# Broadcast and Consensus in Stochastic Dynamic Networks with Byzantine Nodes and Adversarial Edges

**Antoine El-Hayek**
Institute of Science and Technology Austria (ISTA), Klosterneuburg, Austria

**Monika Henzinger**
Institute of Science and Technology Austria (ISTA), Klosterneuburg, Austria

**Stefan Schmid**
TU Berlin, Germany
Fraunhofer SIT, Berlin, Germany

## Abstract

Broadcast and Consensus are most fundamental tasks in distributed computing. These tasks are particularly challenging in dynamic networks where communication across the network links may be unreliable, e.g., due to mobility or failures. Over the last years, researchers have derived several impossibility results and high time complexity lower bounds for these tasks. Specifically for the setting where in each round of communication the adversary is allowed to choose one rooted tree along which the information is disseminated, there is a lower as well as an upper bound that is linear in the number $n$ of nodes for Broadcast and for $n \geq 3$ the adversary can guarantee that Consensus never happens. This setting is called the *oblivious message adversary for rooted trees*. Also note that if the adversary is allowed to choose a graph that does not contain a rooted tree, then it can guarantee that Broadcast and Consensus will never happen.

However, such deterministic adversarial models may be overly pessimistic, as many processes in real-world settings are stochastic in nature rather than worst-case.

This paper studies Broadcast on *stochastic* dynamic networks and shows that the situation is very different to the deterministic case. In particular, we show that if information dissemination occurs along random rooted trees and directed Erdős–Rényi graphs, Broadcast completes in $O(\log n)$ rounds of communication with high probability. The fundamental insight in our analysis is that key variables are mutually independent.

We then study two adversarial models, (a) one with Byzantine nodes and (b) one where an adversary controls the edges. (a) Our techniques without Byzantine nodes are general enough so that they can be extended to Byzantine nodes. (b) In the spirit of smoothed analysis, we introduce the notion of *randomized oblivious message adversary*, where in each round, an adversary picks $k \leq 2n/3$ edges to appear in the communication network, and then a graph (e.g. rooted tree or directed Erdős–Rényi graph) is chosen uniformly at random among the set of all such graphs that include these edges. We show that Broadcast completes in a finite number of rounds, which is, e.g., $O(k + \log n)$ rounds in rooted trees.

We then extend these results to All-to-All Broadcast, and Consensus, and give lower bounds that show that most of our upper bounds are tight.

**2012 ACM Subject Classification** Theory of computation → Distributed algorithms; Networks → Network algorithms

**Keywords and phrases** Broadcast, Smoothed Analysis, Stochastic Networks, Dynamic Networks

**Digital Object Identifier** 10.4230/LIPIcs.DISC.2024.21

**Related Version** *Full Version*: `arXiv:2302.11988v2`

**Funding** *Antoine El-Hayek*: This project has received funding from the Austrian Science Fund (FWF) grant DOI 10.55776/P33775 with additional funding from the netidee SCIENCE Stiftung, 2020–2024.

*Monika Henzinger*: This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (MoDynStruct, No. 101019564) and the Austrian Science Fund (FWF) grant DOI 10.55776/Z422, grant DOI 10.55776/I5982, and grant DOI 10.55776/P33775 with additional funding from the netidee SCIENCE Stiftung, 2020–2024.
*Stefan Schmid*: This project has received funding from the German Research Foundation (DFG), SPP 2378 (project ReNO), 2023-2027.

## 1 Introduction

Broadcast and Consensus are two of most fundamental operations in distributed computing which, in large-scale systems, typically have to be performed over a *network*. These networks are likely to be dynamic and change over time due, e.g., to link failures, interference, or mobility. Understanding how information disseminates in such dynamic networks is hence important for developing and analyzing efficient distributed systems.

Over the last years, researchers have derived several important insights into information dissemination in dynamic networks. A natural and popular model assumes an *oblivious*[1] *message adversary* which controls the information flow between a set of $n$ nodes, by dropping an arbitrary set of messages sent by some nodes in each round [7]. Specifically, the adversary is defined by a set of directed communication graphs, one per round, whose edges determine which node can successfully send a message to which other node in a given round. Based on this set of graphs, the oblivious message adversary chooses a sequence of graphs over time, one per round with repetitions allowed, in such a way that the time complexity of the information dissemination task at hand is maximized. This model is appealing because it is conceptually simple and still provides a highly dynamic network model: The set of allowed graphs can be arbitrary, and the nodes that can communicate with one another can vary greatly from one round to the next. It is, thus, well-suited for settings where significant transient message loss occurs, such as in wireless networks. As information dissemination is faster on dense networks, most literature studies oblivious message adversaries on sparse networks, in particular, on rooted trees [16, 30, 7, 21, 22]. In fact, it is easy to see that rooted trees are a minimal necessary requirement for a successful Broadcast and Consensus: if an adversary may choose a graph that does not contain a rooted tree, then it may forever prevent the dissemination of a piece of information.

Unfortunately, information dissemination can be slow in trees: Broadcast can take time linear in the number of nodes under the oblivious message adversary [16, 30], even for constant-height trees (as we show in the full version); and Consensus can even take super-polynomial time until termination, if it completes at all [7, 21]. Although this is bad news, one may argue that while the deterministic adversary model is useful in malicious environments, in real-word applications, the dynamics of communication networks is often more stochastic in nature. Accordingly, the worst-case model considered in existing literature may be overly conservative.

This motivates us, in this paper, to study information dissemination, and in particular Broadcast and Consensus tasks, in a scenario where the communication network is stochastic. Initially, we study a purely stochastic scenario where in each round, the communication network is chosen uniformly at random among all rooted trees. We then study several

[1] Note that the term oblivious here refers to the property that nodes are oblivious to who their neighbors are. However, our adversary is actually adaptive.

fundamental extensions of this model where the adversary has some limited control. In a first extension, we consider the case where some nodes (up to $\frac{2n}{3}$) may be Byzantine, that is, they may deviate arbitrarily from the protocol (and stop forwarding messages, for example). In a second extension, in the spirit of smoothed analysis, we study a setting where an adversary has some limited control over the communication network; we call this adversary the *randomized oblivious message adversary*. More specifically, we study the setting where first a worst-case adversary chooses $k$ directed edges in the dynamic $n$-node network for some fixed $k$ with $0 \leq k < \frac{2n}{3} - 1$[2], and then a rooted tree is chosen uniformly at random among the set of all rooted trees that include these edges.

We show that Broadcast completes within time $O(\log n)$ with high probability. We then show that this result even holds with Byzantine nodes. Under our randomized oblivious message adversary, Broadcast completes in $O(k + \log n)$ time with high probability.

It is useful to put our model into perspective with the SI (Susceptible-Infectious) model in epidemics [13]: while in the SI model interactions occur on a network that equals a clique, our model revolves around trees which are chosen by an adversary. This tree structure renders the analytical understanding of the information dissemination process harder, due to the lack of independence between the edges in the network in a particular round. A key insight from our paper is that we can prove the independence of a key variable, namely the increase in the number of "informed" nodes, which is crucial for our analysis. Our proof further relies on stochastic dominance, which makes it robust to the specific adversarial objective, and applies to any adversary definition (e.g., whether it aims to maximize the minimum or the expected number of rounds until the process completes).

We then extend our study to adversaries which are not limited to trees. In particular, we are interested in how the time complexity of Broadcast and Consensus depends on the density of the network. To this end, we consider *directed Erdős–Rényi graphs*, a directed version of the classic and well-studied random graphs. This graph family is parameterized by the number of edges $m$ and hence allows us to shed light on the impact of the density. Specifically in this model, in each round the network is formed by sampling $m$ edges. We again study two extensions: in the first extension some nodes behave as Byzantine nodes, while in the second extension, up to $k \leq m$ edges are chosen by an adversary, and then the remaining edges are sampled. While results for this model can be found in some cases where $m$ is chosen so that the graph is an expander w.h.p. in each round by using the results from Augustine et al [2], in the case where $m$ is small, our results are novel.

We show that all our results extend to multiple other problems, namely All-to-All Broadcast, Byzantine Consensus and Reliable Broadcast.

## 1.1 Model

Let $n$ be the number of nodes, and let each node have a unique identifier from $[n]$. Time proceeds in a sequence of rounds $t = 1, 2, \ldots,$ such that in each round $t$ the communication network is chosen according to one of the models defined below. In each round, every honest node sends a message to all of its out-neighbors before receiving one from its in-neighbor. There is no message size restriction. We will study the following models of communication:

[2] We can relax this condition to $k \leq (1 - \epsilon)n$ for a fixed parameter $\epsilon$, which results in a multiplicative factor of $\frac{1}{\epsilon}$ in the running time.

### Uniformly Random Trees

In the *Uniformly Random Trees* model, let $\mathcal{T}_n$ be the set of all directed rooted trees on $n$ nodes (where all edges are pointed away from the root). In each round, the communication network is chosen uniformly at random among graphs in $\mathcal{T}_n$, independently from other rounds. All nodes are honest.

### Uniformly Random Trees with Byzantine Nodes

In the *Uniformly Random Trees with Byzantine Nodes* model, in each round, the communication network is chosen uniformly at random among graphs in $\mathcal{T}_n$, independently from other rounds. We have $n - f$ honest nodes, and $f$ nodes are Byzantine, that is, they might behave arbitrarily (and even coordinate to make the protocol fail). We assume access to cryptographic tools that allow nodes to sign and encrypt messages. We restrict $f \leq \frac{2n}{3} - 1$.

### Uniformly Random Trees with Adversarial Edges

In the *Uniformly Random Trees with Adversarial Edges* model, in each round, the communication network is chosen as follows: A randomized oblivious message adversary chooses $k$ directed edges, then a graph is chosen uniformly at random among all graphs in $\mathcal{T}_n$ that include those $k$ edges, and the choise is independent from other rounds. All nodes are honest. We restrict $k \leq \frac{2n}{3} - 1$.

### Directed Erdős–Rényi graphs

In the *directed Erdős–Rényi graphs* model, let $m \in [n^2]$. In each round, the communication network is chosen by uniformly sampling without replacement $m$ edges out of the possible $n^2$ edges of the graph, independently from other rounds. All nodes are honest.

### Directed Erdős–Rényi graphs with Byzantine Nodes

In the *directed Erdős–Rényi graphs with Byzantine nodes* model, let $m \in [n^2]$. In each round, the communication network is chosen by uniformly sampling without replacement $m$ edges out of the possible $n^2$ edges of the graph, independently from other rounds. We have $n - k$ honest nodes, and $k$ nodes are Byzantine, that is, they might behave arbitrarily (and even coordinate to make the protocol fail). We assume access to cryptographic tools that allow nodes to sign and encrypt messages. We restrict $k < \frac{2n}{3}$.

### Directed Erdős–Rényi graphs with Adversarial Edges

In the *directed Erdős–Rényi graphs with Adversarial Edges* model, let $0 \leq k \leq m \leq n^2$. In each round, the communication network is chosen as follows: A randomized oblivious message adversary chooses $k$ edges, $m - k$ edges are sampled without replacement out of the remaining $n^2 - k$ edges. All nodes are honest. We restrict $k < \frac{3}{4}n^2$.

In those models, we will study the following problems:

### Broadcast

For the *Broadcast*[3] problem, we start by giving a message to *one* (honest) node. Each honest node that received the message will replicate it as many times as needed, and start forwarding it to its neighbors[4]. Then Broadcast *completes* when the message has been forwarded to all other nodes.

### All-to-All Broadcast

In the *All-to-All Broadcast* problem, we start by giving a distinct message to *each* node. Each honest node that received a message will replicate it as many times as needed, and start forwarding it as well. Then All-to-All Broadcast *completes* when each honest node receives a copy of every message. In each round, each honest node forwards all the messages it has received in previous rounds to all its out-neighbors.

### Consensus

In the *Consensus* problem, we start by giving a value $v_p \in \{0, 1\}$ to each node $p$, and Consensus completes when each honest node decided on a value in $\{0, 1\}$. This should satisfy the following conditions:

- **Agreement:** No two honest nodes decide differently.
- **Termination:** Every honest node eventually decides.
- **Validity:** The value the honest nodes agree on should be one of the input values $v_p$.

## 1.2 Our Results

We study Broadcast in the above mentioned models, then apply those results to All-to-All broadcast and Consensus. We prove the following theorems:

▶ **Theorem 1.** *For any $c \geq 1$ and $n \geq 5$, Broadcast on Uniformly Random Trees completes within $32 \cdot c \cdot \ln n$ rounds with probability $p > 1 - \frac{1}{n^c}$.*

We also show that these results are asymptotically tight. Indeed, we cannot hope for a similar probability for a number of rounds that is $o(\ln n)$:

▶ **Theorem 2.** *If $n \geq 2$, then the probability that Broadcast (and All-to-All Broadcast) on Uniformly Random Trees fails to complete within $\log n$ rounds is at least $\frac{1}{4}$.*

We have similar results for all the combinations of model and problem, which we summarize in Table 1.

### Applications

Our results have some interesting applications. In an idea similar to Ghaffari, Kuhn and Su's work [23], All-to-All Broadcast allows us, e.g., to implement algorithms that run on a clique in a synchronous setting in our sparser graphs. Indeed, if All-to-All Broadcast needs

[3] The Broadcast problem can also be seen as computing the *dynamic eccentricity* of the source node. Other flavors of Broadcast have also been studied under the name *dynamic radius* [20].
[4] This is known as “flooding” or “rumor passing”

| | Broadcast | All-to-All Broadcast | Consensus |
|---|---|---|---|
| Uniformly Random Trees (URT) | $O(c \cdot \log n), q \leq n^{-c}$<br>$\Omega(\log n)$ | $O(c \cdot \log n), q \leq n^{1-c}$<br>$\Omega(\log n)$ | $O(c \cdot \log n), q \leq n^{-c}$ |
| URT with Byzantine Nodes | $O(c \cdot \log n), q \leq n^{-c}$<br>$\Omega(\log n)$ | $O(c \cdot \log n), q \leq n^{1-c}$<br>$\Omega(\log n)$ | $O(f \cdot c \cdot \log n), q \leq n^{-c}$ |
| URT with Adversarial Edges | $O(c \cdot (\log n + k)), q \leq n^{-c}$<br>$\Omega(\log n + k)$ | $O(c \cdot (\log n + k)), q \leq n^{1-c}$<br>$\Omega(\log n + k)$ | $O(c \cdot (\log n + k)), q \leq n^{-c}$ |
| Directed Erdős–Rényi graphs (DER) | $O\left(\left\lceil \frac{c}{m/n} \right\rceil \log n\right), q \leq n^{-c} \log n$<br>$O\left(\frac{c \log n}{\log(1+\frac{m}{n})}\right)$ if $\frac{m}{n} \geq \ln n$<br>with $q \leq n^{-c} \log n$<br>$\Omega\left(\frac{\log n}{\log(1+m/n)}\right)$ | $O\left(\left\lceil \frac{c}{m/n} \right\rceil \log n\right), q \leq n^{1-c} \log n$<br>$O\left(\frac{c \log n}{\log(1+\frac{m}{n})}\right)$ if $\frac{m}{n} \geq \ln n$<br>with $q \leq n^{1-c} \log n$<br>$\Omega\left(\frac{\log n}{\log(1+m/n)}\right)$ | $O\left(\left\lceil \frac{c}{m/n} \right\rceil \log n\right), q \leq n^{-c} \log n$<br>$O\left(\frac{c \log n}{\log(1+\frac{m}{n})}\right)$ if $\frac{m}{n} \geq \ln n$<br>with $q \leq n^{-c} \log n$ |
| DER with Byzantine Nodes | $O\left(\left\lceil \frac{c}{m/n} \right\rceil \log n\right), q \leq n^{-c} \log n$<br>$\Omega\left(\frac{\log n}{\log(1+m/n)}\right)$ | $O\left(\left\lceil \frac{c}{m/n} \right\rceil \log n\right), q \leq n^{1-c} \log n$<br>$\Omega\left(\frac{\log n}{\log(1+m/n)}\right)$ | $O\left(f \cdot \left\lceil \frac{c}{m/n} \right\rceil \log n\right), q \leq n^{-c} \log n$ |
| DER with Adversarial Edges | $O\left(\left\lceil \frac{c \cdot (n^2-k)}{(m-k)n} \right\rceil \log n\right)$<br>with $q \leq n^{-c} \log n$<br>$\Omega\left(\frac{\log n}{\log(1+m/n)}\right)$ | $O\left(\left\lceil \frac{c \cdot (n^2-k)}{(m-k)n} \right\rceil \log n\right)$<br>with $q \leq n^{1-c} \log n$<br>$\Omega\left(\frac{\log n}{\log(1+m/n)}\right)$ | $O\left(\left\lceil \frac{c \cdot (n^2-k)}{(m-k)n} \right\rceil \log n\right)$<br>with $q \leq n^{-c} \log n$ |

**Figure 1** Our main results, where $c > 0$ is any constant and $q$ is the failure probability.

$R$ rounds to complete with high probability, then each round of communication of a clique can be simulated by $R$ rounds of Uniformly Random Trees with high probability. Essentially, if an algorithm runs in $T$ rounds, with $T \leq n^{c-1}$, in a clique network, we can implement it with high probability in $R \cdot T$ rounds in the Uniformly Random Trees network, which is essentially a logarithmic overhead. In particular, in the Uniformly Random Trees with Byzantine Nodes model, we have:

▶ **Theorem 3.** *Let $\mathcal{A}$ be a distributed synchronous algorithm that runs on a static clique in $T$ rounds, where $T \leq \alpha n^x$ for some constant $\alpha, x \in \mathbb{R}_+$, and has a probability of success $p$. Assume $\mathcal{A}$ is robust to $f$ Byzantine nodes, and $f \leq \frac{2}{3}n - 1$. Then, assuming standard cryptographic tools[5], there exists a distributed algorithm $\mathcal{A}'$ that runs on Uniformly Random Trees in $T \cdot 144 \cdot \log n \cdot c$ rounds, and has a probability of success $p' \geq p(1 - \alpha n^{1+x-c})$, for any $c \geq 1 + x$. Moreover, $\mathcal{A}'$ is robust to $f$ Byzantine nodes.*

In particular, we can apply known results on reliable Broadcast and Byzantine Consensus to show the following results:

▶ **Corollary 4.** *For any $c \geq 1$, and $f \leq \frac{2}{3}n - 1$, in the Uniformly Random Trees with $f$ Byzantine nodes, there exists an algorithm for Reliable Broadcast, that is robust to $f$ Byzantine nodes, that runs in $(f + 1) \cdot 144 \cdot c \cdot \log n$ rounds, and succeeds with probability $p \geq 1 - n^{2-c}$.*

▶ **Corollary 5.** *For any $c \geq 1$ and $f < \frac{n}{3}$, in the Uniformly Random Trees with $f$ Byzantine nodes, there exists an algorithm for Byzantine Consensus, that is robust to $f$ Byzantine nodes, that runs in $3(f + 1) \cdot 144 \cdot c \cdot \log n$ rounds, and succeeds with probability $p \geq 1 - 2n^{2-c}$.*

Throughout the paper, the filtration of the process is denoted as $\{\mathcal{F}_t\}_{t \in \mathbb{N}}$, that is, $\mathcal{F}_t$ is the amount of information available after timestep $t$.

[5] Specifically, our approach requires authenticated messages. Encryption may also be needed, only if the protocol $\mathcal{A}$ is vulnerable to eavesdropping. Both can be implemented using standard cryptographic tools.

### Organization

The paper is organized as follows. First, we review related work in Section 2. Then, due to space restrictions, we only give a technical overview in Section 3, as further details can be found in the full version of the paper. In this overview we first discuss a new result on the number of rooted trees containing a certain set of edges, then discuss how we analyzed information dissemination in random trees first, and finally in directed Erdős–Rényi graphs.

## 2 Related Work

Information dissemination in general and Broadcasting and Consensus in particular are fundamental topics in distributed computing. In contrast to this paper, most classic literature on network Broadcast as well as on related tasks such as gossiping and Consensus, considers a static setting, e.g., where in each round each node can send information to one neighbor [24, 19].

Especially the Byzantine setting has received much attention in the literature. Important results include Dolev and Strong [12] on reliable Broadcast which is robust to $f$ Byzantine nodes, and runs in $T = f + 1$ rounds, or Berman, Garay and Perry [3] on King's algorithm that solves reliable Broadcast, is robust to $f$ Byzantine nodes, and runs in $T = 3(f + 1)$ rounds. To just name a few.

In terms of dynamic networks, Kuhn, Lynch and Oshman [25] explore the all-to-all data dissemination problem (gossiping) in an undirected setting, where nodes do not know beforehand the total number of nodes and must decide on that number. Dutta, Pandurangan, Rajaraman, Sun and Viola [14] generalize the model to when not all nodes need to forward their message, but only $k$ tokens must be forwarded. Augustine, Pandurangan, Robinson and Upfal [2] show that if the graph is an expander in every round, broadcast is complete within $O(\log n)$ rounds, even if a small enough constant fraction of nodes get churned in each round. Ahmadi, Kuhn, Kutten, Molla and Pandurangan [1] study the message complexity of Broadcast also in an undirected dynamic setting, where the adversary pays up a cost for changing the network.

In dynamic networks, the oblivious message adversary is a commonly considered model, especially for Broadcast and Consensus problems, first introduced by Charron-Bost and Schiper [5]. The Broadcast problem under oblivious message adversaries has been studied for many years. A first key result for this problem was the $n \log n$ upper bound by Zeiner, Schwarz, and Schmid [30] who also gave a $\lceil \frac{3n-1}{2} \rceil - 2$ lower bound. Another important result is by Függer, Nowak, and Winkler [20] who presented an $O(\log \log n)$ upper bound if the adversary can only choose nonsplit graphs; combined with the result of Charron-Bost, Függer, and Nowak [4] that states that one can simulate $n - 1$ rounds of rooted trees with a round of a nonsplit graph, this gives the previous $O(n \log \log n)$ upper bound for Broadcasting on trees. Dobrev and Vrto [10, 9] give specific results when the adversary is restricted to hypercubic and tori graphs with some missing edges. El-Hayek, Henzinger, and Schmid [15, 16] recently settled the question about the asymptotic time complexity of Broadcast by giving a tight $O(n)$ upper bound, also showing the upper bound still holds in more general models. Regarding Consensus, Coulouma, Godard and Peters in [7] presented a general characterization on which dynamic graphs Consensus is solvable, based on Broadcastability. Winkler, Rincon Galeana, Paz, Schmid, and Schmid [21] recently presented an explicit decision procedure to determine if Consensus is possible under a given adversary, enabling a time complexity analysis of Consensus under oblivious message adversaries, both for a centralized decision procedure as well as for solving distributed Consensus. They also showed that reaching Consensus under an oblivious message adversary can take exponentially longer than Broadcasting.

In contrast to the above works, in this paper we study a more randomized message adversary, considering a stochastic model where adversarial graphs are partially chosen uniformly at random. While a randomized perspective on dynamic networks is natural and has been considered in many different settings already, existing works on random dynamic communication networks, e.g., on the radio network model [17], on rumor spreading [6], as well as on epidemics [13], do not consider oblivious message adversaries. Note, however, that the information dissemination considered in this paper is similar to the SI model for virus propagation, with results having implications in both directions [18]. For example, Doerr and Fouz [11] introduced an information dissemination protocol inspired by epidemics. More generally, randomized information dissemination protocols can be well-understood from an epidemiological point-of-view, and are very similar to the SI model which has been very extensively studied. In contrast to the typical SI models considered in the literature [28], however, our model in this paper revolves around tree communication structures which introduce additional technical challenges. Furthermore, existing literature often provides results in expectation, while we in this paper provide tail bounds.

Many papers have tried to bridge the gap between the deterministic and random case, using smoothed analysis. In [27], Meir, Paz and Schwartzman study the broadcast problem in noisy networks, under different definitions on noise. In particular, if in each round the graph given by the adversary is replaced by a graph chosen uniformly at random among graphs at hamming distance at most $k$ from the original graph, in the case where the adversary can suggest any connected graph, then Broadcast is reduced from $n$ rounds to $O(\min\{n, n\sqrt{\frac{\log n}{k}}\})$ rounds, in the case of an adaptive adversary. If the adversary is oblivious, then Dinitz, Fineman, Gilbert and Newport [8] showed that it is further reduced to $O(n^{2/3}/k^{1/3} \times \log n)$.

## 3 Technical Overview

Our paper contains a conceptional contribution, namely the extension of the notion of oblivious message adversary in a natural way to a randomized setting that limits the power of the adversary, as well as two technical contributions. We explained already the conceptional contribution in the introduction, and we sketch in this section now the main technical contributions of our paper. They are of graph theoretical as well as algorithmic nature. (1) On the graph theoretic side, we show a new result on the number of rooted trees that satisfy a certain property. (2) On the algorithmic side we show how to use this result to give an upper bound on the number of rounds for the models introduced in the introduction. Note that we study both the conventional as well as the Byzantine setting, where faulty nodes can stop forwarding, send wrong messages, and even coordinate to make the protocol fail. However, we assume access to cryptographic tools so that is used by each node to sign its messages. Thus, when receiving a message, nodes can be confident about the sender of each message and its content.

### 3.1 Counting rooted trees

Given a graph consisting of $n$ vertices together with a directed rooted forest $F$ of $e$ edges on them, Pitman [29] showed in 1999 that there are $n^{n-1-e}$ many directed rooted trees over these vertices that contain $F$. While useful, this result is not sufficient for our purposes as we need to count the number of trees with a given node $v$ as root.

Thus, we show the following extended result:

▶ **Theorem 6.** *Let us be given a directed rooted forest $F$ on $n$ vertices, let $v \in [n]$ be the root of a component in $F$, and $f$ be the number of vertices of that component (note that we can have $f = 1$ if $v$ is an isolated vertex). Then the number of directed rooted trees $T$ on $n$ vertices, such that $F$ is contained in $T$, and such that $v$ is the root of $T$, is $fn^{n-2-|E|}$.*

Note that our result implies the prior result.

To show our result, we develop techniques which differ significantly from Pitman's proof. Indeed, Pitman relies on the symmetry of the vertices in the rooted tree. However, for our result, the symmetry is broken as one vertex is different from the other with the new requirement that it is the root. We hence make use of another type of symmetry in the trees in our analysis that is based on group actions.

We first ignore the orientations of the edges in $F$ and find the set $A_F$ of all undirected trees that contain $F$. We can compute the cardinality of that set with a result by Lu, Mohr and Székely [26]. We then root each of those trees at $v$. This will give a direction to every edge that might or might not agree with its direction in $F$. We now want to partition $A_F$ into subsets such that all subsets have the same size and only one tree from each subset has edges that agree with the direction of $F$. The number we are looking for is then the number of subsets, which is the ratio between the cardinality of $A_F$ and the size of the subsets.

To create the subsets, we introduce a specific group tailored to $F$, and an action of that group on $A_F$. It is known that the set of all orbits of the action partition $A_F$, and we show that exactly one element in each orbit has edges in the same direction as $F$. To see unicity, we take an element $T$ of $A_F$ that has edges in the same direction as $F$, and take an element $T' \neq T$ in its orbit, that is there exists a nontrivial group element $g$ such that $T'$ is obtained from $T$ by applying the action of $g$ to $T$. We show that this action must change the direction of at least one edge of $F$, and thus $T'$ does not have edges in the same direction as $F$. For existence, we show that for every $T \in A_F$, we can find a group element $g$ such that, if applied to $T$, yields a tree that has edges in the same direction as $F$. We then show how to compute the size of each orbit. This allows us to deduce the number of orbits, which equals the number of trees that we want to count.

## 3.2 Analysis of the information dissemination

The main technical challenge is to analyze Broadcast in uniformly random trees (URTs) and in directed Erdős–Rényi graphs (DERs). Our techniques for both types of graphs are general and can be extended to adversarial settings, i.e. Byzantine nodes or adversarial edges, as well as to all-to-all Broadcast and Consensus. We only discuss Broadcast in this overview and give the technical details for all models in the subsequent sections.

### Random Trees

Our analysis for URTs proceeds in steps. (A) First we analyze the uniformly random tree model, i.e., the model where the adversary controls none of the edges. (B) Second we allow adversarial, i.e., Byzantine, nodes in the uniformly random tree model. (C) Third we analyze the randomized oblivious message adversary with parameter $k$.

We next sketch the main challenges and how to overcome them. We use $n$ to denote the number of nodes, $I_t$, resp. $S_t$ to denote the set of informed, resp. uninformed nodes after round $t$, and set $N_t = |I_t|$.

(A) When choosing a rooted tree $T$ uniformly at random, there is a high dependence between the events that indicate whether an edge belongs to $T$ or not. Assume that nodes 1, 2, and 3 as well as the edges (1,2) and (2,3) belong to $T$. Then the edge (1,3) cannot belong to $T$. Still, we are able to show that for every node $i \in S_t$ the probability that it is informed in round $t$ is $N_t/n$, independently of whether other nodes are informed or not in round $t$, using the tree counting results discussed before, i.e., $\Delta_t := N_{t+1} - N_t$ follows a binomial distribution with parameters $(n - N_t, N_t/n)$:

▶ **Lemma 7.** *For any $t > 0$, conditioned on $N_t$ $N_{t+1} - N_t$ follows a binomial distribution with parameters $\left(n - N_t, \frac{N_t}{n}\right)$.*

Thus, in expectation, $\Delta_t$ is $(n - N_t)N_t/n$. Now assume for the moment that each round would perform according to its expectation. Then as long as $N_t \leq n/2$, $(n-N_t)N_t/n \geq N_t/2$, i.e., the number of informed nodes increases by a multiplicative factor of at least 3/2 in each round and, thus, there are $O(\log n)$ many rounds. As soon as $N_t > n/2$ then $(n - N_t)N_t/n \geq (n - N_t)/2$, i.e., the distance between the maximum number $n$ and the current number $N_t$ of informed nodes is halved, and, thus, there are at most $O(\log n)$ many rounds.

However, $N_t$ will not increase in every round according to its expectation. Thus, to make this intuition formal we define a random variable $X_t$ for each round $t$ with $X_0 = 1$ that increases by $(n - X_t)X_t/n$ if $\Delta_t$ is at least by its expected value (such a round is called an *increasing* round) and $X_t$ remains unchanged otherwise. It follows from the definition of $X_t$ that it increases monotonically, never reaches $n$, and always lower bounds $N_t$. The number of increasing rounds needed for $X_t$ to reach a value larger than $n - 1$ is at most $2 \ln n$, by a similar argument to the one above. It remains to show that $X_t$ increases frequently. We show that the probability that $X_t$ increases in a round is larger than 1/4, as the binomial variable $\Delta_t$ has a probability larger than 1/4 to be at least at its expectation. Then, Hoeffding's inequality for binomial distributions shows that with probability at least $1 - n^{-c}$ there are more than $2 \ln n$ increasing rounds within the first $32c \ln n$ rounds giving the desired upper bound:

▶ **Theorem 1.** *For any $c \geq 1$ and $n \geq 5$, Broadcast on Uniformly Random Trees completes within $32 \cdot c \cdot \ln n$ rounds with probability $p > 1 - \frac{1}{n^c}$.*

We also show that the bound is asymptotically tight by proving that with constant probability at least $\log n$ rounds are needed. To do so let $Z_t := X_{u_t}$, where $u_t$ is the number of increasing rounds up to round $t$. Thus, intuitively $Z_t$ is $X_t$ with non-increasing rounds omitted. We first show inductively that $\mathbb{E}[N_t] \leq Z_t$. The intuitive reason is that initially $Z_0 = N_0 = \mathbb{E}[N_t]$ and, inductively, in each round $Z_t$ increases by at least as much as $\mathbb{E}[N_t]$. Then we show by induction that $Z_t = n(1 - (n - 1/n))^{2^t}$, which implies that $Z_{\log n} = n(1-(n-1/n))^n \leq n(1-1/4) = 3n/4$. Thus, $\mathbb{E}[N_{\log n}] \leq 3n/4$ and the lower bound follows by applying Markov's inequality:

▶ **Theorem 2.** *If $n \geq 2$, then the probability that Broadcast (and All-to-All Broadcast) on Uniformly Random Trees fails to complete within $\log n$ rounds is at least $\frac{1}{4}$.*

(B) We extend the above model by allowing $f < 2n/3$ *Byzantine nodes* that might forward wrong or no messages, and that can coordinate to make the protocol fail. The process that chooses the communication network, i.e., the random tree, does not know which nodes are Byzantine and, thus, they are part of the network as before, i.e., the tree still consists of $n$ nodes. Furthermore, we assume access to cryptographic tools so that every node can be

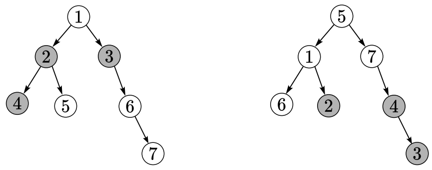


**Figure 2** Shaded nodes are informed nodes. The adversary will choose the right tree over the left tree.

confident about the sender of each message and its content. Here the goal is to inform all $n - f$ honest nodes, i.e., it does not matter whether the Byzantine nodes are informed or not. Almost the same argument as for (A) shows that $N_{t+1} - N_t$ follows a binomial distribution with parameters $(n - f - N_t, N_t/n)$ and also the rest of the analysis, including the lower bound go through.

▶ **Theorem 8.** *For any $c \geq 1$, and $f \leq \frac{2}{3}n - 1$, Broadcast on Uniformly Random Trees with $f$ Byzantine nodes completes within $144 \cdot c \cdot \log n$ rounds with probability $p > 1 - \frac{1}{n^c}$.*

(C) In the *uniformly random trees with adversarial edges model* an adversary chooses first up to $k$ directed edges and then a random tree containing these edges is selected. As before we want to show that the probability that an uninformed node $i$ is informed in round $t$ is independent from other uninformed nodes being informed. This, however, is only true if the adversary uses a *specific optimal* strategy. For an example where the probabilities are not independent, consider a graph with 4 nodes, 2 informed and 2 uninformed. If the adversary introduces an edge from each uninformed node to a different informed node, then for each uninformed node the probability that it is informed in the tree of this round is 1/4. However, the probability that both uninformed nodes are informed in the tree of this round is zero, as only one random edge can be added, which will cause at most one uninformed node to become informed.

(C1) Thus, we first determine the optimal strategy for the adversary: Recall that the adversary wants to maximize the number of rounds. As we show, this implies that a greedy strategy, where the adversary minimizes the increase of $N_t$ in each round $t$, is an optimal strategy for the adversary. To do so, we use a coupling argument comparing the number of informed nodes of the greedy strategy to a non-greedy strategy and showing that a greedy strategy informs all $n$ nodes no later than a non-greedy strategy.

▶ **Lemma 9** (Distribution Domination). *Let $t$ be a round. Let $E_1, E_2$ be two sets of edges the adversaries could choose for round $t$. Let $N_t^{(1)}$ (resp. $I_t^{(1)}$) be the number (resp. set) of informed nodes after round $t$ if $E_1$ is chosen, and $N_t^{(2)}$ (resp. $I_t^{(2)}$) if $E_2$ is chosen. Then if $\mathbb{P}(N_t^{(1)} \geq m) \geq \mathbb{P}(N_t^{(2)} \geq m)$ for every $m \in \mathbb{N}$ (that is, if $N_t^{(1)}$ stochastically dominates $N_t^{(2)}$), then choosing $E_2$ is a better strategy for the adversary than choosing $E_1$.*

Next we analyze what edges are selected by a greedy strategy using three steps: (a) As an edge from an informed node to an uninformed node causes the uninformed node to be informed, the greedy strategy will never put such an edge. Thus, the adversary will only construct trees that do not contain such edges, which we call *non-increasing* trees. This is illustrated in Figure 2.

(b) We show that there is no advantage for the adversary to choose multiple trees. To show this we use a carefully chosen merge operation between any two non-increasing trees that guarantees that the resulting tree is non-increasing together with our new counting theorem for rooted trees. Thus, we can assume that the greedy strategy that is chosen always chooses just one non-increasing rooted tree, which we call $U$. This is illustrated in Figure 3.

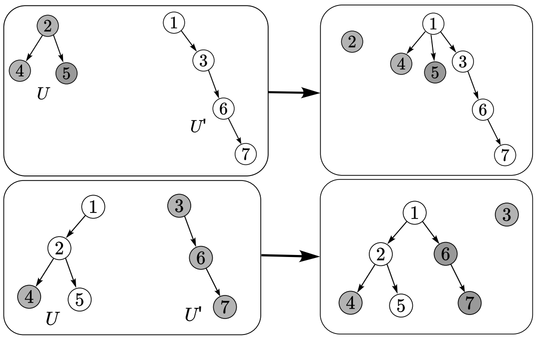


**Figure 3** Merging examples. The adversary will always choose the right option over the left one.

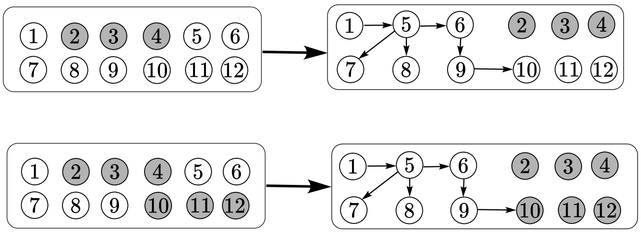

**Figure 4** The best strategy for the adversary $A$, with $k = 6$. Shaded nodes are informed nodes. In the top example, nodes $5, 6, 7, 8, 9$ and 10 are safe from being informed, whereas node 1 can still be informed. In the bottom example, nodes $5, 6, 7, 8,$ and 9 are safe, whereas node 1 can still be informed. However, node 1 is safe from being informed by node 10.

(c) We then argue that $U$ should contain as many uniformed nodes as possible. The basic intuition is that if an uninformed node is the child of another uninformed node, it cannot become informed in this round, i.e., it is "protected". Given $k$ edges, the adversary "protects" as many uninformed nodes as possible by building $U$ from $\min(|S_t|, k+1)$ uninformed and $\max(k+1-|S_t|, 0)$ informed nodes. The fact that $U$ is also non-increasing implies that the root of $U$ is an uninformed node. This gives the optimal strategy, illustrated in Figure 4. We show all the above steps using stochastic dominance.

(C2) Finally we show that with this specific optimal strategy, the adversary can only force $O(k + \log n)$ many rounds with high probability. It follows that with high probability, the adversary cannot achieve more rounds with any other - optimal or non-optimal - strategy. To do so, we break the rounds into phases: (i) The first phase consists of all rounds where $|S_k| \geq k+1$. In this case, the adversary forms one tree with $k+1$ uninformed nodes and there are $|S_k| - k - 1 = n - N_t - k - 1$ additional uniformed singleton nodes, as well as $N_t$ informed singleton nodes in the forest. Thus, we can apply exactly the same argument as in (A) to show that $N_{t+1} - N_t$ follows a binomial distribution with parameters $(n - k - N_t, N_t/n)$. (ii) The second phase consists of all rounds where $|S_k| \leq k$. Thus $U$ consists of all uninformed nodes and at least one informed node. Thus, $N_{t+1} - N_t$ can increase by at most 1, namely if the root $s$ of $U$ receives a parent in the tree, and, using our new counting theorem for rooted trees, we show that the probability of that is $(N_t - (k + 1 - |S_t|))/n = (n - k - 1)/n$, i.e. $N_{t+1} - N_t$ is a binomial distribution with parameters $(1, (n - k - 1)/n)$. Using Hoeffding's inequality for binomial distributions similar to (A) we then show the result:

▶ **Theorem 10.** *If the adversary controls $k$ edges in each round, for $k \leq \frac{2}{3}n - 1$, then for any $c \geq 1$, with probability $p \geq 1 - n^{-c}$, Broadcast completes within $O(k + \log n)$ rounds.*

### Directed Erdős–Rényi graphs

Directed Erdős–Rényi graphs consist of $m$ edges chosen uniformly at random among the $n^2$ potential edges. Intuitively they have less structure than uniformly random trees, which makes the analysis of Broadcast simpler. We present the main ideas below. Note that we also analyze Byzantine nodes and adversarial edges in that model, but omit these extensions in this overview.

Sampling a directed Erdős–Rényi graph is equivalent to choosing $m$ edges *without* replacement from the set of all possible edges. We call that Scheme 1. Then we observe, using a coupling argument, that Scheme 1 requires no more rounds than Scheme 2, where in each round $m$ edges are chosen *with* replacement. Finally, to analyze Scheme 2, we basically partition the sequence of rounds of Scheme 2 into $2\lceil(\log n)/2\rceil$ phases, such that for each of the first $\lceil(\log n)/2\rceil$ phases the number of informed nodes doubles in each phase and for each of the last $\lceil(\log n)/2\rceil$ phases the number of uninformed nodes halves in each phase. Note that Broadcast completes after the last phase. Using Hoeffding's inequality for binomial distributions we show that phase $i$ for $1 \leq i \leq \lceil\log n/2\rceil$ requires with high probability at most $O(\max\{\log n, 2^{i-1}\}n/2^{i-1})$ sampled edges, and, thus, $O(\lceil\max\{\log n, 2^{i-1}\}/(2^{i-1}m/n)\rceil)$ rounds, and for $\lceil\log n/2\rceil + 1 \leq i \leq 2\lceil\log n/2\rceil$ phase $i$ requires with high probability at most $O(\max\{\log n, 2^{j-2}\}n/2^{j-1})$ sampled edges with with $j := 2\lceil\log n/2\rceil - i$, and, thus, $O(\lceil\max\{\log n, 2^{j-1}\}/(2^{j-1}m/n)\rceil)$ rounds. Summed over all phases this shows that with high probability $O(\lceil n/m\rceil \log n)$ rounds suffice for Scheme 2 to reach Broadcast. Note that the analysis extends to the setting when the graph in each round contains *at least* $m$ edges. We also show that a lower bound that implies that this upper bound is tight for $m \leq n$. We also give somewhat different analysis where the number of informed resp. uninformed nodes does not double, but increases by $(1 + m/n)$ that is tight for $m \geq n \ln n$. Our results can thus be summarized by the following theorems:

▶ **Theorem 11.** *For any $c \geq 1$, in scheme 2, and therefore scheme 1, Broadcast completes within $O\left(\left\lceil\frac{cn}{m}\right\rceil \log n\right)$ rounds with probability $p \geq 1 - n^{-c}\log n$.*

▶ **Theorem 12.** *For any $c \geq 1$ and $m \in [n^2]$ such that $m/n \geq \ln n$, in scheme 2 and in scheme 1, Broadcast completes within $O\left(\frac{c\cdot\log n}{\log(1+m/n)}\right)$ rounds with probability $p \geq 1 - n^{-c}\log n$.*

▶ **Theorem 13.** *In scheme 1, and thus in scheme 2, Broadcast fails to complete within $\frac{\log(n)-1}{\log(1+m/n)}$ rounds with probability at least $\frac{1}{2}$.*

## References


1 Mohamad Ahmadi, Fabian Kuhn, Shay Kutten, Anisur Rahaman Molla, and Gopal Pandurangan. The communication cost of information spreading in dynamic networks. In *39th IEEE International Conference on Distributed Computing Systems, ICDCS 2019, Dallas, TX, USA, July 7-10, 2019*, pages 368–378. IEEE, 2019. doi:10.1109/ICDCS.2019.00044.

2 John Augustine, Gopal Pandurangan, Peter Robinson, and Eli Upfal. Towards robust and efficient computation in dynamic peer-to-peer networks. In Yuval Rabani, editor, *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 551–569. SIAM, 2012. doi:10.1137/1.9781611973099.47.

3 Piotr Berman, Juan A. Garay, and Kenneth J. Perry. Towards optimal distributed consensus (extended abstract). In *30th Annual Symposium on Foundations of Computer Science, Research Triangle Park, North Carolina, USA, 30 October - 1 November 1989*, pages 410–415. IEEE Computer Society, 1989. doi:10.1109/SFCS.1989.63511.

4 Bernadette Charron-Bost, Matthias Függer, and Thomas Nowak. Approximate consensus in highly dynamic networks: The role of averaging algorithms. In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part II*, volume 9135 of *Lecture Notes in Computer Science*, pages 528–539. Springer, 2015. doi:10.1007/978-3-662-47666-6_42.

5 Bernadette Charron-Bost and André Schiper. The heard-of model: computing in distributed systems with benign faults. *Distributed Comput.*, 22(1):49–71, 2009. doi:10.1007/s00446-009-0084-6.

6 Andrea E. F. Clementi, Pierluigi Crescenzi, Carola Doerr, Pierre Fraigniaud, Francesco Pasquale, and Riccardo Silvestri. Rumor spreading in random evolving graphs. *Random Struct. Algorithms*, 48(2):290–312, 2016. doi:10.1002/rsa.20586.

7 Étienne Coulouma, Emmanuel Godard, and Joseph G. Peters. A characterization of oblivious message adversaries for which consensus is solvable. *Theor. Comput. Sci.*, 584:80–90, 2015. doi:10.1016/j.tcs.2015.01.024.

8 Michael Dinitz, Jeremy T. Fineman, Seth Gilbert, and Calvin Newport. Smoothed analysis of information spreading in dynamic networks. In Christian Scheideler, editor, *36th International Symposium on Distributed Computing, DISC 2022, October 25-27, 2022, Augusta, Georgia, USA*, volume 246 of *LIPIcs*, pages 18:1–18:22. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICS.DISC.2022.18.

9 Stefan Dobrev and Imrich Vrto. Optimal broadcasting in hypercubes with dynamic faults. *Inf. Process. Lett.*, 71(2):81–85, 1999. doi:10.1016/S0020-0190(99)00093-9.

10 Stefan Dobrev and Imrich Vrto. Optimal broadcasting in tori with dynamic faults. *Parallel Process. Lett.*, 12(1):17–22, 2002. doi:10.1142/S0129626402000781.

11 Benjamin Doerr and Mahmoud Fouz. Asymptotically optimal randomized rumor spreading. In *International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 502–513. Springer, 2011. doi:10.1007/978-3-642-22012-8_40.

12 Danny Dolev and H. Raymond Strong. Authenticated algorithms for byzantine agreement. *SIAM J. Comput.*, 12(4):656–666, 1983. doi:10.1137/0212045.

13 Rick Durrett and Dong Yao. Susceptible–infected epidemics on evolving graphs. *Electronic Journal of Probability*, 27:1–66, 2022.

14 Chinmoy Dutta, Gopal Pandurangan, Rajmohan Rajaraman, Zhifeng Sun, and Emanuele Viola. On the complexity of information spreading in dynamic networks. In Sanjeev Khanna, editor, *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 717–736. SIAM, 2013. doi:10.1137/1.9781611973105.52.

15 Antoine El-Hayek, Monika Henzinger, and Stefan Schmid. Brief announcement: Broadcasting time in dynamic rooted trees is linear. In *Proc. ACM Symposium on Principles of Distributed Computing (PODC)*, 2022.

16 Antoine El-Hayek, Monika Henzinger, and Stefan Schmid. Asymptotically tight bounds on the time complexity of broadcast and its variants in dynamic networks. In *14th Innovations in Theoretical Computer Science (ITCS)*, 2023.

17 Faith Ellen, Barun Gorain, Avery Miller, and Andrzej Pelc. Constant-length labeling schemes for deterministic radio broadcast. *ACM Trans. Parallel Comput.*, 8(3):14:1–14:17, 2021. doi:10.1145/3470633.

18 Patrick T Eugster, Rachid Guerraoui, A-M Kermarrec, and Laurent Massoulié. Epidemic information dissemination in distributed systems. *Computer*, 37(5):60–67, 2004. doi:10.1109/MC.2004.1297243.

19 Pierre Fraigniaud and Emmanuel Lazard. Methods and problems of communication in usual networks. *Discret. Appl. Math.*, 53(1-3):79–133, 1994. doi:10.1016/0166-218X(94)90180-5.

**20** Matthias Függer, Thomas Nowak, and Kyrill Winkler. On the radius of nonsplit graphs and information dissemination in dynamic networks. *Discret. Appl. Math.*, 282:257–264, 2020. doi:10.1016/j.dam.2020.02.013.

**21** Hugo Rincon Galeana, Ami Paz, Stefan Schmid, Ulrich Schmid, and Kyrill Winkler. The time complexity of consensus under oblivious message adversaries. In *14th Innovations in Theoretical Computer Science (ITCS)*, 2023.

**22** Hugo Rincon Galeana, Ulrich Schmid, Kyrill Winkler, Ami Paz, and Stefan Schmid. Topological characterization of consensus solvability in directed dynamic networks. *arXiv preprint arXiv:2304.02316*, 2023. doi:10.48550/arXiv.2304.02316.

**23** Mohsen Ghaffari, Fabian Kuhn, and Hsin-Hao Su. Distributed MST and routing in almost mixing time. In Elad Michael Schiller and Alexander A. Schwarzmann, editors, *Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC 2017, Washington, DC, USA, July 25-27, 2017*, pages 131–140. ACM, 2017. doi:10.1145/3087801.3087827.

**24** Juraj Hromkovič, Ralf Klasing, Burkhard Monien, and Regine Peine. Dissemination of information in interconnection networks (broadcasting & gossiping). In *Combinatorial network theory*, pages 125–212. Springer, 1996.

**25** Fabian Kuhn, Nancy A. Lynch, and Rotem Oshman. Distributed computation in dynamic networks. In Leonard J. Schulman, editor, *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 513–522. ACM, 2010. doi:10.1145/1806689.1806760.

**26** Linyuan Lu, Austin Mohr, and László Székely. Quest for negative dependency graphs. In *Recent Advances in Harmonic Analysis and Applications*, pages 243–258. Springer, 2012.

**27** Uri Meir, Ami Paz, and Gregory Schwartzman. Models of smoothing in dynamic networks. In Hagit Attiya, editor, *34th International Symposium on Distributed Computing, DISC 2020, October 12-16, 2020, Virtual Conference*, volume 179 of *LIPIcs*, pages 36:1–36:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPICS.DISC.2020.36.

**28** James D Murray et al. Mathematical biology i: an introduction, 2002.

**29** Jim Pitman. Coalescent random forests. *Journal of Combinatorial Theory, Series A*, 85(2):165–193, 1999.

**30** Martin Zeiner, Manfred Schwarz, and Ulrich Schmid. On linear-time data dissemination in dynamic rooted trees. *Discret. Appl. Math.*, 255:307–319, 2019. doi:10.1016/j.dam.2018.08.015.