

Information Spreading in Dynamic Networks *

Chinmoy Dutta¹, Gopal Pandurangan², Rajmohan Rajaraman¹,
and Zhifeng Sun¹

1 College of Computer and Information Science,
Northeastern University
Boston, MA 02115, USA
{chinmoy, rraj, austin}@ccs.neu.edu

2 Division of Mathematical Sciences
Nanyang Technological University
Singapore 637371

and

Department of Computer Science
Brown University
Providence RI 02912, USA
gopalpandurangan@gmail.com

Abstract

We study the fundamental problem of information spreading (also known as gossip) in dynamic networks. In gossip, or more generally, k -gossip, there are k pieces of information (or tokens) that are initially present in some nodes and the problem is to disseminate the k tokens to all nodes. The goal is to accomplish the task in as few rounds of distributed computation as possible. The problem is especially challenging in dynamic networks where the network topology can change from round to round and can be controlled by an on-line adversary.

The focus of this paper is on the power of token-forwarding algorithms, which do not manipulate tokens in any way other than storing and forwarding them. We first consider a worst-case adversarial model first studied by Kuhn, Lynch, and Oshman [33] in which the communication links for each round are chosen by an adversary, and nodes do not know who their neighbors for the current round are before they broadcast their messages. Our main result is an $\Omega(nk/\log n)$ lower bound on the number of rounds needed for any deterministic token-forwarding algorithm to solve k -gossip. This resolves an open problem raised in [33], improving their lower bound of $\Omega(n \log k)$, and matching their upper bound of $O(nk)$ to within a logarithmic factor.

Our result shows that one cannot obtain significantly efficient (i.e., subquadratic) token-forwarding algorithms for gossip in the adversarial model of [33]. We next show that token-forwarding algorithms can achieve subquadratic time in the offline version of the problem, where the adversary has to commit all the topology changes in advance at the beginning of the computation. We present two polynomial-time offline token-forwarding algorithms to solve k -gossip: (1) an $O(\min\{nk, n^{1.5}\sqrt{\log n}\})$ round algorithm, and (2) an $(O(n^\epsilon), \log n)$ bicriteria approximation algorithm, for any $\epsilon > 0$, which means that if L is the number of rounds needed by an optimal algorithm, then our approximation algorithm will complete in $O(n^\epsilon L)$ rounds and the number of tokens transmitted on any edge is $O(\log n)$ in each round. Our results are a step towards understanding the power and limitation of token-forwarding algorithms in dynamic networks.

1998 ACM Subject Classification F.2 [ANALYSIS OF ALGORITHMS AND PROBLEM COMPLEXITY]: Non-numerical Algorithms and Problems—computations on discrete structures; G.2.2

* This work supported in part by the following grants: Nanyang Technological University grant M58110000, Singapore Ministry of Education (MOE) Academic Research Fund (AcRF) Tier 2 grant MOE2010-T2-2-082, US NSF grants CCF-1023166 and CNS-0915985, and a grant from the US-Israel Binational Science Foundation (BSF).



© C. Dutta, G. Pandurangan, R. Rajaraman and Z. Sun;
licensed under Creative Commons License NC-ND

Conference title on which this volume is based on.

Editors: Billy Editor, Bill Editors; pp. 1–15



Leibniz International Proceedings in Informatics

LIPIC Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

[Discrete Mathematics]: Graph Theory—graph algorithms G.2.2 [Discrete Mathematics]: Graph Theory—network problems; C.2.4 [COMPUTER COMMUNICATION NETWORKS]: Distributed Systems

Keywords and phrases Dynamic networks, Distributed Computation, Information Spreading, Gossip, Lower Bounds

Digital Object Identifier 10.4230/LIPIcs.xxx.yyy.p

1 Introduction

In a dynamic network, nodes (processors/end hosts) and communication links can appear and disappear at will over time. Emerging networking technologies such as ad hoc wireless, sensor, and mobile networks, overlay and peer-to-peer (P2P) networks are inherently dynamic, resource-constrained, and unreliable. This necessitates the development of a solid theoretical foundation to design efficient, robust, and scalable distributed algorithms and to understand the power and limitations of distributed computing on such networks. Such a foundation is critical to realize the full potential of these large-scale dynamic communication networks.

As a step towards understanding the fundamental computation power of dynamic networks, we investigate dynamic networks in which the network topology changes arbitrarily from round to round. We first consider a worst-case model that was introduced by Kuhn, Lynch, and Oshman [33] in which the communication links for each round are chosen by an online adversary, and nodes do not know who their neighbors for the current round are before they broadcast their messages. (Note that in this model, only edges change and nodes are assumed to be fixed.) The only constraint on the adversary is that the network should be connected at each round. Unlike prior models on dynamic networks, the model of [33] does not assume that the network eventually stops changing and requires that the algorithms work correctly and terminate even in networks that change continually over time.

In this paper, we study the fundamental problem of information spreading (also known as gossip). In gossip, or more generally, k -gossip, there are k pieces of information (or tokens) that are initially present in some nodes and the problem is to disseminate the k tokens to all nodes. (By just gossip, we mean n -gossip, where n is the network size.) Information spreading is a fundamental primitive in networks which can be used to solve other problem such as broadcasting and leader election. Indeed, solving n -gossip, where the number of tokens is equal to the number of nodes in the network, and each node starts with exactly one token, allows any function of the initial states of the nodes to be computed, assuming that the nodes know n [33].

1.1 Our results

The focus of this paper is on the power of *token-forwarding* algorithms, which do not manipulate tokens in any way other than storing and forwarding them. Token-forwarding algorithms are simple, often easy to implement, and typically incur low overhead. In a key result, [33] showed that under their adversarial model, k -gossip can be solved by token-forwarding in $O(nk)$ rounds, but that any deterministic online token-forwarding algorithm needs $\Omega(n \log k)$ rounds. They also proved an $\Omega(nk)$ lower bound for a special class of token-forwarding algorithms, called knowledge-based algorithms. Our main result is a new lower bound on *any* deterministic online token-forwarding algorithm for k -gossip.

- We show that every deterministic online token-forwarding algorithm for the k -gossip problem takes $\Omega(nk/\log n)$ rounds. Our result applies even to centralized (deterministic) token-forwarding algorithms that have a global knowledge of the token distribution.

This result resolves an open problem raised in [33], significantly improving their lower bound, and matching their upper bound to within a logarithmic factor. Our lower bound also enables a better comparison of token-forwarding with an alternative approach based on network coding due to [27, 28], which achieves a $O(nk/\log n)$ rounds using $O(\log n)$ -bit messages (which is not significantly better than the $O(nk)$ bound using token-forwarding), and $O(n+k)$ rounds with large message sizes (e.g., $\Theta(n \log n)$ bits). It thus follows that for large token and message sizes there is a factor $\Omega(\min\{n, k\}/\log n)$ gap between token-forwarding and network coding. We note that in our model we allow only one token per edge per round and thus our bounds hold regardless of the token size.

Our lower bound indicates that one cannot obtain efficient (i.e., subquadratic) token-forwarding algorithms for gossip in the adversarial model of [33]. Furthermore, for arbitrary token sizes, we do not know of any algorithm that is significantly faster than quadratic time. This motivates considering other weaker (and perhaps, more realistic) models of dynamic networks. In fact, it is not clear whether one can solve the problem significantly faster even in an offline setting, in which the network can change arbitrarily each round, but the entire evolution is known to the algorithm in advance. Our next contribution takes a step in resolving this basic question for token-forwarding algorithms.

- We present a polynomial-time offline token-forwarding algorithm that solves the k -gossip problem on an n -node dynamic network in $O(\min\{nk, n^{1.5}\sqrt{\log n}\})$ rounds.
- We also present a polynomial-time offline token-forwarding algorithm that solves the k -gossip problem in a number of rounds within an $O(n^\epsilon)$ factor of the optimal, for any $\epsilon > 0$, assuming the algorithm is allowed to transmit $O(\log n)$ tokens per round.

The above upper bounds show that in the offline setting, token-forwarding algorithms can achieve a time bound that is within $O(\sqrt{n \log n})$ of the information-theoretic lower bound of $\Omega(n+k)$, and that we can approximate the best token-forwarding algorithm to within a $O(n^\epsilon)$ factor, given logarithmic extra bandwidth per edge.

1.2 Related work

Information spreading (or dissemination) in networks is one of the most basic problems in computing and has a rich literature. The problem is generally well-understood on static networks, both for interconnection networks [34] as well as general networks [35, 6]. In particular, the k -gossip problem can be solved in $O(n+k)$ rounds on any n -static network [38]. There also have been several papers on broadcasting, multicasting, and related problems in static heterogeneous and wireless networks (e.g., see [5, 15, 14, 22]).

Dynamic networks have been studied extensively over the past three decades. Some of the early studies focused on dynamics that arise out of faults, i.e., when edges or nodes fail. A number of fault models, varying according to extent and nature (e.g., probabilistic vs. worst-case) and the resulting dynamic networks have been analyzed (e.g., see [6, 35]). There have been several studies on models that constrain the rate at which changes occur, or assume that the network eventually stabilizes (e.g., see [1, 25, 26]).

There also has been considerable work on general dynamic networks. Some of the earliest studies in this area include [2, 13] which introduce general building blocks for communication protocols on dynamic networks. Another notable work is the local balancing approach of [11] for solving routing and multicommodity flow problems on dynamic networks. Algorithms based on the local balancing approach continually balance the packet queues across each edge

of the network and drain packets that have reached their destination. The local balancing approach has been applied to achieve near-optimal throughput for multicast, anycast, and broadcast problems on dynamic networks as well as for mobile ad hoc networks [10, 12, 29].

Modeling general dynamic networks has gained renewed attention with the recent advent of heterogeneous networks composed out of ad hoc, and mobile devices. To address the unpredictable and often unknown nature of network dynamics, [33] introduce a model in which the communication graph can change completely from one round to another, with the only constraint being that the network is connected at each round. The model of [33] allows for a much stronger adversary than the ones considered in past work on general dynamic networks [11, 10, 12]. In addition to results on the k -gossip problem that we have discussed earlier, [33] consider the related problem of counting, and generalize their results to the T -interval connectivity model, which includes an additional constraint that any interval of T rounds has a stable connected spanning subgraph. The survey of [32] summarizes recent work on dynamic networks.

We note that the model of [33], as well as ours, allow only edge changes from round to round while the nodes remain fixed. Recently, the work of [7] introduced a dynamic network model (motivated by P2P networks) where both nodes and edges can change by a large amount (up to a linear fraction of the network size). They show that stable almost-everywhere agreement can be efficiently solved in such networks even in adversarial dynamic settings.

Recent work of [27, 28] presents information spreading algorithms based on network coding [4]. As mentioned earlier, one of their important results is that the k -gossip problem on the adversarial model of [33] can be solved using network coding in $O(n + k)$ rounds assuming the token sizes are sufficiently large ($\Omega(n \log n)$ bits). For further references to using network coding for gossip and related problems, we refer to the recent works of [27, 28, 8, 17, 23, 36] and the references therein.

Our offline approximation algorithm makes use of results on the Steiner tree packing problem for directed graphs [21]. This problem is closely related to the directed Steiner tree problem (a major open problem in approximation algorithms) [19, 39] and the gap between network coding and flow-based solutions for multicast in arbitrary directed networks [3, 37].

Finally, we note that there are also a number of studies that solve k -gossip and related problems using *gossip-based* processes. In a local gossip-based algorithm, each node exchanges information with a small number of randomly chosen neighbors in each round. Gossip-based processes have recently received significant attention because of their simplicity of implementation, scalability to large network size, and their use in aggregate computations, e.g., [16, 24, 31, 20, 30, 36, 18] and the references therein. All these studies assume an underlying static communication network, and do not apply directly to the models considered in this paper. A related recent work on dynamic networks is [9] which analyzes the cover time of random walks on dynamic networks.

2 Model and problem statement

In this section, we formally define the k -gossip problem, the online and offline models, and token-forwarding algorithms.

The k -gossip problem. In this problem, k different tokens are assigned to a set V of $n \geq k$ nodes, where each node may have any subset of the tokens, and the goal is to disseminate all the k tokens to all the nodes.

The online model. Our online model is the worst-case adversarial model of [33]. Nodes communicate with each other using anonymous broadcast. We assume a synchronized

communication. At the beginning of round r , each node in V decides what message to broadcast based on its internal state and coin tosses (for a randomized algorithm); at the same time and independently, the adversary chooses the set of edges that forms the communication network G_r over V for round r . The only constraint on G_r is that it be connected; this is the same as the 1-interval connectivity model of [33]. As observed in [33], deterministic algorithms choose a message based only on their internal state, so this is equivalent to the adversary knowing the messages to be sent in round r before choosing the edges for round r . We do not place any bound on the size of the messages, but require for our lower bound that each message contains at most one token.

The offline model. In the offline model, we are given a sequence of networks $\langle G_r \rangle$ where G_r is a connected communication network for round r . As in the online model, we assume that in each round at most one token is broadcast by any node. It can be easily seen that the k -gossip problem can be solved in nk rounds in the offline model; so we may assume that the given sequence of networks is of length at most nk .

Token-forwarding algorithms. Informally, a token-forwarding algorithm is one that does not combine or alter tokens, only stores and forwards them. Formally, we call an algorithm for k -gossip a token-forwarding algorithm if for every node v , token t , and round r , v contains t at the start of round r of the algorithm if and only if either v has t at the start of the algorithm or v received a message containing t prior to round r .

Finally, several of our arguments are probabilistic. We use the term “with high probability” to mean with probability at least $1 - 1/n^c$, for a constant c that can be made sufficiently high by adjusting related constant parameters.

3 Lower bound for deterministic online token-forwarding algorithms

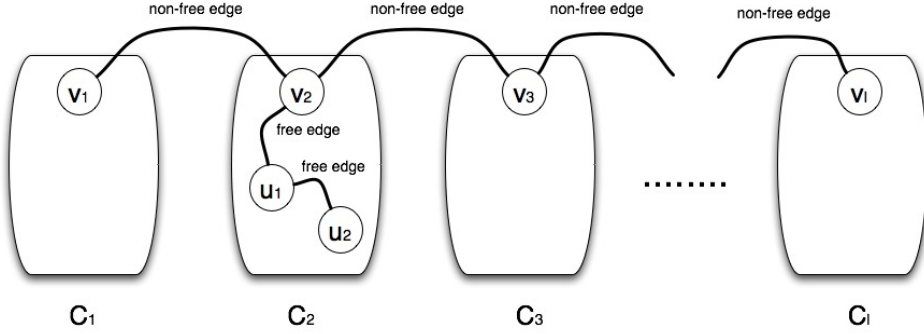
In this section, we give an $\Omega(kn/\log n)$ lower bound on the number of rounds needed by any deterministic online token-forwarding algorithm for the k -gossip problem. Our lower bound applies to even centralized algorithms and a large class of initial token distributions. We first describe the adversary strategy.

Adversary: The strategy of the adversary is simple. We use the notion of *free edge* introduced in [33]. In a given round r , we call an edge (u, v) to be a free edge if at the start of round r , u has the token that v broadcasts in the round and v has the token that u broadcasts in the round¹; an edge that is not free is called *non-free*. Thus, if (u, v) is a free edge in a particular round, neither u nor v can gain any new token through this edge in the round. Since we are considering deterministic token-forwarding algorithms, at the start of each round, the adversary knows for each node v , the token (if any) that v will broadcast in that round. In round r , the adversary constructs the communication graph G_r as follows. First, the adversary adds all the free edges to G_r . Let C_1, C_2, \dots, C_l denote the connected components thus formed. The adversary then guarantees the connectivity of the graph by selecting an arbitrary node in each connected component and connecting them in a line. Figure 1 illustrates the construction.

The network G_r thus constructed has exactly $l - 1$ non-free edges, where l is the number of connected components formed by the free edges of G_r . If (u, v) is a non-free edge in G_r , then u , v , or both will gain at most new token through this edge. We refer to such a token exchange on a non-free edge as a *useful token exchange*.

¹ For convenience, when a node does not broadcast any token we will view it as broadcasting a special *empty* token that every node has. This allows us to avoid treating the empty broadcast as a special case.

We bound the running-time of any token-forwarding algorithm by identifying a critical structure that quantifies the progress made in each round. We say that a sequence of nodes v_1, v_2, \dots, v_k is *half-empty* in round r with respect to a sequence of tokens t_1, t_2, \dots, t_k if the following condition holds at the start of round r : for all $1 \leq i, j \leq k$, $i \neq j$, either v_i is missing t_j or v_j is missing t_i . We then say that $\langle v_i \rangle$ is half-empty with respect to $\langle t_i \rangle$ and refer to the pair $(\langle v_i \rangle, \langle t_i \rangle)$ as a half-empty configuration of size k .



■ **Figure 1** The network constructed by the adversary in a particular round. Note that if node v_i broadcasts token t_i , then the $\langle v_i \rangle$ forms a half-empty configuration with respect to $\langle t_i \rangle$ at the start of this round.

► **Lemma 1.** *If m useful token exchanges occur in round r , then there exists a half-empty configuration of size at least $m/2 + 1$ at the start of round r .*

Proof. Consider the network G_r in round r . Each non-free edge can contribute at most 2 useful token exchanges. Thus, there are at least $m/2$ non-free edges in the communication graph. Based on the adversary we consider, no useful token exchange takes place within the connected components induced by the free edges. Useful token exchanges can only happen over the non-free edges between connected components. This implies there are at least $m/2 + 1$ connected components in the subgraph of G_r induced by the free edges. Let v_i denote an arbitrary node in the i th connected component in this subgraph, and let t_i be the token broadcast by v_i in round r . For $i \neq j$, since v_i and v_j are in different connected components, (v_i, v_j) is a non-free edge in round r ; hence, at the start of round r , either v_i is missing t_j or v_j is missing t_i . Thus, the sequence $\langle v_i \rangle$ of nodes of size at least $m/2 + 1$ is half-empty with respect to the sequence $\langle t_i \rangle$ at the start of round r . ◀

An important point to note about the definition of a semi-anti-clique is that it only depends on the token distribution; it is independent of the broadcast in any round. This allows us to prove the following easy lemma.

► **Lemma 2.** *If a sequence $\langle v_i \rangle$ of nodes is half-empty with respect to $\langle t_i \rangle$ at the start of round r , then $\langle v_i \rangle$ is half-empty with respect to $\langle t_i \rangle$ at the start of round r' for any $r' \leq r$.*

Proof. The lemma follows immediately from the fact that if a node v_i is missing a token t_j at the start of round r , then v_i is missing token t_j at the start of every round $r' < r$. ◀

Lemmas 1 and 2 suggest that if we can identify a token distribution in which all half-empty configurations are small, we can guarantee small progress in each round. We now show that there are many token distributions with this property, thus yielding the desired lower bound.

► **Theorem 3.** *From an initial token distribution in which each node has each token independently with probability $3/4$, any deterministic online token-forwarding algorithm will need $\Omega(kn/\log n)$ rounds to complete with high probability.*

Proof. We first note that if the number of tokens k is less than $5 \log n$, then the number of rounds needed is at least $\Omega(n)$ because even to disseminate one token will take $\Omega(n)$ rounds in the worst-case. The $\Omega(kn/\log n)$ lower bound is trivially true in this case. Thus, in the following proof, we focus on the case where $k \geq 5 \log n$.

Let E_l denote the event that there exists a half-empty configuration of size l at the start of the first round. For E_l to hold, we need l nodes v_1, v_2, \dots, v_l and l tokens t_1, t_2, \dots, t_l such that for all $i \neq j$ either v_i is missing t_j or v_j is missing t_i . For a pair of nodes u and v , by union bound, the probability that u is missing t_v or v is missing t_u is at most $1/4 + 1/4 = 1/2$. Thus, the probability of E_l can be bounded as follows.

$$\Pr[E_l] \leq \binom{n}{l} \cdot \frac{k!}{(k-l)!} \cdot \left(\frac{1}{2}\right)^{\binom{l}{2}} \leq n^l \cdot k^l \frac{1}{2^{l(l-1)/2}} \leq \frac{2^{2l \log n}}{2^{l(l-1)/2}}$$

When $l \geq 5 \log n$, $\Pr[E_l] \leq 1/n^2$. Thus, the largest semi-anti-clique at the start of the first round is of size at most $5 \log n$. It follows that the largest half-empty configuration at the start of any round is of size at most $5 \log n$. By Lemma 1, we thus obtain that the number of useful token exchanges in each round is at most $10 \log n$, with probability at least $1 - 1/n^2$.

Let M_i be the number of tokens that node i is missing in the initial distribution. Then M_i is a binomial random variable with $\mathbb{E}[M_i] = 3k/4$. By a straightforward Chernoff bound, we have the probability that node i misses less than $k/8$ tokens is

$$\Pr\left[M_i \geq \frac{3k}{4} + \frac{k}{8}\right] = \Pr\left[M_i \geq \left(1 + \frac{1}{6}\right) \cdot \mathbb{E}[M_i]\right] \leq e^{-\mathbb{E}[M_i] \left(\frac{1}{6}\right)^2 \frac{1}{2}} = e^{-\frac{3k}{4} \left(\frac{1}{6}\right)^2 \frac{1}{2}} = \frac{1}{e^{96k}}$$

Therefore, the total number of tokens missing in the initial distribution is at least $n \cdot k/8 = \Omega(kn)$ with probability at least $1 - n/e^{96k} \geq 1 - 1/n^2$ ($k \geq 5 \log n$). Since the number of useful tokens exchanged in each round is at most $10 \log n$, the number of rounds needed to complete k -gossip is $\Omega(kn/\log n)$ with high probability. ◀

Theorem 3 does not apply to certain natural initial distributions, such as one in which each token resides at exactly one node. While this class of token distributions has far fewer tokens distributed initially, the argument of Theorem 3 does not rule out the possibility that an algorithm, when starting from a distribution in this class, avoids the problematic configurations that arise in the proof. The following theorem extends the lower bound to this class of distributions.

► **Theorem 4.** *From any distribution in which each token starts at exactly one node, any deterministic online token-forwarding algorithm for k -gossip needs $\Omega(kn/\log n)$ rounds.*

Proof. We consider an initial distribution C where each token is at exactly one node, and no node has more than one token. Our argument can be easily extended to the case where nodes may have multiple tokens (we defer its proof to the full paper). Let C^* be an initial token distribution from which any deterministic online algorithm needs $\Omega(kn/\log n)$ rounds. The existence of C^* follows from Theorem 3. We construct a bipartite graph on two copies of V , V_1 and V_2 . A node $v \in V_1$ is connected to a node $u \in V_2$ if in C^* u has all the tokens that v has in C . We will show below that this bipartite graph has a perfect matching.

Given a perfect matching M , we can complete the proof as follows. For $v \in V_2$, let $M(v)$ denote the node in V_1 that got matched to v . If there is an algorithm A that runs in T

rounds from starting state C , then we can construct an algorithm A^* that runs in the same number of rounds from starting state C^* as follows. First every node v deletes all its tokens except for those which $M(v)$ has in C . Then algorithm A^* runs exactly as A . Thus, the lower bound of Theorem 3, which applies to A^* , also applies to A .

It remains to prove that the above bipartite graph has a perfect matching. This follows from an application of Hall's Theorem. Consider a set of m nodes in V_2 . We want to show their neighborhood in the bipartite graph is of size at least m . We show this condition holds by the following 2 cases. If $m < 3n/5$, let X_i denote the neighborhood size of node i . We know $\mathbb{E}[X_i] \geq 3n/4$. Then by Chernoff bound

$$\Pr[X_i < m] \leq \Pr[X_i < 3n/5] \leq e^{-\frac{(1/5)^2 \mathbb{E}[X_i]}{2}} = e^{-\frac{3n}{200}}$$


By union bound with probability at least $1 - n \cdot e^{-3n/200}$ the neighborhood size of every node is at least m . Therefore, the condition holds in the first case. If $m \geq 3n/5$, we argue the neighborhood size of any set of m nodes is V_1 with high probability. Consider a set of m nodes, the probability that a given token t is missing in all these m nodes is $(1/4)^m$. Thus the probability that any token is missing in all these nodes is at most $n(1/4)^m \leq n(1/4)^{3n/5}$. There are at most 2^n such sets. By union bound, with probability at least $1 - 2^n \cdot n(1/4)^{3n/5} = 1 - n/2^{n/5}$, the condition holds in the second case. ◀

4 Subquadratic time offline token-forwarding algorithms

In this section, we give two centralized algorithms for the k -gossip problem in the offline model. We present an $O(\min\{n^{1.5}\sqrt{\log n}, nk\})$ round algorithm in Section 4.1. Then we present a bicriteria $(O(n^\epsilon), \log n)$ -approximation algorithm in Section 4.2, which means if L is the number of rounds needed by an optimal algorithm where one token is broadcast by every node per round, then our approximation algorithm will complete in $O(n^\epsilon L)$ rounds and the number of tokens broadcast by any node is $O(\log n)$ in any given round. Both of these algorithms uses a directed capacitated leveled graph constructed from the sequence of communication graphs which we call the *evolution graph*.

Evolution graph: Let V be the set of nodes. Consider a dynamic network of l rounds numbered 1 through l and let G_i be the communication graph for round i . The evolution graph for this network is a directed capacitated graph G with $2l + 1$ levels constructed as follows. We create $2l + 1$ copies of V and call them V_0, V_2, \dots, V_{2l} . V_i is the set of nodes at level i and for each node v in V , we call its copy in V_i as v_i . For $i = 1, \dots, l$, level $2i - 1$ corresponds to the beginning of round i and level $2i$ corresponds to the end of round i . Level 0 corresponds to the network at the start. Note that the end of a particular round and the start of the next round are represented by different levels. There are three kinds of edges in the graph. First, for every round i and every edge $(u, v) \in G_i$, we place two directed edges with unit capacity each, one from u_{2i-1} to v_{2i} and another from v_{2i-1} to u_{2i} . We call these edges *broadcast edges* as they will correspond to broadcasting of tokens; the unit capacity on each such edge will ensure that only one token can be sent from a node to a neighbor in one round. Second, for every node v in V and every round i , we place an edge with infinite capacity from $v_{2(i-1)}$ to v_{2i} . We call these edges *buffer edges* as they ensure tokens can be stored at a node from the end of one round to the end of the next. Finally, for every node $v \in V$ and every round i , we also place an edge with unit capacity from $v_{2(i-1)}$ to v_{2i-1} . We call these edges as *selection edges* as they correspond to every node selecting a token out of those it has to broadcast in round i ; the unit capacity ensures that in a given round a

node must send the same token to all its neighbors. Figure 2 illustrates our construction, and Lemma 10 (moved to the appendix due to space constraints) explains its usefulness.



■ **Figure 2** An example of how to construct the evolution graph from a sequence of communication graphs.

4.1 An $O(\min\{n^{1.5}\sqrt{\log n}, nk\})$ round algorithm

Our algorithm is given in Algorithm 1 and analyzed in Lemma 5 and 6, whose proofs are deferred to the appendix due to space constraints.

► **Lemma 5.** *Let there be $k \leq n$ tokens at given source nodes and let v be an arbitrary node. Then, all the tokens can be sent to v using broadcasts in $O(n)$ rounds.*

Algorithm 1 $O(\min\{n^{1.5}\sqrt{\log n}, nk\})$ round algorithm in the offline model

```

1: if  $k \leq \sqrt{n \log n}$  then
2:   for each token  $t$  do
3:     For the next  $n$  rounds, let every node who has token  $t$  broadcast the token.
4:   end for
5: else
6:   Choose a set  $S$  of  $\sqrt{n \log n}$  random nodes.
7:   for each vertex in  $v \in S$  do
8:     Send each of the  $k$  tokens to vertex  $v$  in  $O(n)$  rounds.
9:   end for
10:  for each token  $t$  do
11:    For the next  $\sqrt{n \log n}$  rounds, let every node who has token  $t$  broadcast the token.
12:  end for
13: end if

```

► **Theorem 6.** *Algorithm 1 solves the k -gossip problem using $O(\min\{n^{1.5}\sqrt{\log n}, nk\})$ rounds with high probability in the offline model.*

Algorithm 1 can be derandomized using the standard technique of conditional expectations. We defer the proof of this claim to the full paper.

4.2 An $(O(n^\epsilon), \log n)$ -approximation algorithm

Here we introduce an $(O(n^\epsilon), \log n)$ -approximation algorithm for the k -gossip problem in the offline model. This means, if the k -gossip problem can be solved on any n -node dynamic network in L rounds, then our algorithm will solve the k -gossip problem on any dynamic network in $O(n^\epsilon L)$ rounds, assuming each node is allowed to broadcast $O(\log n)$ tokens, instead of one, in each round. Our algorithm is an LP based one, which makes use of the evolution graph defined earlier. The following is a straightforward corollary of Lemma 10.

► **Corollary 7.** *The k -gossip problem can be solved in l rounds if k directed Steiner trees can be packed in the corresponding evolution graph, where for each token, the root of its Steiner tree is a source node at level 0, and the terminals are all the nodes at level $2l$.*

Packing Steiner trees in general directed graphs is NP-hard to approximate even within $\Omega(m^{1/3-\epsilon})$ for any $\epsilon > 0$ [21], where m is the number of edges in the graph. Thus, our algorithm focuses on solving Steiner tree packing problem with relaxation on edge capacities, allowing the capacity to blow up by a factor of $O(\log n)$. First, we write down the LP for the Steiner tree packing problem (maximizing the number of Steiner trees packed with respect to edge capacities). Let \mathcal{T} be the set of all possible Steiner trees, and c_e be the capacity of edge e . For each Steiner tree $T \in \mathcal{T}$, we associate a variable x_T with it. If $x_T = 1$, then Steiner tree T is in the optimal solution; if $x_T = 0$, it's not. After relaxing the integral constraints on x_T 's, we have the following LP, referred to as \mathcal{P} henceforth. Let $F(\mathcal{P})$ denote the optimal fractional solution for \mathcal{P} .

$$\begin{aligned} \max \quad & \sum_{T \in \mathcal{T}} x_T \\ \text{s.t.} \quad & \sum_{T: e \in T} x_T \leq c_e \quad \forall e \in E \\ & x_T \geq 0 \quad \forall T \in \mathcal{T} \end{aligned}$$

► **Lemma 8** ([21]). *There is an $O(n^\epsilon)$ -approximation algorithm for the fractional maximum Steiner tree packing problem in directed graphs.*

Let L be the number of rounds that an optimal algorithm uses with every node broadcasting at most one token per round. We give an algorithm that takes $O(n^\epsilon L)$ rounds with every node broadcasting $O(\log n)$ tokens per round. Thus ours is an $(O(n^\epsilon), O(\log n))$ bicriteria approximation algorithm, shown in Algorithm 2.

► **Theorem 9.** *Algorithm 2 achieves an $O(n^\epsilon)$ approximation to the k -gossip problem while broadcasting $O(\log n)$ tokens per round per node, with high probability.*

5 Conclusion and open questions

In this paper, we studied the power of token-forwarding algorithms for gossip in dynamic networks. We showed a lower bound of $\Omega(nk/\log n)$ rounds for any deterministic online token forwarding algorithm which matches the known upper bound of $O(nk)$ up to a logarithmic factor. This leaves us with an important open question: what is the complexity of *randomized* online token-forwarding algorithms? In fact, for small token sizes (e.g., $O(\log n)$ bits) even the best online algorithm we know based on network coding takes $O(nk/\log n)$ rounds [28]. In contrast, we show that in the offline setting there exist centralized token-forwarding algorithms that run in $O(n^{1.5}\sqrt{\log n})$ time. An interesting open problem is to obtain tight bounds on offline token-forwarding algorithms.

Algorithm 2 ($O(n^\epsilon), O(\log n)$)-approximation algorithm

- 1: Initialize the set of Steiner trees $\mathcal{S} = \emptyset$.
 - 2: **for** $i = 1 \rightarrow 2n^\epsilon$ **do**
 - 3: Find L^* such that with the evolution graph G constructed from level 0 to level $2L^*$, the approximate value for $F(\mathcal{P})$ is k/n^ϵ . In this step, we use the algorithm of [21] to approximate $F(\mathcal{P})$.
 - 4: Let x_T^* be the value of the variable x_T in the solution from step 3. The number of non-zero x_T^* 's is polynomial with respect to k . Using randomized rounding, with probability x_T^* include T in the solution, $\mathcal{S} = \mathcal{S} \cup \{T\}$. Otherwise, don't include T .
 - 5: Remove communication graphs G_1, G_2, \dots, G_{L^*} from the sequence, and reduce the remaining graphs' indices by L^* .
 - 6: **end for**
 - 7: Use Corollary 7 to convert the set of Steiner trees \mathcal{S} into a token dissemination schedule.
-

References

- 1 Yehuda Afek, Baruch Awerbuch, and Eli Gafni. Applying static network protocols to dynamic networks. In *FOCS'87*, pages 358–370, 1987.
- 2 Yehuda Afek, Eli Gafni, and Adi Rosen. The slide mechanism with applications in dynamic networks. In *ACM PODC*, pages 35–46, 1992.
- 3 A. Agarwal and M. Charikar. On the advantage of network coding for improving network throughput. In *Information Theory Workshop*, 2004.
- 4 R. Ahlswede, N. Cai, S. Li, and R. Yeung. Network information flow. *Transactions on Information Theory*, 46(4):1204–1216, 2000.
- 5 N. Alon, A. Bar-Noy, N. Linial, and D. Peleg. A lower bound for radio broadcast. *Journal of Computer and System Sciences*, 43:290–298, 1991.
- 6 Hagit Attiya and Jennifer Welch. *Distributed Computing: Fundamentals, Simulations and Advanced Topics (2nd edition)*. John Wiley Interscience, March 2004.
- 7 John Augustine, Gopal Pandurangan, Peter Robinson, and Eli Upfal. Towards robust and efficient computation in dynamic peer-to-peer networks. In *ACM-SIAM SODA*, 2012.
- 8 Chen Avin, Michael Borokhovich, Keren Censor-Hillel, and Zvi Lotker. Order optimal information spreading using algebraic gossip. In *ACM PODC*, pages 363–372, 2011.
- 9 Chen Avin, Michal Koucký, and Zvi Lotker. How to explore a fast-changing world (cover time of a simple random walk on evolving graphs). In *ICALP*, pages 121–132, 2008.
- 10 B. Awerbuch, P. Berenbrink, A. Brinkmann, and C. Scheideler. Simple routing strategies for adversarial systems. In *IEEE FOCS*, pages 158–167, 2001.
- 11 B. Awerbuch and F. T. Leighton. Improved approximation algorithms for the multi-commodity flow problem and local competitive routing in dynamic networks. In *ACM STOC*, pages 487–496, May 1994.
- 12 Baruch Awerbuch, André Brinkmann, and Christian Scheideler. Anycasting in adversarial systems: Routing and admission control. In *ICALP'03*, pages 1153–1168, 2003.
- 13 Baruch Awerbuch, Boaz Patt-Shamir, David Peleg, and Michael E. Saks. Adapting to asynchronous dynamic networks. In *STOC'92*, pages 557–570, 1992.
- 14 Amotz Bar-Noy, Sudipto Guha, Joseph Naor, and Baruch Schieber. Message multicasting in heterogeneous networks. *SIAM J. Comput.*, pages 347–358, 2000.
- 15 R. Bar-Yehuda, O. Goldreich, and A. Itai. On the time-complexity of broadcast in radio networks: an exponential gap between determinism and randomization. In *ACM PODC*, pages 98–108, 1987.

- 16 Petra Berenbrink, Jurek Czyzowicz, Robert Elsässer, and Leszek Gasieniec. Efficient information exchange in the random phone-call model. In *ICALP*, pages 127–138, 2010.
- 17 Michael Borokhovich, Chen Avin, and Zvi Lotker. Tight bounds for algebraic gossip on graphs. In *IEEE ISIT*, 2010.
- 18 S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Trans. on Infor. Theory*, 52(6):2508–2530, 2006.
- 19 M. Charikar, C. Chekuri, T. Cheung, Z. Dai, A. Goel, and S. Guha. Approximation algorithms for directed steiner problems. *Journal of Algorithms*, 1998.
- 20 J. Chen and G. Pandurangan. Optimal gossip-based aggregate computation. In *SPAA*, pages 124–133, 2010.
- 21 J. Cheriyan and M. Salavatipour. Hardness and approximation results for packing steiner trees. *Algorithmica*, 2006.
- 22 Andrea E. F. Clementi, Angelo Monti, and Riccardo Silvestri. Distributed multi-broadcast in unknown radio networks. In *PODC'01*, pages 255–264, 2001.
- 23 S. Deb, M. Médard, and C. Choute. Algebraic gossip: a network coding approach to optimal multiple rumor mongering. *IEEE/ACM Trans. Netw.*, 14:2486–2507, June 2006.
- 24 Alan Demers, Dan Greene, Carl Hauser, Wes Irish, John Larson, Scott Shenker, Howard Sturgis, Dan Swinehart, and Doug Terry. Epidemic algorithms for replicated database maintenance. In *PODC*, pages 1–12, 1987.
- 25 Shlomi Dolev. *Self-stabilization*. MIT Press, Cambridge, MA, USA, 2000.
- 26 E. Gafni and B. Bertsekas. Distributed algorithms for generating loop-free routes in networks with frequently changing topology. *IEEE Trans. Comm.*, 29(1):11–18, 1981.
- 27 Bernhard Haeupler. Analyzing network coding gossip made easy. In *ACM STOC*, pages 293–302, 2011.
- 28 Bernhard Haeupler and David Karger. Faster information dissemination in dynamic networks via network coding. In *ACM PODC*, pages 381–390, 2011.
- 29 L. Jia, R. Rajaraman, and C. Scheideler. On local algorithms for topology control and routing in ad hoc networks. In *ACM SPAA*, pages 220–229, June 2003.
- 30 R. M. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized rumor spreading. In *FOCS*, pages 565–574, 2000.
- 31 D. Kempe and J. Kleinberg. Protocols and impossibility results for gossip-based communication mechanisms. In *IEEE FOCS*, 2002.
- 32 F. Kuhn and R. Oshman. Dynamic networks: Models and algorithms. *SIGACT News*, 42(1):82–96, 2011.
- 33 Fabian Kuhn, Nancy Lynch, and Rotem Oshman. Distributed computation in dynamic networks. In *ACM STOC*, pages 513–522, 2010.
- 34 F. T. Leighton. *Introduction to Parallel Algorithms and Architectures: Arrays, Trees, and Hypercubes*. Morgan-Kaufmann, San Mateo, CA, 1991.
- 35 N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann, San Francisco, CA, 1996.
- 36 Damon Mosk-Aoyama and Devavrat Shah. Computing separable functions via gossip. In *PODC*, pages 113–122, 2006.
- 37 Peter Sanders, Sebastian Egner, and Ludo Tolluizen. Polynomial time algorithms for network information flow. In *ACM SPAA*, pages 286–294, 2003.
- 38 Donald M. Topkis. Concurrent broadcast for information dissemination. *IEEE Trans. Softw. Eng.*, 11:1107–1112, October 1985.
- 39 L. Zosin and S. Khuller. On directed Steiner trees. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms*, January 2002.


A Proofs for the offline model

► **Lemma 10.** *Let there be k tokens, each with a source node where it is present in the beginning and a set of destination nodes to whom we want to send it. It is feasible to send all the tokens to all of their destination nodes in a dynamic network using l rounds, where in each round a node can broadcast only one token to all its neighbors, if and only if k directed Steiner trees can be packed in the corresponding evolution graph with $2l + 1$ levels respecting the edge capacities, one for each token with its root being the copy of the source node at level 0 and its terminals being the copies of the destination nodes at level $2l$.*

Proof. Assume that k tokens can be sent to all of their destinations in l rounds and fix one broadcast schedule that achieves this. We will construct k directed Steiner trees as required by the lemma based on how the tokens reach their destinations and then argue that they all can be packed in the evolution graph respecting the edge capacities. For a token i , we construct a Steiner tree T^i as follows. For each level $j \in \{0, \dots, 2l\}$, we define a set S_j^i of nodes at level j inductively starting from level $2l$ backwards. S_{2l}^i is simply the copies of the destination nodes for token i at level $2l$. Once $S_{2(j+1)}^i$ is defined, we define S_{2j}^i (respectively S_{2j+1}^i) as: for each $v_{2(j+1)} \in S_{2(j+1)}^i$, include v_{2j} (respectively nothing) if token i has reached node v after round j , or include a node u_{2j} (respectively u_{2j+1}) such that u has token i at the end of round j which it broadcasts in round $j + 1$ and (u, v) is an edge of G_{j+1} . Such a node u can always be found because whenever v_{2j} is included in S_{2j}^i , node v has token i by the end of round j which can be proved by backward induction starting from $j = l$. It is easy to see that S_0^i simply consists of the copy of the source node of token i at level 0. T^i is constructed on the nodes in $\bigcup_{j=0}^{j=2l} S_j^i$. If for a vertex v , $v_{2(j+1)} \in S_{2(j+1)}^i$ and $v_{2j} \in S_{2j}^i$, we add the buffer edge $(v_{2j}, v_{2(j+1)})$ in T^i . Otherwise, if $v_{2(j+1)} \in S_{2(j+1)}^i$ but $v_{2j} \notin S_{2j}^i$, we add the selection edge (u_{2j}, u_{2j+1}) and broadcast edge $(u_{2j+1}, v_{2(j+1)})$ in T^i , where u was the node chosen as described above. It is straightforward to see that these edges form a directed Steiner tree for token i as required by the lemma which can be packed in the evolution graph. The argument is completed by noting that any unit capacity edge cannot be included in two different Steiner trees as we started with a broadcast schedule where each node broadcasts a single token to all its neighbors in one round, and thus all the k Steiner trees can be simultaneously packed in the evolution graph respecting the edge capacities.

Next assume that k Steiner trees as in the lemma can be packed in the evolution graph respecting the edge capacities. We construct a broadcast schedule for each token from its Steiner tree in the natural way: whenever the Steiner tree T_i corresponding to token i uses a broadcast edge (u_{2j-1}, v_{2j}) for some j , we let the node u broadcast token i in round j . We need to show that this is a feasible broadcast schedule. First we observe that two different Steiner trees cannot use two broadcast edges starting from the same node because every selection edge has unit capacity, thus there are no conflicts in the schedule and each node is asked to broadcast at most one token in each round. Next we claim by induction that if node v_{2j} is in T^i , then node v has token i by the end of round j . For $j = 0$, it is trivial since only the copy of the source node for token i can be included in T^i from level 0. For $j > 0$, if v_{2j} is in T^i , we must reach there by following the buffer edge $(v_{2(j-1)}, v_{2j})$ or a broadcast edge (u_{2j-1}, v_{2j}) . In the former case, by induction node v has token i after round $j - 1$ itself. In the latter case, node u which had token i after round $j - 1$ by induction was the neighbor of node v in G_j and u broadcast token i in round j , thus implying node v has token i after round j . From the above claim, we conclude that whenever a node is asked to broadcast a token in round j , it has the token by the end of round $j - 1$. Thus the schedule we constructed is a feasible broadcast schedule. Since the copies of all the destination nodes

of a token at level $2l$ are the terminals of its Steiner tree, we conclude all the tokens reach all of their destination nodes after round l . ◀



■ **Figure 3** An example of building directed Steiner tree in the evolution graph G based on token dissemination process. Token t starts from node B . Thus, the Steiner tree is rooted at B_0 in G . Since B_0 has token t , we include the infinite capacity buffer edge (B_0, B_2) . In the first round, node B broadcasts token t , and hence we include the selection edge (B_0, B_1) . Nodes A and C receive token t from B in the first round, so we include edges (B_1, A_2) , (B_1, C_2) . Now A_2 , B_2 , and C_2 all have token t . Therefore we include the edges (A_2, A_4) , (B_2, B_4) , and (C_2, C_4) . In the second round, all of A , B , and C broadcast token t , we include edges (A_2, A_3) , (B_2, B_3) , (C_2, C_3) . Nodes D and E receive token t from C . So we include edges (C_3, D_4) and (C_3, E_4) . Notice that nodes A and B also receive token t from C , but they already have token t . Thus, we don't include edges (C_3, B_4) or (C_3, A_4) .

Proof of Lemma 5: By lemma 10, we will be done in $n + k$ rounds if we can show that k paths, one from every source vertex at level 0 to $v_{2(n+k)}$, can be packed in the corresponding evolution graph with $2(n + k) + 1$ levels respecting the edge capacities. For this, we consider the evolution graph and add to it a special vertex v_{-1} at level -1 and connect it to every source at level 0 by an edge of capacity 1. (Multiple edges get fused with corresponding increase in capacity if multiple tokens have the same source.) We claim that the value of the min-cut between v_{-1} and $v_{2(n+k)}$ is at least k . Before proving this, we complete the proof of the claim assuming this. By the max flow min cut theorem, the max flow between v_{-1} and $v_{2(n+k)}$ is at least k . Since we connected v_{-1} with each of the k token sources at level 0 by a unit capacity edge, it follows that unit flow can be routed from each of these sources at level 0 to $v_{2(n+k)}$ respecting the edge capacities. It is easy to see that this implies we can pack k paths, one from every source vertex at level 0 to $v_{2(n+k)}$, respecting the edge capacities.

To prove our claimed bound on the min cut, consider any cut of the evolution graph separating v_{-1} from $v_{2(n+k)}$ and let S be the set of the cut containing v_{-1} . If S includes no vertex from level 0, we are immediately done. Otherwise, observe that if $v_{2j} \in S$ for some $0 \leq j < (n + k)$ and $v_{2(j+1)} \notin S$, then the value of the cut is infinite as it cuts the buffer edge of infinite capacity out of v_{2j} . Thus we may assume that if $v_{2j} \in S$, then $v_{2(j+1)} \in S$. Also observe that since each of the communication graphs G_1, \dots, G_{n+k} are connected, if

the number of vertices in S from level $2(j+1)$ is no more than the number of vertices from level $2j$ and not all vertices from level $2(j+1)$ are in S , we get at least a contribution of 1 in the value of the cut. But since the total number of nodes is n and $v_{2(n+k)} \notin S$, there must be at least k such levels, which proves the claim. \blacktriangleleft

Proof of Theorem 6: It is trivial to see that if $k \leq \sqrt{n \log n}$, then the algorithm will end in nk rounds and each node receives all the k tokens. Assume $k > \sqrt{n \log n}$. By Lemma 5, all the tokens can be sent to all the nodes in S using $O(n^{1.5} \sqrt{\log n})$ rounds. Now fix a node v and a token t . Since token t is broadcast for $\sqrt{n \log n}$ rounds, there is a set S_v^t of at least $\sqrt{n \log n}$ nodes from which v is reachable within those rounds. It is clear that if S intersects S_v^t , v will receive token t . Since the set S was picked uniformly at random, the probability that S does not intersect S_v^t is at most

$$\frac{\binom{n-2\sqrt{n \log n}}{2\sqrt{n \log n}}}{\binom{n}{2\sqrt{n \log n}}} < \left(\frac{n-2\sqrt{n \log n}}{n} \right)^{2\sqrt{n \log n}} = o\left(\frac{1}{n^2}\right).$$

Thus every node receives every token with probability $1 - o(1)$. It is also clear that the algorithm finishes in $O(n^{1.5} \sqrt{\log n})$ rounds. \blacktriangleleft

Proof of Theorem 9: We show the following three claims: (i) In Step 7, $|\mathcal{S}| \geq k$ with probability at least $1 - 1/e^{k/4}$. This is the correctness of Algorithm 2, saying it can find the schedule to disseminate all k tokens. (ii) The number of rounds in the schedule produced by Algorithm 2 is at most $O(n^\epsilon)$ times the optimal one. (iii) In the token dissemination schedule, the number of tokens sent over an edge is $O(\log n)$ in any round with high probability.

First, we prove claim (i). Let X_i denote the sum of non-zero x_T^* 's in iteration i . $X = \sum_{i=1}^{2n^\epsilon} X_i$. We know $\mathbb{E}[X_i] = k/n^\epsilon$. Thus, $\mathbb{E}[X] = 2n^\epsilon k/n^\epsilon = 2k$, which is the expected number of Steiner trees in set \mathcal{S} . By Chernoff bound, we have

$$\Pr[X \leq k] = \Pr\left[X \leq \left(1 - \frac{1}{2}\right) \mathbb{E}[X]\right] \leq e^{-\frac{(1/2)^2 \mathbb{E}[X]}{2}} = e^{-\frac{(1/2)^2 \cdot 2k}{2}} = \frac{1}{e^{k/4}}$$

Thus, $|\mathcal{S}| \geq k$ with probability at least $1 - 1/e^{k/4}$ in Step 7.

Next we prove claim (ii). Let L denote the number of rounds needed by an optimal algorithm. Since in Step 3 we used the $O(n^\epsilon)$ -approximation algorithm in [21] to solve $F(\mathcal{P})$, we know $L^* \leq L$. There are $2n^\epsilon$ iterations. Thus, the number of rounds needed by Algorithm 2 is at most $2n^\epsilon L^* \leq 2n^\epsilon L$, which is an $O(n^\epsilon)$ -approximation on the number of rounds.

Lastly we prove claim (iii). When Algorithm 2 does randomized rounding in Step 4, some constraint $\sum_{T:e \in T} x_T \leq c_e$ in \mathcal{P} may be violated. In the evolution graph, $c_e = 1$. Let Y denote the sum of x_T^* 's in this constraint. We have $\mathbb{E}[Y] \leq c_e = 1$. By Chernoff bound,

$$\begin{aligned} \Pr[Y \geq \mathbb{E}[Y] + \log n] &= \Pr\left[Y \geq \left(1 + \frac{\log n}{\mathbb{E}[Y]}\right) \mathbb{E}[Y]\right] \\ &\leq e^{-\mathbb{E}[Y] \left[\left(1 + \frac{\log n}{\mathbb{E}[Y]}\right) \ln \left(1 + \frac{\log n}{\mathbb{E}[Y]}\right) - \frac{\log n}{\mathbb{E}[Y]} \right]} \leq \frac{1}{n^{\log \log n}} \end{aligned}$$

Thus, the number of tokens sent over a given edge is $O(\log n)$ with probability at least $1 - 1/n^{\log \log n}$. Since there are only polynomial number of edges, no edge will carry more than $O(\log n)$ tokens in a single round with high probability. \blacktriangleleft