

Order Optimal Information Spreading Using Algebraic Gossip

Chen Avin* Michael Borokhovich* Keren Censor-Hillel† Zvi Lotker*

Abstract

In this paper we study gossip based information spreading with bounded message sizes. We use algebraic gossip to disseminate k distinct messages to all n nodes in a network. For arbitrary networks we provide a new upper bound for uniform algebraic gossip of $O((k + \log n + D)\Delta)$ rounds with high probability, where D and Δ are the diameter and the maximum degree in the network, respectively. For many topologies and selections of k this bound improves previous results, in particular, for graphs with a constant maximum degree it implies that uniform gossip is *order optimal* and the stopping time is $\Theta(k + D)$.

To eliminate the factor of Δ from the upper bound we propose a non-uniform gossip protocol, TAG, which is based on algebraic gossip and an arbitrary spanning tree protocol \mathcal{S} . The stopping time of TAG is $O(k + \log n + d(\mathcal{S}) + t(\mathcal{S}))$, where $t(\mathcal{S})$ is the stopping time of the spanning tree protocol, and $d(\mathcal{S})$ is the diameter of the spanning tree. We provide two general cases in which this bound leads to an order optimal protocol. The first is for $k = \Omega(n)$, where, using a simple gossip broadcast protocol that creates a spanning tree in at most linear time, we show that TAG finishes after $\Theta(n)$ rounds for any graph. The second uses a sophisticated, recent gossip protocol to build a fast spanning tree on graphs with large weak conductance. In turn, this leads to the optimality of TAG on these graphs for $k = \Omega(\text{polylog}(n))$. The technique used in our proofs relies on queuing theory, which is an interesting approach that can be useful in future gossip analysis.

Michael Borokhovich is a full-time student at Ben Gurion University and is principally responsible for the paper's contributions.

*Department of Communication Systems Engineering, Ben Gurion University, Beer-Sheva, Israel. E-mail: {avin, borokhom, zviloo}@cse.bgu.ac.il.

†Computer Science and Artificial Intelligence Laboratory, MIT. E-mail: ckeren@csail.mit.edu. Supported by the Simons Postdoctoral Fellows Program.

1 Introduction

One of the most basic information spreading applications is that of disseminating information stored at a subset of source nodes to a set of sink nodes. Here we consider the *k-dissemination* case: k initial messages ($k \leq n$) located at some nodes (a node can hold more than one initial message) need to reach all n nodes. The *all-to-all communication* – each of n nodes has an initial value that is needed to be disseminated to all nodes – is a special case of k -dissemination. The goal is to perform this task in the lowest possible number of time steps when messages have *limited* size (i.e., a node may not be able to send all its data in one message).

Gossiping, or rumor-spreading, is a simple stochastic process for dissemination of information across a network. In a synchronous *round* of gossip, *each* node chooses a *single* neighbor as the *communication partner* and takes an action. In an asynchronous time model a single node wake-ups and chooses the *communication partner* and n consecutive steps are considered as one *round*. The *gossip communication model* defines how to select this neighbor, e.g., *uniform* gossip is when the communication partner is selected uniformly at random from the set of all neighbors. We then consider three possible actions: either the node pushes information to the partner (**PUSH**), pulls information from the partner (**PULL**), or does both (**EXCHANGE**), but here we mostly present results about **EXCHANGE**.

A *gossip protocol* uses a gossip communication model in conjunction with the choice of the particular content that is exchanged. Due to their distributed nature, gossip protocols have gained popularity in recent years and have found applications both in communication networks (for example, updating database replicated at many sites [9, 16], computation of aggregate information [17] and multicast via network coding [8], to name a few) as well as in social networks [18, 6].

In the current work we analyze *algebraic gossip* which is a type of network coding known as random linear coding (RLNC) [20, 19] that uses gossip algorithms for all-to-all communication and k -dissemination. In algebraic gossip the content of messages is the random linear combination of all messages stored at a sender. Once a node has received enough independent messages (independent linear equations) it can solve the system of linear equations and discover all the initial values of all other nodes. It has been proved [14] that network coding can improve the throughput of the network by better sharing of the network resources. Note, however, that in gossip protocols, nodes select a single partner, so for k -dissemination to succeed each node needs to receive at least k messages (of bounded size), hence at least a total of kn messages need to be sent and received. This immediately leads to a trivial lower bound of $\Omega(k)$ rounds for k -dissemination.

We study uniform and non-uniform algebraic gossip both in the synchronous and the asynchronous time models on arbitrary graph topologies. The stopping time obviously depends on the protocol, the gossip communication model, the graph topology, but also on the time model, as shown in other cases [11]. We now give an overview of our results followed by a discussion of previous work.

1.1 Overview of Our Results

Our first set of results is about the stopping time of uniform algebraic gossip. In [2] we have shown a tight bound of $\Theta(n)$ for all-to-all communication for graphs with constant maximum degree. To prove this, we used a reduction of gossip to a network of queues and analyzed the waiting times in the queues. Bounding the general k -dissemination case is significantly harder, despite some similarity in the tools used. Unless explicitly stated, all our results are for gossip using **EXCHANGE**

and are with high probability¹.

We provide a novel upper bound for uniform algebraic gossip of $O((k + \log n + D)\Delta)$ where D is the diameter and Δ is the maximum degree in the graph. For graphs with constant maximum degree this leads to a bound of $O(k + D)$. For the synchronous case we have a matching lower bound of $\Omega(k + D)$ which makes uniform algebraic gossip an order **optimal** gossip protocol for these graphs. We conjecture that the optimality holds for the asynchronous time model as well, but only show it when $k = \Omega(D)$.

However, there are topologies for which uniform algebraic gossip performs badly, e.g., in the barbell graph (two cliques connected with a single edge) it takes $\Omega(n^2)$ rounds to perform all-to-all communication [2]. This is usually the result of bottlenecks that exist in the graph and lead to low conductance. For such "bad" topologies we propose here a modification of the uniform algebraic gossip called *Tree based Algebraic Gossip* (TAG). The basic idea of the protocol is that it operates in two phases: first, using a gossip protocol \mathcal{S} it generates a spanning tree in which each node in the tree has a single parent. In the next phase, algebraic gossip is performed on the tree where each node does **EXCHANGE** with its parent. Let $t(\mathcal{S})$ and $d(\mathcal{S})$ be the stopping time of \mathcal{S} and the diameter of the tree generated by \mathcal{S} , respectively. For any spanning tree gossip protocol \mathcal{S} we prove for TAG an upper bound of: $O(k + \log n + d(\mathcal{S}) + t(\mathcal{S}))$ for the *synchronous* and the *asynchronous* time models. As a special case of a spanning tree protocol, one can use a gossip broadcast (or 1-dissemination) protocol \mathcal{B} – a protocol in which a single message originated at some node should be disseminated to all nodes. Interestingly, using a gossip broadcast for the spanning tree construction in TAG, eliminates the dependence on the diameter of the spanning tree in the synchronous time model, i.e., if we use \mathcal{B} as \mathcal{S} , we obtain the bound of $O(k + \log n + t(\mathcal{B}))$ rounds. For a general spanning tree protocol \mathcal{S} , it follows directly that if $k = \Omega(\max(\log n, d(\mathcal{S}), t(\mathcal{S})))$, TAG is an order **optimal** with a stopping time of $\Theta(k)$. We provide two examples of this scenario: the first example leads to the most significant result of the paper. Using a simple round-robin-based broadcast we show that TAG is an order optimal gossip protocol for *k-dissemination* in any topology when $k = \Omega(n)$. This imply, somewhat surprisingly, that for **any graph**, if $k = \Omega(n)$, TAG finishes in $\Theta(n)$ rounds. In the barbell graph mentioned above, TAG leads to a speedup ratio of n compare to the uniform algebraic gossip. The second example makes use of a recent non-uniform information dissemination protocol from [5] that works well on graphs G with large *weak conductance* denoted by $\Phi_c(G)$ for a parameter c (see Section 6). We provide sufficient conditions on k , c and $\Phi_c(G)$ that make TAG order optimal when using the protocol of [5] as a spanning tree protocol. Table 1 summarizes our main results of the paper and next, we discuss previous results.

1.2 Related Work

Uniform algebraic gossip was first proposed by Deb *et al.* in [8]. The authors studied uniform algebraic gossip using PULL and PUSH on the *complete graph* and showed a tight bound of $\Theta(k)$, for the case of $k = \omega(\log^3(n))$ messages. Boyd *et al.* [3, 4] studied the stopping time of a gossip protocol for the *averaging problem* using the **EXCHANGE** algorithm. They gave a bound for symmetric networks that is based on the second largest eigenvalue of the transition matrix or, equally, the mixing time of a random walk on the network, and showed that the mixing time captures the behavior of the protocol. Mosk-Aoyama and Shah [23] used a similar approach to [3, 4] to first analyze algebraic gossip on arbitrary networks. They consider symmetric stochastic matrices that

¹An event occurs with high probability (*w.h.p.*) if its probability is of at least $1 - O(\frac{1}{n})$.

Protocol	Graph	Synchronous	Asynchronous
Uniform AG	any graph	$O((k + \log n + D)\Delta)$	
	constant max degree	$\Theta(\mathbf{k} + \mathbf{D})$	$O(k + D)$ (*)
TAG	any graph	$O(k + \log n + d(\mathcal{S}) + t(\mathcal{S}))$	
		$O(k + \log n + t(\mathcal{B}))$	$O(k + \log n + d(\mathcal{B}) + t(\mathcal{B}))$
	$k = \Omega(n)$, any graph	$\Theta(\mathbf{n})$	
	$c = O(\log^p(n))$ $k = \Omega(\log^{2p+3}(n))$	$\Theta(\mathbf{k})$	$O(k + d(IS))$ (**)

Table 1: Overview of the main results of the paper. **Bold text** and Θ indicate order optimal result. (*) we prove an upper bound but conjecture it is optimal. (**) we prove the upper bound but conjecture is should be $\Theta(k)$. \mathcal{S} is a spanning tree protocol, \mathcal{B} is a broadcast protocol, and IS is an information dissemination gossip protocol from [5].

(may) lead to a non-uniform gossip and gave an upper bound for the PULL algorithm that is based on a measure of conductance of the network. As the authors mentioned, the offered bound is not tight, which indicates that their conductance-based measure does not capture the full behavior of the protocol.

In [2], we used queuing theory as a novel approach for analyzing algebraic gossip. We then gave an upper bound of $O(n\Delta)$ rounds for any graph for the case of all-to-all communication, where Δ is the maximum degree in the graph. In addition, a lower bound of $\Omega(n^2)$ was obtained for the barbell graph – the worst case graph for algebraic gossip. The bounds (upper and lower) in [2] were tight in the sense that they matched each other for the worst case scenario. The parameter Δ is simple and convenient to use, but, it does not fully capture the behavior of algebraic gossip. While it gives optimal ($\Theta(n)$) result for any constant-degree graphs (e.g., line, grid), it fails to reflect the stopping time of algebraic gossip on the complete graph, for example, by giving the $O(n^2)$ bound instead of $O(n)$.

A recent (yet, unpublished) work of Haeupler [13] is the most related to our work. Haeupler’s paper makes a significant progress in analyzing the stopping time of algebraic gossip. While all previous works on algebraic gossip used the notion of *helpful message/node* to look at the rank evaluation of the matrices each node maintains (this approach was initially proposed by [8]), Haeupler used a completely different approach. Instead of looking on the growth of the node’s subspace (spanned by the linear equations it has), he proposed to look at the orthogonal complement of the subspace and then analyze the process of its disappearing. This elegant and powerful approach led to very impressive results. First, a tight bound of $\Theta(n/\gamma)$ was proposed for all-to-all communication, where γ is a min-cut measure of the a related graph. This bound perfectly captures algebraic gossip behavior for any network topology. For the case of k -dissemination, the author gives a conjecture that the upper bound is of the form of $O(k + T)$ where T is the time to disseminate a single message to all the nodes. But formally, the bound that is proved is $O(k/\gamma + \log^2 n/\lambda)$ where λ is a conductance-based measure of the graph. The work in [13] implicitly considered the uniform algebraic gossip, but could be extend to non-uniform cases. It is therefore hard to compare TAG to the results of [13], nevertheless, our bounds for the uniform algebraic gossip are better for certain families of graphs. Table 2 presents few such examples.

Graph	$O(k/\gamma + \log^2 n/\lambda)/n$ [13]	$O((k + \log n + D)\Delta)$ [here]	Improvement factor
Line	$O(k + n \log^2 n)$	$O(k + n)$	$\log^2 n$
Grid	$O(k + \sqrt{n} \log^2 n)$	$O(k + \sqrt{n})$	$\log^2 n$ for $k = O(\sqrt{n})$
Binary Tree	$O(k + n \log^2 n)$	$O(k + \log n)$	$\Omega(\frac{n \log n}{k})$

Table 2: Comparison of our results with [13]

To give a quick summary of our results and previous work, the two main contributions of the paper are i) we prove that for graphs with constant maximum degree uniform algebraic gossip is order optimal for k -dissemination in the synchronous time model and ii) we offer a new non-uniform algebraic gossip protocol, TAG, that is order optimal for large selections of graphs and k . The rest of the paper is organized as follows: in Section 2 we give definitions. Section 3 proves results for uniform algebraic gossip and Section 4 presents the TAG protocol and its general bound. Sections 5 and 6, then, discuss cases where TAG is optimal.

2 Preliminaries

We model the communication network by a connected undirected graph $G_n = G_n(V, E)$, where V is the set of vertices and E is the set of edges. Number of vertices in the graph is $|V| = n$. Let $N(v) \subseteq V$ be a set of neighbors of node v and $d_v = |N(v)|$ its degree, let $\Delta = \max_v d_v$ be the maximum degree of G_n , and let D be the diameter of the graph.

We consider two time models: asynchronous and synchronous. In the *asynchronous* time model at every **timeslot**, **one node** selected independently and uniformly at random, takes an action and a single pair of nodes communicates². We consider n consecutive timeslots as one *round*. In the *synchronous* time model at every **round**, **every node** takes an action and selects a single communication partner. It is assumed that the information received in the current round will be available to a node for sending only at the beginning of the next round. A **Gossip communication model** (sometimes called gossip algorithm) defines the way information is spread in the network. In the gossip communication model, a node that wakes up (according to the time model) can initiate communication only with a single neighbor³ (i.e., communication partner). The model describes how the communication partner is chosen and in which direction (to – PUSH, from – PULL, or both – EXCHANGE) the message is sent. In this work we use the following communication models:

Definition 1 (Uniform Gossip). *Uniform gossip is a gossip in which a communication partner is chosen randomly and uniformly among all the neighbors.*

Definition 2 (Round-Robin (\mathcal{RR}) Gossip). *In round-robin gossip, the communication partner is chosen according to a fixed, cyclic list, of the nodes' neighbors. This list dictates the order in which neighbors are being contacted. If the initial partner is chosen at random, this gossip communication model is known as the quasirandom rumor spreading model[1, 10].*

²Alternatively, this model can be seen as each node having a clock which ticks at the times of a rate 1 Poisson process and there is a total n clock ticks per round [3].

³Note that this implies that in the synchronous model a node can communicate with more than a single neighbor, if other nodes initiate communication with it.

Gossip Protocols define the task and the message content. In turn, a gossip protocol can use any of the gossip communication models defined above (and others). We will use two types of gossip protocols here. The first is **STP Gossip** – protocols whose task is to create a *spanning tree* of the graph. The goal of a Gossip STP protocol \mathcal{S} is that every node, except a node which is the *root*, will have a single neighbor called the *parent*. Note that one simple way to generate a spanning tree is by using a 1-dissemination protocol, namely a broadcast protocol.

The second protocol, is a k -dissemination protocol called **Algebraic Gossip**. In algebraic gossip, every message sent by a node is sent according to the random linear coding (RLNC) technique which is described next. As mentioned, algebraic gossip can use any of the communication models presented above.

Random Linear Network Coding (RLNC) The random linear network coding approach is used in algebraic gossip for building outgoing messages to achieve fast information dissemination. Let \mathbb{F}_q be a field of size q . There are $k \leq n$ initial messages (x_1, \dots, x_k) that are represented as vectors in \mathbb{F}_q^r . We can represent every message as an integer value bounded by M , and therefore, $r = \lceil \log_q(M) \rceil$. All transmitted messages have a fixed length and represent linear equations over \mathbb{F}_q . The variables (unknowns) of these equations are the initial values $x_i \in \mathbb{F}_q^r$, $1 \leq i \leq k$ and a message contains the coefficients of the variables and the result of the equation; therefore the length of each message is: $r \log_2 q + k \log_2 q$ bits (and it is usually assumed that $r \gg n$ [8]). A message is built as a random linear combination of all messages stored by the node and the coefficients are drawn uniformly at random from \mathbb{F}_q . A received message will be appended to the node's stored messages only if it is independent of all linear equations (messages) that are already stored by the node and otherwise it is ignored. Nodes store messages (linear equations) in a matrix form and once the dimension (or rank) of the matrix becomes k , a node can solve the linear system and discover all the k messages.

The following definition is necessary for understanding the concept of helpfulness in the analysis of algebraic gossip.

Definition 3 (Helpful node and helpful message). *We say that a node x is a **helpful node** to a node y if and only if a random linear combination constructed by x can be linearly independent with all equations (messages) stored in y . We call a message a **helpful message** if it increases the dimension (or rank) of the node (i.e., the rank of the matrix in which the node stores the messages).*

3 k -dissemination with Uniform Algebraic Gossip

The main result of this section is that uniform algebraic gossip is order optimal k -dissemination for graphs with constant maximum degree and for any selection of k . It is formally stated in Theorem 3 and is an almost direct result of the following general bound for uniform algebraic gossip:

Theorem 1. *For any connected graph G_n , the stopping time of the uniform algebraic gossip protocol with k messages is $O((k + \log n + D)\Delta)$ rounds for synchronous and asynchronous time models w.h.p.*

The idea of the proof relies on the queuing networks technique we presented in [2]. The major steps of the proof are:

- Perform a Breath First Search (BFS) on G_n starting at an arbitrary node v . The search results in a directed shortest path spanning tree T_n rooted at v . The maximum depth l_{\max} of the tree T_n rooted at v is at most D .

- Reduce the problem of algebraic gossip on a tree T_n to a simple system of queues Q_n^{tree} rooted at v , where at each node we assume an infinite queue with a single server. Every initial message becomes a customer in the queuing system. The root v finishes once all the customers arrive at it.
- Show that the stopping time of the tree topology queuing system – Q_n^{tree} , is $O((k + \log n + l_{\max})n\Delta)$ timeslots *w.h.p.* So, we obtain the stopping time for the node v .
- Use union bound to obtain the result for all the nodes in G_n .

Just before we start the formal proof of Theorem 1, we present an interesting theorem related to queuing theory. The theorem gives the stopping time of the feedforward queuing system [7] arranged in a tree topology. Consider the following scenario: n identical M/M/1 queues arranged in a tree topology. There are no external arrivals, and there are k customers arbitrarily distributed in the system. In the feedforward network, a customer can not enter the same queue more than once, thus, customers eventually leave the system via the queue at the root of the tree. We ask the following question: how much time will it take for the last customer to leave the system?

Theorem 2. *Let Q_n^{tree} be a network of n nodes arranged in a tree topology, rooted at the node v . The depth of the tree is l_{\max} . Each node has an infinite queue, and a single exponential server with parameter μ . The total amount of customers in the system is k and they are initially distributed arbitrarily in the network. The time by which all the customers leave the network via the root node v is $t(Q_n^{tree}) = O((k + l_{\max} + \log n)/\mu)$ timeslots with probability of at least $1 - \frac{2}{n^2}$.*

The main idea of the proof is to show that the stopping time of the network Q_n^{tree} (i.e., the time by which all the customers leave the network) is stochastically⁴ smaller or equal to the stopping time of the systems of l_{\max} queues arranged in a line topology – $Q_{l_{\max}}^{line}$. Then, we make the system $Q_{l_{\max}}^{line}$ stochastically slower by moving all the customers out of the system and make them enter back via the farthest queue with the rate $\lambda = \mu/2$. Finally, we use Jackson’s Theorem for open networks to find the stopping time of the system. See Fig. 1 for the illustration. The full proof of the above theorem can be found in the appendix. We can now prove Theorem 1.

Proof of Theorem 1. We start the analysis of the uniform algebraic gossip with k messages and the asynchronous time model. First, we perform a Breath First Search (BFS) on G_n starting at an arbitrary node v . The search results in a directed shortest path spanning tree T_n rooted at v . The depth of T_n is l_{\max} , and since T_n is the shortest path tree, $l_{\max} \leq D$, where D is the diameter of the graph. On the tree T_n , consider a message flow towards the root v from all other nodes. Once k *helpful messages* arrive at v , it will reach rank k and finish the algebraic gossip protocol. We ignore messages that are not sent in the direction of v . Ignoring part of messages can only increase the stopping time of the algebraic gossip protocol.

We define a queuing system Q_n^{tree} by assuming an infinite queue with a single server at each node. The root of Q_n^{tree} is the node v . Customers of our queuing network are *helpful messages*, i.e., messages that increase the rank of a node they arrive at. This means that every customer arriving at some node increases its rank by 1. When a customer leaves a node, it arrives at the parent node. The queue length of a node represents a measure of *helpfulness* of the node to its parent, i.e., the number of *helpful messages* it can generate for it.

⁴For completeness, stochastic dominance is formally defined in appendix.

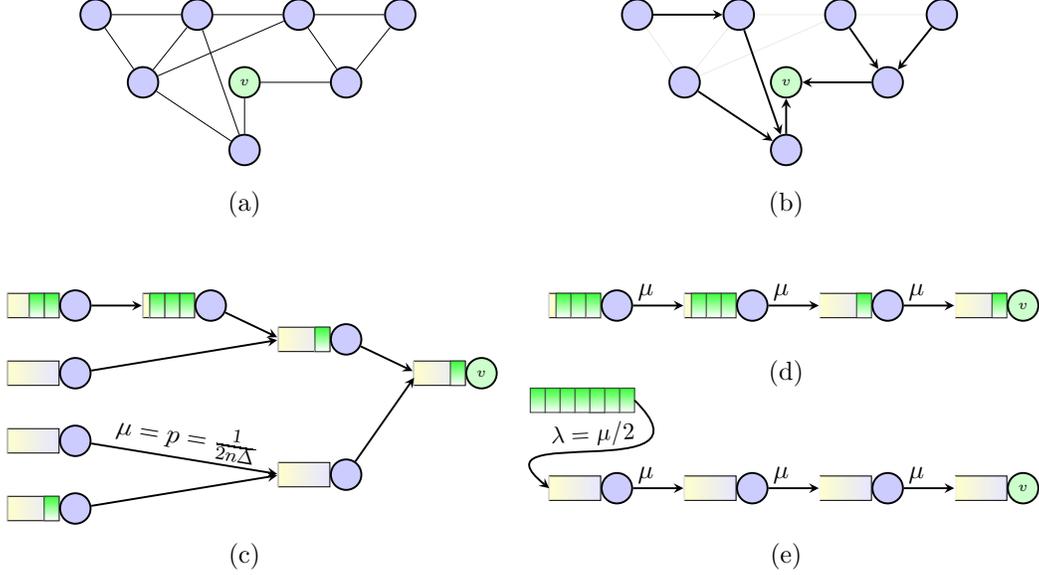


Figure 1: Reduction of AG to a system of queues. (a) – Initial graph G_n . (b) – Spanning tree T_n . (c) – System of queues Q_n^{tree} . (d) – System of queues $Q_{l_{\max}}^{line}$. Stopping time of $Q_{l_{\max}}^{line}$ is larger than of Q_n^{tree} . (e)–Taking all customers out of the system and use Jackson theorem for open networks.

The service procedure at a node is a transmission of a *helpful message* towards the node v (from a node to its parent). Lemma 2.1 in [8] gives a lower bound for the probability of a message sent by a *helpful node* to be a *helpful message*, which is: $1 - \frac{1}{q}$. In the uniform gossip communication model, the communication partner of a node is chosen randomly among all the node’s neighbors in the original graph G_n . The degree of each node in G_n is at most Δ . Thus, in the asynchronous time model, in a given timeslot, a *helpful message* will be sent over the edge in a specific direction with probability of at least $(1 - \frac{1}{q})/n\Delta$, where $\frac{1}{n}$ is the probability that a given node wakes up in a given timeslot, $\frac{1}{\Delta}$ is the minimal probability that a specific partner (the parent of the node) will be chosen, and $1 - \frac{1}{q}$ is the minimal probability that the message will be *helpful*. Thus, we can consider that the service time in our queuing system is geometrically distributed with parameter $p \geq (1 - \frac{1}{q})/n\Delta$, and since $q \geq 2$, we can assume the worst case: $p = \frac{1}{2n\Delta}$.

Lemma 2 in [2] shows that we can model the service time of each server as an exponential random variable with parameter $\mu = p$, since in this case, exponential servers are stochastically *slower* than geometric. Such an assumption can only increase the stopping time.

Theorem 2 with $\mu = p$ gives us an upper bound for the stopping time of the node v , $t_v = O((k + l_{\max} + \log n)2n\Delta)$ timeslots with probability of at least $1 - \frac{2}{n^2}$. Since the depth of every BFS tree is bounded by the diameter D , using a union bound we obtain the upper bound (in timeslots) for all the nodes in G_n :

$$\Pr \left(\bigcap_{v \in V} t_v = O((k + \log n + D)2n\Delta) \right) > 1 - \frac{2}{n}. \quad (1)$$

Thus we obtain the upper bound for uniform algebraic gossip: $O((k + \log n + D)\Delta)$ rounds. Next, we show that this bound holds also for the synchronous time model. The proof for the synchronous time model is almost the same as in the asynchronous case, except for the following

change. Instead of dividing time into timeslots, we measure it by rounds (1 round = n timeslots). In a given **round**, a *helpful message* will be sent over the edge in a specific direction with probability $p \geq (1 - \frac{1}{q})/\Delta$, where the $\frac{1}{\Delta}$ is the minimal probability that a specific partner (the parent of the node) will be chosen, and $1 - \frac{1}{q}$ is the minimal probability that the message will be *helpful*. Since $q \geq 2$, we can assume the worst case: $p = \frac{1}{2\Delta}$. The difference from the asynchronous model is the factor of n in p , since in the synchronous model, every node wakes up exactly once in a each round. Moreover, in the synchronous case (and in the EXCHANGE gossip variation) there is a possibility to receive 2 messages from the same node in one round (in the asynchronous time model it was impossible to receive 2 messages from the same node in one timeslot). We assume that if a node receives 2 messages from the same node at the same round, it will discard the second one. Such an assumption can only increase the stopping time of the protocol, and will make our analysis simpler. From that point on, the analysis is exactly the same as in the asynchronous case since Theorem 2 does not depend on the time model. \square

3.1 Optimality for Constant Maximum Degree Graphs and Synchronous Time

Following Theorem 1 we can state the main results of the section:

Theorem 3. *For any connected graph G_n with constant maximum degree, the stopping time of the uniform algebraic gossip protocol with k messages is $\Theta(k + D)$ in the synchronous time and $O(k + D)$ in the asynchronous time w.h.p.*

Proof. To show the upper bound the following simple claim is proved in the appendix:

Claim 1. *For any connected graph G_n with a constant maximum degree ($\Delta = O(1)$), the diameter of G_n is $\Omega(\log n)$.*

Now, using Claim 1 and fact the the maximum degree is constant the upper bound follows. For the lower bound note that in order to disseminate k messages to n nodes, at least kn transmissions should occur in the network. In synchronous time model, kn transmissions require at least $k/2$ rounds, since every round at most $2n$ messages are sent (2 transmissions per communication pair). In the asynchronous time model, kn transmissions require at least $kn/2$ timeslots, since at each timeslot at most 2 nodes transmit (due to EXCHANGE). Thus, in both time models, $\Omega(k)$ rounds are required. Moreover, in the synchronous time model, dissemination of a single message will take at least $D/2$ rounds, since in this model, a message can travel at most one hop in a single round. So, for the synchronous time model, the bound $\Theta(k + D)$ is tight and optimal. \square

4 TAG: k -dissemination with Tree-based Algebraic Gossip

We now describe the protocol TAG (Tree based Algebraic Gossip), which is a k -dissemination gossip protocol that exploits algebraic gossip in conjunction with a spanning tree gossip protocol \mathcal{S} (see Sec. 2). Given a connected network of n nodes and k messages x_1, \dots, x_k that are initially located at some nodes, the goal of the protocol TAG is to disseminate all the k messages to all the n nodes. The protocol consists of two phases. Both phases are performed simultaneously in the following way: if a node wakes up when the total number of its wakeups until now is even, it acts according to *Phase 1* of the protocol. If the node wakes up when the total number of its wakeups until now is odd, it acts according to *Phase 2* of the protocol.

Protocol TAG Pseudo code for node v . Example for asynchronous time model.

Require: $N(v)$, k , gossip spanning tree protocol \mathcal{S}

Initialize: $parent = null$

On odd wakeup: // Phase 1: EXCHANGE gossip spanning tree protocol \mathcal{S}
 1: choose partner $u \in N(v)$ and exchange messages with it according to \mathcal{S}
 2: according to \mathcal{S} decide if $parent = u$

On even wakeup: // Phase 2: EXCHANGE algebraic gossip
 3: **if** obtained $parent$ during the protocol \mathcal{S} **then**
 4: exchange messages with $parent$ according to algebraic gossip (RLNC)

On contact from other node $w \in N(v)$:
 5: **if** w performs Phase 1 **then**
 6: exchange messages with w according to \mathcal{S}
 7: according to \mathcal{S} decide if $parent = w$
 8: **else**(w performs Phase 2)
 9: exchange messages with w according to algebraic gossip (RLNC)

- In Phase 1, a node performs a spanning tree gossip protocol \mathcal{S} . Once a node becomes a part of the spanning tree, it obtains a **parent**.
- In Phase 2, a node is idle until it obtains a **parent** in Phase 1. From now on, in Phase 2, the node will perform an EXCHANGE algebraic gossip protocol with a fixed communication partner – its **parent**.

The following theorem gives an upper bound on the stopping time of the protocol TAG.

Theorem 4. Let $t(\mathcal{S})$ be the stopping time of the gossip spanning tree protocol \mathcal{S} performed at Phase 1, and let $d(\mathcal{S})$ be the diameter of the spanning tree created by \mathcal{S} . For any connected graph G_n , the stopping time of the k -dissemination protocol TAG, is:

$$t(\text{TAG}) = O(k + \log n + d(\mathcal{S}) + t(\mathcal{S})) \text{ rounds} \quad (2)$$

for synchronous and asynchronous time models, and w.h.p.

Proof. In order to prove this theorem, we will find the time needed to finish TAG, after Phase 1 is completed. Once Phase 1 is completed, every node knows its parent and thus, in Phase 2, we have the algebraic gossip EXCHANGE protocol on the spanning tree T_n , where communication partners of the nodes are their parents. The following lemma gives an upper bound on the stopping time of such a setting.

Lemma 1. Let T_n be a tree with n nodes, rooted at the node r , with depth l_{\max} . There are k initial messages located at some nodes in the tree. Consider algebraic gossip EXCHANGE protocol with the following communication model: the communication partner of a node is fixed to be its parent in T_n during the whole protocol. Then, the time needed for **all the nodes** to learn all the k messages is $O(k + \log n + l_{\max})$ rounds for the synchronous and asynchronous time models, with probability of at least $1 - \frac{2}{n}$.

The proof of Lemma 1 is very similar to the proof of Theorem 1, and relies on reducing the problem of algebraic gossip to a simple system of queues. The service time is geometrically distributed with a worst-case parameter $p = \frac{1}{2n}$. The Δ is eliminated from p since each node chooses now a single communication partner. Then, using Theorem 2 we obtain the stopping time of algebraic gossip with on the tree T_n . Detailed proof of Lemma 1 can be found in appendix.

Since for every choice of the tree root, the depth of the tree T_n (which was created using protocol $t(\mathcal{S})$) is bounded by its diameter, we can replace the l_{\max} in the bound $O(k + \log n + l_{\max})$ with $d(\mathcal{S})$. Now, we just add the stopping time of Phase 1 (the spanning tree time $- t(\mathcal{S})$) and the stopping time of Phase 2 (after Phase 1 has finished), and obtain that the number of rounds needed to complete the protocol TAG is $O(k + \log n + d(\mathcal{S}) + t(\mathcal{S}))$ *w.h.p.* \square

4.1 TAG protocol using 1-dissemination as a spanning tree protocol

The spanning tree task can be successfully performed by a simple gossip broadcast (or 1-dissemination) protocol. When a node receives for the first time the message, it marks the sending node as its parent. In such a way we obtain a spanning tree rooted at the node that initiated the broadcast protocol. Let us denote a gossip 1-dissemination protocol as \mathcal{B} . Clearly, the result of Theorem 4 can be rewritten as: $t(\text{TAG}) = O(k + \log n + d(\mathcal{B}) + t(\mathcal{B}))$. An interesting observation regarding the broadcast protocol \mathcal{B} , is that for synchronous time model the depth of the broadcast tree cannot be larger than the broadcast time (measured in rounds), i.e., $t(\mathcal{B}) \geq d(\mathcal{B})$. The last is true since a message can not travel more than one hop in a single round. Thus, for the synchronous time model we obtain that the number of rounds needed to complete the TAG protocol *w.h.p.* is:

$$t(\text{TAG}) = O(k + \log n + t(\mathcal{B})). \quad (3)$$

5 Optimal All-to-all Dissemination Using TAG

In this section we propose to use the TAG protocol in conjunction with a 1-dissemination (or broadcast) gossip protocol $\mathcal{B}_{\mathcal{R}\mathcal{R}}$ for spanning tree construction. For the case where $k = \Omega(n)$ messages need to be disseminated, TAG with $\mathcal{B}_{\mathcal{R}\mathcal{R}}$ achieves order optimal performance. For the case $k = \Omega(n)$ the lower bound of any gossip dissemination protocol is $\Omega(n)$ rounds. The bound from Theorem 4 gives $t(\text{TAG}) = O(k + \log n + d(\mathcal{S}) + t(\mathcal{S}))$, and if $k = n$ we obtain $O(n + t(\mathcal{S}))$. Thus, all we need to show is the existence of a gossip spanning tree protocol that finishes after $O(n)$ rounds *w.h.p.* on any graph.

Theorem 5. *For any connected graph G_n , the stopping time of the broadcast protocol with the round-robin communication model $- \mathcal{B}_{\mathcal{R}\mathcal{R}}$ is $O(n)$ rounds. In the asynchronous time model, this result holds with probability of at least $1 - n(2/e)^{3n}$, and in the synchronous time model, with probability 1.*

In order to prove Theorem 5 we need the following lemma which is proved in the appendix.

Lemma 2. *For any connected graph G_n with n nodes, the sum of the degrees of the nodes along any shortest path between any two nodes v and u is at most $3n$.*

Proof of Theorem 5. In this proof we assume the PUSH gossip variation, but it is clear that the result holds also for EXCHANGE. Without loss of generality, assume that the message that needs

to be disseminated is initially located at the node v . In the *round-robin* gossip, when a node is scheduled to transmit, it transmits a message to its neighbor according to the *round robin* scheme.

Consider a shortest path between v and some other node u . On the shortest path of length l there is exactly one node at the distance i from v , where $i \in [0, \dots, l]$, and $l \leq n - 1$. Let d_i be the degree of the node at distance i from v . In order to guarantee the delivery of the message from v to u , we need $\sum_{i=0}^l d_i$ transmissions in the following order: first, we need d_0 transmissions of the node v , then d_1 transmissions of the next node in the path $v \rightarrow u$, and so on until the message is delivered to u . From Theorem 2, $\sum_{i=0}^l d_i \leq 3n$.

In the asynchronous model, a node transmits at a given *timeslot* with probability $\frac{1}{n}$. So, the number of timeslots until some specific node transmits is a geometric random variable with parameter $\frac{1}{n}$. We define this geometric random variable as X , i.e., $X \sim \text{Geom}(\frac{1}{n})$.

The number of timeslots until $3n$ specific transmissions occur, is the sum of $3n$ independent geometric random variables. Using a Chernoff bound we obtain $O(n^2)$ timeslots (or $O(n)$ rounds) with exponential high probability. The last allows us to perform union bound for shortest paths to all other nodes in G , thus obtaining the $O(n)$ bound for the broadcast time. We omit here the formal part of the proof. The full proof can be found in the appendix.

It is easy to see that in the synchronous time model, $3n$ specific transmissions will occur exactly after $3n$ communication rounds. E.g., after d_0 rounds, v will perform d_0 transmissions – each one to different neighbor (according to the round-robin scheme). Thus, the message will be delivered to u after at most $3n$ rounds with probability 1. \square

Using Theorems 4 and 5 we obtain the upper bound on the stopping time of TAG with \mathcal{BR} as a spanning tree construction protocol: $O(k + \log n + d(\mathcal{S}) + n)$ which is $\Theta(n)$ for $k = \Omega(n)$.

6 Graphs with a Large Weak Conductance

For values of k which are smaller than n we use the information spreading protocol (hereafter, IS) of [5], which requires only a polylogarithmic number of rounds for broadcast on graphs with large *weak conductance*. Roughly speaking, the weak conductance is a value in $[0, 1]$ that measures the connectivity of subsets of nodes of a graph. It has been used to analyze the time required for *partial* information spreading, where each message is only required to reach some fraction of the nodes. This, in turn, has been applied in the analysis of the IS protocol to show that the running time for full information spreading inversely depends on the weak conductance. The graphs with large weak conductance, for which the IS protocol is fast, form a broad family of graphs, including graphs that exhibit some (though not too many) communication bottlenecks. A simple example is the barbell graph, consisting of two cliques of $n/2$ nodes, connected by a single edge, which corresponds to a bottleneck since information must pass along it, but the probability of randomly choosing it is small due to large node degrees. The IS protocol overcomes this and runs in a logarithmic number of synchronous rounds on the barbell.

We describe this result for both the synchronous and asynchronous time models considered. Although the IS protocol is designed to disseminate n messages originating one at each node, we will only use it for obtaining a spanning tree of our communication graph, while the actual information dissemination is done using algebraic gossip (i.e., we use the TAG protocol with IS as the spanning tree construction protocol). This is since the IS protocol sends large messages, while the goal of algebraic gossip is to address bandwidth concerns. The spanning tree is constructed as

follows. The information sent by a node v is an n -bit string, characterizing the nodes from which v heard from, whether directly or indirectly. This corresponds to empty initial inputs, and initially the n -bit string of node v is a unit vector, characterizing only the empty input of the node v itself. The n -bit string maintained and sent by a node v is monotone, in the sense that as time passes, its entries can only change from zero to one. The spanning tree that is created corresponds to each node v declaring its parent as the first node u from which it received a message that caused its most significant bit to change from zero to one. This means that this node received the input of the node w corresponding to the most significant bit (recall that the input itself is an empty string).

The following theorem characterizes the time required for the IS protocol to complete.

Theorem 6 ([5, Theorem 4.1]). *For every $c > 1$ and every $\delta \in (0, 1/3c)$, the IS protocol obtains full information spreading after at most $O(c(\frac{\log(n) + \log(\delta^{-1})}{\Phi_c(G)} + c))$ rounds, with probability at least $1 - 3c\delta$.*

In the synchronous model we can use the IS protocol in the TAG protocol, directly obtaining the following theorem, which shows optimality of TAG for certain families of parameters.

Theorem 7. *Let $c = O(\log^p(n))$ for some $p \geq 0$, let G be a graph with weak conductance $\Phi_c = \Omega(\frac{1}{\log^p(n)})$, and let $k = \Omega(\log^{2p+1}(n))$. With probability at least $1 - \frac{1}{n}$, the time for disseminating k messages using protocol TAG in conjunction with the IS protocol is $\Theta(k)$ synchronous rounds.*

We show that the IS protocol works in the asynchronous model as well. While this is not a direct usage of the protocol due to some subtleties, we nevertheless show how to obtain our result as for the synchronous model. Our analysis induces an overhead of $O(\log^2(n))$ rounds.

We do not change the protocol itself to cope with asynchrony, but rather analyze the time required using additional techniques. Roughly speaking, the outline of our analysis is showing that segments of the asynchronous execution simulate synchronous rounds. This allows us to use the original analysis of the protocol for the simulated rounds, which gives our result, as stated in the following theorem, and proved in the appendix.

Theorem 8. *Let $c = O(\log^p(n))$ for some $p \geq 0$, let G be a graph with weak conductance $\Phi_c = \Omega(\frac{1}{\log^p(n)})$, and let $k = \Omega(\log^{2p+3}(n))$. With probability at least $1 - \frac{1}{n}$, the time for disseminating k messages using protocol TAG in conjunction with the IS protocol is $O(k + l_{\max})$ rounds for the asynchronous time model, where l_{\max} is the depth of the spanning tree induced by the IS protocol.*

For completeness, we note that, in IS, during the even-numbered steps of each node the choice of neighbor to contact is randomized. For these steps alone, adapting the analysis Mosk-Aoyama and Shah [22] for the asynchronous case to our protocol, implies that the extra $\log(n)$ time slots can be avoided for the purpose of partial information spreading alone (as used in the proof of the information spreading protocol (see [5, Theorem 2.2])). However, as this cost is required anyhow to argue about the deterministic choices, made during the odd-numbered steps, we omit going through this adjustment.

References

- [1] S. Angelopoulos, B. Doerr, A. Huber, and K. Panagiotou. Tight bounds for quasirandom rumor spreading. *The Electronic Journal of Combinatorics*, 16(1):R102,1–R102,19, 2009.

- [2] M. Borokhovich, C. Avin, and Z. Lotker. Tight bounds for algebraic gossip on graphs. In *2010 IEEE International Symposium on Information Theory Proceedings (ISIT)*, pages 1758–1762, jun. 2010.
- [3] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 52(6):2508–2530, June 2006.
- [4] S. P. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Gossip algorithms: design, analysis and applications. In *IEEE International Conference on Computer Communications (INFOCOM)*, pages 1653–1664, 2005.
- [5] K. Censor-Hillel and H. Shachnai. Fast Information Spreading in Graphs with a Large Weak Conductance. To appear in the *22nd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2011.
- [6] A. Chaintreau, P. Fraigniaud, and E. Lebar. Opportunistic spatial gossip over mobile social networks. In *WOSP '08: Proceedings of the first workshop on Online social networks*, pages 73–78, New York, NY, USA, 2008. ACM.
- [7] H. Chen and D. Yao. *Fundamentals of Queueing Networks: Performance, Asymptotics, and Optimization*, volume 46 of *Applications of Mathematics*. Springer-Verlag, New York, first edition, 2001.
- [8] S. Deb, M. Médard, and C. Choute. Algebraic gossip: a network coding approach to optimal multiple rumor mongering. *IEEE Transactions on Information Theory*, 52(6):2486–2507, 2006.
- [9] A. J. Demers, D. H. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. E. Sturgis, D. C. Swinehart, and D. B. Terry. Epidemic algorithms for replicated database maintenance. *Operating Systems Review*, 22(1):8–32, 1988.
- [10] B. Doerr, T. Friedrich, and T. Sauerwald. Quasirandom rumor spreading. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, SODA '08, pages 773–781, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics.
- [11] C. Georgiou, S. Gilbert, R. Guerraoui, and D. R. Kowalski. On the complexity of asynchronous gossip. In *PODC '08: Proceedings of the twenty-seventh ACM symposium on Principles of distributed computing*, pages 135–144, New York, NY, USA, 2008. ACM.
- [12] G. R. Grimmett and D. R. Stirzaker. *Probability and random processes*. Oxford University Press, New York, third edition, 2001.
- [13] B. Haeupler. Analyzing Network Coding Gossip Made Easy. *ArXiv e-prints*, Oct. 2010.
- [14] T. Ho, R. Koetter, M. Medard, D. R. Karger, and M. Effros. The benefits of coding over routing in a randomized setting. In *IEEE International Symposium on Information Theory (ISIT)*, page 442, 2003.
- [15] R. V. D. Hofstad. Random graphs and complex networks. 2010.
- [16] R. M. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized rumor spreading. In *Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 565–574, 2000.

- [17] D. Kempe, A. Dobra, and J. Gehrke. Gossip-based computation of aggregate information. In *Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 482–491, 2003.
- [18] D. Kempe, J. Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *KDD '03: Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146, New York, NY, USA, 2003. ACM.
- [19] S.-Y. R. Li, R. W. Yeung, and N. Cai. Linear network coding. *IEEE Transactions on Information Theory*, 49(2):371–381, 2003.
- [20] M. Médard and R. Koetter. Beyond routing: An algebraic approach to network coding. In *IEEE International Conference on Computer Communications (INFOCOM)*, pages 122–130, 2002.
- [21] M. Mitzenmacher and E. Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, New York, NY, USA, 2005.
- [22] D. Mosk-Aoyama and D. Shah. Computing separable functions via gossip. In *PODC '06: Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*, pages 113–122, New York, NY, USA, 2006. ACM.
- [23] D. Mosk-Aoyama and D. Shah. Information dissemination via network coding. In *IEEE International Symposium on Information Theory Proceedings (ISIT)*, pages 1748–1752, 2006.
- [24] C.-H. Ng and S. Boon-Hee. *Queueing Modelling Fundamentals: With Applications in Communication Networks*. Wiley Publishing, 2008.

Appendix

Table of notations

n	Number of nodes
k	Number of messages needed to be disseminated
G_n	Connected graph with n nodes
T_n	Connected Tree graph with n nodes
Q_n^{tree}	Network of n queues arranged in a tree topology
$Q_{l_{\max}}^{line}$	Network of l_{\max} queues arranged in a line topology
D	Diameter of a graph
$N(v)$	Set of neighbors of the node v
d_v	Degree of the node v ($d_v = N(v) $)
Δ	Maximum degree of the graph ($\Delta = \max_v d_v$)
timeslot	Unit of time in the asynchronous time model
round	Unit of time in the synchronous time model (1 round = n timeslots)
\mathcal{S}	Some spanning tree gossip protocol
\mathcal{B}	Some broadcast (1-dissemination) gossip protocol
$d(\mathcal{S}), d(\mathcal{B})$	Diameter of the spanning tree created by the protocol
\mathcal{RR}	Round-robin communication model
$\mathcal{B}_{\mathcal{RR}}$	Broadcast gossip algorithm based on the round-robin communication model
l_{\max}	Depth of the tree created by a broadcast protocol
TAG	k -dissemination protocol that uses algebraic gossip and a spanning tree protocol
$t(\text{TAG}), t(\mathcal{S}), t(\mathcal{B})$	Stopping time of a protocol
$t(Q_n^{tree})$	Stopping time of a queuing system – time by which all customers leave the system

Table 3: Table of notations

Proof of Lemma 1

Lemma 1 (restated). *Let T_n be a tree with n nodes, rooted at the node r , with depth l_{\max} . There are k initial messages located at some nodes in the tree. Consider algebraic gossip **EXCHANGE** protocol with the following communication model: the communication partner of a node is fixed to be its parent in T_n during the whole protocol. Then, the time needed for **all the nodes** to learn all the k messages is $O(k + \log n + l_{\max})$ rounds for the synchronous and asynchronous time models, with probability of at least $1 - \frac{2}{n}$.*

Proof. The proof is very similar to the proof of Theorem 1, and relies on reducing the problem of algebraic gossip to a simple system of queues.

On T_n , consider a message flow towards an arbitrary node v (*not necessary the root of T_n*) from all other nodes. Once k *helpful messages* arrive at v , it will reach the rank k and finish the algebraic gossip protocol. Due to the proposed communication model, every node in T_n has a fixed

communication partner – its parent, so, each edge e in the tree has at least one node which will issue, on its wakeup, a bidirectional communication (EXCHANGE) over e . Thus, from every node, a message can be sent towards v . We ignore messages that are not sent in the direction of v . Ignoring part of messages can only increase the stopping time of the algebraic gossip protocol.

As in the proof of Theorem 1, we define a queuing system Q_n^{tree} by assuming an infinite queue with a single server at each node. The root of Q_n^{tree} will be an arbitrary node v , and let l_{\max}^v be the depth of the tree Q_n^{tree} .

The service procedure at a node is a transmission of a *helpful message* towards the node v . In our communication model, the communication partner of a node *is always its parent* in the tree. Thus, in the EXCHANGE gossip variation, in the asynchronous time model, in a given timeslot, a *helpful message* will be sent over the edge in a specific direction with probability of at least $(1 - \frac{1}{q})/n$, where $\frac{1}{n}$ is the probability that a given node wakes up in a given timeslot, and $1 - \frac{1}{q}$ is the minimal probability that the message will be *helpful*. Thus, we can consider that the service time in our queuing system is geometrically distributed with parameter $p \geq (1 - \frac{1}{q})/n$, and since $q \geq 2$, we can assume the worst case: $p = \frac{1}{2n}$.

Using Theorem 2 for the tree T_n rooted at v , with $\mu = p$, we get an upper bound for the stopping time of the node v , $t_v = O((k + l_{\max}^v + \log n)2n)$ timeslots with probability of at least $1 - \frac{2}{n^2}$, where the l_{\max}^v is the depth of the tree T_n rooted at v . Since $l_{\max}^v \leq 2l_{\max}$ (where l_{\max} is the depth of T_n rooted at r), we can replace the l_{\max}^v with $2l_{\max}$.

So, using union bound, we obtain the upper bound (measured in timeslots) for all the nodes in T_n :

$$\Pr \left(\bigcap_{v \in V} t_v = O((k + \log n + l_{\max})2n) \right) > 1 - \frac{2}{n}. \quad (4)$$

As in the proof of Theorem 1, in the synchronous time model, the service time distribution parameter p will be larger by a factor of n , and the time will be measured in rounds instead of timeslots. Thus, using the same arguments as in the proof of Theorem 1, we obtain the upper bound of $O(k + \log n + l_{\max})$ rounds for the synchronous time model. Thus, the lemma holds for both time models. \square

Proof of Lemma 2

Lemma 2 (restated). *For any connected graph G_n with n nodes, the sum of the degrees of the nodes along any shortest path between any two nodes v and u is at most $3n$.*

Proof. Without loss of generality, consider a BFS spanning tree of G rooted at some node v , and some arbitrary leaf u . We will find the maximum degree of the node located on the path $(v \rightarrow u)$ at distance i from the root v . Clearly, such a node can be connected only to the following nodes:

- Nodes that are located at distance $i - 1$ from the root. (It can not be connected to the nodes that are closer to the root (than $i - 1$) since then, its distance from the root would be $i - 1$ which contradicts the given BFS execution.)
- Nodes that are at the same distance i from the root.

- Nodes that are located at distance $i + 1$ from the root. (It can not be connected to the nodes that are farther from the root (than $i + 1$) since then, their distance from the root would be $i + 1$ which contradicts the given BFS execution.)

Let us define m_i as the number of nodes at distance i from the root. Clearly, $\sum_{i=0}^{n-1} m_i = n$. (The node at distance 0 is the root v). The degree of a node (at distance i from the root) can be at most: $d_i \leq (m_{i-1} + m_i + m_{i+1})$. Thus, the sum of degrees on a path of length l from the root to a leaf is at most: $d = \sum_{i=0}^l d_i$. Since $l \leq n - 1$, $d = \sum_{i=0}^l d_i \leq \sum_{i=0}^{n-1} d_i = \sum_{i=0}^{n-1} (m_{i-1} + m_i + m_{i+1}) \leq 3n$. \square

Proof of Claim 1

Claim 1 (restated). *For any connected graph G_n with a constant maximum degree ($\Delta = O(1)$), the diameter of G_n is $\Omega(\log n)$.*

Proof. Let us sum up all the n vertices of G_n in the following way. We start with an arbitrary node v and count it as 1. Then we split the sum of n vertices into D parts, where D is the diameter of G_n . Each part represents number of vertices located at the distance i ($i \in [0, \dots, D]$) from the node v . Since we are interested in the lower bound on D , we can assume the maximum degree for every node (so, the number of parts in the sum will be minimal). We define n_i ($i \in [0, \dots, D]$) as the number of vertices located at the distance i from the node v . Thus we obtain:

$$n_0 + n_1 + n_2 + \dots + n_D = n \tag{5}$$

$$1 + \Delta + \Delta^2 + \dots + \Delta^D \geq n \tag{6}$$

$$\frac{\Delta^{D+2} - 1}{\Delta - 1} \geq n \tag{7}$$

$$\Delta^{D+2} \geq n \tag{8}$$

$$D + 2 \geq \log_{\Delta} n \tag{9}$$

$$D = \Omega(\log n) \tag{10}$$

\square

Stochastic Dominance

Definition 4 (Stochastic dominance, stochastic ordering [15, 12]). *We say that a random variable X is stochastically less than or equal to a random variable Y if and only if $\Pr(X \leq t) \geq \Pr(Y \leq t)$, and such a relation is denoted as: $X \preceq Y$.*

Claim 2. *If for $i \in \{1, 2\}$, $X_i \preceq Y_i$, X_i are independent, and Y_i are independent, then: $\max_i X_i \preceq \max_i Y_i$.*

Proof.

$$\begin{aligned} \Pr(\max_i X_i \leq t) &= \bigcap_i \Pr(X_i \leq t) = \prod_i \Pr(X_i \leq t) \\ &\geq \prod_i \Pr(Y_i \leq t) = \Pr(\max_i Y_i \leq t). \end{aligned}$$

Hence:

$$\max_i X_i \preceq \max_i Y_i.$$

□

Claim 3. *If for $i \in \{1, 2\}$, $X_i \preceq Y_i$, X_i are independent, and Y_i are independent, then: $\sum_i X_i \preceq \sum_i Y_i$.*

Proof.

$$\Pr(X_1 + X_2 \leq t) = \int_{-\infty}^t f_{X_1+X_2}(s) ds,$$

where $f_{X_1+X_2}(s) = f_{X_1}(s) * f_{X_2}(s)$.

Thus:

$$\Pr(X_1 + X_2 \leq t) = \int_{-\infty}^t \int_{-\infty}^{\infty} f_{X_1}(\tau) f_{X_2}(s - \tau) d\tau ds \quad (11)$$

$$= \int_{-\infty}^{\infty} f_{X_1}(\tau) \Pr(X_2 \leq t - \tau) d\tau \quad (12)$$

$$\geq \int_{-\infty}^{\infty} f_{X_1}(\tau) \Pr(Y_2 \leq t - \tau) d\tau \quad (13)$$

$$= \int_{-\infty}^t \int_{-\infty}^{\infty} f_{X_1}(\tau) f_{Y_2}(s - \tau) d\tau ds \quad (14)$$

$$= \int_{-\infty}^t \int_{-\infty}^{\infty} f_{Y_2}(\tau) f_{X_1}(s - \tau) d\tau ds \quad (15)$$

$$= \int_{-\infty}^{\infty} f_{Y_2}(\tau) \Pr(X_1 \leq t - \tau) d\tau \quad (16)$$

$$\geq \int_{-\infty}^{\infty} f_{Y_2}(\tau) \Pr(Y_1 \leq t - \tau) d\tau \quad (17)$$

$$= \int_{-\infty}^t \int_{-\infty}^{\infty} f_{Y_2}(\tau) f_{Y_1}(s - \tau) d\tau ds \quad (18)$$

$$= \Pr(Y_1 + Y_2 \leq t). \quad (19)$$

Hence:

$$\sum_{i=1}^2 X_i \preceq \sum_{i=1}^2 Y_i.$$

□

Later arrivals yield later departures

Consider an infinite FCFS queue with a single exponential server. We define a_i as the time of arrival number i to the queue, and d_i as time of the departure number i from the queue. Let X_i

be the exponential random variable representing the service time of the arrival i . For all i , X_i 's are *i.i.d.*

Let a_i be a sequence of m arrival times to the queue, and d_i be a sequence of m departure times from the queue.

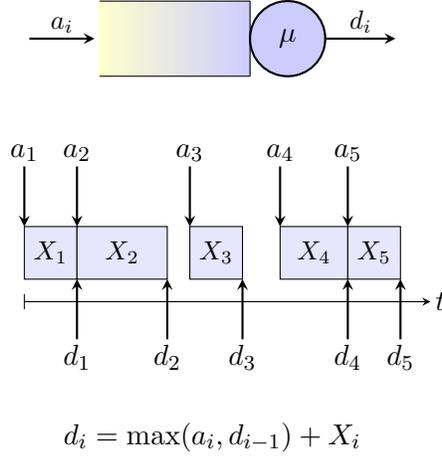


Figure 2: Arrival and departure times.

Lemma 3. *If the sequence a_i is replaced with another sequence of m arrivals \hat{a}_i , such that: $\hat{a}_i \succeq a_i \forall i \in [1, \dots, m]$, then, the resulting sequence of m departures will be such that: $\hat{d}_i \succeq d_i \forall i \in [1, \dots, m]$. I.e., if every new arrival occurred, stochastically, at the same time or later than the old arrival, then, every new departure from the queue will occur, stochastically, at the same time or later than the old departure.*

Proof. The proof is by induction on the arrival index j , $j \in [1, \dots, m]$.

- Induction basis: $\hat{d}_1 \succeq d_1$. Follows since $d_1 = a_1 + X_1$, $\hat{d}_1 = \hat{a}_1 + X_1$, and $\hat{a}_1 \succeq a_1$.
- Induction assumption: $\forall i < j : \hat{d}_i \succeq d_i$.
- Induction step: we need to show that $\hat{d}_j \succeq d_j$.

If the j 's arrival occurred when the server was busy, then $d_j = d_{j-1} + X_j$. If the server was idle when the j 's arrival occurred, then $d_j = a_j + X_j$. Thus, we can write:

$$d_j = \max(d_{j-1}, a_j) + X_j, \quad (20)$$

$$\text{and } \hat{d}_j = \max(\hat{d}_{j-1}, \hat{a}_j) + X_j. \quad (21)$$

Since from induction assumption: $\hat{d}_{j-1} \succeq d_{j-1}$, and $\hat{a}_j \succeq a_j$, using Claims 2 and 3, we obtain $\hat{d}_j \succeq d_j$. \square

Proof of Theorem 2

Theorem 2 (restated). *Let Q_n^{tree} be a network of n nodes arranged in a tree topology, rooted at the node v . The depth of the tree is l_{\max} . Each node has an infinite queue, and a single exponential*

server with parameter μ . The total amount of customers in the system is k and they are initially distributed arbitrarily in the network. The time by which all the customers leave the network via the root node v is $t(Q_n^{tree}) = O((k + l_{\max} + \log n)/\mu)$ timeslots with probability of at least $1 - \frac{2}{n^2}$.

Proof. We denote the nodes of the queuing system Q_n^{tree} as Z_j^l , where l ($l \in [1, \dots, l_{\max}]$) is the level of the node in the tree, and j is the node's index in the level l . The root of the Q_n^{tree} tree is the node Z_1^1 . All servers in the Q_n^{tree} network are ON all the time (work-conserving scheduling), i.e., servers work whenever they have customers to serve. There are no external arrivals to the system. Once a customer is serviced on the level l , it enters the appropriate queue at the level $l - 1$. When a customer is serviced by the root Z_1^1 , it leaves the network.

Now, let us define the auxiliary queuing systems: \hat{Q}_n^{tree} and $Q_{l_{\max}}^{line}$.

Definition 5 (Network \hat{Q}_n^{tree}). \hat{Q}_n^{tree} is the same network as Q_n^{tree} with the following change in the servers' scheduling:

At any given moment, only one server at every level l ($l \in [1, \dots, l_{\max}]$) is ON. Once a customer leaves level l , a server that will be scheduled (turned ON) at the level l , is the server which has in its queue a customer that has earliest arrival time to a queue at the level l among all the current customers at the level l . If there are customers that initially reside at the level l , they will be serviced by the order of their IDs (we assume for analysis that every customer has a unique identification number).

Definition 6 (Network of queues $Q_{l_{\max}}^{line}$). $Q_{l_{\max}}^{line}$ is the the following modification of the network Q_n^{tree} , that results in a network of l_{\max} queues arranged in a line topology.

For all $l \in [1, \dots, l_{\max}]$, we merge all the nodes at the level l to a single node (a single queue with a single server). We name this single node at the level l as the first node in Q_n^{tree} at the level l , i.e., Z_1^l . The customers that initially reside at level l will be placed in a single queue in the order of their IDs. This modification results in $Q_{l_{\max}}^{line}$ – a network of l_{\max} queues arranged in a line topology: $Z_1^{l_{\max}} \rightarrow Z_1^{l_{\max}-1} \rightarrow \dots \rightarrow Z_1^1$.

Definition 7 (Network of queues $\hat{Q}_{l_{\max}}^{line}$). $\hat{Q}_{l_{\max}}^{line}$ – is the same system as $Q_{l_{\max}}^{line}$ with the following modification. We take the last customer at some node Z_1^m ($m \in [1, \dots, l_{\max} - 1]$) and place it at the head of the queue of the node Z_1^{m+1} . I.e., we move one customer, one queue backward in the line of queues.

Definition 8 (Network of queues $\hat{Q}_{l_{\max}}^{line}$). $\hat{Q}_{l_{\max}}^{line}$ – is the same system as $Q_{l_{\max}}^{line}$ with the following modification. We move all the customers to the queue $Z_1^{l_{\max}}$. I.e., all the customers have to traverse now through all the l_{\max} queues in the line.

We summarize the queuing systems defined above in the short Table 4.

The proof of Theorem 2 consists of showing the following relations between the stopping times of the queuing systems:

$$t(Q_n^{tree}) \preceq t(\hat{Q}_n^{tree}) \approx t(Q_{l_{\max}}^{line}) \preceq t(\hat{Q}_{l_{\max}}^{line}) \preceq t(\hat{Q}_{l_{\max}}^{line}) = O((k + \log n + l_{\max})/\mu).$$

Stopping time of a queuing system $t(Q)$, is the time by which the last customer leaves the system (via the node Z_1^1). In order to compare the stopping times of queuing systems, we define the following ordered set (or sequence) of departure time from a server Z in a queuing system Q : $d(Z, Q) = (d_1(Z, Q), d_2(Z, Q), \dots, d_i(Z, Q), \dots)$, where $d_i(Z, Q)$ is the time of the departure number i from the node (server) Z .

Q_n^{tree}	Original system of n queues arranged in a tree topology. Fig. 3 (a).
\hat{Q}_n^{tree}	System of n queues arranged in a tree topology. Only one server is active at each level at a given time. Fig. 3 (b).
$Q_{l_{\max}}^{line}$	System of l_{\max} queues arranged in a line topology. Fig. 3 (c).
$\hat{Q}_{l_{\max}}^{line}$	System of l_{\max} queues arranged in a line topology. One customer is moved one queue backward.
$\hat{Q}_{l_{\max}}^{line}$	System of l_{\max} queues arranged in a line topology. All customers are moved backward to the queue $Z_1^{l_{\max}}$.

Table 4: Queuing systems used in the proof.

First, we want to show that the stopping time of Q_n^{tree} is at most the stopping time of the system \hat{Q}_n^{tree} , i.e., $t(Q_n^{tree}) \leq t(\hat{Q}_n^{tree})$.

Lemma 4. *In \hat{Q}_n^{tree} , every departure from the system (via Z_1^1) will occur, stochastically, at the same time or later than in Q_n^{tree} :*

$$d_i(Z_1^1, \hat{Q}_n^{tree}) \succeq d_i(Z_1^1, Q_n^{tree}) \quad \forall i \in [1, \dots, k]. \quad (22)$$

Thus, in \hat{Q}_n^{tree} , the last customer will leave the system, stochastically, at the same time or later than in Q_n^{tree} , or: $t(Q_n^{tree}) \leq t(\hat{Q}_n^{tree})$.

Proof. The proof is by induction on the tree level l , $l \in [1, \dots, l_{\max}]$.

- Induction basis: $\forall i, j : d_i(Z_j^{l_{\max}}, \hat{Q}_n^{tree}) \succeq d_i(Z_j^{l_{\max}}, Q_n^{tree})$. This is true since in \hat{Q}_n^{tree} , the nodes do not work all the time, and thus the departures will occur, stochastically, at the same time or later than in Q_n^{tree} . If there is a single node at the level l_{\max} , in \hat{Q}_n^{tree} it will be ON all the time as in Q_n^{tree} , and thus, the departures will occur, stochastically, at the same time in both systems.
- Induction assumption: for all $l > m$ ($m \geq 1$), $\forall i, j : d_i(Z_j^l, \hat{Q}_n^{tree}) \succeq d_i(Z_j^l, Q_n^{tree})$.
- Induction step: we need to show that: $\forall i, j : d_i(Z_j^m, \hat{Q}_n^{tree}) \succeq d_i(Z_j^m, Q_n^{tree})$.

By induction assumption, for $l = m + 1$: $\forall i, j : d_i(Z_j^{m+1}, \hat{Q}_n^{tree}) \succeq d_i(Z_j^{m+1}, Q_n^{tree})$. Now let us take a look at the departures from a node Z_j^m . There are two cases: Z_j^m is a leaf, and Z_j^m is not a leaf. If Z_j^m is a leaf, we can use the same argument as in the induction basis: in \hat{Q}_n^{tree} , the node Z_j^m does not work all the time, and thus the departures from it in \hat{Q}_n^{tree} cannot occur earlier than in Q_n^{tree} . If Z_j^m is not a leaf, it has input/inputs of arrivals from the level $m + 1$. Since the arrivals from the level $m + 1$ in \hat{Q}_n^{tree} occur, stochastically, at the same time or later than in Q_n^{tree} (by induction assumption), even if the node Z_j^m would work all the time (as in Q_n^{tree}), we would obtain from Lemma 3: $\forall i, j : d_i(Z_j^m, \hat{Q}_n^{tree}) \succeq d_i(Z_j^m, Q_n^{tree})$. Moreover, in \hat{Q}_n^{tree} , the node Z_j^m does not work all the time (unless it is the only node at the level m), thus the departure times in \hat{Q}_n^{tree} can be even larger. \square

Lemma 5. *In $Q_{l_{\max}}^{line}$, every departure from the system (via Z_1^1) will occur, stochastically, at the same time as in \hat{Q}_n^{tree} . Thus, in $Q_{l_{\max}}^{line}$, the last customer will leave the system, stochastically, at the same time as in \hat{Q}_n^{tree} .*

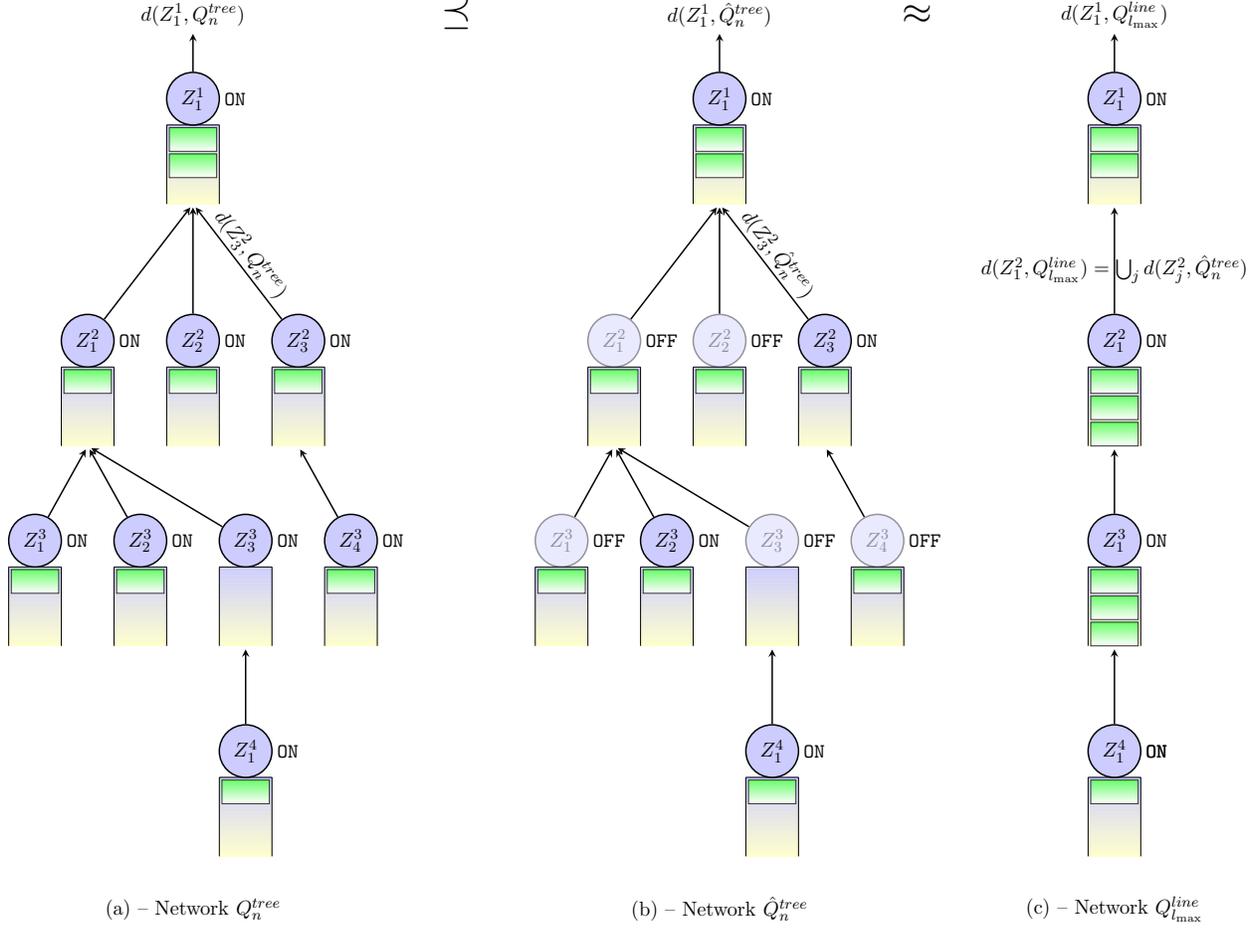


Figure 3: (a) – Network Q_n^{tree} , where all the servers work all the time. (b) – Network \hat{Q}_n^{tree} , where only one server at each level works at a given time. (c) – Network $Q_{l_{max}}^{line}$.

Proof. Consider the two following facts regarding the network \hat{Q}_n^{tree} . First, a customer entering the level l will be serviced after all the customers that arrived to the level l before it, are serviced. Second, at any given moment, only one customer is being serviced at the level l (if there is at least one customer at the nodes Z_j^l). These facts are true due to the scheduling of the servers in \hat{Q}_n^{tree} (Definition 5).

Clearly, the same facts are true for the network $Q_{l_{max}}^{line}$. First, any customer entering to the level l will be serviced after all the customers that arrived to the level l before it, are serviced. Second, at any given moment, only one customer is being serviced at the level l (if there is at least one customer in the node Z_1^l). These facts are true since in $Q_{l_{max}}^{line}$, at every level, there is a single queue with a single server (Definition 6).

So, the departure times of every customer from every level l ($l \in [1, \dots, l_{max}]$) are, stochastically, the same in both systems. The departures from level $l = 1$ are the departures from the node Z_1^1 , and thus the lemma holds. \square

Now we are going to move one customer, one queue backward and will show that the resulting

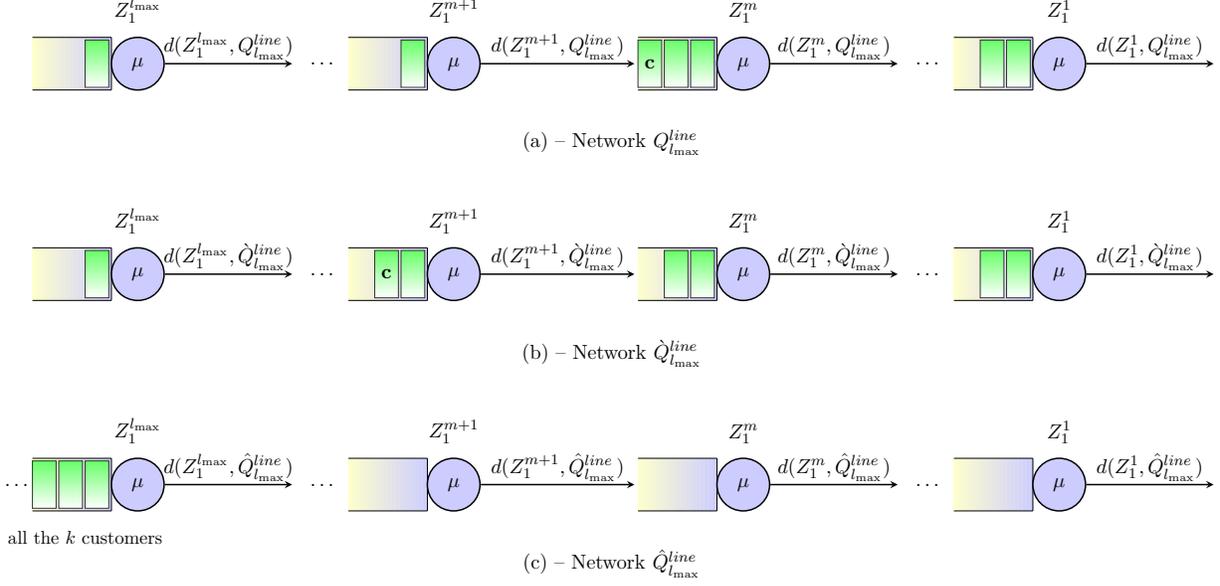


Figure 4: (a) – Network $Q_{l_{max}}^{line}$. (b) – Network $\hat{Q}_{l_{max}}^{line}$, where one customer is moved one queue backward. (c) – Network $\hat{Q}_{l_{max}}^{line}$, where all the customers are at the last queue.

system will have stochastically larger (or the same) stopping time.

Lemma 6. Consider a network $Q_{l_{max}}^{line}$. Let m be a level index: $m \in [1, \dots, l_{max} - 1]$. We take the last customer at the node Z_1^m and place it at the head of the queue of the node Z_1^{m+1} , and call the resulting network – $\hat{Q}_{l_{max}}^{line}$ (Fig. 4 (b)). Then:

$$d_i(Z_1^1, Q_{l_{max}}^{line}) \leq d_i(Z_1^1, \hat{Q}_{l_{max}}^{line}) \quad \forall i \in [1, \dots, k]. \quad (23)$$

Thus, in $\hat{Q}_{l_{max}}^{line}$, the last customer will leave the system, stochastically, at the same time or later than in $Q_{l_{max}}^{line}$, or: $t(Q_{l_{max}}^{line}) \leq t(\hat{Q}_{l_{max}}^{line})$.

Proof. We call the customer that was moved – customer c . Let us take a look at the times of arrivals to the node Z_1^m in $Q_{l_{max}}^{line}$ and in $\hat{Q}_{l_{max}}^{line}$. Since the customer c is already located in the queue of Z_1^m in $Q_{l_{max}}^{line}$, its arrival time can be considered as 0. In $\hat{Q}_{l_{max}}^{line}$, the arrival time of c is at least 0 (it should be serviced at Z_1^{m+1} before arriving at Z_1^m). Each one of the rest customers that should arrive at Z_1^m will arrive in $\hat{Q}_{l_{max}}^{line}$, stochastically, at the same time or later than in $Q_{l_{max}}^{line}$, since in $\hat{Q}_{l_{max}}^{line}$ the server Z_1^{m+1} should first service the customer c , and only then will start servicing the rest customers. Thus, $d_i(Z_1^{m+1}, \hat{Q}_{l_{max}}^{line}) \geq d_i(Z_1^{m+1}, Q_{l_{max}}^{line})$. Using Lemma 3 we obtain that: $d_i(Z_1^m, \hat{Q}_{l_{max}}^{line}) \geq d_i(Z_1^m, Q_{l_{max}}^{line})$. Iteratively applying Lemma 3 to the nodes Z_1^l , $l \in [m - 1, \dots, 1]$, we obtain the result: $d_i(Z_1^1, \hat{Q}_{l_{max}}^{line}) \geq d_i(Z_1^1, Q_{l_{max}}^{line})$. \square

Corollary 1. Consider a network $\hat{Q}_{l_{max}}^{line}$ (Definition 8) which is identical to the network $Q_{l_{max}}^{line}$ with the following change. In $\hat{Q}_{l_{max}}^{line}$, all the k customers are located at the node $Z_1^{l_{max}}$ (Fig. 4 (c)).

Then:

$$d_i(Z_1^1, Q_{l_{\max}}^{line}) \preceq d_i(Z_1^1, \hat{Q}_{l_{\max}}^{line}) \quad \forall i \in [1, \dots, k]. \quad (24)$$

Thus, in $\hat{Q}_{l_{\max}}^{line}$, the last customer will leave the system, stochastically, at the same time or later than in $Q_{l_{\max}}^{line}$, or: $t(Q_{l_{\max}}^{line}) \preceq t(\hat{Q}_{l_{\max}}^{line})$.

Proof. Given the network $Q_{l_{\max}}^{line}$ we take one customer from the tail of some queue (except the queue of the node $Z_1^{l_{\max}}$) and place it at the head of the queue of the preceding node in the $Q_{l_{\max}}^{line}$. According to the Lemma 6, we get a network in which every customer leaves via Z_1^1 , stochastically, not earlier than in $Q_{l_{\max}}^{line}$. Iteratively moving customers (one customer and one queue at a time) backwards we get finally the network $\hat{Q}_{l_{\max}}^{line}$ in which all the k customers are located at the node $Z_1^{l_{\max}}$. Since at each step, according to Lemma 6, the departure times from Z_1^1 could only get, stochastically, larger, the lemma holds. \square

Corollary 2. *The time it will take the last customer to leave the network of n queues arranged in a tree topology is, stochastically, the same or smaller than in the network of n queues arranged in a line topology where all the k customers are located at the farthest queue, i.e., $t(Q_n^{tree}) \preceq t(\hat{Q}_{l_{\max}}^{line})$.*

Proof. This corollary is a direct consequence of the Lemmas 4, 5, and the Corollary 1. \square

Now we are ready for the last step of the proof. We will find the stopping time of a system of queues arranged in a line topology and with all the customers located at the last queue.

Lemma 7. *The time it will take to the last customer to leave the system $\hat{Q}_{l_{\max}}^{line}$ (l_{\max} MM1 queues arranged in a line topology) is $O((k + \log n + l_{\max})/\mu)$ with probability of at least $1 - \frac{1}{n^2}$.*

Proof. Initially, all the customers (from now we will call them *real* customers) are located in the last ($Z_1^{l_{\max}}$) queue. We now take all the *real* customers out of this queue and will make them enter the system (via the $Z_1^{l_{\max}}$) from outside. We define the *real* customers' arrivals as a Poisson process with rate $\lambda = \frac{\mu}{2}$. So, $\rho = \frac{\lambda}{\mu} = \frac{1}{2} < 1$ for all the queues in the system. Clearly, such an assumption only increases the stopping time of the system (stopping time is the time until the last customer leaves the system). According to Jackson's theorem, which proof can be found in [7], there exists an equilibrium state. So, we need to ensure that the lengths of all queues at time $t = 0$ are according to the equilibrium state probability distribution. We add *dummy* customers to all the queues according to the stationary distribution. By adding additional *dummy* customers to the system, we make the *real* customers wait longer in the queues, thus increasing the stopping time.

We will compute the stopping time $t(\hat{Q}_{l_{\max}}^{line})$ in two phases: Let us denote this time as $t_1 + t_2$, where t_1 is the time needed for the k 'th customer to arrive at the first queue, and t_2 is the time needed for the k 'th customer to pass through all the l_{\max} queues in the system.

From Jackson's Theorem, it follows that the number of customers in each queue is independent, which implies that the random variables that represent the waiting times in each queue are independent. To continue with the proof we need the following lemmas; the first is a classical result from queuing theory, the proof of the second lemma is omitted.

Lemma 8 ([24], section 4.3). *Time needed to cross one M/M/1 queue in the equilibrium state has an exponential distribution with parameter $\mu - \lambda$.*

Lemma 9. *Let Y be the sum of n independent and identically distributed exponential random variables. Then, for $\alpha > 1$:*

$$\Pr(Y < \alpha E[Y]) > 1 - (2e^{-\alpha/2})^n. \quad (25)$$

The random variable t_1 is the sum of k independent random variables distributed exponentially with parameter $\mu/2$. From Lemma 8 we obtain that t_2 is the sum of l_{\max} independent random variables distributed exponentially with parameter $\mu - \lambda = \mu/2$. $E[t_1] = \sum_{i=1}^k 2/\mu = 2k/\mu$, and by taking $\alpha = 2 + 4\frac{\ln n}{k}$, we obtain:

$$(26)$$

$$\Pr(t_1 < (4k + 8 \ln n)/\mu) > 1 - (2e^{-(2+4\frac{\ln n}{k})/2})^k \quad (27)$$

$$= 1 - \left(\frac{2}{e}\right)^k e^{-2 \ln n} \quad (28)$$

$$\geq 1 - e^{-2 \ln n} \quad (29)$$

$$\geq 1 - \frac{1}{n^2}. \quad (30)$$

In a similar way we obtain:

$$\Pr(t_2 < (4l_{\max} + 8 \ln n)/\mu) > 1 - \frac{1}{n^2}. \quad (31)$$

$t(\hat{Q}_{l_{\max}}^{line}) = t_1 + t_2$, thus, using union bound:

$$\Pr(t_1 + t_2 < (4k + 4l_{\max} + 16 \ln n)/\mu) > 1 - \frac{2}{n^2} \quad (32)$$

and thus:

$$t(\hat{Q}_{l_{\max}}^{line}) = O((k + l_{\max} + \log n)/\mu) \quad (33)$$

w.p. of at least $1 - \frac{2}{n^2}$.

□

From Claim 2 we obtain that $t(Q_n^{tree}) \preceq t(\hat{Q}_{l_{\max}}^{line})$ and thus: $t(Q_n^{tree}) = O((k + l_{\max} + \log n)/\mu)$ w.p. of at least $1 - \frac{2}{n^2}$. □

Proof of Theorem 5

Theorem 5 (restated). *For any connected graph G_n , the stopping time of the broadcast protocol with the round-robin communication model $\mathcal{B}_{\mathcal{RR}}$ is $O(n)$ rounds. In the asynchronous time model, this result holds with probability of at least $1 - n(2/e)^{3n}$, and in the synchronous time model, with probability 1.*

For the proof we need the following lemma:

Lemma 10. *Let X be a sum of m independent and identically distributed geometric random variables (each one with parameter $p > 0$) and $E[X] = \frac{m}{p}$. Then, for $\alpha > 1$:*

$$\Pr(X \leq \alpha E[X]) > 1 - (\alpha e^{1-\alpha})^m. \quad (34)$$

Proof. First, we will define Y as the sum of k independent Bernoulli random variables, i.e., $Y = \sum_{i=1}^k Y_i$, where $Y_i \sim \text{Bernoulli}(p)$. Let us notice that:

$$\Pr(X \leq k) = \Pr(Y \geq m) \quad (35)$$

The last is true since the event of observing at least m successes in a sequence of k Bernoulli trials implies that the sum of m independent geometric random variables is no more than k . On the other hand, if the sum of m independent geometric random variables is no more than k it implies that m successes occurred no later than the k -th trial and thus $Y \geq m$.

Now we will use a Chernoff bound for the sum of independent Bernoulli random variables presented in [21]: For any $0 < \delta < 1$ and $\mu = \mathbb{E}[Y]$:

$$\Pr(Y \leq (1 - \delta)\mu) \leq \left(\frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \right)^\mu. \quad (36)$$

Since $\mu = \mathbb{E}[Y] = kp$, and by letting $\delta = \frac{kp-m}{kp}$ we obtain:

$$\Pr(Y \leq (1 - \delta)\mu) = \Pr(Y \leq m) \leq \left(\frac{m}{e^{\frac{m-kp}{m}} kp} \right)^{-m}. \quad (37)$$

$$\Pr(Y \geq m) > 1 - \left(\frac{m}{e^{\frac{m-kp}{m}} kp} \right)^{-m} \quad (38)$$

$$\Pr(X \leq k) > 1 - \left(\frac{m}{e^{\frac{m-kp}{m}} kp} \right)^{-m} \quad (39)$$

By substituting $k = \alpha \frac{m}{p} = \alpha \mathbb{E}[X]$ (where $\alpha > 1$) we obtain:

$$\Pr(X \leq \alpha \mathbb{E}[X]) > 1 - \left(\frac{e^\alpha}{e\alpha} \right)^{-m} \quad (40)$$

□

Proof of Theorem 5. In this proof we assume the PUSH gossip variation, but it is clear that the result holds also for EXCHANGE.

Without loss of generality, assume that the message that needs to be disseminated is initially located at the node v . In the *round-robin* gossip, when a node is scheduled to transmit, it transmits a message to its neighbor according to the *round robin* scheme. I.e, at every transmission a message is sent to a different neighbor.

Consider a shortest path between v and some other node u . On the shortest path of length l there is exactly one node at the distance i from v , where $i \in [0, \dots, l]$, and $l \leq n - 1$. Let d_i be the degree of a node at the distance i from v . In order to guarantee the delivery of the message from v to u , we need $\sum_{i=0}^l d_i$ transmissions in the following order: first, we need at d_0 transmissions of the node v , then d_1 transmissions of the next node in the path $v \rightarrow u$, and so on until the message is delivered to u . From Theorem 2, $\sum_{i=0}^l d_i \leq 3n$.

In the asynchronous model, a node transmits at a given *timeslot* with probability $\frac{1}{n}$. So, the number of timeslots until some specific node transmits is a geometric random variable with parameter $\frac{1}{n}$. We define this geometric random variable as X , i.e., $X \sim \text{Geom}\left(\frac{1}{n}\right)$.

The number of timeslots until $3n$ specific transmissions occur, is the sum of $3n$ independent geometric random variables. Using Lemma 10 we obtain the bound of $O(n^2)$ timeslots (or $O(n)$ rounds) with exponential high probability. The last allows us to perform union bound for shortest paths to all other nodes in G , thus obtaining the $O(n)$ bound for the broadcast time.

Now, let us write the above more formally:

We define $t_{v \rightarrow u}$ as the time it takes to guarantee the delivery of the message from the node v to an arbitrary node u . As we showed above, $t_{v \rightarrow u}$ is the number of timeslots until $3n$ specific transmissions occur, so:

$$t_{v \rightarrow u} = \sum_{i=1}^{3n} X_i, \quad (41)$$

$X_i \forall i \in [1, \dots, 3n]$ are i.i.d. and distributed as X ,
where $X \sim \text{Geom}\left(\frac{1}{n}\right)$.

Thus,

$$\mathbb{E}[t_{v \rightarrow u}] = \mathbb{E}\left[\sum_{i=1}^{3n} X_i\right] = \sum_{i=1}^{3n} \mathbb{E}[X_i] = \sum_{i=1}^{3n} n = 3n^2. \quad (42)$$

From Lemma 10 with $\alpha = 2$:

$$\Pr(t_{v \rightarrow u} \leq 2\mathbb{E}[t_{v \rightarrow u}]) > 1 - (2/e)^{3n}, \quad (43)$$

or

$$\Pr(t_{v \rightarrow u} \leq 6n^2) > 1 - (2/e)^{3n}. \quad (44)$$

Now, we will apply a union bound on probabilities of the events: $t_{v \rightarrow u'} > 6n^2$, where $u' \in V$. Notice, that $\mathbb{E}[t_{v \rightarrow u}] = 3n^2$ for all $u' \in V$.

$$\Pr\left(\bigcup_{u' \in V} (t_{v \rightarrow u'} > 6n^2)\right) \leq \sum_{u' \in V} \Pr(t_{v \rightarrow u'}), \quad (45)$$

so,

$$\Pr\left(\bigcup_{u' \in V} (t_{v \rightarrow u'} > 6n^2)\right) \leq n(2/e)^{3n}. \quad (46)$$

Thus,

$$\Pr\left(\bigcap_{u' \in V} (t_{v \rightarrow u'} \leq 6n^2)\right) > 1 - n(2/e)^{3n}. \quad (47)$$

So, we obtain the result of $O(n^2)$ timeslots, or $O(n)$ rounds.

Easy to see that in the synchronous time model, $3n$ specific transmissions will occur exactly after $3n$ communication rounds. E.g., after d_0 rounds, v will perform d_0 transmissions – each one to different neighbor (according to the round-robin scheme). Thus, the message will be delivered to u after at most $3n$ rounds with probability 1. \square

Proof of Theorem 8

Theorem 8 (restated). *Let $c = O(\log^p(n))$ for some $p \geq 0$, let G be a graph with weak conductance $\Phi_c = \Omega(\frac{1}{\log^p(n)})$, and let $k = \Omega(\log^{2p+3}(n))$. With probability at least $1 - \frac{1}{n}$, the time for disseminating k messages using protocol TAG in conjunction with the IS protocol is $O(k + l_{\max})$ rounds for the asynchronous time model, where l_{\max} is the depth of the spanning tree induced by the IS protocol.*

Proof. To simulate one round of a synchronous protocol, we consider the execution of the protocol for $O(n \log(n))$ time slots, which is $\log(n)$ asynchronous rounds. With high probability, $1 - \delta$ for some small δ , each node takes at least one step. This follows from a standard coupon collector's argument, as steps of each node correspond to a different coupon. The crucial property of the information spreading protocol that allows our analysis to go through is its *monotone* nature, that is, the information collected and sent by a node is an n -bit string whose entries can only turn from zero to one as time passes. This implies that whenever each node took at least one step, the strings obtained can only contain more one entries than the strings obtained by one round of the synchronous model (recall that the goal is for all nodes to obtain a string of ones). Hence, after $O(T \log(n))$ asynchronous rounds, the information the nodes have is at least the information that they have after T rounds in the synchronous model. This does not yet conclude the proof, for the following reason. The analysis of the synchronous protocol goes through in this simulation except for one argument [5, Claim 1], which bounds the size of the deterministic list of subset of neighbors that is maintained by a node v . This size is bounded by the number of steps taken by v . On one hand, we need the number of steps taken by v in each $O(n \log(n))$ time slots to be at least one to argue the simulation, but on the other hand, it may be that a node takes a larger number of steps. This would imply that its list is larger than in the corresponding synchronous case. However, returning to the coupon collector's problem, we have that actually no node takes more than $O(\log(n))$ steps in each $O(n \log(n))$ time slots (within the same high probability). This implies that the size of the list a node maintains is at most a multiplicative factor of $O(\log(n))$ larger than its size after T synchronous rounds. Since this size is used, in turn, to bound the number of synchronous rounds required, we have to add an additional $O(\log(n))$ factor to the number of rounds in the asynchronous model.

Finally, we note that the probability of failure of the coupon collector's argument (in either the lower or upper bound on the number of steps per node) needs to be added up for all simulated rounds. When this number of rounds T is polylogarithmic in n , we have that using a union bound we remain with a high probability for the entire argument. \square