Theoretical Computer Science 806 (2020) 489–508

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

On simple back-off in unreliable radio networks ☆

Seth Gilbert [a], Nancy Lynch [b], Calvin Newport [c], Dominik Pajak [d,*]

[a] National University of Singapore, Singapore
[b] Massachusetts Institute of Technology, United States of America
[c] Georgetown University, United States of America
[d] Wroclaw University of Science and Technology, Tooploox, Poland

ARTICLE INFO

Article history:
Received 21 March 2019
Received in revised form 23 July 2019
Accepted 20 August 2019
Available online 26 August 2019

Keywords:
Radio networks
Broadcast
Unreliable links
Algorithm
Stochastic adversary

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.

1. Introduction

In this paper, we study upper and lower bounds for efficient broadcast in the dual graph radio network model [4,15,16,3,7,6,9,8,18,10], a dynamic network model that describes 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 [16,9,18] 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

☆ Definitions and preliminary results concerning the local broadcast problem appeared in the brief announcement [11], published in the Proceedings of 32nd International Symposium on DIStributed Computing (DISC) 2018. Extended abstract appeared in [12], published in the Proceedings of 22nd International Conference on Principles of Distributed Systems (OPODIS) 2018. D. Pajak was partially supported by Polish National Science Center grant UMO-2018/29/B/ST6/02969.

* Corresponding author.
E-mail addresses: seth.gilbert@comp.nus.edu.sg (S. Gilbert), lynch@csail.mit.edu (N. Lynch), cnewport@cs.georgetown.edu (C. Newport), dominik.pajak@pwr.edu.pl (D. Pajak).

https://doi.org/10.1016/j.tcs.2019.08.027
0304-3975/

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.

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 τ 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 Δ, 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 τ 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. Traditional and well-known fading distributions include Rayleigh fading [21], Ricean fading [22], or the Nakagami distribution [19].

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., [23]), 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 these 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

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 τ is the stability factor constraining the adversary.

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

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] [1,17,20].

In [15,16], 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 [9] proves that $\Omega(n/\log n)$ rounds are necessary for both problems with an online adaptive adversary. As also proved in [9]: 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 [18], 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 τ (known 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 sequence 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 a single receiver 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 Δ is an upper bound on local neighborhood size, and $\bar{\tau}=\min\{\tau,\log\Delta\}$. With multiple receivers the bound changes to $O\left(\frac{\Delta^{1/\bar{\tau}}\cdot\bar{\tau}^2}{\log\Delta}\cdot\log(n/\epsilon)\right)$

Notice, for a single receiver, 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 τ 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 $\Omega(\Delta^{1/\bar{\tau}}\bar{\tau}/\log\Delta)$ rounds to solve the problem with probability $1/2$ for a given τ. 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 $O((D+\log(n/\epsilon))\cdot\Delta^{1/\bar{\tau}}\bar{\tau}^2/\log\Delta)$ 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 $O((D+\log(n/\varepsilon))\cdot\log\Delta\log(1/\epsilon))$, matching the near optimal result from the standard radio network model. As with local broadcast, we also prove the degradation of this performance as τ 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 τ rounds. It turns out that in the case of arbitrary correlations any simple algorithm needs time $\Omega(\sqrt{\Delta}/l^2)$ if this algorithm cycles through a sequence of probabilities of length l. In particular, any of our previous algorithms would require time $\Omega(\sqrt{\Delta}/\log^2\Delta)$ 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

[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 independently in [14,5] 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.

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 τ 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 τ 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 edges 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 Δ 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 means 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 τ 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 extreme, if $\tau = \infty$, it must stick with the same distribution for every round. Note that this does not mean that the communication graph is static but that the adversary cannot change the distribution over the unreliable edges. 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, Δ and τ (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 is not the case: uniform algorithms for local and global broadcast match lower bounds that hold for all algorithms (cf., discussion in [20]).

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 τ reaches a threshold of $\log\Delta$. As τ decreases toward a minimum value of 1, this efficiency degrades rapidly. Our lower bounds capture that this degradation for small τ 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):

```
1 Procedure: Uniform(p_1, p_2, ..., p_k)
2 for i = 1, 2, ..., k do
3     if has message then
4         with probability p_i Transmit otherwise Listen
5     else Listen // without a message always listen
```

```
1 Algorithm: RLB(r, τ̄)
2 for i ← 1 to τ̄ do p_i ← Δ^{−i/τ̄}
3 repeat r times
4     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 event $R_t^{(v)}$ to happen exactly one among $d_t(v)$ neighbors of v has to transmit and v must not transmit in step t. Node v does not transmit with probability $1 - p$ if it has the message and clearly with probability 1 if it does not have the message. Hence the probability that exactly one neighbor transmits equals $pd_t(v)(1-p)^{d_t(v)-1}$ and the probability that v does not transmit is at least $1 - p$. Denote by $\alpha = p \cdot d_t(v)$. We have

$$\mathbf{Pr}\left[R_t^{(v)}\right] \ge pd_t(v) \cdot (1-p)^{d_t(v)} = \alpha \cdot \left(1 - \frac{\alpha}{d_t(v)}\right)^{d_t(v)}$$
$$= \alpha \left(\left(1 - \frac{\alpha}{d_t(v)}\right)^{d_t(v)/\alpha - 1} \cdot (1-p)\right)^{\alpha} \ge \alpha(e^{-1}(1-p))^{\alpha} \ge \frac{\alpha}{(2e)^{\alpha}}. \quad \square$$

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

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

Proof. If the algorithm uses probability p in step t then $\mathbf{Pr}\left[R_t^{(v)} \right] \geq p d_t(v)(1-p)^{d_t(v)}$. Function $f(x) = px(1-p)^x$ for $p \in [0,1]$ has a single maximum for $x = 1/(\ln(1/(1-p)))$. 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 $\epsilon > 0$, $RLB(2\lceil \ln(1/\epsilon) \rceil \cdot \lceil 4e \cdot \Delta^{1/\bar{\tau}} \rceil, \bar{\tau})$ 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 τ must also work for any larger τ 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 some of the neighbors of v belong to B and execute the procedure $\texttt{Uniform}(\tau, p_1, \ldots, p_\tau)$ twice, then v 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 2τ consecutive time slots out of which, by the definition of our model, at least τ 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 τ 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 define $q_1 = \mathbf{Pr}\left[1 \leq d_t(v) \leq \Delta^{i/\tau} \right]$, and for each i from 2 to τ we define $q_i = \mathbf{Pr}\left[\Delta^{(i-1)/\tau} < d_t(v) \leq \Delta^{i/\tau} \right]$, for any $t \in T$. Observe that the edge sets are chosen independently from the same distribution for $t \in T$ hence probabilities q_i do not depend on t during the considered τ rounds. Moreover since $v \in R$, then v is connected via a reliable edge to at least one node in B. This means that $\mathbf{Pr}[\, d_t(v) = 0 \,] = 0$, thus:

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

We would like to bound from below the probability that v receives a message in step t_i for $i = 1, 2 \ldots, \tau$:

$$\mathbf{Pr}\left[R_{t_i}^{(u)} \right] \geq \frac{q_i}{2e\Delta^{1/\tau}}. \tag{2}$$

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. Let Q_1 denote the event that $1 \leq d_{t_1}(v) \leq \Delta^{1/\tau}$ and for $i \geq 2$ let Q_i denote the event that $\Delta^{(i-1)/\tau} < d_{t_i}(v) \leq \Delta^{i/\tau}$. We have $p_i \leq 1/2$ hence we can use Lemmas 1 and 2 to get for $i > 1$

$$\begin{aligned}
\mathbf{Pr}\left[R_{t_i}^{(u)} \right] &\geq \mathbf{Pr}\left[R_{t_i}^{(u)} \mid Q_i \right] \cdot \mathbf{Pr}[\, Q_i \,] \\
&\geq q_i \min\left\{ \mathbf{Pr}\left[R_{t_i}^{(u)} \mid d_{t_i}(u) = \left\lfloor \Delta^{(i-1)/\tau} \right\rfloor + 1 \right], \mathbf{Pr}\left[R_{t_i}^{(u)} \mid d_{t_i}(v) = \Delta^{i/\tau} \right] \right\} \\
&\geq q_i \min\left\{ \mathbf{Pr}\left[R_{t_i}^{(u)} \mid d_{t_i}(u) = \Delta^{(i-1)/\tau} \right], \mathbf{Pr}\left[R_{t_i}^{(u)} \mid d_{t_i}(v) = \Delta^{i/\tau} \right] \right\} \\
&\geq q_i \min\left\{ \frac{\Delta^{-1/\tau}}{(2e)^{\Delta^{-1/\tau}}}, \frac{1}{2e} \right\} \geq \frac{q_i}{2e\Delta^{1/\tau}}.
\end{aligned}$$

For $i = 1$ we have:

$$\begin{aligned}
\mathbf{Pr}\left[R_{t_1}^{(u)} \right] &\geq \mathbf{Pr}\left[R_{t_1}^{(u)} \mid Q_1 \right] \cdot \mathbf{Pr}[\, Q_1 \,] \\
&\geq q_1 \min\left\{ \mathbf{Pr}\left[R_{t_1}^{(u)} \mid d_{t_1}(u) = 1 \right], \mathbf{Pr}\left[R_{t_i}^{(u)} \mid d_{t_1}(v) = \Delta^{1/\tau} \right] \right\} \\
&\geq q_1 \min\left\{ \frac{1}{\Delta^{1/\tau} \cdot (2e)^{\Delta^{-1/\tau}}}, \frac{1}{2e} \right\} \geq \frac{q_1}{2e\Delta^{1/\tau}}.
\end{aligned}$$

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:

$$\begin{aligned}
\mathbf{Pr}\left[\bigwedge_{i=1}^{\tau}\left(\neg R_{t_i}^{(u)}\right)\right] &= \prod_{i=1}^{\tau}\mathbf{Pr}\left[\neg R_{t_i}^{(u)}\right] \le \prod_{i=1}^{\tau}\left(1-\frac{q_i}{2e\Delta^{1/\tau}}\right) && \text{by independence and Equation (2)}\\
&\le 1-\sum_{i=1}^{\tau}\frac{q_i}{2e\Delta^{1/\tau}}+\sum_{1\le i<j\le\tau}\frac{q_iq_j}{4e^2\Delta^{2/\tau}} && \text{by Lemma 3}\\
&\le 1-\frac{\sum_{i=1}^{\tau}q_i}{2e\Delta^{1/\tau}}+\frac{\left(\sum_{i=1}^{\tau}q_i\right)^2}{4e^2\Delta^{2/\tau}}\\
&\le 1-\frac{1}{2e\Delta^{1/\tau}}+\frac{1}{4e^2\Delta^{2/\tau}} \le 1-\frac{1}{4e\Delta^{1/\tau}} && \text{by Equation (1).}
\end{aligned}$$

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 τ 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: $\left(1-\frac{1}{4e\Delta^{1/\tau}}\right)^{\lceil\ln(1/\epsilon)\rceil\cdot\lceil 4e\Delta^{1/\tau}\rceil} \le e^{-\lceil\ln(1/\epsilon)\rceil} < \epsilon$ □

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 α/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 in the case where $x=1$ (a single transmitter) we get in i-th step $\alpha=\Delta^{-i/\tau}$ which implies that the sum of probabilities of success in τ consecutive steps equals approximately $\Delta^{-1/\tau}$. This means that RLB needs $\Omega(\tau\Delta^{-1/\tau})$ steps. In order to improve the time complexity we need to find a way to get a higher success probability. The formula $\alpha/(2e)^{-\alpha}$ shows that the success probability depends on α linearly if $\alpha<1$ ("too small" probability) and depends exponentially on α 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.

Algorithm: FRLB($r, \bar{\tau}$)
for $i \leftarrow 1$ **to** $\bar{\tau}$ **do** $p_i \leftarrow \Delta^{-i/\bar{\tau}}\cdot\log_{2e}\Delta/\bar{\tau}$
repeat r times
  Uniform$(\bar{\tau}, p_1, p_2, \ldots, p_{\bar{\tau}})$

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 similar to the proof of Lemma 4. (Proof of Lemma 5 is deferred to Appendix.)

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

1. *FRLB$(2\lceil\ln(1/\epsilon)\rceil\cdot\lceil 4e^2\Delta^{1/\bar{\tau}}\bar{\tau}/\log_{2e}\Delta\rceil, \bar{\tau})$ 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 $\frac{\log_{2e}\Delta}{4e^2\Delta^{1/\bar{\tau}}\bar{\tau}}$.*

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 ϵ, 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 ϵ/n. Applying a union bound, it follows that the probability that any node from R fails to receive a message is less than ϵ. 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 $1-\epsilon$.*

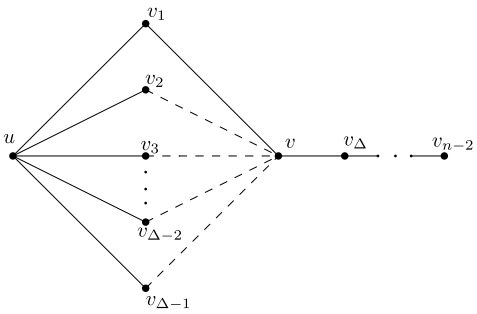

Fig. 1. A graph used in proofs of Theorems 7 and 8. Solid lines correspond to edges in E and dashed lines correspond to edges in $E' \setminus E$ (unreliable edges).

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 τ 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 Δ 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 Δ and fading adversary with stability τ. It follows that $f(\Delta, \tau) \in \Omega(\Delta^{1/\tau} \tau / \log \Delta)$.*

Proof. Consider the dual graph $G = (V, E)$ and $G' = (V, E')$, defined as follows: $V = \{v, u, v_1, \ldots, v_{n-2}\}$ and $E = \{\{u, v_i\}, i \in \{1, 2, \ldots, \Delta - 1\}\} \cup \{\{v_1, v\}, \{v, v_\Delta\}\} \cup \{\{v_i, v_{i+1}\}, i \in \{\Delta, \ldots, n-3\}\}$ and $E' = E \cup \{\{v_i, v\}, i \in \{2, 3, \ldots, \Delta - 1\}\}$ (see Fig. 1). We will study local broadcast in this dual graph with $B = \{u, v_1, v_2, \ldots, v_{\Delta-1}\}$ and $R = \{v\}$. Observe that the maximum degree of any node is indeed Δ and the number of nodes is n. Nodes $v_\Delta, v_{\Delta+1}, \ldots, v_{n-2}$ do not belong to $B \cup R$ hence are not relevant in our analysis.

Using the sequence of probabilities $p_1, p_2, \ldots$ used by algorithm $\mathcal{A}$ we will define a sequence of distributions over the edges that will cause a long delay until node v will receive a message. The adversary we define is allowed to change the distribution every τ steps. Accordingly, we partition the rounds into *phases* of length τ, which we label $1, 2, 3, \ldots$. Phase k consists of time steps $I_k = \{1 + k \cdot \tau, 2 + k \cdot \tau, \ldots, (k+1) \cdot \tau\}$. For each phase $k \geq 1$, the adversary will use a distribution $\mathcal{D}_k$ that's defined with respect to the probabilities used by $\mathcal{A}$ during the rounds in phase k. In particular, let sequence $P_k = \{p_i\}_{i \in I_k}$ be the τ probabilities used by $\mathcal{A}$ during phase k.

We use P_k to define the distribution $\mathcal{D}_k$ as follows. Define $\dot{\Delta} = \Delta - 1$ and let N represent $\lfloor \log \dot{\Delta} \rfloor$ urns labeled with numbers from 1 to $\lfloor \log \dot{\Delta} \rfloor$. Into these urns we place balls with numbers $\lceil \log(1/p_j) \rceil$ and $\lfloor \log(1/p_j) \rfloor$ for all $j \in I_k$. Ball with number i is placed into the bin with the same number. Any probability lower than $1/\Delta$, we treat as equal to $1/\Delta$. Note that with this we never decrease the probability of success of algorithm $\mathcal{A}$. With this procedure for each j, we place two balls in adjacent bins if $\lceil \log(1/p_j) \rceil \neq \lfloor \log(1/p_j) \rfloor$ and a single ball in the opposite case. We arrange the bins in a circular fashion i.e., bins $\lfloor \log \dot{\Delta} \rfloor$ and 1 are consecutive and we want to find the longest sequence of consecutive empty bins. Observe that since for each j we put either a single ball or two balls into adjacent bins we have at most τ sequences of consecutive empty bins. Moreover, since at most 2τ bins contain a ball then there exists a sequence of consecutive empty bins of length at least $\frac{\lfloor \log \dot{\Delta} \rfloor - 2\tau}{\tau}$. Knowing that τ is an integer and that $\tau \leq \log \dot{\Delta}/16$ we can show that: $\frac{\lfloor \log \dot{\Delta} \rfloor - 2\tau}{\tau} \geq \left\lfloor \frac{\log \dot{\Delta}}{\tau} \right\rfloor - 2$. We define:

$$x = \lfloor \log \dot{\Delta}/\tau \rfloor - 3 - \lfloor \log(\lfloor \ln \dot{\Delta}/\tau \rfloor) \rfloor,$$
$$y = \lfloor \log(\lfloor \ln \dot{\Delta}/\tau \rfloor) \rfloor + 1.$$

We observe that for $\tau \leq \log\dot{\Delta}/16$ we have $\log(\lfloor \ln\dot{\Delta}/\tau \rfloor) \geq 4$ hence x and y are both positive integers and moreover $x+y = \lfloor \log\dot{\Delta}/\tau \rfloor - 2$. Hence we already showed that there exists a sequence of consecutive empty bins of length at least $x+y$. Now, we pick the label of $(y+1)$-st bin in this sequence (the order of bins is according to the circular arrangement i.e., 1 comes after $\lfloor \log\dot{\Delta} \rfloor$) and call it a_k. Let $A_k = \{\log(1/p_j) : j \in I_k\}$. This set contains logarithms of all the estimates "tried" by the algorithm in k-th phase. Now we split A_k into elements that are larger and that are smaller than a_k: $A_k = A_k^{(\geq)} \cup A_k^{(<)}$, $A_k^{(\geq)} = \{a \in A_k : a \geq a_k\}$, $A_k^{(<)} = \{a \in A_k : a < a_k\}$. We observe that if $a \in A_k^{(<)}$ then $a \leq a_k - y$ because there are y empty bins between bin a_k and the bin containing ball $\lceil a \rceil$. Symmetrically if $a \in A_k^{(\geq)}$ then $a \geq a_k + x - 1$ because there are $x-1$ empty bins between bin a_k and the bin containing ball $\lfloor a \rfloor$.

In our distribution $\mathcal{D}_k$ in phase k, we include all edges from E, plus a subset of size $2^{a_k} - 1$ selected uniformly from $E' \setminus E$. This is possible since the adversary can choose to activate any subset of links among the set $\{\{v_i, v\}, i \in \{2, \ldots, \dot{\Delta}\}\}$. With this choice, the degree of v is 2^{a_k} in phase k hence we can bound the probability that a successful transmission occurs in phase k.

Having chosen the distribution of the edges between v and $\{v_1, v_2, \ldots, v_{\dot{\Delta}}\}$ we can now bound the probability of a successful transmission in any step t in the considered phase. Recall that the event of a successful reception of a message by v in step t is denoted by $R_t^{(v)}$. For this event to happen exactly one of the 2^{a_k} nodes among $\{v_1, v_2, \ldots, v_{\dot{\Delta}}\}$ that are connected to v need to transmit. We have:

$$\mathbf{Pr}\left[R_t^{(v)} \right] = 2^{a_k} p_t \cdot (1-p_t)^{2^{a_k}-1}.$$

Take any step t and the corresponding probability p_t used by the algorithm. We know that a_k is chosen so that $a_k \geq \log(1/p_t) + y$ or $a_k \leq \log(1/p_t) - x$. We consider these cases separately:

Case 1: $a_k \leq \log(1/p_t) - x + 1$

$$\begin{aligned}
\mathbf{Pr}\left[R_t^{(v)} \right] &= 2^{a_k} \cdot p_t \cdot (1-p_t)^{2^{a_k}-1} \leq 2^{a_k - \log(1/p_t)} \leq 2^{-x+1} \\
&= 2^{-\lfloor \log\dot{\Delta}/\tau \rfloor + 4 + \lfloor \log(\lfloor \ln\dot{\Delta}/\tau \rfloor) \rfloor} \leq 2^{-\log\dot{\Delta}/\tau + \log(\lfloor \ln\dot{\Delta}/\tau \rfloor) + 5} \\
&\leq \frac{32\lfloor \ln\dot{\Delta}/\tau \rfloor}{\dot{\Delta}^{1/\tau}} \\
&\leq \frac{32\ln\dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau}.
\end{aligned}$$

Case 2: $a_k \geq \log(1/p_t) + y$

$$\mathbf{Pr}\left[R_t^{(v)} \right] = 2^{a_k} \cdot p_t \cdot (1-p_t)^{2^{a_k}-1} = \frac{2^{a_k} p_t}{1-p_t}(1-p_t)^{2^{a_k}} \leq \frac{2^{a_k} p_t}{1-p_t} e^{-2^{a_k} p_t}.$$

We know that $a_k \geq \log(1/p_t) + y$ and $a_k \leq \log\dot{\Delta}$ thus $p_t \leq 2^y/\dot{\Delta}$ hence since $\dot{\Delta} \geq 9$ we get $1/(1-p_t) \leq 2$. Moreover since $2^{a_k} p_t \geq 2^y \geq 4$ we have $e^{-2^{a_i} p_t/2} < 1/(2^{a_i} p_t)$ (because $e^{x/2} > x$ for all x). Which gives in this case:

$$\mathbf{Pr}\left[R_t^{(v)} \right] < 2e^{-2^{a_k} p_t/2} \leq 2e^{-2^{y-1}} \leq 2e^{-\lfloor \ln\dot{\Delta}/\tau \rfloor} \leq 2\dot{\Delta}^{-1/\tau} \leq \frac{32\ln\dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau}.$$

We have just shown that the probability that v receives a message in any step t of our fixed phase k is at most $\frac{32\ln\dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau}$. We can repeat this construction of distribution and get that in every step in every phase the success probability is at most $\frac{32\ln\dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau}$. To conclude the proof, we apply a union bound to show that the probability that v receives a message in at least one of $\dot{\Delta}^{1/\tau}\tau/(64\ln\dot{\Delta}) - 1$ steps, is strictly less than $1/2$:

$$\mathbf{Pr}\left[\bigcup_{t=1}^{\dot{\Delta}^{1/\tau}\tau/(64\ln\dot{\Delta})-1} R_t^{(v)} \right] \leq \sum_{t=1}^{\dot{\Delta}^{1/\tau}\tau/(64\ln\dot{\Delta})-1} \mathbf{Pr}\left[R_t^{(v)} \right] < \frac{1}{2}. \quad \square$$

In our next theorem, we refine the argument used in Theorem 7 for the case where τ is a non-trivial amount smaller than the $\log\Delta$ threshold. We will argue that for smaller τ, 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 τ 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 Δ and fading adversary with stability τ. It follows that $f(\Delta,\tau) \in \Omega(\Delta^{1/\tau}\tau^2/\log\Delta)$.*

Proof. In this proof we will use the same graph as in Theorem 7. Let $G = (V, E)$ and $G' = (V, E')$. Let $V = \{v, u, v_1, \ldots, v_{n-2}\}$ and let $R = \{u, v_1, v_2, \ldots, v_{\Delta-1}\}$ and $E = \{\{u, v_i\}, i \in \{1, 2, \ldots, \Delta\}\} \cup \{\{v_1, v\}, \{v, v_\Delta\}\} \cup \{\{v_i, v_{i+1}\}, i \in \{\Delta, \ldots, n-3\}\}$ and $E' = E \cup \{\{v_i, v\}, i \in \{2, 3, \ldots, \Delta\}\}$ (see Fig. 1).

Let $p_1, p_2, \ldots$ be the fixed sequence of broadcast probabilities used by nodes in B running $\mathcal{A}$. Using this sequence we will define a sequence of distributions over the edges that will cause a long time for this algorithm until node v will receive a message.

The adversary is allowed to change the distribution once every τ steps. Therefore we will define the k-th distribution $\mathcal{D}_k$ based on sequence $P_k = (p_{(k-1)\tau+1}, p_{(k-1)\tau+2}, \ldots, p_{k\tau})$ of probabilities and distribution $\mathcal{D}_k$ will be used in rounds $(k-1)\tau + 1, (k-1)\tau + 2, \ldots, k\tau$. Consider intervals of τ time steps (call such interval a *phase*) and the corresponding probabilities $p_{j+i\cdot\tau}$ $(j < \tau)$. Let us fix any phase k and consider values $l_i = \log(1/p_{i+(k-1)\tau})$, for $i = 1, 2, \ldots, \tau$. We denote $\dot{\Delta} = \Delta - 1$. As an adversary we are allowed to define an integer value $l^* \in [1, 2, \ldots, \lfloor \log \dot{\Delta} \rfloor]$ based on the l-values and define a distribution for phase k in which there are always 2^{l^*} active links between nodes $v_1, v_2, \ldots, v_n$ and v. The success probability in i-th step of the considered phase is then

$$s_{i+(k-1)\tau} = p_{i+(k-1)\tau} \cdot 2^{l^*} \cdot (1 - p_{i+(k-1)\tau})^{2^{l^*}-1}.$$

Our goal as an adversary is to find in every phase k, such l^* that minimizes $\sum_{i=1}^{\tau} s_{i+(k-1)\tau}$. We will show that it is always possible to find such l^* that $\sum_{i=1}^{\tau} s_{i+(k-1)\tau} = O(\dot{\Delta}^{-1/\tau} \log \dot{\Delta}/\tau) = \Theta(\Delta^{-1/\tau} \log \Delta/\tau)$. This will give us that $\Omega(\dot{\Delta}^{1/\tau}\tau/\log \dot{\Delta})$ phases of τ steps hence in total $\Omega(\Delta^{1/\tau}\tau^2/\log \Delta)$ steps are needed to complete local broadcast with constant probability.

Consider the first phase (analysis of any subsequent phase is analogical). Assume by contradiction that there exists a choice of $l_1, l_2, \ldots, l_\tau$ such that for any choice of l^* we have that $\sum_{i=1}^{\tau} s_i \geq c\frac{\log \dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau}$, where $c = 2409$. We fix this choice of l-values $l_1, l_2, \ldots, l_\tau$ and we denote:

$$\begin{aligned} x &= \log \dot{\Delta}/\tau + \log \tau - \log \ln \dot{\Delta}, \\ y &= \log(\ln \dot{\Delta}/\tau), \\ x' &= \log \dot{\Delta}/\tau + 2\log \tau - \log \ln \dot{\Delta}, \\ y' &= \log(\ln \dot{\Delta}/\tau + 2\ln \tau). \end{aligned}$$

First note that since $\tau \leq \ln(\Delta - 1)/(12 \log\log(\Delta - 1))$, then:

$$x' \geq 12 \log\log \dot{\Delta} - \log \ln \dot{\Delta} \geq 11 \log\log \dot{\Delta} \geq y' \tag{3}$$

Since $\tau < \ln \dot{\Delta}/(12 \log\log \dot{\Delta})$ we have that $y \geq 3$. Observe that $x + y = \log \dot{\Delta}/\tau$ and $x' + y' \leq \log \dot{\Delta}/\tau + 2\log \tau - \log \ln \dot{\Delta} + \log(\ln \dot{\Delta}/\tau + 2\ln \tau) = \log \dot{\Delta}/\tau + 2\log \tau + \log(1/\tau + 2\ln \tau/\ln \dot{\Delta}) \leq \log \dot{\Delta}/\tau + 2\log \tau + \log 3$. Let Δ^* denote the number of active links between v and $v_1, v_2, \ldots, v_{\dot{\Delta}}$ in the considered phase and $l^* = \log \Delta^*$. In any step if Δ^* is such that $l^* \geq y'$ then if we also have $p_i \geq \frac{2}{3}$ then we get:

$$s_i = p_i 2^{l^*}(1 - p_i)^{2^{l^*}-1} \leq \frac{2^{l^*}}{3^{2^{l^*}-1}} \leq \frac{3(\ln \dot{\Delta}/\tau + 2\ln \tau)}{e^{\ln \dot{\Delta}/\tau + 2\ln \tau}} \leq \frac{9 \ln \dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau^2}, \tag{4}$$

where the last inequality is true because $\tau \geq 1$ and $\ln \tau \leq \ln \dot{\Delta}$. This shows that the sum of all such s_i is at most $\frac{9 \ln \dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau}$. Consider now only steps with $p_i < 2/3$. Then:

$$s_i = 2^{l^*-l_i}\left(1 - \frac{1}{2^{l_i}}\right)^{2^{l^*}-1} = 2^{l^*-l_i} \cdot \frac{\left(1 - \frac{1}{2^{l_i}}\right)^{2^{l^*}}}{1 - p_i} \leq 3 \cdot 2^{l^*-l_i} \cdot e^{-2^{l^*-l_i}}$$

Recall that $R_i^{(v)}$ denotes the event that node v receives the message in i-th step of the considered phase (*i.e.*, a single node out of $v_1, v_2, \ldots, v_{\Delta-1}$ transmits in this round).

$$\mathbf{Pr}\left[R_i^{(v)} \mid l_i \geq l^* + x - 3\right] \leq 3 \cdot 2^{-x+3} \leq \frac{24 \ln \dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau} \tag{5}$$

$$\mathbf{Pr}\left[R_i^{(v)} \mid l_i \geq l^* + x' - 1\right] \leq 3 \cdot 2^{-x'+1} \leq \frac{6 \ln \dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau^2} \tag{6}$$

$$\mathbf{Pr}\left[R_i^{(v)} \mid l_i \leq l^* - y\right] \leq \frac{3 \cdot 2^y}{e^{2^y}} = \frac{3 \ln \dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau} \tag{7}$$

$$\mathbf{Pr}\left[R_i^{(v)} \mid l_i \leq l^* - y'\right] \leq \frac{3 \cdot 2^{y'}}{e^{2^{y'}}} = \frac{3(\ln \dot{\Delta}/\tau + 2\log \tau)}{\dot{\Delta}^{1/\tau}\tau^2} \leq \frac{9 \ln \dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau^2} \tag{8}$$

Observe that for a fixed value of l^*, for any i such that $l_i \notin [l^* - y', l^* + x' - 1]$ we have $s_i \leq \frac{9 \ln \dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau^2}$ (by Equations (6), (8)). Hence the sum of all such values s_i is at most $\frac{9 \ln \dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau}$. Hence we only need to find such l^* that the sum of the values s_i for which the corresponding $l_i \in [l^* - y', l^* + x']$ is less than $\frac{(c-9)\ln \dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau}$.

We denote the smallest and the largest l-values: $l_{sm} = \min_{i\in\{1,2,\ldots,\tau\}}\{l_i\}$ and $l_{lg} = \max_{i\in\{1,2,\ldots,\tau\}}\{l_i\}$. We will prove two following claims about l_{sm} and l_{lg}:

$l_{sm} \leq x'$ Observe that otherwise we can choose $l^* = 0$ (Δ^* is then equal to 1 which corresponds to exactly one active link between $\{v_1, v_2, \ldots, v_{\dot{\Delta}}\}$ and v) and then by Equation (6) under this choice of l^* all values s_i would satisfy $s_i \leq \frac{6\dot{\Delta}^{-1/\tau}\ln \dot{\Delta}}{\tau^2}$ (because if $l_i \geq x'$ then $p_i < 2/3$).

$l_{lg} \geq \log \dot{\Delta} - y'$ If it is not the case, we choose $l^* = \log \dot{\Delta}$ and by Equation (8) we have that if $p_i < 2/3$ then $s_i \leq \frac{9 \ln \dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau^2}$ and by Equation (4) that if $p_i \geq 2/3$ then $s_i \leq \frac{9 \ln \dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau^2}$. And the sum of all values of s_i is at most $\frac{9 \ln \dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau}$ which contradicts our assumption.

Consider now interval $\Gamma_1 = [l_{sm}, l_{lg}]$. Two previous claims showed that $|\Gamma_1| \geq \log \dot{\Delta} - x' - y'$. We can now consider the placement of values l_i on Γ_1 and analyze gaps between the adjacent values. Gap g_i is the difference between the $(i+1)$-st smallest and i-th smallest value out of all values l_j. We want to show the following:

$$\max_i g_i \leq x' + y' \tag{9}$$

Assume on the contrary that such a gap between l_i and l_j exists. Then we pick $l^* = \lceil l_i + y' \rceil$ and observe that l^* is an integer, and $l^* \geq y'$ hence for all i such that $p_i \geq 2/3$ by Equation (4) we have $s_i \leq \frac{9 \log \dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau^2}$. If $p_i < 2/3$ then since l^* is at least y' larger than each smaller l-value and at least $x' - 1$ smaller than each larger l-value, we have by Equations (6) (8) that $s_i \leq \frac{9 \ln \dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau^2}$. Thus if any gap has value at least $x' + y'$ then $\sum_{i=0}^{\tau} s_i \leq \frac{9 \ln \dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau}$ which contradicts our assumption.

We know that there are at most $\tau - 1$ gaps and that they cover area of at least $\log \dot{\Delta} - x' - y'$. Hence we can lower bound the average value of a gap:

$$\begin{aligned} d_1 &= \frac{\log \dot{\Delta} - x' - y'}{\tau - 1} \geq \frac{\log \dot{\Delta} - \log \dot{\Delta}/\tau - 2\log\tau - \log 3}{\tau - 1} \\ &= \frac{\log \dot{\Delta}(1 - 1/\tau)}{\tau(1 - 1/\tau)} - \frac{2\log\tau + \log 3}{\tau} \geq \frac{\log \dot{\Delta}}{\tau} - 2. \end{aligned}$$

Thus there exists a gap with value at least d_1. Knowing that $y \geq 3$ we have $d_1 \geq x + y - 2 \geq 1$ and inside this gap we can find an integer value l_1^* that is at least y larger than the closest smaller l-value and at least $x - 3$ smaller than the closest larger l-value. Consider values of s_i with this choice of l^*. By Equations (5) (7) if $l^* = l_1^*$, each s_i is at most $\frac{24\dot{\Delta}^{-1/\tau}\ln \dot{\Delta}}{\tau}$. Consider now interval $I_1 = [l_1^* - y', l_1^* + x']$. By Equations (6) and (8) for all i such that $l_i \notin I_1$ and $p_i < 2/3$ we have $s_i \leq \frac{9\dot{\Delta}^{-1/\tau}\ln \dot{\Delta}}{\tau^2}$. If $p_i \geq 2/3$ then also $s_i \leq \frac{9\dot{\Delta}^{-1/\tau}\ln \dot{\Delta}}{\tau^2}$ because $l_1^* \geq l_{sm} \geq x' \geq y'$. Thus the sum of all s_i for which $l_i \notin I_1$ or $p_i \geq 2/3$ is at most $\frac{9\dot{\Delta}^{-1/\tau}\ln \dot{\Delta}}{\tau}$. Since by the assumption, the sum of all s_i is at least $\frac{c \ln \dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau}$ then the sum of all s_i for which $l_i \in I_1$ and $p_i < 2/3$ has to be at least $\frac{2400 \ln \dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau}$. By the choice of l_1^*, each s_i for which $l_i \in I_1$ and $p_i \geq 2/3$ is at most $\frac{24 \ln \dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau}$ hence we must have at least 100 such l-values. We have shown that there are at least 100 l-values inside interval I_1.

We find the smallest and the largest l-values inside I_1.

$$\begin{aligned} l_{sm}^{(1)} &= \min_{i\in\{1,2,\ldots,\tau\}}\{l_i : l_i \in I_1\} \\ l_{lg}^{(1)} &= \max_{i\in\{1,2,\ldots,\tau\}}\{l_i : l_i \in I_1\} \end{aligned}$$

We consider interval $\Gamma_2 = \Gamma_1 \setminus (l_{sm}^{(1)}, l_{lg}^{(1)})$ (we remove the interior of the interval $[l_{sm}^{(1)}, l_{lg}^{(1)}]$ keeping the endpoints). We know that we removed at least 98 l-values. Since the l-values have to work for any l^* we can now argue about the average value of a gap inside Γ_2 and locate a different value l_2^* in the remaining interval and identify 100 l-values close to l_2^*. But we need to make sure that $|l_1^* - l_2^*| \geq x' + y'$ since otherwise we would count the same l-values twice.

We extend the interval I_1 to $I_1' = [l_1^* - (x' + y'), l_1^* + (x' + y')]$ and we find the smallest l-value larger than any l-value inside I_1' (call it $l_{sm}'^{(1)}$) and the largest l-value smaller than any l-value inside I_1' (call it $l_{lg}'^{(1)}$). If both these values exist, we consider interval $\Gamma_2^* = \Gamma_1 \setminus (l_{sm}'^{(1)}, l_{lg}'^{(1)})$ (we remove the interior of the interval $[l'_{sm}, l'_{lg}]$ keeping the endpoints). If l'_{sm} does not exist (there is no such l-value), we define $\Gamma_2^* = \Gamma_1 \setminus [l_{sm}, l_{lg}'^{(1)})$ and symmetrically if $l_{lg}'^{(1)}$ does not exist we set $\Gamma_2^* = \Gamma_1 \setminus (l_{sm}'^{(1)}, l_{lg}]$.

Now we want to show that $\Gamma_2^* \geq |\Gamma_1| - 5(x'+y')$. It is because $|I_1'| = 2(x'+y')$ and by Equation (9), value of any gap is at most $x'+y'$ hence distance between $l_1^* - (x'+y')$ and l'_{sm} is at most $x'+y'$ (similarly between $l_1^* + (x'+y')$ and l'_{lg}). If l'_{sm} or l'_{lg} does not exist we remove additionally no more than $x'+y'$ because the smallest l-value that is most x' and the largest is at least $\log\dot{\Delta} - y'$. This shows that we remove the total area of at most $5(x'+y')$.

Now we consider the average value of a gap in Γ_2^*. We removed at least 98 l-values (because I_1^* contains I_1') and area of at most $5(x'+y') \leq 5\log\dot{\Delta}/\tau + 10\log\tau + 5\log 3$. Hence the average value of a gap in Γ_2^* is:

$$d_2 \geq \frac{\log\dot{\Delta} - 6(x'+y')}{\tau - 98} \geq \frac{\log\dot{\Delta}}{\tau}.$$

We pick a gap with value at least $\frac{\log\dot{\Delta}}{\tau}$ and find an integer l_2^* at least y larger than the closest smaller l-value and at least $x-1$ smaller than the closest larger l-value. Observe moreover that $|l_2^* - l_1^*| \geq x'+y'$ because l_2^* does not belong to the interior of interval I_1'. We define $I_2 = [l_2^* - y', l_2^* + x']$. Observe that I_1 and I_2 are disjoint (except possibly their endpoints). We can now argue that I_2 also contains 100 l-values similarly as I_1. And moreover at most one l-value can be shared between I_1 and I_2 (because the interiors of the intervals are disjoint). And now we extend I_2 to I_2^*, construct Γ_3 and Γ_3^* and repeat the whole procedure. In every iteration of this procedure we identify and remove at least 98 unique l-values hence it can last for at most $\tau/98$ iterations. But we remove at most $5\log\dot{\Delta}/\tau + 10\log\tau + 5\log 3$ area (out of the original at least $\log\dot{\Delta} - x' - y'$) per iteration. Since we assumed $\tau \leq \ln\dot{\Delta}/\log\log\dot{\Delta}$, then $5\log\dot{\Delta}/\tau + 10\log\tau + 5\log 3 \leq 10\log\dot{\Delta}/\tau$. This leads to contradiction since there are only τ l-values and in each iteration we find 98 unique ones.

Hence for any choice of $l_1, \ldots, l_\tau$ there exists l^* such that $\sum_{i=1}^{\tau} s_i < c\frac{\ln\dot{\Delta}}{\dot{\Delta}^{1/\tau}\tau}$. Thus by the union bound the algorithm needs to run for at least $\dot{\Delta}^{1/\tau}\tau/(c\ln\dot{\Delta})$ phases to accumulate the total probability of success of $1/2$. Knowing that each phase lasts for τ rounds the total number of steps needed is $\Omega(\Delta^{1/\tau}\tau^2/\ln\Delta)$ □

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 τ shrinks toward 1. Our lower bound will establish that this degradation 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).

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

Theorem 9. *Fix an error bound $\epsilon > 0$. It follows that algorithm RGB(ϵ) 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. Similarly to the analysis of the local broadcast algorithms, we consider only the case of $\tau \leq \log\Delta$ since for any larger τ we use the algorithm for $\tau = \log\Delta$. Take any station u and assume that some positive number of neighbors of u in E execute in parallel procedure FRLB(2). Then by Lemma 5 station u receives a message from some neighbor with probability at least $\frac{\ln\Delta}{4\Delta^{1/\tau}\tau}$. Note that the same number of neighbors of u have to execute both procedures Uniform of FRLB(2) and at least one of these neighbors has to be connected to u by a reliable link. This is true since after receiving the message, a station waits until a time slot that is a multiple of 2τ (Line 3 in the pseudocode). Hence we can treat each execution of FRLB(2) as a single *phase*.

Let o denote the originator of the message. Fix any tree $\mathcal{T}$ of shortest paths on graph G (e.g., BFS Tree) on edges from E (reliable) rooted at o. We would like to bound the progress of the message on tree $\mathcal{T}$. For any station u we can denote

by $p(u)$ the parent of u in tree T. For any station u we can define the earliest time step $T(u)$ in which $p(u)$ receives the message. We set $T(u) = \infty$ if the message does not reach $p(u)$. If $T(u) < \infty$ we consider $\lceil \ln(2n/\epsilon) \rceil \cdot \lceil 4\Delta^{1/\tau}\tau / \log \Delta \rceil$ phases that follow step $T(u)$. A phase is called *successful* if it succeeds in delivering the message to u and *unsuccessful* otherwise. Note that (assuming that $T(u) < \infty$) the probability that all phases are unsuccessful for any fixed u is at most

$$\left(1 - \frac{\ln \Delta}{4\Delta^{1/\tau}\tau}\right)^{\lceil \ln(2n/\epsilon) \rceil \cdot \lceil 4\Delta^{1/\tau}\tau / \log \Delta \rceil} \leq e^{-\lceil \ln(2n/\epsilon) \rceil} = \frac{\epsilon}{2n}. \tag{10}$$

Let us denote by S the event that $T(u) < \infty$ for all stations u. And by S_i we denote the event that $T(u) < \infty$ for all stations at distance at most i from the root in tree $\mathcal{T}$. If d_i denotes the number of stations at distance i from the root in tree $\mathcal{T}$ we get:

$$\begin{aligned}
\mathbf{Pr}[\, S \,] &= \mathbf{Pr}[\, S_D \,] \geq \mathbf{Pr}[\, S_D | S_{D-1} \,] \mathbf{Pr}[\, S_{D-1} \,] \\
&\geq \mathbf{Pr}[\, S_1 \,] \prod_{i=2}^{D} \mathbf{Pr}[\, S_i | S_{i-1} \,] \\
&\geq \prod_{i=1}^{D} \left(1 - \frac{\epsilon d_i}{2n}\right) && \text{by (10) and union bound} \\
&\geq 1 - \sum_{i=1}^{D} \frac{\epsilon d_i}{2n} && \text{by Lemma 3} \\
&= 1 - \frac{\epsilon}{2}.
\end{aligned}$$

If event S takes place, the message reaches all the nodes of the network. Clearly it can reach node u not necessarily from its parent $p(u)$ in tree $\mathcal{T}$, but this would only help in our analysis (it will cause the message to arrive at u faster). Now we want to bound the number of phases it takes for the message to traverse a path in the tree. Fix any station u and let $\mathcal{P} = (o, v_1, v_2, \ldots, v_{D'-1}, u)$ denote the path from o to u in tree T (note that $D' \leq D$). We denote by R_i the round in which v_i receives the message ($R_{D'}$ denotes the round in which u receives the message) and introduce random variables $\Delta_i = \max\{0, R_i - R_{i-1}\}$. Conditioning on event S, variables Δ_i are stochastically dominated by independent geometric random variables with success probability $\frac{\ln \Delta}{4\Delta^{1/\tau}\tau}$. We have D' such variables and the probability that sum T of them exceeds $L = 4(D' + \ln(2n/\epsilon)) \cdot \frac{7\Delta^{1/\tau}\tau}{\log \Delta} = \mathbf{E}[\, T \,] \cdot 4(1 + \ln(2n/\epsilon)/D')$ can be bounded using [13, Theorem 2.3]. Denote $\lambda = 4(1 + \ln(2n/\epsilon)/D')$ and observe that $(\lambda - 1)/2 \geq \ln \lambda$ is true since $\lambda > 4$. We get:

$$\begin{aligned}
\mathbf{Pr}[\, T \geq L \,] &= \mathbf{Pr}[\, T \geq \mathbf{E}[\, T \,] \cdot \lambda \,] \\
&\leq \frac{1}{\lambda} \cdot \left(1 - \frac{\ln \Delta}{4\Delta^{1/\tau}\tau}\right)^{(\lambda - 1 - \ln \lambda)\mathbf{E}[\, T \,]} \\
&\leq \frac{1}{\lambda} \left(1 - \frac{\ln \Delta}{4\Delta^{1/\tau}\tau}\right)^{\mathbf{E}[\, S \,](\lambda - 1)/2} \\
&\leq \frac{1}{\lambda} e^{-\frac{3}{2}D' - 2\ln(2n/\epsilon)} \leq \frac{\epsilon^2}{4n^2},
\end{aligned}$$

and by taking the union bound over all stations u, we get that with probability at least $1 - \epsilon^2/(4n)$ the message reaches all nodes within time $4(D + \ln(2n/\epsilon)) \cdot \frac{4\Delta^{1/\tau}\tau}{\log \Delta}$, conditioned on S. Since S takes place with probability at least $1 - \epsilon/2$ and since each phase takes 2τ time steps, this shows that the algorithm works within time $8(D + \ln(2n/\epsilon)) \cdot \frac{4\Delta^{1/\tau}\tau^2}{\log \Delta}$ with probability at least $(1 - \epsilon/2)(1 - \epsilon^2/(4n)) \geq 1 - \epsilon$. □

4.2. Lower bound

The global broadcast lower bound of $\Omega(D \log(n/D))$, proved by Kushilevitz and Mansour [17] 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].

[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 Δ as well for our purposes here.

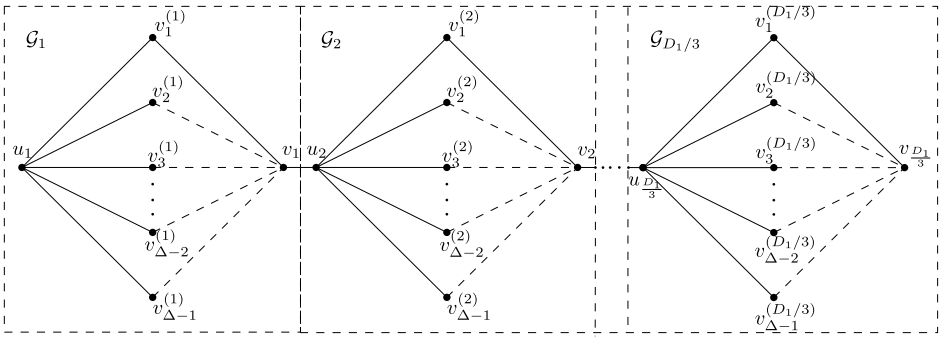

Fig. 2. A graph used in proof of Theorem 10.

For smaller τ, 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 τ is 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 τ, 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 Δ and fading adversary with stability τ. 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. We assume first that D is divisible by 3 (if it is not we can decrease D by one or two nodes to make it divisible by 3, without impacting the asymptotic bounds). We construct the dual graph G, G' by connecting together $D/3$ gadgets, $\mathcal{G}_1, \mathcal{G}_2, \ldots, \mathcal{G}_{D/3}$, as shown in Fig. 2. In particular, each gadget $\mathcal{G}_i$ is the same graph structure used to prove our local broadcast lower bound. Formally, for each $i = 1, 2, \ldots, D/3$, gadget $\mathcal{G}_i$ is a dual graph $G_i = (V_i, E_i)$, $G_i' = (V_i, E_i')$ where $E_i = \{\{u_i, v_j^{(i)}\} : j = 1, 2, \ldots, \Delta-1\} \cup \{\{v_1^{(i)}, v_i\}\}$, $E_i' = E_i \cup \{\{v_j^{(i)}, v_i\}, j = 2, 3, \ldots, \Delta-1\}$. We denote the set of edges connecting the gadgets by $E_c = \{\{v_i, u_{i+1}\} : i = 1, 2, \ldots, D/3-1\}$. Finally we can define the total set of nodes and edges in the complete dual graph $G = (V, E)$ and $G' = (V, E')$ as follows: $V = \bigcup_{i=1}^{D/3} V_i$, $E = E_c \cup \bigcup_{i=1}^{D/3} E_i$, $E' = \bigcup_{i=1}^{D/3} E_i'$. We will show statement 1 by applying Theorem 8 to each gadget, statement 2 can be shown using the same proof by applying Theorem 7.

We bound the dissemination of a broadcast message in this graph originating at node u_1. We can view the progression of the message through the chain of gadgets $\mathcal{G}_1, \mathcal{G}_2, \ldots, \mathcal{G}_{D/3}$ as a sequence of local broadcasts. When the message arrives at a node u_i, it is propagated to nodes $v_1^{(i)}, v_2^{(i)}, \ldots, v_{\Delta-1}^{(i)}$ and at this step delivering the message to v_i is exactly the local broadcast problem considered in Theorem 8. In this theorem we constructed a sequence of distributions that yields a high running time. The distribution changes every exactly τ steps i.e., we have a distribution $\mathcal{D}_k$ for steps $1 + (k-1)\tau, 2 + (k-1)\tau, \ldots, k\tau$. We cannot immediately apply the result for local broadcast because the adversary might not be allowed to change the distribution immediately when the message arrives in a gadget. Moreover in the global broadcast problem, stations are allowed to delay the transmissions for some number of steps. We can easily solve this problem by keeping the "first" distribution $\mathcal{D}_1$ in each gadget until the message reaches the gadget, at which point the adversary can start the sequence of changes specified by the local broadcast lower bound.

More precisely, we denote sequence $p_1, p_2, \ldots$ of probabilities used by algorithm $\mathcal{A}$ and we denote subsequences $P_k = (p_{1+(k-1)\tau}, p_{2+(k-1)\tau}, \ldots, p_{k\tau})$. We want to use distributions $\mathcal{D}_1, \mathcal{D}_2, \ldots$ from Theorem 8 in such a way that if i is the furthest gadget reached by the message and its nodes are in phase k (i.e., are using probabilities from sequence P_k) then distribution in gadget $\mathcal{G}_i$ is $\mathcal{D}_k$. If the message has not reached the gadget yet, distribution in the gadget is $\mathcal{D}_1$. Finally if the message already reached node v_i in gadget $\mathcal{G}_i$ for any i we do not change the distribution in this gadget any more. We need to show that with this construction we do not need to change the distribution more frequently than once per τ steps. This is true because we only change the distribution in the furthest gadget (call it $\mathcal{G}_i$) reached by the message and moreover we change it from $\mathcal{D}_k$ to $\mathcal{D}_{k+1}$ only after the stations $v_1^{(i)}, v_2^{(i)}, \ldots, v_{\Delta-1}^{(i)}$ have finished transmitting with probabilities P_k, which takes at least τ steps (it might take more because stations might delay transmitting with probabilities P_1).

Let us define random variables X_i for $i = 1, 2, \ldots, D/3$ as the number of time steps it takes for nodes $v_1^{(i)}, v_2^{(i)}, \ldots, v_{\Delta-1}^{(i)}$ to deliver the message to v_i. More precisely it is the number of steps between the first step when the stations $v_1^{(i)}, v_2^{(i)}, \ldots, v_{\Delta-1}^{(i)}$ transmit with probability p_1 and the step in which the first successful transmission delivers the message to v_i (including the step of the successful transmission). Note that the steps during which the stations $v_1^{(i)}, v_2^{(i)}, \ldots, v_{\Delta-1}^{(i)}$ delay transmitting until the beginning of the next probability cycle are not counted in variable X_i. The steps counted by

variable X_i can be seen as local broadcast. By Theorem 8 we have that $\mathbf{Pr}\left[X_i \leq \dot{\Delta}^{1/\tau}\tau^2/(c \ln \dot{\Delta}) \right] \leq 1/2$ for some constant $c > 1$. Moreover variables X_i are independent because choices of the stations in each gadget are independent. We can define binary variables Y_i taking value 1 if $X_i > \dot{\Delta}^{1/\tau}\tau^2/(c \ln \dot{\Delta})$ and 0 otherwise. We can use Chernoff bound to lower bound $Y = \sum_{i=1}^{D/3} Y_i$. By the assumptions on D we have that $\mathbf{E}[\, Y \,] \geq 4$ (because $\mathbf{E}[\, Y_i \,] \geq 1/2$), hence

$$\mathbf{Pr}[\, Y \leq \mathbf{E}[\, Y \,]/2 \,] \leq e^{-\mathbf{E}[\, Y \,]/8} \leq \frac{2}{3}$$

Observe that $Y \cdot \dot{\Delta}^{1/\tau}\tau^2/(c \ln \dot{\Delta})$ lower bounds the time of the global broadcast. This shows that the global broadcast needs $\Omega(D\Delta^{1/\tau}\tau^2/\ln \Delta)$ steps with probability at least 1/3.

If D is not divisible by 3 we construct our graph with diameter $3\lfloor D/3 \rfloor$ and attach a path of $D - 3\lfloor D/3 \rfloor$ (one or two) vertices to node $v_{D_1/3}$. This cannot decrease the time of broadcast hence we get the bound $\Omega((D-2)\Delta^{1/\tau}\tau^2/\log \Delta) = \Omega(D\Delta^{1/\tau}\tau^2/\log \Delta)$. □

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 τ 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 τ 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 $\Omega(\sqrt{n}/l^2)$ 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 expectation 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 $\log^2 \Delta$ probabilities requires time $\Omega(\sqrt{\Delta}/\log^2 \Delta)$ if arbitrary correlations are permitted.

Proposition 1. *Any uniform local broadcast algorithm that repeats a procedure consisting of l probabilities requires expected time $\Omega(\sqrt{\Delta}/l^2)$ 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 first star has arms $v_1, v_2, \ldots, v_\Delta$ and center at u. In the first 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

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

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. This means that depending on the probabilities used by the algorithm, either all the unreliable connections are active or none of them. 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 $\Omega(\sqrt{\Delta}/l^2)$. □

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 τ 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 τ steps.

Now we present algorithm RLBC (Robust Local Broadcast with Correlations). We first show that the algorithm works under (l, τ)-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 τ 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 τ and l. Restricted fading adversary can change the distribution arbitrarily once every at most τ 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/e_1)
11         Uniform(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. Assume that $\tau \leq \lceil \log_{2e}\Delta/2\rceil$ and note that in this case $\bar{\tau} = \tau$. In the opposite case we use the algorithm for $\tau = \lceil \log_{2e}\Delta/2\rceil$ which works also for any larger τ. Denote $k = \lfloor \Delta^{1/(\tau-2a)}\rfloor$, $l = k\tau/2$ and observe that for $\tau \geq 1000$ we have $a > 200$ and $\tau - 2a \geq \tau/2$ and $k \geq 2$. We divide the time into intervals of length τ, called *cycles*. In each interval algorithm RLBC repeats the same probabilities. In the first $\tau - 2a$ steps of the cycle it uses probabilities $p_i = k^{-i}$ for $i = 1, 2, \ldots, \tau - 2a$, in the next $2a$ steps it uses probabilities $1/e_1$ and $1/e_2$. We take two consecutive cycles and note that in each such a pair of cycles we can find τ consecutive steps in which the distribution over the unreliable edges is the same (since global changes can happen at most once every τ steps) and moreover the algorithm uses all the probabilities from a cycle. Let us call a sequence of these steps $T = [t, t+1, \ldots, t+\tau-1]$. Note that in this sequence we have either one full procedure RLB$(1, \tau - 2a)$ or parts of two procedures RLB$(1, \tau - 2a)$ (call them R_1 and R_2). In the second case sequence T contains some suffix of R_1 and some prefix of R_2. Connect these steps together into a procedure R, which contains all steps of procedure RLB$(1, \tau - 2a)$ executed in a possibly different order. Fix a receiver v and assume that at least one reliable neighbor of v tries to transmit a message to v. We want to show that in each such a pair of cycles v receives the message independently with probability at least $p_s = \frac{1}{8ek}$.

We know by the definition of the adversary that the effective degree cannot change by too much between steps in the same cycle: $|d_{t+i}(v) - d_{t+i+1}(v)| \leq l$ for any $i = 0, 1, \ldots, \tau - 2$. We can consider two cases depending on the effective degree in the first considered step t_1:

Case 1: $d_t(v) \geq 2l^2$

Here we want to show that procedure R is successful with probability at least p_s. Observe that here since $l \geq \tau$ we have $d_t(v) \geq d_t(v)/2 + l^2 \geq d_t(v)/2 + l\tau$. Thus for any $i \in \{1, 2 \ldots, \tau - 1\}$ we have $d_{t+i}(v) \geq d_t(v) - l\tau \geq d_t(v)/2$ and $d_{t+i}(v) \leq d_t(v) + l\tau \leq 2d_t(v)$ thus the effective degree in the whole considered sequence of steps can change by a factor of at most 2. Recall from the definition of RLB that it uses probabilities $p_i = k^{-i}$. Consider the smallest i such that $1/p_i \geq 2d_t(v)$ by the minimality of i we have that $1/(kp_i) \leq 2d_t(v)$. Probability p_i is used in some step of sequence T. Call this step t^*. We have:

$$1/p_i \geq 2d_t(v) \geq d_{t^*}(v) \geq d_t(v)/2 \geq 1/(4kp_i).$$

Thus by Lemmas 1 and 2:

$$\mathbf{Pr}\left[R^{(v)}_{t^*} \right] \geq \min\left\{ (2e)^{-1}, \frac{(2e)^{-1/(4k)}}{4k} \right\} \geq \frac{1}{8ek} = p_s.$$

Case 2: $d_t(v) < 2l^2$

Here we want to show that a successful transmission occurs with probability at least p_s in one of the $2a$ additional steps (see lines $7 - 11$ of the pseudocode).

Note that since $d_t(v) < 2l^2$ then for any $i \in \{1, \ldots, \tau - 1\}$ we have $d_{t+i}(v) \leq d_t(v) + l\tau \leq 4l^2$. Pick two consecutive steps $t+i, t+1$ such that in step $t+i$ the algorithm uses probability $1/e_1$ and in $t+i+1$ it uses $1/e_2$. Note that in the considered sequence we have at least $a - 1$ such pairs.

Case 2.1: $d_{t+i}(v) \leq 2l$ Here the probability is $1/e_1$ and the degree is within interval $[1, 2l]$ hence we have that:

$$\frac{\tau}{a} = \frac{2l}{e_1} \geq \frac{d_{t_i}}{e_1} \geq \frac{1}{e_1}.$$

By Lemmas 1 2:

$$\mathbf{Pr}\left[R^{(v)}_{t+i} \right] \geq \min\left\{ \frac{2l}{e_1} e^{-2l/e_1}, 1/e_1 e^{-1/e_1} \right\} \geq \min\left\{ \frac{\ln \tau}{\tau}, \frac{1}{eka} \right\} \geq \frac{1}{eka}.$$

Case 2.2: $4l^2 \geq d_{t+i}(v) > 2l$ Note that in this case $d_{t+i+1}(v) \in [l, 4l^2]$ and the probability used in this step is $1/e_2$ hence:

$$\frac{\tau}{a} = \frac{4l^2}{e_2} \geq \frac{d_{i+1}(v)}{e_2} \geq \frac{2l}{e_2} = \frac{1}{e_1},$$

and we can use Lemmas 1 and 2 exactly the same as in the previous case and obtain $\mathbf{Pr}\left[R^{(v)}_{t+i+1} \right] \geq \frac{1}{eka}$.

In each pair the stations are making independent choices hence the probability of failure in all the pairs is by Lemma 3 at most:

$$\left(1 - \frac{1}{eka}\right)^{a-1} \leq 1 - \frac{a-1}{eka} + \binom{a-1}{2} \frac{1}{e^2 k^2 a^2} \leq 1 - \frac{1}{2ek},$$

where in the last inequality we used the fact that $a > 20$. Thus also in this case with probability at least $1/(2ek) \geq p_s$ node v receives a message during this cycle.

The two considered cases showed that any full two cycles deliver the message with probability at least p_s. If we perform at least $2r = 2\lceil \ln(1/\epsilon)/p_s \rceil = O(\Delta^{1/\tau})$ cycles then the probability that v does not receive a message is at most $(1 - p_s)^{\ln(1/\epsilon)/p_s} \leq \epsilon$. □

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. Fix any receiver v. We know that RLBC$(8e\lceil \ln(1/\epsilon)\Delta^{1/\tau} \rceil, \tau)$ solves local broadcast in the presence of $(l\tau, \tau)$-deterministic adversary. But in the case with arbitrary correlations we can still bound the probability that the degree of v does not change too much. Take any two consecutive steps $t, t+1$. We have by Markov Inequality:

$$\mathbf{Pr}[\, |d_t(v) - d_{t+1}(v)| > 2\tau l \,] \leq 1/(2\tau)$$

If we pick τ steps like in the proof of Theorem 11 then by the union bound with probability at least $1/2$ in each of these steps the degree changes by at most $2l\tau$. From now on we can use the same analysis as in Theorem 11 and we obtain only a constant slowdown compared to the case with deterministic adversary. Hence RLBC$(16e\lceil \ln(1/\epsilon)\Delta^{1/\tau} \rceil, \tau)$ solves local broadcast with restricted fading adversary with probability at least $1 - \epsilon$. □

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

Many thanks to William Kuszmaul and Zachary Newman and to the anonymous reviewers for their very helpful comments on earlier versions of this manuscript.

Appendix. Proof of Lemma 5

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

We want to show that if some of the neighbors of v belong to B and execute the procedure $\texttt{Uniform}(\tau, p_1, \ldots, p_\tau)$ twice, then v receives some message with probability at least $\log \Delta/(4\Delta^{1/\tau}\tau)$. Since we execute `Uniform` twice, we have a total of 2τ consecutive time slots out of which, by the definition of our model, at least τ consecutive slots have the same distribution of the edges in $E' \setminus E$ 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 τ 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 define $q_1 = \mathbf{Pr}\left[\, 1 \leq d_t(v) \leq \Delta^{1/\tau} \,\right]$, and for each i from 2 to τ we define $q_i = \mathbf{Pr}\left[\, \Delta^{(i-1)/\tau} < d_t(v) \leq \Delta^{i/\tau} \,\right]$, for any $t \in T$. Observe that the edge sets are chosen independently from the same distribution for $t \in T$ hence probabilities q_i do not depend on t during the considered τ rounds. Moreover since $v \in R$, then v is connected via a reliable edge to at least one node in B. This means that $\mathbf{Pr}[\, d_t(v) = 0 \,] = 0$, thus:

$$\sum_{i=1}^{\tau} q_i = 1. \tag{11}$$

We would like to lower bound the probability that v receives a message in step t_i for $i = 1, 2, \ldots, \tau$:

$$\mathbf{Pr}\left[\, R_{t_i}^{(v)} \,\right] \geq \frac{q_i \log_{2e} \Delta}{2\Delta^{1/\tau}\tau}. \tag{12}$$

In t_i-th slot each station with a message transmits independently with probability is $p_i = \Delta^{-i/\tau} \cdot \log_{2e} \Delta/\tau$ and the transmission choices done by the stations are independent from the choice of edges E_{t_i} active in round t_i. Let Q_i denote the event that $\Delta^{(i-1)/\tau} < d_{t_i}(v) \leq \Delta^{i/\tau}$. We have $p_i \leq 1/2$ hence we can use Lemmas 1 and 2 and get for $i > 1$:

$$\begin{aligned}
\mathbf{Pr}\left[\, R_{t_i}^{(v)} \,\right] &\geq \mathbf{Pr}\left[\, R_{t_i}^{(v)} \mid Q_i \,\right] \cdot \mathbf{Pr}[\, Q_i \,] \\
&\geq q_i \cdot \min\left\{ \mathbf{Pr}\left[\, R_{t_i}^{(v)} \mid d_{t_i}(v) = \left\lfloor \Delta^{(i-1)/\tau} \right\rfloor + 1 \,\right], \mathbf{Pr}\left[\, R_{t_i}^{(v)} \mid d_{t_i}(v) = \Delta^{-i/\tau} \,\right] \right\} \\
&\geq q_i \cdot \min\left\{ \mathbf{Pr}\left[\, R_{t_i}^{(v)} \mid d_{t_i}(v) = \Delta^{(i-1)/\tau} \,\right], \mathbf{Pr}\left[\, R_{t_i}^{(v)} \mid d_{t_i}(v) = \Delta^{-i/\tau} \,\right] \right\} \\
&\geq q_i \cdot \min\left\{ \frac{\Delta^{-1/\tau} \log_{2e} \Delta/\tau}{(2e)^{\Delta^{-1/\tau} \log_{2e} \Delta/\tau}}, \frac{\log_{2e} \Delta/\tau}{(2e)^{\log_{2e} \Delta/\tau}} \right\}.
\end{aligned}$$

Note that $\Delta^{-1/\tau} \log_{2e} \Delta/\tau \leq 1$ hence:

$$\mathbf{Pr}\left[\, R_{t_i}^{(v)} \,\right] \geq \frac{q_i \log_{2e} \Delta}{2e\Delta^{1/\tau}\tau}.$$

For $i = 1$ we get:

$$\begin{aligned}
\mathbf{Pr}\left[\, R_{t_1}^{(v)} \,\right] &\geq \mathbf{Pr}\left[\, R_{t_1}^{(v)} \mid Q_1 \,\right] \cdot \mathbf{Pr}[\, Q_1 \,] \\
&\geq q_1 \cdot \min\left\{ \frac{\Delta^{-1/\tau} \log_{2e} \Delta/\tau}{(2e)^{\Delta^{-1/\tau} \log_{2e} \Delta/\tau}}, \frac{\log_{2e} \Delta/\tau}{(2e)^{\log_{2e} \Delta/\tau}} \right\} \geq \frac{q_1 \log_{2e} \Delta}{2\Delta^{1/\tau}\tau}.
\end{aligned}$$

Since the edge sets are chosen independently in each step and the choices of the stations are also independent we have:

$$
\begin{aligned}
\mathbf{Pr}\left[\bigwedge_{i=1}^{\tau}\left(\neg R_{t_i}^{(v)}\right)\right] &= \prod_{i=1}^{\tau}\mathbf{Pr}\left[\neg R_{t_i}^{(v)}\right] && \text{by independence}\\
&\le \prod_{i=1}^{\tau}\left(1-\frac{q_i\log_{2e}\Delta}{2e\Delta^{1/\tau}\tau}\right) && \text{by Equation (12)}\\
&\le 1-\sum_{i=1}^{\tau}\frac{q_i\log_{2e}\Delta}{2e\Delta^{1/\tau}\tau}+\sum_{1\le i<j\le n}\frac{q_iq_j\log_{2e}^2\Delta}{4e^2\Delta^{2/\tau}\tau^2} && \text{by Lemma 3}\\
&\le 1-\frac{\log_{2e}\Delta}{2e\Delta^{1/\tau}\tau}\sum_{i=1}^{\tau}q_i+\frac{\log_{2e}^2\Delta}{4e^2\Delta^{2/\tau}\tau^2}\left(\sum_{i=1}^{n}q_i\right)^2\\
&= 1-\frac{\log_{2e}\Delta}{2e\Delta^{1/\tau}\tau}+\frac{\log_{2e}^2\Delta}{4e^2\Delta^{2/\tau}\tau^2}, && \text{by Equation (11)}\\
&\le 1-\frac{\log_{2e}\Delta}{4e^2\Delta^{1/\tau}\tau},
\end{aligned}
$$

where the last inequality is true since if we denote $\tau=(\log_{2e}\Delta)/\alpha$ (for $\alpha\ge 1$) then we have $\Delta^{1/\tau}\tau=(2e)^{\alpha}\log\Delta/(2\alpha)\ge\log\Delta$ hence $\frac{\log_{2e}^2\Delta}{4e^2\Delta^{2/\tau}\tau^2}\le\frac{\log_{2e}\Delta}{4e^2\Delta^{1/\tau}\tau}$. This completes proof of 2. To prove 1 we observe that if we execute the procedure for $2\tau\lceil\ln(1/\epsilon)\rceil\cdot\lceil 4e^2\cdot\Delta^{1/\tau}\tau/\log_{2e}\Delta\rceil$ time steps we have at least $\lceil\ln(1/\epsilon)\rceil\cdot\lceil 4e^2\cdot\Delta^{1/\tau}\tau/\log_{2e}\Delta\rceil$ sequences of τ 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 $1-\log_{2e}\Delta/(4e^2\Delta^{1/\tau}\tau)$ hence the probability that all the procedures fail is at most:

$$
\left(1-\frac{\tau\log_{2e}\Delta}{4e^2\cdot\Delta^{1/\tau}\tau}\right)^{\lceil\ln(1/\epsilon)\rceil\cdot\lceil 4e^2\cdot\Delta^{1/\tau}\tau/\log_{2e}\Delta\rceil}\le e^{-\lceil\ln(1/\epsilon)\rceil}<\epsilon\quad\square
$$

References

[1] N. Alon, A. Bar-Noy, N. Linial, D. Peleg, A lower bound for radio broadcast, J. Comput. Syst. Sci. 43 (2) (1991) 290–298.
[2] Reuven Bar-Yehuda, Oded Goldreich, 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) (1992) 104–126, https://doi.org/10.1016/0022-0000(92)90042-H.
[3] Keren Censor-Hillel, Seth Gilbert, Fabian Kuhn, Nancy Lynch, Calvin Newport, Structuring unreliable radio networks, Distrib. Comput. 27 (1) (2014) 1–19.
[4] Andrea E.F. Clementi, Angelo Monti, Riccardo Silvestri, Round robin is optimal for fault-tolerant broadcasting on wireless networks, J. Parallel Distrib. Comput. 64 (1) (2004) 89–96.
[5] Artur Czumaj, Wojciech Rytter, Broadcasting algorithms in radio networks with unknown topology, J. Algorithms 60 (2) (2006) 115–143, https://doi.org/10.1016/j.jalgor.2004.08.001.
[6] 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.
[7] Mohsen Ghaffari, Bernhard Haeupler, Nancy Lynch, Calvin Newport, Bounds on contention management in radio networks, in: Marcos K. Aguilera (Ed.), Distributed Computing: 26th International Symposium, DISC 2012, Salvador, Brazil, October, 2012, in: Lecture Notes in Computer Science, vol. 7611, Springer, 2012, pp. 223–237.
[8] Mohsen Ghaffari, Erez Kantor, Nancy Lynch, Calvin Newport, Multi-message broadcast with abstract MAC layers and unreliable links, in: Proceedings of the 33rd Annual ACM Symposium on Principles of Distributed Computing, PODC'14, Paris, France, July 2014, pp. 56–65.
[9] Mohsen Ghaffari, Nancy Lynch, 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, Montreal, Canada, July 2013, pp. 345–354.
[10] Mohsen Ghaffari, Calvin Newport, A leader election in unreliable radio networks, in: Proceedings of the International Colloquium on Automata, Languages, and Programming, ICALP, 2016.
[11] Seth Gilbert, Nancy A. Lynch, Calvin Newport, 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, 2018, pp. 48:1–48:3, 2018.
[12] Seth Gilbert, Nancy A. Lynch, Calvin Newport, Dominik Pajak, On simple back-off in unreliable radio networks, in: 22nd International Conference on Principles of Distributed Systems, OPODIS 2018, December 17–19, 2018, Hong Kong, China, 2018, pp. 27:1–27:17.
[13] Svante Janson, Tail bounds for sums of geometric and exponential variables, Stat. Probab. Lett. 135 (2018) 1–6.
[14] Dariusz R. Kowalski, Andrzej Pelc, Broadcasting in undirected ad hoc radio networks, Distrib. Comput. 18 (1) (2005) 43–57, https://doi.org/10.1007/s00446-005-0126-7.
[15] Fabian Kuhn, Nancy Lynch, 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.
[16] Fabian Kuhn, Nancy Lynch, Calvin Newport, Rotem Oshman, Andrea Richa, Broadcasting in unreliable radio networks, in: Proceedings of the 29th ACM Symposium on Principles of Distributed Computing, PODC, Zurich, Switzerland, July 2010, pp. 336–345.
[17] E. Kushilevitz, Y. Mansour, An $\Omega(D\log(N/D))$ lower bound for broadcast in radio networks, SIAM J. Comput. 27 (3) (1998) 702–712.
[18] Nancy Lynch, Calvin Newport, A (truly) local broadcast layer for unreliable radio networks, in: Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC, 2015.
[19] Minoru Nakagami, The m-distribution—a general formula of intensity distribution of rapid fading, in: Statistical Methods in Radio Wave Propagation, Elsevier, 1960, pp. 3–36.
[20] Calvin Newport, Lower bounds for radio networks made easy, in: Proceedings of the International Symposium on Distributed Computing, DISC, 2014.

[21] John G. Proakis, Digital Communications, McGraw-Hill Inc., New York, 1995.
[22] Stephen O. Rice, Mathematical analysis of random noise, Bell Syst. Tech. J. 24 (1) (1945) 46–156.
[23] Mike Willis, Propagation tutorial: fading, http://www.mike-willis.com/Tutorial/PF15.htm, 2007, May 5.