

Asymptotically Tight Bounds on the Time Complexity of Broadcast and its Variants in Dynamic Networks *

Antoine El-Hayek
Faculty of Computer Science
University of Vienna, Austria

Monika Henzinger
Faculty of Computer Science
University of Vienna, Austria

Stefan Schmid
TU Berlin, Germany
Fraunhofer SIT, Germany

Abstract

Data dissemination is a fundamental task in distributed computing. This paper studies *broadcast problems* in various innovative models where the communication network connecting n processes is dynamic (e.g., due to mobility or failures) and controlled by an adversary.

In the first model, the processes transitively communicate their ids in synchronous rounds along a rooted tree given in each round by the adversary whose goal is to maximize the number of rounds until *at least one id is known by all processes*. Previous research has shown a $\lceil \frac{3n-1}{2} \rceil - 2$ lower bound and an $O(n \log \log n)$ upper bound. We show the first linear upper bound for this problem, namely $\lceil (1 + \sqrt{2})n - 1 \rceil \approx 2.4n$.


We extend these results to the setting where the adversary gives in each round k -disjoint forests and their goal is to maximize the number of rounds until there is a set of k ids such that *each process knows of at least one of them*. We give a $\lceil \frac{3(n-k)}{2} \rceil - 1$ lower bound and a $\frac{\pi^2+6}{6}n + 1 \approx 2.6n$ upper bound for this problem.

Finally, we study the setting where the adversary gives in each round a directed graph with k roots and their goal is to maximize the number of rounds until *there exist k ids that are known by all processes*. We give a $\lceil \frac{3(n-3k)}{2} \rceil + 2$ lower bound and a $\lceil (1 + \sqrt{2})n \rceil + k - 1 \approx 2.4n + k$ upper bound for this problem.

For the two latter problems no upper or lower bounds were previously known.

1 Introduction

Data dissemination is one of the most fundamental tasks in distributed systems. This paper studies data dissemination in an innovative model where the communication network connecting n processes is *dynamic*. In particular, we consider a worst-case perspective and assume that the information flow between the processes is controlled by an *oblivious message adversary* which may drop an arbitrary set of messages sent by some processes in each round. This results in a

*This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 101019564) . This work was further supported by the Austrian Science Fund (FWF) and netIDEE SCIENCE project P 33775-N, as well as the FWF project I 4800-N (ADVISE).

sequence of directed communication graphs, whose edges tell which process can successfully send a message to which other process in a given round. The oblivious message adversary 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 subject to interference, jamming, or mobility.

We look into three data dissemination problems in dynamic networks: *broadcast*, *cover* and *k-broadcast*. These problems come in many flavors and feature intriguing connections to other classic problems such as leader(s) election, regular and *k*-set consensus (also known as *k*-set agreement), for which our problems' time complexity is typically a lower bound.

In particular, we assume that each process has a unique id and in every message each process communicates all the ids it knows of so far. We first study a fundamental model where communication happens along arbitrary rooted trees, chosen by an adversary who aims to maximize the *broadcast time*, which is the number of rounds it takes until there exists a process that everyone knows of. We then extend our investigations to sparser networks, considering *k*-forests (a union of *k* rooted trees). Here the adversary will maximize the *cover time*, which is the number of rounds it takes until there exists *k* process such that everyone knows of at least one of them. Moreover, in more highly connected networks, we study *k*-rooted dynamic networks (directed graphs with *k* roots), where the adversary aims at maximizing the *k-broadcast time*, which is the number of rounds it takes until there exist *k* processes that everyone knows of. Before presenting our results, we introduce our model more formally.

Model Let n be the number of processes and let each process have a unique identifier from $[n]$. Let \mathcal{G} be a fixed set of directed networks with n nodes such that each node has a unique identifier from $[n]$. There will be a sequence of rounds $t = 1, 2, \dots$, such that, in each round t , an adversary chooses a network \tilde{G}_t from \mathcal{G} , which determines the communication links in round t as follows: Initially every process *knows* (or *has heard*) of its own identifier. During round t process i sends all identifiers it knows of to its out-neighbors in \tilde{G}_t . The rounds stop whenever the objective – broadcast, *k*-broadcast or cover of size k – is attained. The goal of the adversary is to maximize the number of rounds.

To model the information propagation we use graph products:

Definition 1.1. *If $A = ([n], E_1)$ and $B = ([n], E_2)$ are two directed networks on n nodes, then the product graph $A \circ B$ is the network on n nodes, with edge set E , where $(x, y) \in E$ if and only if there exist a node z such that $(x, z) \in E_1$ and $(z, y) \in E_2$.*

Consider round t and let $G(t)$ be the *product graph* $G(t) = G_1 \circ \dots \circ G_t$, where G_i is created from \tilde{G}_i by adding a self-loop to every node. Note that in $G(t)$ the in-neighbors of a process x are exactly the processes that x knows of after t rounds, and its out-neighbors are the processes it has sent information to. We added a self-loop to every node in G_i to capture the fact that no process “forgets” any piece of information in any round. An example is given in Figure 1.

We will consider three different models, each of which has a different objective, detailed below. Figure 2 summarizes the 3 models, gives examples and states the results.

Broadcasting on Trees

Definition 1.2. *Let \mathcal{T}_n be the set of all rooted trees with a self-loop added at every node.*

In this variant, the networks given by the adversary are restricted to trees in \mathcal{T}_n and we analyze the *broadcasting time* t^* , which is the smallest round t such that there exists a node




Figure 1: Example of information propagation. Self-loops are omitted, but are present on every node in every round. $x : yz$ means that y was an in-neighbor of x before the current round (and still is), whereas z is a new in-neighbor of x . We omit mentioning x in its own in-neighbors as it is always the case.

in $G(t)$ with an out-edge to every other node. Note that this corresponds to a process such that every other process has heard of its identifier. We will say that *the node has broadcast (its identifier to everyone)* or simply *broadcast has happened*. Trees being rooted ensure that broadcast happens in a finite number of rounds¹.

Definition 1.3. *The broadcast time $t^*(G_1, G_2, \dots)$ of a sequence of graphs G_1, G_2, \dots , is*

$$t^*(G_1, G_2, \dots) = \min\{t \in \mathbb{N} : \exists x \in [n], \forall y \in [n], (x, y) \in G_1 \circ \dots \circ G_t\}$$

Definition 1.4. *The broadcast time $t^*(\mathcal{G})$ of an adversary \mathcal{G} , is defined as follows:*

$$t^*(\mathcal{G}) = \max\{t^*(G_1, G_2, \dots) : \forall i \in \mathbb{N}, G_i \in \mathcal{G}\}$$

We will give tight asymptotic bounds on $t^*(\mathcal{T}_n)$. Note that even in the simple case where the adversary gives the same directed tree in each round, the broadcast time can be as large as $n - 1$, namely if the tree is simply a path. Conversely, in each round, it is easy to see that at least one new edge appears in the product graph², and thus the broadcast time is at most n^2 . This raises the question of how large the broadcast time can be made if in each round a different directed tree can be used.

This has been an open question for several years. Results from Charron-Bost and Schiper in 2009 [4] and Charron-Bost, Függer, and Nowak in 2015 [3] imply an $n \log n$ upper bound. In 2019, Zeiner, Schwarz, and Schmid [18] gave a linear upper bound when the adversary is restricted to trees with either a constant number of leaves or a constant number of inner nodes. They also gave a $\lceil \frac{3n-1}{2} \rceil - 2$ lower bound. In 2020, Függer, Nowak, and Winkler [12] improved the general upper bound to $2n \log \log n + O(n)$. So far, it has been an open conjecture [18] whether the broadcast time is linear for arbitrary sequences of rooted trees.

Covering on k -forests

Definition 1.5. *We define \mathcal{F}_n^k to be the set of all forests over n processes which are the union of k rooted trees and a self-loop is added at every node.*

In this variant, the networks given by the adversary are restricted to networks from \mathcal{F}_n^k , and we analyze the *cover time* t_k^c , which is the smallest round t such that there exists a set $I \subset [n]$, $|I| \leq k$, such that every node x of $G(t)$ has at least one in-neighbor from I . Said differently, there exists a set of k nodes such that every node knows of *at least one* of them.

¹If the adversary can choose a non-rooted graph, it could repeat this graph indefinitely, preventing broadcast.

²In each round, the identifier of the root reaches someone new.

Definition 1.6. The cover time $t_k^c(G_1, G_2, \dots)$ of a sequence of graphs G_1, G_2, \dots , is

$$t_k^c(G_1, G_2, \dots) = \min\{t \in \mathbb{N} : \exists x_1, \dots, x_k \in [n], \forall y \in [n], \exists i \in [k], (x_i, y) \in G_1 \circ \dots \circ G_t\}$$

Definition 1.7. The cover time $t_k^c(\mathcal{G})$ of an adversary \mathcal{G} , is defined as follows:

$$t_k^c(\mathcal{G}) = \max\{t_k^c(G_1, G_2, \dots) : \forall i \in \mathbb{N}, G_i \in \mathcal{G}\}$$

We will give tight asymptotic bounds on $t_k^c(\mathcal{F}_n^k)$. To the best of our knowledge, there is no prior work that gives upper or lower bounds on $t_k^c(\mathcal{F}_n^k)$.

k -Broadcasting on k -rooted Networks

Definition 1.8. Let \mathcal{R}_n^k be the set of all (directed) networks over n nodes that (1) have k roots, that is k different processes r_1, \dots, r_k such that there exists a directed path from any r_i to any process in $[n]$, and (2) have a self-loop at every node.

In this variant, the networks given by the adversary are restricted to networks from \mathcal{R}_n^k , and we analyze the k -broadcasting time t_k^* , which is the smallest round t such that there exist k nodes in $G(t)$ with an out-edge to every other node. Said differently, there exists a set of k nodes such that every node knows of *all* of them.

Definition 1.9. The k -broadcast time $t_k^*(G_1, G_2, \dots)$ of a sequence of graphs G_1, G_2, \dots , is

$$t_k^*(G_1, G_2, \dots) = \min\{t \in \mathbb{N} : \exists x_1, \dots, x_k \in [n], \forall i \neq j, x_i \neq x_j \\ \wedge \forall y \in [n], \forall i \in [k], (x_i, y) \in G_1 \circ \dots \circ G_t\}$$

Definition 1.10. The k -broadcast time $t_k^*(\mathcal{G})$ of an adversary \mathcal{G} , is defined as follows:

$$t_k^*(\mathcal{G}) = \max\{t_k^*(G_1, G_2, \dots) : \forall i \in \mathbb{N}, G_i \in \mathcal{G}\}$$

We will give tight asymptotic bounds on $t_k^*(\mathcal{R}_n^k)$. To the best of our knowledge, there is no prior work that gives upper or lower bounds on $t_k^*(\mathcal{R}_n^k)$.

Contribution We give asymptotically tight bounds for all three settings, see also Figure 2.

(1) First we settle the open problem about time complexity of broadcast in dynamic trees, by showing that it is *linear*. Hence, Zeiner et al.'s [18] conjecture is true. In particular, we present an upper bound of $\lceil (1 + \sqrt{2})n \rceil \approx 2.4n$, which complements their $\lceil \frac{3n-1}{2} \rceil - 2$ lower bound.

(2) We further show that covering on k -forests also takes linear time, by giving a $\frac{\pi^2+6}{6}n + 1 \approx 2.6n$ upper bound and a $\lceil \frac{3(n-k)}{2} \rceil - 1$ lower bound.

(3) Finally, we show that k -broadcasting on k -rooted networks is linear as well, by giving an upper bound of $\lceil (1 + \sqrt{2})n \rceil + k - 1 \approx 2.4n + k$, and a lower bound of $\lceil \frac{3(n-3k)}{2} \rceil + 2$.

Organization The remainder of this paper is organized as follows. Section 2 introduces some basic tools that will be useful throughout the paper, and presents first insights. In Section 3, we give the linear upper bound for broadcast time in our model. Section 4 and Section 5 respectively showcase our results for the cover on k -forests and the k -broadcast on k -rooted networks. After reviewing related work in Section 6, we conclude and provide some future research directions in Section 7. All the lower bound results can be found in appendix A, as well as other basic proofs.

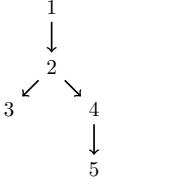
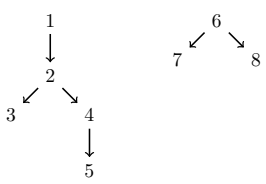
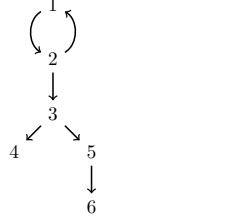
| Model: | Broadcasting on Trees | Covering on k -Forests | k -Broadcasting on k -rooted Networks |
|--------------|---|--|---|
| Adversary: |  <p>Rooted trees over n nodes. Example for $n = 5$</p> |  <p>Forests over n processes, composed of k rooted trees. Example for $n = 8$ and $k = 2$</p> |  <p>Networks over n processes and k roots. Example for $n = 6$ and $k = 2$</p> |
| Objective: | <p>1 : <u>12</u></p> <p>2 : <u>23</u></p> <p>3 : <u>23</u></p> <p>4 : <u>234</u></p> <p>Broadcast Example for $n = 4$</p> | <p>1 : <u>12</u></p> <p>2 : <u>24</u></p> <p>3 : <u>123</u></p> <p>4 : <u>34</u></p> <p>Cover of size k Example for $n = 4$ and $k = 2$</p> | <p>1 : <u>124</u></p> <p>2 : <u>24</u></p> <p>3 : <u>234</u></p> <p>4 : <u>124</u></p> <p>k-broadcast of k processes Example for $n = 4$ and $k = 2$</p> |
| Lower Bound: | [18]: $\lceil \frac{3n-1}{2} \rceil - 2$ | $\lceil \frac{3(n-k)}{2} \rceil - 1$ | $\lceil \frac{3(n-3k)}{2} \rceil + 2$ |
| Upper Bound: | [12]: $O(n \log \log n)$ $\lceil (1 + \sqrt{2})n \rceil$ | $\frac{\pi^2+6}{6}n + 1$ | $\lceil (1 + \sqrt{2})n \rceil + k - 1$ |

Figure 2: Summary of models and results. The adversary chooses a sequence of graphs as stated in order to delay the objective as much as possible. In the objective examples, $x : y$ means that y is an in-neighbour of x at the round the objective is attained. Underlines are for emphasis purposes only.

2 Basic Tools

In this section, we define generalizations of the in- and out-neighborhoods in the product graph, and some basic properties these tools follow. The proofs of these properties being basic, we defer them to the appendix.

Definition 2.1. *The in-neighborhood of process x between rounds t and t' , $t, t' \geq 0$, denoted by $\mathcal{I}_t^{t'}(x)$, is defined as follows: If $t \leq t'$, $\mathcal{I}_t^{t'}(x)$ is the in-neighborhood of process x in the graph $G_t \circ \dots \circ G_{t'}$. If $t = t' + 1$, set $\mathcal{I}_t^{t'}(x) = \{x\}$. Otherwise, $\mathcal{I}_t^{t'}(x) = \emptyset$.*

Definition 2.2. *The out-neighborhood of process x between rounds t and t' , $t, t' \geq 0$, denoted by $\mathcal{O}_t^{t'}(x)$, is defined as follows: If $t \leq t'$, $\mathcal{O}_t^{t'}(x)$ is the out-neighborhood of process x in the graph $G_t \circ \dots \circ G_{t'}$. If $t = t' + 1$, set $\mathcal{O}_t^{t'}(x) = \{x\}$. Otherwise, $\mathcal{O}_t^{t'}(x) = \emptyset$.*

We prove in Appendix A that these generalized neighborhoods behave as expected:

Lemma 2.3. *Let $t, t' \geq 0, x, y \in [n]$. Then $x \in \mathcal{O}_t^{t'}(y) \Leftrightarrow y \in \mathcal{I}_t^{t'}(x)$.*

Lemma 2.4 (Transitivity). *Let $t, t', t'' \geq 0$, and $x, y, z \in [n]$. We have the following properties:*

- i. *If $y \in \mathcal{O}_t^{t'}(x)$ and $z \in \mathcal{O}_{t'+1}^{t''}(y)$, then $z \in \mathcal{O}_t^{t''}(x)$.*
- ii. *If $x \in \mathcal{I}_t^{t'}(y)$ and $z \in \mathcal{O}_{t'+1}^{t''}(y)$, then $z \in \mathcal{O}_t^{t''}(x)$.*
- iii. *If $x \in \mathcal{I}_t^{t'}(y)$ and $y \in \mathcal{I}_{t'+1}^{t''}(z)$, then $z \in \mathcal{O}_t^{t''}(x)$.*

And finally, we show that these neighborhoods only grow over time:

Lemma 2.5 (Monotonicity). *If in each round, all nodes have a self-loop, then for any $t_1 \leq t_2$ and $t_3 \leq t_4$, for any process x we have:*

- i. $\mathcal{I}_{t_2}^{t_3}(x) \subseteq \mathcal{I}_{t_1}^{t_4}(x)$.
- ii. $\mathcal{O}_{t_2}^{t_3}(x) \subseteq \mathcal{O}_{t_1}^{t_4}(x)$.

3 Broadcasting on Trees

In this section, we focus on the fundamental problem of broadcasting on dynamic trees. We give an upper bound for the problem, before recalling a lower bound.

3.1 The Upper Bound

We will show that the key to understand how information propagate is to consider what the root knows – or the in-neighbors of the root – before the beginning of every round. We will show that the root must either have a lot of in-neighbors that were roots in previous rounds, or many in-neighbors in general. We will then show that any in-neighbor of the root before a round has at least one more out-neighbor after the round than before it. We will finally show that any adversary that tries to balance these two facts will fail to prevent broadcast for a time longer than linear.

Definition 3.1. *Let t be a round. We denote by r_t the root of G_t , and call it the root of the round t .*

Harnessing the fact that there always exists a path from a root to any other process in a network, we give the two following lemmas:

Lemma 3.2. *Let t, t' be rounds such that $t \leq t'$. We have that:*

i *If x is a process such that $r_t \notin \mathcal{I}_{t+1}^{t'}(x)$, then $|\mathcal{I}_t^{t'}(x)| > |\mathcal{I}_{t+1}^{t'}(x)|$.*

ii *If x is a process such that $r_{t'} \in \mathcal{O}_t^{t'-1}(x)$, then $|\mathcal{O}_t^{t'}(x)| > |\mathcal{O}_t^{t'-1}(x)|$, unless $|\mathcal{O}_t^{t'-1}(x)| = n$.*

Proof. i. We will show that there exists a process $y \in \mathcal{I}_{t+1}^{t'}(x)$ that has an in-neighbor z in G_t such that $z \notin \mathcal{I}_{t+1}^{t'}(x)$. Then, by Transitivity (Lemma 2.4), $z \in \mathcal{I}_t^{t'}(x)$. By Monotonicity (Lemma 2.5), we have $\mathcal{I}_{t+1}^{t'}(x) \subset \mathcal{I}_t^{t'}(x)$, this will show that $|\mathcal{I}_t^{t'}(x)| > |\mathcal{I}_{t+1}^{t'}(x)|$.

Let us now find such a z . Consider the path from r_t to x in G_t . Since $r_t \notin \mathcal{I}_{t+1}^{t'}(x)$, and trivially $x \in \mathcal{I}_{t+1}^{t'}(x)$, this path must include an edge (z, y) such that $z \notin \mathcal{I}_{t+1}^{t'}(x)$, $y \in \mathcal{I}_{t+1}^{t'}(x)$.

ii. Let us look at the case $|\mathcal{O}_t^{t'-1}(x)| < n$. We will show that there exists a process $y \in \mathcal{O}_t^{t'-1}(x)$ that has an out-neighbor z in $G_{t'}$ such that $z \notin \mathcal{O}_t^{t'-1}(x)$. Then, by Transitivity (Lemma 2.4), $z \in \mathcal{O}_t^{t'}(x)$. By Monotonicity (Lemma 2.5), we have $\mathcal{O}_t^{t'-1}(x) \subset \mathcal{O}_t^{t'}(x)$, this will show that $|\mathcal{O}_t^{t'}(x)| > |\mathcal{O}_t^{t'-1}(x)|$.

Let us now find such a z . Since $|\mathcal{O}_t^{t'-1}(x)| < n$, there exists a process a such that $a \notin \mathcal{O}_t^{t'-1}(x)$. Consider the path from $r_{t'}$ to a in $G_{t'}$. Since $r_{t'} \in \mathcal{O}_t^{t'-1}(x)$, and $a \notin \mathcal{O}_t^{t'-1}(x)$, this path must include an edge (y, z) such that $z \notin \mathcal{O}_t^{t'-1}(x)$, $y \in \mathcal{O}_t^{t'-1}(x)$. \square

The following lemma will link the number of in-neighbors a node has to the number of in-neighbors it has among the roots of the preceding rounds:

Lemma 3.3. *Let x be a process, and t_1, t_2 be rounds such that $t_1 \leq t_2$. Then:*

$$|\{t : t_1 \leq t \leq t_2, r_t \notin \mathcal{I}_{t_1}^{t_2}(x)\}| + 1 \leq |\mathcal{I}_{t_1}^{t_2}(x)|$$

This will be proven using Lemma 3.2.i, since every time $r_t \notin \mathcal{I}_{t_1}^{t_2}(x)$, $\mathcal{I}_{t_1}^{t_2}(x)$ gets larger:

Proof. Let $A = \{t : t_1 \leq t \leq t_2, r_t \notin \mathcal{I}_{t_1}^{t_2}(x)\}$. Then, in particular, for any $t \in A$, we have $r_t \notin \mathcal{I}_{t+1}^{t_2}(x)$. Then, for all $t \in A$, applying Lemma 3.2.i, we have $|\mathcal{I}_t^{t_2}(x)| > |\mathcal{I}_{t+1}^{t_2}(x)|$. Let $A = \{t^1, \dots, t^k\}$, with $t^i < t^{i+1}$ for any $1 \leq i < k$. Then:

$$|\mathcal{I}_{t_1}^{t_2}(x)| \geq |\mathcal{I}_{t^1}^{t_2}(x)| > |\mathcal{I}_{t^1+1}^{t_2}(x)| \geq |\mathcal{I}_{t^2}^{t_2}(x)| > |\mathcal{I}_{t^2+1}^{t_2}(x)| \geq \dots > |\mathcal{I}_{t^k+1}^{t_2}(x)| \geq 1$$

Where the non-strict inequalities derive by Monotonicity(Lemma 2.5), and the last one from the fact that $t^k + 1 \leq t_2 + 1 \Rightarrow x \in \mathcal{I}_{t^k+1}^{t_2}(x)$. There are $k = |A|$ strict inequalities over integers, which concludes the proof. \square

We now define the *rounds graph*, which will keep track of the information – the in-neighbors – the root of each round has:

Definition 3.4. *We define the rounds graph as follows:*

The graph has $2n + \lceil \sqrt{2n} \rceil$ nodes: one node representing each process, and one node for each of the first $\lceil (1 + \sqrt{2})n \rceil$ rounds.

And it has the directed edges: one edge from a process p to a round t if $p \in \mathcal{I}_1^{t-1}(r_t)$, and one edge from a round $t < \lceil \sqrt{2n} \rceil$ to a round $t' > t$ if $r_t \in \mathcal{I}_1^{t'-1}(r_{t'})$.

We will now show that there is at least a node of out-degree n in that graph, which will translate into a process that has broadcast its piece of information to everyone:

Lemma 3.5. *In the rounds graph, there is at least a node of out-degree n .*

Proof. Let us look at round t . If $t \leq \lceil \sqrt{2n} \rceil$, then t has in-degree at least t . Indeed, it has in-degree:

$$\begin{aligned} & |\{t' : 1 \leq t' \leq t-1, r_{t'} \in \mathcal{I}_1^{t-1}(r_t)\}| + |\mathcal{I}_1^{t-1}(x)| \\ & \geq 1 + |\{t' : 1 \leq t' \leq t-1, r_{t'} \in \mathcal{I}_1^{t-1}(r_t)\}| + |\{t' : 1 \leq t' \leq t-1, r_{t'} \notin \mathcal{I}_1^{t-1}(r_t)\}| = t \end{aligned}$$

where we used Lemma 3.3 for the inequality. Similarly, if $t > \lceil \sqrt{2n} \rceil$, then t has in-degree at least $\lceil \sqrt{2n} \rceil$. Indeed, it has in-degree:

$$\begin{aligned} & |\{t' : 1 \leq t' < \lceil \sqrt{2n} \rceil, r_{t'} \in \mathcal{I}_1^{t-1}(r_t)\}| + |\mathcal{I}_1^{t-1}(r_t)| \\ & \geq 1 + |\{t' : 1 \leq t' < \lceil \sqrt{2n} \rceil, r_{t'} \in \mathcal{I}_1^{t-1}(r_t)\}| + |\{t' : 1 \leq t' \leq t-1, r_{t'} \notin \mathcal{I}_1^{t-1}(r_t)\}| \\ & \geq 1 + |\{t' : 1 \leq t' < \lceil \sqrt{2n} \rceil, r_{t'} \in \mathcal{I}_1^{t-1}(r_t)\}| + |\{t' : 1 \leq t' < \lceil \sqrt{2n} \rceil, r_{t'} \notin \mathcal{I}_1^{t-1}(r_t)\}| \\ & = \lceil \sqrt{2n} \rceil \end{aligned}$$

where we used Lemma 3.3 for the first inequality.

Summing the in-degrees over all the rounds, we get that the number of edges $|E|$ is at least:

$$|E| \geq \sum_{t=1}^{\lceil \sqrt{2n} \rceil} t + n \times \lceil \sqrt{2n} \rceil = \frac{\lceil \sqrt{2n} \rceil (\lceil \sqrt{2n} \rceil + 1)}{2} + n \lceil \sqrt{2n} \rceil > \lceil (1 + \sqrt{2})n \rceil n$$

But only the n nodes representing the processes and the nodes representing the first $\lceil \sqrt{2n} \rceil - 1$ rounds have out edges. The pigeonhole principle asserts then that one of those nodes has an out-degree of at least n . \square

Theorem 3.6. $t^*(\mathcal{T}_n) \leq \lceil (1 + \sqrt{2})n \rceil$

Proof. Assume it is not the case, that is, for every process $x \in [n]$, for every round $t \leq \lceil (1 + \sqrt{2})n \rceil$, $|\mathcal{O}_1^t(x)| < n$. We know that in the rounds graph, there is a node y of degree at least n . Define z as follows: if y represents a process, let z be that process. If y represents a round, let z be the root of that round. We will show that z must have broadcast before $\lceil (1 + \sqrt{2})n \rceil$ rounds.

Let $t_1 < \dots < t_n$ be the rounds y has out-edges to. By definition, this means that $z \in \mathcal{I}_1^{t_i-1}(r_{t_i}) \Leftrightarrow r_{t_i} \in \mathcal{O}_1^{t_i-1}(z)$ for every $i \in [n]$. By Lemma 3.2ii, we thus have, for every $i \in [n]$, that $|\mathcal{O}_1^{t_i}(z)| > |\mathcal{O}_1^{t_i-1}(z)|$. Then, using Monotonicity (Lemma 2.5) for non-strict inequalities:

$$|\mathcal{O}_1^{t_n}(z)| > |\mathcal{O}_1^{t_n-1}(z)| \geq |\mathcal{O}_1^{t_{n-1}}(z)| > |\mathcal{O}_1^{t_{n-1}-1}(z)| \geq \dots \geq |\mathcal{O}_1^{t_1}(z)| > |\mathcal{O}_1^{t_1-1}(z)| \geq 0$$

We have n strict inequalities over non-negative integers, the largest one must be at least n , which is a contradiction. \square

3.2 The Lower Bound

A lower bound for this problem has been given by Zeiner, Schwarz, and Schmid [18]:

Theorem 3.7. $t^*(\mathcal{T}_n) \geq \lceil \frac{3n-1}{2} - 2 \rceil$

A figure of that lower bound can be found in Appendix A.

4 Covering on k -Forests

In this section, we study an adversary that has to choose a communication network in each round that is a union of k trees. In this setting, we cannot ensure broadcast, so we look at the time when there exists a cover of size k : k processes such that any other process has heard of at least one of them.

4.1 The Upper Bound

Even though the problem is very similar to broadcasting on trees, the proofs of Section 3 do not translate in a straightforward way into an upper bound for covering on k -forests. We thus have a completely different proof for this problem.

The intuition of our approach is as follows: We will start with a cover of size n at some time $t' = t^{(n)}$ that is large enough, and then go back in time until we find a process that can reach two other processes, say, x and y , of that cover. Calling this process p_{n-1} and the corresponding time $t^{(n-1)}$, we thus have $p_{n-1} \in \mathcal{I}_{t^{(n-1)}}^{t^{(n)}}(x)$ and $p_{n-1} \in \mathcal{I}_{t^{(n-1)}}^{t^{(n)}}(y)$. When repeating the process, we then remove x and y from our set of processes to cover, add p_{n-1} , and start over, until the cover has size k . We need to be careful to guarantee that rounds do not overlap.

Indeed, to remove p_{n-1} from our set of processes to cover, we have to reach it before round $t^{(n-1)}$, otherwise we will not be guaranteed to reach x or y at time $t^{(n)}$. Thus, we will store with each process x of the cover the corresponding round t such that x has to be reached by round t by the process that replaces x in the cover. More specifically, we model the cover by a series of sets $(A_u)_{k \leq u \leq n}$, where each A_u is a collection of u pairs (p, t) , where p is a process, and t is

a round. To compute A_{u-1} from A_u , we have to find a process that can reach two processes p_1, p_2 by rounds t_1, t_2 , such that $(p_1, t_1), (p_2, t_2) \in A_u$ and then we replace these two pairs by a new pair, creating the cover A_{u-1} .

In this section we first state the definitions and results of this section, before giving the full proof to each of our claims. We first define what it means for a set A of pairs (p, t) to be a cover.

Definition 4.1. A set $A = \{(a_1, t_1), \dots, (a_s, t_s)\}$ of s (process, round) pairs is a cover of a set B of processes for round $t \geq \max_s \{t_s\}$ if for every $b \in B$, there exists an $i \in [s]$ such that $b \in \mathcal{O}_{t_i+1}^t(a_i)$.

We next couple the cover property of set A with *strictness*, which indicates that we did not (yet) go back enough in time to find a process that reaches two different processes in A .

Definition 4.2. A set $A = \{(a_1, t_1), \dots, (a_s, t_s)\}$ of s (process, round) pairs is strict at round t if there exists no process $p \in [n]$ and $i, j \in [s], i \neq j$ such that $p \in \mathcal{I}_t^{t_i}(a_i)$ and $p \in \mathcal{I}_t^{t_j}(a_j)$.

As we consider earlier and earlier rounds, the sets $\mathcal{I}_t^{t_i}(a_i)$ will get larger and larger, and A will lose its strictness. Thus, we then define the following sequence of covers of $[n]$, and analyze their strictness over time. We carefully choose our set A_s so that it has cardinality s . This means that A_k has cardinality k , which is our goal.

Definition 4.3. Let t' be a large enough round. For every $k \leq s \leq n$, we define a sequence of strict sets A_s and rounds $t^{(s)}$ as follows:

Define $A_n = \{(i, t') : i \in [n]\}$, $t^{(n)} = t'$.

Define $t^{(s)} = \max_{j \in \mathbb{N}} \{A_{s+1} \text{ is not strict at round } j\}$. As A_{s+1} is not strict at round $t^{(s)}$, there exist $i, j \in [s+1]$ and a process p_s such that $p_s \in \mathcal{I}_{t^{(s)}}^{t_i}(a_i) \cap \mathcal{I}_{t^{(s)}}^{t_j}(a_j)$. If $(p_s, t^{(s)} - 1) \in A_{s+1}$, we define $A_s = A_{s+1} \setminus \{(a_i, t_i)\}$, else we define $A_s = (A_{s+1} \setminus \{(a_i, t_i), (a_j, t_j)\}) \cup \{(p_s, t^{(s)} - 1)\}$.

Define $\Delta_s = t^{(s)} - t^{(s-1)}$ for $k+1 \leq s \leq n$.

Recall that our goal is to upper bound t' , which can be done if we upper bound $\sum_{s=k+1}^n \Delta_s$. The strictness of a set is a key notion as a strict set has the following very useful property:

Lemma 4.4 (Strict Increments). Let $s \geq k+1$ and let $A = \{(a_1, t_1), \dots, (a_s, t_s)\}$ be a strict set of size s at round t . Let $I \subseteq \{i \in [s] : t \leq t_i + 1\}$. Then there exists a set of indices $J \subseteq I$, $|J| \geq |I| - k$, such that for every $i \in J$, $|\mathcal{I}_{t-1}^{t_i}(a_i)| > |\mathcal{I}_t^{t_i}(a_i)|$.

Proof. Consider a root r of round $t-1$. As A is strict at round t it follows that $r \in \mathcal{I}_t^{t_i}(a_i)$ for at most one a_i with $i \in I$. As there are at most k roots in round $t-1$, it follows that there are at least $|I| - k$ values $i \in I$ such that none of the roots of round $t-1$ is in $\mathcal{I}_t^{t_i}(a_i)$. Let J be the set of all such values of i . Let us denote for any $p \in [n]$ the root of the tree that contains p in round t by the value $r_t(p)$. It follows that in particular $r_{t-1}(a_i) \notin \mathcal{I}_t^{t_i}(a_i)$. Since $t \leq t_i + 1$, we have that $a_i \in \mathcal{I}_t^{t_i}(a_i)$. Now for each such i the path from $r_{t-1}(a_i)$ to a_i in G_{t-1} must contain an edge (x, y) such that $x \notin \mathcal{I}_t^{t_i}(a_i)$, and $y \in \mathcal{I}_t^{t_i}(a_i)$. By Transitivity (Lemma 2.4), it holds that $(\mathcal{I}_t^{t_i}(a_i) \cup \{x\}) \subseteq \mathcal{I}_{t-1}^{t_i}(a_i)$, which implies that $|\mathcal{I}_{t-1}^{t_i}(a_i)| > |\mathcal{I}_t^{t_i}(a_i)|$. \square

It is not hard to show that all a_i of A_s fulfill $t^{(s)} \leq t_i + 1$. The following lemma helps us find more sets that satisfy the conditions of the Strict Increments Lemma. This essentially follows from the fact that, by construction, at least $2s - u$ elements from A_u are shared with A_s .

Lemma 4.5. Let $s, u \in \{k, \dots, n\}$ such that $s \leq u$. Let $A_s = \{(a_1, t_1), \dots, (a_s, t_s)\}$. Let $I = \{i \in [s] : t^{(u)} \leq t_i + 1\}$. Then it holds that $|I| \geq 2s - u$.

We now define the strict rounds graph, which we use to analyze the values of Δ_s as s varies from n to k . A depiction of that graph can be seen in Figure 3.




Figure 3: The strict rounds graph. Vertex u 's weight represents $\Delta_u = t^{(u)} - t^{(u-1)}$. The weight of an edge (u, s) represents how much a round $t \in [t^{(u-1)} + 1, t^{(u)}]$ contributes at least to S_s . Because of the strictness of A_s , the sum of the weights of the in-neighbors (multiplied by the edges' weight) of a node s , which is smaller than S_s , should not exceed n .

Definition 4.6. *The strict rounds graph (V, E) consists of $n - k$ vertices, labeled from $k + 1$ to n , where each vertex s has weight Δ_s . There exists a directed edge from vertex u to vertex s if $s \leq u \leq \min\{2s - k - 1, n\}$, and its weight is $w(u, s) = 2s - k - u$.*

The next lemma is the crucial lemma: To bound t' we first bound the following “weighted volumes”. If we define α_s to be the weighted out-degree of a node s , and the *volume* of s to be α_s multiplied by its own weight, then in the strict rounds graph, the cumulative sum over the volumes of the first j vertices is at most $j \cdot n$.

Lemma 4.7. *Let $u \in \{k + 1, \dots, n\}$. Then $\sum_{s \leq u} \Delta_s \alpha_s \leq (u - k) \cdot n$.*

We briefly sketch the proof of this lemma. We first prove that for every $k + 1 \leq w \leq u$ the sum $\sigma_w := \sum_{v=w}^{\min\{2w-k-1, n\}} (2w - k - v) \Delta_v$ is at most n as follows. Since for every $k + 1 \leq w \leq u$ the set A_w is strict at all rounds $t \geq t^{(w-1)} + 1$, it follows that $S_w := \sum_{(a_i, t_i) \in A_w} \left| \mathcal{I}_{t^{(w-1)}+1}^{t_i}(a_i) \right|$ is at most n . We then lower bound S_w by $\sum_{v=w}^{\min\{2w-k-1, n\}} (2w - k - v) \Delta_v$ in two steps: First, Lemma 4.5 allows us to find a set I of cardinality larger than $2w - v$ for each round in the interval $[t^{(v-1)} + 1, t^{(v)}]$ that fits the Strict Increments Lemma's conditions. We then use the Strict Increments Lemma to show that S_w increases by $2w - v - k$ in each of those rounds, which results in a lower bound of σ_w for S_w , and thus the upper bound $\sigma_w \leq n$. Summing this inequality over all $u - k$ nodes w with $w \leq u$ gives $\sum_{w \leq u} \sigma_w \leq (u - k) \cdot n$.

Next note that σ_w is the weighted in-degree of node w in the strict rounds graph where each edge is weighted by the product of its edge weight and the weight of its tail. By definition, the tail of every (directed) edge has a higher label than its head and, thus, every outgoing edge of a node $s \leq u$ is also an incoming edge of a node $w \leq u$. This allows us to argue that $\sum_{w \leq u} \sigma_w$ is at least $\sum_{s \leq u} \Delta_s \alpha_s$, which leads to the final result.

We use Lemma 4.7 to bound t' as follows. We first show that any vertex weight distribution on the strict rounds graph following the volume bound must fulfill the following property.

Lemma 4.8. *Let $\{\delta_s\}_{k+1 \leq s \leq n}$ with $\delta_s \in \mathbb{R}$ be a vertex weight distribution over the strict rounds graph such that for every $u \in \{k + 1, \dots, n\}$, $\sum_{s \leq u} \delta_s \alpha_s \leq (u - k) \cdot n$. Then $\sum_{s=k+1}^n \delta_s \leq \frac{\pi^2+6}{6} n$.*

Lemma 4.7 and Lemma 4.8 give the desired bound on $\sum_{s=k+1}^n \Delta_s$, which in turn bounds t' .

Corollary 4.9. $\sum_{s=k+1}^n \Delta_s \leq \frac{\pi^2+6}{6} n$

Theorem 4.10. $t_k^c(\mathcal{F}_n^k) \leq \frac{\pi^2+6}{6}n + 1$

The rest of this section is dedicated to proving in detail all of those claims, as well as introducing all concepts, stating and proving any intermediate lemmata that would be necessary. First we analyze in detail the sets defined in Definition 4.3, making sure they have the desired cardinality and proving that they form a cover of $[n]$ from round 1 to t' . Note that it directly follows from Definition 4.3 that the size of A_s is s and we will in the following always use (a_i, t_i) for $1 \leq i \leq s$ to denote the elements of the set A_s .

Lemma 4.11. *For every $s \in \{k, \dots, n\}$, $|A_s| = s$.*

Proof. It is true by reverse induction. Indeed, it is trivially true for $s = n$. Suppose it is true for some $s \in \{k+1, \dots, n\}$ and consider two cases: If $(p_{s-1}, t^{(s-1)} - 1) \in A_s$, then $|A_s| - |A_{s-1}| = 1$, as only (a_i, t_i) is removed from A_s . If $(p_{s-1}, t^{(s-1)} - 1) \notin A_s$, then both (a_i, t_i) and (a_j, t_j) are removed from A_s and $(p_{s-1}, t^{(s-1)} - 1)$ is added, which implies that $|A_s| - |A_{s-1}| = 1$. Thus, in both cases $|A_s| - |A_{s-1}| = 1$. As by induction $|A_s| = s$, it follows that $|A_{s-1}| = s - 1$. \square

In Definition 4.3, we offset $t^{(s)}$ by one when introducing p_s in A_s , which guarantees that A_s is a cover of $[n]$ for round t' , as shown below.

Lemma 4.12. *For every $s \in \{k, \dots, n\}$, we have that A_s is a cover of $[n]$ for round t' .*

Proof. We show the claim by reverse induction. It holds trivially for $s = n$. Assume now that A_{s+1} is a cover of $[n]$ for round t' . We will show that this implies that A_s is also a cover of $[n]$ for round t' . By the fact that A_{s+1} is a cover, it follows for each $x \in [n]$ that there exists a $h \in [s+1]$ such that $x \in \mathcal{O}_{t_h+1}^{t'_h}(a_h)$. If $(a_h, t_h) \in A_s$, then there is nothing to do. If however $(a_h, t_h) \notin A_s$, then $p_s \in \mathcal{I}_{t^{(s)}}^{t'_h}(a_h)$, and by Transitivity (Lemma 2.4) it follows that $x \in \mathcal{O}_{t^{(s)}}^{t'_h}(p_s)$. \square

Now that we are assured that the cover property will hold throughout, we give the intuition of what follows. As defined above, we have $\Delta_u = t^{(u)} - t^{(u-1)}$ for $k < u \leq n$. We will then look at A_s for some $s \in \{k, \dots, n\}$. Since A_s is strict at round $t^{(s-1)} + 1$, it follows that $(\mathcal{I}_{t^{(s-1)}+1}^{t_i}(a_i))_{i \in [s]}$ are pairwise disjoint subsets of $[n]$, which implies that $S_s = \sum_{(a_i, t_i) \in A_s} |\mathcal{I}_{t^{(s-1)}+1}^{t_i}(a_i)|$ is at most n . We will find a lower bound for that value that depends on the values of Δ_u .

To do so, we will use the Strict Increments Lemma (Lemma 4.4). Indeed we will first prove that for every $i \in [s]$, it holds that $t^{(s)} \leq t_i + 1$. This will allow us to use the Strict Increments Lemma (Lemma 4.4) between rounds $t^{(s-1)} + 2$ and $t^{(s)}$ and so each of those rounds contributes an additive $s - k$ to the lower bound of S_s . Of course, such a contribution only happens if $t^{(s-1)} + 2 \leq t^{(s)}$.

We will further prove that there exist $s - 1$ elements $i \in [s]$ such that $t^{(s+1)} \leq t_i + 1$. This will follow from the fact that $|A_s \cap A_{s+1}| \geq s - 1$, and that for every $(a, t) \in A_{s+1}$, we have that $t^{(s+1)} \leq t + 1$. This allows us to use the Strict Increments Lemma (Lemma 4.4) between rounds $t^{(s)} + 1$ and $t^{(s+1)}$, so that each of those rounds contributes at least an additive $s - 1 - k$ to the lower bounds of S_s in addition to the contribution of rounds $[t^{(s-1)} + 2, t^{(s)}]$.

In fact, we will generalize this analysis to all values of $u \geq s$ with $A_u \cap A_s \neq \emptyset$, and for each u every round in $[t^{(u-1)} + 1, t^{(u)}]$ contributes at least $2s - k - u$ to the lower bound of S_s , leading to a contribution at at least $(2s - k - u)\Delta_u$ for u .

Lemma 4.13. *A_s is strict at round $t^{(s)} + 1$, for any $s \in \{k, \dots, n\}$.*

Proof. Let $s \in \{k, \dots, n\}$. A_{s+1} is strict at round $t^{(s)} + 1$. Assume by contradiction that A_s is not strict at round $t^{(s)} + 1$, that is, there exists a $x \in [n]$ and $i, j \in [s]$ such that $x \in \mathcal{I}_{t^{(s)}+1}^{t_i}(a_i)$

and $x \in \mathcal{I}_{t^{(s)+1}}^{t_j}(a_j)$. Since $\mathcal{I}_{t^{(s)+1}}^{t^{(s)}-1}(p_s) = \emptyset$ because $t^{(s)} + 1 > (t^{(s)} - 1) + 1$, we have that $x \notin \mathcal{I}_{t^{(s)+1}}^{t^{(s)}-1}(p_s)$. This implies that $(a_i, t_i), (a_j, t_j) \in A_{s+1}$, which contradicts the strictness of A_{s+1} . This concludes the proof. \square

Corollary 4.14. $t^{(s)} \leq t^{(s+1)} \quad \forall s \in \{k, \dots, n-1\}$

Corollary 4.15. *For every $s \in \{k, \dots, n\}$, we have that for every $i \in [s]$, $t^{(s)} \leq t_i + 1$.*

Proof. By reverse induction, it is true for $s = n$, as for every $i \in [n]$, $t_i = t^{(n)}$. Assume it is true for some $s + 1$ for $s \in \{k, \dots, n-1\}$, and let us look at A_s . Then by Corollary 4.14 and the induction hypothesis, for every $i \in [s]$ we have that either $(a_i, t_i) \in A_{s+1}$ and thus $t^{(s-1)} \leq t^{(s)} \leq t_i + 1$, or that $(a_i, t_i) = (p_s, t^{(s)} - 1)$ and trivially the inequality holds. \square

Lemma 4.16. *Let $s, u \in \{k, \dots, n\}$ such that $s \leq u$. Then $A_s \cap A_u \geq 2s - u$.*

Proof. By reverse induction, it is trivially true for $s = u$. Assume we have $|A_{s+1} \cap A_u| \geq 2(s+1) - u$ for some s such that $k \leq s \leq u-1$.

We have that, by distribution of the intersection over the union operator:

$$A_{s+1} \cap A_u = ((A_{s+1} \cap A_s) \cup (A_{s+1} \setminus A_s)) \cap A_u = (A_{s+1} \cap A_s \cap A_u) \cup ((A_{s+1} \setminus A_s) \cap A_u)$$

Hence:

$$|A_s \cap A_u| \geq |A_{s+1} \cap A_s \cap A_u| \geq |A_{s+1} \cap A_u| - |(A_{s+1} \setminus A_s) \cap A_u| \geq 2(s+1) - u - 2 = 2s - u$$

Where we used that $|(A_{s+1} \setminus A_s) \cap A_u| \leq |A_{s+1} \setminus A_s| \leq 2$, because we removed at most two elements from A_{s+1} to get A_s . \square

Lemma 4.5. *Let $s, u \in \{k, \dots, n\}$ such that $s \leq u$. Let $A_s = \{(a_1, t_1), \dots, (a_s, t_s)\}$. Let $I = \{i \in [s] : t^{(u)} \leq t_i + 1\}$. Then it holds that $|I| \geq 2s - u$.*

Proof. It is a direct implication of Corollary 4.15 and Lemma 4.16.

Indeed, $|I| = |\{(a, t) \in A_s : t^{(u)} \leq t + 1\}|$, and we have $\{(a, t) \in A_s : t^{(u)} \leq t + 1\} \supseteq \{(a, t) \in A_s \cap A_u : t^{(u)} \leq t + 1\} = \{(a, t) \in A_s \cap A_u\} = A_s \cap A_u$, where the penultimate equality holds by Corollary 4.15 applied to u . Lemma 4.16 states that $|A_s \cap A_u| \geq 2s - u$. \square

Lemma 4.17. *Let $\Delta_s = t^{(s)} - t^{(s-1)}$. Then, for every $s \in \{k+1, \dots, n\}$:*

$$\sum_{u=s}^{\min\{2s-k-1, n\}} (2s - k - u) \Delta_u \leq n$$

An illustration of the proof can be found in Figure 4.

Proof. Let $s \in \{k+1, \dots, n\}$. If for every u such that $s \leq u \leq \min\{2s-k-1, n\}$, we have $\Delta_u = 0$, then the result is immediate. If not, let $x = \min\{u : s \leq u \leq \min\{2s-k-1, n\} \wedge \Delta_u > 0\}$. In the rest of the proof, let u be restricted to fulfill $x \leq u \leq \min\{2s-k-1, n\}$ and let $I_s^u = \{i \in [s] : t^{(u)} \leq t_i + 1\}$. By Lemma 4.5, we have that $|I_s^u| \geq 2s - u$. To prove the claim we will give upper and lower bounds for $S_s = \sum_{i \in [s]} \left| \mathcal{I}_{t^{(s-1)+1}}^{t_i}(a_i) \right|$. The upper bound directly follows from the fact that $t^{(s-1)} + 1$ is the smallest round at which A_s is strict, which implies that $S_s \leq n$.

We next fix a round t with $t^{(u-1)} + 1 \leq t \leq t^{(u)}$. We have that $t^{(s-1)}$ is the largest round A_s is not strict at (by Definition 4.3), so it is strict at $t \geq t^{(u-1)} + 1 > t^{(s-1)}$. By the Strict Increments Lemma (Lemma 4.4), there exists a set $J_s^u(t) \subseteq I_s^u$ with $|J_s^u(t)| \geq |I_s^u| - k \geq 2s - u - k$ such




Figure 4: Illustration of proof of Lemma 4.17, with $s = 7$ and $k = 2$. Each column represents a round, each row represents an element (a_i, t_i) of A_s , ordered according to t_i . Remember that every t_i directly precedes a $t^{(u)}$. We add a square if $t \leq t_i$, and a circle otherwise. The entry is black if $|\mathcal{I}_t^{t_i}(a_i)| > |\mathcal{I}_{t+1}^{t_i}(a_i)|$ and white otherwise. We have one black circle per row. Indeed, we have that $|\mathcal{I}_{t_i+1}^{t_i}(a_i)| = |\{a_i\}| = 1 > 0 = |\mathcal{I}_{t_i+2}^{t_i}(a_i)|$. Lemma 4.5 gives a lower bound on the number of squares for each column. The Strict Increments Lemma asserts that there are at most k white squares per column. A lower bound for S_s is thus the number of black entries.

that for every $j \in J_s^u(t)$, $|\mathcal{I}_{t-1}^{t_j}(a_j)| > |\mathcal{I}_t^{t_j}(a_j)|$. As every $j \in J_s^u(t)$ belongs to I_s^u it follows that $t \leq t^{(u)} \leq t_j + 1$. Thus, for every $(j, t) \in [s] \times \mathbb{N}$ with $t > t_j + 1$ it holds that it does not belong to $J_s^u(t)$, or, equivalently, $1(j \in J_s^u(t)) = 0$.

We can then write:

$$S_s := \sum_{i \in [s]} |\mathcal{I}_{t^{(s-1)+1}}^{t_i}(a_i)| = \sum_{i \in [s]} \left(|\mathcal{I}_{t_i+1}^{t_i}(a_i)| + \sum_{t=t^{(s-1)+1}}^{t_i} (|\mathcal{I}_t^{t_i}(a_i)| - |\mathcal{I}_{t+1}^{t_i}(a_i)|) \right)$$

We first separate the second sum according to which interval $[t^{(u-1)} + 1, t^{(u)}]$ round t belongs to, and furthermore add a third sum that is equal to 0, since $1(j \in J_s^u(t)) = 0$ for every $(j, t) \in [s] \times \mathbb{N}$ such that $t > t_j + 1$. For the consecutive inequality note that by the definition of $J_s^u(t)$, we have that in the second sum $|\mathcal{I}_t^{t_i}(a_i)| - |\mathcal{I}_{t+1}^{t_i}(a_i)| \geq 1(i \in J_s^u(t+1))$.

$$\begin{aligned} S_s &\geq \sum_{i \in [s]} \left(|\mathcal{I}_{t_i+1}^{t_i}(a_i)| + \sum_{t=t^{(s-1)+1}}^{t_i} \sum_{u=s}^{\min\{2s-k-1, n\}} 1(t^{(u-1)} + 1 \leq t+1 \leq t^{(u)}) (|\mathcal{I}_t^{t_i}(a_i)| - |\mathcal{I}_{t+1}^{t_i}(a_i)|) \right. \\ &\quad \left. + \sum_{t=t_i+2}^{t^{(\min\{2s-k-1, n\})}} \sum_{u=s}^{\min\{2s-k-1, n\}} 1(t^{(u-1)} + 1 \leq t \leq t^{(u)}) 1(i \in J_s^u(t)) \right) \\ &\geq \sum_{i \in [s]} \left(|\mathcal{I}_{t_i+1}^{t_i}(a_i)| + \sum_{t=t^{(s-1)+1}}^{t_i} \sum_{u=s}^{\min\{2s-k-1, n\}} 1(t^{(u-1)} + 1 \leq t+1 \leq t^{(u)}) 1(i \in J_s^u(t+1)) \right. \\ &\quad \left. + \sum_{t=t_i+2}^{t^{(\min\{2s-k-1, n\})}} \sum_{u=s}^{\min\{2s-k-1, n\}} 1(t^{(u-1)} + 1 \leq t \leq t^{(u)}) 1(i \in J_s^u(t)) \right) \\ &= \sum_{i \in [s]} \left(|\mathcal{I}_{t_i+1}^{t_i}(a_i)| + \sum_{t=t^{(s-1)+2}}^{t_i+1} \sum_{u=s}^{\min\{2s-k-1, n\}} 1(t^{(u-1)} + 1 \leq t \leq t^{(u)}) 1(i \in J_s^u(t)) \right. \\ &\quad \left. + \sum_{t=t_i+2}^{t^{(\min\{2s-k-1, n\})}} \sum_{u=s}^{\min\{2s-k-1, n\}} 1(t^{(u-1)} + 1 \leq t \leq t^{(u)}) 1(i \in J_s^u(t)) \right) \end{aligned}$$

$$S_s \geq \sum_{i \in [s]} \left(\left| \mathcal{I}_{t_{i+1}}^{t_i}(a_i) \right| + \sum_{t=t^{(s-1)+2}}^{\max\{t_i+1, t^{\min\{2s-k-1, n\}}\}} \sum_{u=s}^{\min\{2s-k-1, n\}} 1(t^{(u-1)} + 1 \leq t \leq t^{(u)}) 1(i \in J_s^u(t)) \right) =: X_1$$

We next invert the two sum symbols, and delete terms where $1(t^{(u-1)} + 1 \leq t \leq t^{(u)}) = 0$:

$$\begin{aligned} X_1 &\geq \sum_{i \in [s]} \left(1 + \sum_{t=t^{(x-1)+2}}^{t^{(x)}} 1(i \in J_s^x(t)) + \sum_{u=x+1}^{\min\{2s-k-1, n\}} \sum_{t=t^{(u-1)+1}}^{t^{(u)}} 1(i \in J_s^u(t)) \right) \\ &\geq s + \sum_{t=t^{(x-1)+2}}^{t^{(x)}} \sum_{i \in [s]} 1(i \in J_s^x(t)) + \sum_{u=x+1}^{\min\{2s-k-1, n\}} \sum_{t=t^{(u-1)+1}}^{t^{(u)}} \sum_{i \in [s]} 1(i \in J_s^u(t)) =: X_2 \end{aligned}$$

Using that $\sum_{i \in [s]} 1(i \in J_s^u(t)) = |J_s^u(t)| \geq 2s - k - u$ and $\Delta_u = t^{(u)} - t^{(u-1)}$:

$$X_2 \geq s + (\Delta_x - 1)(2s - k - x) + \sum_{u=x+1}^{\min\{2s-k-1, n\}} (2s - k - u)\Delta_u \geq \sum_{u=s}^{\min\{2s-k-1, n\}} (2s - k - u)\Delta_u,$$

where the last inequality follows since $x \geq s$ which implies that $2s - k - x \leq 2s - k - s \leq s$ and that $\Delta_u = 0$ for $u < x$.

As $\sum_{i \in [s]} \left| \mathcal{I}_{t_{(s-1)+1}}^{t_i}(a_i) \right| = S_s \leq n$, the claim follows. \square

We now recall the definition of the strict rounds graph, which is built to harness the previous result.

Definition 4.6. *The strict rounds graph (V, E) consists of $n - k$ vertices, labeled from $k + 1$ to n , where each vertex s has weight Δ_s . There exists a directed edge from vertex u to vertex s if $s \leq u \leq \min\{2s - k - 1, n\}$, and its weight is $w(u, s) = 2s - k - u$.*

The graph is represented in Figure 3. Each node s with $k + 1 \leq s \leq (n + k)/2$ has $s - k + 2$ incoming edges, namely $(s, s), (s + 1, s), \dots, (2s - k - 1, s)$, with weights $s - k, s - k - 1, \dots, 1$, respectively. Each node s with $(n + k)/2 < s \leq n$ has $n + 1 - s$ incoming edges, namely $(s, s), (s + 1, s), \dots, (n, s)$, with weights $s - k, s - k - 1, \dots, 2s - k - n$, respectively.

Each node s has $s + 1 - \lfloor \frac{s+k}{2} \rfloor$ outgoing edges, namely $(s, s), (s, s - 1), \dots, (s, \lfloor \frac{s+k}{2} \rfloor + 1)$ of weight $s - k, s - k - 2, \dots$, respectively. If $s - k$ is even, all outgoing edges of s have even weight and the edge $(s, \lfloor \frac{s+k}{2} \rfloor + 1)$ has weight 2. If $s - k$ is odd, all outgoing edges of s have odd weight and the edge $(s, \lfloor \frac{s+k}{2} \rfloor + 1)$ has weight 1.

Lemma 4.18. *Let $u \in [n]$, and define $E_u^{out} := \{(s, v) \in E : s \leq u\}$. In the strict rounds graph, we have that*

$$\sum_{(s,v) \in E_u^{out}} \Delta_s w(s, v) \leq (u - k) \cdot n$$

Proof. We are summing over all the out-edges of all the vertices with a label at most u . Let $E_u^{in} = \{(s, v) \in E : v \leq u\}$. Every outgoing edge of a node $s \leq u$ must end at a node v with $v \leq s \leq u$. Thus, it is contained in the set of incoming edge of all nodes v with $v \leq u$. Said differently, $E_u^{out} \subseteq E_u^{in}$.

Since all the terms are positive, we have:

$$\begin{aligned} \sum_{(s,v) \in E_u^{out}} \Delta_s w(s,v) &\leq \sum_{(s,v) \in E_u^{in}} \Delta_s w(s,v) = \sum_{v \leq u} \sum_{s:(s,v) \in E} \Delta_s w(s,v) \\ &\leq \sum_{v \leq u} \sum_{s:(s,v) \in E} \Delta_s (2v - k - s) \leq \sum_{v \leq u} \sum_{s=v}^{\min\{n, 2v-s-1\}} \Delta_s (2v - k - s) \end{aligned}$$

Applying Lemma 4.17, the claim follows. \square

Lemma 4.19. *Let $s \geq k + 1$. If $s - k$ is odd, we have $\sum_{v:(s,v) \in E} w(s,v) = (\frac{s-k+1}{2})^2$. If it is even, we have $\sum_{v:(s,v) \in E} w(s,v) = (\frac{s-k}{2})^2 + \frac{s-k}{2}$.*

Basically, as discussed above and seen in Figure 3, if $s - k$ is odd, we have that $\sum_{v:(s,v) \in E} w(s,v) = 1 + 3 + 5 + \dots + s - k = (\frac{s-k+1}{2})^2$, while if it is even, $\sum_{v:(s,v) \in E} w(s,v) = 2 + 4 + 6 + \dots + s - k = (\frac{s-k}{2})^2 + \frac{s-k}{2}$. The full proof can be found in Appendix A.

Definition 4.20. *We define the following numbers: $\beta = \frac{\pi^2+6}{6}$, $\alpha_s = (\frac{s-k+1}{2})^2$ if $s - k$ is odd, $\alpha_s = (\frac{s-k}{2})^2 + \frac{s-k}{2}$ if $s - k$ is even.*

Note that α_s is an integer for all values of $s \geq k$. By applying this definition to the bounds of Lemmata 4.18 and 4.19 we achieve the following result.

Lemma 4.7. *Let $u \in \{k + 1, \dots, n\}$. Then $\sum_{s \leq u} \Delta_s \alpha_s \leq (u - k) \cdot n$.*

Lemma 4.8. *Let $\{\delta_s\}_{k+1 \leq s \leq n}$ with $\delta_s \in \mathbb{R}$ be a vertex weight distribution over the strict rounds graph such that for every $u \in \{k + 1, \dots, n\}$, $\sum_{s \leq u} \delta_s \alpha_s \leq (u - k) \cdot n$. Then $\sum_{s=k+1}^n \delta_s \leq \frac{\pi^2+6}{6} n$.*

Proof. We will show that $\sum_{s=k+1}^n \delta_s \alpha_s$ is maximized when $\delta_s = \alpha_s^{-1} \cdot n$ for every s . Let $\{\delta_s\}_{k+1 \leq s \leq n}$ be such that $\sum_{s=k+1}^n \delta_s \alpha_s$ is maximized. Assume by contradiction that there exists a $k + 1 \leq s \leq n$ such that $\delta_s < \alpha_s^{-1} \cdot n$. Let v be the smallest such s and set $\epsilon = n - \delta_v \alpha_v > 0$. We can then build a different solution γ by setting $\gamma_s = \delta_s$ for every $s \notin \{v, v + 1\}$, and setting $\gamma_v = \delta_v + \epsilon \alpha_v^{-1} = \alpha_v^{-1} \cdot n > \delta_v$ and $\gamma_{v+1} = \delta_{v+1} - \epsilon \alpha_{v+1}^{-1}$. Note that for every $u \in \{k + 1, \dots, n\}$, $\sum_{s \leq u} \gamma_s \alpha_s \leq (u - k) \cdot n$ as

- (1) for $u < v$ it holds that $\sum_{s \leq u} \gamma_s \alpha_s = (u - k) \cdot n$,
- (2) for $u = v$ it holds that $\sum_{s \leq u} \gamma_s \alpha_s = (\epsilon \alpha_v^{-1}) \alpha_v + \delta_v \alpha_v + \sum_{s < u} \delta_s \alpha_s = (n - \delta_v \alpha_v) + \delta_v \alpha_v + \sum_{s < u} \delta_s \alpha_s = n + (u - 1 - k) \cdot n = (u - k) \cdot n$, and
- (3) for $u \geq v + 1$ it holds that $\sum_{s \leq u} \gamma_s \alpha_s = \sum_{s \leq v-1} \delta_s \alpha_s + \gamma_v \alpha_v + \gamma_{v+1} \alpha_{v+1} + \sum_{v+2 \leq s \leq u} \delta_s \alpha_s = \sum_{s \leq v-1} \delta_s \alpha_s + \delta_v \alpha_v + \epsilon \alpha_{v+1} \alpha_{v+1}^{-1} + \delta_u \alpha_u^{-1} - \epsilon \alpha_u \alpha_u^{-1} + \sum_{v+2 \leq s \leq u} \delta_s \alpha_s = \sum_{s \leq u} \delta_s \alpha_s \leq (u - k) \cdot n$, and

Now note that $\sum_{s=k+1}^n \gamma_s \alpha_s > \sum_{s=k+1}^n \delta_s \alpha_s$ as the α_s values are decreasing in s , which gives a contradiction to the assumption that δ maximizes $\sum_{s=k+1}^n \delta_s \alpha_s$. It now follows that

$$\sum_{s=k+1}^n \delta_s \leq \sum_{s=k+1}^n \alpha_s^{-1} n \leq \sum_{s=k+1}^{\infty} \alpha_s^{-1} n \leq n \sum_{\ell \in \mathbb{N}} \frac{1}{\ell^2} + n \sum_{\ell \in \mathbb{N}} \frac{1}{\ell^2 + \ell} = \beta n.$$

where the last equation follows by the fact that $\sum_{\ell \in \mathbb{N}} \frac{1}{\ell^2} = \pi^2/6$ and $\sum_{\ell \in \mathbb{N}} \frac{1}{\ell^2 + \ell} = \sum_{\ell \in \mathbb{N}} \frac{1}{\ell} - \frac{1}{\ell+1} = 1$. \square

Corollary 4.9. $\sum_{s=k+1}^n \Delta_s \leq \frac{\pi^2+6}{6} n$

Theorem 4.10. $t_k^c(\mathcal{F}_n^k) \leq \frac{\pi^2+6}{6}n + 1$

Proof. We have $\sum_{s=k+1}^n \Delta_s = \sum_{s=k+1}^n t^{(s)} - t^{(s+1)} = t' - t^{(k)}$, which is at most $\frac{\pi^2+6}{6}n$ by Corollary 4.9. Setting $t^{(k)} = 1$, we get that $t' \leq \frac{\pi^2+6}{6}n + 1$. Lemma 4.12 ensures that A_k is a cover of $[n]$ for t' , which means that for every $x \in [n]$, there exists a $i \in [k]$ such that $x \in \mathcal{O}_{i+1}^{t'}(a_i)$. Since $\min_i \{t_i\} = t^{(k)} - 1 = 0$, this means that every $x \in [n]$ belongs to $\mathcal{O}_1^{t'}(a_i)$, which implies that $t_k^c(\mathcal{F}_n^k) \leq t'$. \square

4.2 The Lower Bound

We build a lower bound example based on the lower bound for broadcasting on trees. A figure and analysis of that example can be found in Appendix A, which yields the following result:

Theorem 4.21. $t_k^c(\mathcal{T}_n^k) \geq \lceil \frac{3n-3k}{2} - 1 \rceil$

5 k -Broadcasting on k -rooted Networks

In this section, we study the case where the adversary has to choose a communication network that has k roots at least in each round. This allows us to enforce not only broadcast, but rather k -broadcast.

The problem is very similar to the one studied in the Section 3, and indeed a slight modification of the proofs there work for this problem.

5.1 The Upper Bound

By using a technique very similar to Section 3, we will show that whenever we have a set $A \subset [n]$ of size at most $k - 1$, we can find a process $p \notin A$ that has broadcast, as long as we have waited for a large enough number of rounds.

Definition 5.1. Let t be a round. We denote by R_t the set of the roots of G_t .

Recall that $|R_t| \geq k$.

The following two lemmata, very similar to Section 3, can be proven by looking at paths going from a root to a carefully chosen node.

Lemma 5.2. Let $t \leq t'$ be rounds, and let x and r be processes such that, $r \in R_t$ and $r \notin \mathcal{I}_{t+1}^{t'}(x)$. Then $|\mathcal{I}_t^{t'}(x)| > |\mathcal{I}_{t+1}^{t'}(x)|$.

Proof. We will show that there exists a process $y \in \mathcal{I}_{t+1}^{t'}(x)$ that has an in-neighbor z in G_t such that $z \notin \mathcal{I}_{t+1}^{t'}(x)$. Then, by Transitivity (Lemma 2.4), $z \in \mathcal{I}_t^{t'}(x)$. By Monotonicity (Lemma 2.5), we have that $\mathcal{I}_{t+1}^{t'}(x) \subset \mathcal{I}_t^{t'}(x)$, this will show that $|\mathcal{I}_t^{t'}(x)| > |\mathcal{I}_{t+1}^{t'}(x)|$.

Let us now find such a z . Consider the path from r to x in G_t . Since $r \notin \mathcal{I}_{t+1}^{t'}(x)$, and trivially $x \in \mathcal{I}_{t+1}^{t'}(x)$, this path must include an edge (z, y) such that $z \notin \mathcal{I}_{t+1}^{t'}(x)$, $y \in \mathcal{I}_{t+1}^{t'}(x)$. \square

Lemma 5.3. Let $t \leq t'$ be rounds, and let x and r be processes such that $r \in R_{t'}$, and $r \in \mathcal{O}_t^{t'-1}(x)$. Then $|\mathcal{O}_t^{t'}(x)| > |\mathcal{O}_t^{t'-1}(x)|$, unless $|\mathcal{O}_t^{t'-1}(x)| = n$.

Proof. Let us look at the case $|\mathcal{O}_t^{t'-1}(x)| < n$. We will show that there exists a process $y \in \mathcal{O}_t^{t'-1}(x)$ that has an out-neighbor z in $G_{t'}$ such that $z \notin \mathcal{O}_t^{t'-1}(x)$. Then, by Transitivity (Lemma 2.4), $z \in \mathcal{O}_t^{t'}(x)$. By Monotonicity, we obviously have $\mathcal{O}_t^{t'-1}(x) \subset \mathcal{O}_t^{t'}(x)$, which implies that $|\mathcal{O}_t^{t'}(x)| > |\mathcal{O}_t^{t'-1}(x)|$.

Let us now find such a z . Since $|\mathcal{O}_t^{t'-1}(x)| < n$, there exists a process a such that $a \notin \mathcal{O}_t^{t'-1}(x)$. Consider the path from r to a in $G_{t'}$. Since $r \in \mathcal{O}_t^{t'-1}(x)$, and $a \notin \mathcal{O}_t^{t'-1}(x)$, this path must include an edge (y, z) such that $z \notin \mathcal{O}_t^{t'-1}(x)$, $y \in \mathcal{O}_t^{t'-1}(x)$. \square

We now give a lemma that links the number of in-neighbors of a node to the number of previous rounds whose root are in-neighbors of said node:

Lemma 5.4. *Let x be a process, let $t_1 \leq t_2$ be rounds, and let r_t denote a root of R_t for every $t \in \{t_1, \dots, t_2\}$, then it holds that*

$$|\{t : t_1 \leq t \leq t_2, r_t \notin \mathcal{I}_{t_1}^{t_2}(x)\}| + 1 \leq |\mathcal{I}_{t_1}^{t_2}(x)|$$

Proof. Let $A = \{t : t_1 \leq t \leq t_2, r_t \notin \mathcal{I}_{t_1}^{t_2}(x)\}$. Then, in particular, for any $t \in A$, we have $r_t \notin \mathcal{I}_{t+1}^{t_2}(x)$. Then, for all $t \in A$, applying Lemma 5.2, we have $|\mathcal{I}_t^{t_2}(x)| > |\mathcal{I}_{t+1}^{t_2}(x)|$. Let $k = |A|$ and let $A = \{t^{(1)}, \dots, t^{(k)}\}$, with $t^{(i)} < t^{(i+1)}$ for $1 \leq i < k$. Then by Monotonicity (Lemma 2.5) it follows that

$$|\mathcal{I}_{t_1}^{t_2}(x)| \geq |\mathcal{I}_{t^{(1)}}^{t_2}(x)| > |\mathcal{I}_{t^{(1)}+1}^{t_2}(x)| \geq |\mathcal{I}_{t^{(2)}}^{t_2}(x)| > |\mathcal{I}_{t^{(2)}+1}^{t_2}(x)| \geq \dots > |\mathcal{I}_{t^{(k)}+1}^{t_2}(x)| \geq 1$$

There are $k = |A|$ inequalities over integers, which concludes the proof. \square

We now define the rounds graph avoiding some set A , which is the equivalent of the rounds graph of Section 3, this time being very careful not to choose any process from A being used as a root for a round.

Definition 5.5. *Let A be an arbitrary set of processes with $|A| < k$. We define the rounds graph avoiding set A as follows: It contains $2n + \lceil \sqrt{2n} \rceil + |A|$ nodes, namely*

1. one process node representing each process.
2. one round node for each of the first $\lceil (1 + \sqrt{2})n \rceil + |A|$ rounds.

For each round t , let r_t be a vertex of $R_t \setminus A$. The graph contains

1. a directed edge to every round node t from each process node p with $p \in \mathcal{I}_1^{t-1}(r_t)$.
2. a directed edge from every round node $t < \lceil \sqrt{2n} \rceil + |A|$ to a round node $t' > t$ if $r_t \in \mathcal{I}_1^{t'-1}(r_{t'})$.

We now find a node in the rounds graph that has out-degree n , that is not a node from A . This node will represent a process that has broadcast.

Lemma 5.6. *For any set A of processes with $|A| < k$, the rounds graph avoiding set A contains at least one node of out-degree n that is not a node representing a process of A .*

Proof. As in the proof of Lemma 3.5 we apply the pigeonhole principle on the edges, however, this time we will ignore out-edges from nodes representing processes from A . Let us look at round t . If $t \leq \lceil \sqrt{2n} \rceil + |A|$, then the round node t has in-degree at least t . Indeed, it has in-degree:

$$\begin{aligned} & |\{t' : 1 \leq t' \leq t-1, r_{t'} \in \mathcal{I}_1^{t-1}(r_t)\}| + |\mathcal{I}_1^{t-1}(r_t)| \\ & \geq 1 + |\{t' : 1 \leq t' \leq t-1, r_{t'} \in \mathcal{I}_1^{t-1}(r_t)\}| + |\{t' : 1 \leq t' \leq t-1, r_{t'} \notin \mathcal{I}_1^{t-1}(r_t)\}| = t \end{aligned}$$

where we used Lemma 5.4 for the inequality. Therefore, it has at least $t - |A|$ in-neighbors not in A . Similarly, if $t > \lceil \sqrt{2n} \rceil + |A|$, then the round node t has in-degree at least $\lceil \sqrt{2n} \rceil + |A|$. Indeed, it has in-degree:

$$\begin{aligned} & |\{t' : 1 \leq t' < \lceil \sqrt{2n} \rceil + |A|, r_{t'} \in \mathcal{I}_1^{t-1}(r_t)\}| + |\mathcal{I}_1^{t-1}(r_t)| \\ & \geq 1 + |\{t' : 1 \leq t' < \lceil \sqrt{2n} \rceil + |A|, r_{t'} \in \mathcal{I}_1^{t-1}(r_t)\}| + |\{t' : 1 \leq t' \leq t-1, r_{t'} \notin \mathcal{I}_1^{t-1}(r_t)\}| \\ & \geq 1 + |\{t' : 1 \leq t' < \lceil \sqrt{2n} \rceil + |A|, r_{t'} \in \mathcal{I}_1^{t-1}(r_t)\}| \\ & \quad + |\{t' : 1 \leq t' < \lceil \sqrt{2n} \rceil + |A|, r_{t'} \notin \mathcal{I}_1^{t-1}(r_t)\}| = \lceil \sqrt{2n} \rceil + |A| \end{aligned}$$

where we used Lemma 5.4 for the first inequality. Therefore, it has at least $\lceil \sqrt{2n} \rceil$ in-neighbors not elements of A .

Summing the in-degrees over all the rounds, we get that the number of edges $|E_{\bar{A}}|$ with no endpoint in A is at least:

$$|E_{\bar{A}}| \geq \sum_{t=|A|}^{\lceil \sqrt{2n} \rceil + |A|} (t - |A|) + n \times \lceil \sqrt{2n} \rceil = \frac{\lceil \sqrt{2n} \rceil (\lceil \sqrt{2n} \rceil + 1)}{2} + n \lceil \sqrt{2n} \rceil > \lceil (1 + \sqrt{2})n \rceil n$$

But only the $n - |A|$ nodes representing the processes and the nodes representing the first $\lceil \sqrt{2n} \rceil - 1 + |A|$ rounds have out-edges among those counted. The pigeonhole principle asserts then that one of those nodes has an out-degree of at least n . \square

Lemma 5.7. *For any set A of processes with $|A| < k$, there exists a process $z \notin A$ that has broadcast its identifier to everyone after $\lceil (1 + \sqrt{2})n \rceil + |A|$ rounds.*

Proof. Build the rounds graph avoiding A . We know that in that graph, there is a node y of degree at least n , that is not a process node representing A . Define z as follows: (1) If y represents a process, let z be that process. It follows that $z \notin A$. (2) If y represents a round t , let z be the root r_t , which is not in A by definition of the graph. We will show in either case that z must have broadcast before $\lceil (1 + \sqrt{2})n \rceil + |A|$ rounds.

Let $t_1 < \dots < t_n$ be the rounds y has out-edges to. By definition, and in both cases (1) and (2), this means $z \in \mathcal{I}_1^{t_i-1}(r_{t_i})$ which is equivalent to $r_{t_i} \in \mathcal{O}_1^{t_i-1}(z)$ for every $i \in [n]$. By Lemma 5.3, we thus have, for every $i \in [n]$, that $|\mathcal{O}_1^{t_i}(z)| > |\mathcal{O}_1^{t_i-1}(z)|$. Then:

$$|\mathcal{O}_1^{t_n}(z)| > |\mathcal{O}_1^{t_n-1}(z)| \geq |\mathcal{O}_1^{t_{n-1}}(z)| > |\mathcal{O}_1^{t_{n-1}-1}(z)| \geq \dots \geq |\mathcal{O}_1^{t_1}(z)| > |\mathcal{O}_1^{t_1-1}(z)| \geq 0$$

We have n strict inequalities over non-negative integers, the largest one must be at least n , which implies that z has broadcast. \square

This naturally leads to an upper bound for the k -broadcasting on k -rooted networks:

Theorem 5.8. $t_k^*(\mathcal{R}_n^k) \leq \lceil (1 + \sqrt{2})n \rceil + k - 1$

Proof. By contradiction, assume that the set A of elements that have broadcast after $\lceil (1 + \sqrt{2})n \rceil + k - 1$ rounds is smaller than k and apply Lemma 5.7 to find a process not in A that has broadcast in time less than $\lceil (1 + \sqrt{2})n \rceil + |A|$ rounds, contradicting the assumption. \square

5.2 The Lower Bound

We build a lower bound example based on the lower bound for broadcasting on trees. A figure and analysis of that example can be found in Appendix A, which yields the following result:

Theorem 5.9. $t_k^*(\mathcal{R}_n^k) \geq \lceil \frac{3n-9k}{2} + 2 \rceil$

6 Related Work

Broadcasting, gossiping, and other information dissemination problems have been studied by the distributed computing community for decades already [13]. Most classic literature on network broadcast considers a static setting, e.g., where in each round each node can send information to one neighbor [14]. This model has also been explored in the context of gossiping, e.g., by Fraignaud and Lazard [11]. Kuhn, Lynch and Oshman [15] explore the all-to-all data dissemination problem (gossiping) in an undirected dynamic network, where processes do not know beforehand the total number of processes and must decide on that number. Ahmadi, Kuhn, Kutten, Molla and Pandurangan [2] study the message complexity of broadcast in an undirected dynamic setting, where the adversary pays up a cost for changing the network. Broadcast has also been studied in dynamic communication networks which evolve randomly, e.g., by Clementi et al. [5], and in the radio network model [10], just to give a few examples.

A closely related yet different problem to broadcasting is the consensus problem. Our model builds up on the heard-of model first introduced by Charron-Bost and Schiper [4], where authors prove results for the solvability of consensus also considering oblivious message adversaries. Among other results, they give a $\log n$ upper bound for nonsplit graphs, which are graphs for which every pair of nodes has a common in-neighbor. This would result in an $n \log n$ upper bound for rooted trees when combining it with the result of Charron-Bost, Függer and Nowak [3]. Függer, Nowak, and Winkler [12] prove that the time complexity of broadcast is a lower bound for consensus time. A general characterization of oblivious message adversaries on which consensus is solvable, based on broadcastability, has been presented by Coulouma, Godard and Peters in [6]. A time complexity analysis has further been studied by Winkler, Rincon Galeana, Paz, Schmid, and Schmid [17]. Another similar problem is agreement, considered by Santoro and Widmayer [16], where only a k -majority should agree on a value, as opposed to everyone for consensus. Afek, Gafni, Rajsbaum, Raynal and Travers [1] studied a generalization to consensus that is k -set consensus, where each node has to decide on a value such that all the nodes together do not decide on more than k different values.

In this paper, we have studied the broadcasting problem on directed dynamic networks, with an adversary that can choose the communication network at each round among rooted trees. Zeiner, Schwarz, and Schmid [18] give a $n \log n$ upper bound to our exact problem by using graph-theoretic reasoning. They also give a $\lceil \frac{3n-1}{2} \rceil - 2$ lower bound by providing an explicit example. They further show that under an adversary that can only choose rooted trees with a fixed number of leaves or internal nodes, broadcast time is linear.

There has also been interest in a problem variant which only differs in the pool of networks the adversary can choose a network from for each communication round. Függer, Nowak, and Winkler [12] give 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 [3] 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 [8, 7] give specific results when the adversary is restricted to hypercubic and tori graphs with some missing edges.

Bibliographic note: an announcement of this work has been presented at ACM PODC 2022 [9].

7 Conclusion

In this paper, we considered an innovative version of the classic broadcast problem where processes communicate across a dynamically changing network, as it often arises in practice (e.g., due to interference). Like in the static setting, the broadcast problem on dynamic networks is related to consensus and leader election: broadcast is a prerequisite for consensus, and hence, the time complexity of broadcast is a lower bound for the consensus and leader election complexity.

Our main contribution is a proof that the broadcast time is at most linear in this setting, which is asymptotically optimal. We further presented several natural generalizations of our model and result.

Our work opens several avenues for future research. In particular, it will be interesting to study the broadcast time also in non-adversarial environments where graphs evolve according to a random process (e.g., due to random node movements). It will also be interesting to further explore the implications of our methods on the closely related consensus problem.

References

- [1] Yehuda Afek, Eli Gafni, Sergio Rajsbaum, Michel Raynal, and Corentin Travers. The k -simultaneous consensus problem. *Distributed Comput.*, 22(3):185–195, 2010.
- [2] 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.
- [3] 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.
- [4] Bernadette Charron-Bost and André Schiper. The heard-of model: computing in distributed systems with benign faults. *Distributed Comput.*, 22(1):49–71, 2009.
- [5] 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.
- [6] 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.

- [7] Stefan Dobrev and Imrich Vrto. Optimal broadcasting in hypercubes with dynamic faults. *Inf. Process. Lett.*, 71(2):81–85, 1999.
- [8] Stefan Dobrev and Imrich Vrto. Optimal broadcasting in tori with dynamic faults. *Parallel Process. Lett.*, 12(1):17–22, 2002.
- [9] Antoine El-Hayek, Monika Henzinger, and Stefan Schmid. Brief announcement: Broadcasting time in dynamic rooted trees is linear. In Alessia Milani and Philipp Woelfel, editors, *PODC '22: ACM Symposium on Principles of Distributed Computing, Salerno, Italy, July 25 - 29, 2022*, pages 54–56. ACM, 2022.
- [10] 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.
- [11] Pierre Fraigniaud and Emmanuel Lazard. Methods and problems of communication in usual networks. *Discret. Appl. Math.*, 53(1-3):79–133, 1994.
- [12] 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.
- [13] Sandra Mitchell Hedetniemi, Stephen T. Hedetniemi, and Arthur L. Liestman. A survey of gossiping and broadcasting in communication networks. *Networks*, 18(4):319–349, 1988.
- [14] 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.
- [15] 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.
- [16] Nicola Santoro and Peter Widmayer. Time is not a healer. In Burkhard Monien and Robert Cori, editors, *STACS 89, 6th Annual Symposium on Theoretical Aspects of Computer Science, Paderborn, FRG, February 16-18, 1989, Proceedings*, volume 349 of *Lecture Notes in Computer Science*, pages 304–313. Springer, 1989.
- [17] Kyrill Winkler, Hugo Rincon Galeana, Ami Paz, Stefan Schmid, and Ulrich Schmid. The time complexity of consensus under oblivious message adversaries. In *14th Innovations in Theoretical Computer Science (ITCS)*, 2023.
- [18] Martin Zeiner, Manfred Schwarz, and Ulrich Schmid. On linear-time data dissemination in dynamic rooted trees. *Discret. Appl. Math.*, 255:307–319, 2019.

A Omitted Proofs

A.1 Basic Tools

Lemma 2.3. *Let $t, t' \geq 0, x, y \in [n]$. Then $x \in \mathcal{O}_t^{t'}(y) \Leftrightarrow y \in \mathcal{I}_t^{t'}(x)$.*

Proof. If $t > t' + 1$, the equivalence is empty. If $t = t' + 1$, then $x \in \mathcal{O}_t^{t'}(y) \Leftrightarrow x = y \Leftrightarrow y \in \mathcal{I}_t^{t'}(x)$. If $t \leq t'$, $x \in \mathcal{O}_t^{t'}(y) \Leftrightarrow (y, x) \in G_t \circ \dots \circ G_{t'} \Leftrightarrow y \in \mathcal{I}_t^{t'}(x)$. \square

Lemma 2.4 (Transitivity). *Let $t, t', t'' \geq 0$, and $x, y, z \in [n]$. We have the following properties:*

- i. *If $y \in \mathcal{O}_t^{t'}(x)$ and $z \in \mathcal{O}_{t'+1}^{t''}(y)$, then $z \in \mathcal{O}_t^{t''}(x)$.*
- ii. *If $x \in \mathcal{I}_t^{t'}(y)$ and $z \in \mathcal{O}_{t'+1}^{t''}(y)$, then $z \in \mathcal{O}_t^{t''}(x)$.*
- iii. *If $x \in \mathcal{I}_t^{t'}(y)$ and $y \in \mathcal{I}_{t'+1}^{t''}(z)$, then $z \in \mathcal{O}_t^{t''}(x)$.*

Proof. ii. and iii. will follow from i. and Lemma 2.3, so we only prove i..

Since $y \in \mathcal{O}_t^{t'}(x)$, this means $\mathcal{O}_t^{t'}(x) \neq \emptyset$. Therefore, $t \leq t' + 1$. If $t = t' + 1$, we have that $\mathcal{O}_t^{t'}(x) = \{x\} \Rightarrow y = x$, and thus $z \in \mathcal{O}_t^{t''}(x)$. We will then only need to consider the case $t \leq t'$.

Similarly, since $z \in \mathcal{O}_{t'+1}^{t''}(y)$, we have that $\mathcal{O}_{t'+1}^{t''}(y) \neq \emptyset$ and thus $t' + 1 \leq t'' + 1$. If $t' = t''$, we have that $\mathcal{O}_{t'+1}^{t''}(y) = \{y\}$ thus $z = y$ and therefore $z \in \mathcal{O}_t^{t''}(x)$. Hence we only need to consider the case $t \leq t' \leq t'' - 1$.

In that last case, we know that (x, y) is an edge in $G_t \circ \dots \circ G_{t'}$ and (y, z) is an edge in $G_{t'+1} \circ \dots \circ G_{t''}$. By definition of a product graph, this means that (x, z) is an edge in $G_t \circ \dots \circ G_{t''}$, and thus that $z \in \mathcal{O}_t^{t''}(x)$. \square

Lemma 2.5 (Monotonicity). *If in each round, all nodes have a self-loop, then for any $t_1 \leq t_2$ and $t_3 \leq t_4$, for any process x we have:*

- i. $\mathcal{I}_{t_2}^{t_3}(x) \subseteq \mathcal{I}_{t_1}^{t_4}(x)$.
- ii. $\mathcal{O}_{t_2}^{t_3}(x) \subseteq \mathcal{O}_{t_1}^{t_4}(x)$.

Proof. Let us first consider the case $t_2 > t_3 + 1$. Then $\mathcal{I}_{t_2}^{t_3}(x) = \mathcal{O}_{t_2}^{t_3}(x) = \emptyset$ and the result is trivial. Similarly, if $t_2 = t_3 + 1$, then $\mathcal{I}_{t_2}^{t_3}(x) = \mathcal{O}_{t_2}^{t_3}(x) = \{x\}$, and then either $t_1 = t_4 + 1$ which means $\mathcal{I}_{t_1}^{t_4}(x) = \mathcal{O}_{t_1}^{t_4}(x) = \{x\}$, or $t_1 \leq t_4$, and then x has a self-loop in G_t for every $t_1 \leq t \leq t_4$, which implies $x \in \mathcal{O}_{t_1}^{t_4}(x)$ and $x \in \mathcal{I}_{t_1}^{t_4}(x)$.

Let's now consider the case $t_2 \leq t_3$. Let $y \in \mathcal{I}_{t_2}^{t_3}(x)$ (respectively $z \in \mathcal{O}_{t_2}^{t_3}(x)$). Then edge (y, x) ((x, z)) is in graph $G_{t_2} \circ \dots \circ G_{t_3}$. Since all rounds have self-loops, we clearly have that edge (y, y) ((x, x)) is in the graph $G_{t_1} \circ \dots \circ G_{t_2}$. Similarly, edge (x, x) ((z, z)) is in graph $G_{t_3} \circ \dots \circ G_{t_4}$. This shows that (y, x) ((x, z)) is in graph $G_{t_1} \circ \dots \circ G_{t_4}$ and thus $y \in \mathcal{I}_{t_1}^{t_4}(x)$ ($z \in \mathcal{O}_{t_1}^{t_4}(x)$). \square

A.2 Broadcasting on Trees

A.2.1 The Lower Bound

The lower bound for this problem, due to Zeiner, Schwarz and Schmid [18] can be seen in Figure 5. The sequence of graphs in that example do not achieve broadcast before $\lceil \frac{3n-1}{2} \rceil - 2$ rounds. This examples proves the theorem:

Theorem 3.7. $t^*(\mathcal{T}_n) \geq \lceil \frac{3n-1}{2} \rceil - 2$




Figure 5: The lower bound example for the broadcasting on trees problem [18]. $[G]_a^b$ means that network G is the communication graph for all rounds between a and b (inclusive).




Figure 6: Main idea of the proof for the upper bound on covering on k -forests. An edge from process x in column t_1 to process y in column t_2 means that $x \in \mathcal{I}_{t_1}^{t_2-1}(y)$ (unless $t_2 = t'$ where we don't require the offset). We start with a cover of size n at t' then go back in time, finding a smaller cover at every step. After the first step, we don't need to cover 1 and 3 anymore, because if we reach p_{n-1} before round $t^{(n-1)}$, we also reach 1 and 3 before round t' , by going through p_{n-1} .

A.3 Covering on k -Forests

A.3.1 The Upper Bound

Lemma 4.19. *Let $s \geq k + 1$. If $s - k$ is odd, we have $\sum_{v:(s,v) \in E} w(s, v) = (\frac{s-k+1}{2})^2$. If it is even, we have $\sum_{v:(s,v) \in E} w(s, v) = (\frac{s-k}{2})^2 + \frac{s-k}{2}$.*

Proof. If $s - k$ is odd, let ℓ be such that $s - k = 2\ell + 1$:

$$\sum_{v:(s,v) \in E} w(s, v) = \sum_{v:(s,v) \in E} 2v - k - s = \sum_{v:v \leq s \leq 2v-k} 2v - k - s = \sum_{(s+k)/2 \leq v \leq s} 2v - k - s$$

but then $s - k \equiv s - k + 2k \pmod{2}$, so $s + k$ is also odd, and $\lceil \frac{s+k}{2} \rceil = \ell + k + 1$, therefore:

$$\begin{aligned} \sum_{v:(s,v) \in E} w(s, v) &= \sum_{\ell+k+1 \leq v \leq 2\ell+k+1} 2v - k - 2\ell - k - 1 = \sum_{0 \leq v \leq \ell} 2v + 1 \\ &= 2 \frac{\ell(\ell+1)}{2} + \ell + 1 = \frac{s-k-1}{2} \frac{s-k+1}{2} + \frac{s-k+1}{2} \end{aligned}$$

If $s - k$ is even, let ℓ be such that $s - k = 2\ell$:

$$\sum_{v:(s,v) \in E} w(s, v) = \sum_{v:(s,v) \in E} 2v - k - s = \sum_{v:v \leq s \leq 2v-k} 2v - k - s = \sum_{(s+k)/2 \leq v \leq s} 2v - k - s$$

Therefore:

$$\sum_{v:(s,v) \in E} w(s, v) = \sum_{\ell+k \leq v \leq 2\ell+k} 2v - k - 2\ell - k = \sum_{0 \leq v \leq \ell} 2v = 2 \frac{\ell(\ell+1)}{2} = \frac{s-k}{2} \frac{s-k+2}{2}$$

□

A.3.2 The Lower Bound

Our lower bound for this problem is very similar to the lower bound in Figure 5. The lower bound specific for this problem can be found in Figure 7. This lower bound is essentially “trapping” $k - 1$ vertices each in its own 1-vertex tree, and repeating the strategy for the lower bound with i vertices for broadcasting on trees on the last tree. Since none of the first $k - 1$ vertices can communicate with the others, cover is only achieved when broadcast is achieved on the last tree. Setting $i = n - k + 1$, we get a lower bound for the Covering on k -forests of $\lceil \frac{3n-3k}{2} - 1 \rceil$:

Theorem 4.21. $t_k^c(\mathcal{T}_n^k) \geq \lceil \frac{3n-3k}{2} - 1 \rceil$

A.4 k -Broadcasting on k -Rooted Networks

A.4.1 The Lower Bound

The lower bound for this problem is based on the lower bound for broadcasting on trees, in the Figure 5.

Theorem 5.9. $t_k^*(\mathcal{R}_n^k) \geq \lceil \frac{3n-9k}{2} + 2 \rceil$




Figure 7: The lower bound example for the Covering on k -Forests problem. $[G]_a^b$ means that network G is the communication graph for all rounds between a and b (inclusive). Each of the vertices $i + 1, \dots, i + k - 1$ is isolated in every round.




Figure 8: The lower bound example for the Covering on k -Forests problem. $[G]_a^b$ means that network G is the communication graph for all rounds between a and b (inclusive). K_k^a is the complete graph that replaces vertex a from the example in Figure 5. $K_k^a \rightarrow y$ (respectively $y \rightarrow K_k^a$) means that an edge is inserted from every $x \in K_k^a$ to y (respectively from y to every $x \in K_k^a$). Similarly, $K_k^a \rightarrow K_k^b$ means that an edge is inserted from every $x \in K_k^a$ to every $y \in K_k^b$.

Proof. The graph for this lower bound can be found in Figure 8. The idea is to reduce the k -broadcasting problem to the broadcast problem. More specifically, we use the sequence of networks from the lower bound of Figure 5 with i vertices for broadcasting on trees, while replacing each of the 3 vertices that act as root in the three networks of Figure 5 by k fully connected vertices, i.e. k vertices such that everyone points to the other $k - 1$. Then k -cover is only achieved when broadcast is achieved on the original networks. Setting $i = n - 3k + 3$, we get a lower bound for the k -broadcasting on k -rooted networks of $\lceil \frac{3n-9k}{2} + 2 \rceil$. \square