Information Dissemination via Network Coding

Damon Mosk-Aoyama
Department of Computer Science
Stanford University
Stanford, CA 94305, USA
damonma@cs.stanford.edu

Devavrat Shah
Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
devavrat@mit.edu

***Abstract*— We study distributed algorithms, also known as *gossip* algorithms, for information dissemination in an arbitrary connected network of nodes. Distributed algorithms have applications to peer-to-peer, sensor, and ad hoc networks, in which nodes operate under limited computational, communication, and energy resources. These constraints naturally give rise to "gossip" algorithms: schemes in which nodes repeatedly communicate with randomly chosen neighbors, thus distributing the computational burden across all the nodes in the network and making the computation robust against node failures.**

Information dissemination based on network coding was introduced by Deb and Médard. They showed the virtue of coding by analyzing a coding algorithm for a complete graph. Although their scheme generalizes to arbitrary graphs, the analysis does not. We present analysis of this algorithm for arbitrary graphs. Specifically, we find that the information dissemination time is naturally related to the spectral properties of the underlying network graph. Our results provide insight into how the graph topology affects the performance of the coding-based information dissemination algorithm.

I. Introduction

With the development of peer-to-peer, sensor, and wireless ad hoc networks, there has been recent interest in distributed algorithms for information dissemination and fault-tolerant computation. Motivated by this, we consider randomized gossip algorithms for communication. Gossip algorithms impose a spatial restriction on the information possessed by a node: since a node can communicate only with its neighbors in the network, it has a local view of the state of the system at any time. To obtain the global state of the network, a node must repeatedly communicate with its neighbors. Through communication across links in the network, the global state diffuses to each individual node in the network.

Network coding has been studied in a number of recent papers, such as [1], [2], [3], [4]. In the context of multicasting, network coding has been able to provide significant performance improvements. More recently, Deb and Médard [5] showed that, in a complete graph on n nodes, a coding-based gossip algorithm for information dissemination transmits n messages to all the nodes in $O(n)$ time with probability $1-O(1/n)$. This provides an improvement over the $\Theta(n \log n)$ time required for a sequential dissemination of n messages using the randomized gossip algorithm of [6] in a complete graph. The algorithm of [5] easily generalizes to arbitrary graphs. However, the method of analysis does not extend.

In this paper, we study the problem of information dissemination (or information spreading) through the use of network coding in the gossip setting for arbitrary graphs. The information dissemination time of the coding-based gossip algorithm depends on the evolution of the "dimension of the subspace" spanned by the messages at the various nodes during the course of the algorithm. The lack of symmetry in the topology of an arbitrary graph, in contrast to the case of a complete graph, leaves one with the task of studying the evolution of a rather complicated process whose state evolves in a very large space (exponential in the number of graph nodes). This makes such an analysis rather non-trivial.

The gossip algorithm's dependence on network coding makes its analysis very different from the analysis of a sequential information dissemination algorithm that we studied in our recent work [7]. As such, the method utilized as well as the precise quantitative results of this paper are very different from that of [7]. In this paper, our main contribution is an upper bound on the running time of the coding-based gossip algorithm in terms of spectral properties (or sparse cuts) of the graph. Our result provides insight into how the graph topology affects the performance of the algorithm.

A. Setup and model

Consider an arbitrary connected network, represented by an undirected graph $G = (V, E)$, with $|V| = n$ nodes. We assume that the nodes are numbered arbitrarily so that $V = \{1, \ldots, n\}$. Each node $i \in V$ has a message m_i. We seek a communication protocol that can be used to disseminate all of the messages to each of the n nodes.

In the networks in which we are interested, it is useful to have distributed protocols, in which nodes must obtain global information through local communication. This notion is captured by the communication graph G. Specifically, two nodes i and j in the network can communicate with each other if and only if $(i, j) \in E$.

To model some of the resource constraints on the nodes, we impose a *transmitter gossip* constraint on node communication. Each node is allowed to contact at most one other node at a particular time for communication. However, a node can be contacted by multiple nodes simultaneously.

A time model determines when nodes in the network communicate with each other. In this work, we consider both a synchronous and an asynchronous time model. These models are defined as follows.

- *Synchronous*: Time is measured in time slots or rounds, which are common to all nodes in the network. In any time slot, each node contacts one neighbor to initiate a communication. The choice by a node i of which node to contact can be made randomly, but any random choice must be independent of the choices made by all other nodes $j \neq i$. The gossip constraint governs the simultaneous communication among the nodes.
- *Asynchronous*: In this model, time is discretized according to the ticks of various clocks. Each node has an independent clock that ticks according to a Poisson process of rate 1. When a node's clock ticks, it chooses one neighbor (possibly at random), and contacts that neighbor.
  Equivalently, there is a global clock that ticks according to a Poisson process of rate n. Let R_k denote the time corresponding to the kth clock tick. Then, the inter-clock-tick times $\{R_{k+1} - R_k\}$ are i.i.d. exponential random variables of rate n. On each tick of the global clock, a node a_k in the network is chosen uniformly at random, and we consider the global clock tick to be a tick of the clock at the node a_k.

We measure the running times of algorithms in this paper in absolute time, which is the number of time slots in the synchronous time model, and is (on average) the number of global clock ticks divided by n in the asynchronous time model. The relationship between clock ticks in the asynchronous model and absolute time is further characterized by the following lemma and corollary.

Lemma 1: For any $k \geq 1$, let $X_1, \ldots, X_k$ be i.i.d. exponential random variables with rate λ. Let $U_k = \frac{1}{k}\sum_{i=1}^{k} X_i$. Then, for any $\epsilon \in (0, 1/2)$,

$$\Pr\left(\left|U_k - \frac{1}{\lambda}\right| \geq \frac{\epsilon}{\lambda}\right) \quad \leq \quad 2\exp\left(-\frac{\epsilon^2 k}{3}\right). \qquad (1)$$

A direct implication of Lemma 1 is the following corollary.

Corollary 2: For $k \geq 1$, $E[R_k] = k/n$. Further, for any $\epsilon \in (0, 1/2)$,

$$\Pr\left(\left|R_k - \frac{k}{n}\right| \geq \frac{\epsilon k}{n}\right) \quad \leq \quad 2\exp\left(-\frac{\epsilon^2 k}{3}\right). \qquad (2)$$

To measure the performance of a gossip protocol, we now define a quantity, the information spreading time, as follows. For any node $i \in V$, and any time t, let $M_i(t)$ be the set of messages that node i can decode using the information that it has at time t. Let $\mathcal{D}$ be an information spreading algorithm.

Definition 1: For any $\delta > 0$, the δ-information-spreading time of the algorithm $\mathcal{D}$, denoted $T^{\mathrm{spr}}_{\mathcal{D}}(\delta)$, is defined as

$$T^{\mathrm{spr}}_{\mathcal{D}}(\delta) = \inf\left\{t : \Pr\left(\cup_{i=1}^{n}\{|M_i(t)| < n\}\right) \leq \delta\right\}. \qquad (3)$$

B. Our contribution

We characterize the performance of the coding-based information dissemination algorithm in an arbitrary connected graph in terms of properties of cuts in the graph. Given the graph $G = (V, E)$ with n nodes, an $n \times n$ non-negative matrix P is said to conform to the graph G if, for $i \neq j$, $P_{ij} = 0$ whenever $(i, j) \notin E$. For such a matrix P, we define the k-conductance of P as follows.

Definition 2: The k-conductance of P, denoted Φ^k_P, is defined as

$$\Phi^k_P = \min_{S \subset V, 0 < |S| \leq k} \frac{\sum_{i \in S, j \notin S} P_{ij}}{|S|}. \qquad (4)$$

Each stochastic matrix that conforms to G gives rise to a coding-based information dissemination algorithm, whose δ-information-spreading time we denote by $T^{\mathrm{spr}}_P(\delta)$. The algorithm (described in detail in Section II) performs as follows.

Theorem 3: Consider the gossip algorithm based on Random Linear Coding (over the finite field $\mathbf{F}_q$, $q \geq n$), using a matrix P that is non-negative, stochastic, symmetric, and conforms to G. Suppose $\delta > 0$ is given and n is large enough. Let $\hat{\mu} = \sum_{k=1}^{n-1} \frac{k}{\Phi^k_P}$. Then, in the asynchronous time model,

$$T^{\mathrm{spr}}_P(\delta) = O\left(\frac{\hat{\mu}}{n}\left(1 + \frac{\log \delta^{-1}}{n}\right)\right), \qquad (5)$$

while in the synchronous time model,

$$T^{\mathrm{spr}}_P(\delta) = O\left(\frac{\hat{\mu}}{n}\log \delta^{-1}\right). \qquad (6)$$

Note. Theorem 3 implies that the δ-information-spreading time when $\delta = 1/n$ for complete graphs, constant-degree expanders, and ring graphs scales as $O(n \log n)$, $O(n \log n)$, and $O(n^2)$, respectively[1]. The bound for the complete graph is weaker than that of [5] due to the generality of the result. Specifically, a potential topic for future research is to improve the lower bound of Lemma 5, which would lead to tighter time bounds.

C. Organization

The rest of the paper is organized as follows. In Section II, we describe the network-coding-based information dissemination algorithm. We prove Theorem 3 in Section III, which consists of analysis of the information dissemination algorithm in the synchronous and asynchronous models. Finally, we present our conclusions in Section IV.

II. Coding-based gossip algorithm

The coding-based gossip algorithm for information dissemination consists of two components: the gossip *mechanism*, which determines how a node chooses a neighbor to contact when it initiates a communication; and the gossip *protocol*, which specifies the message transmitted by a node to its communication partner during a communication. Recall that each node starts the communication protocol with its unique message, and the goal is to spread all of the messages to all of the nodes. We now describe the gossip mechanism and the random coding-based gossip protocol.

[1]These bounds are for the asynchronous time model. Our bounds for the synchronous time model have an additional $\log n$ factor, though we suspect that they can be improved to match the asynchronous bounds.

Gossip Mechanism. We study a simple gossip mechanism. When node i initiates a communication, it contacts node j with probability P_{ij}, independently of all other communication events. As such, the matrix P containing the entries P_{ij} is a stochastic matrix. The restriction that communication can occur only across edges in the graph G corresponds to the requirement that P conform to the graph G. We will restrict our attention to symmetric matrices P, i.e., $P_{ij} = P_{ji}$.

In this paper, we assume that nodes transmit data according to the pull mechanism. That is, when node i contacts node j, it receives data from node j but does not send data to node j. Another popular gossip mechanism is the push mechanism, in which node i sends data to node j when node i contacts node j. We shall restrict our attention to the pull mechanism here. However, it will be clear from the result of the paper that a similar analysis applies to the push mechanism.

The data transmitted from one node to another during a communication are determined by the random linear coding (RLC) protocol explained below. When a node has received "enough" coded messages, it can decode them (see below) to obtain all n original messages.

Random Linear Coding (RLC) Protocol. This is exactly the same setup as in [5]. Each message is a vector over a finite field $\mathbf{F}_q$ of size $q \geq n$. Let each message be a vector of size $r \in \mathbf{Z}$. In particular, let the initial message at node i be $m_i \in \mathbf{F}_q^r$, for $1 \leq i \leq n$, and let $M = \{m_1, \ldots, m_n\}$ denote the set of the n message vectors. We assume that all the n messages in M are linearly independent. During the execution of the gossip algorithm, each node collects linear combinations of message vectors in M. When each node has n linearly independent such vectors, it can recover all the messages in M successfully.

Now, consider a certain instant t, during the execution of the gossip algorithm, when node i becomes active and contacts j. Let $S_i(t) = \{f_1, \ldots, f_{|S_i(t)|}\}$ and $S_j(t) = \{g_1, \ldots, g_{|S_j(t)|}\}$ be the sets of all the coded messages at nodes i and j, respectively, at time t. By definition, for $g_l \in S_j(t)$, $1 \leq l \leq |S_j(t)|$, $g_l \in \mathbf{F}_q^r$ and $g_l = \sum_{u=1}^n a_{l_u} m_u$, $a_{l_u} \in \mathbf{F}_q$. The protocol ensures that node j knows the coefficients a_{l_u} (see [5] for details). An analogous situation holds for $S_i(t)$.

When a node i contacts node j, it receives a message from node j. This message is a random coded message with payload $e_{ji} \in \mathbf{F}_q^r$, where

$$e_{ji} = \sum_{g_l \in S_j(t)} \beta_l g_l, \quad \beta_l \in \mathbf{F}_q; \;\; \Pr(\beta_l = \beta) = \frac{1}{q}, \quad \forall \beta \in \mathbf{F}_q.$$

The message e_{ji} can be re-written as follows.

$$\begin{aligned} e_{ji} &= \sum_{g_l \in S_j(t)} \beta_l g_l = \sum_{g_l \in S_j(t)} \beta_l \sum_{u=1}^{n} a_{l_u} m_u \\ &= \sum_{u=1}^{n} \left(\sum_{l=1}^{|S_j(t)|} \beta_l a_{l_u} \right) m_u = \sum_{u=1}^{n} \theta_u m_u, \end{aligned} \tag{7}$$

where $\theta_u = \sum_{l=1}^{|S_j(t)|} \beta_l a_{l_u} \in \mathbf{F}_q$. For the purpose of decoding, along with e_{ji}, node j transmits the coefficients $(\theta_1, \ldots, \theta_n)$ to node i. We now recall the following key result.

Lemma 4 (Lemma 2.1, [5]): Let $S_i(t)^-$ and $S_j(t)^-$ denote the subspaces spanned by the code vectors in $S_i(t)$ and $S_j(t)$, respectively. Let $S_i(t)^+$ be the subspace spanned by the code vectors in $S_i(t) \cup \{e_{ji}\}$. Then,

$$\Pr\left(\dim(S_i(t)^+) > \dim(S_i(t)^-) \mid S_j(t)^- \not\subseteq S_i(t)^-\right) \geq 1 - \frac{1}{q}.$$

III. Analysis of Gossip Algorithm

The performance of the gossip algorithm presented in the previous section is described by Theorem 3, which we prove here. We first prove the upper bound involving the asynchronous time model, and then the upper bound regarding the synchronous time model. Before proceeding towards separate treatments based on the time model, we first present some definitions and notation that are common to both time models. To this end, let t denote a certain instant of time when some nodes are communicating ($t \in \mathbf{R}_+$ for the asynchronous model and $t \in \mathbf{Z}_+$ for the synchronous model).

Message space. The subspaces spanned by the coded messages at node i before and after the communication at time t, respectively, are denoted by $S_i(t)^-$ and $S_i(t)^+$. We refer to the dimension of the subspace $S_i(t)^-$ as the dimension of the node i. In the synchronous model, $S_i(t)^+ = S_i(t+1)^-$.

Type. Two nodes i and j are said to be of the same type at time t if $S_i(t)^- = S_j(t)^-$, i.e., the subspaces spanned by the messages at nodes i and j are identical. For example, if both nodes have enough messages to decode all n original messages, then the subspaces spanned by both of them will be the same, so they are of the same type.

Maximum type-size. Under the definition of type, all of the nodes are partitioned into different equivalence classes, which we refer to as type classes. At time t, let $A(t)$ be the size of the largest type class (the type class containing the most nodes), also referred to as the maximum type-size.

Dimension increase. When a node j transmits a random linear code to a node i such that $S_j(t)^- \not\subseteq S_i(t)^-$, from Lemma 4, $\dim(S_i(t)^+) \geq \dim(S_i(t)^-) + 1$ with probability at least $1 - 1/q$. Now, suppose that, at time t, two nodes i and j are not of the same type. Then it must be that either (a) $S_i(t)^- \not\subseteq S_j(t)^-$ or (b) $S_j(t)^- \not\subseteq S_i(t)^-$. Thus, if the nodes i and j are of different types, then the dimension of at least one of the two nodes will increase with probability at least $1 - 1/q$ when it pulls a coded message from the other node.

Stopping condition. Since a node can decode all of the messages when the dimension of its subspace is n, the information will be disseminated to all of the nodes at the time $\min\{t : \dim(S_i(t)^-) = n, \;\; \forall i \in V\}$. Initially, at $t = 0$ we have $\dim(S_i(0)^-) = 1$ for all i. Thus, the information spreads to all the nodes when the overall dimension increase among all the nodes is $n(n-1)$. Let $D(t) = \sum_{i=1}^n \dim(S_i(t)^-) - n$ be the total dimension increase at time t. By definition, $D(0) = 0$ and the information has spread to all of the nodes when $D(t) = n(n-1)$. Now, define $t_k = \min\{t : A(t) \geq k\}$ and $Y_k = D(t_k)$. In words, t_k is the first time when any type class has at least k nodes, and Y_k is the total dimension increase at time t_k. By definition, $t_1 = Y_1 = 0$. The following result provides a lower bound on Y_k.

Lemma 5: For any $1 \le k \le n$, $Y_k \ge k(k-1)$.

We note that $Y_n = D(t_n) = n(n-1)$, and t_n is the time when all nodes have received enough coded messages to decode the original messages.

A. Asynchronous model

Preliminaries. Consider a sequence of independent geometric random variables $G_1, \dots, G_k$ with parameters $p_1, \dots, p_k$, where $0 < p_i < 1$ for $i = 1, \dots, k$. Now consider independent exponential random variables $X_1, \dots, X_k$, where X_i is of rate $\lambda_i = \ln(1-p_i)^{-1}$. It is straightforward to see that $X_i + 2$ stochastically dominates G_i, and G_i stochastically dominates $X_i - 1$. Define $U_k = \frac{1}{k}\sum_{i=1}^k G_i$ and $\hat{U}_k = 2 + \frac{1}{k}\sum_{i=1}^k X_i$. Then, $\hat{U}_k$ stochastically dominates U_k. Thus, to obtain bounds on $\Pr(U_k > x)$, it suffices to obtain bounds on $\Pr(\hat{U}_k > x)$.

The following result can be proven using properties of independent exponential random variables and inequalities based on Taylor series expansions. Let $\lambda^* = \min_{i=1}^k \lambda_i$.

Lemma 6: For $\hat{U}_k$ as defined above, let $\hat{\mu}_k = E[\hat{U}_k]$. By definition, $\hat{\mu}_k = 2 + \frac{1}{k}\sum_{i=1}^k \frac{1}{\lambda_i}$. Then, for any $\epsilon > 0$,

$$\Pr(\hat{U}_k > (2+\epsilon)\hat{\mu}_k) \le \exp\left(-\frac{k\epsilon\lambda^*\hat{\mu}_k}{2}\right).$$

We now present a straightforward corollary of Lemma 6.

Corollary 7: For U_k as defined above, let $\mu_k = E[U_k]$. Then, for any $\epsilon > 0$,

$$\Pr(U_k > (2+\epsilon)(\mu_k + 3)) \le \exp\left(-\frac{k\epsilon\lambda^*\mu_k}{2}\right).$$

Probability of dimension increase. Consider a time t when the global clock ticks (according to a Poisson process of rate n). At this instant, only one node receives a coded message from another node, so the total dimension increase is at most 1. We want to obtain a lower bound on the probability of increase. To this end, suppose there are $b \le n$ types, $C_1, \dots, C_b$. Let C^i denote the type class of a node i.

For a pair of nodes i, j, let X_{ij} be an indicator random variable that is 1 if node i contacts node j at time t and the dimension of i increases as a result of the communication, and is 0 otherwise. The node i becomes active with probability $1/n$ and contacts j with probability P_{ij}. Similarly, j contacts i with net probability $P_{ji}/n = P_{ij}/n$. If $C^i = C^j$, then there will be no increase in total dimension if i and j communicate. As noted before, however, if i and j belong to different type classes, then the dimension of at least one of the two nodes will increase with probability at least $1 - 1/q$ if it contacts the other node. This implies that whenever $C^i \ne C^j$, $E[X_{ij}] + E[X_{ji}] \ge (1-1/q)P_{ij}/n$. Using this inequality, we obtain a lower bound on the probability of dimension increase, denoted $\hat{p}$.

$$\begin{aligned} \hat{p} &= \sum_{i \in V} \sum_{j \notin C^i, j > i} (E[X_{ij}] + E[X_{ji}]) \\ &\ge \sum_{i \in V} \sum_{j \notin C^i, j > i} \left(1 - \frac{1}{q}\right)\frac{P_{ij}}{n} \\ &= \frac{1}{2n}\left(1 - \frac{1}{q}\right)\sum_{i \in V}\sum_{j \notin C^i} P_{ij}. \end{aligned} \tag{8}$$

Here, we have used the fact that P is symmetric. Now, we rewrite the sum in (8) in terms of the type classes.

$$\begin{aligned} \hat{p} &\ge \frac{1}{2n}\left(1 - \frac{1}{q}\right)\sum_{a=1}^b \sum_{i \in C_a, j \notin C_a} P_{ij} \\ &= \frac{1}{2n}\left(1 - \frac{1}{q}\right)\sum_{a=1}^b |C_a| \frac{\sum_{i \in C_a, j \notin C_a} P_{ij}}{|C_a|}. \end{aligned} \tag{9}$$

Suppose that $t \in [t_k, t_{k+1})$. Then, by definition, $|C_a| \le k$ for all $1 \le a \le b$. Using the definition of Φ_P^k and (9), we obtain

$$\hat{p} \ge \frac{1}{2n}\left(1 - \frac{1}{q}\right)\sum_{a=1}^b |C_a|\Phi_P^k = \frac{\Phi_P^k}{2}\left(1 - \frac{1}{q}\right). \tag{10}$$

Thus, in the time interval $[t_k, t_{k+1})$, the number of clock ticks required for a unit dimension increase can be stochastically bounded from above by a geometric random variable with parameter $p_k \triangleq \left(1 - \frac{1}{q}\right)\frac{\Phi_P^k}{2}$.

When the total dimension increase is $n(n-1)$, each node has received enough coded messages to obtain the original messages. As such, the number of global clock ticks W required for all nodes to decode all the original messages can be stochastically upper bounded as $W \le \sum_{d=1}^{n(n-1)} G_d$, where the G_d are independent geometric random variables with parameter p_k when $t \in [t_k, t_{k+1})$. By definition, p_k is monotonically non-increasing in k. Hence, the smaller the t_k values are, the worse this stochastic upper bound on W is. From Lemma 5, the worst stochastic upper bound on W is

$$W \le \sum_{k=1}^{n-1}\sum_{l=1}^{2k} G_l^k \triangleq \hat{W}, \tag{11}$$

where the G_l^k are independent geometric random variables with parameter p_k. From (11), it is straightforward that for $q \ge n \ge 2$,

$$E[W] \le E[\hat{W}] = \frac{4}{1 - \frac{1}{q}}\sum_{k=1}^{n-1}\frac{k}{\Phi_P^k} = \Theta(\hat{\mu}). \tag{12}$$

To obtain the bound with probability $1 - \delta/2$, we use Corollary 7. Let $p^* = \min_{k=1}^{n-1}\ln(1-p_k)^{-1} \ge \min_{k=1}^{n-1} p_k = \min_{k=1}^{n-1}\left(1 - \frac{1}{q}\right)\frac{\Phi_P^k}{2}$. By definition, Φ_P^k is monotonically non-increasing in k. Hence, by the definition of $\hat{\mu}$,

$$p^* = \left(1 - \frac{1}{q}\right)\frac{\Phi_P^{n-1}}{2} = \Omega\left(\frac{n}{\hat{\mu}}\right). \tag{13}$$

Now, from Corollary 7, for $\epsilon > 0$,

$$\begin{aligned} &\Pr\left(\hat{W} > (2+\epsilon)(E[\hat{W}] + 3n(n-1))\right) \\ &\le \exp\left(-\frac{\epsilon p^* E[\hat{W}]}{2}\right). \end{aligned} \tag{14}$$

Let $B = (2+\epsilon)(E[\hat{W}] + 3n(n-1))$ when $\epsilon = \frac{2\ln(2/\delta)}{p^* E[\hat{W}]}$. Since $E[\hat{W}] = \Omega(n^2)$, we have

$$B = O\left((1+\epsilon)E[\hat{W}]\right) = O\left(\hat{\mu}\left(1 + \frac{\log\delta^{-1}}{n}\right)\right).$$

Substituting for ϵ in the inequality in (14), we obtain $\Pr(\hat{W} > B) \leq \delta/2$. This provides an upper bound on the number of clock ticks required for every node to obtain every message.

To extend the bound to absolute time, we apply Corollary 2, which implies that the probability that $B = \Omega(\log \delta^{-1})$ clock ticks do not occur by absolute time $O(B/n)$ is at most $\delta/2$. We conclude from the union bound, (12), and (13) that

$$T_P^{\text{spr}}(\delta) = O\left(\frac{\hat{\mu}}{n}\left(1 + \frac{\log \delta^{-1}}{n}\right)\right).$$

B. Synchronous model

We begin as in the analysis of the asynchronous model. Suppose that at time $t \in [t_k, t_{k+1})$ there are b type classes, $C_1, \ldots, C_b$. As before, X_{ij} is an indicator random variable specifying whether node i contacts node j and the dimension of i increases in round t. Let $L(t) = D(t+1) - D(t)$ denote the total dimension increase of all nodes in this round, so that

$$L(t) = \sum_{i \in V} \sum_{j \in V} X_{ij}. \tag{15}$$

Repeating the argument for the asynchronous model, we consider two nodes i and j of different classes $C^i \neq C^j$. In the synchronous model, we have $E[X_{ij}] + E[X_{ji}] \geq P_{ij}(1 - 1/q)$ (the factor of $1/n$ in the asynchronous case is not present because all nodes are simultaneously active in the synchronous model). We use (15) and this lower bound to obtain a lower bound on $E[L(t)]$.

$$\begin{aligned} E[L(t)] &= \sum_{i \in V} \sum_{j \notin C^i, j > i} (E[X_{ij}] + E[X_{ji}]) \\ &\geq \frac{1}{2}\left(1 - \frac{1}{q}\right) \sum_{a=1}^{b} |C_a| \frac{\sum_{i \in C_a, j \notin C_a} P_{ij}}{|C_a|} \\ &\geq \frac{n\Phi_P^k}{2}\left(1 - \frac{1}{q}\right) \triangleq np_k. \end{aligned} \tag{16}$$

This provides a lower bound on the expected total dimension increase during any round in the period $[t_k, t_{k+1})$. Note that this lower bound holds for any $t \in [t_k, t_{k+1})$ uniformly. Define

$$Z^k(t) = \sum_{v=t_k}^{t-1} (L(v) - np_k)\mathbf{1}_{\{v < t_{k+1}\}}, \tag{17}$$

where $Z^k(t_k) = 0$. For $t \geq t_k$, $Z^k(t)$ is a submartingale, i.e.,

$$E[Z^k(t+1) \mid Z^k(t)] \geq Z^k(t). \tag{18}$$

The quantity t_{k+1} is a stopping time with respect to the history of the algorithm. It is easy to show that $E[t_{k+1}] < \infty$ via a stochastic upper bound using a certain geometric random variable with positive probability. Moreover, the submartingale $Z^k(t)$ has bounded increments. A stopped submartingale is a submartingale, and hence we obtain

$$E[Z^k(t_{k+1})] \geq E[Z^k(t_k)] = 0. \tag{19}$$

Now, from the definitions of $t_k, t_{k+1}, Y_k, Y_{k+1}$, and (19), we obtain

$$E[Y_{k+1} - Y_k] \geq np_k E[t_{k+1} - t_k]. \tag{20}$$

Recall that t_n is the time when all nodes can decode all the messages. Summing the inequality in (20) for all $1 \leq k \leq n-1$ yields

$$E[t_n] \leq \sum_{k=1}^{n-1} \frac{E[Y_{k+1} - Y_k]}{np_k}. \tag{21}$$

From Lemma 5 and the fact that p_k is monotonically non-increasing in k, the quantity in the right-hand side of the inequality in (21) is maximized when $Y_k = k(k-1)$. Hence,

$$E[t_n] \leq \sum_{k=1}^{n-1} \frac{2k}{np_k} = \frac{2\hat{\mu}}{n}. \tag{22}$$

By Markov's inequality, the inequality in (22) implies that $\Pr(t_n > 4\hat{\mu}/n) < 1/2$.

Now, for the purpose of analysis, consider dividing time into epochs of length $4\hat{\mu}/n$, and executing the information dissemination algorithm from the initial state in each epoch, independently of the other epochs. The probability that, after $\log \delta^{-1}$ epochs, some execution of the algorithm has run to completion in its epoch is greater than $1-\delta$. Using the running time of this virtual process as a stochastic upper bound on the running time of the actual algorithm, we can conclude that $T_P^{\text{spr}}(\delta) = O((\hat{\mu}/n) \log \delta^{-1})$.

IV. Conclusion

In this paper, we considered the question of information dissemination via gossip algorithms. Specifically, we studied the information dissemination time for a gossip algorithm based on network coding. The use of coding was shown to be beneficial by Deb and Médard for information dissemination in the context of a complete graph. The main question that remained open was whether this coding-based gossip algorithm can help improve the performance of information dissemination in an arbitrary graph.

Motivated by this question, we analyzed the performance of an information dissemination algorithm based on coding for arbitrary graphs. We found that the performance of the algorithm is closely related to spectral properties of the graph.

References

[1] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Transactions on Information Theory*, vol. 46, no. 4, pp. 1204–1216, 2000.

[2] S.-Y. R. Li, R. W. Yeung, and N. Cai, "Linear network coding," *IEEE Transactions on Information Theory*, vol. 49, no. 2, pp. 371–381, 2003.

[3] R. Koetter and M. Médard, "An algebraic approach to network coding," *IEEE/ACM Transactions on Networking*, vol. 11, no. 5, pp. 782–795, 2003.

[4] ——, "Beyond routing: An algebraic approach to network coding," in *Proceedings of IEEE INFOCOM 2002*, 2002, pp. 122–130.

[5] S. Deb and M. Médard, "Algebraic gossip: A network coding approach to optimal multiple rumor mongering," in *Proceedings of the 42nd Annual Allerton Conference on Communication, Control, and Computing*, 2004.

[6] R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking, "Randomized rumor spreading," in *Proceedings of the 41st Annual IEEE Symposium on Foundations of Computer Science*, 2000, pp. 565–574.

[7] D. Mosk-Aoyama and D. Shah, "Computing separable functions via gossip," in *25th Annual ACM Symposium on Principles of Distributed Computing*, 2006, to appear.