# On Simple Back-Off in Unreliable Radio Networks*


## Seth Gilbert
National University of Singapore, Singapore
seth.gilbert@comp.nus.edu.sg

## Nancy Lynch
MIT, Cambridge, MA, USA
lynch@csail.mit.edu

## Calvin Newport
Georgetown University, Washington, DC, USA
cnewport@cs.georgetown.edu

## Dominik Pajak
MIT, Cambridge, MA, USA
pajak@csail.mit.edu



## Abstract

In this paper, we study local and global broadcast in the dual graph model, which describes communication in a radio network with both reliable and unreliable links. Existing work proved that efficient solutions to these problems are impossible in the dual graph model under standard assumptions. In real networks, however, simple back-off strategies tend to perform well for solving these basic communication tasks. We address this apparent paradox by introducing a new set of constraints to the dual graph model that better generalize the slow/fast fading behavior common in real networks. We prove that in the context of these new constraints, simple back-off strategies now provide efficient solutions to local and global broadcast in the dual graph model. We also precisely characterize how this efficiency degrades as the new constraints are reduced down to non-existent, and prove new lower bounds that establish this degradation as near optimal for a large class of natural algorithms. We conclude with an analysis of a more general model where we propose an enhanced back-off algorithm. These results provide theoretical foundations for the practical observation that simple back-off algorithms tend to work well even amid the complicated link dynamics of real radio networks.



**2012 ACM Subject Classification** Theory of computation → Design and analysis of algorithms → Distributed algorithms, Networks → Network types → Ad hoc networks

**Keywords and phrases** radio networks; broadcast; unreliable links; distributed algorithm; robustness

**Digital Object Identifier** 10.4230/LIPIcs.OPODIS.2018.0

**Related Version** The full version is available at `https://arxiv.org/abs/1803.02216`


## 1 Introduction

In this paper, we study upper and lower bounds for efficient broadcast in the dual graph radio network model [4, 12, 13, 3, 6, 5, 8, 7, 15, 9], a dynamic network model that describes

* Definitions and preliminary results concerning the local broadcast problem appeared in the brief announcement [10], published in the Proceedings of 32nd International Symposium on DIStributed Computing (DISC) 2018.



22nd International Conference on Principles of Distributed Systems (OPODIS 2018).
Editors: Jiannong Cao, Faith Ellen, Luis Rodrigues, and Bernardo Ferreira; Article No. 0; pp. 0:1–0:16
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

wireless communication over both reliable and unreliable links. As argued in previous studies of this setting, including unpredictable link behavior in theoretical wireless network models is important because in real world deployments radio links are often quite dynamic.

**The Back-Off Paradox.** Existing papers [13, 8, 15] proved that it is impossible to solve standard broadcast problems efficiently in the dual graph model without the addition of strong extra assumptions (see related work). In real radio networks, however, which suffer from the type of link dynamics abstracted by the dual graph model, simple back-off strategies tend to perform quite well. These dueling realities seem to imply a dispiriting gap between theory and practice: basic communication tasks that are easily solved in real networks are impossible when studied in abstract models of these networks.

*What explains this paradox?* This paper tackles this fundamental question.

As detailed below, we focus our attention on the *adversary* entity that decides which unreliable links to include in the network topology in each round of an execution in the dual graph model. We introduce a new type of adversary with constraints that better generalize the dynamic behavior of real radio links. We then reexamine simple back-off strategies originally introduced in the standard radio network model [2] (which has only reliable links), and prove that for reasonable parameters, these simple strategies *now do* guarantee efficient communication in the dual graph model combined with our new, more realistic adversary.

We also detail how this performance degrades toward the existing dual graph lower bounds as the new constraints are reduced toward non-existent, and prove lower bounds that establish these bounds to be near tight for a large and natural class of back-off strategies. Finally, we perform investigations of even more general (and therefore more difficult) variations of this new style of adversary that continue to underscore the versatility of simple back-off strategies.

We argue that these results help resolve the back-off paradox described above. When unpredictable link behavior is modeled properly, predictable algorithms prove to work surprisingly well.

**The Dual Graph Model.** The dual graph model describes a radio network topology with two graphs, $G = (V, E)$ and $G' = (V, E')$, where $E \subseteq E'$, $V$ corresponds to the wireless devices, $E$ corresponds to reliable (high quality) links, and $E' \setminus E$ corresponds to unreliable (quality varies over time) links. In each round, all edges from $E$ are included in the network topology. Also included is an additional subset of edges from $E' \setminus E$, chosen by an *adversary*. This subset can change from round to round. Once the topology is set for the round, the model implements the standard communication rules from the classical radio network model: a node $u$ receives a message broadcast by its neighbor $v$ in the topology if and only if $u$ decides to receive and $v$ is its only neighbor broadcasting in the round.

We emphasize that the abstract models used in the sizable literature studying distributed algorithms in wireless settings do not claim to provide high fidelity representations of real world radio signal communication. They instead each capture core dynamics of this setting, enabling the investigation of fundamental algorithmic questions. The well-studied radio network model, for example, provides a simple but instructive abstraction of message loss due to collision. The dual graph model generalizes this abstraction to also include network topology dynamics. Studying the gaps between these two models provides insight into the hardness induced by the types of link quality changes common in real wireless networks.

**The Fading Adversary.** Existing studies of the dual graph model focused mainly on the information about the algorithm known to the model adversary when it makes its edge choices. In this paper, we place additional constraints on how these choices are generated.

| Problem | Time | Prob. | Remarks | Ref. |
|---|---|---|---|---|
| Local broadcast | $O\left(\frac{\Delta^{1/\bar{\tau}}\cdot\bar{\tau}^2}{\log\Delta}\cdot\log(1/\epsilon)\right)$ | $1-\epsilon$ | $\bar{\tau}=\min\{\tau,\log\Delta\}$ | Thm 6 |
| | $\Omega\left(\frac{\Delta^{1/\tau}\tau}{\log\Delta}\right)$ | $\frac{1}{2}$ | $\tau\in O(\log\Delta)$ | Thm 7 |
| | $\Omega\left(\frac{\Delta^{1/\tau}\tau^2}{\log\Delta}\right)$ | $\frac{1}{2}$ | $\tau\in O(\log\Delta/\log\log\Delta)$ | Thm 8 |
| Global broadcast | $O\left((D+\log(n/\epsilon))\cdot\frac{\Delta^{1/\bar{\tau}}\bar{\tau}^2}{\log\Delta}\right)$ | $1-\epsilon$ | $\bar{\tau}=\min\{\tau,\log\Delta\}$ | Thm 9 |
| | $\Omega\left(D\cdot\frac{\Delta^{1/\tau}\tau}{\log\Delta}\right)$ | $\frac{1}{2}$ | $\tau\in O(\log\Delta)$ | Thm 10 |
| | $\Omega\left(D\cdot\frac{\Delta^{1/\tau}\tau^2}{\log\Delta}\right)$ | $\frac{1}{2}$ | $\tau\in O(\log\Delta/\log\log\Delta)$ | Thm 10 |

**Table 1** A summary of the upper and lower bounds proved in this paper, along with pointers to the corresponding theorems. In the following, $n$ is the network size, $\Delta\leq n$ is an upper bound on local neighborhood size, $D$ is the (reliable link) network diameter, and $\tau$ is the stability factor constraining the adversary.

In more detail, in each round, the adversary independently draws the set of edges from $E'\setminus E$ to add to the topology from some probability distribution defined over this set. We do not constrain the properties of the distributions selected by the adversary. Indeed, it is perfectly valid for the adversary in a given round to use a point distribution that puts the full probability mass on a single subset, giving it full control over its selection for the round. We also assume the algorithm executing in the model has no advance knowledge of the distributions used by the adversary.

We do, however, constrain how often the adversary can change the distribution from which it selects these edge subsets. In more detail, we parameterize the model with a *stability factor*, $\tau\geq 1$, and restrict the adversary to changing the distribution it uses at most once every $\tau$ rounds. For $\tau=1$, the adversary can change the distribution in every round, and is therefore effectively unconstrained and behaves the same as in the existing dual graph studies. On the other extreme, for $\tau=\infty$, the adversary is now quite constrained in that it must draw edges independently from the same distribution for the entire execution. As detailed below, we find $\tau\approx\log\Delta$, for local neighborhood size $\Delta$, to be a key threshold after which efficient communication becomes tractable.

Notice, these constraints do not prevent the adversary from inducing large amounts of changes to the network topology from round to round. For non-trivial $\tau$ values, however, they do require changes that are nearby in time to share some underlying stochastic structure. This property is inspired by the general way wireless network engineers think about unreliability in radio links. In their analytical models of link behavior (used, for example, to analyze modulation or rate selection schemes, or to model signal propagation in simulation), engineers often assume that in the short term, changes to link quality come from sources like noise and multi-path effects, which can be approximated by independent draws from an underlying distribution (Gaussian distributions are common choices for this purpose). Long term changes, by contrast, can come from modifications to the network environment itself, such as devices moving, which do not necessarily have an obvious stochastic structure, but unfold at a slower rate than short term fluctuations.

In our model, the distribution used in a given round captures short term changes, while the adversary's arbitrary (but rate-limited) changes to these distributions over time capture long term changes. Because these general types of changes are sometimes labeled *short/fast fading* in the systems literature (e.g., [17]), we call our new adversary a *fading adversary*.

**Our Results and Related Work.** In this paper, we study both *local* and *global* broadcast. The local version of this problems assumes some subset of devices in a dual graph network are provided broadcast messages. The problem is solved once each receiver that neighbors a broadcaster in $E$ receives at least one message. The global version assumes a single broadcaster starts with a message that it must disseminate to the entire network. Below we summarize the relevant related work on these problems, and the new bounds proved in this paper. We conclude with a discussion of the key ideas behind these new results.

*Related Work.* In the standard radio network model, which is equivalent to the dual graph model with $E = E'$, Bar-Yehuda et al. [2] demonstrate that a simple randomized back-off strategy called *Decay* solves local broadcast in $O(\log^2 n)$ rounds and global broadcast in $O(D \log n + \log^2 n)$ rounds, where $n = |V|$ is the network size and $D$ is the diameter of $G$. Both results hold with high probability in $n$, and were subsequently proved to be optimal or near optimal[1] [1, 14, 16].

In [12, 13], it is proved that global broadcast (with constant diameter), and local broadcast require $\Omega(n)$ rounds to solve with reasonable probability in the dual graph model with an offline adaptive adversary controlling the unreliable edge selection, while [8] proves that $\Omega(n/\log n)$ rounds are necessary for both problems with an online adaptive adversary. As also proved in [8]: even with the weaker oblivious adversary, local broadcast requires $\Omega(\sqrt{n}/\log n)$ rounds, whereas global broadcast *can* be solved in an efficient $O(D \log(n/D) + \log^2 n)$ rounds, but only if the broadcast message is sufficiently large to contain enough shared random bits for all nodes to use throughout the execution. In [15], an efficient algorithm for local broadcast with an oblivious adversary is provided given the assumption of geographic constraints on the dual graphs, enabling complicated clustering strategies that allow nearby devices to coordinate randomness.

*New Results.* In this paper, we turn our attention to local and global broadcast in the dual graph model with a fading adversary constrained by some stability factor $\tau$ (unknown to the algorithm). We start by considering upper bounds for a simple back-off style strategy inspired by the *decay* routine from [2]. This routine has broadcasters simply cycle through a fixed set of broadcast probabilities in a synchronized manner (all broadcasters use the same probability in the same round). We prove that this strategy solves local broadcast with probability at least $1 - \epsilon$, in $O\left(\frac{\Delta^{1/\bar{\tau}} \cdot \bar{\tau}^2}{\log \Delta} \cdot \log(1/\epsilon)\right)$ rounds, where $\Delta$ is an upper bound on local neighborhood size, and $\bar{\tau} = \min\{\tau, \log \Delta\}$.

Notice, for $\tau \geq \log \Delta$ this bound simplifies to $O(\log \Delta \log(1/\epsilon))$, matching the optimal results from the standard radio network model.[2] This performance, however, degrades toward the polynomial lower bounds from the existing dual graph literature as $\tau$ reduces from $\log \Delta$ toward a minimum value of 1. We show this degradation to be near optimal by proving that *any* local broadcast algorithm that uses a fixed sequence of broadcast probabilities requires $\Omega(\Delta^{1/\tau}\tau/\log \Delta)$ rounds to solve the problem with probability $1/2$ for a given $\tau$. For $\tau \in O(\log \Delta / \log \log \Delta)$ , we refine this bound further to $\Omega(\Delta^{1/\tau}\tau^2/\log \Delta)$, matching our upper bound within constant factors.

We next turn our attention to global broadcast. We consider a straightforward global broadcast algorithm that uses our local broadcast strategy as a subroutine. We prove that this algorithm solves global broadcast with probability at least $1 - \epsilon$, in $O(D + \log(n/\epsilon)) \cdot$

---

[1] The broadcast algorithm from [2] requires $O(D \log n + \log^2 n)$ rounds, whereas the corresponding lower bound is $\Omega(D \log(n/D) + \log^2 n)$. This gap was subsequently closed by a tighter analysis of a natural variation of the simple *Decay* strategy used in [2]

[2] To make it match exactly, set $\Delta = n$ and $\epsilon = 1/n$, as is often assumed in this prior work.

$\Delta^{1/\bar{\tau}}\bar{\tau}^2/\log\Delta$) rounds, where $D$ is the diameter of $G$, and $\bar{\tau} = \min\{\tau, \log\Delta\}$. Notice, for $\tau \geq \log\Delta$ this bound reduces to $O(D\log\Delta + \log\Delta\log(1/\epsilon))$, matching the near optimal result from the standard radio network model. As with local broadcast, we also prove the degradation of this performance as $\tau$ shrinks to be near optimal. (See Table 1 for a summary of these results and pointers to where they are proved in this paper.)

Finally we consider the generalized model when we allow correlation between the distributions selected by the adversary within a given stable period of $\tau$ rounds. It turns out that in the case of arbitrary correlations any simple algorithm needs time $\Omega(\sqrt{\Delta}/l)$ if it uses only cycles of length $l$. In particular any our previous algorithms would require time $\Omega(\sqrt{\Delta}/\log\Delta)$ in the model with arbitrary correlations. The adversary construction in this lower bound requires large changes in the degree of a node in successive steps. Such changes are unlikely in real networks thus we propose a restricted version of the adversary. We assume that the expected change in the degree of any node can be at most $\Delta^{1/(\bar{\tau}(1-o(1))}$. With such restriction it is again possible to propose a simple, but slightly enhanced, back-off strategy (with a short cycle of probabilities) that works efficiently in time $O\left(\Delta^{1/\bar{\tau}} \cdot \bar{\tau} \cdot \log(1/\epsilon)\right)$.

*Technique Discussion.* Simple back-off strategies can be understood as experimenting with different *guesses* at the amount of contention afflicting a given receiver. If the network topology is static, this contention is fixed, therefore so is the *right* guess. A simple strategy cycling through a reasonable set of guesses will soon arrive at this right guess—giving the message a good chance of propagating.

The existing lower bounds in the dual graph setting deploy an adversary that changes the topology in each round to specifically thwart that round's guess. In this way, the algorithm never has the right guess for the current round so its probability of progress is diminished. The fading adversary, by contrast, is prevented from adopting this degenerate behavior because it is required to stick with the same distribution for $\tau$ consecutive rounds. An important analysis at the core of our upper bounds reveals that any fixed distribution will be associated with a right guess defined with respect to the details of that distribution. If $\tau$ is sufficiently large, our algorithms are able to experiment with enough guesses to hit on this right guess before the adversary is able to change the distribution.

More generally speaking, the difficulty of broadcast in the previous dual graph studies was *not* due to the ability of the topology to change dramatically from round to round (which can happen in practice), but instead due to the model's ability to precisely tune these changes to thwart the algorithm (a behavior that is hard to motivate). The dual graph model with the fading adversary preserves the former (realistic) behavior while minimizing the latter (unrealistic) behavior.

## 2 Model and Problem

We study the dual graph model of unreliable radio networks. This model describes the network topology with two graphs $G = (V, E)$ and $G' = (V, E')$, where $E \subseteq E'$. The $n = |V|$ vertices in $V$ correspond to the wireless devices in the network, which we call *nodes* in the following. The edge in $E$ describe reliable links (which maintain a consistently high quality), while the edges in $E' \setminus E$ describe unreliable links (which have quality that can vary over time). For a given dual graph, we use $\Delta$ to describe the maximum degree in $G'$, and $D$ to describe the diameter of $G$.

Time proceeds in synchronous rounds that we label $1, 2, 3...$ For each round $r \geq 1$, the network topology is described by $G_r = (V, E_r)$, where $E_r$ contains all edges in $E$ plus a subset of the edges in $E' \setminus E$. The subset of edges from $E' \setminus E$ are selected by an *adversary*. The graph $G_r$ can be interpreted as describing the high quality links during round $r$. That is, if $\{u, v\} \in E_r$, this mean the link between $u$ and $v$ is strong enough that $u$ could deliver a message to $v$, or garble another message being sent to $v$ at the same time.

With the topology $G_r$ established for the round, behavior proceeds as in the standard radio network model. That is, each node $u \in V$ can decide to transmit or receive. If $u$ transmits, it learns nothing about other messages transmitted in the round (i.e., the radios are half-duplex). If $u$ receives and exactly one neighbor $v$ of $u$ in $E_r$ transmits, then $u$ receives $v$'s message. If $u$ receives and two or more neighbors in $E_r$ transmit, $u$ receives nothing as the messages are lost due to collision. If $u$ receives and no neighbor transmits, $u$ also receives nothing. We assume $u$ does not have collision detection, meaning it cannot distinguish between these last two cases.

**The Fading Adversary.** A key assumption in studying the dual graph model are the constraints placed on the adversary that selects the unreliable edges to include in the network topology in each round. In this paper, we study a new set of constraints inspired by real network behavior. In more detail, we parameterize the adversary with a *stability factor* that we represent with an integer $\tau \geq 1$. In each round, the adversary must draw the subset of edges (if any) from $E' \setminus E$ to include in the topology from a distribution defined over these edges. The adversary selects which distributions it uses. Indeed, we assume it is *adaptive* in the sense that it can wait until the beginning of a given round before deciding the distribution it will use in that round, basing its decision on the history of the nodes' transmit/receive behavior up to this point, including the previous messages they send, but not including knowledge of the nodes' private random bits.

The adversary is constrained, however, in that it can change this distribution at most once every $\tau$ rounds. On one extreme, if $\tau = 1$, it can change the distribution in every round and is effectively unconstrained in its choices. On the other other extreme, if $\tau = \infty$, it must stick with the same distribution for every round. For most of this paper, we assume the draws from these distributions are independent in each round. Toward the end, however, we briefly discuss what happens when we generalize the model to allow more correlations.

As detailed in the introduction, because these constraints roughly approximate the fast/slow fading behavior common in the study of real wireless networks, we call a dual graph adversary constrained in this manner a *fading adversary.*

**Problem.** In this paper, we study both the *local* and *global* broadcast problems. The local broadcast problem assumes a set $B \subseteq V$ of nodes are provided with a message to broadcast. Each node can receive a unique message. Let $R \subseteq V$ be the set of nodes in $V$ that neighbor at least one node in $B$ in $E$. The problem is solved once every node in $R$ has received at least one message from a node in $B$. We assume all nodes in $B$ start the execution during round 1, but do not require that $B$ and $R$ are disjoint (i.e., broadcasters can also be receivers). The global broadcast problem, by contrast, assumes a single source node in $V$ is provided a broadcast message during round 1. The problem is solved once all nodes have received this message. Notice, local broadcast solutions are often used as subroutines to help solve global broadcast.

**Uniform Algorithms.** The broadcast upper and lower bounds we study in this paper focus on *uniform algorithms*, which require nodes to make their probabilistic transmission decisions according to a predetermined sequence of broadcast probabilities that we express as a repeating cycle, $(p_1, p_2, ..., p_k)$ of $k$ probabilities in synchrony. In studying global broadcast, we assume that on first receiving a message, a node can wait to start making probabilistic transmission decisions until the cycle resets. We assume these probabilities can depend on $n$, $\Delta$ and $\tau$ (or worst-case bounds on these values).

In uniform algorithms in the model with fading adversary an important parameter of any node $v$ is its *effective degree* in step $t$ denoted by $d_t(v)$ and defined as the number of nodes $w$ such that $(v, w) \in E_t$ and $w$ has a message to transmit (i.e., will participate in step $t$).

As mentioned in the introduction, uniform algorithms, such as the *decay* strategy from [2], solve local and global broadcast with optimal efficiency in the standard radio network model. A major focus of this paper is to prove that they work well in the dual graph model as well, if we assume a fading adversary with a reasonable stability factor.

The fact that our lower bounds assume the algorithms are uniform technically weaken the results, as there might be non-uniform strategies that work better. In the standard radio network model, however, this does not prove to be the case: uniform algorithms for local and global broadcast match lower bounds that hold for all algorithms (c.f., discussion in [16]).

## 3 Local broadcast

We begin by studying upper and lower bounds for the local broadcast problem. Our upper bound performs efficiently once the stability factor $\tau$ reaches a threshold of $\log \Delta$. As $\tau$ decreases toward a minimum value of 1, this efficiency degrades rapidly. Our lower bounds capture that this degradation for small $\tau$ is unavoidable for uniform algorithms. In the following we use the notation $\bar{\tau} = \min\{\tau, \lceil \log \Delta \rceil\}$. By $\log n$ we will always denote logarithm at base 2 and by $\ln n$ the natural logarithm.

### 3.1 Upper Bound

All uniform local broadcast algorithms behave in the same manner: the nodes in $B$ repeatedly broadcast according to some fixed cycle of $k$ broadcast probabilities. We formalize this strategy with algorithm `RLB` (Robust Local Broadcast) described below (we break out `Uniform` into its own procedure as we later use it in our improved `FRLB` local broadcast algorithm as well):

**Procedure:** `Uniform`$(k, p_1, p_2, \dots, p_k)$

```
for i = 1, 2, ..., k do
    if has message then
        with probability p_i Transmit otherwise Listen
    else
        Listen // without a message always listen
```

**Algorithm:** `RLB`$(r, \bar{\tau})$

```
for i ← 1 to τ̄ do p_i ← Δ^(−i/τ̄)
repeat r times
    Uniform (τ̄, p_1, p_2, ..., p_τ̄)
```

Before we prove the complexity of `RLB` we will show two useful properties of any uniform algorithm. Let $R_t^{(v)}$ denote the event that node $v$ receives a message from some neighbor in step $t$.

▶ **Lemma 1.** *For any uniform algorithm and any node $v$ and step $t$ if $d_t(v) > 0$ and the algorithm uses in step $t$ probability $p \le 1/2$, then* $\mathbf{Pr}\left[R_t^{(v)}\right] \ge \frac{p \cdot d_t(v)}{(2e)^{p \cdot d_t(v)}}$.

**Proof.** For this to happen exactly one among $d_t(v)$ neighbors of $v$ has to transmit and $v$ must not transmit. Node $v$ does not transmit with probability $1 - p$ if it has the message and clearly with probability 1 if it has the message. Denote by $\alpha = p \cdot d_t(v)$. We have

$$\mathbf{Pr}\left[R_t^{(v)}\right] \ge p d_t(v) \cdot (1-p)^{d_t(v)} = \alpha \cdot \left(1 - \frac{\alpha}{d_t(v)}\right)^{d_t(v)}$$

$$= \alpha \left(\left(1 - \frac{\alpha}{d_i(v)}\right)^{d_t(v)/\alpha - 1} \cdot (1-p)\right)^{\alpha} \ge \alpha (e^{-1}(1-p))^{\alpha} \ge \frac{\alpha}{(2e)^{\alpha}}.$$

◀

▶ **Lemma 2.** *For any uniform algorithm, node $v$ and step $t$ if $d_t(v) > 0$:*

$$\mathbf{Pr}\left[R_t^{(v)} \mid d_t(v) \in [d_1, d_2]\right] \geq \min\left\{\mathbf{Pr}\left[R_t^{(v)} \mid d_t(v) = d_1\right], \mathbf{Pr}\left[R_t^{(v)} \mid d_t(v) = d_2\right]\right\}.$$

**Proof.** If the algorithm uses probability $p$ in step $t$ then $\mathbf{Pr}\left[R_t^{(v)}\right] = pd_t(v)(1-p)^{d_t(v)}$. Seeing this expression as a function of $d_t(v)$ we can compute the derivative and obtain that this function has a single maximum in $d_t(v) = 1/(\ln(1/(1-p)))$. Hence if we restrict $d_t(v)$ to be within a certain interval, then value of the function is lower bounded by the minimum at the endpoints of the interval. ◀

Our upper bound analysis leverages the following useful lemma which can be shown by induction on $n$ (the left side is also known as the Weierstrass Product Inequality):

▶ **Lemma 3.** *For any $x_1, x_2, \ldots, x_n$ such that $0 \leq x_i \leq 1$:*

$$1 - \sum_{i=1}^{n} x_i \leq \prod_{i=1}^{n}(1 - x_i) \leq 1 - \sum_{i=1}^{n} x_i + \sum_{1 \leq i < j \leq n} x_i x_j.$$

To begin our analysis, we focus on the behavior of our algorithm with respect to a single receiver when we use the transmit probability sequence $p_1, p_2, ..., p_{\bar{\tau}}$, where $\bar{\tau} = \min\{\tau, \lceil \log \Delta \rceil\}$, and $p_i = \Delta^{-i/\bar{\tau}}$.

▶ **Lemma 4.** *Fix any receiver $u \in R$ and error bound $\epsilon > 0$. It follows: $\mathtt{RLB}(2\lceil \ln(1/\epsilon) \rceil \cdot \lceil 4e \cdot \Delta^{1/\bar{\tau}} \rceil, \bar{\tau})$ delivers a message to $u$ with probability at least $1 - \epsilon$ in time $O(\Delta^{1/\bar{\tau}} \bar{\tau} \log(1/\epsilon))$.*

**Proof.** It is sufficient to prove the claim for $\tau \leq \log \Delta$. For $\tau > \log \Delta$ we use the algorithm for $\tau = \log \Delta$. Note that any algorithm that is correct for some $\tau$ must also work for any larger $\tau$ because the adversary may not choose to change the distribution as frequently as it is permitted to. In the case where $\tau \leq \log \Delta$ we get that $\Delta^{1/\tau} \geq 2$.

We want to show that if the nodes from $N_u \cap B$ execute procedure $\mathtt{Uniform}(\tau, p_1, \ldots, p_\tau)$ twice, then $u$ receives some message with probability at least $\log \Delta / (2e\Delta^{1/\tau}\tau)$. Every time we execute `Uniform` twice, we have a total of $2\tau$ consecutive time slots out of which, by the definition of our model, at least $\tau$ consecutive slots have the same distribution of the additional edges and moreover stations try all the probabilities $p_1, p_2, \ldots, p_\tau$ (not necessarily in this order). Let $T$ denote the set of these $\tau$ time slots and for $i = 1, 2, \ldots, \tau$ let $t_i \in T$ be the step in which probability $p_i$ is used. We also denote the distribution used in steps from set $T$ by $\mathcal{E}^{(\mathcal{T})}$. Hence we can denote the edges between $u$ and its neighbors that have some message by $E_{part} = \{(u, b) : b \in B\} \cap E'$. We know that the edge sets are chosen independently from the same distribution: $E_t \sim \mathcal{E}^{(\mathcal{T})}$ for $t \in T$. Let us denote by $X_t = |E_t \cap E_{part}|$ the random variable being the number of neighbors that are connected to $u$ in step $t$ and belong to $B$. For each $i$ from 1 to $\tau$ we define $q_i = \mathbf{Pr}\left[\Delta^{(i-1)/\tau} < X_t \leq \Delta^{i/\tau}\right]$, for any $t \in T$. Observe that probabilities $q_i$ do not depend on $t$ during the considered $\tau$ rounds. Moreover since $u \in R$ then $u$ is connected via a reliable edge to at least one node in $B$, thus $E \cap E_{part} \neq \emptyset$, hence $\mathbf{Pr}[X_t = 0] = 0$ thus:

$$\sum_{i=1}^{\tau} q_i = 1, \tag{1}$$

Let $S_i$ denote the indicator random variable being 1 if in $t_i$-th round if exactly one neighbor of $u$ transmits and $u$ is not transmitting in round $t$ and 0 otherwise. Clearly if $S_i = 1$ in some round $t$, then $u$ receives some message in round $t$. Then we would like to show for each $i = 1, 2, \ldots, \tau$ that:

$$\mathbf{Pr}[S_i = 1] \geq \frac{q_i}{2e\Delta^{1/\tau}}. \tag{2}$$

In $t_i$-th slot the transmission probability is $p_i = \Delta^{-i/\tau}$ and the transmission choices done by the stations are independent from the choice of edges $E_{t_i}$ active in round $t_i$. Note that $u$ might also belong $R$ and try to transmit. But since $p_i \leq 1/2$ then $u$ is not transmitting with probability at least $1/2$. If $Q_i$ denotes the event that $\Delta^{(i-1)/\tau} < X_{t_i} \leq \Delta^{i/\tau}$ then:

$$\begin{aligned}
\mathbf{Pr}[\,S_i = 1\,] &\geq \mathbf{Pr}[\,S_i = 1|Q_i\,] \cdot \mathbf{Pr}[\,Q_i\,] \\
&\geq p_i(\Delta^{(i-1)/\tau} + 1) \cdot (1-p_i)^{\Delta^{(i-1)/\tau}} \cdot \frac{1}{2} \cdot q_i \\
&\geq p_i \Delta^{(i-1)/\tau} \cdot (1-p_i)^{\Delta^{i/\tau}-1} \cdot \frac{1}{2} \cdot q_i \\
&\geq \Delta^{-1/\tau} \cdot \left(1 - \frac{1}{\Delta^{i/\tau}}\right)^{\Delta^{i/\tau}-1} \cdot \frac{q_i}{2} \geq \frac{q_i}{2e\Delta^{1/\tau}},
\end{aligned}$$

because inequality $(1-1/x)^{x-1} \geq e^{-1}$ holds for all $x > 0$. Since the edge sets are chosen independently in each step and the random choices of the stations whether to transmit or not are also independent from each other we have:

$$\begin{aligned}
\mathbf{Pr}\left[\bigwedge_{i=1}^{\tau}(S_i = 0)\right] &= \prod_{i=1}^{\tau} \mathbf{Pr}[\,S_i = 0\,] \leq \prod_{i=1}^{\tau}\left(1 - \frac{q_i}{2e\Delta^{1/\tau}}\right) && \text{by Equation (2)} \\
&\leq 1 - \sum_{i=1}^{\tau} \frac{q_i}{2e\Delta^{1/\tau}} + \sum_{1\leq i<j\leq\tau} \frac{q_i q_j}{4e^2\Delta^{2/\tau}} && \text{by Lemma 3} \\
&\leq 1 - \frac{\sum_{i=1}^{\tau} q_i}{2e\Delta^{1/\tau}} + \frac{\left(\sum_{i=1}^{\tau} q_i\right)^2}{4e^2\Delta^{2/\tau}} \\
&\leq 1 - \frac{1}{2e\Delta^{1/\tau}} + \frac{1}{4e^2\Delta^{2/\tau}} \leq 1 - \frac{1}{4e\Delta^{1/\tau}} && \text{by Equation (1)}
\end{aligned}$$

Hence if we execute the procedure for $2\tau\lceil\ln(1/\epsilon)\rceil \cdot \lceil 4e \cdot \Delta^{1/\tau}\rceil$ time steps, we have at least $\lceil\ln(1/\epsilon)\rceil \cdot \lceil 4e \cdot \Delta^{1/\tau}\rceil$ sequences of $\tau$ consecutive time steps in which the distribution over the unreliable edges is the same and the algorithm tries all the probabilities $\{p_1, p_2, \ldots, p_\tau\}$. Each of these procedures fails independently with probability at most $1 - 1/(4e\Delta^{1/\tau})$ hence the probability that all the procedures fail is at most: $\left(1 - \frac{1}{4e\Delta^{1/\tau}}\right)^{\lceil\ln(1/\epsilon)\rceil \cdot \lceil 4e\Delta^{1/\tau}\rceil} \leq e^{-\lceil\ln(1/\epsilon)\rceil} < \epsilon$ ◀

On closer inspection of the analysis of Lemma 4, it becomes clear that if we tweak slightly the probabilities used in our algorithm, we require fewer iterations. In more detail, the probability of a successful transmission in the case where each of the $x$ transmitters broadcasts independently with probability $\alpha/x$ is approximately $\alpha/(2e)^\alpha$. In the previous algorithm we were transmitting in successive steps with probabilities $\Delta^{-1/\tau}, \Delta^{-2/\tau}, \ldots$. Thus if $x = 1$ we would get in $i$-th step $\alpha = \Delta^{-i/\tau}$ and approximately the sum of probabilities of success in $\tau$ consecutive steps would be $\Delta^{-1/\tau}$. The formula $\alpha/(2e)^{-\alpha}$ shows that the success probability depends on $\alpha$ linearly if $\alpha < 1$ ("too small" probability) and depends exponentially on $\alpha$ if $\alpha > 1$ ("too large" probability). In the previous theorem we intuitively only use the linear term. In the next one we would like to also use, to some extent, the exponential term. If we shift all the probabilities by multiplying them by a factor of $\beta > 1$, the total success probability would be approximately $\beta\Delta^{-1/\tau}$ if $x = 1$ and $\beta(2e)^{-\beta}$ if $x = \Delta$. Thus by setting $\beta = \log_{2e}\Delta/\tau$ we maximize both these values.

The following lemma makes this above intuition precise and gains a log-factor in performance in algorithm `FRLB` (Fast Robust Local Broadcast) compared to `RLB`. As part of this analysis, we add a second statement to our lemma that will prove useful during our subsequent analysis of global broadcast. The correctness of this second lemma is a straightforward consequence of the analysis.

```
Algorithm: FRLB(r, τ̄)
for i ← 1 to τ̄ do p_i ← Δ^{-i/τ̄} · log_{2e} Δ/τ̄
repeat r times
  Uniform (τ̄, p_1, p_2, ..., p_τ̄)
```

▶ **Lemma 5.** *Fix any receiver $u \in R$ and error bound $\epsilon > 0$. It follows:*

1. *FRLB*$(2\lceil \ln(1/\epsilon) \rceil \cdot \lceil 4\Delta^{1/\bar{\tau}}\bar{\tau}/\log_{2e}\Delta \rceil, \bar{\tau})$ *completes local broadcast with a single receiver in time* $O\left(\frac{\Delta^{1/\bar{\tau}}\cdot\bar{\tau}^2}{\log\Delta}\cdot\log(1/\epsilon)\right)$ *with probability at least $1-\epsilon$, for any $\epsilon > 0$,*
2. *FRLB*$(2, \bar{\tau})$ *completes local broadcast with a single receiver with probability at least* $\frac{\log_{2e}\Delta}{4\Delta^{1/\bar{\tau}}\bar{\tau}}$.

**Proof Idea.** The proof is similar to the one of Lemma 4. We define the probabilities $q_i$ and events $Q_i$ in the same way. The key difference is in the evaluation of the probability of success in round $t_i$ conditioned on $Q_i$ ($\mathbf{Pr}[\,S_i = 1 \mid Q_i\,]$). Event $Q_i$ restricts the number of neighbors connected to $u$ to some interval. We prove that the success probability $\mathbf{Pr}[\,S_i = 1 \mid Q_i\,]$ is lower bounded by the minimum of the values at the endpoints of this interval. This is true because when $x$ stations transmit with probability $p$ to a common neighbor then the probability of a successful transmission seen as a function of $x$ has a single maximum at $x = 1/p$ hence its value at any point of some fixed interval is lower bounded by the minimum of the values at the endpoints. ◀

In Lemmas 4 and 5 we studied the fate of a single receiver in $R$ during an execution of algorithms RLB and FRLB. Here we apply this result to bound the time for all nodes in $R$ to receive a message, therefore solving the local broadcast problem. In particular, for a desired error bound $\epsilon$, if we apply these lemmas with error bound $\epsilon' = \epsilon/n$, then we end up solving the single node problem with a failure probability upper bounded by $\epsilon/n$. Applying a union bound, it follows that the probability that any node from $R$ fails to receive a message is less than $\epsilon$. Formally:

▶ **Theorem 6.** *Fix an error bound $\epsilon > 0$. It follows that algorithm FRLB*$(2\lceil \ln(n/\epsilon) \rceil \cdot \lceil 4\Delta^{1/\bar{\tau}}\bar{\tau}/\log\Delta \rceil)$ *solves local broadcast in* $O\left(\frac{\Delta^{1/\bar{\tau}}\cdot\bar{\tau}^2}{\log_{2e}\Delta}\cdot\log(n/\epsilon)\right)$ *rounds, with probability at least $1-\epsilon$.*

## 3.2 Lower bound

Observe that for $\tau = \Omega(\log\Delta)$, FRLB has a time complexity of $O(\log\Delta\log n)$ rounds for $\epsilon = 1/n$, which matches the performance of the optimal algorithms for this problem in the standard radio model. This emphasizes the perhaps surprising result that even large amounts of topology changes do not impede simple uniform broadcast strategies, so long as there is independence between nearby changes.

Once $\tau$ drops below $\log\Delta$, however, a significant gap opens between our model and the standard radio network model. Here we prove that gap is fundamental for any uniform algorithm in our model.

In the local broadcast problem, a receiver from set $R$ can have between 1 and $\Delta$ neighbors in set $B$. The neighbors should optimally use probabilities close to the inverse of their number. But since the number of neighbors is unknown, the algorithm has to check all the values. If we look at the logarithm of the inverse of the probabilities (call them *log-estimates*) used in Lemma 4 we get $i \log \Delta/\tau$, for $i = 1, 2, \ldots, \tau$—which are spaced equidistantly on the interval $[0, \log \Delta]$. The goal of the algorithm is to minimize the maximum gap between two adjacent log-estimates placed on this interval since this maximizes the success probability in the worst case. With this in mind, in the proof of the following lower bound, we look at the dual problem. Given a predetermined sequence of probabilities used by an arbitrary uniform algorithm, we seek the largest gap between adjacent log-estimates, and then select edge distributions that take advantage of this weakness.

▶ **Theorem 7.** *Fix a maximum degree $\Delta \geq 10$, stability factor $\tau \leq \log(\Delta - 1)/16$, and uniform local broadcast algorithm $\mathcal{A}$. Assume that $\mathcal{A}$ guarantees with probability at least $1/2$ to solve local broadcast in $f(\Delta, \tau)$ rounds when executed in any dual graph network with maximum degree $\Delta$ and fading adversary with stability $\tau$. It follows that $f(\Delta, \tau) \in \Omega(\Delta^{1/\tau}\tau/\log \Delta)$.*

**Proof Idea.** In this proof we use a star with $\Delta$ arms out of which only one is reliable – all other arms are controlled by the adversary. The single receiver $u$ is the center of the star. For any uniform algorithm we divide the probabilities $p_i$ into sequences of length $\tau$ and find a distribution in which the degree of $u$ is "hard" for each sequence. The algorithm places $\tau$ log-estimates on interval $[0, \log \Delta]$ we, as an adversary, can clearly find a largest gap between adjacent log-estimates of length approximately $\log \Delta/\tau$. We choose the degree $d$ of $u$ such that its logarithm is inside this gap (in correct distances from both its endpoints). With this choice we can upper bound the probability of a successful transmission in any step during these $\tau$ steps, because the distance between the log-estimate and the logarithm of the degree of $u$ gives us lower bound on $dp_i$ if $p_i > 1/d$ or of $1/(dp_i)$ if $p_i < 1/d$ which in turn upper bounds the probability of a successful transmission. ◀

In our next theorem, we refine the argument used in Theorem 7 for the case where $\tau$ is a non-trivial amount smaller than the $\log \Delta$ threshold. We will argue that for smaller $\tau$, the complexity is $\Omega(\Delta^{1/\tau}\tau^2/\log \Delta)$, which more exactly matches our best upper bound. We are able to trade this small amount of extra wiggle room in $\tau$ for a stronger lower bound because it simplifies certain probabilistic obstacles in our argument. Combined with our previous theorem, the below result shows our upper bound performance is asymptotically optimal for uniform algorithms for all but a narrow range of stability factors, for which it is near tight.

▶ **Theorem 8.** *Fix a maximum degree $\Delta \geq 10$, stability factor $\tau \leq \ln(\Delta - 1)/(12 \log\log(\Delta - 1))$, and uniform local broadcast algorithm $\mathcal{A}$. Assume that $\mathcal{A}$ guarantees with probability at least $1/2$ to solve local broadcast in $f(\Delta, \tau)$ rounds when executed in any dual graph network with maximum degree $\Delta$ and fading adversary with stability $\tau$. It follows that $f(\Delta, \tau) \in \Omega(\Delta^{1/\tau}\tau^2/\log \Delta)$.*

**Proof Idea.** The proof is similar to proof of Theorem 7. Here we also find a gap of length $\log \Delta/\tau$ and then we argue that in a "proximity" of each such a large gap there has to exist a large number of log-estimates. The proximity is defined so that all log-estimates outside of it are (almost) irrelevant, give a very small probability of success, if we choose the logarithm of the degree of $u$ to be inside the considered gap. This in turn implies that in the remaining part of the interval the "density" of log-estimates is lower hence there must exist another large gap. By repeating this argument we can derive a contradiction with the assumed time complexity. The reason why we need to restrict $\tau$ is that our defined proximity must be of the same order as $\log \Delta/\tau$ which is no longer true for $\tau$ being close to $\log \Delta$. ◀

# 4 Global Broadcast

We now turn our attention to the global broadcast problem. Our upper bound will use the same broadcast probability sequence as our best local broadcast algorithm from before. As with local broadcast, for $\tau \geq \log\Delta$, our performance nearly matches the optimal performance in the standard radio network model, and then degrades as $\tau$ shrinks toward 1. Our lower bound will establish that this degredation is near optimal for uniform algorithms in this setting. In this section we also use the notation $\bar{\tau} = \min\{\tau, \lceil \log\Delta \rceil\}$.

## 4.1 Upper Bound

A uniform global broadcast algorithm requires each node to cycle through a predetermined sequence of broadcast probabilities once it becomes *active* (i.e., has received the broadcast message). The only slight twist in our algorithm's presentation is that we assume that once a node becomes active, it waits until the start of the next probability cycle to start broadcasting. To implement this logic in pseudocode, we use the variable $Time$ to indicate the current global round count. We detail this algorithm below (notice, the FRLB(2) is the local broadcast algorithm analyzed in Lemma 5).

```
Algorithm: RGB(ϵ)
Wait until receiving the message
Wait until (Time mod 2τ̄) = 0
repeat ⌈ln(2n/ϵ)⌉ · ⌈4Δ^{1/τ̄} τ̄ / log Δ⌉ times
  FRLB(2)
```

▶ **Theorem 9.** *Fix an error bound $\epsilon > 0$. It follows that algorithm RGB($\epsilon$) completes global broadcast in time $O\left((D + \log(n/\epsilon)) \cdot \frac{\Delta^{1/\bar{\tau}}\bar{\tau}^2}{\log\Delta}\right)$, with probability at least $1-\epsilon$.*

**Proof Idea.** Here we use the same idea as in the proof of [2, Theorem 4]. There a local broadcast algorithm (*Decay*) is used as a black box in a global broadcast algorithm. We use a different local broadcast algorithm (FRLB) but the same analysis applies. ◀

## 4.2 Lower Bound

The global broadcast lower bound of $\Omega(D\log(n/D))$, proved by Kushilevitz and Mansour [14] for the standard radio network model, clearly still holds in our setting, as the radio network model is a special case of the dual graph model where $E' = E$. Similarly, the $\Omega(\log n \log\Delta)$ lower bound proved by Alon *et al.* [1] also applies.[3] It follows that for $\tau \geq \log\Delta$, we almost match the optimal bound for the standard radio network model, and do match the time of the seminal algorithm of Bar-Yehuda et al. [2].

For smaller $\tau$, this performance degrades rapidly. Here we prove this degradation is near optimal for uniform global broadcast algorithms in our model. We apply the obvious approach of breaking the problem of global broadcast into multiple sequential instances of local broadcast (though there are some non-obvious obstacles that arise in implementing this idea). As with our local broadcast lower bounds, we separate out the case where $\tau$ is

[3] This bound is actually stated as $\Omega(\log^2 n)$, but $\Delta = \Theta(n)$ in the lower bound network, so it can be expressed in terms of $\Delta$ as well for our purposes here.

at least a $1/\log\log\Delta$ factor smaller than our $\log\Delta$ threshold, as we can obtain a slightly stronger bound under this assumption.

▶ **Theorem 10.** *Fix a maximum degree $\Delta \geq 10$, stability factor $\tau$, diameter $D \geq 24$ and uniform global broadcast algorithm $\mathcal{A}$. Assume that $\mathcal{A}$ solves global broadcast in expected time $f(\Delta, D, \tau)$ in all graphs with diameter $D$, maximum degree $\Delta$ and fading adversary with stability $\tau$. It follows that:*

1. *if $\tau < \ln(\Delta-1)/(12\log\log(\Delta-1))$ then $f(\Delta, D, \tau) \in \Omega(D\Delta^{1/\tau}\tau^2/\log\Delta)$,*
2. *if $\tau < \ln(\Delta-1)/16$ then $f(\Delta, D, \tau) \in \Omega(D\Delta^{1/\tau}\tau/\log\Delta)$.*

**Proof Idea.** In this proof we connect together $\Omega(D)$ gadgets used in the proof of Theorem 7 (and 8) and lower bound the time the message spends in each of the gadgets. The only problem in this approach is that after the message enters to the next gadget, the adversary might not be allowed to change the distribution for some number of steps. We solve this by keeping a distribution that is "hard" for the first $\tau$ probabilities of the algorithm in each of the gadgets that has not been reached by the message yet. ◀

## 5 Correlations

Here we explore a promising direction for the study of broadcast in realistic radio network models. In particular, the fading adversary studied above assumes that the distribution draws are independent. As we will show, interesting results are still possible when considering the even more general case where the marginal distributions in each step are not necessarily independent in each round. More precisely, in this case, the adversary chooses a distribution over sequences of length at least $\tau$ of the sets of unreliable edges. A sequence from this distribution is used to determine which unreliable edges are active in successive steps. The adversary after a least $\tau$ steps can decide to change the distribution. In this model, we first show a simple lower bound that any uniform algorithm using a short list of probabilities of length $l$ (our algorithms in previous sections always used list of length $\min\{\tau, \log\Delta\}$) needs time $\Omega(\sqrt{n}/l)$ for some graphs. Our lower bound uses distributions over sequences of graphs in which the degrees of nodes change by a large number in successive steps. Such large changes in degree turn out to be crucial as we show that if in the sequence taken from the distribution chosen by the adversary, in every step in expectancy only $O(\Delta^{1/(\tau-o(\tau))})$ edges adjacent to each node can be changed then we can get an algorithm working in time $O(\Delta^{1/\tau}\tau\log(1/\epsilon))$ with probability at least $1-\epsilon$ and using list of probabilities of length $O(\min\{\tau, \log\Delta\})$.

### 5.1 A Lower Bound for Correlated Distributions

The following lower bound shows that any simple back-off algorithm, similar to the ones presented in Section 3, that uses at most $\log\Delta$ probabilities requires time $\Omega(\sqrt{\Delta}/\log\Delta)$ if arbitrary correlations are permitted.

▶ Proposition 1. Any uniform local broadcast algorithm that repeats a procedure consisting of $l$ probabilities requires expected time $\Omega(\sqrt{\Delta}/l)$ in some graph with $\Delta = n-2$ even if $\tau = \infty$.

**Proof.** Denote the procedure that is being used by the algorithm by $\mathcal{P}$. Assume for simplicity that $\sqrt{\Delta}$ is a natural number. We take as a graph a connected pair of stars (a similar graph was used in Theorem 7).

The fist star has arms $v_1, v_2, \ldots, v_\Delta$ and center at $u$. In the fist star, arms $v_1, v_2, \ldots, v_\Delta$ are connected to center $u$ by reliable edges. The second star has arms $v_1, v_2, \ldots, v_\Delta$ and center at $v$. In the second star, connection from $v_1$ to $v$ is reliable and all other connections are unreliable. Note that by such construction, graph $G$ is connected. All nodes, except $v$, are initially holding a message.

The single distribution is defined in the following way. Let $e_i = \min\{1/p_i, \Delta\}$ for $i = 1, 2, \ldots, l$ be the estimates used by procedure $\mathcal{P}$. Let

$$\bar{e}_i = \begin{cases} 1 \text{ if } e_i \geq \sqrt{\Delta}, \\ n \text{ if } e_i < \sqrt{\Delta}. \end{cases}$$

Let $s$ be a number chosen uniformly at random from $\{1, 2, \ldots, l\}$. In our distribution, the degree of $v$ in step $t$ is $d_t = \bar{e}_{1+r_t}$, where $r_t$ is the remainder of $t+s$ modulo $l$. More precisely, in step $t$ in the distribution exactly $d_t - 1$ edges chosen at random among edges between $v$ and $v_2, v_3, \ldots, v_\Delta$ are activated. Observe that before the algorithm starts, the distribution of the degree of node $v$ in each step is simply a uniform number from multiset $\{\bar{e}_1, \bar{e}_2, \ldots, \bar{e}_l\}$. But after step 1 the sequence of degrees of $v$ becomes deterministic and depends only on the value $s$ of the shift. The dependencies are designed in such a way that if $s = l$ (which happens with probability $1/l$) then in any step $t$ of the algorithm, the probability $p_t$ used by the algorithm satisfies either $p_t \cdot d_t \geq \sqrt{\Delta}$ or $p_t \cdot d_t < 1/\sqrt{\Delta}$. This means by Lemma 1 that the success probability is at most $1/\sqrt{\Delta}$ in each step and hence by the union bound the success probability in the whole procedure is at most $l/\sqrt{\Delta}$. Thus with probability at least $1/l$ the algorithm has to repeat procedure $\mathcal{P}$ at least $\sqrt{\Delta}/(2l)$ times to get a constant probability of success. Hence the expected time is $\Omega(\sqrt{\Delta}/l)$. ◀

## 5.2 Locally Limited Changes

The previous section shows that under an adversary that is allowed to use arbitrary correlations then any simple procedure need polynomial time in the worst case.

In this section we want to consider the adversary that can use correlations but cannot change the degree too much in successive steps. Of course once every at most $\tau$ steps the adversary is allowed to define a completely new distribution over the unreliable edges. We want to argue that it is possible to build a simple algorithm resistant to such an adversary. Intuitively the changes of the degree are problematic only if the changes are by a large (non-constant) factor. Note by Lemma 1 that if we perturb the effective degree by only a constant factor then the bound also changes only by a constant factor. Hence in order to design an algorithm that is immune to such changes we should add more "coverage" to the small-degree nodes. We do this by enhancing each phase of algorithm `RLB` with additional steps in which we assume that the effective degree of a node is small. The adversary may try to avoid the successful transmission in these steps by changing the degree (the adversary knows the probabilities used by the algorithm). But having the restriction on the distance the adversary can move the degree allows us to define overlapping "zones" such that in two consecutive steps we are sure to find the degree in one of the zones. We also have to make sure that the whole phase of the new algorithm fits into $\tau$ steps.

Now we present algorithm `RLBC` (Robust Local Broadcast with Correlations). We first show that the algorithm works under $(l, \tau)$–deterministic adversary that can change at most $l$ edges adjacent to each node per round and all the edges from $E' \setminus E$ once every at most $\tau$ rounds. Our algorithm will be resistant to deterministic adversary that can change at most $\tau\Delta^{1/(\tau - o(\tau))}$ edges adjacent to each node in every step.

Then we show that it also works under restricted fading adversary with parameters $\tau$ and $l$. Restricted fading adversary can change the distribution arbitrarily once every at most $\tau$ steps, if the distribution is not changed then the expected change of the degree of any node can be at most $l$. Under these restrictions, the adversary can design arbitrary correlations between successive steps. We show that `RLBC` works with restricted fading adversary with $l$ of at most $\Delta^{1/(\tau - o(\tau))}$.

```
1  Algorithm: RLBC(r, τ)
2  τ̄ = min{⌈log_{2e} Δ/2⌉, τ}
3  a ← ⌈τ̄/ log_{2e} τ̄⌉
4  k ← ⌈Δ^{1/(τ−2a)}⌉
5  e_1 ← k · a
6  e_2 ← k^2 · τ · a
7  repeat 2r times
8  |   RLB(1, τ̄ − 2a)
9  |   repeat a times
10 |   |   Uniform (1, 1/e_1)
11 |   |   Uniform (1, 1/e_2)
```

▶ **Theorem 11.** *If* $\tau \geq 1000$ *Algorithm RLBC*$(8e\lceil\ln(1/\epsilon)\Delta^{1/\tau}\rceil, \tau)$ *solves local broadcast in the presence of* $\left(\left\lfloor \Delta^{\frac{1}{\tau - 2\lceil \tau/\log_{2e}\tau\rceil}} \right\rfloor \tau/2, \tau\right)$*-deterministic adversary in time* $O(\Delta^{1/\tau}\tau\log(1/\epsilon))$ *with probability at least* $1 - \epsilon$.

**Proof Idea.** For a fixed receiver $v$ we want to show that the probability that $v$ receives the message in one of the $r$ cycles (each 2 iterations of loop in Lines $7 - 11$ is one cycle) is at least $p_s = \frac{1}{8ek}$. We do it by separately considering two cases depending on degree $d_t(v)$, where $t$ is the first step of the considered cycle. If $d_t(v) \geq 2l^2$ we can show that the degree cannot change in total in this cycle by more than a factor of 2 (here we use the restriction on the adversary) in which case we can show that in one of the steps of procedure `RLB` the probability of success is at least $p_s$. For smaller degrees $d_t(v) < 2l^2$ we pick $a$ pairs of steps such that in the first step of the pair the algorithm uses probability $1/e_1$ and in the second it uses $1/e_2$. Then we observe that either in the first step of the pair the degree is at most $2l$ in which case broadcasting with probability $1/e_1$ gives probability $p_s/a$ of success. In the opposite case the degree is at least $l$ (here we use the restriction on the adversary) in the second step and broadcasting with probability $1/e_2$ gives probability $p_s/a$ of success. Since we have $a$ such pairs the claim follows. ◀

The case with deterministic adversary can be generalized to stochastic restricted adversary.

▶ **Theorem 12.** *If* $\tau \geq 1000$ *Algorithm RLBC*$(16e\lceil\ln(1/\epsilon)\Delta^{1/\tau}\rceil, \tau)$ *solves local broadcast in the presence of $l$-restricted fading adversary using correlations with* $l = \left\lfloor \Delta^{\frac{1}{\tau(1-1/\log_{2e}\tau)}} \right\rfloor /4$ *in time* $O(\Delta^{1/\tau}\tau\log(1/\epsilon))$ *with probability at least* $1 - \epsilon$.

**Proof Idea.** We show that if an algorithm works with $2l\tau$-deterministic adversary then it also works with $l$-stochastic adversary with correlations. We note that by Markov's inequality with probability at least $1/(2\tau)$ the degree of the receiver changes by at most $2l\tau$. By the union bound with probability at least $1/2$, the degree does not change by more then $2l\tau$ throughout the whole cycle of length $\tau$. For such cycles, the analysis of the deterministic case gives us probability $p_s$ of success. Thus in the stochastic case the probability of success in each cycle is at least $p_s/2$. ◀

## 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. Reuven Bar-Yehuda, Oded Goldreich, and Alon Itai. On the time-complexity of broadcast in multi-hop radio networks: An exponential gap between determinism and randomization. *J. Comput. Syst. Sci.*, 45(1):104–126, 1992. doi:10.1016/0022-0000(92)90042-H.
3. Keren Censor-Hillel, Seth Gilbert, Fabian Kuhn, Nancy Lynch, and Calvin Newport. Structuring unreliable radio networks. *Distributed Computing*, 27(1):1–19, 2014.
4. Andrea E. F. Clementi, Angelo Monti, and Riccardo Silvestri. Round robin is optimal for fault-tolerant broadcasting on wireless networks. *J. Parallel Distrib. Comput.*, 64(1):89–96, 2004.
5. Mohsen Ghaffari. Bounds on contention management in radio networks. Master's thesis, Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, February 2013.
6. Mohsen Ghaffari, Bernhard Haeupler, Nancy Lynch, and Calvin Newport. Bounds on contention management in radio networks. In Marcos K. Aguilera, editor, *Distributed Computing: 26th International Symposium (DISC 2012), Salvador, Brazil, October, 2012*, volume 7611 of *Lecture Notes in Computer Science*, pages 223–237. Springer, 2012.
7. Mohsen Ghaffari, Erez Kantor, Nancy Lynch, and Calvin Newport. Multi-message broadcast with Abstract MAC layers and unreliable links. In *Proceedings of the 33nd Annual ACM Symposium on Principles of Distributed Computing (PODC'14)*, pages 56–65, Paris, France, July 2014.
8. Mohsen Ghaffari, Nancy Lynch, and Calvin Newport. The cost of radio network broadcast for different models of unreliable links. In *Proceedings of the 32nd Annual ACM Symposium on Principles of Distributed Computing*, pages 345–354, Montreal, Canada, July 2013.
9. Mohsen Ghaffari and Calvin Newport. A leader election in unreliable radio networks. In *Proceedings of the International Colloquium on Automata, Languages, and Programming (ICALP)*, 2016.
10. Seth Gilbert, Nancy A. Lynch, Calvin Newport, and Dominik Pajak. Brief announcement: On simple back-off in unreliable radio networks. In *32nd International Symposium on Distributed Computing, DISC 2018, New Orleans, LA, USA, October 15-19, 2018*, pages 48:1–48:3, 2018.
11. Svante Janson. Tail bounds for sums of geometric and exponential variables. *arXiv preprint arXiv:1709.08157*, 2017.
12. Fabian Kuhn, Nancy Lynch, and Calvin Newport. Brief announcement: Hardness of broadcasting in wireless networks with unreliable communication. In *Proceedings of the 28th Annual ACM Symposium on the Principles of Distributed Computing (PODC 2009)*, Calgary, Alberta, Canada, August 2009.
13. Fabian Kuhn, Nancy Lynch, Calvin Newport, Rotem Oshman, and Andrea Richa. Broadcasting in unreliable radio networks. In *Proceedings of the 29th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 336–345, Zurich, Switzerland, July 2010.
14. 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.
15. Nancy Lynch and Calvin Newport. A (truly) local broadcast layer for unreliable radio networks. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, 2015.
16. Calvin Newport. Lower bounds for radio networks made easy. In *Proceedings of the International Symposium on Distributed Computing (DISC)*, 2014.
17. Mike Willis. Propagation tutorial: Fading. http://www.mike-willis.com/Tutorial/PF15.htm. May 5, 2007.