

## FAST INFORMATION SPREADING IN GRAPHS WITH LARGE WEAK CONDUCTANCE\*

KEREN CENSOR-HILLEL<sup>†</sup> AND HADAS SHACHNAI<sup>‡</sup>

**Abstract.** Gathering data from nodes in a network is at the heart of many distributed applications, most notably while performing a global task. We consider *information spreading* among  $n$  nodes of a network, where each node  $v$  has a message  $m(v)$  which must be received by all other nodes. The time required for information spreading has been previously upper-bounded with an inverse relationship to the conductance of the underlying communication graph. This implies high running time bounds for graphs with small conductance. The main contribution of this paper is an information spreading algorithm which overcomes communication bottlenecks and thus achieves fast information spreading for a wide class of graphs, despite their small conductance. As a key tool in our study we use the recently defined concept of *weak conductance*, a generalization of classic graph conductance which measures how well-connected the components of a graph are. Our hybrid algorithm, which alternates between random and deterministic communication phases, exploits the connectivity within components by first applying *partial information spreading*, in which information is exchanged within well-connected components, and then sending messages across bottlenecks, thus spreading further throughout the network. This yields substantial improvements over the best known running times of algorithms for information spreading on any graph that has large weak conductance, from a polynomial to a polylogarithmic number of rounds.

**Key words.** distributed computing, randomized algorithms, weak conductance, information spreading

**AMS subject classifications.** 68Q85, 68W15, 68W05, 68R10

**DOI.** 10.1137/110845380

**1. Introduction.** Collecting data of all nodes in a network is required by many distributed applications which perform global tasks. The goal of an *information spreading* algorithm is to distribute the messages sent by each of  $n$  nodes in a network to all other nodes. We consider the synchronous push/pull model of communication along with the *transmitter gossip constraint* [31], where each node contacts in each round *one* neighbor to exchange information with, though a node can be contacted by multiple neighbors.

Intuitively, the time required for achieving information spreading depends on the structure of the communication graph or, more precisely, on how well-connected it is. The notion of *conductance*, defined by Sinclair [37], gives a measure for the connectivity of a graph. Roughly speaking, the conductance of a graph  $G$ , denoted by  $\Phi(G)$ , is a value in  $[0, 1]$ : This value is large for graphs that are well-connected (e.g., cliques) and small for graphs that are not (i.e., graphs which have communication bottlenecks). It has been shown that the time required for information spreading can be bounded from above based on the conductance of the underlying communica-

---

\*Received by the editors August 23, 2011; accepted for publication (in revised form) September 17, 2012; published electronically November 15, 2012. A preliminary version of this paper appeared in *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2011, pp. 440–448.

<http://www.siam.org/journals/sicomp/41-6/84538.html>

<sup>†</sup>CSAIL, MIT, Cambridge, MA 02139 (ckeren@csail.mit.edu). Supported by the Simons Postdoctoral Fellows Program and NSF award 0939370-CCF. Part of this work was done while the author was a Ph.D. student at the Department of Computer Science, Technion, Haifa 32000, Israel and supported in part by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities and by the Israel Science Foundation (grant 953/06).

<sup>‡</sup>Department of Computer Science, Technion, Haifa 32000, Israel (hadas@cs.technion.ac.il).

tion graph [4, 11, 12, 22, 31]. In particular, Giakkoupis [22] shows that information spreading can be achieved in  $O(\frac{\log n}{\Phi(G)})$  rounds with probability at least  $1 - \delta$ , where  $\delta = \frac{1}{n^d}$  for some constant  $d$ . This holds when each node randomly chooses a neighbor to contact in every round.

Some graphs have small conductance, implying that they are not well-connected, and therefore the above approach may require many rounds of communication for information spreading. This led us to define in [9] *weak conductance*,  $\Phi_c(G)$ , of a graph  $G$ , which measures connectivity among *subsets* of nodes, whose sizes depend on the parameter  $c \geq 1$ . The paper shows that a relaxed requirement of *partial information spreading*, where each node needs to receive only some of the messages, can be solved fast with high probability in graphs with large weak conductance, although they may have small conductance. As shown in [9], partial information spreading is a sufficient guarantee for some applications.<sup>1</sup>

In this paper we return to the question of achieving *full information spreading*, where each node must receive every message. We present an algorithm that obtains full information spreading on connected graphs and runs fast with high probability on graphs with large weak conductance, independent of their conductance.<sup>2</sup> This expands the known family of graphs for which fast information spreading can be guaranteed, since the weak conductance of a graph is always lower bounded by its conductance and is significantly larger for many graphs.

More generally, for graphs with large weak conductance, our algorithm induces fast solutions for tasks which can be solved using full information spreading, such as leader election, achieving consensus, and computation of aggregation functions.

It has been long known that the conductance itself is insufficient as a lower bound for information spreading. For example, Feige et al. [19] show that information spreading on the hypercube can be obtained in  $O(\log n)$  rounds, despite its conductance being  $O(1/\log n)$ . Our results refine this observation, showing that weak conductance is a more accurate measure for full information spreading.

**1.1. Our contribution.** The main contribution of this paper is an algorithm which achieves fast information spreading with high probability for graphs with large weak conductance. Formally, for any  $c > 1$  and  $\delta = \frac{1}{n^d}$  for some constant  $d$ , our algorithm achieves full information spreading in  $O(c(\frac{\log n}{\Phi_c(G)} + c))$  rounds with probability at least  $1 - 3c\delta$ . This yields substantial improvements in the best known running times of algorithms for information spreading, in particular on graphs that have small conductance but large weak conductance from polynomial to polylogarithmic number of rounds.<sup>3</sup>

Since the best known running times of algorithms for full information spreading inversely depend on the conductance, which may be small due to communication bottlenecks, a natural direction toward speeding up information spreading is to identify such bottlenecks and choose these links with higher probability compared to other neighboring links. However, detecting bottlenecks may not be easy. One approach for separating bottlenecks from other neighbors is to show that a node receives messages from nodes across a bottleneck only with small probability. This seems to

<sup>1</sup>The results in [9] use an alternative definition of conductance, proposed by [31]. Here we return to the original definition of Sinclair [37] and show that the results for partial information spreading still hold.

<sup>2</sup>We consider an algorithm to be *fast* if it runs in a *scalable* number of rounds, i.e., in  $O(\log n)$  or  $O(\text{polylog}(n))$  rounds.

<sup>3</sup>Consider, for example, the class of graphs with conductance  $O(1/n)$  but constant weak conductance. We give specific examples in section 2.

reduce to finding lower bounds for information spreading, a direction which has not proved fruitful so far. Instead, we develop an algorithm that does not detect bottlenecks or formally define their properties, which also appears to be a challenging task. Nonetheless, our algorithm successfully *cope*s with bottlenecks and guarantees fast information spreading despite their presence throughout the network. Roughly speaking, our approach ensures that each node exchanges information quickly with all of its neighbors, perhaps indirectly, and despite the possibility of large degrees, as long as the weak conductance of the graph is large.

We propose a hybrid approach for choosing the neighbor to contact in a given round, which interleaves random choices and deterministic ones. As in the case of random choices, selecting neighbors only in a deterministic manner may require a number of rounds (at least) proportional to the degree of the node, which may be large.<sup>4</sup> Our approach combines random and deterministic techniques using a framework where each node carefully maintains a diminishing list of its neighbors to contact deterministically and alternates between selections from this list and random choices from the set of all neighbors. The lists maintained by the nodes assure that the information spreads across bottlenecks. A main challenge overcome by our algorithm is the tradeoff imposed by managing the lists, namely, inducing a connected subgraph while having scalable list sizes that allow contacting each of the neighbors in them within a small number of rounds.

This constitutes our second contribution: obtaining a connected scalable-degree subgraph in a distributed network of unbounded degree. We believe that finding such subgraphs can be useful in other applications, e.g., in obtaining scalable-degree *spanners* [32, 33]—fundamental subgraphs that preserve distances between nodes up to some stretch. We elaborate about this in the discussion with the technical details of our algorithm in hand.

**1.2. Related work.** Information spreading algorithms have been extensively studied, starting with the work of Demers et al. [14] for replicated database maintenance. Additional research uses information spreading for computation of global functions [28, 31].

Communication models vary in different studies. For example, Karp et al. [26] consider the *random phone-call* model, where in each round every node chooses a random node to communicate with, assuming the communication graph is complete. Much attention has been given to this model, such as bounding the number of calls [15], bounding the number of bits sent [21], or bounding the amount of randomization used [23]. Chen and Pandurangan [10], Jelasity, Montresor, and Babaoglu [25], and Kashyap et al. [27] address the problem of computing aggregate functions in this model, mainly by constructing efficient overlay structures in the graph. Berenbrink et al. [5] give a lower bound trade-off between message size and time complexity for gossip, and Alistarh et al. [1] study spreading of information while tolerating failures in the network. We emphasize that the above papers assume that every node can contact every other node in the network.

As opposed to the random phone-call model, our results hold for arbitrary communication graphs. For such graphs, much work has been done to analyze the time it takes for information spreading using the uniform randomized algorithm (each neighbor has the same probability of being chosen). The most recent papers are by Mosk-Aoyama and Shah [31] and Chierichetti, Lattanzi, and Panconesi [11], and there is

---

<sup>4</sup>Indeed, this is due to the fact that a neighbor adjacent to a bottleneck link may be contacted last.

the recent result of Giakkoupis [22], who showed the best bound of  $O(\frac{\log n}{\Phi(G)})$  for information spreading using this algorithm.

Additional work includes distance-based bounds that were given for nodes placed with uniform density in  $R^d$  [29, 30]. These papers also address gossip-based solutions to specific problems such as resource location and minimum spanning tree. Bradonjić et al. [8] analyze information spreading in random geometric graphs. Georgiou et al. [24] study information spreading in asynchronous networks. Chierichetti, Lattanzi, and Panconesi [12] and Doerr, Fouz, and Friedrich [16] consider information spreading in social networks. Sauerwald and Stauffer [36] study information spreading on regular graphs with respect to vertex expansion, a notion similar to conductance which addresses vertices rather than edges. Sarwate and Dimakis [35] and Boyd et al. [7] study the problem in wireless sensor networks, while Pettarin et al. [34] consider sparse mobile networks.

The *quasi-random* model for information spreading has been introduced by Doerr, Friedrich, and Sauerwald [17] and studied subsequently in [18, 20], as an approach to reduce the amount of randomness. In this model each node has a (cyclic) list of its neighbors, in which it chooses randomly a starting position. It then accesses its list sequentially using this starting position. The paper [17] shows that this model behaves essentially the same as the standard rumor-spreading model.

Recently, our algorithm was used in [3] for obtaining a spanning tree, by each node marking its parent as the neighbor from which it obtained the message of the node with the smallest ID. The spanning tree was then used for information spreading through *algebraic gossip* [13], which is similar to random linear network coding and is used when message sizes are bounded. In fact, our algorithm can be used to obtain a spanning subgraph with additional properties, as described in section 5.

**1.3. Organization.** The rest of the paper is organized as follows. We give an overview of the notions of conductance, weak conductance, and partial information spreading in section 2. In section 3, we present our algorithm for obtaining full information spreading. The analysis of the number of rounds required by our algorithm is given in section 4. Finally, a summary and discussion of the results appear in section 5.

**2. Preliminaries.** The notion of graph conductance was introduced as a measure of how well-connected a graph is. We follow the definition of Sinclair [37] but extend it, so we can measure the conductance of *subsets of nodes* in a graph. This is different than the conductance of the induced subset, as will be explained in the proof of Theorem 1.

For  $S, T \subseteq V$ , the *cut* between  $S$  and  $T$  is defined as  $cut(S, T) = \{\{u, v\} \in E \mid u \in S, v \in T\}$ . For  $S \subseteq V$ , the *volume* of  $S$  is defined as  $vol(S) = \sum_{v \in S} d_v$ , where  $d_v$  is the degree of  $v$  in  $G$ . Notice that although the cut considers only edges between nodes in  $S$  and in  $T$ , the volume is defined with respect to the entire graph  $G$  (namely, it does not correspond to the induced subgraph).

The conductance of a cut is defined as

$$(1) \quad \varphi(S, T) = \frac{|cut(S, T)|}{\min\{vol(S), vol(T)\}},$$

and the conductance of the graph  $G$  is then defined to be the minimal such value, taken over all cuts:

$$\Phi(G) = \min_{S \subseteq V, |S| \leq vol(V)/2} \varphi(S, V \setminus S).$$

Notice that the conductance satisfies  $0 \leq \Phi(G) \leq 1$ , since the number of edges from  $S$  to  $T$  is at most the total number of edges leaving nodes in  $S$ .

Full information spreading is the condition that each node receives the information of all other nodes. For some applications, partial information spreading, where this condition is relaxed, suffices. Formally, for some values  $c \geq 1$  and  $\delta \in (0, 1)$ , we require that with probability at least  $1 - \delta$  every message reaches at least  $n/c$  nodes, and every node receives at least  $n/c$  messages. An algorithm that satisfies this requirement is called  $(\delta, c)$ -spreading. Indeed, the special case where  $c = 1$  corresponds to full information spreading.

Since only a relaxed spreading guarantee is required, the concept of weak conductance can be used in order to analyze partial information spreading. While conductance provides a measure for the connectivity of the whole graph, weak conductance measures the *best* connectivity among subsets that include each node. Formally, for an integer  $c \geq 1$ , the weak conductance of a graph  $G = (V, E)$  is defined as

$$\Phi_c(G) = \min_{i \in V} \left\{ \max_{\substack{V_i \subseteq V, \\ i \in V_i, \\ |V_i| \geq \frac{n}{c}}} \left\{ \min_{\substack{S \subseteq V_i, \\ |S| \leq \frac{\text{vol}(V_i)}{2}}} \varphi(S, V_i \setminus S) \right\} \right\},$$

where  $\varphi(S, T)$  is defined in (1). Indeed, in the special case where  $c = 1$ , the weak conductance of  $G$  is equal to its conductance, namely,  $\Phi_1(G) = \Phi(G)$ . Moreover, this definition implies that the weak conductance of a graph is a monotonically increasing function of  $c$ ; therefore, the weak conductance of a graph  $G$  is at least as large as its conductance.

The following theorem bounds the number of rounds required for  $(\delta, c)$ -spreading.

**THEOREM 1** (see [9, Theorem 3]). *For  $\delta = \frac{1}{n^d}$  for some constant  $d$ , the number of rounds required for  $(\delta, c)$ -spreading is  $O\left(\frac{\log n}{\Phi_c(G)}\right)$ .*

Instead of the proof given in [9], we give a proof that corresponds to the definition of conductance we use here.

*Proof.* For a subset of nodes of the graph, the conductance measure we need is not the conductance of the induced subgraph, since our analysis of spreading within a certain subset still has to account for the probability of choosing edges that leave that subset (and are therefore not part of the induced subgraph).

Instead, when analyzing spreading within a subset, the conductance of the subset can be calculated by considering any edge that leaves the subset as a self-loop,<sup>5</sup> thereby paying for contacting a neighbor over that edge without obtaining information from within the subset.

Therefore, we can use the result of [22] for information spreading within  $V_i$  in time  $O\left(\frac{\log n}{\Phi(G(V_i))}\right)$ , where  $G(V_i)$  is the graph on nodes in  $V_i$  with edges in  $V_i$  taken as they are and edges leaving  $V_i$  taken as self-loops. Together with the definition of weak conductance the theorem follows. We note that the result of [22] holds also for graphs with self-loops (although this is not mentioned explicitly).  $\square$

We note that the proof of Theorem 1 actually gives a much stronger result. For any vertex  $v \in V$ , let  $V_v$  be a component realizing the definition of the weak conductance. Then the above proof shows that this is also a bound on the number of rounds required for every node  $v$  to obtain the message  $m(u)$  of every  $u \in V_v$  and for every  $u$  to obtain  $m(v)$ .

<sup>5</sup>Recall that the volume of a set is measured with respect to the entire graph.




FIG. 1. The  $c$ -barbell graph is a path of  $c$  equal-sized cliques. It is an example of a graph with small conductance and large weak conductance.

**THEOREM 2.** For  $\delta = \frac{1}{n^d}$  for some constant  $d$  with probability at least  $1 - \delta$ , the number of rounds required for every node  $v$  to obtain the message  $m(u)$  of every  $u \in V_v$  and for every  $u \in V_v$  to obtain  $m(v)$  is  $O(\frac{\log n}{\Phi_c(G)})$ .

Therefore, for all nodes  $v \in V$ , the sets  $V_v$  will be central to our analysis. These sets depend on the value of  $c$ , which is omitted to avoid excessive notation.

We proceed by giving examples of the conductance and weak conductance of different graphs. A *clique* has constant conductance. Its weak conductance is equal to its conductance, since for every node  $i$  the best subset  $V_i$  is  $V$  itself. The conductance of a *path* is  $\frac{1}{n}$ , while its weak conductance improves only to  $\frac{c}{n}$ . For these two examples, the weak conductance is in the same order as the conductance for some constant  $c \geq 1$ .

The  $c$ -barbell graph is an example of a graph with very small conductance (for which, to the best of our knowledge, it was not known up to this work how to achieve fast information spreading) but large weak conductance. The  $c$ -barbell graph, which is a generalization of the *barbell* graph, consists of a path of  $c$  cliques, where each contains  $n/c$  nodes (see Figure 1). While the conductance of the  $c$ -barbell graph is  $\Theta(\frac{c}{n^2})$ , the weak conductance is a constant. For any constant  $c \geq 1$ , this implies conductance of  $\Theta(1/n^2)$  while the weak conductance is  $\Theta(1)$ . Indeed, the barbell graph has been studied before (see, e.g., [2, 6]) as a graph for which information spreading requires a large number of rounds. In [2] the context is random walks, which are closely related, since the path of a message can be viewed as a random walk on the graph.

There are additional families of graphs that have a similar property of small conductance and large weak conductance. Examples include rings of cliques and other structures with  $c$  equal-sized well-connected components that are connected by only a few edges.

**3. A fast information spreading algorithm.** Our algorithm for full information spreading applies several phases of partial information spreading, interleaved with our deterministic spreading on a scalable subgraph (whose number of edges can be significantly smaller). We emphasize that the only initial information a node has is  $n$ , the size of the network, and the set of its neighbors. No value of  $c$  is given to the nodes nor do they aim to obtain partial information spreading;  $c$  is only used in the analysis. Moreover, spreading information requires no extra headers, only the messages  $m(v)$  of different nodes.

We consider a synchronous system with  $n$  nodes  $V = \{v_1, v_2, \dots, v_n\}$ , represented by a graph  $G = (V, E)$ . In each round  $r$ , every node contacts one of its neighbors, as explained below, and exchanges information with it. For the analysis, we consider each round  $r$  as a sequence of  $n$  events of information exchange, ordered by the ID of the node that initiated the exchange. (If both nodes choose each other, we consider the node with the smaller ID as the initiator.) We number these events by an *event time*  $t$ , such that  $rn \leq t < (r + 1)n$ . However, these exchanges occur in parallel, which means that a node sends information it had by the end of the last round without additional information it may have received in the current round.

For each node  $v$ , let  $N(v)$  be the set of all neighbors of  $v$ . At every event time  $t$ , node  $v$  maintains a cyclic list  $B_t(v)$  of suspected bottlenecks among its neighbors, where  $B_0(v)$  is initialized to be  $N(v)$ , in an arbitrary order. To exchange information with its neighbors, each node  $v$  alternates between choosing a random neighbor from  $N(v)$  and choosing the next neighbor from  $B_t(v)$ . During this procedure,  $v$  removes neighbors from  $B_t(v)$  according to the following policy.

*Neighbor removal policy for node  $v$ .* Let  $u$  be a node in  $N(v)$ . At event time  $t$  in which  $v$  exchanges information with another node,  $v$  removes any node  $u$  whose message is received for the first time, unless it is received from  $u$  itself and  $v$  is the initiator of this information exchange.

We emphasize that a node  $v$  which removes a node  $u$  from  $B_t(v)$  can still contact node  $u$  if it happens to be its random choice in an even-numbered round  $r$ .

Each node  $v$  also maintains a buffer  $M_t(v)$  of received messages, initialized to consist only of its own message  $m(v)$ . When  $v$  has all  $n$  messages it returns the buffer  $M_t(v)$  as its output.

At every event time  $t$ , for every node  $v$  we define a partition of  $N(v)$  into three sets, as follows:

- $White_t(v) = B_0(v) \setminus B_t(v)$ : The set of nodes that have been removed from  $B_0(v)$ ,
- $Black_t(v) = \{u \in N(v) \mid m(u) \in M_t(v) \text{ and } u \notin White_t(v)\}$ : The set of nodes at event time  $t$  guaranteed to never be removed from  $B_0(v)$ ,
- $Gray_t(v) = N(v) \setminus (Black_t(v) \cup White_t(v))$ : The rest of the nodes, which may or may not be removed in later event times.

We illustrate this partition on the directed graph  $G_t = (V, E)$ , which is the same as the communication graph but with colors associated with each edge at event time  $t$ : if  $u \in White_t(v)$ , then  $(v, u)$  is colored white; if  $u \in Black_t(v)$ , then  $(v, u)$  is colored black; otherwise  $(v, u)$  is colored gray.

The pseudocode for a node  $v$  appears in Algorithm 1.

---

ALGORITHM 1. Full information spreading code for node  $v$ .

---

Initially  $M_0(v) = \{m(v)\}$ ,  $B_0(v) = N(v)$ ,  $r = 0$ ,  $t = 0$

```

1: Repeat:
2:   if  $r$  is even
3:      $w$  = a random neighbor from  $N(v)$ 
4:   else
5:      $w$  = the next neighbor from  $Gray_t(v)$  if it is not empty, and the next from
        $B_t(v)$  otherwise
6:   Contact  $w$  and exchange information
7:   Exchange information with every  $w'$  that contacts  $v$ 
8:   for  $i = 1$  to  $n$ 
9:     if  $v = v_i$  or  $v$  is contacted by  $w' = v_i$ 
10:      Add new messages to  $M_t(v)$ 
11:      for every node  $u \in N(v)$ 
12:        if  $v$  receives  $m(u)$  for the first time from  $w$  and  $w \neq u$ 
           or  $v$  receives  $m(u)$  for the first time from  $w'$ 
13:           $B_{t+1}(v) = B_t(v) \setminus \{u\}$ 
14:       $t = t + 1$ 
15:       $r = r + 1$ 

```

---

We start the analysis with four simple claims regarding the colors of the edges of  $G_t$ . First, we claim that initially all edges are gray. Second, if a node  $v$  has the message of a neighbor  $u$ , then the edge  $(v, u)$  cannot be gray. The third claim restricts the possibility of having symmetric edges change their color at the same event time  $t$ : if this occurs, then these two nodes are the ones that exchange information in this event time. Finally, it is not hard to see that two symmetric edges cannot turn white at the same event time  $t$ , which is our fourth claim.

LEMMA 3. *The following four claims hold:*

- (i) *In  $G_0$  all edges are colored gray.*
- (ii) *For any nodes  $v$  and  $u$ , if  $m(u) \in M_t(v)$  and  $u \in N(v)$ , then  $(v, u)$  is not gray in  $G_t$ .*
- (iii) *For any event time  $t$ , if for some nodes  $v$  and  $u$  both edges  $(v, u)$  and  $(u, v)$  change their color at  $t$ , then  $v$  and  $u$  are the pair of nodes that exchange information in this event time.*
- (iv) *If  $(v, u)$  and  $(u, v)$  are both gray, they cannot both turn white at the same time step.*

*Proof.*

- (i) By definition, at event time 0 no edges are colored white, since no node  $u$  has been removed from any set  $B_0(v)$ . Moreover, no buffer  $M_0(v)$  contains any message other than  $m(v)$ , therefore no edge is colored black. This implies that all edges in  $G_0$  are gray.
- (ii) If  $v$  has the message  $m(u)$  by event time  $t$  and  $u \notin \text{White}_t(v)$ , then by definition,  $u \in \text{Black}_t(v)$ , that is,  $(v, u)$  is black in  $G_t$ .
- (iii) By the code of Algorithm 1, an edge  $(v, u)$  changes its color at event time  $t$  only if  $v$  received a message at that time. If  $u$  is not one of the two nodes that exchange information at event time  $t$ , then  $(u, v)$  cannot change its color at time  $t$ .
- (iv) By (iii), two neighbors can change the color of the edges connecting them at the same time step only if they exchange information at this time step. Assume without loss of generality that  $v$  is the initiator of this exchange. Then  $v$  does not remove  $u$  from  $B_t(v)$  at this step because neither of the conditions in line 12 is satisfied.  $\square$

We are now ready to prove a key lemma in our analysis, which shows that whenever a node  $v$  removes a neighbor  $u$  from  $B_t(v)$ , that is, the edge  $(v, u)$  is colored white, there is an undirected path of edges between  $v$  and  $u$  that are guaranteed to never be removed, i.e., a black path.

LEMMA 4. *For any event time  $t \geq 0$ , if  $(v, u)$  turns white, then there is a path  $v = a_0, a_1, \dots, a_{\ell-1}, a_\ell = u$  such that for all  $0 \leq i \leq \ell - 1$ , either  $(a_i, a_{i+1})$  is black or  $(a_{i+1}, a_i)$  is black.*

*Proof.* The proof is by induction on the time step  $t$ , where the base case for  $t = 0$  is the initial coloring of the graph. By Lemma 3(i), at this time all the edges are colored gray, therefore the lemma holds vacuously. For the induction step, assume that the lemma is true for every edge that turns white in a step  $t' < t$ . We prove the lemma for any edge  $(v, u)$  that turns white at time  $t$ . If this happens, then one of the conditions in line 12 of the algorithm is satisfied. We distinguish between two cases.

*Case 1.* Node  $v$  receives  $m(u)$  for the first time from node  $w \neq u$  and removes node  $u$  from  $B_t(v)$ . The fact that  $w$  has the message  $m(u)$  at this time implies that  $m(u)$  traveled a path  $u = a_0, a_1, \dots, a_{\ell'} = w$  by time  $t - 1$ . For all  $0 \leq i < \ell'$ , there has been information exchange between  $a_i$  and  $a_{i+1}$ , which implies that both nodes have






FIG. 2. Proof of Lemma 4. Dashed arrows represent white edges, while solid arrows represent black edges. If  $v$  removes  $u$  from  $B_t(v)$  by receiving  $m(u)$  from  $w$ , then there is an undirected black path between  $u$  and  $w$ .

each other’s message. Lemma 3(ii) implies that both edges  $(a_i, a_{i+1})$  and  $(a_i, a_{i+1})$  are not gray.

Assume that for some  $i$ , both edges are white. (If no such  $i$  exists, then we are done.) By Lemma 3(iv), these two edges could not have turned white at the same time step. Let  $t'$  be the time step in which the second edge turned white. By the induction hypothesis, at time  $t'$  there was an undirected path  $a_i = b_{i,0}, b_{i,1}, \dots, b_{i,\ell_i} = a_{i+1}$  colored black. Since this is true for all  $i = 0, \dots, \ell' - 1$ , we have that there is an undirected black path between  $u$  and  $w$  (see Figure 2).

It remains to be shown that there is a black path between  $v$  and  $w$ . Since  $v$  and  $w$  exchange information, by Lemma 3(ii), both edges  $(v, w)$  and  $(w, v)$  are colored. Assume that they are both colored white. (Otherwise we are done.) Then  $(v, w)$  does not turn white in event time  $t$  because  $v$  is the initiator. This implies that  $(v, w)$  turns white at some time  $t' < t$ . By the induction hypothesis, there is an undirected black path between  $v$  and  $w$  at that time, which completes the proof of this case.

Case 2. Node  $v$  receives  $m(u)$  for the first time from node  $w'$  which was the initiator of the information exchange and removes node  $u$  from  $B_t(v)$ . The proof follows the lines of Case 1 to show that there is an undirected black path between  $u$  and  $w'$ . It then remains to be shown that there is a black path between  $v$  and  $w'$ . A similar argument to that of Case 1, replacing the initiator  $v$  with  $w'$ , proves that there is such a path.  $\square$

Lemma 4 guarantees that after removing elements from the sets  $B_t(v)$  of different nodes, the nodes always remain connected by black and gray edges, even though white edges in the communication graph are ignored in line 5 of the algorithm. Again, recall that a node  $v$  can still contact a neighbor  $u$  it removed from  $B_t(v)$  by choosing it in line 3 of the algorithm. Finally, we note that for every round  $r$  and node  $v$ , only one edge can be colored black by  $v$  through all  $n$  event times of this round, since a node  $u$  joins  $Black_t(v)$  only if  $u$  is the unique node  $w$  with whom  $v$  initiated information exchange at event time  $t$ .

CLAIM 5. For every round  $r$  and every node  $v$  we have that  $|Black_{n(r+1)}(v)| - |Black_{nr}(v)| \leq 1$ .

**4. Analysis.** We now claim that there are not *too many* black edges outgoing from any node  $v$ . This guarantees that every neighbor remaining in  $B_t(v)$  will be eventually contacted after a *small* number of rounds. The precise measures of these amounts will be defined later.

For the rest of this section, we fix a value  $c > 1$  and a value  $\delta = \frac{1}{n^c}$  for some  $d$ , as guaranteed by [22]. However, we emphasize that these values are not used in the algorithm and are therefore unknown to the nodes. They are only used for the analysis,<sup>6</sup> and eventually we will choose a value  $c$  that minimizes the number of rounds.

<sup>6</sup>This is formalized in Theorem 12.

Let  $T = O(\frac{\log n}{\Phi_c(G)})$ , the number of rounds obtained in Theorem 2. We consider *phases* of the algorithm, each phase consisting of  $2T$  rounds. The outline of our analysis is as follows. Recall that for any  $v \in V$ ,  $V_v$  is the component realizing the definition of the weak conductance. Roughly speaking, Theorem 2 shows that with high probability after one phase a node  $v$  has the messages of all nodes  $u$  in its component  $V_v$ , since the even-numbered rounds comprise of *regular* information spreading. We then show that after three phases, a node  $v$  has the messages of all nodes that are either in its component or in an intersecting component. Finally, we show that after  $c(6T + 2c)$  rounds, a node  $v$  has the messages of all nodes. This strongly relies on the connectivity argument in Lemma 4 and a careful bookkeeping of the number of edges in  $B_t(v)$  throughout these phases. In addition, we need to keep track of the probability of failure in every phase.

We begin by using Theorem 2 to show that starting from *any* round  $r_0$ , after one phase of the algorithm we have spread the information of  $M_{nr_0}(v)$  inside the component  $V_v$  (instead of just  $m(v)$  if  $r_0 = 0$ ). We emphasize that the probability of success is for *all* nodes  $v$  to satisfy the requirements.

LEMMA 6. *Let  $r_0$  be a round number. After round  $r = r_0 + 2T$  with probability at least  $1 - \delta$ , for every node  $v$  we have  $\bigcup_{u \in V_v} M_{nr_0}(u) \subseteq M_{nr}(v)$  and  $M_{nr_0}(v) \subseteq M_{nr}(u)$  for every  $u \in V_v$ .*

*Proof.* Consider the state of the buffers  $M_{nr}(v)$  after  $r = r_0 + 2T$  rounds. By Theorem 2 and since  $2T$  rounds contain  $T$  even-numbered rounds, with probability  $1 - \delta$  every node  $v$  has all messages  $M_{nr_0}(u)$  of every  $u \in V_v$  and every  $u \in V_v$  has  $M_{nr_0}(v)$ .  $\square$

We consider the progress of the algorithm after three phases. For every node  $v$  we define the set  $I_v = \{u \in V \mid V_u \cap V_v \neq \emptyset\}$  of nodes whose component intersects the component of  $v$ . Further, for every node  $v$  let  $A_v = \bigcup_{u \in I_v} V_u$ . Notice that  $V_v \subseteq A_v$ .

LEMMA 7. *Let  $r_0$  be a round number. After round  $r = r_0 + 6T$  with probability at least  $1 - 3\delta$ , for every node  $v$  we have  $\bigcup_{x \in A_v} M_{nr_0}(x) \subseteq M_{nr}(v)$ .*

*Proof.* Consider the algorithm after  $6T$  rounds. Let  $v$  be a node and  $u$  a node not in  $V_v$ . If  $V_v \cap V_u \neq \emptyset$ , then there is a node  $w \in V_v \cap V_u$ . Lemma 6 implies that after  $4T$  rounds  $w$  has all messages  $M_{nr_0}(x)$  of nodes  $x$  in  $V_u$  (one phase for  $u$  to receive  $M_{nr_0}(x)$  of every  $x \in V_u$  and another phase for this information to reach  $w \in V_u$ ). Applying Lemma 6 again gives that  $w$  spreads these messages to  $v$  in  $2T$  additional rounds. That is,  $v$  has the messages  $M_{nr_0}(x)$  of all nodes  $x$  in  $\bigcup_{u \in I_v} V_u$ . All three phases of information spreading need to succeed for the above to happen. A simple union bound on the probability that either fails gives that with probability at least  $1 - 3\delta$  all three phases succeed.  $\square$

After each node  $v$  has all the messages of nodes in  $A_v$ , it takes only  $2c$  rounds for all remaining gray edges to turn white or black. The reason we handled nodes in  $A_v$  separately in Lemma 7 is that nodes that are outside  $A_v$  have components that do not intersect  $V_v$ . This implies that if  $v$  has to spend a round contacting a node outside  $A_v$ , then it receives at least  $n/c$  new messages of nodes in  $V_v$ . Hence,  $v$  does not need to contact these nodes directly. In other words, this allows us to claim that soon there are no more gray edges in the graph (see Figure 3).

LEMMA 8. *After  $r = 6T + 2c$  rounds with probability at least  $1 - 3\delta$ , for every node  $v$  we have that  $\{m(u) \mid u \in N(v)\} \subseteq M_{nr}(v)$ .*

*Proof.* Let  $S_v^i$  be the set of nodes  $u$  such that  $v$  receives the message  $m(u)$  after  $r_i = 6T + 2i$  rounds. We claim that with probability at least  $1 - 3\delta$ , for every node  $v$  and every  $i$ ,  $1 \leq i \leq c$ , after round  $r_i$  the buffer  $M_{nr_i}(v)$  of messages received by  $v$  either contains  $m(u)$  of all nodes in  $N(v)$  or there are  $i$  different nodes




FIG. 3. Illustration of a node  $v$ , its component  $V_v$ , and the set  $A_v$  of intersecting components.

$u_1, \dots, u_i \in N(v)$  such that  $V_{u_j} \cap V_{u_k} = \emptyset$  for all  $1 \leq j < k \leq i$ , and for every  $1 \leq j \leq i$  we have  $A_{u_j} \subseteq S_v^i$ .

We prove this claim by induction. For the base case  $i = 1$ , by Lemma 7 we have that with probability at least  $1 - 3\delta$  after  $6T \leq r_1$  rounds each node  $v$  has  $m(u)$  of all nodes  $u$  in  $A_v$ , therefore we choose  $u_1 = v$ .

We next assume that the claim holds up to  $i - 1$  and prove it for  $i$ . By the induction hypothesis, with probability at least  $1 - 3\delta$  for every node  $v$  we have after round  $r_{i-1} = 6T + 2(i-1)$  that there are  $i-1$  nodes  $u_1, \dots, u_{i-1}$  such that  $V_{u_j} \cap V_{u_k} = \emptyset$  for all  $1 \leq j < k \leq i-1$ , and  $A_{u_j} \subseteq S_v^{i-1}$  for every  $1 \leq j \leq i-1$ .

If  $S_v^{i-1}$  contains all nodes in  $N(v)$ , then we are done. Otherwise, there is a node  $u \in N(v)$  such that the edge  $(v, u)$  is gray at the beginning of round  $r_i$ . In the next odd-numbered round  $v$  contacts such a node  $u$ . Since  $u \notin S_v^{i-1}$ , by the induction hypothesis we have  $u \notin A_{u_j}$  for every  $1 \leq j \leq i-1$ . This implies that  $V_u \cap V_{u_j} = \emptyset$  for every  $1 \leq j \leq i-1$ . Moreover, by Lemma 7  $A_u \subseteq S_v^i$ . This completes the proof of our claim. The claim implies that after  $6T + 2c$  rounds, either  $v$  has  $m(u)$  for every  $u \in N(v)$  or there are  $c$  different nodes  $u_1, \dots, u_c$  such that  $V_{u_j} \cap V_{u_k} = \emptyset$  for all  $1 \leq j < k \leq c$ , and for every  $1 \leq j \leq c$  we have  $A_{u_j} \subseteq S_v^c$ . In particular,  $v$  has the messages of all nodes of the pairwise disjoint sets  $V_{u_j}$  for  $1 \leq j \leq c$ , all of which are of size  $n/c$ , which implies that  $v$  has all messages.  $\square$

Having the messages of all components of the neighbors of a node immediately implies that there are no more gray edges. In addition, by Claim 5 we can bound the number of out-going black edges for each node.

**COROLLARY 9.** *After  $6T + 2c$  rounds, i.e., for  $t = (6T + 2c)n$  with probability at least  $1 - 3\delta$ , for every node  $v$  we have  $B_t(v) = \text{Black}_t(v)$  and  $|B_t(v)| \leq 6T + 2c$ .*

We are now ready to prove our main lemma, which bounds the complexity of the algorithm. Roughly speaking, the argument follows the line of proof of Lemma 8, but instead of considering gray edges, it relies on having connectivity among the black edges.

**LEMMA 10.** *With probability at least  $1 - 3c\delta$  after at most  $r = 2c(6T + 2c)$  rounds, for every node  $v$  we have  $M_{nr}(v) = \{m(u) \mid u \in V\}$ .*

*Proof.* Let  $r_0$  be a round number, and let  $S_v^i$  be the set of nodes  $u$  for which  $M_{nr_0}(u) \subseteq M_{nr_i}(v)$  after  $r_i = r_0 + r'_i$  rounds, where  $r'_i = 2i(6T + 2c)$ . We use in the proof the following.

**CLAIM 11.** *With probability at least  $1 - 3i\delta$ , for every node  $v$  and  $1 \leq i \leq c$ , after round  $r_i$  the buffer  $M_{nr_i}(v)$  of messages received by  $v$  either contains messages of all nodes or there are  $i$  different nodes  $u_1, \dots, u_i$  such that  $V_{u_j} \cap V_{u_k} = \emptyset$  for all  $1 \leq j < k \leq i$ , and for every  $1 \leq j \leq i$  we have  $A_{u_j} \subseteq S_v^i$ .*

*Proof.* We prove the claim by induction. For the base case  $i = 1$ , by Lemma 7 with probability at least  $1 - 3\delta$ , after  $r_0 + 6T \leq r_1$  rounds, for every node  $v$  we have  $\bigcup_{\{x \in A_v\}} M_{nr_0}(x) \subseteq M_{nr_1}(v)$ . Therefore,  $A_v \subseteq S_v^1$  so we choose  $u_1 = v$ .

Next, we assume that the claim holds up to  $i - 1$  and prove it for  $i$ . By the induction hypothesis with probability at least  $1 - 3(i - 1)\delta$ , for every node  $v$  we have after round  $r_{i-1} = r_0 + 2(i - 1)(6T + 2c)$  that there are  $i - 1$  nodes  $u_1, \dots, u_{i-1}$  such that  $V_{u_j} \cap V_{u_k} = \emptyset$  for all  $1 \leq j < k \leq i - 1$ , and  $A_{u_j} \subseteq S_v^{i-1}$  for every  $1 \leq j \leq i - 1$ .

If  $S_v^{i-1}$  contains all nodes, then we are done. Otherwise, by Lemma 4 and Corollary 9, by this time there is an undirected black path between any two nodes of the graph, and specifically there is a node  $u \notin S_v^{i-1}$  connected by a black edge to some node  $w \in S_v^{i-1}$ . Since after  $6T + 2c$  rounds all nodes  $v$  have at most  $6T + 2c$  nodes in  $B_t(v)$ , after at most  $6T + 2c$  additional rounds each node contacts each of its black neighbors. Therefore, after round  $r' = r_0 + 2(6T + 2c)$  the node  $w$  has the messages of all nodes in  $S_u^1$ , that is,  $\bigcup_{x \in S_u^1} M_{nr_0}(x) \subseteq M_{nr'}(w)$ .

By the induction hypothesis with  $r'$  replacing  $r_0$ , after  $r'_{i-1}$  additional rounds  $v$  has them as well with another factor of  $3\delta$  added to the probability of failure. Formally,  $\bigcup_{x \in S_u^1} M_{nr_0}(x) \subseteq M_{nr'}(w) \subseteq M_{nr''}(v)$ , where  $r'' = r' + r'_{i-1} = r_0 + 2(6T + 2c) + 2(i - 1)(6T + 2c) = r_i$ .

We now prove that taking  $u_i = u$  satisfies the requirements of our claim. Since  $u \notin S_v^{i-1}$ , by the induction hypothesis we have  $u \notin A_{u_j}$  for every  $1 \leq j \leq i - 1$ . This implies that  $V_u \cap V_{u_j} = \emptyset$  for every  $1 \leq j \leq i - 1$ . Moreover,  $A_u \subseteq S_v^i$  since  $A_u \subseteq S_u^1$ . This completes the proof.  $\square$

By Claim 11, after  $2c(6T + 2c)$  rounds either  $v$  has  $m(u)$  for every  $u \in N(v)$  or there are  $c$  different nodes  $u_1, \dots, u_c$  such that  $V_{u_j} \cap V_{u_k} = \emptyset$  for all  $1 \leq j < k \leq c$ , and for every  $1 \leq j \leq c$  we have  $A_{u_j} \subseteq S_v^c$ . In particular,  $v$  has the messages of all nodes of the pairwise disjoint sets  $V_{u_j}$  for  $1 \leq j \leq c$ , all of which are of size  $n/c$ . This implies that  $v$  has all messages.  $\square$

Rephrasing Lemma 10 gives our main theorem for full information spreading.

**THEOREM 12.** *For every  $c > 1$  and  $\delta = \frac{1}{n^d}$  for some constant  $d$ , Algorithm 1 obtains full information spreading after at most  $O(c(\frac{\log n}{\Phi_c(G)} + c))$  rounds with probability at least  $1 - 3c\delta$ .*

Looking at some specific values of the parameters in the above theorem we get that Algorithm 1 is fast for graphs with scalable weak conductance (for a scalable value of  $c$ ). For example, if  $c = \text{polylog}(n)$  and  $\Phi_c = 1/\text{polylog}(n)$ , then our algorithm requires a polylogarithmic number of rounds. The probability of failure is  $3c\delta$ , which is  $O(\text{polylog}(n)/n^d)$  since  $\delta = 1/n^d$ , and it is  $O(1/n^d)$  if  $c$  is a constant; in both cases it is  $o(1)$ .

Since the bound holds for every value of  $c$ , we can also state our result as follows.

**THEOREM 13.** *Algorithm 1 obtains full information spreading after at most  $\min_c \{O(c(\frac{\log n}{\Phi_c(G)} + c))\}$  rounds with probability at least  $1 - 3c'\delta$ , where  $c'$  is the value that realizes the minimum in this expression and  $\delta = \frac{1}{n^d}$  for some constant  $d$ .*

**5. Discussion.** This paper studies information spreading, presenting a hybrid algorithm, which interleaves random neighbor choices with deterministic ones for exchange of information. Our algorithm is fast on graphs which have large weak conductance. For graphs which also have small conductance, it substantially improves upon the running times of previously known algorithms, from a polynomial to a polylogarithmic number of rounds.

Our algorithm was used in [3] for obtaining a spanning tree by each node marking its parent as the neighbor from which it obtained the message of the node with the smallest ID. In fact, a by-product of our algorithm is the maintenance of a connected scalable-degree subgraph, which we believe will find additional applications.

Specifically, it may be possible to obtain scalable-degree spanners with low stretch by applying similar techniques. The connection between our algorithm and spanners is as follows. Consider the subgraph that is induced by the partial lists of neighbors obtained during the algorithm. The number of edges in this subgraph is bounded by the size of the lists  $Black_t(v)$  at the end of the algorithm, therefore if the weak conductance is large, then the subgraph is sparse. In addition, a node can communicate with all of its neighbors using this subgraph. Hence, if the weak conductance is large, the number of steps required to obtain information from all neighbors is small, which corresponds to the stretch of the subgraph.

An intriguing open question is whether there is a nontrivial lower bound on the number of rounds required for information spreading as a function of the weak conductance of the underlying graph. Another avenue for future research is to adapt our algorithm to failure-prone environments, as resilience to faults is typically required in practical scenarios.

Finally, we note that our model allows messages of unbounded size. Bounding the size of messages is another direction for further research.

**Acknowledgments.** The authors thank Hagit Attiya for helpful suggestions and comments on an earlier version of this paper. We also thank Chen Avin and the anonymous reviewers of SODA 2011 for useful comments on a previous version of this paper. Finally, we thank the anonymous referees for invaluable suggestions.

## REFERENCES

- [1] D. ALISTARH, S. GILBERT, R. GUERRAOU, AND M. ZADIMOGHADDAM, *How efficient can gossip be? (On the cost of resilient information exchange)*, in Proceedings of the 37th International Colloquium on Automata, Languages and Programming (ICALP), 2010, pp. 115–126.
- [2] N. ALON, C. AVIN, M. KOUCKY, G. KOZMA, Z. LOTKER, AND M. R. TUTTLE, *Many random walks are faster than one*, in Proceedings of the 20th Annual Symposium on Parallelism in Algorithms and Architectures (SPAA), 2008, pp. 119–128.
- [3] C. AVIN, M. BOROKHOVICH, K. CENSOR-HILLEL, AND Z. LOTKER, *Order optimal information spreading using algebraic gossip*, in Proceedings of the 30th ACM Symposium on Principles of Distributed Computing (PODC), 2011, pp. 363–372.
- [4] C. AVIN AND G. ERCAL, *Bounds on the mixing time and partial cover of ad-hoc and sensor networks*, in Proceedings of the Second European Workshop on Wireless Sensor Networks, 2005, pp. 1–12.
- [5] P. BERENBRINK, J. CZYZOWICZ, R. ELSÄSSER, AND L. GASIENIEC, *Efficient information exchange in the random phone-call model*, in Proceedings of the 37th International Colloquium on Automata, Languages and Programming (ICALP), 2010, pp. 127–138.
- [6] M. BOROKHOVICH, C. AVIN, AND Z. LOTKER, *Tight bounds for algebraic gossip on graphs*, in Proceedings of the 2010 IEEE International Symposium on Information Theory (ISIT), 2010, pp. 1758–1762.
- [7] S. BOYD, A. GHOSH, B. PRABHAKAR, AND D. SHAH, *Randomized gossip algorithms*, IEEE/ACM Trans. Netw., 14 (2006), pp. 2508–2530.
- [8] M. BRADONJIC, R. ELSÄSSER, T. FRIEDRICH, T. SAUERWALD, AND A. STAUFFER, *Efficient broadcast on random geometric graphs*, in Proceedings of the 21st ACM-SIAM Symposium on Discrete Algorithms (SODA), 2010, pp. 1412–1421.
- [9] K. CENSOR HILLEL AND H. SHACHNAI, *Partial information spreading with application to distributed maximum coverage*, in Proceedings of the 27th ACM Symposium on Principles of Distributed Computing (PODC 2010), 2010, pp. 161–170.
- [10] J. Y. CHEN AND G. PANDURANGAN, *Optimal gossip-based aggregate computation*, in Proceedings of the 22nd Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), 2010, pp. 124–133.
- [11] F. CHERICHETTI, S. LATTANZI, AND A. PANCONESI, *Almost tight bounds for rumour spreading with conductance*, in Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC), 2010, pp. 399–408.

- [12] F. CHERICHETTI, S. LATTANZI, AND A. PANCONESI, *Rumour spreading and graph conductance*, in Proceedings of the 21st ACM-SIAM Symposium on Discrete Algorithms (SODA), 2010, pp. 1657–1663.
- [13] S. DEB, M. MÉDARD, AND C. CHOUTÉ, *Algebraic gossip: A network coding approach to optimal multiple rumor mongering*, IEEE/ACM Trans. Netw., 14 (2006), pp. 2486–2507.
- [14] 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 6th Annual ACM Symposium on Principles of Distributed Computing (PODC), 1987, pp. 1–12.
- [15] B. DOERR AND M. FOUZ, *Asymptotically optimal randomized rumor spreading*, in Proceedings of the 38th International Colloquium on Automata, Languages and Programming (ICALP), 2011, pp. 502–513.
- [16] B. DOERR, M. FOUZ, AND T. FRIEDRICH, *Social networks spread rumors in sublogarithmic time*, in Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC), 2011, pp. 21–30.
- [17] B. DOERR, T. FRIEDRICH, AND T. SAUERWALD, *Quasirandom rumor spreading*, in Proceedings of the 19th annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2008, pp. 773–781.
- [18] B. DOERR, T. FRIEDRICH, AND T. SAUERWALD, *Quasirandom rumor spreading: Expanders, push vs. pull, and robustness*, in Proceedings of the 36th International Colloquium on Automata, Languages and Programming (ICALP), Part 1, 2009, pp. 366–377.
- [19] U. FEIGE, D. PELEG, P. RAGHAVAN, AND E. UPFAL, *Randomized broadcast in networks*, in Proceedings of the SIGAL International Symposium on Algorithms, 1990, pp. 128–137.
- [20] N. FOUNTOLAKIS AND A. HUBER, *Quasirandom rumor spreading on the complete graph is as fast as randomized rumor spreading*, SIAM J. Discrete Math., 23 (2009), pp. 1964–1991.
- [21] P. FRAIGNAUD AND G. GIAKKOUPIS, *On the bit communication complexity of randomized rumor spreading*, in Proceedings of the 22nd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), ACM, New York, 2010, pp. 134–143.
- [22] G. GIAKKOUPIS, *Tight bounds for rumor spreading in graphs of a given conductance*, in Proceedings of the 28th International Symposium on Theoretical Aspects of Computer Science (STACS), 2011, pp. 57–68.
- [23] G. GIAKKOUPIS AND P. WOELFEL, *On the randomness requirements of rumor spreading*, in Proceedings of the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA), 2011, pp. 449–461.
- [24] C. GEORGIU, S. GILBERT, R. GUERRAOU, AND D. R. KOWALSKI, *On the complexity of asynchronous gossip*, in Proceedings of the 27th ACM Symposium on Principles of Distributed Computing (PODC), 2008, pp. 135–144.
- [25] M. JELASITY, A. MONTRESOR, AND Ö. BABAÖGLU, *Gossip-based aggregation in large dynamic networks*, ACM Trans. Comput. Syst., 23 (2005), pp. 219–252.
- [26] R. KARP, C. SCHINDELHAUER, S. SHENKER, AND B. VOCKING, *Randomized rumor spreading*, in Proceedings of the 41st Annual Symposium on Foundations of Computer Science (FOCS), IEEE Computer Society, Washington, DC, 2000, pp. 565–574.
- [27] S. R. KASHYAP, S. DEB, K. V. M. NAIDU, R. RASTOGI, AND A. SRINIVASAN, *Efficient gossip-based aggregate computation*, in Proceedings of the 25th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS), 2006, pp. 308–317.
- [28] D. KEMPE, A. DOBRA, AND J. GEHRKE, *Gossip-based computation of aggregate information*, in Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science (FOCS), IEEE Computer Society, 2003, Washington, DC, pp. 482–491.
- [29] D. KEMPE, J. KLEINBERG, AND A. DEMERS, *Spatial gossip and resource location protocols*, in Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC), ACM, New York, 2001, pp. 163–172.
- [30] D. KEMPE AND J. M. KLEINBERG, *Protocols and impossibility results for gossip-based communication mechanisms*, in Proceedings of the 43rd Symposium on Foundations of Computer Science (FOCS), IEEE Computer Society, Washington, DC, 2002, pp. 471–480.
- [31] D. MOSK-AOYAMA AND D. SHAH, *Computing separable functions via gossip*, in Proceedings of the 25th annual ACM Symposium on Principles of Distributed Computing (PODC), ACM, New York, 2006, pp. 113–122.
- [32] D. PELEG AND A. A. SCHÄFFER, *Graph spanners*, J. Graph Theory, 13 (1989), pp. 99–116.
- [33] D. PELEG AND J. D. ULLMAN, *An optimal synchronizer for the hypercube*, SIAM J. Comput., 18 (1989), pp. 740–747.
- [34] A. PETTARIN, A. PIETRACAPRINA, G. PUCCI, AND E. UPFAL, *Tight bounds on information dissemination in sparse mobile networks*, in Proceedings of the 30th ACM Symposium on Principles of Distributed Computing (PODC), 2011, pp. 355–362.

- [35] A. D. SARWATE AND A. G. DIMAKIS, *The impact of mobility on gossip algorithms*, in 28th IEEE International Conference on Computer Communications (INFOCOM), 2009, pp. 2088–2096.
- [36] T. SAUERWALD AND A. STAUFFER, *Rumor spreading and vertex expansion on regular graphs*, in Proceedings of the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA), 2011, pp. 462–475.
- [37] A. SINCLAIR, *Algorithms for random generation and counting: A Markov chain approach*, Birkhäuser-Verlag, Basel, Switzerland, 1993.