The Unreasonable Effectiveness of Decay-Based Broadcasting in Radio Networks

Calvin Newport
Georgetown University

ADGA : October 15th, 2018

The Radio Network Model

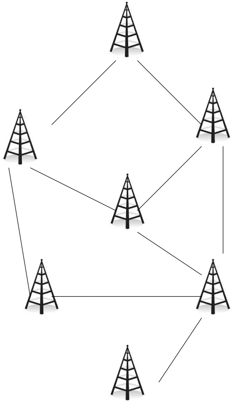

Definition

- Network topology described by undirected graph $G = (V, E)$

The Radio Network Model

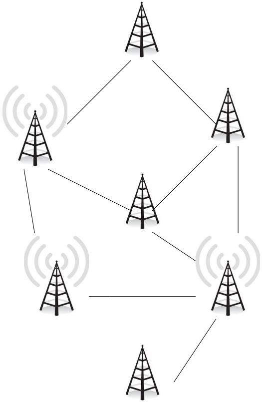

Definition

- Network topology described by undirected graph $G = (V, E)$
- Time is divided into synchronous rounds

The Radio Network Model

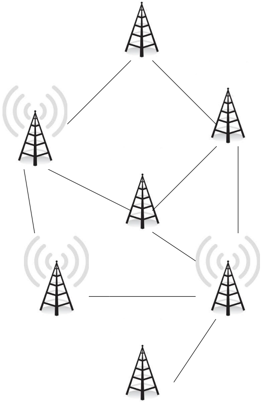

Definition

- Network topology described by undirected graph $G = (V, E)$
- Time is divided into synchronous rounds
- In each round, each station can `transmit` or `listen`

The Radio Network Model

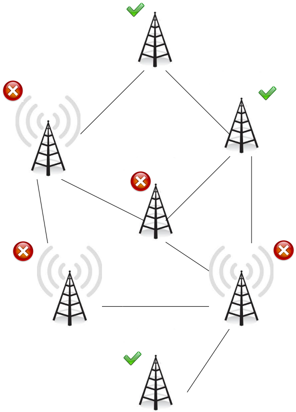

Definition

- Network topology described by undirected graph $G = (V, E)$
- Time is divided into synchronous rounds
- In each round, each station can `transmit` or `listen`
- Station $u \in V$ successfully receives message m if station u `listens`, **exactly one** neighbor `transmits`, and it transmits m.

Backstory

1973: ALOHA

Norman Abramson publishes on ALOHA, one of the very first computer networks. It happens to be a multihop radio network (limited to a star topology).

THE ALOHA SYSTEM

by

Norman Abramson
University of Hawaii

ABSTRACT

This report provides a status report and description of THE ALOHA SYSTEM research project at the University of Hawaii. THE ALOHA SYSTEM involves the analysis and construction of advanced methods of random access communications in large computer-communication systems.

Backstory

1973 to 1985: Systems Research

The decade that followed the development of the ALOHA network generated much systems research on multihop radio networks.

The DARPA Packet Radio Network Protocols

JOHN JUBIN AND JANET D. TORNOW, ASSOCIATE, IEEE

Invited Paper

In this paper we describe the current state of the DARPA packet radio network. Fully automated algorithms and protocols to organize, control, maintain, and move traffic through the packet radio network have been designed, implemented, and tested. By means of protocols, networks of about 50 packet radios with some degree of nodal mobility can be organized and maintained under a fully distributed mode of control. We have described the algorithms and illustrated how the PRNET provides highly reliable network transport and datagram service, by dynamically determining optimal routes, effectively controlling congestion, and fairly allocating the channel in the face of changing link conditions, mobility, and varying traffic loads.

scribe the algorithms used to route a packet through the packet radio communications subnet. In Section V, we examine the protocols for transmitting packets. In Section VI, we describe some of the hardware capabilities of the packet radio that strongly influence the design and characteristics of the PRNET protocols. We conclude by looking briefly at some applications of packet radio networks and by summarizing the state of the current technology.

Backstory

1985: Formalizing the Model

Chlamtac and Kutten described these radio networks with a formal abstraction suitable for algorithm and complexity analysis. This is the first appearance of the **radio network model** and the first formal analysis of the one-to-all **broadcast** problem.

1240 IEEE TRANSACTIONS ON COMMUNICATIONS, VOL. COM-33, NO. 12, DECEMBER

On Broadcasting in Radio Networks—Problem Analysis and Protocol Design

IMRICH CHLAMTAC AND SHAY KUTTEN

Abstract—In this paper we develop a graph-oriented model for dealing with broadcasting in radio networks. Using this model, optimality in broadcasting protocols is defined, and it is shown that the problem of finding an optimal protocol is NP-hard. A polynomial time algorithm is proposed under which a channel is assigned to nodes from global, multiple-source broadcasting considerations. In particular, nodes participating in the broadcast do not interfere with each other's transmissions, but otherwise simultaneous channel reuse is permitted. Protocol implementations of this approach by frequency division and by time division are given. It is shown that, using these protocols, bounded delay for broadcasted messages can be guaranteed.

I. INTRODUCTION

BROADCASTING a message to all network nodes is an important activity in computer networks [2], [10], [11].

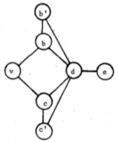

Fig. 1. Alternative broadcast routing procedures.

Backstory

1987: Solving Broadcast

Chlamtac and Weinstein describe a polynomial time centralized algorithm that creates a **broadcast schedule** of length $\mathbf{O}(\mathbf{D} \log^2 \mathbf{n})$ for diameter D and network size n. They discuss how to make solution distributed given a special control channel.

The Wave Expansion Approach to Broadcasting in Multihop Radio Networks

Imrich Chlamtac, *Senior Member, IEEE*, and Orly Weinstein

***Abstract*—In this paper we propose an algorithm for efficient communication between neighbors in multihop radio networks. The algorithm guarantees a bound on the transmission efficiency in a radio channel for arbitrary topology. The algorithm can be embedded in protocols for solving basic network problems such as broadcast, multicast, leader election, or finding shortest paths. We then specifically address the problem of bounded-time broadcasting utilizing the proposed algorithm. In the presented polynomial solution, we view the process of spreading information over the network as the expansion of a wave caused by a point of disturbance. The broadcast is originated at a source node, and is accomplished in repeated transmission periods, emulating a wave progressing away from the source. Using the proposed algorithm, a subset of potential transmitters is selected in each period so that a tightly bound proportion of potential receivers receive the transmission without collisions, guaranteeing a high level of spatial-reuse in the broadcast process. The spatial-reuse and the collision-free transmission properties of the algorithm, allow us to develop a broadcasting protocol with bounded delays shown to be better than in other currently known solutions. Specifically, the presented protocol gives a bound of $r \ln^2 N/r$ in a network consisting of N nodes with radius r. This result is at most logarithmically worse that the optimum given by at least r.**

also be obtained by multihop oriented routing protocols [15], [16]. Among broadcast based solutions, in slotted ALOHA type networks broadcasting was proposed by "propagating" the transmission along the frontier dividing the network into nodes which have and have not yet received a message [10]. The objective in [10] was to maximize throughput in each slot, assuming the existence of a collision-free reservation channel in dynamic topologies. Additional probabilistic, randomized protocols for broadcasting were given in [13], [14]. To obtain bounded delays, several collision-free broadcasting protocols have been proposed in the past [5], [6], [11]. A collision-free protocol concentrating on the spatial-reuse aspect of the broadcasting problem was first presented in [6]. The algorithm given in [6] allows several stations to transmit simultaneously without collisions by taking advantage of their distribution in the network. Specifically, it allows multiple nodes to be assigned the same transmission phase, while assuring that no neighboring nodes, or nodes with a common neighbor, will have the same phase. The

Backstory

1987: Solving Distributed Broadcast

Bar-Yehuda, Goldreich and Itai publish a PODC paper that describes and analyzes a simple randomized distributed algorithm that solves broadcast with probability at least $1 - \epsilon$ in $\mathbf{O}\left(\left(\mathbf{D} + \log\left(\mathbf{n}/\epsilon\right)\right)\log\mathbf{n}\right)$ rounds.

On the Time-Complexity of Broadcast in Radio Networks: An Exponential Gap Between Determinism and Randomization

(Extended Abstract)

Reuven Bar-Yehuda *Oded Goldreich* *Alon Itai*
Department of Computer Science
Technion - Israel Institute of Technology
Haifa 32000, ISRAEL

ABSTRACT

The time-complexity of deterministic and randomized protocols for achieving broadcast (distributing a message from a source to all other nodes) in arbitrary multi-hop radio networks is investigated. In these networks communication takes place in synchronous time-slots. A processor receives a message at a certain time-slot if exactly one if its neighbors transmits at that time-slot.

1. INTRODUCTION

Our model of radio communication consists of an arbitrary multi-hop (undirected) network, with processors communicating in synchronous time-slots subject to the following rules. In each time-slot, each processor acts either as a *transmitter* or as a *receiver*. A processor acting as a receiver is said to receive a message in time-slot t if exactly one of its neighbors transmits in time-slot t. The message received is the

Backstory

The Introduction of Decay

The key to Bar-Yehuda, Goldreich and Itai's surprisingly effective broadcast algorithm was a simple distributed strategy they called **Decay**.

We now present a precise description of the procedure as executed by each processor. By $a \in_R S$ we mean that a is randomly selected from the set A with uniform probability distribution.

procedure *Decay* (k, m);
repeat at most k times
    send m to all neighbors;
    flip *coin* $\in_R \{0,1\}$
until *coin* $= 0$.

The Original Decay Strategy

Decay(k, m)

Repeat k times:

    `transmit` m

    $x \leftarrow 0$ or 1 with equal probability

    if $x = 0$ **break**

The Original Decay Strategy

Decay(k, m)

Repeat k times:

    `transmit` m

    $x \leftarrow 0$ or 1 with equal probability

    if $x = 0$ **break**

Intuition: As we *decay* a set of sending stations down to 0 there is a good chance that we first arrive at a round with exactly 1 sending station left active.

Analyzing Decay

Theorem 1.

Let y be a vertex of G. Assume $d \geq 2$ neighbors of y execute *Decay*(k, m) during the interval $[0, k)$ starting at time 0. Then $P(k, d)$, the probability that y receives a message by time k, satisfies:

1. $\lim_{k \to \infty} P(k, d) \geq \frac{2}{3}$;
2. for $k \geq 2 \cdot \lceil \log d \rceil$, $P(k, d) > \frac{1}{2}$.

Analyzing Decay

Theorem 1.

Let y be a vertex of G. Assume $d \geq 2$ neighbors of y execute $Decay(k, m)$ during the interval $[0, k)$ starting at time 0. Then $P(k, d)$, the probability that y receives a message by time k, satisfies:

1. $\lim_{k \to \infty} P(k, d) \geq \frac{2}{3}$;
2. for $k \geq 2 \cdot \lceil \log d \rceil$, $P(k, d) > \frac{1}{2}$.

Core technical argument: defines $P(k, d)$ as a recurrence and bounds inductively; $\frac{2}{3}$ is strong bound.

Analyzing Decay

Theorem 1.

Let y be a vertex of G. Assume $d \geq 2$ neighbors of y execute $Decay(k, m)$ during the interval $[0, k)$ starting at time 0. Then $P(k, d)$, the probability that y receives a message by time k, satisfies:

1. $\lim_{k \to \infty} P(k, d) \geq \frac{2}{3}$;
2. for $k \geq 2 \cdot \lceil \log d \rceil$, $P(k, d) > \frac{1}{2}$.

Lose a little probability to cover rare event of non-termination in $2 \cdot \lceil \log d \rceil$ rounds.

Uniform Decay

Uniform-Decay(k, m, n)

Repeat k times:

$r \leftarrow$ global round

$p \leftarrow 2^{-1-(r \mod \lceil \log n \rceil)}$

`transmit` m with prob. p

Uniform Decay

Uniform-Decay(k, m, n)

Repeat k times:

$r \leftarrow$ global round

$p \leftarrow 2^{-1-(r \mod \lceil \log n \rceil)}$

`transmit` m with prob. p

Induces fixed transmit probability schedule:
$\frac{1}{2}, \frac{1}{4}, \ldots, \frac{2}{n}, \frac{1}{n}, \frac{1}{2}, \frac{1}{4} \cdots$

Uniform Decay

- First introduced: [Jurdziński, Stachowiak: ISAAC 2002]

Uniform-Decay(k, m, n)

Repeat k times:

$r \leftarrow$ global round

$p \leftarrow 2^{-1-(r \mod \lceil \log n \rceil)}$

`transmit` m with prob. p

Induces fixed transmit probability schedule:

$\frac{1}{2}, \frac{1}{4}, \ldots, \frac{2}{n}, \frac{1}{n}, \frac{1}{2}, \frac{1}{4} \cdots$

Uniform Decay

- First introduced: [Jurdziński, Stachowiak: ISAAC 2002]
- For one cycle, original *Decay* has similar basic dynamics: transmit in i^{th} round with probability 2^{-i}.

Uniform-Decay(k, m, n)

Repeat k times:

$r \leftarrow$ global round

$p \leftarrow 2^{-1-(r \mod \lceil \log n \rceil)}$

`transmit` m with prob. p

Induces fixed transmit probability schedule:
$\frac{1}{2}, \frac{1}{4}, \ldots, \frac{2}{n}, \frac{1}{n}, \frac{1}{2}, \frac{1}{4} \cdots$

Uniform Decay

- First introduced: [Jurdziński, Stachowiak: ISAAC 2002]
- For one cycle, original *Decay* has similar basic dynamics: transmit in i^{th} round with probability 2^{-i}.
- Introduced because algorithm is *uniform*—transmit probabilities fixed in advance—simplifying lower bound arguments.

Uniform-Decay(k, m, n)

Repeat k times:

$r \leftarrow$ global round

$p \leftarrow 2^{-1-(r \mod \lceil \log n \rceil)}$

`transmit` m with prob. p

Induces fixed transmit probability schedule: $\frac{1}{2}, \frac{1}{4}, \ldots, \frac{2}{n}, \frac{1}{n}, \frac{1}{2}, \frac{1}{4} \cdots$

Uniform Decay

- First introduced: [Jurdziński, Stachowiak: ISAAC 2002]
- For one cycle, original *Decay* has similar basic dynamics: transmit in i^{th} round with probability 2^{-i}.
- Introduced because algorithm is *uniform*—transmit probabilities fixed in advance—simplifying lower bound arguments.
- Easy to analyze if happy with smaller constant.

Uniform-Decay(k, m, n)

Repeat k times:

$r \leftarrow$ global round

$p \leftarrow 2^{-1-(r \mod \lceil \log n \rceil)}$

`transmit` m with prob. p

Induces fixed transmit probability schedule:
$\frac{1}{2}, \frac{1}{4}, \ldots, \frac{2}{n}, \frac{1}{n}, \frac{1}{2}, \frac{1}{4} \cdots$

Analyzing Uniform Decay

Theorem 2.

Assume a $d \leq n$ stations call *Uniform-Decay* with $k = \lceil \log n \rceil$. The probability one station transmits alone is at least $1/8$.

Analyzing Uniform Decay

Theorem 2.

Assume a $d \leq n$ stations call *Uniform-Decay* with $k = \lceil \log n \rceil$. The probability one station transmits alone is at least $1/8$.

Fix some round r s.t. $p = 2^{-1-(r \mod \lceil \log n \rceil)} \in (\frac{1}{2d}, \frac{1}{d}]$ (requires $k \geq \lceil \log n \rceil$.)

Analyzing Uniform Decay

Theorem 2.

Assume a $d \leq n$ stations call *Uniform-Decay* with $k = \lceil \log n \rceil$. The probability one station transmits alone is at least $1/8$.

Fix some round r s.t. $p = 2^{-1-(r \mod \lceil \log n \rceil)} \in (\frac{1}{2d}, \frac{1}{d}]$ (requires $k \geq \lceil \log n \rceil$.)

$$\Pr\{\text{isolated station in } r\} = \sum_{j=1}^{d} p\,(1-p)^{d-1}$$

Analyzing Uniform Decay

Theorem 2.

Assume a $d \leq n$ stations call *Uniform-Decay* with $k = \lceil \log n \rceil$.
The probability one station transmits alone is at least $1/8$.

Fix some round r s.t. $p = 2^{-1-(r \mod \lceil \log n \rceil)} \in (\frac{1}{2d}, \frac{1}{d}]$ (requires $k \geq \lceil \log n \rceil$.)

$$\begin{aligned} \Pr\{\text{isolated station in } r\} &= \sum_{j=1}^{d} p\,(1-p)^{d-1} \\ &> \sum_{j=1}^{d} p\,(1-p)^{d} \end{aligned}$$

Analyzing Uniform Decay

Theorem 2.

Assume a $d \le n$ stations call *Uniform-Decay* with $k = \lceil \log n \rceil$. The probability one station transmits alone is at least $1/8$.

Fix some round r s.t. $p = 2^{-1-(r \mod \lceil \log n \rceil)} \in (\frac{1}{2d}, \frac{1}{d}]$ (requires $k \ge \lceil \log n \rceil$.)

$$\begin{aligned}
\Pr\{\text{isolated station in } r\} &= \sum_{j=1}^{d} p\,(1-p)^{d-1} \\
&> \sum_{j=1}^{d} p\,(1-p)^{d} \\
&\overset{p \in (\frac{1}{2d}], \frac{1}{d})}{\ge} \sum_{j=1}^{d} \frac{1}{2d}\left(1 - \frac{1}{d}\right)^{d}
\end{aligned}$$

Analyzing Uniform Decay

Theorem 2.

Assume a $d \leq n$ stations call *Uniform-Decay* with $k = \lceil \log n \rceil$. The probability one station transmits alone is at least $1/8$.

Fix some round r s.t. $p = 2^{-1-(r \mod \lceil \log n \rceil)} \in (\frac{1}{2d}, \frac{1}{d}]$ (requires $k \geq \lceil \log n \rceil$.)

$$\begin{aligned}
\Pr\{\text{isolated station in } r\} &= \sum_{j=1}^{d} p\,(1-p)^{d-1} \\
&> \sum_{j=1}^{d} p\,(1-p)^{d} \\
&\overset{p \in (\frac{1}{2d}], \frac{1}{d})}{\geq} \sum_{j=1}^{d} \frac{1}{2d} \left(1 - \frac{1}{d}\right)^{d} \\
&\overset{(1-p) \geq \frac{1}{4}^{p}}{\geq} \frac{d}{2d} (1/4)^{\frac{d}{d}} = \frac{1}{8}
\end{aligned}$$

The Optimality of Decay

Upper Bounds	**Lower Bounds**
Local Broadcast	**Local Broadcast**
Global Broadcast	**Global Broadcast**

(Results shown for high probability of success.)

The Optimality of Decay

Upper Bounds	**Lower Bounds**
Local Broadcast	**Local Broadcast**
$O(\log^2 n)$ [BGI: PODC 1987], [JS: ISAAC 2002]	
Global Broadcast	**Global Broadcast**

(Results shown for high probability of success.)

The Optimality of Decay

Upper Bounds	Lower Bounds
Local Broadcast	**Local Broadcast**
$O(\log^2 n)$ [BGI: PODC 1987], [JS: ISAAC 2002]	$\Omega(\log^2 n)$ [Farach-Colton: LATIN 2006], [Newport: DISC 2014]
Global Broadcast	**Global Broadcast**

(Results shown for high probability of success.)

The Optimality of Decay

Upper Bounds	Lower Bounds
Local Broadcast	**Local Broadcast**
$O(\log^2 n)$ [BGI: PODC 1987], [JS: ISAAC 2002]	$\Omega(\log^2 n)$ [Farach-Colton: LATIN 2006], [Newport: DISC 2014]
Global Broadcast	**Global Broadcast**
$O(D \log n + \log^2 n)$ [BGI: PODC 1987]	

(Results shown for high probability of success.)

The Optimality of Decay

Upper Bounds	Lower Bounds
Local Broadcast	**Local Broadcast**
$O(\log^2 n)$ [BGI: PODC 1987], [JS: ISAAC 2002]	$\Omega(\log^2 n)$ [Farach-Colton: LATIN 2006], [Newport: DISC 2014]
Global Broadcast	**Global Broadcast**
$O(D \log n + \log^2 n)$ [BGI: PODC 1987]	$\Omega(D \log (n/D) + \log^2 n)$ [Alon, bar-Noy, Linial, Peleg: Journal of Comp. and Sys. Sciences, 1991], [Kushilevitz and Mansour: SIAM J. Comp. 1998], [Newport, PODC 2013]

(Results shown for high probability of success.)

Diving Deeper on Decay

The Big Question

Does *Decay*'s surprising effectiveness depend on the stability of the underlying radio network topology?

The Dual Graph Model [Kuhn, Lynch, Newport: PODC 2009]

Adding Unreliable Edges to Radio Network Model

- Network described by two graphs: $G = (V, E)$, $G' = (V, E')$, where $E \subseteq E'$.
- Edges from E are always present.
- In each round, an **adversary** adds subset of edges from $E' \setminus E$.

The Dual Graph Model [Kuhn, Lynch, Newport: PODC 2009]

Adding Unreliable Edges to Radio Network Model

- Network described by two graphs: $G = (V, E)$, $G' = (V, E')$, where $E \subseteq E'$.
- Edges from E are always present.
- In each round, an **adversary** adds subset of edges from $E' \setminus E$.

Two types of adversary

online adaptive: can use the current state of the network (but not future random choices) to choose edges to add.

oblivious: must make all edge decisions at beginning of execution.

Diving Deeper on Decay Revisited

Refined Big Question

Is *Decay* still effective when the radio network topology can change; i.e., in the dual graph model?

A Negative Answer

Theorem 3. [Ghaffari, Lynch and Newport: PODC 2013]

There exists a constant diameter dual graph network in which both local and global broadcast require $\Omega(n/\log n)$ rounds to solve w.h.p. with an online adaptive adversary.

A Negative Answer

Theorem 3. [Ghaffari, Lynch and Newport: PODC 2013]

There exists a constant diameter dual graph network in which both local and global broadcast require $\Omega(n/\log n)$ rounds to solve w.h.p. with an online adaptive adversary.

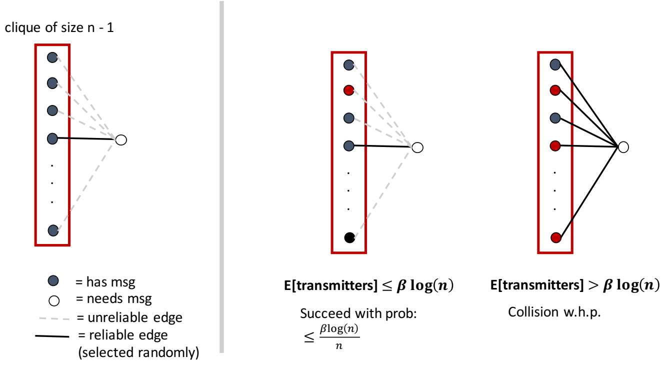

A Negative Answer

Theorem 3. [Ghaffari, Lynch and Newport: PODC 2013]

There exists a constant diameter dual graph network in which both local and global broadcast require $\Omega(n/\log n)$ rounds to solve w.h.p. with an online adaptive adversary.

Not Quite Right: If $E[transmitter]$ is small and receiver sends, then endpoint of single reliable edge learns its identity. To fix requires we replace 1 receiver with clique of $n/2$ receivers.

A Negative Answer

Theorem 4. [Ghaffari, Lynch and Newport: PODC 2013]

There exists a dual graph network in which local broadcast requires $\Omega(\sqrt{n}/\log n)$ rounds to solve w.h.p. with an oblivious adversary.

A Negative Answer

Theorem 4. [Ghaffari, Lynch and Newport: PODC 2013]

There exists a dual graph network in which local broadcast requires $\Omega(\sqrt{n}/\log n)$ rounds to solve w.h.p. with an oblivious adversary.

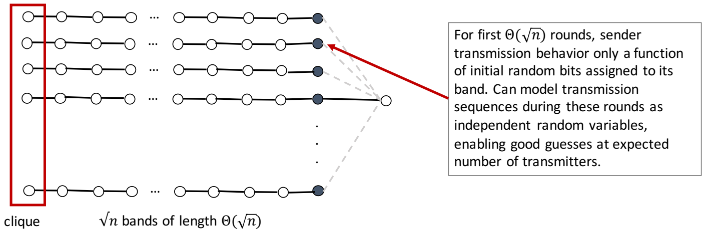

Some Good News

Positive Results for Oblivious Adversary [Ghaffari, Lynch and Newport: PODC 2013]

- Global broadcast can be solved in $O(D \log(n/D) + \log^2 n)$ rounds in the dual graph model with oblivious adversary and a broadcast message that can fit lots of random bits.
  - **Idea.** Randomly permute the order of probabilities used by *Decay* after execution begins.
  - **Note.** There *exist* good random number generators for this purpose that require only a $O(\text{polylog}(n))$-size seed. [Newport: PODC 2016]

Diving Even Deeper on Decay

Further Refined Big Question

Is *Decay* still effective when the radio network topology can change in a *realistic* manner?

The Fading Adversary

Intuition

Two types of changes to the topology:

- Short-term “fades” are **stochastic** (e.g., background noise).
- Long-term “fades” are **adversarial** (e.g., radio moves).

The Fading Adversary

Intuition

Two types of changes to the topology:

- Short-term "fades" are **stochastic** (e.g., background noise).
- Long-term "fades" are **adversarial** (e.g., radio moves).

Fading Adversary [Gilbert, Lynch, Newport and Pajak: 2018]

The Fading Adversary

Intuition

Two types of changes to the topology:

- Short-term "fades" are **stochastic** (e.g., background noise).
- Long-term "fades" are **adversarial** (e.g., radio moves).

Fading Adversary [Gilbert, Lynch, Newport and Pajak: 2018]

- In each round: adversary draws unreliable edges to include in the topology from an arbitrary **distribution** over $\mathcal{P}(E' \setminus E)$.

The Fading Adversary

Intuition

Two types of changes to the topology:

- Short-term "fades" are **stochastic** (e.g., background noise).
- Long-term "fades" are **adversarial** (e.g., radio moves).

Fading Adversary [Gilbert, Lynch, Newport and Pajak: 2018]

- In each round: adversary draws unreliable edges to include in the topology from an arbitrary **distribution** over $\mathcal{P}(E' \setminus E)$.
- The adversary can **change the distribution** at most once every τ steps. Note: $\tau = 1$ reduces to standard adversaries and $\tau = \infty$ is purely stochastic.

The Fading Adversary

Intuition

Two types of changes to the topology:

- Short-term "fades" are **stochastic** (e.g., background noise).
- Long-term "fades" are **adversarial** (e.g., radio moves).

Fading Adversary [Gilbert, Lynch, Newport and Pajak: 2018]

- In each round: adversary draws unreliable edges to include in the topology from an arbitrary **distribution** over $\mathcal{P}(E' \setminus E)$.
- The adversary can **change the distribution** at most once every τ steps. Note: $\tau = 1$ reduces to standard adversaries and $\tau = \infty$ is purely stochastic.
- Allows lots of short-term stochastic changes and arbitrary long term changes.

The Fading Adversary

Intuition

Two types of changes to the topology:

- Short-term "fades" are **stochastic** (e.g., background noise).
- Long-term "fades" are **adversarial** (e.g., radio moves).

Fading Adversary [Gilbert, Lynch, Newport and Pajak: 2018]

- In each round: adversary draws unreliable edges to include in the topology from an arbitrary **distribution** over $\mathcal{P}(E' \setminus E)$.
- The adversary can **change the distribution** at most once every τ steps. Note: $\tau = 1$ reduces to standard adversaries and $\tau = \infty$ is purely stochastic.
- Allows lots of short-term stochastic changes and arbitrary long term changes.

In basic case, the samples in each round are independent.

Dynamic Decay

Dynamic-Decay($k, \hat{\tau}, m, \Delta$)

Repeat k times:

$r \leftarrow$ *global round* mod $\hat{\tau}$

$p \leftarrow \Delta^{-\frac{(r+1)}{\hat{\tau}}}$

`transmit` m with prob. p

Dynamic Decay

Dynamic-Decay($k, \hat{\tau}, m, \Delta$)

Repeat k times:

$r \leftarrow$ *global round* mod $\hat{\tau}$

$p \leftarrow \Delta^{-\frac{(r+1)}{\hat{\tau}}}$

`transmit` m with prob. p

Note: Δ is upper bound on maximum degree, and $\hat{\tau}$ is minimum of $\log \Delta$ and lower bound on τ.

Dynamic Decay

- For $\hat{\tau} = \log \Delta$: *Dynamic-Decay* reduces to *Uniform-Decay*.

Dynamic-Decay($k, \hat{\tau}, m, \Delta$)

Repeat k times:

$r \leftarrow$ *global round* mod $\hat{\tau}$

$p \leftarrow \Delta^{-\frac{(r+1)}{\hat{\tau}}}$

`transmit` m with prob. p

Note: Δ is upper bound on maximum degree, and $\hat{\tau}$ is minimum of $\log \Delta$ and lower bound on τ.

Dynamic Decay

- For $\hat{\tau} = \log \Delta$: *Dynamic-Decay* reduces to *Uniform-Decay*.
- For smaller $\hat{\tau}$ the samples of probabilities in range $1/\Delta$ to $1/2$ becomes sparser.

Dynamic-Decay($k, \hat{\tau}, m, \Delta$)

Repeat k times:

$r \leftarrow$ *global round* mod $\hat{\tau}$

$p \leftarrow \Delta^{-\frac{(r+1)}{\hat{\tau}}}$

`transmit` m with prob. p

Note: Δ is upper bound on maximum degree, and $\hat{\tau}$ is minimum of $\log \Delta$ and lower bound on τ.

Dynamic Decay Performance

Theorem 5.

When called with $k = \Theta(\log(1/\epsilon)\Delta^{\frac{1}{\hat{\tau}}}\hat{\tau})$, *Dynamic-Decay* solves local broadcast with probability $1 - \epsilon$.

Dynamic Decay Performance

Theorem 5.

When called with $k = \Theta(\log(1/\epsilon)\Delta^{\frac{1}{\hat{\tau}}}\hat{\tau})$, *Dynamic-Decay* solves local broadcast with probability $1-\epsilon$. Thus: for $\tau \geq \log\Delta$, *Uniform-Decay* solves local broadcast with a fading adversary with high probability in $O(\log^2 n)$ rounds, matching result for standard radio network model.

Dynamic Decay Performance

Theorem 5.

When called with $k = \Theta(\log(1/\epsilon)\Delta^{\frac{1}{\hat{\tau}}}\hat{\tau})$, *Dynamic-Decay* solves local broadcast with probability $1-\epsilon$. Thus: for $\tau \geq \log\Delta$, *Uniform-Decay* solves local broadcast with a fading adversary with high probability in $O(\log^2 n)$ rounds, matching result for standard radio network model.

Intuition:

- Every distribution must put a "reasonable" amount of probability mass on receiver degrees near at least one *target d* (counting argument).
- This target stays the same for *stable phases* of length $\hat{\tau}$ rounds.
- *Dynamic-Decay* samples $\hat{\tau}$ evenly spread guesses of d during stable phase. If $\hat{\tau}$ is big, one of these guesses will be within constant factor of d.

Dynamic Decay Performance: Optimized Bounds

[Results from GLNP: 2018]...

Problem	Time	Prob.	Remarks
Local bcast	$\mathcal{O}\left(\frac{\Delta^{1/\bar{\tau}} \cdot \bar{\tau}^2}{\log \Delta} \cdot \log(1/\epsilon)\right)$	$1-\epsilon$	$\bar{\tau} = \min\{\tau, \log \Delta\}$
	$\Omega\left(\frac{\Delta^{1/\tau}\tau}{\log \Delta}\right)$	$\frac{1}{2}$	$\tau \in O(\log \Delta)$
	$\Omega\left(\frac{\Delta^{1/\tau}\tau^2}{\log \Delta}\right)$	$\frac{1}{2}$	$\tau \in O\left(\frac{\log \Delta}{\log\log \Delta}\right)$

Dynamic Decay Performance: Optimized Bounds

[Results from GLNP: 2018]...

Problem	Time	Prob.	Remarks
Local bcast	$\mathcal{O}\left(\frac{\Delta^{1/\bar{\tau}} \cdot \bar{\tau}^2}{\log\Delta} \cdot \log(1/\epsilon)\right)$	$1-\epsilon$	$\bar{\tau} = \min\{\tau, \log\Delta\}$
	$\Omega\left(\frac{\Delta^{1/\tau}\tau}{\log\Delta}\right)$	$\frac{1}{2}$	$\tau \in O(\log\Delta)$
	$\Omega\left(\frac{\Delta^{1/\tau}\tau^2}{\log\Delta}\right)$	$\frac{1}{2}$	$\tau \in O\left(\frac{\log\Delta}{\log\log\Delta}\right)$
Global bcast	$\mathcal{O}\left((D + \log(n/\epsilon)) \cdot \frac{\Delta^{1/\bar{\tau}}\bar{\tau}^2}{\log\Delta}\right)$	$1-\epsilon$	$\bar{\tau} = \min\{\tau, \log\Delta\}$
	$\Omega\left(D \cdot \frac{\Delta^{1/\tau}\tau}{\log\Delta}\right)$	$\frac{1}{2}$	$\tau \in O(\log\Delta)$
	$\Omega\left(D \cdot \frac{\Delta^{1/\tau}\tau^2}{\log\Delta}\right)$	$\frac{1}{2}$	$\tau \in O\left(\frac{\log\Delta}{\log\log\Delta}\right)$

Note: *Slightly optimized* Dynamic-Decay *considered here multiples probabilities by* $\Theta(\log\Delta/\hat{\tau})$ *factor. This still reduces to* Uniform-Decay *for* $\hat{\tau} = \log\Delta$, *but performs slightly better for small* $\hat{\tau}$.

Conclusion

Some Surprising Observations...

Conclusion

Some Surprising Observations...

- **Thirty years ago:** Bar-Yehuda, Goldreich, and Itai proved that a basic distributed strategy of cycling through transmission probabilities solved radio network broadcast better than the best-known centralized solution. This is surprising.

Conclusion

Some Surprising Observations...

- **Thirty years ago:** Bar-Yehuda, Goldreich, and Itai proved that a basic distributed strategy of cycling through transmission probabilities solved radio network broadcast better than the best-known centralized solution. This is surprising.
- This simple *decay* strategy turned out to be provably optimal for both local and global distributed broadcast tasks in the radio network model. This is surprising.

Conclusion

Some Surprising Observations...

- **Thirty years ago:** Bar-Yehuda, Goldreich, and Itai proved that a basic distributed strategy of cycling through transmission probabilities solved radio network broadcast better than the best-known centralized solution. This is surprising.
- This simple *decay* strategy turned out to be provably optimal for both local and global distributed broadcast tasks in the radio network model. This is surprising.
- This simple strategy continues to efficiently and optimally solve these problems in *highly dynamic* versions of the radio network model (so long as the dynamism is realistically modeled). This is surprising.

Conclusion

Some Surprising Observations...

- **Thirty years ago:** Bar-Yehuda, Goldreich, and Itai proved that a basic distributed strategy of cycling through transmission probabilities solved radio network broadcast better than the best-known centralized solution. This is surprising.
- This simple *decay* strategy turned out to be provably optimal for both local and global distributed broadcast tasks in the radio network model. This is surprising.
- This simple strategy continues to efficiently and optimally solve these problems in *highly dynamic* versions of the radio network model (so long as the dynamism is realistically modeled). This is surprising.
- **Conclusion:** This strategy is unreasonably effective!