Network Coding Based Information Spreading in Dynamic Networks With Correlated Data

Citation: Cohen, Asaf, Bernhard Haeupler, Chen Avin, and Muriel Medard. "Network Coding Based Information Spreading in Dynamic Networks With Correlated Data." IEEE Journal on Selected Areas in Communications 33, no. 2 (February 2015): 213–224. doi:10.1109/jsac.2014.2384293.

As Published: http://dx.doi.org/10.1109/JSAC.2014.2384293

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Persistent URL: http://hdl.handle.net/1721.1/100956

Version: Original manuscript: author's manuscript prior to formal peer review

Network Coding Based Information Spreading in Dynamic Networks with Correlated Data

Bernhard Haeupler*, Asaf Cohen*, *Member, IEEE,* Chen Avin, *Member, IEEE,* and Muriel Médard, *Fellow, IEEE*

***Abstract*—In this work, we design and analyze information spreading algorithms for dynamic networks with correlated data. In these networks, either the data to be distributed, the data already available at the nodes, or both, are correlated. Moreover, nodes' availability and connectivity is dynamic – a scenario typical for wireless networks.**

Our contribution is twofold. First, although coding schemes for correlated data have been studied extensively, the focus has been on characterizing the rate region in static networks. In an information spreading scheme, however, nodes may communicate by continuously exchanging packets according to some underlying communication model. The main figure of merit is the stopping time – the time required until nodes can successfully decode. While information spreading schemes, such as gossip, are practical, distributed and scalable, they have only been studied for uncorrelated data. We close this gap by providing techniques to analyze network-coded information spreading in dynamic networks with correlated data.

Second, we give a clean framework for oblivious dynamic network models that in particular applies to a multitude of wireless network and communication scenarios. We specify a general setting for the data model, and give tight bounds on the stopping times of network coded protocols in this wide range of settings. En route, we analyze the *capacities seen by nodes* under a network-coded information spreading protocol, a previously unexplored question.

We conclude with extensive simulations, clearly validating the key trends and phenomena predicted in the analysis.

***Index Terms*—Dynamic Networks; Information Dissemination; Spreading; Gossip; Network Coding; Correlated Data.**

I. Introduction

From last-mile connectivity to smart sensing, wireless networks play a key role in the communication infrastructure today, and will constitute an even larger part in the future. Although various communication protocols for such networks are available, it is clear much is still required to harness the full potential in wireless networks while coping with the many challenges they pose. A primary technique in fulfilling this potential is network coding, and, specifically, *network-coded gossip* appears to be a favorable scheme for (all-to-all) information dissemination (spreading) in wireless networks.

B. Haeupler is with the School of Computer Science, Carnegie Mellon University, e-mail: haeupler@cs.cmu.edu.

A. Cohen and C. Avin are with the Department of Communication Systems Engineering, Ben-Gurion University of the Negev, e-mail: {coasaf,avin}@bgu.ac.il.

M. Médard is with the Department of Electrical Engineering and the Research Laboratory of Electronics, MIT, e-mail: medard@mit.edu.

Parts of this work appeared at the IEEE International Symposium on Information Theory, ISIT 2012.

This work is partially supported by DARPA under Contract No. N66001-11-C-4003 and by RESCUE, the Israeli Chief Scientist.

* A. Cohen and B. Haeupler contributed equally to this work.

In a gossip scheme, nodes communicate by continuously exchanging messages among them, according to some underlying communication model. Gossip schemes are local, structure-free, and distributed. Moreover, they are efficient and scalable. These are highly desirable key properties in dynamic networks in general and wireless networks in particular.

In this work, our main focus is on networks with correlated data. In such networks, either the data to be distributed, the data already available at the nodes, or both, are correlated. One example to keep in mind throughout this work can be a wireless sensor network which distributes measurements such as temperature readings. If sensors also have position readings they might be able to exploit correlations between (known) positions of nodes and their measurements or any correlation between other nodes measurements, such as the correlation between their temperature and the temperature of other close-by sensors. In such scenarios, it is clear that *a cross-layer design* which uses the available data in the network, or the correlation in the data to be sent (across multiple nodes) can reduce the required rates and dissemination times. This model is especially useful in situations where communication has a higher cost compared to local computation.

The current literature in information theory includes several coding schemes for correlated data. Yet, the main focus in these works is on characterizing the rate region – the set of achievable rates. Moreover, the topologies discussed are mainly static, with memory-free and error-free links. Thus, there is still a gap to distributed schemes for large or more dynamic networks, in which the need for decentralized models dominate. On the other hand, recent work in the networking literature offers a multitude of efficient, decentralized and address-oblivious schemes for information dissemination (e.g., randomized gossip). Unfortunately, these schemes treat the packets to be delivered as uncorrelated and ignore the possibility of correlated data or side information at the receiving nodes. The focus of this paper is thus to close this gap by suggesting gossip-based schemes for networks with side information or correlated data and give the required tools to analyze them.

Furthermore, while analysis of gossip schemes usually focuses on one figure of merit, the *stopping time*, which is the time required to disseminate all messages to all nodes, such an analysis is too coarse when correlated data or side information is present. In such cases, one has to ensure sufficient throughput from *several locations* in the network (e.g., at least the conditional entropy from each source, and additional sum-rate constraints). Thus, to successfully analyze gossip schemes in this setting, we formally define the concept of *oblivious*

networks, which captures channel quality, connectivity and packet loss issues under a unified model, and give accurate measures of the *actual capacities available between sets of nodes* under a network-coded gossip scheme in oblivious networks. Such an analysis has implications beyond networks with correlated data, and can be used to understand network-coded gossip schemes in general.

A. Main Contributions and Paper Structure

To the best of our knowledge, this paper is the first to combine the two strains of research and analyze gossip protocols in networks with correlated data. Our contributions are manifold. First, we define a setting for a correlated environment in Section III and give a clean and general framework for oblivious network models in Section IV. In this general setting, we extend the analysis of [1] by making a connection between the coefficient vectors a node knows and the amount and type of information it has learned. This results in direct and self-contained proofs of tight bounds on the stopping time in the canonical models of one source and side information at the receivers, as well as two correlated sources.

In Section VI we provide tight bounds on the time required for any set of (fractional) capacities to be induced by the (random) packet exchanges generated in a network-coded gossip session on an oblivious network model. These capacity bounds are interesting in their own right and have the potential to be useful in a multitude of information dissemination problems. For example, they could be used to prove the fault-tolerance and robustness of gossip networks based on establishing bounds on their connectivities. In Section VII we harness the results of Section VI to transform results on the rate region of static memory-free networks (e.g., [2], [3]) into bounds on the stopping times of gossip-based algorithms for the general scenario of multiple correlated sources and side information. Section VIII includes several simulation results illustrating the key findings in this paper.

II. Related Work

In this section, we briefly review the literature in the context of distributed source coding, gossip and network coding.

A. Distributed Source Coding

Distributed source coding has been studied in the information theoretic literature through mainly small, canonical problems. In their seminal work [4], Slepian and Wolf considered the problem of separately encoding two correlated sources and joint decoding. See Figure 1 (a) for the canonical model. A key finding was that one can achieve the fundamental bound on the sum-rate, the joint entropy, even without a cooperation between the two encoders. In [5], Ahlswede and Körner considered a seemingly very similar problem, where the decoder is interested in only one of the sources (the other being a *helper*). Results showed that in general, the sum-rate in such cases may be higher than the entropy of the required source. Further extensions appeared in [6], [7] and more recent works such as [8]. However, these works

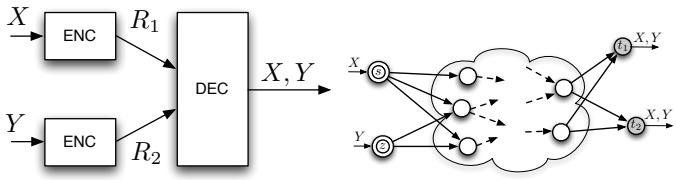

(a) Slepian-Wolf (b) Ho *et al.*

Fig. 1: (a) The Slepian-Wolf network [4]. Upper encoder has the source X, lower one has the side information Y, which is, in general, correlated with X. We are interested in the set of rates (R_1, R_2) such that both X and Y can be reconstructed at the decoder. (b) The network model of [9]. Sources X and Y are available at s and z, respectively, and demanded at t_1, t_2.

consider specific network topologies which are usually very small (e.g. three-node networks). In [9], Ho et al. considered the multicast problem with correlated sources (Figure 1 (b)), and completely characterized the *rate region* for this problem: the set of required link capacities to support the multicast. In a way, [9] can be viewed as extending the Slepian-Wolf problem to arbitrary networks through *network coding*. Further extensions also appeared in [10] and [3]. In [11], the authors considered the model where a node sends the information X to a few designated receivers, while a helper node has correlated data Y and can thus aid in the transmission. The objective discussed therein was again the rate region. Thus, the above works deal with computing rate regions under a designated set of receivers, and not the *decentralized*, dissemination model of gossip schemes. An alternative approach was taken in [12] by Haupt et al. which considered a compressed sensing scheme for compression of network data.

B. Gossip

Gossip schemes were first introduced in [13] as a simple and decentralized way to disseminate information in a network. A detailed analysis of a class of these algorithms was given in [14]. In these schemes, nodes communicate in rounds. In each round, a random node in the network chooses a single communication partner according to the gossip *algorithm* (e.g., selecting a random neighbor). Once a partner is chosen, a limited amount of data is transferred, as defined by the gossip *protocol*. The main figure of merit is the dissemination time: the average time for all nodes to receive all messages. Such randomized gossip-based protocols are attractive due to their locality, simplicity, and structure-free nature, and have been offered in the literature for various tasks, such as ensuring database consistency and computing aggregate information and other functions of the data [14], [15], [16].

In [17], Deb et al. introduced *algebraic gossip*, a coding-based gossip protocol where nodes exchange *linear combinations* of their available messages. The advantages in terms of average dissemination time were clear: for example, a node in a complete graph of size n would require $O(n)$ messages to receive all data with high probability. [18] and [19] extended the analysis to arbitrary graphs. Haeupler [1] proved a tight bound for the stopping time of algebraic gossip for various

graph models, including dynamically changing graphs. The methods in [1] will play a key role in this work. More results in this direction were given in [20], [21].

A key assumption in all of the above gossip schemes, is the fact that the data to be distributed among nodes is uncorrelated. That is, each node carries its own data it wishes to disseminate and this data is independent of the data other nodes wish to distribute or might have in their possession. In this work, we explore algorithms for data dissemination which use this correlation to achieve a faster dissemination time, at the price of decoding complexity at the nodes (first steps towards the above goals were given by the authors in the extended abstract [22], yet with a wire-line oriented nature and no proofs).

In fact, at the basis of this paper stands the assumption that while compression schemes for point-to-point links have been studied extensively, there is still a huge gap, yet to be exhausted, in distributed compression schemes for dynamic and highly decentralized network models.

C. Network Coding

A key concept which facilitated the analysis and design of coding problems for large networks, for both uncorrelated and correlated data, is *network coding*. First introduced by Ahlswede et al. in [23], network coding deals with various coding operations that can be performed at *intermediate nodes* in the network in order to achieve certain rate goals. The theory of network coding includes, for example, algebraic and random coding [24], [2].

Linear network coding concepts played a key role in information dissemination problems. In [25], the problem of disseminating M packets in a hyper-graph setting was considered. The suggested solutions, however, were *centralized* and assumed the source node can both choose which hyper-arcs should be active as well as which data to send. A distributed, *algebraic* gossip scheme was suggested in [17]. A testbed implementation including mobile nodes was presented in [26]. The information model therein assumed one source wishing to distribute its data to all other nodes. Transmissions were based on UDP broadcast and random linear network coding. The benefit of broadcasting coded packets over TCP unicast was clearly visible for both indoor and outdoor. Algebraic (network-coded) gossip will also play a key role in this paper.

The pioneering work in [9] combined network coding and distributed source coding. Correlated sources were also considered in [27], where *analog* random linear network coding and compressive sensing decoders were used. The scheme in [27], however, considered only Gaussian channels and a joint source-channel-network coding solution.

III. Network and Information Model

We first state the problem rigorously and define the network and communication model. We then define the information model: the *sources and side information*. The oblivious network models for which the results in this paper will apply are defined rigorously in Section IV.

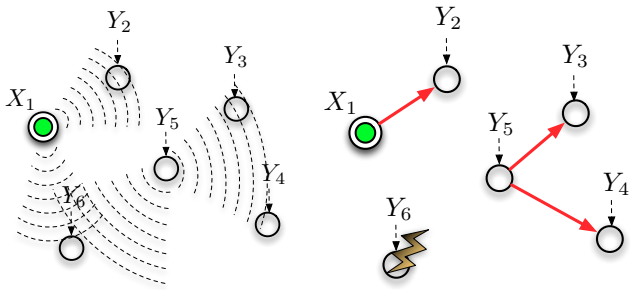

(a) (b)

Fig. 2: Network model. (a) One source and five terminals, each with its corresponding side information. During a single round, several nodes may transmit simultaneously. (b) The message sent from the source was received by one node. The message sent from node 5 was received by two. Node 6, however, did not receive a message due to, e.g., a collision. The red links constitute the active edge set E_t for the specific round.

A. Network and Communication Model

The network consists of a fixed set V of n nodes. Communication takes place in synchronous rounds. Each round, each node $v \in V$ decides on a packet p_v of s bits to send (possibly using randomness). Given the current state of the network, the (possibly randomized) network model then specifies which packet will be received by which node in the current round. Such a scheme corresponds to having a probability distribution over the directed edges, where an edge from v to u means that p_v is successfully received by u. Nodes are assumed to have sufficient memory to store all received packets (limited memory was considered in [28]). We denote the *active edge set* of directed edges chosen for round t with E_t. Figure 2 depicts a simple wireless model and a possible active edge set. At a certain round, several nodes may transmit simultaneously. However, due to interference, collisions or low SNR not all packets are received by all nodes. A probability distribution specifies which packets were received at a certain round. In this example, E_t includes only the three red edges; Nodes 3 and 4 might not have heard the message from the source due to its low received power, yet managed to receive the message from node 5, while node 6 did not hear either messages due to a collision (low SINR from both transmitters).

B. Source and Side Information

We assume random variables $\{X_i\}_{i=1}^{k} \cup \{Y_v\}_{v \in V}$ are arbitrarily correlated according to some known memoryless joint distribution. The messages and side information at the nodes are generated by taking l i.i.d. samples from each variable. The l-length message vectors are denoted $x_1, \ldots, x_k$ while the side information vectors are denoted as y_v, $v \in V$. The k messages are initially distributed to the source nodes such that each vector is available at at least one source node. Each node $v \in V$ is given the vector y_v as side information (nodes without side information can be modeled as nodes having Y_v uncorrelated with the sources). Given this initial information

distribution, we are interested in the stopping time: the time when all nodes are *able to decode* $x_1, \dots, x_k$ *based on their side information and the packets exchanged* with other nodes.

C. The Encoding and Decoding Schemes

For a given field size q and slack $\delta > 0$ we assume throughout that nodes employ the following coding scheme: Prior to communication, source nodes perform random binning, that is, for every $1 \leq i \leq k$ each node receiving the message vector x_i applies the same random mapping into $2^{l(H(X_i)+\delta)}$ bins. The resulting bin *indices* (the same for every node initialized with x_i) are interpreted as vectors of length

$$h = \left\lceil \frac{l}{\log q}(H(X_i) + \delta) \right\rceil$$

over the finite field $\mathcal{F}_q$. These vectors are split into $\frac{h \log q}{s}$ blocks of $\frac{s}{\log q}$ symbols in $\mathcal{F}_q$ each, for a total of s bits per block.

During the communication phase nodes send out random linear combinations over F_q of these blocks as packets. This is the standard random linear network coding over $\mathcal{F}_q$. That is, each packet contains $\frac{s}{\log q}$ symbols in $\mathcal{F}_q$, each resulting from a random linear combination of the symbols. We say that each packet contains $\frac{s}{\log q}$ *equations*. To keep track of the linear combination contained in a packet one coefficient for each block of each message is introduced and sent in the header of each packet. As in all prior works on distributed network coding (e.g., [17], [18], [19], [1], [28], [29]), we assume that $\frac{s}{\log q}$ is sufficiently large compared to the number of coefficients. This renders the overhead of the header negligible, leaving a packet size close to s bits as desired.

Each node collects independent linear equations on the blocks. We denote with S_v the subspace spanned by all coefficient vectors received at node v. Note that S_v depends only on the messages received at v, and is independent of the side information Y_v available at v. We also use the following notion of knowledge from [1]:

Definition 1. Node v *knows a coefficient vector* μ iff S_v is not orthogonal to μ, i.e, there exists $c \in S_v$ such that $\langle c, \mu \rangle \neq 0$.

Lastly, we will make use of the following lemma.

Lemma 1 ([3], [11])**.** *Let $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$ be two arbitrarily correlated random variables and let x, y be two vectors that are created by taking l i.i.d. samples from their joint distribution. Suppose, for some $\epsilon > 0$ and for some constant $\delta > 0$, that all possible sequences in $\mathcal{X}^l$ are randomly and uniformly divided to at least $2^{l(H(X|Y)+\delta)}$ bins. Then joint typicality decoding correctly decodes x with probability at least $1 - \epsilon$ for $l = \Omega(\log \epsilon^{-1})$ using y and any $\lceil (H(X|Y) + \delta)l \rceil$ bits of information on the bin index of the bin in which the true x resides.*

In particular, Lemma 1 asserts that having access to the side information y, the message vector x can be decoded with high probability using *any* $\left\lceil \frac{l}{s}(H(X|Y) + \delta) \right\rceil$ *linearly independent equations* on the blocks describing the bin index of x.

Finally, note that the decoding error probability ϵ is not the result of the network models herein or in [3], [11], but is inherent to distributed encoding of correlated sources [4] and vanishes only for infinite block length. However, as the error probability decays *exponentially* with this length, for non-asymptotic statements, $l = \Omega(\log \epsilon^{-1})$ suffices to achieve an error probability less than ϵ.

IV. Oblivious Network Models

In this section, we introduce the definition of an oblivious network model. This gives a clean and very general framework capturing a wide variety of communication and network settings. The flexibility of the oblivious network framework is particularly suited to model complex network behaviors as found in wireless or dynamic networks. While this was already somewhat implicit in [1], we restrict ourselves to networks without adaptive adversarial behavior. This greatly facilitates the much cleaner framework presented herein.

Definition 2. A network model is *oblivious* if the active edge set E_t of time t only depends on t, $E_{t'}$ for any $t' < t$ and some randomness. An oblivious network model is i.i.d. if the active edge set E_t is sampled independently for every t from the same distribution.

It is important to note that in an oblivious network model, the active edge set *does not* depend on the previously or currently sent data. The following are simple and somewhat classical examples that fall under this wide umbrella of oblivious (and i.i.d.) network models:

1) (Uniform) Gossip in General Graphs [18], [19], [1]: Let G be a (weighted) directed graph. In each round each node picks a (uniformly) random neighbor and sends a message to it (PUSH) or requests a message (PULL) or both (EXCHANGE).
2) Rumor Mongering or Random Phone Calls [14], [17]: This is a well-studied special case of uniform gossip in which the underlying graph is complete and unweighted, that is, in every round a node chooses a random other node for a PUSH, PULL or EXCHANGE.
3) Wired Networks with Random Packet Losses: Let G be a directed graph with a weight $w_e \in [0, 1]$ on each edge e. In each round, an edge e is active independently at random with probability $1 - w_e$.
4) (Edge-)Markovian Evolving Graphs [30], [31], [32]: These are interesting examples of non-i.i.d. oblivious network models in which the packet loss probability of edges evolves over time according to a Markov chain.

While this demonstrates that our model covers and to some extent generalizes existing models, it is equally important to see how this general notion can be applied to much more complex and more realistic settings of interest such as wireless and/or dynamic networks. These settings might not have network topologies which can be summarized by a graph. Instead, omni-directional broadcasts in some geometry, power constraints, half-duplex transmissions, collisions, and random (correlated) packet losses that depend on SINR and other characteristics are typical properties of radio networks one would like to capture. In a packet network, however, these complex considerations still lead to nodes either receiving a

message sent from a node in a specific round or not. The outcome of a round can therefore be perfectly captured by a set of directed edges.

A natural random process associated with an oblivious network model M, is the random *flooding process* $F(M, p, S)$. Informally speaking, this random process tracks the information spreading over time: which nodes are informed at each round if initially only nodes in S were informed and on every time step t, every informed node informs all its communication partners specified by M. However, an important modification to this standard infectious process is the fault parameter p, which adds an independent probability $1-p$ for each transmission to be overheard. This is since when investigating the spreading of a single coefficient vector μ, a linear combination sent by a single node is either orthogonal to μ or not. In the former case, the message will not inform *any node* about μ, while in the latter all nodes successfully receiving the linear combination will be informed. The probabilities for those two cases are captured by the probabilities $(p, 1-p)$. Formally, we have the following definition.

Definition 3. Let M be an oblivious network model, p be a probability of fault and $S \subseteq V$ be a starting set of nodes. We define the *flooding process* $F(M, p, S)$ to be the random process $S_1 \subseteq S_2 \subseteq \dots$ that is characterized by $S_1 = S$ and for every time t we define S_{t+1} by selecting a subset $S'_t \subseteq S_t$ by selecting each node in S_t independently with probability $1-p$ and then consider the edge set E_t specified by M for time t to set $S_{t+1} = S_t \cup \{v \in V \setminus S_t \mid \exists u \in S'_t : (u, v) \in E_t\}$.

Figure 3 gives an example of a simple flooding process. Note that Definition 3 is only well defined if M is oblivious. Under the above definition, it is not hard to see that the following claim holds.

Claim 1. *Assume M is an oblivious and i.i.d. network model. Then $F(M, p, S)$ is a time-homogeneous Markov chain. Furthermore, if $p \neq 1$, $S \neq \phi$ and for every time t the union over the edges in M from t to infinity $\cup_{t'=t}^{\infty} E_{t'}$ is almost surely connected then $F(M, p, S)$ is absorbing with the unique absorbing state V.*

We say the flooding process F stops if it reaches this absorbing state and we denote the time this happens with the random variable S_F. The next definition pairs this flooding time with a throughput parameter α that corresponds to the exponent of the flooding process tail probability. The reason for this definition and its connection to the multi-message throughput behavior of network coding becomes apparent in the statement and proof of Theorem 1 below.

Definition 4. We say an oblivious network model M on a node set V *floods in time T with throughput α* if there exists a prime power q such that for every $v \in V$ and every $k > 0$ we have $P[S_{F(M,1/q,\{v\})} \geq T + k] < q^{-\alpha k}$.

In other words, the network floods in time T if the probability of *not disseminating the message to all nodes* decays exponentially after T.

It is not hard to give illustrative examples of flooding times for a few interesting communication models. For example,

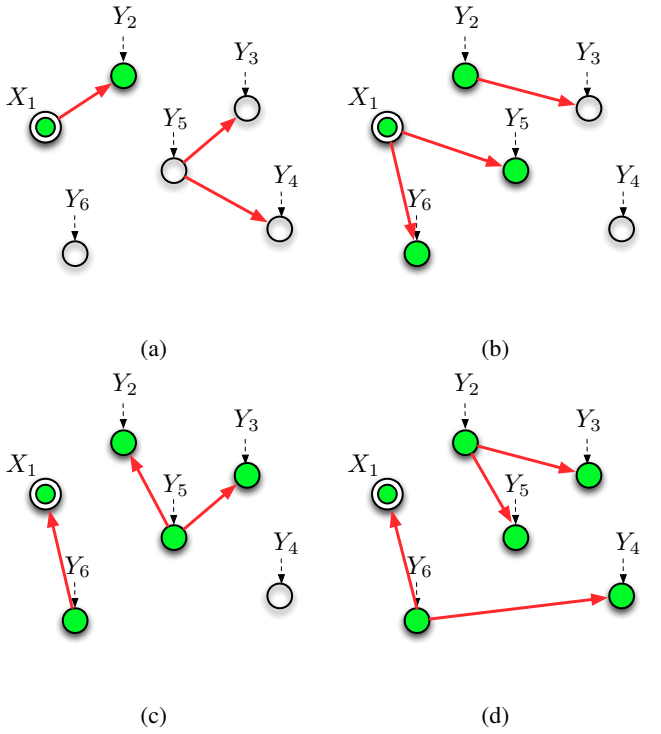

Fig. 3: *"Rumor... She flourishes as she flies, gains strength by mere motion. Small at first and in fear, she soon rises to heaven..." (Virgil, The Aeneid).* A flooding process in an oblivious network. Consider a specific coefficient vector μ: (a) Node 1 knows μ, and successfully informs node 2 about it. (b) Although node 2 knowns μ, and a directed link to node 3 is in the active edge set, node 3 is not informed in this round as node 2 was not selected in this flooding round (that is, the random coefficient it chose were perpendicular to μ. (c) Node 5 was selected and the link to node 3 is in the active edge set, hence node 3 is informed as well. (d) The flooding process is completed.

in the random phone call model with n nodes, the flooding time is $\Theta(\log n)$ with constant throughput. The lower bound is clear: even if informed nodes miraculously contact only uninformed nodes, each round the number of informed nodes will only double (assuming each node contacts a single node each round), hence one will need $\Omega(\log n)$ rounds. Analysis of the upper bound is similar to that in [1, Proof of Lemma 4]. We include it here for completeness. If the number of informed nodes, i, is at most $n/2$, then there are at least $n/2$ uninformed nodes. Each of these is contacted with probability $1-(1-1/n)^i$ by an informed node. Thus, on average, at least $n/2\left(1-(1-1/n)^i\right) \geq i/4$ are now newly informed. This means an exponential growth of informed nodes until $n/2$. If the number of informed nodes is larger than $n/2$, then each uninformed node has a chance of at least $1-(1-1/n)^{n/2}$ to be informed in a round, hence a constant fraction of the uninformed will be informed. This means the number of uninformed, on average, is cut down by a constant fraction each round, resulting in $O(\log n)$ stopping time.

Another example is that of uniform gossip on a connected,

bounded degree (i.e., constant) graph G. In this case, the network floods in time $\Theta(D)$ and with constant throughput where D is the diameter of G. This is easily seen noting that on the one hand, since D is the diameter, there exists a placement of the initial message such that at least one node will required D rounds. On the other hand, for any path from the source to a destination and for any node on such a path, since the degree of the node is constant, with high probability after a constant number of rounds the message will proceed in the right direction (towards the destination), hence will arrive after $O(D)$ rounds with high probability.

Finally, we note that one may also use the wide variety of results already available in the literature on flooding times or global network outreach [33], [30], [34], [35]. Such bounds may come in handy when using the results in this paper to extend the various models to multiple messages, correlated data and side information.

With this framework for oblivious network models in place, we can give a cleaner restatement of Theorem 3 in [1]. The concepts in the proof will also play a role in the remainder of the paper.

Theorem 1 (Theorem 3 of [1])**.** *Suppose M is an oblivious network model that floods in time T with throughput α. Then, for any k, random linear network coding in the network model M spreads k arbitrarily distributed messages to all nodes with probability $1-\epsilon$ after $T' = T + \frac{1}{\alpha}(k + \log \epsilon^{-1})$ rounds.*

Theorem 1 depicts the *perfect pipelining property* of random linear network coding in dynamic networks as well: once the first message arrives after T rounds (with high probability), the proceeding messages arrive (namely, can be decoded) "one-by-one", with no delay, as if there was a direct link to the destinations with throughout α. For example, in a random phone call model, where it takes a message $\Theta(\log n)$ rounds to disseminate, Theorem 1 asserts that with random linear network coding any number of additional messages can be disseminated with an additional number of rounds which is linear in the number of messages. The multiplicative factor can be analyzed, and it is determined by the topology and the communication model of the network.

Proof: The random linear network coding protocol analyzed herein will use the same field size that achieves the parameters T and α for M in Definition 4. Fix a coefficient vector $\mu \in F_q^k$. Using the definition of *knowing* from Definition 1, this vector is initially known to a non-empty subset S of nodes. Furthermore, a node learns about μ if and only if it receives a packet from a node that knows μ which furhtermore chooses a non perpendicular coefficient vector. It is easy to see that the later happens with probability $1-1/q$ and that this probability is independent for every node in S. This implies that knowledge of μ spreads through the network at least as fast the flooding process $F(M, 1/q, S)$ in which also only nodes that got an incoming directed edge from a node in S which was furthermore independently selected with probability exactly $1-p$. Now, using the assumption, Definition 4, and the monotonicity of $S_{F(M,1/q,S)}$ in the initial node set S asserts that the probability that a vector $\mu \in F_q^k$ is not known to all vectors after T' steps is at most $q^{-(k+\log \epsilon^{-1})}$. A union bound over all q^k vectors shows that with probability at least

$$1 - q^{-\log \epsilon^{-1}} \geq 1 - \epsilon$$

all nodes will know about all vectors and it is easy to see that this implies that all nodes are able to decode all messages. ■

Note that the added term $\log \epsilon^{-1}$, which is the result of the exponential bound guaranteed through Definition 4, can be made negligible if the number of messages, k, is large enough. This will also be true in the remainder of the paper, where side information and correlated sources are considered, as while one increases the block length l (the equivalent of the number of messages in this case), this term increases only logarithmically with l. Still, to asses its effect in the non-asymptotic regime, simulations for the probability of not being able to decode all messages, ϵ, as a function of the number of rounds are given in Section VIII. The trends predicted in the analysis are clearly visible therein.

Finally, we mention that a key strength of the oblivious network model suggested herein, is the ability to *abstract over issues such as MAC protocols, interference and losses*, and only consider the time T it takes one message to disseminate through the network with high probability. Of course, MAC protocols affect the value of T. One protocol may allow for only one node to transmit per time slot, while another may intelligently schedule multiple nodes together as long as they do not interfere with each other. An even more sophisticated scheme may use multi-user coding and decoding schemes to allow interfering nodes to transmit as well. Each protocol may result in a different flooding time T. The strength of the results herein is in computing the time it takes to disseminate multiple messages, correlated messages, or messages to nodes with correlated data *given* T. While it is not easy to compute T for any network model and MAC protocol, it is possible for several interesting scenarios (e.g., random phone call model or bounded degree graphs). Moreover, any meaningful bound on T would result in meaningful bounds on the stopping time for correlated scenarios as well.

V. Simple and Direct Proofs for Tight Stopping Times

In this section we give a simple, direct derivation of tight stopping time bounds for gossip with one source and side information at the nodes and gossip with two correlated sources. Our two main results in this section are the following.

Theorem 2. *Suppose M is an oblivious network that floods in time T with throughput α. For any error probability $\epsilon > 0$, constant $\delta > 0$, $l = \Omega(\log \epsilon^{-1})$, packet size s and distributions on $X, \{Y_v\}_{v \in V}$, suppose the network is initialized with a single message x generated from X and side information y_w generated from Y_w at every node w. Then, every node v will correctly decode x with probability at least $1-\epsilon$ after $T' = T + \frac{1}{\alpha}\left(\left\lceil \frac{l}{s}(H(X|Y_v) + \delta)\right\rceil + \log \epsilon^{-1} + 3\right)$ rounds.*

Theorem 3. *Suppose M is an oblivious network that floods in time T with throughput α. For any error probability $\epsilon > 0$, constant $\delta > 0$, $l = \Omega(\log \epsilon^{-1})$, packet size s and distributions on X_1, X_2, suppose the network is initialized with two messages x_1, x_2 generated from X_1, X_2*

and nodes have no side information. Then, every node v will correctly decode x with probability at least $1-\epsilon$ after $T+\frac{1}{\alpha}\left(\left\lceil\frac{l}{s}(H(X_1,X_2)+2\delta)\right\rceil+\log\epsilon^{-1}+3\right)$ rounds.

At the heart of the proofs of Theorem 2 and Theorem 3 is a generalization of the main observation in [1], which states that the question of when a node can decode is equivalent to determining when this node knows (see Definition 1) enough coefficient vectors. The proof of Theorem 1 shows that flooding or spreading of knowledge of vectors can be analyzed using a union bound. This implies that only the number of vectors needed is of importance. In the case with *uncorrelated* sources and no side information essentially knowledge of *all coefficient vectors* is necessary. In the correlated scenario, however, we want to relate the number of vectors a node v needs to know to the conditional entropy $H(X|Y_v)$. Lemma 1 helps in this respect. It asserts that in order to decode, a node with side information Y does not necessarily need $i=\lceil(H(X|Y)+\delta)l\rceil$ *specific bits*, but rather, assuming joint typicality decoding, it requires only *any sufficient amount* of information about the index of the bin in which x resides. This is achieved by *any* i/s packets containing independent equations on the bin index. We can thus focus on the knowledge a node is required to obtain in order for its coefficient matrix to have rank at least i/s, rather than the requirement to receive certain bits. This viewpoint also concurs with the lack of network-source coding separation for more than two terminals [36], which states that in general, optimal (cut-set achieving) rates cannot be achieved with separate source and network coding.

We note here that while the *gossip protocol* we use, algebraic gossip, winds up by a node sending a random linear combination of its data, analysis is inherently different than that of random linear network coding. For example, unlike the progression of the *rank* at a node when receiving independent linear equations, it is possible that a node *knows* many vectors *without having a large rank*. This is at the heart of the projection analysis herein. In fact, upon reception of the *first packet* (assuming it is non-zero) a node gets to know *all but a $1/q$ fraction* of all vectors. Then, the more independent packets arrive, the *lesser* is the marginal knowledge gained by the node. This is different than the linear scaling of the rank as a function of the number of independent packets. On the other hand, in order to prove faster stopping times we argue that the knowledge of only *an exponentially small* fraction of all vectors suffices for decoding when the node has side information or data is correlated. Rigorously, this is summarized by the following lemma, which states that indeed only q^l *specific coefficient vectors* suffice to guarantee that at least l independent coefficient vectors were received:

Lemma 2. *Let $V=\mathcal{F}_q^k$ be a k-dimensional vector space over a finite field $\mathcal{F}_q$. For every $0\leq h<k$ there exist $w=2^h+1$ vectors $v_1,\ldots,v_w\in V$ such that any subspace K of V for which for every i the vector $v_i\notin K^\perp$ has dimension greater than h.*

Proof: It suffices to show that for the vectors $v_1,\ldots,v_{2^h}$ one can choose any subspace W of dimension h and for the last vector v_w any additional vector. We will assume, without loss of generality, that v_w is perpendicular to W. To prove that this is indeed so, we take any subspace K of dimension at most h and show that its orthogonal complement $K^\perp$ intersects with $W\cup v_w$. Note that $K^\perp$ has dimension at least $k-h$. Thus if $K^\perp$ and W do not intersect one can find a set of k orthogonal basis vectors for $\mathcal{F}_q^k$ of which h span W and $k-h$ span $K^\perp$. Since v_w is orthogonal to W, it is spanned by the $k-h$ vectors in $K^\perp$ and thus $v_w\in K^\perp$ – a contradiction. ∎

It is now possible to prove the two main results of this section. Their proofs are self-contained and involve only random binning (that is, Lemma 1) and Lemma 2.

Proof (Theorem 2): Fix field size q which achieves parameters T and α in Definition 4. Choose l large enough so the decoding error probability is smaller than $\epsilon/2$.

Let x be the message vector, $\hat{x}$ the reconstructed source and $i(x)$ be the bin index of x. By Lemma 1, any node v, having access to y_v and $\left\lceil\frac{l}{s}(H(X|Y_v)+\delta)\right\rceil$ independent equations on the blocks of $i(x)$, can decode with high probability, namely, $P[\hat{x}\neq x]\leq\epsilon/4$. By Lemma 2, there exists a set Y of $2^{\left\lceil\frac{l}{s}(H(X|Y_v)+\delta)\right\rceil}+1$ coefficient vectors such that if v has knowledge on these vectors, it has $\left\lceil\frac{l}{s}(H(X|Y_v)+\delta)\right\rceil$ independent equations on the blocks of $i(x)$. Knowledge of any coefficient vector in Y spreads like a flooding process. Hence, the probability that any of the coefficient vectors is not known after $T+\frac{1}{\alpha}(|Y|+\log\epsilon^{-1})$ rounds is smaller than ϵ. Consequently, after $T'=T+\frac{1}{\alpha}\left(\left\lceil\frac{l}{s}(H(X|Y_v)+\delta)\right\rceil+\log\epsilon^{-1}+1\right)$ rounds the probability that all nodes received all vectors in Y is at least $1-\epsilon/2$ and we have:

$$\begin{aligned}
P[\hat{x}\neq x] &= P\left[\hat{x}\neq x|dim(S_v)<\left\lceil\frac{l}{s}(H(X|Y_v)+\delta)\right\rceil\right]\\
&\quad\cdot P\left[dim(S_v)<\left\lceil\frac{l}{s}(H(X|Y_v)+\delta)\right\rceil\right]\\
&\quad+P\left[\hat{x}\neq x|dim(S_v)\geq\left\lceil\frac{l}{s}(H(X|Y_v)+\delta)\right\rceil\right]\\
&\quad\cdot P\left[dim(S_v)\geq\left\lceil\frac{l}{s}(H(X|Y_v)+\delta)\right\rceil\right]\\
&\leq P\left[dim(S_v)<\left\lceil\frac{l}{s}(H(X|Y_v)+\delta)\right\rceil\right]\\
&\quad+P\left[\hat{x}\neq x|dim(S_v)\geq\left\lceil\frac{l}{s}(H(X|Y_v)+\delta)\right\rceil\right].
\end{aligned}$$

This results in

$$\begin{aligned}
P[\hat{x}\neq x] &\leq P\left[\cup_{\mu\in Y}\mu\perp S_v\right]+\epsilon/4\\
&\leq (2^{\left\lceil\frac{l}{s}(H(X|Y_v)+\delta)\right\rceil}+1)\\
&\quad\cdot 2^{\alpha\left[-\frac{1}{\alpha}\left(\left\lceil\frac{l}{s}(H(X|Y_v)+\delta)\right\rceil+\log\epsilon^{-1}+1\right)\right]}+\frac{\epsilon}{4}\\
&< \epsilon.
\end{aligned}$$

∎

Before we prove Theorem 3, note that in the single source model of Theorem 2, for each terminal there is only one rate constraint: $r\geq H(X|Y_v)+\delta$. Via Lemma 2, it is translated into a *rank constraint*, i.e., $dim(S_v)\geq\left\lceil\frac{l}{s}(H(X|Y_v)+\delta)\right\rceil$. For more than one source, however, the rate region is defined by multiple constraints. Fortunately, for two sources and no side

information at the nodes, satisfying them can be guaranteed using a single rank constraint.

Proof (Theorem 3): Let x_1, x_2 be blocks of l symbols from the sources X_1 and X_2. We assume x_1 is available at source node s_1 and x_2 is available at source node s_2. At source node s_j, $j \in \{1,2\}$, we randomly bin all x_j sequences to $2^{\lceil l(H(X_j)+\delta)\rceil}$ bins. Let $i_j(x_j)$ be the index of the bin in which x_j resides. Similar to the proof of Theorem 2, $i_j(x_j)$ is represented as symbols over $\mathcal{F}_q$, where each packet (sent as a message in a round) includes $\frac{s}{\log q}$ equations on these symbols, over $\mathcal{F}_q$. Note, however, that there is a difference in the binning rate compared to the proof of Theorem 2. Herein, we cannot bin at a rate higher than $\lceil l(H(X_j)+\delta)\rceil$, since we want to make sure if a node received enough independent equations on the symbols of both $i_1(x_1)$ and $i_2(x_2)$, these must have included enough equations on the symbols of $i_1(x_1)$, and enough equations on the symbols of $i_2(x_2)$, according to the Slepian-Wolf limits.

The sources now spread uniformly randomly chosen linear combinations on the symbols of $i_1(x_1)$ and $i_2(x_2)$. Throughout the gossip rounds, these linear combinations are, of course, mixed together. We now verify that when a terminal receives sufficiently many independent linear combinations, these include enough linear combinations on each bin index. Assume that a terminal t has received $\left\lceil \frac{l}{s}(H(X_1,X_2)+2\delta)\right\rceil + 1$ independent equations on the blocks describing $i_1(x_1)$ and $i_2(x_2)$. Since at most $\left\lceil \frac{l}{s}(H(X_1)+\delta)\right\rceil$ originated from source s_1, we have

$$\left\lceil \frac{l(H(X_1,X_2)+2\delta)}{s}\right\rceil + 1 - \left\lceil \frac{l(H(X_1)+\delta)}{s}\right\rceil \geq \left\lceil \frac{l(H(X_2|X_1)+\delta)}{s}\right\rceil$$

which means sufficiently many were originated in s_2. As a result, the single constraint to receive $\left\lceil \frac{l}{s}(H(X_1,X_2)+2\delta)\right\rceil + 1$ independent equations suffices.

Now, fix $\epsilon > 0$. Let $\hat{x_j}$ denote the reconstructed source vector x_j. For the case of multiple sources, Lemma 1 extends trivially, stating that for sufficiently large l, any node t, having access to at least $\left\lceil \frac{l}{s}(H(X_1|X_2)+\delta)\right\rceil$ independent equations on the blocks describing $i_1(x_2)$, at least $\left\lceil \frac{l}{s}(H(X_2|X_1)+\delta)\right\rceil$ independent equations on $i_2(x_2)$, and at least $\left\lceil \frac{l}{s}(H(X_1,X_2)+\delta)\right\rceil$ independent equations on both, one can decode with high probability. By Lemma 2, there exists a set Y of $2^{\left\lceil \frac{l}{s}(H(X_1,X_2)+2\delta)\right\rceil+1}+1$ coefficient vectors such that if t has knowledge on these blocks, it indeed has $\left\lceil \frac{l}{s}(H(X_1,X_2)+2\delta)\right\rceil + 1$ independent equations, hence, after $T + \frac{1}{\alpha}\left(\left\lceil \frac{l}{s}(H(X_1,X_2)+2\delta)\right\rceil + 1 + \log(2/\epsilon)\right)$ rounds, we have,

$$\begin{aligned} P[\hat{x_1},\hat{x_2} \neq x_1,x_1] \leq & P\left[dim(S_v) < \left\lceil \frac{l}{s}(H(X_1,X_2)+2\delta)\right\rceil + 1\right] \\ & + P\Big[\hat{x_1},\hat{x_2} \neq x_1,x_1 \Big| \\ & \quad dim(S_v) \geq \left\lceil \frac{l}{s}(H(X_1,X_2)+2\delta)\right\rceil + 1\Big]. \end{aligned}$$

As a result,

$$\begin{aligned} P[\hat{x_1},\hat{x_2} \neq x_1,x_1] \leq & P\left[\cup_{i=1}^{2^{\left\lceil \frac{l}{s}(H(X_1,X_2)+2\delta)\right\rceil+1}+1} v_i \perp S_v\right] + \epsilon/2 \\ \leq & \left(2^{\left\lceil \frac{l}{s}(H(X_1,X_2)+2\delta)\right\rceil+1}+1\right) \\ & \cdot 2^{\alpha\left[-\frac{1}{\alpha}\left(\left\lceil \frac{l}{s}(H(X_1,X_2)+2\delta)\right\rceil+1+\log(2/\epsilon)\right)\right]} + \epsilon/2, \end{aligned}$$

which is smaller than or equal to ϵ. ∎

For more than two sources, the above results and analysis do not apply directly, and a much more refined analysis is required. The reason is as follows. When only two information sources are available, one can bin the first source at rate $H(X_1)$ and the second at rate $H(X_2)$. We assume, for simplicity, two correlated sources and no side information. We also omit the ϵ, δ notation for brevity. Since $H(X_1,X_2) = H(X_1) + H(X_2|X_1) = H(X_2) + H(X_1|X_2)$, if a terminal received information at rate which is *at least* $H(X_1,X_2)$, it is guaranteed that at least $H(X_1|X_2)$ were received from the first and at least $H(X_2|X_1)$ were received from the second. Thus, one should only guarantee *the sum rate is adequate*, regardless of the constraints on the individual sources. However, for more than two sources, e.g., tree, it is easy to see that ensuring a rate of $H(X_1,X_2,X_3)$ *at the terminal* is not enough to guarantee all the individual and sub-set constraints, e.g., that at least $H(X_1|X_2,X_3)$ is received from the first source and at least $H(X_2,X_3|X_1)$ is received from the second and third together.

To make sure these individual and sub-set constraints are met *under a gossip scheme*, one has to make sure the scheme (implicitly) provides enough capacity for these paths, rather than considering only the total rate. Thus, a refined analysis of the capacities seen under a gossip scheme is required. This is the subject of the next section.

VI. Capacities in Oblivious Network Models

To date, analysis of gossip schemes focused only on the dissemination time – the number of rounds required to gain the complete knowledge in the network. Especially when network coding is discussed, this results in an *all-or-nothing* viewpoint. However, when dynamic networks are analyzed, and, to a greater extent, when correlated data is available, it is essential to gain a more accurate measure of the *actual capacities achievable between sets of nodes*. Namely, to analyze the capacities induced by the gossip packet exchange process. This is an interesting question in its own right, and, in particular, can give a "black-box" method to transfer any results of prior works that bound the rates or capacities needed between sources and sinks in the static memory-free setting to stopping times in oblivious network models.

In this section, we develop such a characterization of the capacities, and apply it to the results of [2] and [3] to obtain stopping times for gossip protocols with an *arbitrary number of correlated sources* and side-information. Besides generalizing Theorem 2 and Theorem 3, this characterization gives a viewpoint on the rates achievable with gossip schemes which can be found useful in a variety of problems. We first introduce the required notation.

Definition 5. Let $T > 0$, node set V and active edges $E_1, E_2, \ldots, E_T$ be given. *A path* P from s to d is a sequence of nodes $P = (v_0, v_1, \ldots, v_T)$ such that $v_0 = s$, $v_T = d$ and for every $t \leq T$ we have $v_{t-1} = v_t$ or $(v_{t-1}, v_t) \in E_t$.

Definition 6. A set of m paths with weights $w_1, \ldots, w_m$ *is valid* if for every $t < T$ and every $(u, v) \in E_t$ the weights of paths using (u, v) sum to at most one. We say a set of valid weighted paths *achieves a capacity* of c between two nodes s and d if the weights of paths from s to d sum up to c.

Quite intuitively, these paths correspond to an information flow through the network from the sources to the sinks; or, alternatively, to a (fractional) network flow in a time-expanded (hyper-)graph or trellis. This intuition was made formal in [37] which proved an explicit equivalence between the algebraic gossip protocol and random linear network coding in the classical memory-free model (e.g., [2]). Roughly speaking, network-coded gossip can be viewed as simple random linear network coding on a time-expanded graph. The directed time-expanded hypergraph G_{PNC} that corresponds to a sequence of active edges in a gossip scheme can be found in [37]. However, we omit the details of this equivalence and instead only recall the facts that are needed in this paper. Specifically, let node set V, the active edges $E_1, \ldots, E_T$, destination $d \in V$ and sources $s_1, s_2, \ldots, s_k \in V$ be given. Then, algebraic gossip on $\{E_i\}_{i=1}^T$ is equivalent to classical random linear network coding in the transformed hypergraph G_{PNC} described in [37]. In particular, we have the following result regarding the actual capacities available from the sources to the destination.

Lemma 3. *Assume an algebraic gossip model with destination $d \in V$ and sources $s_1, s_2, \ldots, s_k \in V$. If for some integers $c_1, \ldots, c_k$, it is possible for every s_i to transmit c_i packets to d, then there exists a sequence of valid paths of weight one and a rate c_i between s_i and d. Conversely, if for some positive reals $c_1, \ldots, c_k$ there is a set of valid weighted paths that achieve a capacity c_i between s_i and d, then the capacities c_i lie in the min-cut region of G_{PNC}.*

Given this setup, we show the first result in this direction:

Lemma 4. *Let M be a network on a node set V that floods in time T with throughput α. For any T', $\epsilon > 0$, destination $d \in V$ and set of k source nodes $s_1, s_2, \ldots, s_k \in V$ with integral capacities $c_1, c_2, \ldots, c_k \geq 1$, suppose $E_1, \ldots, E_{T'}$ is a sequence of active edges on V sampled from M. If $T' \geq T + \frac{1}{\alpha}(\sum_i c_i + \log \epsilon^{-1})$ then with probability at least $1 - \epsilon$ there exists a selection of valid weighted paths that achieve the capacity c_i between s_i and t for every i.*

Proof: We think of putting c_i messages at node s_i and run the standard algebraic gossip protocol for T' rounds using the field size q that achieves the parameters T and α on M. Theorem 1 now shows that with probability $1 - \epsilon$ the sink t can decode the messages. According to Lemma 3 this shows that there are c_i mutually disjoint paths from s_i to d for every i with weight one which achieve the desired capacities. ∎

Note that the above lemma requires the capacities to be integral and thus essentially asks for the time until a certain number of mutually disjoint paths occur. While this is sufficient and optimal in the *uncorrelated* information spreading setting, this requirement can be a severe restriction in the *correlated data model*.

One setting where this makes a drastic difference is when we have k sources and the total capacities needed sum up to less than one. This corresponds to asking for the time until there is a path from each of the sources to the sink – without these paths having to be disjoint. If one considers for example the random phone call model with n nodes and k sources it takes in expectation $\log n + k$ time until a disjoint path between a node and each source appears while merely $\log n + \log k$ rounds are sufficient to get this for non-disjoint paths.

The following lemma generalizes this observation and strengthens Lemma 4 in this direction to give order optimal bounds for any set of fractional capacities:

Lemma 5. *Let M be a network model on a node set V that floods in time T with throughput α. For any T', any $\epsilon > 0$, any sink $d \in V$ and any set of k source nodes $s_1, s_2, \ldots, s_k \in V$ with rates $c_1, c_2, \ldots, c_k > 0$, if $T' \geq T + \frac{1}{\alpha}(\lceil \sum_i c_i \rceil + \log k + \log \epsilon^{-1})$ then with probability at least $1 - \epsilon$ there exists a selection of valid weighted paths that achieve a capacity of c_i between s_i and d for every i.*

Proof: The key idea in the proof is to combine k applications of Lemma 4 using a union bound and capacity sharing. Set the failure probability to ϵ/k and in the ith application of Lemma 4 set the c_i to $\lceil \sum_i c_i \rceil$ while setting all other capacities to zero. As a result, for every i with probability $1 - \epsilon/k$ the number of disjoint paths from s_i to d is at least $\lceil \sum_i c_i \rceil$. Via a union bound, with probability of $1 - \epsilon$ all these paths exist. Yet, while the paths from each source are disjoint, the paths starting at different sources may not be disjoint. We now take the union of these paths while choosing the weight of a paths starting at source s_i to be $\frac{c_i}{\lceil \sum_j c_j \rceil}$. This gives capacity of c_i between s_i and d. Furthermore, since any edge e is used by at most one path going out from each source, we get that the total weight on e summed over all paths is at most $\sum_i \frac{c_i}{\lceil \sum_j c_j \rceil} \leq 1$. ∎

VII. Stopping Times for Networks with Multiple Sources and Side Information

Finally, We use Lemma 5 above to prove our main result about information dissemination with correlated data in oblivious networks. To state our result formally, we first need the following definition.

Definition 7 (Slepian-Wolf region [4])**.** A capacity vector $c = (c_1, \ldots, c_k)$ is sufficient for $v \in V$ if and only if for every index subset $S \in [k]$ we have $\sum_{i \in S} c_i \geq H(X_S | X_{\overline{S}}, Y_v)$.

Putting together Lemma 5, Lemma 3 and applying the results on network coding with correlated data from [2] and [3], we can now directly state our main result which generalizes and encompasses Theorem 3 and Theorem 2.

Theorem 4. *Suppose M is an oblivious network model that floods in time T with throughput α. For any constant $\delta > 0$, error probability $\epsilon > 0$, $l = \Omega(\log \epsilon^{-1})$, joint distribution of*

$X_1, \ldots, X_k$ and the Y_v's, packet size $s > 0$, node v and any capacity vector $(c_1, \ldots, c_k)$ that is sufficient for v, node v will correctly decode $x_1, \ldots, x_k$ with probability at least $1-\epsilon$ after $T + \frac{1}{\alpha}\left(\left\lceil \frac{l}{s}\sum_{i\in[k]} c_i + \delta \right\rceil + \log k + \log \epsilon^{-1} + \delta\right)$ rounds.

Proof: Use Lemma 5 to show that in T' rounds the rate vector $(r_1, \ldots, r_k)$ is achieved by a collection of paths with probability at least $1-\epsilon/2$. Then apply Lemma 3 to show that the rates are min-cut rates in G_{PNC}, the classical network corresponding to the sampled topologies. Finally we can now directly apply the known results on network coding with correlated data from [9] and [3] in a black-box manner to show that an arbitrary small decoding probability can be achieved if l is chosen large enough. The only thing to check here is that the complete coding scheme described here matches the setup of [9], [3]. In particular the field size q used needs to be large enough to make these results applicable. ■

Note that the suggested solutions throughout this work are based on joint network-source coding [36]. That is, the coding at the sources is done using some kind of a binning scheme, either random or structured (e.g., [38]) and the *exchange of linearly-coded packets* performs both the standard network coding and the actual distributed compression. Such joint network-source codes can, of course, be decoded using joint typicality as in the original Spelian-Wolf model [4] or using minimum entropy methods similar to [9]. While both methods are computationally intensive, it is important to note that efficient joint codes exist in several scenarios, e.g., [39], [40]. Still, the general case is computationally expensive and remains a fascinating future research direction.

VIII. Simulations

Simulations of the Gossip process in networks with side information or networks with correlated sources were carried out in Matlab. Unless mentioned otherwise, each point was averaged over 25 runs with the same parameters. Underlying field was $GF(2^2)$.

We first illustrate the flooding process defined in Definition 3. Figure 4 depicts the results for the random phone call model (i.e., a complete graph where in each round, each transmitter selects a random receiver), with varying values of the fault probability p. In this case, $1-p$ is simply the probability of an *informed* node to send a message in a round. The edge set E_t includes an edge with probability $1/n$, namely, on average, each node selected to send a message indeed informs only a single random receiver per round. Due to the logarithmic scale of the horizontal axis, the logarithmic behaviour of the stopping time is clearly visible.

Figure 5 depicts the stopping times for k independent messages distributed randomly on k nodes of a complete graph (random phone call model). Again, the logarithmic behavior is clearly visible. Moreover, it is easy to see that the increase in k results only in an additive factor (hence the parallel lines for different values of k), as predicted by Theorem 1.

Figure 6 gives a different perspective, where the $\log \epsilon^{-1}$ term can be evaluated. Specifically, the probability of *not all nodes receiving all messages* is depicted, as a function of the number of rounds. The setting is similar to Figure 5, with

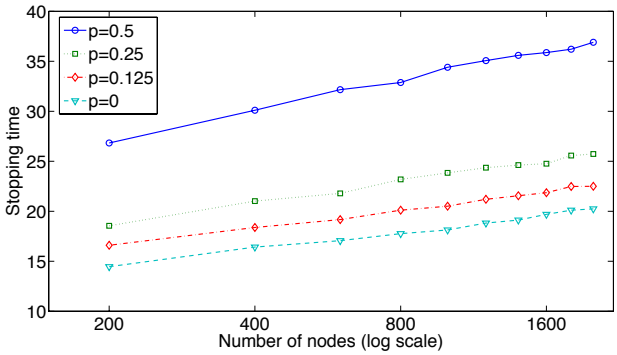

Fig. 4: Stopping times for the flooding process of a single message in the random phone call model, as a function of the number of nodes. The logarithmic increase is clearly visible for any value of the fault probability p.

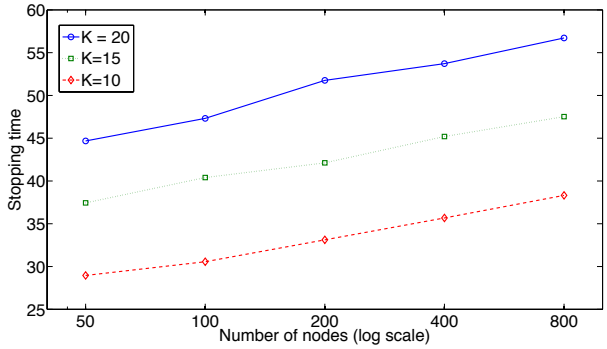

Fig. 5: Stopping times for k independent messages in a complete graph. Increasing the number of messages, k, results in the same increase in the stopping time for all network sizes.

40 runs per point. The steep decrease in ϵ as the number of rounds increases over a certain threshold (the flooding time when k is small) is clearly visible. For example, for all k, ϵ is reduced by a factor of about 2^5 in 8–10 rounds. It is also clear that increasing k increases the number of rounds required by approximately the same amount regardless of ϵ. The exponential decrease is the reason why stopping times are *concentrated around their mean value*.

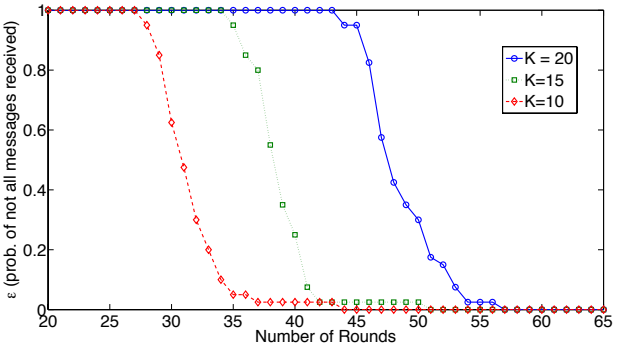

Fig. 6: The probability (ϵ) of *not all messages disseminated to all nodes* as a function of the number of rounds. $n = 100$.

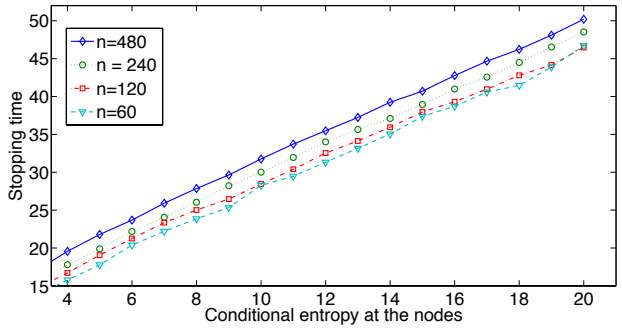

Fig. 7: Stopping times for a single source and side information at the nodes. It is clear that the stopping time increases linearly with the conditional entropy (the weaker the side information is). Larger networks only increase the flooding time T, which raises the graphs by a constant (in $H(X|Y_i)$) factor.

Figure 7 corresponds to the setting in Theorem 2. Namely, a single source with message X and entropy $H(X)$. Node $i \in V$ has side information Y_i. The simulations include the communication process, that is, the network coding based information spreading throughout the network. We assume decoding is successful once a node reached the required rank (see Lemma 1 and the discussion which follows). Hence, we normalize $H(x)$ and $H(X|Y_i)$ such that the information at the source equals 20 independent elements in the field, and the conditional entropy at the nodes varies from 4 (having significant knowledge) to 20 (having no knowledge at all). Points in the plot are averaged over all nodes having the same conditional entropy. The linear increase as a function of the conditional entropy is clear. Note also that when the size of the network, n, increases, this results in a constant increase in the stopping time throughout all values of $H(X|Y_i)$. This is in accordance with Theorem 2, which clearly states that the stopping time is the fixed flooding time T, plus a linear term in $H(X|Y_i)$.

Figure 8 includes the same information model, however, on a dynamic, location-based graph. Nodes are randomly spread in space, and the connectivity is a random function of the distance between the nodes, changing independently for each communication round. Specifically, nodes close by will always be able to communicate, distant nodes have no direct connection, and communication at intermediate distances is sporadic. It is clear that the linear trend in the conditional entropy remains for this dynamic setting as well. However, it is important to note that if nodes' locations are fixed throughout the simulation (diamond-shaped markers), stopping times vary, as a node's location might have a non-negligible effect, while if we average over a few permutations of the nodes' locations (circle markers), stopping times are more concentrated near the expected linear trend.

Finally, Figure 9 corresponds to the setting in Theorem 3. Using the same decoding assumptions as in Figure 7, the network includes two *correlated* sources. When source correlation is maximal (i.e., $H(X_1) = H(X_2) = H(X1, X_2) = 8$), stopping time is low. As the correlation decreases (up to $H(X_1, X_2) = H(X_1) + H(X_2)$), stopping time increases. The

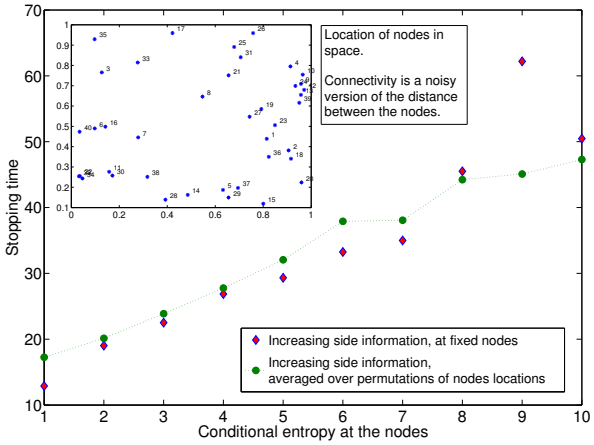

Fig. 8: Stopping times for a single source and side information at the nodes in a dynamic, location-based graph. The sub-plot depicts nodes' locations in space. Diamond-shaped markers depict the case where the amount of side information is fixed throughout the simulations, while circle-shaped markers are for an average over a few simulations, each with newly placed Y_i's.

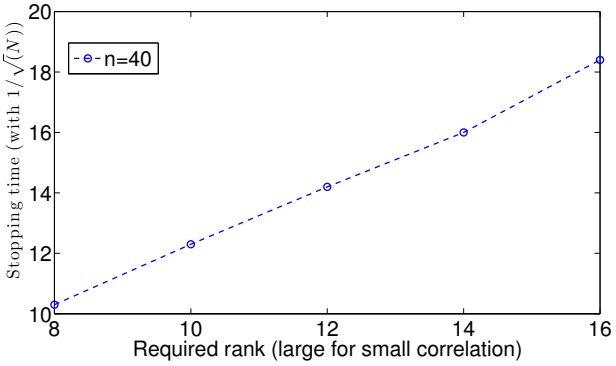

Fig. 9: Stopping time for a complete graph of size 40 with two correlated sources. The horizontal axis is the rank required at a node in order to successfully decode the two sources. When correlation is maximal (identical sources), it is 8. When the sources are independent, it is 16.

plot clearly depict the linear increase in the stopping time, as a function of the rank required for successful decoding. In this simulation, the probability of an edge to be in E_t in each round is $1/\sqrt{n}$, hence the shorter stopping times compared to previous simulations on complete graphs where this probability was $1/n$ (random phone calls).

IX. Conclusion

In this work, we designed and analyzed gossip schemes for networks with correlated data. First, we formally defined the concept of oblivious networks. This network model allows to include many of the dynamics and connectivity features of

wireless networks and other dynamic settings. In particular it allowed us to analyze network-coded gossip schemes in a setting which includes a wide variety of scenarios, including collisions, packets losses, broadcast and many other aspects of wireless networks.

We then defined the concept of partial knowledge required to decode using side information or correlated data in the context of linear network coding. Using the results on network capacities under gossip schemes, this allowed us to give tight bounds on stopping times of network-coded gossip in oblivious networks with correlated data and side information.

Finally, we performed an extensive set of simulations, validating the key trends predicted by the analysis, e.g., the dependence of the stopping time on the number of messages or strength of side information, the exponential decrease of the fault probability ϵ and the robustness of the results even in dynamic, location-based models.

REFERENCES

[1] B. Haeupler, "Analyzing network coding gossip made easy," in *Proceedings of the ACM Symposium on Theory of Computing*. ACM, 2011, pp. 293–302.

[2] T. Ho, M. Médard, R. Koetter, D. R. Karger, M. Effros, J. Shi, and B. Leong, "A random linear network coding approach to multicast," *IEEE Transactions on Information Theory*, vol. 52, no. 10, pp. 4413–4430, 2006.

[3] M. Bakshi and M. Effros, "On achievable rates for multicast in the presence of side information," in *Proceedings of the International Symposium on Information Theory*, 2008, pp. 1661–1665.

[4] D. Slepian and J. Wolf, "Noiseless coding of correlated information sources," *IEEE Trans. Inform. Theory*, vol. 19, no. 4, pp. 471–480, July 1973.

[5] R. Ahlswede and J. Körner, "Source coding with side information and a converse for degraded broadcast channel," *IEEE Trans. Inform. Theory*, vol. 21, no. 6, pp. 629–637, November 1975.

[6] I. Csiszar and J. Körner, "Towards a general theory of source networks," *IEEE Trans. Inform. Theory*, vol. 26, no. 2, pp. 155–165, March 1980.

[7] T. S. Han and K. Kobayashi, "A unified achievable rate region for a general class of multiterminal source coding systems," *IEEE Trans. Inform. Theory*, vol. 26, no. 3, pp. 277–288, May 1980.

[8] W.-H. Gu and M. Effros, "Source coding for a multihop network," in *Proc. DCC*, Snowbird, Utah, March 2005.

[9] T. Ho, M. Médard, M. Effros, and R. Koetter, "Network coding for correlated sources," in *CISS*, 2004.

[10] J. Barros and S. D. Servetto, "Network information flow with correlated sources," *IEEE Trans. Inform. Theory*, vol. 52, no. 1, pp. 155–170, January 2006.

[11] A. Cohen, S. Avestimehr, and M. Effros, "On networks with side information," in *Proceedings of the IEEE International Symposium on Information Theory*, 2009, pp. 1343–1347.

[12] J. Haupt, W.U. Bajwa, M. Rabbat, and R. Nowak, "Compressed sensing for networked data," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 92–101, 2008.

[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 *Proceedings of the ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing*, 1987, p. 12.

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

[15] D. Kempe, A. Dobra, and J. Gehrke, "Gossip-based computation of aggregate information," in *FOCS*, 2003, pp. 482–491.

[16] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Randomized gossip algorithms," *IEEE/ACM Transactions on Networking (TON)*, vol. 14, no. SI, pp. 2508–2530, 2006.

[17] S. Deb, M. Médard, and C. Choute, "Algebraic gossip: a network coding approach to optimal multiple rumor mongering," *IEEE/ACM Trans. Networking*, vol. 14, pp. 2486–2507, June 2006.

[18] D. Mosk-Aoyamam and D. Shah, "Information dissemination via network coding," in *Proceedings of the International Symposium on Information Theory*, 2006, pp. 1748–1752.

[19] M. Borokhovich, C. Avin, and Z. Lotker, "Tight bounds for algebraic gossip on graphs," jun. 2010, pp. 1758 –1762.

[20] Chen Avin, Michael Borokhovich, Keren Censor-Hillel, and Zvi Lotker, "Order optimal information spreading using algebraic gossip," *Distributed computing*, vol. 26, no. 2, pp. 99–117, 2013.

[21] B. Haeupler, "Tighter worst-case bounds on algebraic gossip," *Communications Letters, IEEE*, vol. 16, no. 8, pp. 1274–1276, August 2012.

[22] Bernhard Haeupler, Asaf Cohen, Chen Avin, and Muriel Médard, "Network coded gossip with correlated data," in *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on*. IEEE, 2012, pp. 2616–2620.

[23] R. Ahlswede, N. Cai, S-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Trans. Inform. Theory*, vol. 46, no. 4, pp. 1204–1216, July 2000.

[24] R. Koetter and M. Médard, "An algebraic approach to network coding," *IEEE/ACM Trans. on Networking*, vol. 11, no. 5, pp. 782–794, October 2003.

[25] Daniel Enrique Lucani, Frank HP Fitzek, Muriel Médard, and Milica Stojanovic, "Network coding for data dissemination: it is not what you know, but what your neighbors don't know," in *Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks, 2009. WiOPT 2009. 7th International Symposium on*. IEEE, 2009, pp. 1–8.

[26] Péter Vingelmann, Morten Videbæk Pedersen, Janus Heide, Qi Zhang, and Frank HP Fitzek, "Data dissemination in the wild: A testbed for high-mobility manets," in *Communications (ICC), 2012 IEEE International Conference on*. IEEE, 2012, pp. 291–296.

[27] Soheil Feizi and Muriel Médard, "A power efficient sensing/communication scheme: Joint source-channel-network coding by using compressive sensing," in *Communication, Control, and Computing (Allerton), 2011 49th Annual Allerton Conference on*. IEEE, 2011, pp. 1048–1054.

[28] B. Haeupler and M. Medard, "One packet suffices - highly efficient packetized network coding with finite memory," in *ISIT*, 2011, pp. 1151 –1155.

[29] Chen Avin, Michael Borokhovich, Keren Censor-Hillel, and Zvi Lotker, "Order optimal information spreading using algebraic gossip," in *PODC*, 2011, pp. 363–372.

[30] A. Clementi, C. Macci, A. Monti, F. Pasquale, and R. Silvestri, "Flooding time in edge-markovian dynamic graphs," in *PODC*, 2008, pp. 213–222.

[31] C. Avin, M. Koucký, and Z. Lotker, "How to explore a fast-changing world (cover time of a simple random walk on evolving graphs)," *Proceedings of the International Colloquium on Automata, Languages and Programming*, pp. 121–132, 2008.

[32] A. Clementi, A. Monti, F. Pasquale, and R. Silvestri, "Information spreading in stationary markovian evolving graphs," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 22, no. 9, pp. 1425–1432, 2011.

[33] Konstantinos Oikonomou and Ioannis Stavrakakis, "Performance analysis of probabilistic flooding using random graphs," in *World of Wireless, Mobile and Multimedia Networks, 2007. WoWMoM 2007. IEEE International Symposium on a*. IEEE, 2007, pp. 1–6.

[34] Fabian Kuhn and Rotem Oshman, "Dynamic networks: models and algorithms," *ACM SIGACT News*, vol. 42, no. 1, pp. 82–96, 2011.

[35] Andrea Clementi, Angelo Monti, and Riccardo Silvestri, "Fast flooding over manhattan," *Distributed computing*, vol. 26, no. 1, pp. 25–38, 2013.

[36] A. Ramamoorthy, K. Jain, P. A. Chou, and M. Effros, "Separating distributed source coding from network coding," *IEEE Trans. Inform. Theory*, vol. 52, pp. 2785–2795, June 2006.

[37] B. Haeupler, M. Kim, and M. Medard, "Optimality of Network Coding with Buffers," in *ITW*, 2011, pp. 533–537.

[38] R. Zamir, S. Shamai, and U. Erez, "Nested linear/lattice codes for structured multiterminal binning," *IEEE Transactions on Information Theory*, vol. 48, no. 6, pp. 1250–1276, 2002.

[39] G. Maierbacher, J. Barros, and M. Médard, "Practical source-network decoding," in *Proceedings of the International Symposium on Wireless Communication Systems*, 2009, pp. 283–287.

[40] C. Avin, M. Borokhovich, A. Cohen, and Z. Lotker, "Efficient Distributed Source Coding for Multiple Receivers Via Matrix Sparsification," in *Proceedings of the IEEE International Symposium on Information Theory*, 2011, pp. 2045–2049.