Randomized information dissemination in dynamic environments

Soh, De Wen; Tay, Wee Peng; Quek, Tony Q. S.

2013

Soh, D. W., Tay, W. P., & Quek, T. Q. S. (2013). Randomized information dissemination in dynamic environments. *IEEE/ACM Transactions on Networking*, 21(3), 681-691.
doi:10.1109/TNET.2012.2209676

<https://hdl.handle.net/10356/89751>

<https://doi.org/10.1109/TNET.2012.2209676>

© 2012 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The published version is available at:
<https://doi.org/10.1109/TNET.2012.2209676>

Downloaded on 26 Apr 2023 11:50:07 SGT

Randomized Information Dissemination in Dynamic Environments

De Wen Soh, *Student Member, IEEE*, Wee Peng Tay, *Member, IEEE*,
and Tony Q.S. Quek, *Senior Member, IEEE*,

Abstract—We consider randomized broadcast or information dissemination in wireless networks with switching network topologies. We show that an upper bound for the ϵ -dissemination time consists of the conductance bound for a network without switching, and an adjustment that accounts for the number of informed nodes in each period between topology changes. Through numerical simulations, we show that our bound is asymptotically tight. We apply our results to the case of mobile wireless networks with unreliable communication links, and establish an upper bound for the dissemination time when the network undergoes topology changes and periods of communication link erasures.

Index Terms—Randomized broadcast, gossip, distributed algorithms, sensor networks.

I. INTRODUCTION

The need for data aggregation and dissemination over large networks has fueled interest in distributed consensus and broadcasting algorithms. Distributed consensus has applications in many areas, including data fusion [1]–[4], coordination [5], [6] and parallel computation [7]. Various distributed consensus algorithms [8]–[10] require nodes to exchange or broadcast information to other nodes in the network. Randomized gossip has been investigated as a mechanism for spreading information in a network [9]–[19]. In randomized gossip, each node chooses one of its neighbors randomly at each time step. Information is exchanged with this chosen neighbor, so that both nodes have the same information after the exchange. If a node transmits information to its neighbor, a PUSH is said to occur; if a node extracts information from its neighbor, then a PULL occurs. Given sufficient time, the piece of information from the source node is distributed to all nodes in the network. Randomized gossiping is robust to node and link failures since the failure of some nodes or links in the network does not stop the dissemination process. Randomized gossiping is also simple, asynchronous, and does not require any routing, which can be challenging in an ad hoc network.

Manuscript received March 28, 2011; revised April 21, 2012; accepted June 20, 2012. The material in this paper has been presented in part at the IEEE International Workshop on Signal Processing Advances in Wireless Communications, San Francisco, CA, June 2011. This research was supported, in part, by the MOE AcRF Tier 1 Grant RG25/10, the SRG ISTD 2012 037, and the CAS Fellowships for Young International Scientists.

D. W. Soh was with the Institute for Infocomm Research. He is now with the Department of Electrical Engineering at Yale University, New Haven, CT. (e-mail: dewen.soh@yale.edu).

W. P. Tay is with the Nanyang Technological University, Singapore (e-mail: wptay@ntu.edu.sg).

T. Q. S. Quek is with the Singapore University of Technology and Design, and the Institute for Infocomm Research, Singapore (e-mail: tonyquek@sutd.edu.sg).

Digital Object Identifier XXX

Moreover, randomized gossiping is a distributed algorithm that depends only on local information, since each node only needs to know its set of neighbors in the network.

If the network topology is known, a randomized gossiping algorithm may not perform as well as a deterministic, albeit more complex, algorithm that makes use of knowledge of the network topology. It is therefore of interest to quantify the amount of time required for the source node message to be broadcast in the network. A common measure of performance used is the ϵ -dissemination time, where ϵ is a small positive number. This is the minimum time by which the probability that all the nodes in the network have the message is at least $1 - \epsilon$. Upper bounds for the dissemination time have been established in [10], [12]–[16], [19]–[23]. However, all these results assume an underlying communication network topology that remains constant over time. In a wireless network with mobile nodes (e.g. a vehicular network), the underlying communication network topology varies over time and the communication channels are subject to different fading conditions. Therefore, most of the previous results in the literature do not apply to randomized broadcasting in a wireless network, with some exceptions like [19], [24], where communications between nodes are assumed to fail with a fixed probability. While there have been much work involving average consensus in dynamic topologies [25]–[29], there has been little or no work that considers information broadcast in varying topologies.

In this paper, we consider randomized broadcast in a wireless network. We study the case where the network graph has multiple phases, each with a distinct topology and channel failure distribution. The times at which the network changes from one phase to another are called switching times. Our main contributions are the following.

- (a) In the case where there is a single switching time, and where communications between nodes are assumed to be reliable,
 - we provide an upper bound for the ϵ -dissemination time when, in addition to the network topology switch, there are potentially new nodes joining the network at the switching time; and
 - we establish an upper bound for the ϵ -dissemination time when some informed nodes leave the network at the switching time.
- (b) We generalize the bounds in (a) to the case where there are multiple switching times, and when the number of nodes are nondecreasing at each switching time.

Fig. 1. The topology of the network switches as the dissemination process is occurring. The black vertices represent the informed nodes whereas the white vertices are the uninformed ones. The arrows represent the directions of communication links.

- (c) We quantify the impact of communication erasures on the information exchange between nodes. We establish an upper bound for the ϵ -dissemination time when the network experiences multiple switches in the network topology and the channel erasure probabilities.

The rest of the paper is organized as follows. In Section II, we define some notations, and present our model in detail. In Section III, we investigate the case where the underlying connectivity matrix undergoes a single switch. In Section IV, we generalize to the case where the connectivity matrix undergoes multiple switches. In Section V, we analyze the impact of channel erasures on the dissemination time. In Section VI, we discuss a numerical example demonstrating that our theoretical bounds are asymptotically tight up to a multiplicative constant. Finally, we summarize and provide some concluding remarks in Section VII.

II. SYSTEM MODEL

Let $\mathcal{G} = (\mathcal{V}, P)$ be a weighted, directed graph, where $\mathcal{V} = \{1, 2, \dots, n\}$ is a set of n nodes, and the matrix $P = (P_{ij})$ is a stochastic matrix. There is an arc from node i to node j if and only if (iff) $P_{ij} > 0$. We say that P is the connectivity matrix of the node set $\mathcal{V}$. All the nodes that can be accessed from node i are called the neighbors of node i . In this paper, we focus on doubly stochastic connectivity matrices P .

We study the convergence times of randomized gossip algorithms [10], [12]–[16], under various dynamical changes. Suppose node $v \in \mathcal{V}$ has a piece of information that it wants to disseminate to the other $n - 1$ nodes. Let time be discrete and indexed by $t \in \mathbb{N}$, and let $S(t)$ be the set of nodes that possess the information at time t , with $S(0) = \{v\}$. At each time t , each node i chooses to contact one of its neighbors j with probability P_{ij} . In the case where $i = j$, node i does not contact any of its neighbors. When contact is made, if either node i or j has the information, then both will get the information. This is known as the PUSH-PULL algorithm. PUSH happens when node i chooses node j , and node i has the information while node j does not, so node i “pushes” the information to node j . PULL happens in the opposite scenario: node i “pulls” the information from node j if node j has the information. The dissemination process is said to be complete when the information has been disseminated to all nodes in the network.

For $\epsilon > 0$, let the ϵ -dissemination time be

$$\mathcal{T}_0(\epsilon) = \sup_{v \in \mathcal{V}} \inf \{t : \mathbb{P}(S(t) \neq \mathcal{V} \mid S(0) = \{v\}) \leq \epsilon\}.$$

Related to the dissemination time is the notion of conductance. For $k \geq 1$, the k -conductance of a doubly stochastic matrix P is defined as

$$\Phi_k(P) = \min_{S \subseteq \mathcal{V}: |S| \leq k} \frac{\sum_{i \in S, j \in S^c} P_{ij}}{|S|}.$$

We also define $\Phi_0(P) = \infty$. The conductance $\Phi(P)$ of a $n \times n$ doubly stochastic connectivity matrix P , is defined to be

$$\Phi(P) = \Phi_{n/2}(P).$$

It turns out that the conductance of the connectivity matrix is closely related to the dissemination time. Note that the conductance here is defined according to the connectivity probability matrix P . There are other definitions of conductance that sum up the number of edges in the cut instead of the probabilities of those edges, but here we are looking at the total probability of the edges in the cuts. It was shown in [10] that for a fixed, doubly stochastic connectivity matrix P , the ϵ -dissemination time is

$$\mathcal{T}_0(\epsilon) = O\left(\frac{\log n + \log \epsilon^{-1}}{\Phi(P)}\right). \quad (1)$$

In a dynamic environment, for example when nodes are mobile, the connectivity matrix P changes over time (cf. Figure 1). An example is a group of vehicles, where the inter-vehicular communication network changes dynamically over time as the relative positions of the vehicles change. We wish to analyze how changes in P affect the ϵ -dissemination time. The times at which these changes take place are called switching times. In general, switching times may depend on n , the size of the network. In the example of the vehicular network, the vehicles may be traveling along a road whose width changes at a particular point, thus affecting the topology of the network. If only a small fraction of vehicles have passed through the change point, the network topology remains approximately unchanged. A possible model is to consider the topology to have a switching time only after a significant portion of all the vehicles have passed through the change point. In this paper, we assume that network topology switchings occur on a time scale that is much longer than inter-node communications. Therefore we discretize time into intervals, and investigate the effect of a sequence of switching times on the ϵ -dissemination time.

In Section III-A, we first study the case where there is a single switching time with new nodes (if any) joining the network at the switching time. In Section III-B, we consider the case where nodes are removed at the switching time, i.e., some nodes have moved out of our region of interest. In Section IV, we consider a somewhat generalized problem in which there are multiple switching times, but with nondecreasing number of nodes. Finally in Section V, we incorporate the effect of fading into our models, and study the dissemination times when the network experiences multiple periods of fading.

Let $\tau_1 < \dots < \tau_l$ be increasing switching times, with $\tau_0 = 0$ and $\tau_{l+1} = \infty$. Suppose that for $i = 0, 1, \dots, l$, the network graph is $\mathcal{G}_i = (\mathcal{V}_i, P_i)$ in the time interval $[\tau_i, \tau_{i+1})$, with $|\mathcal{V}_i| = n_i$. At time τ_i , for $0 < i \leq l$, the set of nodes changes from $\mathcal{V}_{i-1}$ to $\mathcal{V}_i$, and the connectivity matrix switches

from P_{i-1} to P_i . We assume that all connectivity matrices are doubly stochastic.

Let $\tau = (\tau_1, \dots, \tau_l)$ be a sequence of switching times, and $\mathcal{G} = (\mathcal{G}_0, \dots, \mathcal{G}_l)$ be the corresponding sequence of graphs in each period between switching times. Let $|\tau| = l$ be the number of switching times.

The ϵ -dissemination time for a system with l switching times is defined as

$$\mathcal{T}_l(\epsilon) = \sup_{v \in \mathcal{V}_0} \inf \{t : \mathbb{P}(S(t) \neq \mathcal{V}_l \mid S(0) = \{v\}) \leq \epsilon\}.$$

Note that from our definition, we are interested in the first time when all the nodes in $\mathcal{V}_l$ are informed with probability at least $1 - \epsilon$. This time may occur before some or all of the switching times if those switching times are large. For example in Section III-B, we study the case where informed nodes may be removed at a switching time. If the topology switches only after all the nodes in the network have received the information, then the change in topology does not affect the dissemination time of the information.

In the same spirit as [10], [13], [14], [30], [31] (c.f. (1)), we are interested in characterizing order upper bounds for $\mathcal{T}_l(\epsilon)$, where $l \geq 1$.

Throughout this paper, we assume, for simplicity, that the connectivity matrices P_i are deterministic. A similar analysis holds if we let P_i be a random matrix that is independent of the filtration up to τ_i , and replace P_i in the order bound with $\mathbb{E}[P_i]$.

A. Notations and Preliminaries

In this section, we define some notations that we frequently use. We also provide some preliminary results that will be useful in our subsequent proofs.

We will make use of certain quantities extensively in our proofs. These quantities serve as markers that divide time into different intervals. Which interval the switching times fall in determines the ϵ -dissemination time. For example if a topology switch happens after the dissemination process has completed, then the change in the connectivity matrix at that switching time does not affect the information spreading.

For $k \leq l$, let

$$\lambda_{k,l}(P) = \frac{8 \left(\log(2n_l \epsilon^{-1}) + \sum_{j=0}^k \log n_j \right)}{\Phi_{n_l/2}(P)}, \quad (2)$$

and

$$\zeta_n(P) = \frac{2 \log n + \log \epsilon^{-1}}{\Phi(P)}, \quad (3)$$

Let the first time when more than half the final number of nodes have the information be $L = \inf\{t : |S(t)| > n_l/2\}$, and the completion time be $T = \inf\{t : S(t) = \mathcal{V}_l\}$. We define

$$M(t_0) = \max\{T - t_0, 0\},$$

as the remaining dissemination time starting from time t_0 .

We will also make use of the following results established in [10] in our proofs. These results involve graphs that do not change over time.

Proposition 1: Suppose that $\mathcal{G}_0 = (\mathcal{V}_0, P)$, $|\mathcal{V}_0| = n_0$, and the network graph $\mathcal{G}_0$ remains static over time.

- (i) Let the first time that the number of informed nodes exceed k be $L_k = \inf\{t : |S(t)| > k\}$. Then, for any time t , we have

$$\mathbb{P}(L_k > t) \leq n_0 \exp\left(-\frac{\Phi_k(P)t}{8}\right). \quad (4)$$

Furthermore, for $L = L_{n_0/2}$, we have

$$\mathbb{P}(L < \lambda_{0,0}(P)) \geq 1 - \frac{\epsilon}{2n_0},$$

$$\mathbb{P}(T < \lambda_{0,0}(P) + \zeta_{n_0}(P)) \geq 1 - \frac{\epsilon}{n_0}.$$

- (ii) Suppose that $|S(t)| \leq n_0/2$. Then

$$\mathbb{E}[|S(t+1)| - |S(t)| \mid S(t)] \geq |S(t)| \frac{\Phi(P)}{2}. \quad (5)$$

- (iii) Suppose that $|S(t_0)| \leq k$. Let $L'_k = L_k - t_0$ and $t' = t - t_0$. Then, for $t' \geq 0$, the random variable¹, defined as

$$Z(t') = \frac{\exp\left(\frac{\Phi_k(P)}{8}(L'_k \wedge t')\right)}{|S(L'_k \wedge t')|},$$

is a supermartingale. As a result, $\mathbb{E}[Z(t') \mid L_k > t_0] \leq \mathbb{E}[Z(0) \mid L_k > t_0]$.

- (iv) Suppose that $|S(t)^c| \leq k$. Then

$$\mathbb{E}[|S(t+1)^c| \mid S(t)] \leq |S(t)^c|(1 - \Phi_k(P)). \quad (6)$$

- (v) Suppose that $|S(t_0)^c| \leq n_0/2$. Then

$$\mathbb{P}(M(t_0) \leq \zeta_{n_0}(P) \mid L_{n_0/2} \leq t_0) \geq 1 - \frac{\epsilon}{2n_0}.$$

We establish the following lemma which we use extensively in our proofs. This result is similar to Proposition 1(i).

Lemma 1: Let t_0 be a fixed time, and suppose that the underlying graph has n nodes and connectivity matrix P for all $t \in [t_0, t_1]$, where t_1 is some arbitrary time greater than t_0 . Then, for all $t \in [t_0, t_1]$, we have

$$\mathbb{P}(L_k > t \mid L_k > t_0) \leq n \exp\left(-\frac{\Phi_k(P)(t_0 - t)}{8}\right).$$

The proof is provided in the Appendix, Section VIII-A.

III. SINGLE SWITCHING TIME

In this section, we consider the case where there is a single switching time, with the network graph switching from $\mathcal{G}_0 = (\mathcal{V}_0, P_0)$ to $\mathcal{G}_1 = (\mathcal{V}_1, P_1)$ at time τ_1 , which can be a function of n . We generalize to multiple switching times in Section IV.

If we restart the information dissemination at time τ_1 , from (1), all nodes will be informed by time

$$O\left(\tau_1 + \frac{\log n_1 + \log \epsilon^{-1}}{\Phi(P_1)}\right) \quad (7)$$

with probability at least $1 - \epsilon$. This is a naïve upper bound for $\mathcal{T}_1(\epsilon)$. However, many nodes may have become informed by time τ_1 if the switching time is large. We therefore seek an upper bound that takes into account the dissemination before the switching time.

¹We use the notation $a \wedge b = \min\{a, b\}$.

A. Nondecreasing Nodes

In this section, we consider the case where $n_1 \geq n_0$, i.e., $n_1 - n_0$ uninformed nodes join the network at time τ_1 . No nodes are allowed to leave the network, but uninformed nodes are allowed to join the network at τ_1 .

The main result of this section is the following theorem. If n_0 is small compared to n_1 , the number of informed nodes at time τ_1 is also small compared to n_1 , since there can be at most n_0 informed nodes at time τ_1 . Therefore, the ϵ -dissemination time is dominated by the events after time τ_1 , and Theorem 1(i) follows from our previous discussion. On the other hand, if $n_0 \geq n_1/2$ and the switching time τ_1 is greater than $c(\log n + \log \epsilon^{-1})/\Phi(P)$ for some constant c , then from (1), with high probability, all n_0 nodes have the information before the topology switches. Therefore the total dissemination time consists of τ_1 and the time required to disseminate information to the remaining $n_1 - n_0$ nodes. This is Theorem 1(ii), whose proof is provided in the Appendix, Section VIII-B. The final and most interesting case is Theorem 1(iii), where $n_0 \geq n_1/2$ but the switching time is not significantly large. In this case, with high probability, the information has disseminated to a strict subset of the n_0 nodes by the switching time τ_1 . We find a probabilistic bound on the time required to disseminate information to the remaining nodes. The proof is provided in the Appendix, Section VIII-C.

Theorem 1: Suppose that $n_1 \geq n_0 \geq 2$.

(i) If $n_0 < n_1/2$, we have

$$\mathcal{T}_1(\epsilon) = O\left(\tau_1 + \frac{\log n_1 + \log \epsilon^{-1}}{\Phi(P_1)}\right).$$

(ii) If $n_0 \geq n_1/2$ and $\tau_1 \geq \lambda_{0,0}(P_0) + \zeta_{n_0}(P_0)$, we have

$$\mathcal{T}_1(\epsilon) = O\left(\tau_1 + \frac{\log(n_1 - n_0 + 1) + \log \epsilon^{-1}}{\Phi_{n_1-n_0}(P_1)}\right).$$

We set $\Phi_0(P) = \infty$.

(iii) If $n_0 \geq n_1/2$ and $\tau_1 < \lambda_{0,0}(P_0) + \zeta_{n_0}(P_0)$, we have

$$\mathcal{T}_1(\epsilon) = O\left(\frac{26 \log n_1 + 9 \log(2\epsilon^{-1})}{\Phi(P_1)} + \tau_1 \left(1 - \frac{\Phi_{n_1/2}(P_0)}{\Phi(P_1)}\right)\right).$$

The order bound in Theorem 1(iii) shows an intuitive dependence on the $n_1/2$ -conductance of P_0 and the conductance of P_1 . We can think of the conductances as an indicator of the speed at which dissemination occurs. The larger the $n_1/2$ -conductance of P_0 , the faster the dissemination occurs before the switching time τ_1 . This in turn reduces the number of uninformed nodes at τ_1 , and therefore it is more probable that dissemination completes in a shorter time. The amount of time reduction compared to the naive bound (7) is a fraction of τ_1 , given by the ratio of $\Phi_{n_1/2}(P_0)$ to $\Phi(P_1) = \Phi_{n_1/2}(P_1)$. If we consider $\Phi_{n_1/2}(P_0)$ as the speed at which dissemination occurs before the switching time τ_1 , then $\tau_1 \Phi_{n_1/2}(P_0)$ (up to a multiplicative constant) is approximately the number of informed nodes by time τ_1 , and $\tau_1 \Phi_{n_1/2}(P_0)/\Phi(P_1)$ is the time taken to inform these nodes if the connectivity matrix is P_1 throughout the dissemination process. We will see that

factors similar to $\Phi_{n_1/2}(P_0)/\Phi(P_1)$ play an important role in characterizing the ϵ -dissemination time for the other scenarios we study in this paper.

B. Node Removal

In this section, we consider the case where $n_1 < n_0$, and at time τ_1 , $n_0 - n_1$ nodes are removed from the network. We aim to find an upper bound for the ϵ -dissemination time when all the remaining n_1 nodes become informed.

We assume that the number of informed nodes that are removed at time τ_1 is upper bounded by $w \leq n_0 - n_1$, and at least one informed node remains at time τ_1 . It can be shown that $\mathcal{T}_1(\epsilon)$ increases with the number of informed nodes removed at τ_1 . To obtain an upper bound on $\mathcal{T}_1(\epsilon)$, we assume in our analysis, the worst case where w informed nodes are removed. Also, if all the nodes are informed before τ_1 , then dissemination is complete. Thus, we assume $\tau_1 < \lambda_{0,0}(P_0) + \zeta_{n_0}(P_0)$.

Let $S^0(t)$ be the set of informed nodes before removal and $S^1(t)$ be the set of informed nodes after time τ_1 , such that $|S^0(t)| \leq w + |S^1(t)|$. Let $L = \inf\{t : |S^1(t)| > n_1/2\}$. We have the following result whose proof is provided in the Appendix, Section VIII-D.

Theorem 2: Suppose that $\tau_1 < \lambda_{0,0}(P_0) + \zeta_{n_0}(P_0)$ and at most w informed nodes are removed at time τ_1 , with at least one informed node remaining in the network. Let $s = w + \frac{n_1}{2}$ and

$$\lambda = \frac{8}{\Phi_s(P_0)} \log(2n_0 n_1 \epsilon^{-1}).$$

If $\tau_1 < \lambda$, then

$$\mathcal{T}_1(\epsilon) = O\left(\frac{8 \log n_0 + 18 \log n_1 + 9 \log(2\epsilon^{-1})}{\Phi(P_1)} + \tau_1 \left(1 - \frac{\Phi_s(P_0)}{\Phi(P_1)}\right)\right),$$

otherwise

$$\mathcal{T}_1(\epsilon) = O\left(\tau_1 + \frac{\log n_1 + \log \epsilon^{-1}}{\Phi(P_1)}\right).$$

The second time bound in Theorem 2 is the naive upper bound. The first time bound on the other hand, depends on the s -conductance of P_0 and the conductance of P_1 , much like the theoretical time bound in Theorem 1(iii). The effect of the total number of informed nodes removed is seen in the quantity s .

In this section, we assume that the number of informed nodes, w , removed at time τ_1 is constant. A possible direction of future work is to study the generalization where w is a random variable.

IV. MULTIPLE SWITCHING TIMES

In this section, we generalize our results to multiple switching times. Consider a system that undergoes $l > 1$ graph switches at times $\tau_1, \tau_2, \dots, \tau_l$. For $i = 0, 1, \dots, l$, the system has node set $\mathcal{V}_i$ with $|\mathcal{V}_i| = n_i$, and connectivity matrix P_i , which is doubly stochastic, in the time interval $[\tau_i, \tau_{i+1})$. We

consider the case where the number of nodes at each switching time is a nondecreasing sequence, i.e., $n_0 \leq n_1 \leq \dots \leq n_l$. We restrict the node change at each switching time to only node addition, that is, nodes can be added at switching times but never removed. This means no node leaves the network. The reason for this restriction is that if nodes are removed, some of these nodes could have already been informed with the message. This greatly increases the complexity of the analysis involved. It is nonetheless a rich area of future research in dynamic systems. Thus $\mathcal{V}_i \subseteq \mathcal{V}_{i+1}$ for all relevant i . In the following, we first provide an upper bound for the ϵ -dissemination time in the case where $n_0 \geq n_l/2$. The proof is in the Appendix, Section VIII-E. Similar to Theorem 1(iii), we include the adjustments $(\tau_{j+1} - \tau_j)\Phi_{n_l/2}(P_j)/\Phi(P_l)$, where $j = 0, 1, \dots, l-1$, to the upper bound that is obtained if there are no topology changes and the connectivity matrix is P_l throughout the dissemination process. Each adjustment is approximately the amount of time required to disseminate information to those nodes that were informed in the time period $[\tau_j, \tau_{j+1})$, if the connectivity matrix is P_l .

Proposition 2: Suppose that $2 \leq n_0 \leq n_1 \leq \dots \leq n_l$, and $n_0 \geq n_l/2$. We have

$$\mathcal{T}_l(\epsilon) \leq \lambda_{l,l}(P_l) + \zeta_{n_l}(P_l) + \sum_{j=0}^{l-1} (\tau_{j+1} - \tau_j) \left(1 - \frac{\Phi_{n_l/2}(P_j)}{\Phi_{n_l/2}(P_l)}\right).$$

The following is the main result in this section, and is a generalization of Theorem 1. The proof is provided in the Appendix, Section VIII-F.

Theorem 3: Suppose that $n_0 \leq n_1 \leq \dots \leq n_l$. Let

$$y = \min\{i : n_i \geq n_l/2\}.$$

Then,

$$\mathcal{T}_l(\epsilon) = O\left(\tau_y + \frac{(8l+18)\log n_l + 9\log(2\epsilon^{-1})}{\Phi(P_l)} + \sum_{j=y}^{l-1} (\tau_{j+1} - \tau_j) \left(1 - \frac{\Phi_{n_l/2}(P_j)}{\Phi(P_l)}\right)\right). \quad (8)$$

If for a particular time τ_i , the number of nodes $n_i \leq n_l/2$, then many more nodes have joined the network by the last switching time τ_l . Therefore, the total number of informed nodes up to time τ_i is insignificant, and we obtain the upper bound for the ϵ -dissemination time by restarting the dissemination process at time τ_y , the first time the graph has comparable number of nodes with the final network. The upper bound in (8) accounts for the number of newly informed nodes in each time interval $[\tau_j, \tau_{j+1})$, where $j \geq y$.

In the following, we present the special case where $n_0 = n_1 = \dots = n_l = n$. In this case, the dissemination may finish before some of the switching times. Those switching times that are greater than the completion time of the algorithm do not affect $\mathcal{T}_l(\epsilon)$, therefore it is natural to exclude them. Note that ξ_i as defined in the following result is an upper bound for the ϵ -dissemination time if the number of switching times is i (cf. Proposition 2).

Corollary 1: Suppose that $n_0 = n_1 = \dots = n_l = n$. For $i = 0, 1, \dots, l$, let

$$\xi_i = \lambda_{i,i}(P_i) + \zeta_n(P_i) + \sum_{j=0}^{i-1} (\tau_{j+1} - \tau_j) \left(1 - \frac{\Phi(P_j)}{\Phi(P_i)}\right).$$

Let $x = \min\{i : \tau_{i+1} \geq \xi_i\}$. Then, we have

$$\mathbb{P}(T \leq \xi_x) > 1 - \epsilon. \quad (9)$$

V. DISSEMINATION WITH ERASURE

In this section, we apply the results from the preceding sections to determine time bounds for information dissemination in a wireless network that goes through periods of channel fading. In a wireless network, sending a message or a signal from a transmitter node to a receiver node is subject to communication errors due to noise and interference. In this section, we investigate the case where message sending fails with some probability for each link in the network.

We first consider the case where the network graph $\mathcal{G}$ remains unchanged. Suppose $\mathcal{G}$ has node set $\mathcal{V} = \{1, \dots, n\}$ and connectivity matrix P . At every time step, the probability that node i chooses to contact one of its neighbors, node j , is P_{ij} . We assume that the communication link from node i to node j fails with probability $1 - Q_{ij}$, independent of all other links in the network. We call the matrix $Q = (Q_{ij})$ the success probability matrix.

Let $P \cdot Q$ denote the element-wise product of the matrices P and Q . The following result shows that the dissemination time is dependent on $\Phi(P \cdot Q)$. If there are no link failures, $Q = \mathbf{1}\mathbf{1}^T$, where $\mathbf{1}$ is an $n \times 1$ vector of all 1's, and the dissemination time reduces to (1). The proof is almost the same as the proof of Lemma 4 in [10], and is omitted here.

Theorem 4: Let P be an irreducible, doubly stochastic connectivity matrix, and Q be the success probability matrix. Then, we have

$$\mathcal{T}_0(\epsilon) = O\left(\frac{\log n + \log \epsilon^{-1}}{\Phi(P \cdot Q)}\right).$$

In particular, if $Q_{ij} = q$ for all i and j , we have

$$\mathcal{T}_0(\epsilon) = O\left(\frac{\log n + \log \epsilon^{-1}}{q\Phi(P)}\right).$$

Consider now a network that undergoes a number of fading periods. Similar to Section IV, let $\tau_1 < \tau_2 < \dots < \tau_l$ be switching times. For $i = 0, 1, \dots, l$, let P_i and Q_i be the connectivity matrix and success probability matrix respectively in the time interval $[\tau_i, \tau_{i+1})$, where $\tau_0 = 0$ and $\tau_{l+1} = \infty$. Suppose that the number of nodes in the network is n , and we assume that no nodes join or leave the network at all switching times. We use Theorem 4 and Corollary 1 to get an upper bound for the ϵ -dissemination time for this network.

Corollary 2: For any given switching times $0 = \tau_0 < \tau_1 < \tau_2 < \dots < \tau_l < \tau_{l+1} = \infty$, doubly stochastic connectivity matrices $P_0, \dots, P_l$, and success probability - matrices $Q_0, \dots, Q_l$, let

$$\xi'_i = \lambda_{i,i}(P_i \cdot Q_i) + \zeta_n(P_i \cdot Q_i) + \sum_{j=0}^{i-1} (\tau_{j+1} - \tau_j) \left(1 - \frac{\Phi(P_j \cdot Q_j)}{\Phi(P_i \cdot Q_i)}\right),$$

and

$$x = \min\{i : \tau_{i+1} \geq \xi'_i\}.$$

Then,

$$\begin{aligned} \mathcal{T}_l(\epsilon) = & O\left(\frac{(8l+18)\log n + 9\log(2\epsilon^{-1})}{\Phi(P_x \cdot Q_x)}\right. \\ & \left. + \sum_{j=0}^{x-1} (\tau_{j+1} - \tau_j) \left(1 - \frac{\Phi(P_j \cdot Q_j)}{\Phi(P_x \cdot Q_x)}\right)\right). \end{aligned}$$

From Corollary 2, we have the following result for a network that suffers r distinct periods of erasure, where the success probability matrix in a non-erasure period is $\mathbf{1}\mathbf{1}^T$, while that within an erasure period is Q . Let the connectivity matrix with no erasure be P . We assume that the network has no erasure at time 0. Then, for $i = 1, \dots, r$, we let τ_{2i-1} and τ_{2i} be the start and end times respectively for the i th erasure period, so that there are a total of $2r$ switching times. For $i = 1, \dots, r$, let

$$\begin{aligned} \xi''_{2i-1} = & \lambda_{2i-1, 2i-1}(P \cdot Q) + \zeta_n(P \cdot Q) \\ & + \sum_{j=0}^{i-1} (\tau_{2j+1} - \tau_{2j}) \left(1 - \frac{\Phi(P)}{\Phi(P \cdot Q)}\right), \end{aligned}$$

and

$$\begin{aligned} \xi''_{2i} = & \lambda_{2i, 2i}(P) + \zeta_n(P) \\ & + \sum_{j=1}^i (\tau_{2j} - \tau_{2j-1}) \left(1 - \frac{\Phi(P \cdot Q)}{\Phi(P)}\right), \end{aligned}$$

The time ξ''_{2i-1} is an upper bound for the ϵ -dissemination time if there are i erasure periods, with the network communications remaining in erasure after the last switching time, i.e., $\tau_{2i} = \infty$. The time ξ''_{2i} is an upper bound for the ϵ -dissemination time in the case where the network makes a final transition to the non-erasure mode.

The following result characterizes the ϵ -dissemination time when there are r phases of erasure. The proof follows directly from Corollary 2.

Corollary 3: Suppose a wireless network has connectivity matrix P , and goes through r phases of erasure, with success probability matrix Q . Let $x = \min\{i : \tau_{i+1} \geq \xi'_i\}$. Then, we have

$$\mathbb{P}(T \leq \xi''_x) > 1 - \epsilon. \quad (10)$$

In particular, if $\tau_{2r} < \infty$ and $x = 2r$, then the ϵ -dissemination time is

$$\begin{aligned} \mathcal{T}_{2r}(\epsilon) = & O\left(\frac{(16r+18)\log n + 9\log(2\epsilon^{-1})}{\Phi(P)}\right. \\ & \left. + t_e \left(1 - \frac{\Phi(P \cdot Q)}{\Phi(P)}\right)\right), \end{aligned}$$

where $t_e = \sum_{i=1}^r (\tau_{2i} - \tau_{2i-1})$ is the total time the network is in erasure.

Fig. 2. Plot of the ratios $T_{th}/\mathcal{T}_1(\epsilon)$ and $T_b/\mathcal{T}_1(\epsilon)$ versus n .

VI. NUMERICAL EXAMPLE

In this section, we present a numerical example that shows that our theoretical bound is asymptotically tight up to a multiplicative constant. We also compare our theoretical bound with the naïve upper bound (7). We demonstrate through this comparison that our theoretical bound is significantly better asymptotically than the trivial bound which is based off what has been proven for the constant connectivity matrix case.

Our example involves a complete graph with n nodes, with a change in the connectivity matrix at time τ_1 . Suppose that the connectivity matrix before τ_1 is P_0 , with $(P_0)_{ij} = 1/n$ for all i, j . This graph has conductance $1/2$ if n is even and $(n+1)/2n$ if n is odd. At the switching time τ_1 , the connectivity matrix becomes to P_1 , with $(P_1)_{ij} = 1/n^2$ for all $i \neq j$, and $(P_1)_{ii} = 1 - (n-1)/n^2$ for all i . The conductance of P_1 is $1/2n$ if n is even, and $(n+1)/2n^2$ if n is odd.

Fix $\epsilon = 0.01$, and let

$$\tau_1 = \frac{1}{2} \left(\frac{\log n}{\Phi(P_0)} - \frac{3}{4} (\log n)^{\frac{1}{10}} \right).$$

From Theorem 1, we have $\mathcal{T}_1(\epsilon) \leq c_1 T_{th}$, where c_1 is a constant and

$$T_{th} = \tau_1 \left(1 - \frac{\Phi(P_0)}{\Phi(P_1)} \right) + \frac{\log n}{\Phi(P_1)}.$$

On the other hand, the naïve upper bound (7), obtained by restarting the dissemination at time τ_1 , is given by $\mathcal{T}_1(\epsilon) \leq c_2 T_b$, where $T_b = \tau_1 + \log n / \Phi(P_1)$, and c_2 is a constant. We simulate the dissemination process to obtain $\mathcal{T}_1(\epsilon)$, and plot the ratios $T_{th}/\mathcal{T}_1(\epsilon)$ and $T_b/\mathcal{T}_1(\epsilon)$ versus n in Figure 2. The ratio with our theoretical bound approaches a constant, showing that our theoretical bound is asymptotically tight for this example. The ratio with the upper bound T_b diverges as n increases. This example shows that it is important to account for the dissemination before the switching time, and that our theoretical bound is asymptotically better than the naïve upper bound.

VII. CONCLUSION

We have analyzed the dissemination time of random gossiping for dynamic network topologies. We established order bounds for the dissemination time when the connectivity matrix undergoes changes at known times. We showed that an upper bound can be found by using the same bound for the dissemination time in a network where the dissemination is restarted at the last switching time, but with adjustments added. In a network with l switching times $\tau_1 < \dots < \tau_l$, the adjustments are $(\tau_{j+1} - \tau_j)\Phi_{n_1/2}(P_j)/\Phi(P_l)$, where $j = 1, \dots, l-1$. Each adjustment is approximately the amount of time required to disseminate information to those nodes that were informed in the time period $[\tau_j, \tau_{j+1})$, if the connectivity matrix is P_l throughout the dissemination process. We also derived an upper bound for the dissemination time when communications between nodes are subject to failures. The current work can be extended in several directions. We have assumed that the connectivity matrices and switching times are known a priori, which is often not the case in a wireless mobile environment. It is of interest to study the expected dissemination time when only the distributional statistics of the connectivity matrix and switching times are available. Our assumption of erasure channels between nodes can also be too restrictive. Part of our future research will investigate the impact of different channel models on the randomized gossiping algorithm.

VIII. APPENDIX

A. Proof of Lemma 1

Proof: Suppose that $L_k > t_0$. Set $L'_k = L_k - t_0$ and $t' = t - t_0$. We define the random variable $Z(t)$ to be

$$Z(t) = \frac{\exp\left(\frac{\Phi_k(P)t}{8}\right)}{|S(t)|}.$$

Since $|S(t)| \leq n$ for all $t \in [t_0, t_1]$, with probability one, we have

$$Z(L'_k \wedge t') \geq \frac{1}{n} \exp\left(\frac{\Phi_k(P)}{8}(L'_k \wedge t')\right)$$

Set $\alpha_k = \exp[\Phi_k(P)(L'_k \wedge t')/8]$. Conditioning on $L_k > t_0$, and taking expectation on both sides, we obtain

$$\mathbb{E}[\alpha_k \mid L_k > t_0] \leq n\mathbb{E}[Z(L'_k \wedge t') \mid L_k > t_0]. \quad (11)$$

Using Proposition 1(iii), we have

$$\begin{aligned} \mathbb{E}[\alpha_k \mid L_k > t_0] &\leq n\mathbb{E}[Z(L'_k \wedge 0) \mid L_k > t_0] \\ &= n\mathbb{E}\left[\frac{1}{|S(t_0)|} \mid L_k > t_0\right] \\ &\leq n. \end{aligned}$$

Applying Markov's inequality, we have

$$\begin{aligned} \mathbb{P}(L_k > t \mid L_k > t_0) &\leq \mathbb{P}(L_k \wedge t \geq t \mid L_k > t_0) \\ &\leq n \exp\left(\frac{\Phi_k(P)(t_0 - t)}{8}\right). \end{aligned}$$

□

B. Proof of Theorem 1(ii)

Proof: If $n_0 = n_1$, then the result trivially holds. In the following, we assume $n_0 < n_1$. Let $\delta = n_1 - n_0$ be the number of new nodes that join the network at time τ_1 . In the following, we condition on the event $A = \{|S(\tau_1)| = n_0\}$. For all $t \geq \tau_1$, we have $|S(t)| \geq n_1/2$, and $|S(t)^c| \leq \delta$. Proposition 1(iv) then yields

$$\mathbb{E}[|S(t)^c| \mid S(t-1), A] \leq |S(t-1)^c|(1 - \Phi_{n_1-n_0}(P_1)). \quad (12)$$

Taking expectations on both sides and recursively applying (12), we obtain

$$\begin{aligned} \mathbb{E}[|S(t)^c| \mid A] &\leq (1 - \Phi_{n_1-n_0}(P_1))^{t-\tau_1} \mathbb{E}[|S(\tau_1)^c| \mid A] \\ &\leq \delta \exp(-\Phi_{n_1-n_0}(P_1)(t - \tau_1)), \end{aligned} \quad (13)$$

since $|S(\tau_1)^c| \leq \delta$, conditioned on A . Applying Markov's inequality, we have

$$\mathbb{P}(|S(t)^c| \geq 1 \mid A) \leq \delta \exp(-\Phi_{n_1-n_0}(P_1)(t - \tau_1)). \quad (14)$$

Substituting $t = \tau_1 + (2 \log(\delta + 1) + \log \epsilon^{-1})/\Phi_{n_1-n_0}(P_1)$ into (14), we obtain

$$\mathbb{P}\left(T \leq \tau_1 + \frac{2 \log(\delta + 1) + \log \epsilon^{-1}}{\Phi_{n_1-n_0}(P_1)} \mid A\right) \geq 1 - \frac{\epsilon}{(\delta + 1)^2},$$

from which we have

$$\begin{aligned} &\mathbb{P}\left(T \leq \tau_1 + \frac{2 \log(\delta + 1) + \log \epsilon^{-1}}{\Phi_{n_1-n_0}(P_1)}\right) \\ &\geq \mathbb{P}\left(T \leq \tau_1 + \frac{2 \log(\delta + 1) + \log \epsilon^{-1}}{\Phi_{n_1-n_0}(P_1)} \mid A\right) P(A) \\ &\geq \left(1 - \frac{\epsilon}{(\delta + 1)^2}\right) \left(1 - \frac{\epsilon}{n_0}\right) \\ &\geq 1 - \left(\frac{1}{(\delta + 1)^2} + \frac{1}{n_0}\right) \epsilon \\ &\geq 1 - \epsilon, \end{aligned} \quad (15)$$

where the second inequality follows from Proposition 1(i), and the last inequality follows because $\delta \geq 1$ and $n_0 \geq 2$. The proof is now complete. □

C. Proof of Theorem 1(iii)

Proof: Our proof methodology is similar to that of Lemma 4 in [10], but is more involved because of the change in the connectivity matrix at τ_1 . A full proof is presented here for completeness. It suffices to show that $\mathcal{T}_1(\epsilon) \leq \xi$, where

$$\xi = \lambda_{1,1}(P_1) + \zeta_{n_1}(P_1) + \tau_1 \left(1 - \frac{\Phi_{n_1/2}(P_0)}{\Phi(P_1)}\right),$$

since $\lambda_{1,1}(P_1) + \zeta_{n_1}(P_1) \leq (26 \log n_1 + 9 \log(2\epsilon^{-1}))/\Phi(P_1)$.

We divide the proof into the two cases where $\tau_1 < \lambda_{0,1}(P_0)$ and $\tau_1 \geq \lambda_{0,1}(P_0)$.

Suppose $\tau_1 < \lambda_{0,1}(P_0)$. Recall that $L = \inf\{t : |S(t)| > n_1/2\} = L_{n_1/2}$. From Proposition 1(i), we have

$$\mathbb{P}(L > \tau_1) \leq n_0 \exp\left(-\frac{\Phi_{n_1/2}(P_0)\tau_1}{8}\right). \quad (16)$$

Using Lemma 1 with $t_0 = \tau_1$, $\mathcal{V} = \mathcal{V}_1$, $P = P_1$, and $k = n_1/2$, we get

$$\mathbb{P}(L > t \mid L > \tau_1) \leq n_1 \exp\left(\frac{\Phi(P_1)(\tau_1 - t)}{8}\right).$$

Let $t_1 = \lambda_{1,1}(P_1) + (\Phi(P_1) - \Phi_{n_1/2}(P_0))\tau_1/\Phi(P_1)$. Notice that $\mathbb{P}(L > t_1 \mid L \leq \tau_1) = 0$ since $t_1 \geq \tau_1$. Thus, together with (16), we have

$$\begin{aligned} & \mathbb{P}(L > t_1) \\ &= \mathbb{P}(L > t_1 \mid L > \tau_1)\mathbb{P}(L > \tau_1) \\ &\leq n_0 n_1 \exp\left(\frac{\Phi(P_1)(\tau_1 - t_1)}{8}\right) \exp\left(-\frac{\Phi_{n_1/2}(P_0)\tau_1}{8}\right) \\ &= \frac{\epsilon}{2n_1}. \end{aligned} \quad (17)$$

From Proposition 1(v), we have

$$\mathbb{P}(M(t_1) \leq \zeta_{n_1}(P_1) \mid L \leq t_1) \geq 1 - \frac{\epsilon}{2n_1}. \quad (18)$$

From both (17) and (18), we obtain

$$\begin{aligned} \mathbb{P}(T \leq t_1 + \zeta_{n_1}(P_1)) &\geq \mathbb{P}(L \leq t_1, M(t_1) \leq \zeta_{n_1}(P_1)) \\ &\geq \left(1 - \frac{\epsilon}{2n_1}\right)^2 \\ &> 1 - \frac{\epsilon}{n_1}, \end{aligned}$$

from which $\mathcal{T}_1(\epsilon) \leq \xi$ follows.

We now suppose $\tau_1 \geq \lambda_{0,1}(P_0)$. From Proposition 1(i), we have

$$\mathbb{P}(L \leq \lambda_{0,1}(P_0)) \geq 1 - \epsilon/2n_1.$$

We now condition on the event that $L \leq \lambda_{0,1}(P_0)$. We reset the clock tick to start at zero at $\lambda_{0,1}(P_0)$. Let $\tilde{S}(t) = S(t + \lambda_{0,1}(P_0))$. We have $|\tilde{S}(0)^c| \leq n_1/2$. Let $\tau' = \tau_1 - \lambda_{0,1}(P_0)$. Thus in the new time frame, the switching time is τ' . Since $|\tilde{S}(0)^c| \leq n_1/2$, we obtain from Proposition 1(iv), for all $t < \tau'$,

$$\mathbb{E}\left[|\tilde{S}(t+1)^c| \mid \tilde{S}(t)\right] \leq |\tilde{S}(t)^c|(1 - \Phi_{n_1/2}(P_0)), \quad (19)$$

and for $t \geq \tau'$,

$$\mathbb{E}\left[|\tilde{S}(t+1)^c| \mid \tilde{S}(t)\right] \leq |\tilde{S}(t)^c|(1 - \Phi(P_1)), \quad (20)$$

Applying inequalities (19) and (20) recursively, for any $t > \tau'$, we have

$$\begin{aligned} & \mathbb{E}\left[|\tilde{S}(t)^c| \mid L \leq \lambda_{0,1}(P_0)\right] \\ &\leq \mathbb{E}\left[\mathbb{E}\left[|\tilde{S}(t)^c| \mid \tilde{S}(t-1)^c, L \leq \lambda_{0,1}(P_0)\right] \mid L \leq \lambda_{0,1}(P_0)\right] \\ &\leq (1 - \Phi(P_1))\mathbb{E}\left[|\tilde{S}(t-1)^c| \mid L \leq \lambda_{0,1}(P_0)\right] \\ &\leq (1 - \Phi(P_1))^{t-\tau'}(1 - \Phi_{n_1/2}(P_0))^{\tau'} \\ &\quad \cdot \mathbb{E}\left[|\tilde{S}(0)^c| \mid L \leq \lambda_{0,1}(P_0)\right] \\ &\leq \exp(-(\Phi(P_1)(t - \tau') + \Phi_{n_1/2}(P_0)\tau')) \left(\frac{n_1}{2}\right). \end{aligned} \quad (21)$$

Let $t_2 = \zeta_{n_1}(P_1) + \tau'(\Phi(P_1) - \Phi_{n_1/2}(P_0))/\Phi(P_1)$. Then, since $\tau_1 < \lambda_{0,0}(P_0) + \zeta_{n_0}(P_0)$, we have $\tau' < \zeta_{n_0}(P_0)$. This implies that $t_2 > \tau'$, and substituting $t = t_2$ in (21), we get

$$\mathbb{E}\left[|\tilde{S}(t_2)^c| \mid L \leq \lambda_{0,1}(P_0)\right] \leq \frac{\epsilon}{2n_1}.$$

Using Markov's Inequality, we have

$$\mathbb{P}\left(|\tilde{S}(t_2)^c| > 0 \mid L \leq \lambda_{0,1}(P_0)\right) \leq \frac{\epsilon}{2n_1}.$$

Therefore in the original time frame, we have

$$\mathbb{P}(M(\lambda_{0,1}(P_0)) \leq t_2 \mid L \leq \lambda_{0,1}(P_0)) > 1 - \frac{\epsilon}{2n_1},$$

which yields

$$\begin{aligned} & \mathbb{P}(T \leq t_2 + \lambda_{0,1}(P_0)) \\ &\geq \mathbb{P}(M(\lambda_{0,1}(P_0)) \leq t_2 \mid L \leq \lambda_{0,1}(P_0))\mathbb{P}(L \leq \lambda_{0,1}(P_0)) \\ &\geq \left(1 - \frac{\epsilon}{2n_1}\right)^2 > 1 - \frac{\epsilon}{n_1}. \end{aligned}$$

Since $t_2 + \lambda_{0,1}(P_0) \leq \xi$, we have $\mathcal{T}_1(\epsilon) \leq \xi$. The theorem is now proved. $\square$

D. Proof of Theorem 2

Proof: Suppose $\tau_1 < \lambda$. It suffices to show that $\mathcal{T}_1(\epsilon) \leq \xi'$, where

$$\xi' = \lambda_{1,1}(P_1) + \zeta_{n_1}(P_1) + \tau_1 \left(1 - \frac{\Phi_s(P_0)}{\Phi(P_1)}\right).$$

From Proposition 1(i), we have

$$\mathbb{P}(L > \tau_1) \leq \mathbb{P}(L_s > \tau_1) \leq n_0 \exp\left(-\frac{\Phi_s(P_0)\tau_1}{8}\right). \quad (22)$$

From Lemma 1 with $t_0 = \tau_1$, $\mathcal{V} = \mathcal{V}_1$, $P = P_1$, and $k = n_1/2$, we have for $t \geq \tau_1$,

$$\mathbb{P}(L > t \mid L > \tau_1) \leq n_1 \exp\left(\frac{\Phi(P_1)(\tau_1 - t)}{8}\right).$$

Let $t'_1 = \lambda_{1,1}(P_1) + \tau_1(1 - \Phi_s(P_0))/\Phi(P_1)$. Since $t'_1 \geq \tau_1$, from (22), we obtain

$$\begin{aligned} & \mathbb{P}(L > t'_1) \\ &= \mathbb{P}(L > t'_1 \mid L > \tau_1)\mathbb{P}(L > \tau_1) \\ &\leq n_0 n_1 \exp\left(\frac{\Phi(P_1)(\tau_1 - t'_1)}{8}\right) \exp\left(-\frac{\Phi_s(P_0)\tau_1}{8}\right) \\ &= \frac{\epsilon}{2n_1}. \end{aligned} \quad (23)$$

From Proposition 1(v), we have

$$\mathbb{P}(M(t'_1) \leq \zeta_{n_1}(P_1) \mid L \leq t'_1) \geq 1 - \frac{\epsilon}{2n_0}. \quad (24)$$

From both (23) and (24), we obtain

$$\begin{aligned} \mathbb{P}(T \leq t'_1 + \zeta_{n_1}(P_1)) &\geq \mathbb{P}(L \leq t'_1, M(t'_1) \leq \zeta_{n_1}(P_1)) \\ &\geq \left(1 - \frac{\epsilon}{2n_1}\right)^2 \\ &> 1 - \frac{\epsilon}{n_1}, \end{aligned}$$

from which $\mathcal{T}_1(\epsilon) \leq \xi'$ follows.

We now suppose $\tau_1 \geq \lambda$. From Proposition 1(i), we have

$$\mathbb{P}(L \leq \lambda) \geq 1 - \epsilon/2n_1.$$

From Proposition 1(v), we also have

$$\mathbb{P}(M(\tau_1) \leq \zeta_{n_1}(P_1) \mid L \leq \lambda \leq \tau_1) \geq 1 - \epsilon/2n_1.$$

Therefore,

$$\begin{aligned} & \mathbb{P}(T \leq \tau_1 + \zeta_{n_1}(P_1)) \\ & \geq \mathbb{P}(M(\tau_1) \leq \zeta_{n_1}(P_1) \mid L < \lambda) \mathbb{P}(L < \lambda) \\ & \geq \left(1 - \frac{\epsilon}{2n_1}\right)^2 > 1 - \frac{\epsilon}{n_1}. \end{aligned}$$

The proof is now complete. $\square$

E. Proof of Proposition 2

For $i = 0, 1, \dots, l$, let $a_i = \tau_{i+1} - \tau_i$, and let

$$\chi_i = \lambda_{i,l}(P_i) + \sum_{j=0}^{i-1} a_j \left(1 - \frac{\Phi_{n_l/2}(P_j)}{\Phi_{n_l/2}(P_i)}\right).$$

The last summation term is zero for $i = 0$. We use the following two lemmas in our proofs. The first lemma is an extension of Proposition 1(i). Recall that $L = \inf\{t : |S(t)| > n_l/2\} = L_{n_l/2}$.

Lemma 2: Suppose $n_0 \geq n_l/2$. For all $t \in (\tau_i, \tau_{i+1}]$, we have

$$\begin{aligned} \mathbb{P}(L > t) & \leq n_i \exp\left(\frac{\Phi_{n_l/2}(P_i)}{8}(\tau_i - t)\right) \\ & \cdot \prod_{j=0}^{i-1} n_j \exp\left(-\frac{\Phi_{n_l/2}(P_j)}{8}a_j\right), \end{aligned} \quad (25)$$

where the product term is one if $i = 0$. In particular,

$$\mathbb{P}(L > \tau_{i+1}) \leq \prod_{j=0}^i n_j \exp\left(-\frac{\Phi_{n_l/2}(P_j)}{8}a_j\right). \quad (26)$$

Proof: We proceed by induction on i . Suppose $i = 1$. From Proposition 1(i), for $t \leq \tau_1$, we have

$$\mathbb{P}(L > t) \leq n_0 \exp\left(-\frac{\Phi_{n_l/2}(P_0)t}{8}\right),$$

which proves the claim for $i = 1$. Suppose the hypothesis holds for i . Let $t \in (\tau_i, \tau_{i+1}]$. From the induction hypothesis, we obtain

$$\mathbb{P}(L > \tau_i) \leq \prod_{j=0}^{i-1} n_j \exp\left(-\frac{\Phi_{n_l/2}(P_j)}{8}a_j\right). \quad (27)$$

Using Lemma 1, with $t_0 = \tau_i$, $t_1 = \tau_i + 1$, $\mathcal{V} = \mathcal{V}_i$, $P = P_i$ and $k = n_l/2$, we get

$$\mathbb{P}(L > t \mid L > \tau_i) \leq n_i \exp\left(\frac{\Phi_{n_l/2}(P_i)(\tau_i - t)}{8}\right). \quad (28)$$

From (27) and (28), we have

$$\begin{aligned} \mathbb{P}(L > t) & = \mathbb{P}(L > t \mid L > \tau_i) \mathbb{P}(L > \tau_i) \\ & \leq n_i \exp\left(\frac{\Phi_{n_l/2}(P_i)}{8}(\tau_i - t)\right) \\ & \cdot \prod_{j=0}^{i-1} n_j \exp\left(-\frac{\Phi_{n_l/2}(P_j)}{8}a_j\right), \end{aligned}$$

and the induction proof is complete. $\square$

The following lemma gives a high probability bound for L .

Lemma 3: Suppose $n_0 \geq n_l/2$. Let

$$r = \min\{i : \tau_{i+1} \geq \chi_i\}$$

Then

$$\mathbb{P}(L \leq \chi_r) > 1 - \frac{\epsilon}{2n_l}. \quad (29)$$

Proof: Since $\chi_{r-1} > \tau_r$ and $\chi_r \leq \tau_{r+1}$, we have

$$\begin{aligned} \chi_r & = \lambda_{r,l}(P_r) + \sum_{j=0}^{r-1} a_j \left(1 - \frac{\Phi_{n_l/2}(P_j)}{\Phi_{n_l/2}(P_r)}\right) \\ & = \lambda_{r,l}(P_r) + \tau_r - \sum_{j=0}^{r-1} \frac{a_j \Phi_{n_l/2}(P_j)}{\Phi_{n_l/2}(P_r)} \\ & = \tau_r + \lambda_{r,l}(P_r) - \tau_r \frac{\Phi_{n_l/2}(P_{r-1})}{\Phi_{n_l/2}(P_r)} \\ & \quad + \sum_{j=0}^{r-2} a_j \frac{\Phi_{n_l/2}(P_{r-1}) - \Phi_{n_l/2}(P_j)}{\Phi_{n_l/2}(P_r)} \\ & > \tau_r + \frac{\Phi_{n_l/2}(P_{r-1})}{\Phi_{n_l/2}(P_r)} \left[\lambda_{r-1,l}(P_{r-1}) \right. \\ & \quad \left. + \sum_{j=0}^{r-2} a_j \left(1 - \frac{\Phi_{n_l/2}(P_j)}{\Phi_{n_l/2}(P_{r-1})}\right) - \tau_r \right] \\ & > \tau_r. \end{aligned}$$

where the second last inequality follows from $\lambda_{r,l}(P_r) > \lambda_{r-1,l}(P_{r-1})\Phi_{n_l/2}(P_{r-1})/\Phi_{n_l/2}(P_r)$. Since $\tau_r < \chi_r \leq \tau_{r+1}$, we obtain from Lemma 2

$$\begin{aligned} \mathbb{P}(L > \chi_r) & \leq n_r \exp\left(\frac{\Phi_{n_l/2}(P_r)}{8}(\tau_r - \chi_r)\right) \\ & \cdot \prod_{j=0}^{r-1} n_j \exp\left(-\frac{\Phi_{n_l/2}(P_j)}{8}a_j\right) \\ & = \frac{\epsilon}{2n_l}, \end{aligned}$$

and the proof is complete. $\square$

We are now ready to prove the proposition.

Proof of Proposition 2: Let $r = \min\{i : \tau_{i+1} \geq \chi_i\}$. Suppose $r = l$. Then, the connectivity matrix of the graph is P_l for all $t \geq \chi_l$. Using Proposition 1(v) with P replaced by P_l , we get

$$\mathbb{P}(M(\chi_l) \leq \zeta_{n_l}(P_l) \mid L \leq \chi_l) \geq 1 - \frac{\epsilon}{2n_l}. \quad (30)$$

Combining (30) with Lemma 3, we have

$$\begin{aligned} & \mathbb{P}(M(\chi_l) \leq \zeta_{n_l}(P_l)) \\ & \geq \mathbb{P}(M(\chi_l) \leq \zeta_{n_l}(P_l) \mid L \leq \chi_l) \mathbb{P}(L \leq \chi_l) \\ & > 1 - \frac{\epsilon}{n_l}. \end{aligned}$$

Therefore, it follows that

$$\mathbb{P}(T \leq \chi_l + \zeta_{n_l}(P_l)) > 1 - \epsilon.$$

Suppose now that $r < l$. We condition on the event that $L \leq \chi_r$, and reset the clock tick after χ_r to be zero. Let $\tilde{S}(t) = S(t + \chi_r)$. We have $|\tilde{S}(0)^c| \leq n_l/2$. Let $b = \tau_{r+1} - \chi_r$, and $d = n_l/2$. Repeatedly applying Proposition 1(iv), we have, for all $t \geq \tau_l - \chi_r$,

$$\begin{aligned} & \mathbb{E} \left[|\tilde{S}(t)^c| \mid L \leq \chi_r \right] \\ & \leq (1 - \Phi_d(P_l)) \mathbb{E} \left[|\tilde{S}(t-1)^c| \mid L \leq \chi_r \right] \\ & \leq (1 - \Phi_d(P_l))^{t - (\tau_l - \chi_r)} (1 - \Phi_d(P_r))^b \\ & \quad \cdot \prod_{j=r+1}^{l-1} (1 - \Phi_d(P_j))^{a_j} \mathbb{E} \left[|\tilde{S}(0)^c| \mid L \leq \chi_r \right] \\ & \leq d \cdot \exp \left(-\Phi_d(P_l)(t - \tau_l + \chi_r) - \Phi_d(P_r)b \right. \\ & \quad \left. - \sum_{j=r+1}^{l-1} \Phi_d(P_j)a_j \right) \\ & = d \cdot \exp \left(-\Phi_d(P_l)t + (\Phi_d(P_l) - \Phi_d(P_r))b \right. \\ & \quad \left. + \sum_{j=r+1}^{l-1} (\Phi_d(P_l) - \Phi_d(P_j))a_j \right). \end{aligned} \quad (31)$$

Let

$$\begin{aligned} \theta_{r,P_l} &= \frac{2 \log n_l + \log \epsilon^{-1}}{\Phi_d(P_l)} + b \left(1 - \frac{\Phi_d(P_r)}{\Phi_d(P_l)} \right) \\ & \quad + \sum_{j=r+1}^{l-1} a_j \left(1 - \frac{\Phi_d(P_j)}{\Phi_d(P_l)} \right). \end{aligned}$$

We have from (31),

$$\mathbb{E} \left[|\tilde{S}(\theta_{r,P_l})^c| \mid L \leq \chi_r \right] \leq \frac{\epsilon}{2n_l}.$$

From Markov's Inequality, we obtain

$$\mathbb{P}(M(\chi_r) \leq \theta_{r,P_l} \mid L \leq \chi_r) \geq 1 - \frac{\epsilon}{2n_l}. \quad (32)$$

Combining (32) and (29), we have

$$\mathbb{P}(M(\chi_r) \leq \theta_{r,P_l}) \geq 1 - \frac{\epsilon}{n_l},$$

and so

$$\mathbb{P}(T \leq \chi_r + \theta_{r,P_l}) \geq 1 - \epsilon. \quad (33)$$

It remains to show that $\chi_r + \theta_{r,P_l} \leq \chi_l + \zeta_{n_l}(P_l)$. We have

$$\begin{aligned} & \chi_r + \theta_{r,P_l} \\ & = \zeta_{n_l}(P_l) + \tau_{r+1} \left(1 - \frac{\Phi_d(P_r)}{\Phi_d(P_l)} \right) \\ & \quad + \sum_{j=r+1}^{l-1} a_j \left(1 - \frac{\Phi_d(P_j)}{\Phi_d(P_l)} \right) + \chi_r \frac{\Phi_d(P_r)}{\Phi_d(P_l)} \\ & = \lambda_{r,l}(P_l) + \zeta_{n_l}(P_l) + \sum_{j=0}^{l-1} a_j \left(1 - \frac{\Phi_d(P_j)}{\Phi_d(P_l)} \right) \\ & \leq \chi_l + \zeta_{n_l}(P_l), \end{aligned} \quad (34)$$

where the last inequality follows from the fact that $\lambda_{r,l}(P_l) \leq \chi_{l,l}(P_l)$. Therefore, from (33), we get

$$\mathbb{P}(T \leq \chi_l + \zeta_{n_l}(P_l)) > 1 - \epsilon.$$

The proof is now complete. $\square$

F. Proof of Theorem 3

Proof: For all $i < y$, $n_i < n_l/2$. Therefore, the dissemination that happens before τ_y is negligible, and we restart the dissemination process at τ_y . Considering only switching times larger than τ_y , the result follows from Proposition 2. $\square$

REFERENCES

- [1] L. Xiao, S. Boyd, and S. Lall, "A scheme for robust distributed sensor fusion based on average consensus," in *Proc. IEEE Inform. Processing in Sensor Networks*, Los Angeles, CA, Apr. 2005.
- [2] W. P. Tay, J. N. Tsitsiklis, and M. Z. Win, "On the impact of node failures and unreliable communications in dense sensor networks," *IEEE Trans. Signal Process.*, vol. 56, no. 6, pp. 2535–2546, Jun. 2008.
- [3] —, "Bayesian detection in bounded height tree networks," *IEEE Trans. Signal Process.*, vol. 57, no. 10, pp. 4042–4051, Oct. 2009.
- [4] T. Q. S. Quek, D. Dardari, and M. Z. Win, "Energy efficiency of dense wireless sensor networks: To cooperate or not to cooperate," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 2, pp. 459–470, Feb. 2007.
- [5] A. Jadbabaie, J. Lin, and S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Trans. Autom. Control*, vol. 48, no. 6, pp. 988–1001, Jun. 2003.
- [6] R. W. Beard and V. Stepanyan, "Synchronization of information in distributed multiple vehicle coordinated control," in *Proc. IEEE Conf. Decision and Control*, Dec. 2003, p. 20292034.
- [7] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Englewood Cliffs, NJ: Prentice Hall, 1989.
- [8] L. Xiao and S. Boyd, "Fast linear iterations for distributed averaging," *Syst. Control Lett.*, vol. 53, no. 1, pp. 65–78, Sep. 2004.
- [9] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Randomized gossip algorithms," *IEEE Trans. Inf. Theory*, vol. 52, no. 6, pp. 2508–2530, Jun. 2006.
- [10] D. Mosk-aoyama and D. Shah, "Fast gossip algorithm for computing separable function," *IEEE Trans. Inf. Theory*, vol. 54, no. 7, pp. 2997–3007, Jul. 2008.
- [11] U. Feige, D. Peleg, P. Raghavan, and E. Upfal, "Randomized broadcast in networks," *Random Structure and Algorithms*, vol. 1, no. 4, pp. 447–460, 1990.
- [12] B. Pittel, "On spreading a rumor," *SIAM J. Appl. Math.*, vol. 47, no. 1, pp. 213–223, Feb. 1987.
- [13] 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 *Proc. ACM Symp. on Principles of Distrib. Comput.*, Vancouver, BC, Canada, Aug. 1987, pp. 1–12.
- [14] R. Karp, C. Schindelhauer, S. Shenker, and B. Vocking, "Randomized rumor spreading," in *Proc. IEEE Symp. Found. Comput. Sci.*, Redondo Beach, CA, Nov. 2000, pp. 565–574.
- [15] D. W. Krumme, G. Cybenko, and K. N. Venkataraman, "Gossiping in minimal time," *SIAM J. Comp.*, vol. 21, no. 1, pp. 111–139, 1992.

- [16] R. Ravi, "Rapid rumor ramification: Approximating the minimum broadcast time," in *Proc. IEEE Symp. Found. Comput. Sci.*, Santa Fe, NM, Nov. 1994, pp. 200–213.
- [17] I. Gaber and Y. Mansour, "Broadcast in radio networks," in *Proc. ACM-SIAM Symp. Discrete Algorithms*, San Francisco, CA, Jan. 1995, pp. 577–585.
- [18] L. Massoulié, A. Twigg, C. Gkantsidis, and P. Rodriguez, "Randomized decentralized broadcasting algorithms," in *Proc. IEEE Joint Conf. of the IEEE Computer and Commun. Societies*, San Francisco, CA, May. 2007, pp. 1073–1081.
- [19] R. Elsasser and T. Sauerwald, "On the runtime and robustness of randomized broadcasting," *Theoretical Comput. Sci.*, vol. 410, no. 36, pp. 3414–3427, Aug. 2009.
- [20] A. M. Frieze and G. R. Grimmett, "The shortest-path problem for graphs with random arc-lengths," *Discrete Appl. Math.*, vol. 10, no. 1, pp. 55–77, Jan. 1985.
- [21] F. Chierichetti, S. Lattanzi, and A. Panconesi, "Rumor spreading in social networks," in *Proc. Int. Colloquium on Automata, Languages and Programming*, Rhodes, Greece, Jul. 2009, pp. 375–386.
- [22] D. Kempe, A. Dobra, and J. Gehrke, "Gossip-based computation of aggregate information," in *Proc. IEEE Symp. Found. Comput. Sci.*, Cambridge, MA, Oct. 2003, pp. 482–491.
- [23] T. Sauerwald, "On mixing and edge expansion properties in randomized broadcasting," *Algorithmica*, vol. 56, no. 1, pp. 51–88, Jan. 2010.
- [24] S. Kar and J. M. F. Moura, "Distributed consensus algorithms in sensor networks with imperfect communication: Link failures and channel noise," *IEEE Trans. Signal Process.*, vol. 57, no. 1, pp. 355–369, Jan. 2009.
- [25] W. Ren and R. Beard, "Consensus seeking in multiagent systems under dynamically changing interaction topologies," *IEEE Trans. Autom. Control*, vol. 50, no. 5, pp. 655–661, May 2005.
- [26] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proc. of IEEE*, vol. 95, pp. 215–233, 2007.
- [27] M. Zhu and S. Martinez, "Dynamic average consensus on synchronous communication networks," in *American Control Conference*, Seattle, WA, Jun. 2008, pp. 4382–4387.
- [28] D. Bauso, L. Giarre, and R. Pesenti, "Distributed consensus in networks of dynamic agents," in *Proc. of the 44th IEEE Conference on Decision and Control*, Seville, Spain, Dec. 2005, pp. 7054–7059.
- [29] Y. Gao and L. Wang, "Consensus of multiple dynamic agents with sampled information," *IET Control Theory and Applications*, vol. 4, pp. 945–956, 2010.
- [30] D. Mosk-aoyama and D. Shah, "Information dissemination via network coding," in *Proc. IEEE Int. Symp. on Inform. Theory*, Seattle, WA, Jul. 2006, pp. 1748–1752.
- [31] F. Chierichetti, S. Lattanzi, and A. Panconesi, "Almost tight bounds for rumour spreading with conductance," in *Proc. ACM Symp. Theory of Comput.*, Cambridge, MA, Jun. 2010, pp. 399–408.

De Wen Soh (S'11) received the B.S. honors degree in Mathematics from Stanford University in 2009. He was a Research Officer with the Institute for Infocomm Research, A*STAR (2009–2011). Currently, he is currently working towards a Ph.D. degree in the Department of Electrical Engineering at Yale University, New Haven, CT.

He received the A*STAR National Science Scholarship (B.S.) in 2006 and the National Science Scholarship (Ph.D.) in 2011. He was awarded the Stanford Undergraduate Research Award in 2009.

His main research interests include distributed network algorithms, message passing, consensus algorithms, compressed sensing, network tomography and algebraic geometry.

communications in ad hoc networks, machine learning and applied probability.

Wee Peng Tay (S'06–M'08) received the B.S. degree in electrical engineering and mathematics, and the M.S. degree in electrical engineering from Stanford University in 2002. He received the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology in 2008. He is currently an Assistant Professor in the School of Electrical and Electronics Engineering at Nanyang Technological University in Singapore. His research interests include distributed decision making, data fusion, distributed algorithms, commu-

Tony Q.S. Quek (S'98–M'08–SM'12) received the B.E. and M.E. degrees in Electrical and Electronics Engineering from Tokyo Institute of Technology, Tokyo, Japan, in 1998 and 2000, respectively. At Massachusetts Institute of Technology, Cambridge, MA, he earned the Ph.D. in Electrical Engineering and Computer Science in Feb. 2008.

Currently, he is an Assistant Professor with the Information Systems Technology and Design Pillar at Singapore University of Technology and Design (SUTD). He is also a Principal Investigator and

Scientist with the Institute for Infocomm Research, A*STAR. His main research interests are the application of mathematical, optimization, and statistical theories to communication, networking, signal processing, information theoretic and resource allocation problems. Specific current research topics include cooperative networks, interference networks, heterogeneous networks, green communications, smart grid, wireless security, localization, compressed sensing, and cognitive radio.

Professor Quek has been actively involved in organizing and chairing sessions, and has served as a member of the Technical Program Committee in a number of international conferences. He served as the Technical Program Chair for the Services & Applications Track for the IEEE WCNC in 2009, the Cognitive Radio & Cooperative Communications Track for the IEEE VTC in Spring 2011, the Wireless Communications Symposium for the IEEE Globecom in 2011, and the Systems, Standards, & Regulations Track for the IEEE WPMC in 2012; as Technical Program Vice-Chair for the IEEE ICUWB in 2011; and as the Workshop Chair for the IEEE Globecom 2010 Workshop on Femtocell Networks, the IEEE ICC 2011 Workshop on Heterogeneous Networks, the IEEE ICC 2012 Workshop on Smart and Green Communications & Networks, and the IEEE Globecom 2012 Workshop on Heterogeneous and Small Cell Networks. He is currently an Editor for the IEEE TRANSACTIONS ON COMMUNICATIONS, the IEEE WIRELESS COMMUNICATIONS LETTERS, the WILEY JOURNAL ON SECURITY AND COMMUNICATION NETWORKS, and the TRANSACTIONS ON EMERGING TELECOMMUNICATIONS TECHNOLOGIES. He was Guest Editor for the JOURNAL OF COMMUNICATIONS AND NETWORKS (Special Issue on Heterogeneous Networks) in 2011, the IEEE COMMUNICATIONS MAGAZINE (Special Issue on Heterogeneous and Small Cell Networks) in May 2013, and for the IEICE TRANSACTIONS ON COMMUNICATIONS (Special Issue on Heterogeneous Networks for Future Cellular Systems) in Jun 2013.

Professor Quek received the Singapore Government Scholarship in 1993, Tokyo Foundation Fellowship in 1998, and the A*STAR National Science Scholarship in 2002. He was honored with the 2008 Philip Yeo Prize for Outstanding Achievement in Research, the IEEE Globecom 2010 Best Paper Award, the 2011 JSPS Invited Fellow for Research in Japan, and the 2012 CAS Fellowship for Young International Scientist.