly, the maximum interference that any single node outside S_i can generate on nodes in S_i is $c_{max}|S_i|P/2^{i\alpha}$, even if they all broadcast at once. (See Claims 1 and 2).

We then use a Chernoff bound to show that the total interference caused by all the nodes outside S_i, since they broadcast with probability p, is at most $c|S_i|P/2^{i\alpha+1}$ with probability at least $1 - e^{-c'|S_i|}$, for some constant c' that depends only on c and c_{max} (i.e., it depends on α). (See Claim 3.) This then immediately implies the lemma.

Claim 1. The total interference experienced by all the nodes in S_i, collectively, is at most $P_{max} = c_{max}|S_i|P/2^{i\alpha}$, for some constant c_{max}.

Proof of Claim 1. Fix a node $u \in S_i$. We now bound the maximum interference at u. In more detail, to sum the interference at u, we sum the possible outside interference from the nodes in every (potentially) non-empty annulus defined with respect to u and d_i; e.g., $A_0^i(u), A_1^i(u), \ldots, A_{\lg R}^i(u)$. Noting that a node in $A_t^i(u)$ is of distance at least $2^i 2^t$ from u, we bound the maximum interference by:

$$\begin{aligned}
\sum_{t=0}^{\lg R} |A_t^i(u)| \frac{P}{(2^i 2^t)^\alpha} &= \frac{P}{2^{i\alpha}} \sum_{t=0}^{\lg R} \frac{|A_t^i(u)|}{2^{t\alpha}} \\
&\le \frac{P}{2^{i\alpha}} \sum_{t=0}^{\lg R} \frac{96 \cdot 2^{t(\alpha-\epsilon)}}{2^{t\alpha}} \\
&= \frac{96 \cdot P}{2^{i\alpha}} \sum_{t=0}^{\lg R} \frac{1}{2^{t\epsilon}} \\
&\overset{\epsilon>0}{<} \frac{96 \cdot P}{2^{i\alpha}} \left(\frac{1}{1 - 1/2^\epsilon} \right)
\end{aligned}$$

The final simplification is taken from a geometric series, observing that 2^ϵ is a constant strictly greater than 1 (because $\epsilon = \alpha/2 - 1 > 0$ for $\alpha > 2$).

Setting $c_{max} = 96(1/(1-1/2^\epsilon))$, we conclude that the total interference at u is at most $c_{max}P/2^{i\alpha}$, and hence the total interference experienced by all the nodes in S_i, collectively, is at most $c_{max}|S_i|P/2^{i\alpha}$. □

Notice that if $c \ge c_{max}$, then we are already done, since we have shown in the proof of Claim 1 that the total interference at each node $u \in S_i$ is at most $c_{max}P/2^{i\alpha}$. For the remainder of the proof, we assume that $c < c_{max}$.

Next, we focus on the nodes broadcasting. For each node u *not* in S_i, let $int(u)$ be the sum total possible interference generated by node u at nodes in S_i, i.e., $int(u) = \sum_{v \in S_i} P/d(u,v)^\alpha$.

Claim 2. For all $u \notin S_i$, $int(u) \le c_{max}P/2^{i\alpha}$, for constant c_{max}.

Proof of Claim 2. Notice that this is actually almost identical to Claim 1, except the symmetric opposite: we are quantifying the interference caused by u instead of the interference caused at u. This claim again follows by summing the possible interference by u in each of the exponential annuli, as before.

However, we cannot depend on node u being good. Instead, we conclude that there cannot be too many nodes in S_i in an annulus because the nodes in S_i are at least distance 2^i apart.

More specifically, for each node in S_i, we can draw a circle of radius 2^{i-1} around it such that no two circles intersect. That is, each point occupies an area of $\pi 2^{2(i-1)}$. The annulus A_t^i has outer radius $2^{t+1}2^i$ and inner radius $2^t 2^i$. Some of the circles drawn around points in S_i may be near the edge of the annulus, and hence we increase its outer radius and decrease its inner radius by 2^{i-1} to account for this. Thus, all the points in S_i in A_t^i are contained in area:

$$\begin{aligned}
\pi(2^{t+1}2^i + 2^{i-1})^2 - \pi(2^t 2^i - 2^{i-1})^2 &\le \pi(2^{t+1}2^i + 2^{i-1})^2 \\
&\le \pi(2^{t+1}2^{i+1})^2 \\
&\le \pi 2^{2(t+i+2)}
\end{aligned}$$

(Note that this is a loose approximation that ignores the inner part that is subtracted off, but it is sufficient.) Since each point occupies $\pi 2^{2(i-1)}$ space, we conclude that there are at most $2^{2(t+3)}$ points in S_i in the annulus A^i_t.

$$\begin{aligned}\sum_{t=0}^{\lg R} |A^i_t(u) \cap S_i| \frac{P}{(2^i 2^t)^\alpha} &= \frac{P}{2^{i\alpha}} \sum_{t=0}^{\lg R} \frac{|A^i_t(u) \cap S_i|}{2^{t\alpha}} \\ &\leq \frac{P}{2^{i\alpha}} \sum_{t=0}^{\lg R} \frac{2^{2t+6}}{2^{t\alpha}} \\ &= \frac{64P}{2^{i\alpha}} \sum_{t=0}^{\lg R} \frac{1}{2^{t(\alpha-2)}} \\ &\overset{\alpha-2>0}{<} \frac{64 \cdot P}{2^{i\alpha}} \left(\frac{1}{1 - 1/2^{\alpha-2}} \right)\end{aligned}$$

Since $(\alpha - 2) \geq (\alpha/2 - 1) = \epsilon$ (and $64 < 96$), we conclude that the maximum interference by any node not in S_i is bounded by $c_{max} P/2^{i\alpha}$. □

We now proceed to show that, with high probability (with respect to $|S_i|$), the total interference when nodes broadcast with some probability p is not too large, i.e., is no greater than $c|S_i|P/2^{i\alpha+1}$.

Claim 3. The total interference caused by outside nodes is $c|S_i|P/2^{i\alpha+1}$ with probability at least $1 - e^{-c'|S_i|}$, for some constant c' that depends only on c and c_{max}.

PROOF OF CLAIM 3. First, if the total interference, even if all the outside nodes broadcast, is already $\leq c|S_i|P/2^{i\alpha+1}$, then we are done. For the remainder of the proof, we assume the maximum total interference created by outside nodes is at least $c|S_i|P/2^{i\alpha+1}$, that is, $\sum_{u \notin S_i \cup T_i} int(u) \geq c|S_i|P/2^{i\alpha+1}$.

We have shown in Claim 2 that for all $u \notin S_i$, $int(u) \leq c_{max}P/2^{i\alpha}$. We define the random variable x_u that equals $int(u)2^{i\alpha}/(c_{max}P)$ when u broadcasts with probability p, and 0 otherwise. Notice that $x_u \in [0, 1]$.

We now fix the probability p that a node broadcasts: choose $p = c/(4c_{max})$ (where the constant c is that given in the statement of Lemma 3). Recall that $c < c_{max}$ (or we would have been done after Claim 1), and so we know that $p \in [0, 1/4]$.

Recall, from Claim 1, we know that $\sum_{u \notin S_i \cup T_i} int(u) \leq c_{max}|S_i|P/2^{i\alpha}$. Hence, we can bound the expected value of the sum of x_u as follows:

$$\begin{aligned}\mathbb{E}\left[\sum_{u \notin S_i \cup T_i} x_u\right] &= \sum_{u \notin S_i} p \cdot int(u)2^{i\alpha}/(c_{max}P) \\ &= p \sum_{u \notin S_i \cup T_i} int(u)2^{i\alpha}/(c_{max}P) \\ &\leq p \cdot (c_{max}|S_i|P/2^{i\alpha})(2^{i\alpha})/(c_{max}P) \\ &\leq p|S_i| \\ &\leq c|S_i|/(4c_{max}) \ .\end{aligned}$$

By the same logic, since we have assumed that $\sum_{u \notin S_i} int(u) \geq c|S_i|P/2^{i\alpha+1}$, we can lower bound the expected value of the sum of x_u:

$$\begin{aligned}\mathbb{E}\left[\sum_{u \notin S_i \cup T_i} x_u\right] &= \sum_{u \notin S_i} p \cdot int(u)2^{i\alpha}/(c_{max}P) \\ &= p \sum_{u \notin S_i \cup T_i} int(u)2^{i\alpha}/(c_{max}P) \\ &\geq p \cdot (c|S_i|P/2^{i\alpha+1})(2^{i\alpha})/(c_{max}P) \\ &\geq p(c/2c_{max})|S_i| \\ &\geq (c^2/8c_{max}^2)|S_i| \ .\end{aligned}$$

We now rely on the standard Chernoff bound which states that for any set of independent random variables $x_1, \ldots, x_n$ where $x_i \in [0, 1]$, we know that $\Pr(\sum_{i=1}^n x_i \geq 2\mu) \leq e^{-\mu/3}$, where $\mu = \mathbb{E}\left[\sum_{i=1}^n x_i\right]$.

Let $\mu = \mathbb{E}\left[\sum_{u \notin S_i \cup T_i}(x_u)\right]$. Then we conclude that:

$$\begin{aligned}\Pr\left(\sum_{u \notin S_i} x_u \geq 2\mu\right) &\leq e^{-\mu/3} \\ &\leq e^{-\left(\frac{c^2}{24c_{max}^2}\right)|S_i|}\end{aligned}$$

Here, we set $c' = \left(\frac{c^2}{24c_{max}^2}\right)$, and conclude that the probability of failure is at most $e^{-c'|S_i|}$, for some constant c' that depends only on c and c_{max}, as stipulated by the claim.

Finally, we observe that when this failure does not occur, i.e., $\sum_{u \notin S_i \cup T_i} x_u < 2\mu$, this implies that the total interference caused by nodes not in S_i can be bounded by:

$$\begin{aligned}\sum_{u \notin S_i \cup T_i} x_u \cdot c_{max}P/2^{i\alpha} &\leq (c_{max}P/2^{i\alpha}) \cdot (2c|S_i|/(4c_{max})) \\ &\leq c|S_i|P/2^{i\alpha+1}\end{aligned}$$

This concludes the proof of the claim. □

From Claim 3, we immediately derive the desired conclusion for Lemma 3: it cannot be the case that there are $|S_i|/2$ nodes in S_i that are subject to more than $cP/2^{i\alpha}$ interference each from outside nodes, since that requires total interference from outside nodes of more than $c|S_i|P/2^{i\alpha+1}$. □

We next consider the interference caused by nodes *in* the set $S_i \cup T_i$ and show that it too is bounded. By leveraging the fact that nodes in S_i are not too close to each other (by definition), we can prove the strong bound that even if all nodes in $S_i \cup T_i$ transmit, the interference this causes at each node in the set is low.

LEMMA 4. *For any constant $c > 0$, there exists a choice of constant $s > 0$ such that the following holds: Let $u \in S_i$ be a good node in S_i, let $v \in T_i$ be u's partner, and let d_i be the link class of u. Then the sum of interference at u from nodes in $S_i \cup T_i \setminus \{v\}$ is no more than $cP/2^{i\alpha}$.*

PROOF. Let $S^i_t(u) = (S_i \cup T_i \setminus \{v\}) \cap A^i_t(u)$ be the set of nodes in S_i in the annulus $A^i_t(u)$. We calculate the interference from the nodes in $S_i \cup T_i \setminus \{v\}$ by summing over each annulus. Every node in S_i is at distance at least $(s+1)2^i$ away from u; hence every node in $S_i \cup T_i$ is at distance at least $s2^i$ away from u, and the smallest non-empty annulus $A^t_i(u)$ is $t = \log s$. Note that nodes in annulus $A^i_t(u)$ are at distance at least $2^i 2^t$ away from u, we rely on the notion of "good" to bound the interference from $S_i \cup T_i \setminus \{v\}$:

$$\sum_{t=\log s}^{\lg R} |A_t^i(u)| \frac{P}{(2^i 2^t)^\alpha} \frac{96 \cdot P}{2^{i\alpha}} \left(\frac{1/s^\epsilon}{1 - 1/2^\epsilon} \right)$$

Setting $s = \left(\frac{96}{c(1-1/2^\epsilon)} \right)^{1/\epsilon}$ we get our needed total interference of $cP/2^{i\alpha}$. □

At this point, we have proved that for each d_i, the set S_i (consisting of well-spaced good nodes) contains many nodes that expect low total interference (both internal and external). We now leverage this property to prove the key corollary: with high probability in $|S_i|$, a constant fraction of the nodes in S_i are knocked out.

Corollary 5. *Let d_i be a non-empty link class. There exists a constant broadcast probability $p > 0$ and constant $c > 0$, such that with probability at least $1 - e^{-c|S_i|}$: at least a constant fraction of the nodes in S_i become inactive.*

Proof. By Lemmas 3 and 4, for any constant c', with probability at least $1 - e^{-|S_i|/c''}$ (for some c''), there exists a subset of nodes $S \subseteq S_i$ such that: S contains at least $|S_i|/2$ nodes and for each node $u \in S$ with partner node $v \in T_i$, even if all nodes in $S_i \cup T_i \setminus \{u, v\}$ broadcast, the interference at node u is upper bounded by $\leq 2c'P/2^{i\alpha}$. For now, assume these two properties hold for S.

Next, consider a node $u \in S$, and let v be its partner. The interference bounds derived in Lemmas 3 and 4 are independent of whether u and v broadcast. Hence, conditioned on S satisfying the above two properties, with probability $p(1-p)$, node v transmits and node u listens.

If this event (v transmits and u listens) occurs, we argue that node u receives this message and becomes inactive: the strength of the signal from v at u is at least least $P/2^{\alpha(i+1)}$, and the total interference at u is at most $2c'P/2^{i\alpha}$, for an arbitrary choice of c'. Thus the condition for u to receive the message from v is:

$$\frac{P/2^{\alpha(i+1)}}{2c'P/2^{i\alpha} + N} > \beta \ .$$

We see that the condition is satisfied as long as: (i) $c' \leq 1/(2^{\alpha+2}\beta)$ (as determined in Lemma 3 by choice of the constant), and (ii) $P > 4BNd(u,v)^\alpha \geq 2\beta N 2^{\alpha(i+1)}$ (as required by the assumption in Section 2 that the network is connected).

Let $S' \subseteq S$ be the subset of S that is knocked out. Each node in S succeeds independently with probability $p(1-p)$ (since no two nodes share the same partner, by the definition of S_i), and hence the expected number of successes is $p(1-p)|S|$. A standard Chernoff bound then implies that: $\mathbb{P}[|S'| < p(1-p)|S|/2] \leq e^{-p(1-p)|S|/8}$. Thus, with probability at least $1 - e^{-|S_i|/16}$, our above two properties on S hold, and with an additional probability $1 - e^{-p(1-p)|S|/8}$ we know that S' is at least a constant fraction of S (and therefore is a constant fraction of S_i). The corollary follows from a union bound of these two error probabilities. □

Bounding the Number of Good Nodes. So far, we have shown that a constant fraction of the good nodes in a given link class will be knocked out with high probability (w.r.t. the number of such nodes in the class). We now prove that we expect there to be *enough* good nodes, in general, that these knock outs can add up to something significant. The remainder of this section is dedicated to proving a single lemma that takes on this challenge. In particular, it argues that if some link class d_i has more nodes than all smaller link classes combined (by some constant factor), then at least a constant fraction of the nodes in d_i are good. Combining this lemma with Corollary 5 implies that big link classes shrink quickly (see Corollary 7 below). This observation is the club that Section 3.3 will later wield to subdue our algorithm's time complexity.

Lemma 6. *There exists a constant fraction δ, $0 < \delta < 1$, such that for every link class d_i, if $n_{<i} \leq \delta \cdot n_i$, then at least half the nodes in V_i are good.*

Proof. There are two types of nodes found in the annuli of a node in d_i those in link classes d_i or larger, and those from smaller link classes. We deal with each case separately. Before doing so, we introduce an extra definition. We define a node u to be ***extra good*** with respect to $V_{\geq i}$ (or $V_{<i}$), if the number of these nodes that fall within each $A_t^i(u)$ is at least a factor of 2 smaller than required simply to be *good* (i.e., replace the 96 with a 48 in the earlier definition). If a node is extra good with respect to *both* $V_{<i}$ and $V_{\geq i}$ simultaneously, then the node is good as per the original definition. We will show: (1) every node in V_i is extra good with respect to $V_{\geq i}$, and (2) at least half the nodes are extra good with respect to $V_{<i}$.

Step #1: Nodes in $V_{\geq i}$. Fix any node $u \in d_i$ and any annulus $A_t^i(u)$. We will show that the number of nodes from larger link classes $V_{\geq i}$ falling within this annulus satisfies the definition of extra good. By the definition of the link class, no disk of radius $2^i/4$ can contain more than one such node. A straightforward packing argument bounds the number of such disks that fit in the annulus, showing that there are at most $48 \cdot 2^{2t} \leq 48 \cdot 2^{t(\alpha-\epsilon)}$ node from link classes $V_{\geq i}$ in $A_t^i(u)$.

Step #2: Nodes in $V_{<i}$. We now consider the more difficult case. Fix an annulus distance t. We will sum $|A_t^i(u) \cap V_{<i}|$ over every $u \in V_i$. To do so, we leverage the first key technical insight in this argument: $\sum_{u \in V_i} |A_t^i(u) \cap V_{<i}| = \sum_{v \in V_{<i}} |A_t^i(v) \cap V_i|$ (i.e., just switch the endpoint from which we consider each relevant link). This insight is important because the latter sum is counting nodes from V_i, which allows us to deploy the same style of density bound that proved useful in the large link class case. In more detail, let Γ_t^i indicate the above sum. We bound this parameter as:

$$\begin{aligned} \Gamma_t^i &= \sum_{u \in V_i} |A_t^i(u) \cap V_{<i}| = \sum_{v \in V_{<i}} |A_t^i(v) \cap V_i| \\ &\leq n_{<i} \cdot 48 \cdot 2^{2t} \leq n_i \cdot \delta \cdot 48 \cdot 2^{2t}. \end{aligned}$$

This derivation uses the following properties: (1) $|A_t^i(v) \cap V_i| \leq 48 \cdot 2^{2t}$, as we established in our argument above for larger link classes; and (2) $n_{<i} \leq \delta \cdot n_i$, as assumed by the lemma statement.

The second key technical insight in this argument is to note that the average number of nodes in $A_t^i(u) \cap V_{<i}$ for a node $u \in V_i$ is $O(2^{2t})$. This amount is asymptotically less than the number of nodes allowed in this annulus for the definition of extra good. (Recall, a good node can tolerate up to $O(2^{t(\alpha-\epsilon)})$ nodes in this annulus, and because $\epsilon =$

$\alpha/2-1$ and $\alpha > 2$, this simplifies to $2^{t(\alpha/2+1)} = \omega(2^{2t})$). It follows the number of nodes in V_i that violate the definition of (extra) good for this annulus distance must be *less* than a constant fraction. To be more precise, suppose that $n_i \cdot \gamma \cdot \frac{1}{2^{t(\alpha-\epsilon-2)}}$ nodes violate the extra good definition for this annulus, for a constant fraction γ, $0 < \gamma < 1$, which we will fix later. At a minimum, these nodes contribute more than the following number of nodes to Γ_t^i:

$$n_i \cdot \gamma \cdot \frac{1}{2^{t(\alpha-\epsilon-2)}} \cdot 48 \cdot 2^{t(\alpha-\epsilon)} = n_i \cdot \gamma \cdot 48 \cdot \frac{2^{t(\alpha-\epsilon)}}{2^{t(\alpha-\epsilon-2)}} = n_i \cdot \gamma \cdot 48 \cdot 2^{2t}.$$

If we fix $\delta < \gamma$, then this sum is already larger than Γ_t^i. It follows, therefore, that the fraction of nodes that are not extra good in this annulus must be less than $\gamma \cdot \frac{1}{2^{t(\alpha-\epsilon-2)}}$.

We now sum up the total number of nodes that might violate the definition of extra good over all $\lg R$ annulus distances. We note that:

$$\sum_{t=0}^{\lg R} n_i \cdot \gamma \frac{1}{2^{t(\alpha-\epsilon-2)}} = \sum_{t=0}^{\lg R} n_i \cdot \gamma \frac{1}{2^{t\epsilon}} = n_i \cdot \gamma \sum_{t=0}^{\lg R} \frac{1}{c^t},$$

for constant $c = 2^\epsilon > 1$.[2] The fact that $c > 1$ is important as this casts the sum as a converging geometric series. In particular, we know that: $\sum_{t=0}^{\lg R} \frac{1}{c^t} < \sum_{t=0}^{\infty} \frac{1}{c^t} \leq \frac{1}{1-(1/c)}$.

To achieve our final lemma, however, we need our sum over all annuli to be upper bounded by $n_i/2$. Unfortunately, $1/(1-1/c)$ may be much larger than $1/2$, i.e., if c is close to 1. Fixing $\gamma = (1-1/c)/2$, we recast our above upper bound sum of the number of *not* extra good nodes overall as:

$$\sum_{t=0}^{\lg R} n_i \cdot \gamma \frac{1}{2^{\epsilon t}} \leq n_i \cdot \gamma \cdot \frac{1}{1-1/c} = n_i/2,$$

as needed. Having fixed γ, we set $\delta = \gamma/2$. □

Combining the main result of both parts of this analysis section we conclude with the final corollary:

COROLLARY 7. *There exists a constant broadcast probability $p > 0$, constant δ, $0 < \delta < 1$, and constant $c > 0$ such that for every link class d_i, if V_i is non-empty, and $n_{<i} \leq \delta \cdot n_i$, then with probability at least $1 - e^{-c|V_i|}$, at least a constant fraction of the active nodes in V_i become inactive.*

PROOF. If we fix δ as required by Lemma 6, we know by this same lemma that (at least) half the nodes in V_i are good. By Lemma 2, we know that S_i contains at least a constant fraction of the good nodes in V_i. By Corollary 5, we know that there exists a broadcast probability p such that with probability at least $1 - e^{-c'|S_i|}$, for a constant $c' > 0$, a constant fraction of the nodes in S_i, and therefore a (smaller) constant fraction of V_i, receive a message and become inactive. The constant factor c in the probability in our corollary statement is simply the constant needed to satisfy the equality $c'|S_i| = c|V_i|$. □

3.3 Round Complexity Analysis

We now leverage the analysis of interference in a single round from Section 3.2 to bound the behavior of our algorithm over multiple rounds. We will conclude with our main result: our algorithm solves contention resolution in $O(\log n + \log R)$ rounds, with high probability. Our strategy is to first define a series of variables that capture the evolution of link class sizes in an ideal execution. We will then prove that these size bounds provide an upper bound (of sorts) on these sizes in all executions.

[2]To establish that $c > 1$, note that $2^\epsilon = 2^{\alpha/2-1}$. Because $\alpha > 2$, c is defined as 2 raised to some small value greater than 0, which implies $c > 1$.

Definitions. We begin by extracting key constants from Corollary 7. In the following, assume the algorithm uses a broadcast probability no larger than the necessary probability identified by this corollary. Let γ be the constant fraction, also provided by the corollary, that describes an upper bound on the fraction of nodes in a link class d_i that will remain (i.e., not be knocked out), with high probability in the link class size $|V_i| = n_i$, given that the size of smaller link classes sum to something sufficiently small; i.e., $n_{<i} \leq \delta \cdot n_i$.

Building on these values, let γ_{slow}, where $\gamma < \gamma_{slow} < 1$, be a constant fraction larger than γ that we will define later. Let $\rho < 1$ be another positive constant that we will define later, and let $\ell = \lceil \log_{\gamma_{slow}} \rho \rceil$. In the following, let $m = \lg R$. Recall, we assumed for simplicity that m is a whole number.

Class Bounds Vectors. We now define a series of m-vectors, $q_0, q_1, q_2, \ldots$, that we call ***class bound vectors***. These vectors will later be used to provide a sequence of upper bounds on the sizes of the link classes. To define these vectors, we first define a *start step* s_i for each $i \in [m]$ as $s_i = i \cdot \ell$. That is, $s_0 = 0$, $s_1 = \ell$, $s_2 = 2\ell$, and so on. Prior to round s_i, we do not require any progress to be made in link class d_i. For each $i \in [m]$, we define the corresponding position in our class bound vector as follows:

$$\forall t \geq 0 : q_t(i) = \begin{cases} n & \text{if } t \leq s_i, \\ \lfloor q_{t-1}(i) \cdot \gamma_{slow} \rfloor & \text{if } t > s_i \end{cases}$$

In other words, the size of each link class reduces by a factor of γ_{slow} in each phase. The smallest link class, $i = 0$, begins this reduction immediately. For each $i > 0$, the reduction begins an additional ℓ rounds later. Viewed differently, because link class $i-1$ has undergone ℓ more reductions in size than class i whenever $t \geq s_i$, we have $q_t(i-1) = \gamma_{slow}^{\ell} q_t(i) \leq \rho q_t(i)$. That is, ρ can be interpreted as (approximately) the constant ratio in size between two consecutive link classes. Given that γ_{slow} and ℓ are both constants, it is straightforward to bound the smallest t for which all positions in the vector q_t are 0:

CLAIM 8. *Let T be the smallest phase number where $\forall i \in [m], q_T(i) = 0$. Then $T = \Theta(\log n + \log R)$.* □

Link Class Size Behavior. To establish progress, we want to identify a time when a given link class bound is achieved. That is, for each position t (equiv., *step* t) in our sequence of class bound vectors, we want to find an execution round such that in this round, for every link class d_i: $n_i \leq q_t(i)$. This task, however, is complicated by the observation that a link class can both shrink *and* grow (i.e., when a node u's nearest neighbor is knocked out, u's next nearest neighbor might upgrade it to a larger link class). Therefore, we refine our goal, seeking a round such that for that round, *and all following rounds*, the link class sizes are bounded by q_t. With this in mind, we define an auxiliary class bound vector that captures an important threshold of link size that implies performance. Specifically, for each step t and class d_i we define:

$$\widehat{q_{t+1}}(i) = q_t(i)\gamma_{slow} - q_t(i)\rho/(1-\rho).$$

In other words, $\widehat{q_{t+1}}(i)$ is a more aggressive bound on $q_{t+1}(i)$ that subtracts an additional factor of $q_t(i)\rho/(1-\rho)$ from the number of nodes allowed in a class. Intuitively, this bound implies permanence. To make this argument more precise, we need the following claim about the original class bound vector definitions, which follows from the fact that $q_t(i)$ decreases geometrically in t:

LEMMA 9. *If $q_{t+1}(i) < n$ for some step t and link class i, then $q_t(<i) = \sum_{j=0}^{i-1} q_t(j) \le q_t(i)\left(\frac{\rho}{1-\rho}\right)$.*

Now assume that in round r, for link class d_i and step t, two things are true: (1) for all $j < i$, $n_j \le q_t(j)$; and (2) $n_i \le \widehat{q_{t+1}(i)}$. Then n_i will remain less than $q_{t+1}(i)$ for all future rounds. To see why, observe that even if all the at most $q_t(<i)$ nodes in smaller link classes jumped to link class d_i, by the definition of $\widehat{q_{t+1}(i)}$, the size of d_i would still be bounded by $q_{t+1}(i)$. (And no node can join a smaller link class.)

With our auxiliary definitions in place, we can continue our effort to bound the evolution of link class sizes with respect to our class bound vectors. To this end, we define event $r(t)$, for step t, to be the earliest point in the execution where, for every size class d_i, $n_i \le q_t(i)$, and this bound holds for the remainder of the execution. That is, $r(t)$ is the point where the link class sizes *permanently* fall below q_t.

To analyze these events, we divide the execution into *segments* each consisting of a constant number of rounds (we will fix the constant in the proof). Conditioned on the fact that event $r(t)$ occurred by the end of segment $k-1$, we calculate the probability that event $r(t+1)$ happens by the end of segment k. We will show that this probability is constant. To do so, we emphasize that success in each round (and therefore also each multi-round segment) is independent, as they consist of nodes making independent random choices.

LEMMA 10. *Consider a segment and assume that by the beginning of the segment, event $r(t)$ has occurred, for some step $t \ge 0$. Then event $r(t+1)$ occurs by the end of this segment, with constant probability.*

PROOF SKETCH. By assumption, at the beginning of the segment, the vector q_t bounds all size classes. We examine each round in the segment and calculate the probability that by the end of the round, q_{t+1} is a permanent upper bound on all link classes. To do so, we will show the stronger property that the link classes sizes drop below the $\widehat{q_{t+1}}$ sizes. We begin by fixing class d_i and focus on a single round r in the segment.

One round analysis: There are three cases for this class in this round. The first two cases are easy to dispose of: if $q_t(i) = 0$ or if $q_{t+1}(i) = n$, then it is easily observed that the condition is satisfied $q_{t+1}(i)$.

The third case is that $q_{t+1}(i) < n$. First we fix γ_{slow} and ρ. Specifically, we choose: $\gamma_{slow} = \gamma + \rho/(1-\rho)$, and we choose ρ sufficiently small that $\frac{\rho}{1-\rho} < \gamma\delta$, where δ is the constant derived in Corollary 7.

We now focus on round r in the current segment, and calculate the probability that n_i falls below $\widehat{q_{t+1}}(i)$. Assume that prior to round r, $n_i \ge \widehat{q_{t+1}}(i)$. (Otherwise, we are already done.) In order to apply Corollary 7, our main tool for reducing link class sizes, we need that $n_{<i} \le \delta n_i$. By assumption $n_{<i} \le q_t(<i)$, and by Lemma 9, $q_t(<i) \le q_t(i)\left(\rho/(1-\rho)\right)$. We also assumed: $n_i > \widehat{q_{t+1}}(i) = q_t(i)\gamma_{slow} - q_t(i)\rho/(1-\rho)$. Therefore, solving for $q_t(i)$ in the $\widehat{q_{t+1}}(i)$ definition and substituting into our $q_t(<i)$ equation, we get $n_{<i} \le \left(\frac{\rho}{1-\rho}\right)\frac{n_i}{\gamma_{slow}-\rho/(1-\rho)}$. To bound $n_{<i}$, therefore, we need to show the following inequality:

$$n_{<i} \le \left(\frac{\rho}{1-\rho}\right)\frac{n_i}{\gamma_{slow}-\rho/(1-\rho)} < \delta n_i \ .$$

By the manner in which we fixed γ_{slow} and ρ, this inequality is true.

We can now apply Corollary 7 to class i and obtain the following result: with probability at least $1-e^{-cn_i}$, a constant fraction of nodes in n_i are knocked out, i.e., after round r, the number of remaining active nodes (not counting new nodes that migrate to d_i from smaller link classes) is at most

$$\gamma n_i \le (\gamma_{slow} - \rho/(1-\rho))q_t(i) = \widehat{q_{t+1}}(i) \ .$$

That is, with probability at least $1-e^{-c|n_i|}$, link class d_i falls below $\widehat{q_{t+1}(i)}$, and by our above arguments, will therefore remain bounded by $q_{t+1}(i)$ permanently—regardless of migrations from smaller link classes.

Concluding: To this point, we have shown that in every round in our fixed segment, with probability at least $1-e^{-cn_i}$, the size of link class d_i drops permanently below the threshold $q_{t+1}(i)$ (i.e., as indicated by dropping below $\widehat{q_{t+1}(i)}$). Until we meet this threshold, this is true independently in each round independently of what happened in previous rounds, as each round's outcome is based on the independent broadcast choices made during the round.

Fix the segment length to $\max\{(2\tau)/(c(1-\gamma_{slow})), 2\tau\} = \Theta(1)$ rounds, where $\tau > 1$ is the constant that satisfies $q_{t+1}(i) = \tau\widehat{q_{t+1}}(i)$. Notice that

$$\tau \le q_{t+1}(i)/\left(q_t(i)\left[\gamma_{slow} - \rho/(1-\rho)\right]\right) \le \gamma_{slow}/\gamma.$$

The probability of class d_i *not* dropping below the threshold is therefore at most:

$$\begin{aligned} e^{-2\tau n_i/(1-\gamma_{slow})} &\le (1-\gamma_{slow})/(2\tau n_i) \\ &\le (1-\gamma_{slow})/(2\tau\widehat{q_{t+1}}(i)) \\ &= (1-\gamma_{slow})/(2q_{t+1}(i)). \end{aligned}$$

Since n_i has not yet dropped below our target threshold, we can substitute n_i with the value $\widehat{q_{t+1}}(i)$.

We then take a union bound over all the link classes d_j where $q_t(j) > 0$ and $q_{t+1}(j) < n$, calculating the probability that even one of them does not drop permanently below the threshold $q_{t+1}(\cdot)$:

$$\sum (1-\gamma_{slow})/(2n_j) \le \frac{(1-\gamma_{slow})}{2}\sum_{j=0}^{j=\infty}\gamma_{slow}^j \le \frac{1}{2}$$

This follows from the geometrically decreasing values of $q_t(\cdot)$, which bounds the n_j. We conclude that, with probability $\ge 1/2$, by the end of the segment, every link class irrevocably crosses the threshold of q_{t+1}. □

We can now conclude the analysis with our main result:

THEOREM 11. *Our algorithm solves contention resolution in* $O\big(\log n + \log R\big)$ *rounds, with high probability.*

PROOF SKETCH. In each segment, with constant probability, the system advances from $r(t)$ to $r(t+1)$ by Lemma 10. Since segments are independent, we conclude by a Chernoff bound that the process completes with high probability in $\Theta(T)$ segments, that is, (via Claim 8) in $\Theta(\log n + \log R)$ rounds. □

4. LOWER BOUND

By leveraging spatial reuse, the algorithm described above can solve high probability contention resolution in $O(\log n)$ rounds in networks of size n with $O(\log n)$ link classes.[3] Here we prove this bound near tight by establishing a $\Omega(\log n)$ lower bound for the contention resolution problem under these assumptions:

THEOREM 12. *Let* $\mathcal{A}$ *be an algorithm that guarantees to solve contention resolution in* $f(n)$ *rounds, with probability at least* $1-\frac{1}{n}$*, when run in a network of size* n *with* $O(\log n)$ *link classes. It follows that* $f(n) \in \Omega(\log n)$.

We emphasize that our lower bound places no restrictions on the behavior of the algorithms we lower bound; e.g., as compared to many lower bounds in the MAC setting that require the assumption that the algorithm can be described as a fixed sequence of broadcast probabilities (see [20] for more discussion). The algorithms we bound, for example, can correlate behavior in one round with the outcome of random choices in a previous round, or have nodes that have successfully communicated correlate their behavior. As noted in the introduction, there are not many non-trivial lower bounds known for distributed computing in fading models, rendering this general strategy—which connects fading to combinatorial hitting games—of standalone interest.

We now continue with the proof details.

Restricted k-Hitting. We begin by introducing an abstract hitting game that was defined and bounded in [20]. We will use this existing result below to generate our bound for the related problem of two-player symmetry breaking. In more detail, the *restricted k-hitting* game is defined for an integer $k > 1$. It is played between a *player* and a *referee*. At the beginning of the game, the referee chooses a target set $T \subseteq \{1, 2, ..., k\}$ such that $|T| = 2$. The game proceeds in rounds. In each round, the player proposes a set $P \subseteq \{1, 2, ..., k\}$. If $|P \cap T| = 1$, the player wins. Otherwise, it moves to the next round, learning no information except that its proposal did not win. We can lower bound solutions to this game as follows:

LEMMA 13 (FROM [20]). *Fix some player* $\mathcal{P}$ *that guarantees, for all* $k > 1$*, to solve the restricted* k*-hitting game in* $f(k)$ *rounds, with probability at least* $1-\frac{1}{k}$*. It follows that* $f(k) \in \Omega(\log k)$.

[3] When we say a network has ℓ link classes, we mean there are ℓ link classes that contain at least one of the $\binom{n}{2}$ possible links in the network. It is straightforward to show that the algorithm analyzed in this paper solves contention resolution in $O(\log n+\ell)$ rounds.

Two-Player Contention Resolution. Consider a *two-player* variant of the contention resolution problem that we parameterize with an integer $k > 1$, and that requires the players to break symmetry with probability at least $1 - \frac{1}{k}$. Notice, with two players, the fading behavior of the channel does not matter as with only two nodes there is no opportunity for spatial reuse. The game is won the first time one player transmits while the other listens: in all previous rounds, no messages are received.

We now lower bound this problem by reducing k-hitting to it and then applying Lemma 13.

LEMMA 14. *Let* $\mathcal{A}$ *be an algorithm that solves two-player contention resolution in* $f(k)$ *rounds with probability* $1 - \frac{1}{k}$*, for every parameter* $k > 1$*. It follows that* $f(k) \in \Omega(\log k)$.

PROOF. Fix some algorithm $\mathcal{A}$ that solves two-player contention resolution in $f(k)$ rounds with probability $1 - \frac{1}{k}$, for every parameter k. We use $\mathcal{A}$ to construct a player $\mathcal{P}_{\mathcal{A}}$ for the restricted k-hitting game that works in this same time with this same probability. It will then follow, by the direct application of Lemma 13, that $f(k) \in \Omega(\log k)$.

In more detail, our player simulates $\mathcal{A}$ on k nodes with unique ids from $\{1, 2, ..., k\}$. Each simulated round corresponds to a round of the restricted hitting game as follows: first, the player proposes the set containing the id of every node that broadcast in the current simulated round; then second, the player completes its simulation of the round by simulating all k nodes receiving nothing.

Let $T = \{i, j\}$ be the target set (unknown to the player) chosen for this instance of the restricted hitting game. Although the player is simulating k processes, we only care about its simulation of the processes with ids i and j. Notice, in each round of its simulation, these processes (like all others in the simulation) receive nothing. We want to argue that this behavior is consistent with an execution where only two nodes, with ids i and j, are present in the network. To see why this is true, notice that there are two relevant cases for what happens in a given round. The first case is that exactly one of the pair broadcasts. Here, it would be incorrect to subsequently simulate both receiving nothing (as in an execution with just i and j, the silent node would hear from the broadcaster). In this case, however, the player's proposal would win the restricted hitting game *before* the player is required to simulate the receive behavior for the current simulated round. The second case is that both nodes are silent or both broadcast. Here, the player is correct to simulate them receiving nothing, as this is consistent with what would happen in an execution with just i and j present, during a round where both did the same behavior.

Because the states of simulated nodes i and j are consistent with an execution in which only nodes i and j are present, it follows that $\mathcal{A}$ must solve two-player contention resolution with respect to these two nodes in $f(k)$ rounds, with probability at least $1-\frac{1}{k}$. It follows, therefore, that our player solves the restricted hitting game in this same time with this same probability—as required by our argument. □

Reducing Two-Player to General Contention Resolution (sketch). To prove Theorem 12, we will prove that an algorithm that solves contention resolution in $f(n)$

rounds, with probability $1 - \frac{1}{n}$, when executed in a network of size n (for any $n > 1$), can be used to solve two-player contention resolution in $f(k)$ rounds, for any parameter k. The high-level idea is to have the two players—whom we will call Alice and Bob in this argument—each simulate the general contention resolution algorithm on half the nodes in a network of size k. In each round, if any of Alice's (resp. Bob's) simulated nodes transmit, then Alice (resp. Bob) transmits. Assuming Alice and Bob keep their local simulations valid, eventually exactly one node in the full simulated network transmits, meaning exactly one of the two simulators transmits—winning the two-player version of the problem.

The tricky part of this argument is defining the network simulated by Alice and Bob in such a way that they can keep their local simulations valid, even though they do not always know the behavior of the nodes simulated by the other player. Recall that we need this network to still maintain the strong assumptions that the network is single hop (for every (a, b), $SINR(a, b, \emptyset) \geq c \cdot \beta$, for some fixed constant $c \geq 1$) and that there are only $O(\log n)$ link classes (recall, our theorem statement only makes assumptions about the algorithm's behavior when the link class count is bounded, therefore the network we use in our simulation must bound the link classes). Our strategy is to carefully construct a network (e.g., identify node positions and valid model parameter values) in which interference from Bob cannot change the outcome of receive behaviors among Alice's nodes, and vice versa. The details of this network construction and the final proof for Theorem 12 are deferred to the full version.

5. REFERENCES

[1] N. Abramson. The ALOHA System: Another Alternative for Computer Communications. In *Proceedings of the Fall Joint Computer Conference*, 1970.

[2] R. Bar-Yehuda, O. Goldreich, and A. Itai. On the Time Complexity of Broadcast in Radio Networks: an Exponential Gap Between Determinism and Randomization. In *Proceedings of the International Symposium on Principles of Distributed Computing*, 1987.

[3] I. Chlamtac and S. Kutten. On Broadcasting in Radio Networks: Problem Analysis and Protocol Design. In *Proceedings of the Conference on Computer Communication*, 1985.

[4] S. Daum, S. Gilbert, F. Kuhn, and C. Newport. Broadcast in the Ad Hoc SINR Model. In *Proceedings of the International Symposium on Distributed Computing*, 2013.

[5] S. Daum, F. Kuhn, and C. Newport. Efficient Symmetry Breaking in Multi-Channel Radio Networks. In *Proceedings of the International Symposium on Distributed Computing*, 2012.

[6] R. Gallager. A Perspective on Multiaccess Channels. *IEEE Transactions on Information Theory*, 31(2):124–142, 1985.

[7] L. Gasieniec, A. Pelc, and D. Peleg. The Wakeup Problem in Synchronous Broadcast Systems. *SIAM Journal on Discrete Mathematics*, 14(2):207–222, 2001.

[8] O. Goussevskaia, T. Moscibroda, and R. Wattenhofer. Local Broadcasting in the Physical Interference Model. In *Proceedings of the International Workshop on the Foundations of Mobile Computing*. ACM, 2008.

[9] A. Greenberg and S. Winograd. A Lower Bound on the Time Needed in the Worst Case to Resolve Conflicts Deterministically in Multiple Access Channels. *Journal of the ACM*, 32(3):589–596, 1985.

[10] B. Hajek and T. van Loon. Decentralized Dynamic Control of a Multiaccess Broadcast Channel. *IEEE Transactions on Automatic Control*, 27(3):559–569, 1982.

[11] M. M. Halldorsson and P. Mitra. Distributed Connectivity of Wireless Networks. In *Proceedings of the ACM Conference on Distributed Computing*, 2012.

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

[13] T. Jurdzinski and D. Kowalski. Distributed Backbone Structure for Algorithms in the SINR Model of Wireless Networks. In *Proceedings of the International Conference on Distributed Computing*, 2012.

[14] T. Jurdzinski, D. R. Kowalski, M. Rozanski, and G. Stachowiak. On the impact of geometry on ad hoc communication in wireless networks. In *Proceedings of the ACM Conference on Distributed Computing*, 2014.

[15] T. Jurdziński and G. Stachowiak. Probabilistic Algorithms for the Wakeup Problem in Single-Hop Radio Networks. In *Algorithms and Computation*, pages 535–549. Springer, 2002.

[16] T. Jurdzinski and G. Stachowiak. The cost of synchronizing multiple-access channels. In *Proceedings of the ACM Conference on Distributed Computing*, 2015.

[17] M. Kaplan. A Sufficient Condition for Non-Ergodicity of a Markov Chain. *IEEE Transactions on Information Theory*, IT-25:470–471, July 1979.

[18] J. Komlos and A. Greenberg. An Asymptotically Nonadaptive Algorithm for Conflict Resolution in Multiple-Access Channels. *IEEE Transactions on Information Theory*, 31(2):302–306, 1985.

[19] T. Moscibroda and R. Wattenhofer. The Complexity of Connectivity in Wireless Networks. In *Proceedings of the IEEE International Conference on Computer Communications*, 2006.

[20] C. Newport. Radio Network Lower Bounds Made Easy. In *Proceedings of the International Symposium on Distributed Computing*, 2014.

[21] L. G. Roberts. ALOHA Packet System with and Without Slots and Capture. *ACM SIGCOMM Computer Communication Review*, 5(2):28–42, 1975.

[22] 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 *Proceedings of the International Symposium on Mobile Ad Hoc Networking and Computing*, 2008.

[23] D. Willard. Log-logarithmic selection resolution protocols in a multiple access channel. *SIAM Journal on Computing*, 15(2):468–477, 1986.