

On the Complexity of Information Spreading in Dynamic Networks

Chinmoy Dutta* Gopal Pandurangan† Rajmohan Rajaraman* Zhifeng Sun*
Emanuele Viola*

Abstract

We study how to spread k tokens of information to every node on an n -node dynamic network, the edges of which are changing at each round. This basic *gossip problem* can be completed in $O(n+k)$ rounds in any static network, and determining its complexity in dynamic networks is central to understanding the algorithmic limits and capabilities of various dynamic network models. Our focus is on token-forwarding algorithms, which do not manipulate tokens in any way other than storing, copying and forwarding them.

We first consider the *strongly adaptive* adversary model where in each round, each node first chooses a token to broadcast to all its neighbors (without knowing who they are), and then an adversary chooses an arbitrary connected communication network for that round with the knowledge of the tokens chosen by each node. We show that $\Omega(nk/\log n + n)$ rounds are needed for any randomized (centralized or distributed) token-forwarding algorithm to disseminate the k tokens, thus resolving an open problem raised in [KLO10]. The bound applies to a wide class of initial token distributions, including those in which each token is held by exactly one node and *well-mixed* ones in which each node has each token independently with a constant probability.

Our result for the strongly adaptive adversary model motivates us to study the *weakly adaptive* adversary model where in each round, the adversary

is required to lay down the network first, and then each node sends a possibly distinct token to each of its neighbors. We propose a simple randomized distributed algorithm where in each round, along every edge (u, v) , a token sampled uniformly at random from the symmetric difference of the sets of tokens held by node u and node v is exchanged. We prove that starting from any well-mixed distribution of tokens where each node has each token independently with a constant probability, this algorithm solves the k -gossip problem in $O((n+k) \log n \log k)$ rounds with high probability over the initial token distribution and the randomness of the protocol. We then show how the above uniform sampling problem can be solved using $\tilde{O}(\log n)$ bits of communication, making the overall algorithm communication-efficient.

We next present a centralized algorithm that solves the gossip problem for every initial distribution in $O((n+k) \log^2 n)$ rounds in the offline setting where the entire sequence of communication networks is known to the algorithm in advance. Finally, we present an $O(n \min\{k, \sqrt{k \log n}\})$ -round centralized offline algorithm in which each node can only broadcast a single token to all of its neighbors in each round.

1 Introduction

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

*College of Computer and Information Science, Northeastern University, Boston, 02115, USA. Email: {chinmoy, rraj, austin, viola}@ccs.neu.edu. Chinmoy Dutta is supported in part by NSF grant CCF-0845003 and a Microsoft grant to Ravi Sundaram; Rajmohan Rajaraman and Zhifeng Sun are supported in part by NSF grant CNS-0915985; Emanuele Viola is supported by NSF grant CCF-0845003.

†Division of Mathematical Sciences, Nanyang Technological University, Singapore 637371 and Department of Computer Science, Brown University, Providence, RI 02912, USA. Email: gopalpandurangan@gmail.com. Supported in part by the following research grants: Nanyang Technological University grant M58110000, Singapore Ministry of Education (MOE) Academic Research Fund (AcRF) Tier 2 grant MOE2010-T2-2-082, and a grant from the US-Israel Binational Science Foundation (BSF).

In this paper, we study a fundamental problem of information spreading, called *k-gossip*, on dynamic networks. This problem was analyzed for static networks by Topkis [Top85], and was first studied on dynamic networks by Kuhn, Lynch, and Oshman [KLO10]. In *k-gossip* (also referred to as *k-token dissemination*), there are k distinct pieces of information (tokens) that are initially present in some nodes and the problem is to disseminate all the k tokens to all the n nodes in the network, under the bandwidth constraint that one token can go through an edge per round, under a synchronous model of communication. This problem is a fundamental primitive for distributed computing; indeed, solving *n-gossip*, where each node starts with exactly one token, allows any function of the initial states of the nodes to be computed, assuming the nodes know n [KLO10].

The dynamic network models that we consider in this paper allow an adversary to choose an arbitrary set of communication links among the nodes for each round, with the only constraint being that the resulting communication graph is connected in each round. Our adversarial models are either the same as or closely related to those adopted in recent studies [AKL08, KLO10, OW05, CFQS10].

The focus of this paper is on the power of *token-forwarding* algorithms, which do not manipulate tokens in any way other than storing, copying, and forwarding them. Token-forwarding algorithms are simple and easy to implement, typically incur low overhead, and have been widely studied (e.g, see [Lei91b, Pel00]). In any n -node static network, a simple token-forwarding algorithm that pipelines token transmissions up a rooted spanning tree, and then broadcasts them down the tree completes *k-gossip* in $O(n+k)$ rounds [Top85, Pel00], which is tight since $\Omega(n+k)$ rounds is a straightforward lower bound due to bandwidth constraints. The central question motivating our study is whether a linear or near-linear bound is achievable for *k-gossip* on dynamic networks.

1.1 Our results Our first result, in Section 2, is a lower bound for *k-gossip* under a worst-case

model due to [KLO10], which we call the *strongly adaptive adversary* model. We now define the model and then state the theorem.

DEFINITION 1.1. (STRONGLY ADAPTIVE ADV.) *In each round of the strongly adaptive adversary model, each node first chooses a token to broadcast to all its neighbors (without knowing who they are), and then the adversary chooses an arbitrary connected communication network for that round with the knowledge of the tokens chosen by each node.*

We note that the choice made by each node may depend arbitrarily on the tokens held by that and other nodes. Hence this model allows for both distributed and centralized algorithms.

THEOREM 1.1. (a) *Any randomized token-forwarding algorithm (centralized or distributed) for *k-gossip* needs $\Omega(nk/\log n + n)$ rounds in the strongly adaptive adversary model starting from any initial token distribution in which each of $k \leq n$ tokens is held by exactly one node.* (b) *In addition, the same bound holds with high probability over an initial token distribution where each of the n nodes receives each of $k \leq n$ tokens independently with probability $3/4$.*

This result resolves an open problem raised in [KLO10], improving their lower bound of $\Omega(n \log n)$ for $k = \omega(\log n \log \log n)$, 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 [Hae11, HK11]. Assuming the size of each message is bounded by the size of a token, network coding completes *k-gossip* in $O(nk/\log n + n)$ rounds for $O(\log n)$ -bit tokens, and $O(n+k)$ rounds for $\Omega(n \log n)$ bit tokens. Thus, for large token sizes, our result *establishes a factor $\Omega(\min\{n, k\}/\log n)$ gap between token-forwarding and network coding*, a significant new bound on the network coding advantage for information dissemination.¹ Furthermore, for small token and message sizes (e.g.,

¹The strongly adaptive adversary model allows each node to broadcast one token in each round, and thus our bounds hold regardless of the token size.

$O(\text{polylog}(n))$ bits), we do not know of any algorithm (network coding, or otherwise) that completes k -gossip against a strongly adaptive adversary in $o(nk/\text{polylog}(n))$ rounds.

Our lower bound for the strongly adaptive adversary model motivates us to study models which restrict the power of the adversary and/or strengthen the capabilities of the algorithm. We would like to restrict the adversary power as little as possible and yet design fast algorithms.

DEFINITION 1.2. (WEAKLY ADAPTIVE ADV.)
In each round of the weakly adaptive adversary model, the adversary is required to lay down the communication network first, before the nodes can communicate. Hence nodes get to know their neighbors and thus each node can send a possibly distinct token to each of its neighbors. Note that the adversary still has full control of the topology in each round.

We propose a simple protocol which we call the *symmetric difference* (SYM-DIFF) protocol.

DEFINITION 1.3. (SYM-DIFF PROTOCOL)
The protocol SYM-DIFF works as follows: in each round, independently along every edge (u, v) , sample a token t uniformly at random from the symmetric difference (i.e., XOR) of the sets of tokens held by node u and node v at the start of the round. Then the node that holds t sends it to the other node.

Our second main result, in Section 3.1, shows that in the weakly adaptive model, the SYM-DIFF protocol beats the lower bound for mixed starting distribution of Theorem 1.1.

THEOREM 1.2. *Starting from any well-mixed distribution of tokens where each of the n nodes has each of the k tokens independently with a positive constant probability, the SYM-DIFF protocol completes k -gossip in $O((n+k)\log n \log k)$ rounds with high probability. The probability is both over the initial assignment of tokens and the randomness of the protocol.*

A communication-efficient implementation of SYM-DIFF hinges on the communication

complexity of sampling a uniform element from the symmetric difference of two sets. As another technical contribution, we give an explicit, communication-efficient protocol for this task in Section 3.2.

THEOREM 1.3. *Let Alice and Bob have two subsets $A \subseteq [k]$ and $B \subseteq [k]$ respectively. There is an explicit, private-coin protocol to sample a random element from the symmetric difference of the two sets, $A \oplus B := (A \setminus B) \cup (B \setminus A)$, such that the sampled distribution is statistically ϵ -close to the uniform distribution on $A \oplus B$ and the protocol uses $O(\log^{3/2}(k/\epsilon))$ bits of communication.*

A recent improvement on pseudorandom generators for combinatorial rectangles [GMR⁺12] implies an improvement in the communication in Theorem 1.3 to $\tilde{O}(\lg k/\epsilon)$. We also note that for SYM-DIFF to be communication-efficient it is important that we work with symmetric difference as opposed to set difference, which might have looked a natural choice. This is because Theorem 1.3 becomes false if we replace symmetric difference $A \oplus B$ with set difference $A \setminus B$. For the latter, communication $\Omega(k)$ is required, due to the lower bounds for disjointness [KS92, Raz92].

Although we have only been able to establish the efficiency of the SYM-DIFF protocol starting from well-mixed distributions as in Theorem 1.2, we conjecture that in fact SYM-DIFF is efficient starting from any token distribution. A priori, however, it is unclear if there is *any* token-forwarding algorithm that solves k -gossip in $\tilde{O}(n+k)$ rounds 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 result, in Section 4.1, resolves this problem.

DEFINITION 1.4. (OFFLINE ALGORITHM) *An offline algorithm for k -gossip takes as input an initial token distribution and a sequence of nk graphs G_1, \dots, G_{nk} , where G_t represents the communication network in round t . The output of the algorithm is a schedule that specifies, for*

each t , each edge e of G_t , a token (if any) sent along e in each direction. The length of the schedule is the largest t for which a token is sent on any edge in round t .

THEOREM 1.4. *There is a polynomial-time randomized offline algorithm that returns, for every k -gossip instance, a schedule of length $O((n+k)\log^2 n)$ with high probability.*

Like SYM-DIFF, the schedule returned by the above offline algorithm allows each node to send a possibly distinct token to each of its neighbors in each round. However, in some applications, e.g., wireless networks, the preferred mode of communication is broadcast. Hence, we also consider offline broadcast schedules where each node can only broadcast a single token to all of its neighbors in each round and show the following result in Section 4.2.

THEOREM 1.5. *There is a polynomial-time randomized offline algorithm that returns, for every k -gossip instance, a broadcast schedule of length $O(n \min\{k, \sqrt{k \log n}\})$, with high probability.*

1.2 Related work Information spreading (or dissemination) in networks is a fundamental problem in distributed computing and has a rich literature. The problem is generally well-understood on static networks, both for interconnection networks [Lei91a] as well as general networks [Lyn96, Pel00, AW04]. In particular, the k -gossip problem can be solved in $O(n+k)$ rounds on any n -node static network [Top85]. There also have been several papers on broadcasting, multicasting, and related problems in static heterogeneous and wireless networks (e.g., see [ABNLP91, BYGI87, BNGNS00, CMPS09]).

Dynamic networks have been studied extensively over the past three decades. Early studies focused on dynamics that arise when edges or nodes fail. A number of fault models, varying according to extent and nature (e.g., probabilistic vs. worst-case) of faults allowed, and the resulting dynamic networks have been analyzed (e.g., see [AW04, Lyn96]). There have been several studies that constrain the rate at which changes

occur, or assume that the network eventually stabilizes (e.g., see [AAG87, Dol00, GB81]).

There also has been considerable work on general dynamic networks. Early studies in this area include [AGR92, APSPS92], which introduce building blocks for communication protocols on dynamic networks. Another notable work is the local balancing approach of [AL94] for solving routing and multicommodity flow problems on dynamic networks, which has also been applied to multicast, anycast, and broadcast problems on mobile ad hoc networks [ABBS01, ABS03, JRS03]. To address highly unpredictable network dynamics, stronger adversarial models have been studied by [AKL08, OW05, KLO10] and others; see the recent survey of [CFQS10] and the references therein. Unlike prior models on dynamic networks, these models and ours do not assume that the network eventually stops changing; the algorithms are required to work correctly and terminate even in networks that change continually over time. The recent work of [CST12], studies the flooding time of *Markovian* evolving dynamic graphs, a special class of evolving graphs. The survey of [KO11] summarizes recent work on dynamic networks. We also note that our model and the ones we have discussed thus far only allow edge changes from round to round; the recent work of [APRU12] studies a dynamic network model where both nodes and edges can change in each round.

Recent work of [Hae11, HK11] presents information spreading algorithms based on network coding [ACLY00]. As mentioned earlier, one of their important results is that the k -gossip problem on the adversarial model of [KLO10] 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 [Hae11, HK11, ABCHL11, BAL10, DMC06, MAS06] and the references therein.

As we show in Section 4.2, the problem of finding an optimal broadcast schedule in the offline setting reduces to the Steiner tree packing problem for directed graphs [CS06]. This problem is closely related to the directed Steiner tree problem (a major open problem in approx-

imation algorithms) [CCC⁺98, ZK02] and the gap between network coding and flow-based solutions for multicast in arbitrary directed networks [AC04, SET03].

Finally, we note that a number of recent studies solve k -gossip and related problems using *gossip-based* processes, in which each node exchanges information with a small number of randomly chosen neighbors in each round, e.g., see [BCEG10, DGH⁺87, KK02, CP12, KSSV00, MAS06, BGPS06] and the references therein. All these studies assume a static communication network, and do not apply directly to the models considered in this paper.

2 Lower bound for the strongly adaptive adversary model

In this section, we prove Theorem 1.1. We first define the adversary used in the proof of Theorem 1.1.

Adversary: The strategy of the adversary is simple. We use the notion of *free edge* introduced in [KLO10]. In a given round r , we call an edge (u, v) *free* if at the start of the round, u has the token that v broadcasts in the round and v has the token that u broadcasts in the round; 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 a strong adversary model, at the start of each round, the adversary knows for each node v , the token 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 will gain at most one new token each through (u, v) . We refer to this exchange on a non-free

edge as a *useful token exchange*.

Our proof proceeds as follows. First, we show that with high probability over the initial assignment of tokens, in every round there are at most $O(\lg n)$ useful token exchanges. Then we note that, again with high probability over the initial assignment of tokens, overall $\Omega(nk)$ useful token exchanges must occur for the protocol to complete.

DEFINITION 2.1. *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 .*

LEMMA 2.1. *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. 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 half-empty configuration is that, in a




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.

given round, it only depends on the tokens held by the nodes; it is independent of the tokens that the nodes broadcast. This allows us to prove the following easy lemma that shows a monotonicity property of half-empty configurations.

LEMMA 2.2. (MONOTONICITY PROPERTY) *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 $r' \leq r$. Hence, the size of the largest half-empty configuration cannot increase with the increase in the number of rounds.*

Proof. The lemma follows by noting 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 2.1 and 2.2 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 a well-mixed distribution satisfies the desired property, establishing part (b) of the theorem.

Proof. [Proof of Theorem 1.1(b)] We first note that if the number of tokens k is less than $100 \log n$, then the $\Omega(n + nk/\log n)$ lower bound is trivially true because even to disseminate one token on a line it takes $\Omega(n)$ rounds². Thus, in the following proof, we focus on the case where $k \geq 100 \log n$.

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.

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

In the above inequality, $\binom{n}{l}$ is the number of ways of choosing the l nodes that form the half-empty configuration, $k!/(k-l)!$ is the number of ways of assigning l distinct tokens, and $(1/2)^{\binom{l}{2}}$ is the upper bound on the probability for each pair $i \neq j$ that either v_i is missing t_j or v_j is missing t_i . For $l \geq 5 \log n$, $\Pr[E_l] \leq 1/n^2$. Thus, the largest half-empty configuration at the start of the first round, and hence at the start of *any* round (by Lemma 2.2), is of size at most $5 \log n$ with probability at least $1 - 1/n^2$. By Lemma 2.1, 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 missing at node i in the initial distribution. Then M_i is a binomial random variable with $\mathbb{E}[M_i] = k/4$. By a Chernoff bound, the probability that node i misses at most $k/8$ tokens is

$$\begin{aligned} \Pr\left[M_i \leq \frac{k}{8}\right] &= \Pr\left[M_i \leq \left(1 - \frac{1}{2}\right) \cdot \mathbb{E}[M_i]\right] \\ &\leq e^{-\frac{\mathbb{E}[M_i] \left(\frac{1}{2}\right)^2}{2}} = e^{-\frac{k}{32}}. \end{aligned}$$

Thus, the total number of tokens missing in the initial distribution is at least $n \cdot k/8 = \Omega(kn)$ with

²The choice of the constant 100 here is arbitrary; we have not optimized the choice of constants in the proof.

probability at least $1 - n/e^{\frac{k}{32}} \geq 1 - 1/n^2$ ($k \geq 100 \log n$). 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.

Part (b) of Theorem 1.1 does not apply to some natural initial distributions, such as one in which each token resides at exactly one node. When starting from a distribution in this class, though there are far fewer tokens distributed initially, the argument above does not rule out the possibility that an algorithm avoids the problematic configurations that arise in the proof. Part (a) of Theorem 1.1 extends the lower bound to this class of distributions. The main idea of the proof is showing that a reduction exists (via the probabilistic method) to an initial well-mixed distribution of Theorem 1.1.

LEMMA 2.3. *From any distribution in which each token starts at exactly one node and no node has more than one token, any online token-forwarding algorithm for k -gossip needs $\Omega(kn/\log n)$ rounds against a strong adversary.*

Proof. We consider an initial distribution C where each token is at exactly one node, and no node has more than one token. Let C^* be an initial token distribution in which each node has each token independently with probability $3/4$. By Theorem 1.1, any online algorithm starting from distribution C^* needs $\Omega(kn/\log n)$ rounds with high probability.

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 first show, using Hall's Theorem, that this bipartite graph has a perfect matching with very high probability. 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

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

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 that the neighborhood size of any set of m nodes from V_2 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.

By applying the union bound, we obtain that with positive probability (in fact, high probability), C^* takes $\Omega(kn/\log n)$ rounds and there is a perfect matching M in the above bipartite graph. By the probabilistic method, thus both C^* and M exist. Given such C^* and M , we 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 1.1, which applies to A^* and C^* , also applies to A and C .

Proof. [Proof of Theorem 1.1(a)] We extend our proof in Lemma 2.3 to the initial distribution C where each token starts at exactly one node, but nodes may have multiple tokens. We consider the following two cases.

The first case is when at least $n/2$ nodes start with some token. This implies that $k \geq n/2$. Let us focus on the $n/2$ nodes with tokens. Each of them has at least one unique token. By the same argument used in Lemma 2.3, disseminating these $n/2$ distinct tokens to n nodes takes $\Omega(n^2/\log n)$ rounds. Thus, in this case the number of rounds needed is $\Omega(kn/\log n)$.

The second case is when less than $n/2$ nodes start with some token. In this case, the adversary can group these nodes together, and treat

them as one super node. There is only one edge connecting this super node to the rest of the nodes. Thus, the number of useful token exchanges provided by this super node is at most one in each round. If there exists an algorithm that can disseminate k tokens in $o(kn/\log n)$ rounds, then the contribution by the super node is $o(kn/\log n)$. And by the same argument used in Lemma 2.3 we know dissemination of k tokens to $n/2$ nodes (those start with no tokens) takes $\Omega(kn/\log n)$ rounds. Thus, the theorem also holds in this case.

3 Upper bound in the weakly adaptive adversary model

In this section, we first analyze the SYM-DIFF protocol starting from a well-mixed distribution of tokens and prove Theorem 1.2 (presented in Section 3.1), and then show how to sample an element from the symmetric difference of two sets efficiently in the two-player communication complexity model (presented in Section 3.2). However, before doing that, we present the following lower bound that shows randomization is crucial for the SYM-DIFF protocol.

THEOREM 3.1. *Consider the protocol DET-SYM-DIFF for k -gossip in the weakly adaptive adversary model which is identical to the SYM-DIFF protocol except for, in each round, the token sent along each edge (u, v) is chosen deterministically from the symmetric difference of the set of tokens held by node u and the set of tokens held by node v . Starting from an initial token distribution where one node has all the k tokens and others have none, a strongly adaptive adversary can force $\Omega(nk)$ rounds for the DET-SYM-DIFF protocol to disseminate the k tokens to the n nodes.*

Proof. Let the node u start with all the tokens and nodes v_1, \dots, v_{n-1} start with no tokens. The adversary can connect u, v_1, \dots, v_{n-1} in a line in the first round thereby guaranteeing only node v_1 gets a token, say t_1 . In the next round, the adversary connects $u, v_2, \dots, v_{n-1}, v_1$ in a line. In this round, node v_2 and v_{n-1} will both get

token t_1 . The adversary can continue this way for $\frac{n-2}{2} + 1$ rounds, at which point all the nodes v_1, v_2, \dots, v_{n-1} will have token t_1 . We can repeat this argument for all the k tokens proving the lower bound of $\Omega(nk)$.

3.1 Analysis of SYM-DIFF starting from well-mixed distributions

For the proof of Theorem 1.2, we will assume that we start from the initial token distribution where each node has each token independently with probability $\frac{1}{2}$. It is easy to extend it to any positive constant probability. We need the following definition. We call a maximal set of nodes that holds the same set of tokens at the start of a round r to be a *group* for round r .

LEMMA 3.2. *In a token distribution where each node has each token independently with probability $\frac{1}{2}$, the union of the set of tokens of any ℓ nodes misses at most $\frac{n+k}{\ell}$ tokens with high probability.*

Proof. There are $\binom{n}{\ell}$ ways of choosing ℓ nodes out of n nodes, and $\binom{k}{\frac{n+k}{\ell}}$ ways of choosing $\frac{n+k}{\ell}$ tokens out of k tokens. Thus the probability that the union of the set of tokens of any ℓ nodes misses more than $\frac{n+k}{\ell}$ tokens is at most

$$\binom{n}{\ell} \binom{k}{\frac{n+k}{\ell}} \left(\frac{1}{2}\right)^{n+k},$$

which is inverse polynomial in both n and k .

Since in any round, no token can be exchanged along an edge between two nodes of the same group, we will consider only the edges that connect two nodes from different groups. We call such edges *inter-group* edges for that round. In fact, we will prove the theorem in a stronger sense where we let the adversary orient the inter-group edges to determine the direction of token movement along all these edges, and the token sent along each of these edges is chosen uniformly at random from the symmetric difference conditioned on this orientation. (The adversary must respect the condition that there can be no token movement from a node u to a node v if the set of nodes held by node u is a subset of that held by node v .) We define one unit of progress in a

round as a node receiving a token in that round that it did not have at the start of the round.

LEMMA 3.3. *With high probability, the following holds for every node v and every round i : If v misses $m > \log n$ tokens at the start of round i and it has $d > \log k$ incoming inter-group edges in that round, then node v makes $\Omega(\min\{m, d\})$ units of progress in round i . Here, the probability is over the initial token distribution and the randomness used in the protocol.*

Proof. First we prove the claim that for some sufficiently small constant $\alpha < 1$, with probability $1 - o(1)$, the following holds for every node v and every round i : If v misses $m > \log n$ tokens at the start of round i and it has $d > \log k$ in-neighbors in that round, then αd of these neighbors each have, at the start of round i , αm tokens that node v misses. Let us compute the probability that the claim is not true for some node v in some round i . The d inter-group in-neighbors can be chosen in at most $\binom{n}{d}$ different ways and the m missing tokens can be chosen in at most $\binom{k}{m}$ different ways. There are at most $\binom{d}{\alpha d}$ ways of choosing the in-neighbors that do not have the claimed number of missing tokens, and for each of them there are at most $\binom{m}{\alpha m}$ ways of choosing which of these tokens they miss. Thus the probability of failure is at most

$$\binom{n}{d} \binom{k}{m} \binom{d}{\alpha d} \binom{m}{\alpha m}^{(1-\alpha)d} \left(\frac{1}{2}\right)^{(1-\alpha)^2 md},$$

which is $o\left(\frac{1}{(nk)^2}\right)$ since $m > \log n$ and $d > \log k$ and α is chosen sufficiently small. Noting that there are at most n choices for d and at most k choices for m , the claim follows. From the above claim, the lemma follows by standard calculations.

Proof. [Proof of Theorem 1.2] We color each of the rounds *red*, *blue*, *green* or *black*. If in a round, there is a node v that misses less than $\log n$ tokens and makes at least one unit of progress in that round, we color the round red. If a round is not colored red, and there is a node that gets a constant fraction of its missing tokens in that round (the same fraction as in Lemma 3.3), we

color it green. If a round is neither colored red nor colored green, we color the round blue.

It is immediate that there can be at most $n \log n$ red rounds since each of the n nodes can be responsible for coloring at most $\log n$ rounds red. Similarly, there can be at most $O(n \log k)$ green rounds since each node can be responsible for coloring at most $O(\log k)$ rounds green. Now let us turn to the blue rounds. Fix a blue round and let there be r groups in that round. Using Lemma 3.2, we infer that there are at most $(n + k)r$ tokens missing in total at the start of this round. We also note that there must be at least $r - 1$ inter-group edges in this round and combining this with Lemma 3.3 and the fact that this round was not colored red or green, we infer that we make $\Omega\left(\frac{r}{\log k}\right)$ units of progress in this round.

We can label each blue round by the smallest number of groups in a blue round seen so far. The sequence of labels is non-increasing and let us say it starts from $s \leq n$. We divide the blue rounds in partitions where the i 'th partition contain those with labels in $[s/2^{i-1}, s/2^i)$. There are at most $\log n$ partitions. From the above argument, we see that there can be at most $O((n + k) \log k)$ blue rounds in each partition, which implies a bound of $O((n + k) \log n \log k)$ for the total number of blue rounds. This completes the proof of the theorem.

3.2 Uniform sampling from symmetric difference

We now restate and prove our result on a communication-efficient protocol to sample from the symmetric difference of two sets.

THEOREM 1.3. *Let Alice and Bob have two subsets $A \subseteq [k]$ and $B \subseteq [k]$ respectively. There is an explicit, private-coin protocol to sample a random element from the symmetric difference of the two sets, $A \oplus B := (A \setminus B) \cup (B \setminus A)$, such that the sampled distribution is statistically ϵ -close to the uniform distribution on $A \oplus B$ and the protocol uses $O(\log^{3/2}(k/\epsilon))$ bits of communication.*

We now explain how we obtain a communication-efficient protocol to sample

from the symmetric difference $A \oplus B$ of two sets $A, B \subseteq [k]$, proving Theorem 1.3.

Our starting point is Nisan and Safra's protocol [Nis93] to determine the least i such that $i \in A \oplus B$. (In [Nis93] the protocol is phrased as deciding if $A > B$, when A and B are viewed as k -bit integers. It is easy to switch between the two.) For uniform sampling from $A \oplus B$, our idea is to first let the parties permute their sets according to a random permutation σ , then run Nisan and Safra's protocol. This results in an explicit protocol for uniform generation from $A \oplus B$ with communication $O(\log k/\epsilon)$ that uses *public coins*. A standard transformation to private coins via [New91] results in a protocol that is not explicit.

To obtain an explicit, private-coin protocol we derandomize the space of random permutations σ . The key idea is that it is sufficient to have a distribution on permutations σ such that, for any set $D = A \oplus B$, any element in D has roughly the same probability of *being the first element in D to appear in the sequence $\sigma(1), \sigma(2), \sigma(3), \dots$* . We then construct such a space of permutations with seed length $O(\lg^{3/2}(k/\epsilon))$ using Lu's pseudorandom generator for combinatorial rectangles [Lu02] (cf. [Nis92, NZ96, INW94, EGL⁺98, ASWZ96, Lu02, Vio11]). Plugging a better pseudorandom generator for combinatorial rectangles in our argument would result in a protocol for uniform sampling from $A \oplus B$ with communication $\tilde{O}(\log k/\epsilon)$ and error ϵ .

As a first step, we have the following simple derandomization of Nisan and Safra's protocol [Nis93], essentially from [Vio13].

LEMMA 3.4. *There is an explicit, private-coin protocol to determine the least $i \in A \oplus B$, where $A, B \subseteq [k]$, with error α and communication $O(\lg(k/\alpha) \lg \lg k) = \tilde{O}(\lg k/\alpha)$.*

Proof. [Proof sketch] Nisan and Safra's protocol amounts to walking for $O(\lg k/\alpha)$ on a certain binary tree. At every node, the two parties just need to determine with error probability, say, $1/100$ if a portion of their inputs are different. This latter task can be achieved using small-bias

generators with public randomness $O(\lg k)$ and communication $O(1)$. [NN93, AGHP92]

The resulting protocol can be seen as a randomized algorithm needing a one-way stream of $R := O(\lg k/\alpha) \lg k$ random bits and using space $S := O(\lg k/\alpha)$ to store the current node.

Nisan's space-bounded generator [Nis92] can reduce the randomness to $S \lg(R/S) = \lg(k/\alpha) \lg \lg k$ with error loss $2^{-S} = \alpha/k$.

The parties start by exchanging a seed for Nisan's generator, and then proceed with the previous protocol.

Specifically, for given k and ϵ as in Theorem 1.3 we set $d = k \log(\frac{3k}{\epsilon})$ and $\alpha := \epsilon/3kd$. Alice then picks a random seed of length $s(k, d, \alpha)$ for a generator that fools every combinatorial rectangle with universe size k and d dimensions with error α . That is, if X is the output of the generator on a random seed, we have, for every set $R := R_1 \times R_2 \times \dots \times R_d \subseteq [k]^d$,

$$|\Pr[X \in R] - |R|/k^d| \leq \alpha.$$

Alice sends the seed to Bob.

Both Alice and Bob expand the seed into a sample X of the generator, and use X to generate a permutation σ as follows. Let the number of distinct elements of $[k]$ that appear in X be t . The permutation σ is constructed by defining $\sigma(i)$ to be the i 'th distinct element of $[k]$ that appears in X as we scan it from the beginning, for $i \leq t$. For every $i > t$, $\sigma(i)$ is defined to be a distinct element not appearing in X in an arbitrary but deterministic way that is fixed before the start of the protocol and both Alice and Bob are aware of it. (For concreteness, it can simply be to assign the elements not appearing in X by order).

To show the correctness of our protocol we need the following lemma.

LEMMA 3.5. *Let $X \in [k]^d$ be the output of a combinatorial rectangle generator with error $\alpha = \epsilon/3kd$, over a uniform seed. Let D be any set, and let j be any element in D . The probability that j appears in a coordinate of X before any other element of D is $\geq \frac{1}{|D|} - \frac{2\epsilon}{3k}$.*

Proof. We note that the desired probability is the union of disjoint rectangles, and then apply the property of the generator:

$$\begin{aligned}
& \Pr \left[X \in \bigcup_{0 \leq t < d} ([k] \setminus D)^t \times \{j\} \times [k]^{d-t-1} \right] \\
&= \sum_{0 \leq t < d} \Pr \left[X \in ([k] \setminus D)^t \times \{j\} \times [k]^{d-t-1} \right] \\
&\geq \sum_{0 \leq t < d} |([k] \setminus D)^t \times \{j\} \times [k]^{d-t-1}| / k^d - \frac{\epsilon}{3k} \\
&= \frac{1}{k} + \left(\frac{k - |D|}{k} \right) \frac{1}{k} + \dots + \left(\frac{k - |D|}{k} \right)^{d-1} \frac{1}{k} \\
&\quad - \frac{\epsilon}{3k} \\
&= \frac{1}{k} \left(1 + \left(1 - \frac{|D|}{k} \right) + \dots + \left(1 - \frac{|D|}{k} \right)^{d-1} \right) \\
&\quad - \frac{\epsilon}{3k} \\
&= \frac{1}{|D|} \left(1 - \left(1 - \frac{|D|}{k} \right)^d \right) - \frac{\epsilon}{3k} \\
&\geq \frac{1}{|D|} \left(1 - e^{-\frac{|D|}{k} k \log(\frac{3k}{\epsilon})} \right) - \frac{\epsilon}{3k} \\
&= \frac{1}{|D|} - \frac{1}{|D|} \left(\frac{\epsilon}{3k} \right)^{|D|} - \frac{\epsilon}{3k} \\
&\geq \frac{1}{|D|} - \frac{2\epsilon}{3k},
\end{aligned}$$

since $|D| \geq 1$.

Now we can complete the proof of Theorem 1.3.

Proof. [Proof of Theorem 1.3] For given k, ϵ , we set $d = k \log(\frac{3k}{\epsilon})$ and $\alpha := \epsilon/3kd$. Alice then picks a random seed of length $s(k, d, \alpha)$.

If σ is chosen such that every element $j \in D$ has probability $\frac{1}{|D|}$ of preceding all other elements of D , then $\sigma(i^*)$ is a uniform random element of D , where i^* is the first position where the permuted A and B differ. Using Lemma 3.5, we immediately see that if σ is chosen as in the first step of the protocol, then the distribution of $\sigma(i^*)$ is at most $(\frac{2\epsilon}{3k})|D| \leq \frac{2\epsilon}{3}$ -far from the uniform distribution on D .

For the second part of the protocol we use Lemma 3.4 with $\alpha := \epsilon/3$.

Overall, the sampled distribution has distance $\leq 2\epsilon/3 + \epsilon/3 = \epsilon$ from the uniform distribution on D .

Using the generator in [Lu02] we have $s(k, d, \alpha) = O(\lg n + \lg d + \lg^{32} 1/\alpha) = O(\lg^{3/2} n/\epsilon)$. So overall the communication is $O(\lg^{3/2} n/\epsilon)$.

4 Offline token-forwarding algorithms

We present two offline algorithms for k -gossip. The first computes an $O((n+k)\log^2 n)$ -round schedule assuming that each node can send at most one token to each neighbor in each round (Section 4.1); the second computes an $O(\min\{n\sqrt{k\log n}, nk\})$ -round broadcast schedule assuming that each node can broadcast at most one token to its neighbors in each round (Section 4.2).

4.1 $O((n+k)\log^2 n)$ -round offline schedule

In this section, we present an algorithm for computing an $O((n+k)\log^2 n)$ round offline schedule. Our bound is tight to within an $O(\log^2 n)$ factor since the dissemination of any k tokens to even a single node of the network requires $\Omega(n+k)$ rounds in the worst case. We begin by defining the notion of an evolution graph that facilitates the design of the offline algorithms.

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 $\widehat{G}[l]$ for this network is a directed capacitated graph G with $l+1$ levels constructed as follows. We create $l+1$ copies of V and call them $V_0, V_1, V_2, \dots, V_l$. 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 $i-1$ corresponds to the beginning of round i and level i corresponds to the end of round i . Level 0 corresponds to the network at the start. There are two kinds of edges in the graph. First, for every node v in V and every round i , we place an edge with infinite capacity from v_{i-1} to v_i . 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. Second, for every round i and every edge $(u, v) \in G_i$, we place

two directed edges with unit capacity each, one from u_{i-1} to v_i and another from v_{i-1} to u_i . We call these edges as *transmit edges* as they correspond to every node transmitting a message to a neighbor in round i ; the unit capacity ensures that in a given round a node can transmit at most one token to each neighbor. Figure 2 illustrates our construction. Lemma 4.1 explains the usefulness of this construction.

LEMMA 4.1. *Let S be a set of source nodes, each with a subset of the k tokens and let T be a subset of sink nodes. Let $\widehat{G}[\ell]$ be an evolution graph over ℓ rounds. Let P denote a set of edge-disjoint paths starting from S and ending at T . If P contains for each sink v and each token i , a distinct path from a source containing i to v , then P yields an ℓ -round schedule for disseminating the k tokens to each node in T .*

Proof. For each sink v , let p_v^i denote the path in P starting at a source containing token i and ending at v . We construct a schedule in the following natural way: for each token i and sink v , p_v^i is the schedule by which i is sent from a source to v . In particular, if (u_t, v_{t+1}) is in p_v^i , then the node u sends token i to v in round t .

We need to show that this is a feasible schedule. First we observe that two different paths in P cannot use the same transmit edge since each such edge has unit capacity. Next we claim by induction that if node v_j is in p_v^i , then node v has token i by the end of round j . For $j = 0$, it is trivial since path p_v^i starts from a source that has token i . For $j > 0$, if v_j is in p_v^i , then the preceding edge is either a buffer edge (v_{j-1}, v_j) or a transmit edge (u_{j-1}, v_j) . 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 sent token i in round j according to p_v^i , thus implying node v has token i after round j . From the above claim, we conclude that whenever a node is asked to transmit a token in round j , it has the token by the end of round $j-1$. Thus the schedule we constructed is feasible. Since k paths terminate at each of the sinks, we conclude all the tokens reach all of the sinks after round ℓ .

Lemma 4.1 provides the foundation for the following randomized algorithm that first gathers all tokens at a random source node and then, in $O(\log n)$ phases, disseminates these tokens to geometrically increasing sets of nodes, until all of the nodes have all tokens.

Algorithm 1 Computing an $O((n+k)\log^2 n)$ -round schedule for k -gossip

Require: A sequence of communication graphs G_1, G_2, \dots

Ensure: Schedule to disseminate k tokens to all nodes

- 1: **Gather:** Send the k tokens to a node v_0 , chosen uniformly at random, in $n+k$ rounds.
 - 2: **for** i from 0 to $\lg n$ (Phase i) **do**
 - 3: Choose a set S_i of 2^i nodes uniformly at random from the collection of all 2^i -size node sets.
 - 4: **Flow:** Send the k tokens to every node in S_i using a maximum flow in an $O((n+k)\log n)$ -round evolution graph from the set $\{v_0\} \cup \bigcup_{j < i} S_j$ of sources to the set S_i of sinks.
-

We first show that the gather step can be completed in $O(n+k)$ rounds.

LEMMA 4.2. *Let k tokens be at given source nodes and v be an arbitrary node. Then, all the tokens can be gathered at v in at most $n+k$ rounds.*

Proof. Following Lemma 4.1, it suffices to show that any evolution graph $\widehat{G}[n+k]$ contains k edge-disjoint paths, each starting from a source node and ending at v . To prove this, we add to $\widehat{G}[n+k]$ 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_{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_{n+k} is at least k . Since we connected v_{-1}

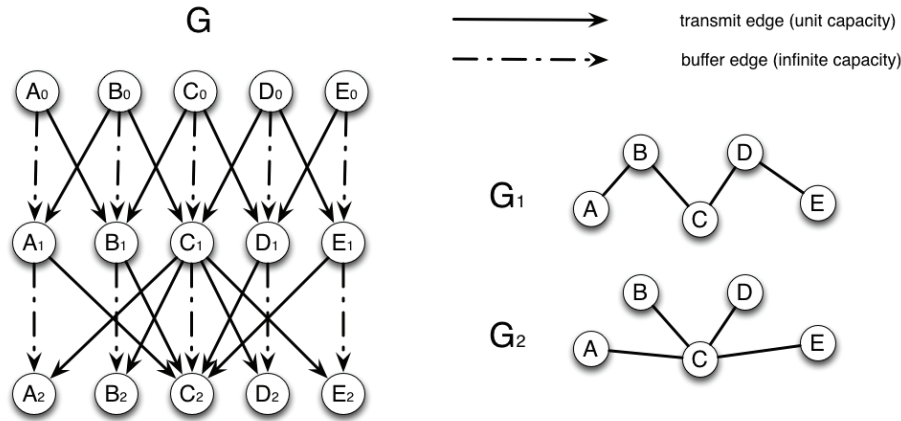


Figure 2: An example of how to construct the evolution graph from a sequence of communication graphs.

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_{n+k} respecting the edge capacities, establishing the desired claim.

To prove our claimed bound on the min cut, consider any cut of the evolution graph separating v_{-1} from v_{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 $v_j \in S$ for some $0 \leq j < (n+k)$ and $v_{j+1} \notin S$, then the value of the cut is infinite as it cuts the buffer edge of infinite capacity out of v_j . Thus we may assume that if $v_j \in S$, then $v_{j+1} \in S$. 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 $j+1$ is no more than the number of vertices from level j and not all vertices from level $j+1$ are in S , we get at least a contribution of 1 in the value of the cut owing to a transmit edge. But since the total number of nodes is n and $v_{n+k} \notin S$, there must be at least k such levels, which proves the claim.

The remainder of the proof concerns the $\lg n$ phases. We first establish an elementary tree decomposition lemma that is critical in showing that there is enough capacity in any $O((n+k) \log n)$ -level evolution graph to complete each

phase.

LEMMA 4.3. *For any n -node tree T and any integer $1 \leq s \leq n$, there exists an edge-disjoint partition of T into subtrees T_1, T_2, \dots such that each T_i has $\Theta(s)$ nodes, every node of T is in some T_i , and for each i , at most one node in T_i is in $\bigcup_{j \neq i} T_j$.*

Proof. The proof is by induction on the size of T . The base case $n = 1$ is trivial. We now consider the induction step. Arbitrarily root the tree T at a node r . For any node v , let T_v denote the subtree rooted at node v ; let $n_v = |T_v|$. Thus, $n_r = n$. Let v denote an arbitrary node such that $n_v \geq s$ and for every child w of v , $n_w < s$. We first consider the case $n_v \leq 2s$. By the induction hypothesis, there exist edge-disjoint subtrees of $T - T_v$ such that each subtree has $\Theta(s)$ edges, every node of $T - T_v$ is in some subtree, and any two subtrees share at most one node. Adding T_v to this collection of subtrees yields the desired claim for T .

We now consider the case where $n_v > 2s$. Here we consider two subcases. The first subcase is where either v is the root or $|T - T_v| \geq s$. We partition the children of v into a set X of groups such that for each group $g \in X$, $s \leq 1 + \sum_{w \in g} n_w \leq 2s$. Let $T(g)$ denote the tree $\{v\} \cup \bigcup_{w \in g} T_w$. All of these subtrees are edge-disjoint and any pair of subtrees share

at most one node (v). If v is the root, then we have established the desired property for T . Otherwise, since $|T - T_v| \geq s$, by the induction hypothesis, there exist edge-disjoint subtrees of $T - T_v$ such that each subtree has $\Theta(s)$ edges, every node of $T - T_v$ is in some subtree, and for any subtree, at most one node in the subtree is in any of the other subtrees. Adding the trees $T(g)$ to this collection of subtrees yields the desired claim for T .

The second subcase is where $0 < |T - T_v| < s$. In this subcase, we make the parent of v as the child of v and proceed to the first subcase, thus establishing the desired claim and completing the induction step.

The set of sources at the start of phase i is $\bar{S}_i = \{v_0\} \cup \bigcup_{j < i} S_j$. We next place a lower bound on the size of \bar{S}_i .

LEMMA 4.4. *For each i , $0 \leq i \leq \lg n$, $|\bar{S}_i|$ is at least $\min\{1, 2^{i-2}\}$ with probability at least $1 - 1/n^3$; furthermore, \bar{S}_i is drawn uniformly at random from the collection of all $|\bar{S}_i|$ -node sets.*

Proof. For $i \leq \lg \lg n$, we calculate the probability, for each v , that there exist more than four values of j for which S_j contains v as at most

$$\binom{\lg n}{5} n \frac{1}{n^5} \leq \frac{1}{n^3}.$$

Thus, the size of the given set is at least $2^i/4 = 2^{i-2}$ with probability at least $1 - 1/n^3$. We now consider the case $i > \lg \lg n$. Let X_v denote the indicator variable for node v to be in the set. Then,

$$\begin{aligned} E[X_v] &= 1 - (1 - 1/n) \prod_{0 \leq j < i} (1 - 2^j/n) \\ &\geq 1 - e^{-1/n - \sum_{j < i} 2^j/n} = 1 - e^{-2^i/n} \\ &\geq 4 \cdot 2^i/(7n). \end{aligned}$$

Thus, the expected size of the set is at least 2^{i-1} . Now, using a Chernoff-type argument (e.g., by using the method of bounded differences and invoking Azuma's inequality), we obtain the size of the set is at least 2^{i-2} whp.

LEMMA 4.5. *Let $r \leq n$ be an arbitrary integer. Let \mathcal{S} denote a set of at least $r/4$ sources and \mathcal{T} a set of r sinks, each set drawn independently and uniformly at random from V . Then, with high probability, the evolution graph $\widehat{G}[\ell]$ with $\ell = \Theta((n+k) \log n)$ contains rk edge-disjoint paths, each path starting from a source and ending at a sink, and each sink having exactly k paths ending at it.*

Proof. We add a super-source having edges of capacity rk to each source and a super-sink with edges of capacity k from each sink. It thus suffices to prove that the maximum flow from the super-source to the super-sink is at least rk . For $r \leq \lg n$, we invoke Lemma 4.2 to obtain that the maximum flow is at least rk . In the remainder of this proof, we assume $r \geq \lg n$. We show that with high probability, the capacity of every cut is at least rk . Note that since there are an exponentially large number of cuts to consider, it may not be sufficient to establish a high probability bound for each cut separately. We address this challenge by identifying an important property that holds for $\widehat{G}[\ell]$ that enables the capacity bound to hold for all cuts simultaneously.

Consider graph G_i with the source and sink sets \mathcal{S} and \mathcal{T} . Recall that \mathcal{S} and \mathcal{T} are drawn uniformly at random from the collection of all $|\mathcal{S}|$ -node and $|\mathcal{T}|$ -node sets, respectively, and \mathcal{T}' is an arbitrary subset of \mathcal{T}' of size r' . By Lemma 4.3 applied to a spanning tree of G_i with parameter $s = (n \log n)/r$, there exist edge-disjoint subtrees T_i^1, T_i^2, \dots , each having $\Theta(s)$ edges from the spanning tree, and together containing all of the nodes in V . Furthermore, for each T_i^j , at most one of its nodes is present in the other subtrees. Since \mathcal{S} and \mathcal{T} are drawn at random and have are of size at least $r/4$ and equal to r , respectively, it follows from a standard Chernoff bound that each of these subtrees has $\Omega(\log n)$ (resp., $\Theta(\log n)$), nodes from \mathcal{S} (resp., \mathcal{T}) whp. In the remainder of the proof, we thus assume that the preceding property holds for each of the graphs in the $\Theta((n+k) \log n)$ levels of $\widehat{G}[\ell]$.

We now argue that every cut $C = (\mathcal{S}, \mathcal{T})$ of

$\widehat{G}[\ell]$ has capacity at least rk . If any of the sources in \mathcal{S} is separated from the super-source, then the capacity of the cut is at least rk since the capacity of the edge connecting the super-source to any source is rk . So in the remainder, we assume that all nodes in S are on the same side of the cut as the super-source. Let \mathcal{T}' denote the set of sinks that are separated from the super-source in C ; let $r' = |\mathcal{T}'|$. All of the edges from $\mathcal{T} - \mathcal{T}'$ to the super-sink cross C and have a total capacity of $(r - r')k$. It thus remains to show that the total capacity of the edges crossing the cut in the intermediate levels 1 through t is at least $r'k$.

Let V_i denote the set of nodes in level i that are in \mathcal{S} . Since every parallel edge has infinite capacity, we have $V_{i+1} \supseteq V_i$. Since each V_i is of size at most n , there are at least $t - n$ levels such that $V_{i+1} = V_i$. For any such level i , C includes all edges that separate \mathcal{S} from \mathcal{T}' in the graph G_i . By the property established above, there exist edge-disjoint partition of a spanning tree of G_i that such that each tree in the partition contains $\Theta(\log n)$ nodes from both \mathcal{S} and \mathcal{T} . Therefore, for any arbitrary subset \mathcal{T}' of size r' , we can find $\Omega(r'/\log n)$ edges that separate \mathcal{T}' from \mathcal{S} . For the number of levels exceeding $\Omega(k \log n)$, it then follows that the total capacity of the edges crossing the cut in the intermediate levels is at least $r'k$. This establishes the desired lower bound on the capacity of the cut, completing the proof of the lemma.

THEOREM 1.4. *There is a polynomial-time randomized offline algorithm that returns, for every k -gossip instance, a schedule of length $O((n + k) \log^2 n)$ with high probability.*

Proof. By Lemma 4.2, the gather step completes in $O(n + k)$ rounds. We now argue that each phase completes in $O((n + k) \log n)$ rounds whp. By Lemma 4.4, the number of sources at the start of phase i is at least 2^{i-2} whp. By Lemmas 4.1 and 4.5, the number of rounds needed for phase i is $O((n + k) \log n)$ whp. Since the number of phases is $\lg n$, the statement of the theorem follows.

4.2 An $O(\min\{n\sqrt{k \log n}, nk\})$ -round broadcast schedule We extend the notion of the evolution graph to the broadcast model. The primary difference is the addition of a new level of nodes and edges for every round that enforces the broadcast constraint.

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 $\tilde{G}[2l + 1]$ with $2l + 1$ levels constructed as follows. We create $2l + 1$ copies of V and call them $V_0, V_1, V_2, \dots, V_{2l}$. 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. Lemma 4.6 explains the usefulness of this construction.

LEMMA 4.6. *Let there be k tokens, each with a source and a set of destinations. It is feasible to send all the tokens to all of their destinations using l rounds, where every node broadcasts only*

one token in each round, iff k directed Steiner trees can be packed in $\tilde{G}[2l + 1]$ levels, one for each token with its root being the copy of the source at level 0 and its terminals being the copies of the destinations 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 $\tilde{G}[2l + 1]$ 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 by 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 $\cup_{j=0}^{j=2l} S_j^i$. 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 $\tilde{G}[2l + 1]$. 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 $\tilde{G}[2l + 1]$ respecting the edge capacities.

Next assume that k Steiner trees as in the lemma can be packed in $\tilde{G}[2l + 1]$ 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 .

Our algorithm is given in Algorithm 2 and analyzed in Lemma 4.2 and Theorem 1.5.

LEMMA 4.7. *Let $k \leq n$ tokens be at given source nodes and v be an arbitrary node. Then, all the tokens can be gathered at v in the broadcast model in at most $n + k$ rounds.*

The proof is analogous to that for the multiport model and is omitted.

THEOREM 1.5. *There is a polynomial-time randomized offline algorithm that returns, for every*

Algorithm 2 $O(\min\{n\sqrt{k \log n}, nk\})$ round algorithm in the offline model

Require: A sequence of communication graphs $G_i, i = 1, 2, \dots$

Ensure: Schedule to disseminate k tokens.

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

k-gossip instance, a broadcast schedule of length $O(n \min\{k, \sqrt{k \log n}\})$, with high probability.

Proof. It is trivial to see that if $k \leq \sqrt{\log n}$, then the algorithm will end in nk rounds and each node receives all the k tokens. Assume $k > \sqrt{\log n}$. By Lemma 4.2, all the tokens can be sent to all the nodes in S using $O(n\sqrt{k \log n})$ rounds. Now fix a node v and a token t . Since token t is broadcast for $2n\sqrt{(\log n)/k}$ rounds, there is a set S_v^t of at least $2n\sqrt{(\log n)/k}$ 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

$$\frac{\binom{n - 2n\sqrt{(\log n)/k}}{2\sqrt{k \log n}}}{\binom{n}{2\sqrt{k \log n}}} < \left(\frac{n - 2n\sqrt{(\log n)/k}}{n} \right)^{2\sqrt{k \log n}} \leq \frac{1}{n^4}.$$

Thus every node receives every token with probability $1 - 1/n^3$. It is also clear that the algorithm finishes in $O(n\sqrt{k \log n})$ rounds.

Algorithm 1 can be derandomized using the technique of conditional expectations, as shown in Algorithm 3 and analyzed in Lemma 4.8.

Algorithm 1 can be derandomized using the standard technique of conditional expectations, as shown in Algorithm 3. Given a sequence of communication graphs, if node u broadcasts token t for Δ rounds and every node that receives token t also broadcasts t during that period, then we say node v is within Δ broadcast distance to u if and only if v receives token t by the end of round Δ . Let S be a set of nodes, and $|S| \leq 2\sqrt{k \log n}$. We use $\Pr[u; S; T]$ to denote the probability that the broadcast distance from node u to set X is greater than $2n\sqrt{(\log n)/k}$, where X is the union of S and a set of $2\sqrt{k \log n} - |S|$ nodes picked uniformly at random from $V \setminus T$, and $P(S, T)$ denotes the sum, over all u in V , of $\Pr[u; S; T]$.

Algorithm 3 Derandomized algorithm for Step 5 in Algorithm 1

Require: A sequence of communication graphs $G_i, i = 1, 2, \dots$, and $k \geq \sqrt{\log n}$

Ensure: A set of $2\sqrt{k \log n}$ nodes S such that the broadcast distance from every node u to S is within $2n\sqrt{(\log n)/k}$.

- 1: Set S and T be \emptyset .
 - 2: **for** each $v \in V$ **do**
 - 3: $T = T \cup \{v\}$
 - 4: **if** $P(S \cup \{v\}, T) \leq P(S, T)$ **then**
 - 5: $S = S \cup \{v\}$
 - 6: Return S
-

LEMMA 4.8. *The set S returned by Algorithm 3 contains at most $2\sqrt{k \log n}$ nodes, and the broadcast distance from every node to S is at most $2n\sqrt{(\log n)/k}$.*

Proof. Let us view the process of randomly selecting $2\sqrt{k \log n}$ nodes as a computation tree. This tree is a complete binary tree of height n . There are $n + 1$ nodes on any root-leaf path. The level of a node is its distance from the root. The computation starts from the root. Each node at the i th level is labeled by $b_i \in \{0, 1\}$, where 0 means not including node i in the final set and 1 means including node i in the set. Thus, each root-leaf path, $b_1 b_2 \dots b_n$, corresponds to a selection of nodes. For a node a in the tree, let

S_a (resp., T_a) denote the sets of nodes that are included (resp., lie) in the path from root to a .

By Theorem 1.5, we know that for the root node r , we have $P(\emptyset, S_r) = P(\emptyset, \emptyset) \leq 1/n^3$. If c and d are the children of a , then $T_c = T_d$, and there exists a real $0 \leq p \leq 1$ such that for each u in V , $\Pr[u; S_a; T_a]$ equals $p \Pr[u; S_c; T_c] + (1-p) \Pr[u; S_d; T_d]$. Therefore, $P(S_a, T_a)$ equals $pP(S_c, T_c) + (1-p)P(S_d, T_d)$. We thus obtain that $\min\{P(S_c, T_c), P(S_d, T_d)\} \leq P(S_a, T_a)$. Since we set S to be X in $\{S_c, S_d\}$ that minimizes $P(X, T_c)$, we maintain the invariant that $P(S, T) \leq 1/n^3$. In particular, when the algorithm reaches a leaf l , we know $P(S_l, V) \leq 1/n^3$. But a leaf l corresponds to a complete node selection, so that $\Pr[u; S_l; V]$ is 0 or 1 for all u , and hence $P(S_l, V)$ is an integer. We thus have $P(S_l, V) = 0$, implying that the broadcast distance from node u to set S_l is at most $2n\sqrt{(\log n)/k}$ for every l . Furthermore, $|S_l|$ is $2k\sqrt{\log n}$ by construction.

Finally, note that Step 4 of Algorithm 3 can be implemented in polynomial time, since for each u in V , $\Pr[u; S; T]$ is simply the ratio of two binomial coefficients with a polynomial number of bits. Thus, Algorithm 3 is a polynomial time algorithm with the desired property.

5 Concluding remarks and open questions

We studied the fundamental k -gossip problem in dynamic networks and showed a lower bound of $\Omega(n + nk/\log n)$ rounds for any token forwarding algorithm against a strongly adaptive adversary, significantly improving over the previous best bound of $\Omega(n \log k)$ [KLO10] for sufficiently large k . Our lower bound matches the known upper bound of $O(nk)$ up to a logarithmic factor, and establishes a near-linear factor separation between token-forwarding and network-coding based algorithms. While our bound rules out significantly faster algorithms in the strongly adaptive adversary model, we complement our lower bound by presenting the SYM-DIFF protocol for a weakly adaptive adversary. We show that SYM-DIFF is near-optimal when the starting distribution is well-mixed. Intuitively, a well-mixed distribution captures the “hard” regime

for information spreading in the adversarial setting, when most nodes have most of the tokens. Perhaps, the most interesting problem left open by our work is the analysis of SYM-DIFF in the weakly adaptive adversary model for an *arbitrary* starting distribution.

We also presented offline algorithms for k -gossip. An important intermediate model between the offline setting and the adaptive adversary models is the oblivious adversary model in which the adversary lays the dynamic network in advance (as in the offline setting), but the changing topology is revealed to the algorithm one round at a time. Finally, this paper has focused on models in which at most one token is sent per edge per round and the network can change every round. Subsequent to the announcement of our lower bound [DPRS11], the argument has been extended to the model where multiple tokens can be broadcast and the dynamic network is required to contain a stable subgraph for multiple rounds [HK].

References

- [AAG87] Y. Afek, B. Awerbuch, and E. Gafni. Applying static network protocols to dynamic networks. In *IEEE FOCS*, 1987.
- [ABBS01] B. Awerbuch, P. Berenbrink, A. Brinkmann, and C. Scheideler. Simple routing strategies for adversarial systems. In *IEEE FOCS*, 2001.
- [ABCHL11] C. Avin, M. Borokhovich, K. Censor-Hillel, and Z. Lotker. Order optimal information spreading using algebraic gossip. In *ACM PODC*, 2011.
- [ABNLP91] 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.
- [ABS03] B. Awerbuch, A. Brinkmann, and C. Scheideler. Anycasting in adversarial systems: Routing and admission control. In *ICALP*, 2003.
- [AC04] A. Agarwal and M. Charikar. On the advantage of network coding for improving network throughput. In *Information Theory Workshop*, 2004.
- [ACLY00] R. Ahlswede, N. Cai, S. Li, and R. Yeung. Network information flow. *Transactions on Information Theory*, 46(4):1204–1216, 2000.

- [AGHP92] Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta. Simple constructions of almost k -wise independent random variables. *Random Structures & Algorithms*, 3(3):289–304, 1992.
- [AGR92] Y. Afek, E. Gafni, and A. Rosen. The slide mechanism with applications in dynamic networks. In *ACM PODC*, 1992.
- [AKL08] C. Avin, M. Koucký, and Z. Lotker. How to explore a fast-changing world (cover time of a simple random walk on evolving graphs). In *ICALP*, 2008.
- [AL94] 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*, 1994.
- [APRU12] J. Augustine, G. Pandurangan, P. Robinson, and E. Upfal. Towards robust and efficient computation in dynamic peer-to-peer networks. In *ACM-SIAM SODA*, 2012.
- [APSPS92] B. Awerbuch, B. Patt-Shamir, D. Peleg, and M. E. Saks. Adapting to asynchronous dynamic networks. In *ACM STOC*, 1992.
- [ASWZ96] Roy Armoni, Michael E. Saks, Avi Wigderson, and Shiyu Zhou. Discrepancy sets and pseudorandom generators for combinatorial rectangles. In *37th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 412–421, 1996.
- [AW04] H. Attiya and J. Welch. *Distributed Computing: Fundamentals, Simulations and Advanced Topics (2nd edition)*. John Wiley Interscience, 2004.
- [BAL10] M. Borokhovich, C. Avin, and Z. Lotker. Tight bounds for algebraic gossip on graphs. In *IEEE ISIT*, 2010.
- [BCEG10] P. Berenbrink, J. Czyzowicz, R. Elsässer, and L. Gasieniec. Efficient information exchange in the random phone-call model. In *ICALP*, 2010.
- [BGPS06] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Trans. on Infor. Theory*, 52(6):2508–2530, 2006.
- [BNGNS00] A. Bar-Noy, S. Guha, J. Naor, and B. Schieber. Message multicasting in heterogeneous networks. *SIAM J. Comput.*, pages 347–358, 2000.
- [BYGI87] 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*, 1987.
- [CCC+98] M. Charikar, C. Chekuri, T. Cheung, Z. Dai, A. Goel, and S. Guha. Approximation algorithms for directed Steiner problems. *Journal of Algorithms*, 1998.
- [CFQS10] Arnaud Casteigts, Paola Flocchini, Walter Quattrociocchi, and Nicola Santoro. Time-varying graphs and dynamic networks. *CoRR*, abs/1012.0009, 2010. Short version in ADHOC-NOW 2011.
- [CMPS09] Andrea E. F. Clementi, Angelo Monti, Francesco Pasquale, and Riccardo Silvestri. Broadcasting in dynamic radio networks. *J. Comput. Syst. Sci.*, 75(4):213–230, 2009.
- [CP12] J. Chen and G. Pandurangan. Optimal gossip-based aggregate computation. *SIAM Journal on Computing*, 41(3):455–483, 2012. Conference version: ACM SPAA, 2012.
- [CS06] J. Cheriyan and M. Salavatipour. Hardness and approximation results for packing Steiner trees. *Algorithmica*, 2006.
- [CST12] Andrea Clementi, Riccardo Silvestri, and Luca Trevisan. Information spreading in dynamic graphs. In *PODC*, 2012.
- [DGH+87] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic algorithms for replicated database maintenance. In *ACM PODC*, 1987.
- [DMC06] S. Deb, M. Médard, and C. Choute. Algebraic gossip: a network coding approach to optimal multiple rumor mongering. *IEEE/ACM Trans. Netw.*, 14, 2006.
- [Dol00] S. Dolev. *Self-stabilization*. MIT Press, 2000.
- [DPRS11] C. Dutta, G. Pandurangan, R. Rajaraman, and Z. Sun. Information spreading in dynamic networks. *CoRR*, abs/1112.0384, 2011.
- [EGL+98] Guy Even, Oded Goldreich, Michael Luby, Noam Nisan, and Boban Velickovic. Efficient approximation of product distributions. *Random Struct. Algorithms*, 13(1):1–16, 1998.
- [GB81] E. Gafni and B. Bertsekas. Distributed algorithms for generating loop-free routes in networks with frequently changing topology. *IEEE Trans. Comm.*, 29, 1981.
- [GMR+12] Parikshit Gopalan, Raghu Meka, Omer Reingold, Luca Trevisan, and Salil Vadhan. Better pseudorandom generators from milder pseudorandom restrictions. In *FOCS*, 2012.
- [Hae11] B. Haeupler. Analyzing network coding gossip made easy. In *ACM STOC*, 2011.
- [HK] B. Haeupler and F. Kuhn. Personal communication.
- [HK11] B. Haeupler and D. Karger. Faster information dissemination in dynamic networks via net-

- work coding. In *ACM PODC*, 2011.
- [INW94] Russell Impagliazzo, Noam Nisan, and Avi Wigderson. Pseudorandomness for network algorithms. In *26th ACM Symp. on the Theory of Computing (STOC)*, pages 356–364, 1994.
- [JRS03] L. Jia, R. Rajaraman, and C. Scheideler. On local algorithms for topology control and routing in ad hoc networks. In *ACM SPAA*, 2003.
- [KK02] D. Kempe and J. Kleinberg. Protocols and impossibility results for gossip-based communication mechanisms. In *IEEE FOCS*, 2002.
- [KLO10] F. Kuhn, N. Lynch, and R. Oshman. Distributed computation in dynamic networks. In *ACM STOC*, 2010.
- [KO11] F. Kuhn and R. Oshman. Dynamic networks: Models and algorithms. *SIGACT News*, 42(1), 2011.
- [KS92] Bala Kalyanasundaram and Georg Schnitger. The probabilistic communication complexity of set intersection. *SIAM J. Discrete Math.*, 5(4):545–557, 1992.
- [KSSV00] R. M. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized rumor spreading. In *IEEE FOCS*, 2000.
- [Lei91a] F. T. Leighton. *Introduction to Parallel Algorithms and Architectures: Arrays, Trees, and Hypercubes*. Morgan-Kaufmann, 1991.
- [Lei91b] F. Thomson Leighton. *Introduction to Parallel Algorithms and Architectures: Arrays, Trees, Hypercubes*. Morgan Kaufmann, 1991.
- [Lu02] Chi-Jen Lu. Improved pseudorandom generators for combinatorial rectangles. *Combinatorica*, 22(3):417–433, 2002.
- [Lyn96] N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1996.
- [MAS06] D. Mosk-Aoyama and D. Shah. Computing separable functions via gossip. In *ACM PODC*, 2006.
- [New91] Ilan Newman. Private vs. common random bits in communication complexity. *Information Processing Letters*, 39(2):67–71, 1991.
- [Nis92] Noam Nisan. Pseudorandom generators for space-bounded computation. *Combinatorica*, 12(4):449–461, 1992.
- [Nis93] Noam Nisan. The communication complexity of threshold gates. In *Combinatorics, Paul Erdős is Eighty, number 1 in Bolyai Society Mathematical Studies*, pages 301–315, 1993.
- [NN93] Joseph Naor and Moni Naor. Small-bias probability spaces: efficient constructions and applications. *SIAM J. Comput.*, 22(4):838–856, 1993.
- [NZ96] Noam Nisan and David Zuckerman. Randomness is linear in space. *J. of Computer and System Sciences*, 52(1):43–52, February 1996.
- [OW05] Regina O’Dell and Roger Wattenhofer. Information dissemination in highly dynamic graphs. In *DIALM-POMC*, pages 104–110, 2005.
- [Pel00] David Peleg. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, 2000.
- [Raz92] Alexander A. Razborov. On the distributional complexity of disjointness. *Theor. Comput. Sci.*, 106(2):385–390, 1992.
- [SET03] P. Sanders, S. Egnér, and L. Tolhuizen. Polynomial time algorithms for network information flow. In *ACM SPAA*, 2003.
- [Top85] Donald M. Topkis. Concurrent broadcast for information dissemination. *IEEE Trans. Softw. Eng.*, 11:1107–1112, 1985.
- [Vio11] Emanuele Viola. Randomness buys depth for approximate counting. In *IEEE Symp. on Foundations of Computer Science (FOCS)*, 2011.
- [Vio13] Emanuele Viola. The communication complexity of addition. In *ACM-SIAM SODA*, 2013.
- [ZK02] L. Zosin and S. Khuller. On directed Steiner trees. In *ACM-SIAM SODA*, 2002.