FIFO Queues are Bad for Rumor Spreading

Article *in* IEEE Transactions on Information Theory · November 2016

DOI: 10.1109/TIT.2016.2632153

Marcos Kiwi
University of Chile

FIFO Queues Are Bad for Rumor Spreading

Marcos Kiwi and Christopher Thraves Caro

***Abstract*— The two most intensively studied communication paradigms for spreading rumors are the so-called PUSH and PULL algorithms. The previous analysis of these protocols assumed that every node could process all such push/pull operations within a single step, which could be unrealistic in practical situations. We propose a new framework for the analysis of rumor spreading accommodating buffers, in which a node can process only few push/pull messages at a time. We develop time complexity upper and lower bounds for randomized rumor spreading in the new framework, and compare the results with analogous ones in the classical setting. Our results highlight that there might be a very significant performance loss if messages are processed at each network node in first-in first-out order.**

***Index Terms*— Randomized rumor spreading, FIFO queues, graph conductance.**

I. INTRODUCTION

RUMOR spreading is a fundamental concept in ad hoc communication, databases and systems: a rumor, initially stored in one node, is delivered to all other nodes in the network by passing it along available links. Some appealing applications of randomized rumor spreading are data aggregation [2], maintenance of replicated databases [8], resource discovery [18] and failure detection [24], among others.

Randomized and distributed spreading algorithms are of special interest due to their simplicity, robustness and locality [13], [19]. A well known such algorithm is PUSH [8] which is executed in each node at discrete time steps. If a node holds the rumor at a given time slot, then it chooses uniformly at random one neighbor with whom to establish a point-to-point communication. If the chosen neighbor does not hold the rumor, then it is *pushed* through and the neighbor becomes a holder of the rumor at the same step. If a node does not hold the rumor, it waits until it becomes a holder. The PUSH algorithm (or just PUSH for short) has a counterpart known as PULL algorithm (henceforth, PULL for short) which is also executed in each node and evolves in steps. Now, however, if a node does not hold the rumor at a given step, then it chooses uniformly at random one neighbor to establish a point-to-point communication. If the chosen neighbor holds the rumor, then the rumor is *pulled* through and the node becomes a holder of the rumor during that same step. A third well studied randomized rumor spreading algorithm is a combination of PUSH and PULL, called PUSH&PULL, in which every node that holds the rumor acts according to PUSH, and every node that does not hold the rumor performs PULL.

The classical framework in which PUSH, PULL and PUSH&PULL have been studied assumes that time evolves in discrete steps. Furthermore, it assumes that at each time step every node can establish many parallel point-to-point communications if many neighbors request one. Hence, it is implicitly assumed that, at every step, each node is able to process every received (pushed) and requested (pulled) message. Therefore, if a node that holds the rumor is contacted by many neighbors at the same time step, the classical framework allows this node to process all messages and answer requests during the same time step. Such a strong assumption fails to capture some important phenomena occurring in practical situations. For instance, if the graph consists of one central node connected to n leaves and the spreading algorithm is PUSH&PULL, then the rumor spreads in at most 2 steps. While even if the central node can send n messages at the same time, assuming processors can service $O(1)$ requests per time step, it may take $\Omega(n)$ steps to process all the incoming messages sent by the n leaves. Therefore, if the source node is one of the leaves, in the worst case, the central node may learn the rumor after $\Omega(n)$ steps (after processing all the n messages sent by the leaves). This fact is not captured by the analyses performed so far concerning randomized rumor spreading processes.

In order to address the aforementioned phenomena, we introduce a new framework to study the time complexity of rumor spreading algorithms, i.e., the number of steps required by the rumor to propagate to all nodes. The model we propose can be understood as an extension of the phone call model [6], [20] where every node has a telephone answering machine (the buffer). If a node receives a call while it is already in a different call, the buffer takes the message. In one time step nodes are allowed to perform one call to one neighbor, and one access to its buffer to retrieve a single messages (if any) and copy it to the buffer's local memory.[1] If the buffer is empty at some step, one incoming call goes directly to the node and the other incoming calls (if any) go to the buffer (the call that goes directly to the node is chosen uniformly at random). In a step where the buffer is not empty, given that the node accesses its buffer to retrieve

Manuscript received February 1, 2016; revised July 25, 2016; accepted November 6, 2016. Date of publication November 23, 2016; date of current version January 18, 2017. M. Kiwi was supported by the Millennium Nucleus Information and Coordination in Networks ICM/FIC RC130003 and CONICYT via Basal in Applied Mathematics. C. Thraves Caro was supported by the Spanish MICINN Grant Juan de la Cierva.

M. Kiwi is with the Depto. de Ingeniería Matemática and Ctr. de Modelamiento Matemático, Fac. de Cs. Físicas y Matemáticas, U. de Chile, Santiago 837-0456, Chile (e-mail: mk@dim.uchile.cl).

C. Thraves Caro is with the Depto. Ing. Matemática, Facultad de Ciencias Físicas y Matemáticas, University of Concepción, Concepción, Chile (e-mail: cthraves@ing-mat.udec.cl).

Communicated by C. E. Koksal, Associate Editor for Communication Networks.

Digital Object Identifier 10.1109/TIT.2016.2632153

[1]Further on, we argue that restricting the nodes to accessing either 1 or $O(1)$ messages from the buffer is irrelevant.

a message, every incoming call goes directly to the buffer. Messages that can not be stored due to insufficient buffer space are dropped. For the model to be unambiguously determined, one needs to fix the *buffer scheduling policy* which specifies, for each buffer, the set of stored messages (if any) that can be accessed by the buffer's processor upon the next request. A typical and popular choice of scheduling policy is First-In First-Out (FIFO). We argue that when messages are locally enqueued, FIFO might not be a good choice of scheduling policy.

1) Main Contributions: An immediate consequence of our proposed model setup will be that rumors will always propagate slower through the network in comparison with the classical setting. This is natural giving that we only add restrictions to how network nodes interact. Thus, the best we can hope is for no slowdown.

We establish similarities and, most importantly, crucial differences between the performance of PUSH and PULL when they are analyzed under the new framework as opposed to the classical one. In particular, we first observe that there is no difference between the classical setting and our new framework when the spreading algorithm is PUSH.

In contrast, when the spreading algorithm is PULL, the time complexity increases considerably, mostly because the amount of messages stored at each buffer may grow rapidly. Specifically, we consider two cases when we analyze the time complexity of PULL. First, we look at the case where all nodes have unbounded buffer size. In this scenario, in the classical model, we show that there is an n-node graph with maximum degree 4 and a starting source node for which PULL informs all nodes in $O(n \log n)$ rounds *with high probability* (abbreviated w.h.p.), i.e., with probability at least $1 - n^{-c}$ for any fixed constant $c > 0$. In contrast, in the new model, it requires $2^{\Omega(n)}$ rounds to inform all nodes. As a more positive counterpart, we show that the number of rounds required by PULL over Δ-regular graphs of diameter D is $O(D\Delta \max\{D, \Delta^2\} \ln n)$.

We then consider the case in which buffers have bounded size, and the size of each buffer is a parameter of the system denoted by B. Now, when a message arrives to a full buffer, the message is dropped. Hence, a second random factor appears in the system that takes into account the probability that a message is successfully enqueued in a buffer. In this context, we show that given a family of graphs of bounded degree (independent of n) whose graph conductance is at least Φ the number of rounds required by PULL to inform all nodes is $O(B\Phi^{-1} \log n)$ w.h.p. (In fact, we state our bound in terms of weighted graph conductance for a specific choice of weights that takes into account the probability that a message is successfully enqueued in a buffer — for restricted graph classes the weighted and standard graph conductance measure are within a $\Theta(1)$ factor of each other). We also establish an essentially matching lower bound that shows that, in general, the unfortunate linear dependency on B in the upper bound for the number of rounds is in fact unavoidable.

2) Related Work: Concerning the model itself, we are aware of only one other model that resembles some aspects of ours; the rumor spreading with bounded in-degree model proposed recently by Daum *et al.* [7]. Specifically, the restricted PULL protocol of Daum *et al.* where the requests to be served among a set of pull requests at a given node is chosen uniformly at random, which corresponds exactly to the case of our model where buffers are of size $B = 1$.

Time complexity of rumor spreading has been extensively studied. For instance, the time complexity of PUSH for the n-node complete graph and in subfamilies of Cayley graphs (including star, pancake, and transposition graphs – see [15], [21]). In [14], it is shown that in Δ-regular graphs, $\Delta \geq 3$, with probability $1 - o(1)$, PUSH informs all n nodes in $O(\ln n)$ rounds, where the hidden constant depends on Δ. Regarding random graphs, hypercubes and bounded degree graphs, Feige *et al.* [13] provide asymptotically optimal upper bounds for the time complexity of PUSH. These bounds are improved in [12] where tight lower and upper bounds are obtained. Sauerwald [22] relate the time complexity of PUSH to the mixing time of random walks. In [23] the relation between time complexity of PUSH and the vertex expansion of a graph is investigated.

Regarding PULL, it has been shown [8], [19] that in complete graphs if a constant fraction of the nodes are informed, then within $O(\log \log n)$ additional rounds every node of the graph becomes informed with probability $1 - o(1)$.

There also is a vast amount of literature concerning PUSH&PULL. In [17] and [23], the time complexity of PUSH&PULL is bounded in terms of the vertex expansion of the underlying network. For complete graphs it is known that PUSH&PULL requires $\Theta(\log n)$ rounds to spread the rumor w.h.p. [8], [9], [11], [19]. Chierichetti *et al.* [5] study rumor spreading in social networks generated according to the Barabassi-Albert preferential attachment model. For the same class of networks, Doerr *et al.* [10] later gave a tight analysis proving that $\Theta(\log n)$ rounds are sufficient, w.h.p., for PUSH&PULL to spread the rumor throughout the network.

Results that connect runtime of rumor spreading with conductance in general graphs were first established in [3] and [4]. Giakkoupis, in [16], presented a tight bound for rumor spreading via conductance. His main result says that for any graph with conductance at least Φ, PUSH&PULL informs all nodes in $O((1/\Phi) \log n)$ rounds w.h.p. This bound is tight, since there exist graphs where PUSH&PULL requires $\Omega((1/\Phi) \log n)$ rounds to spread the rumor to all nodes. Variants of these bounds, but for PUSH, are also given in [16].

3) Organization: In Section II, we define notation and give a detailed description of the proposed model. In Section III, we establish a tight upper bound for the time complexity of PULL in the bounded buffer model via weighted conductance. In Section IV, we study PULL in the unbounded buffer model. We conclude with some final comments in Section V.

II. BUFFER MODEL

We now describe in detail the new model. First, we fix some of the notation we use. All graphs we consider will be undirected and simple. Given a graph $G = (V, E)$ we let $N_W(v)$ denote the set of neighbors in $W \subseteq V$ of node v and denote by Δ_v the degree of v in G, i.e., $\Delta_v = |N_V(v)|$. Also, we let $\Delta_{\max} = \Delta_{\max}(G)$ and $D = D(G)$ denote the maximum degree and diameter of G, respectively. The set $\{1, \ldots, n\}$ will be denoted $[n]$.

A. Network

The network is modeled as an n node graph $G = (V, E)$. The node set V represents the computing entities and the set of edges E the point-to-point bi-directional communication links available. We do not assume any global node or link labeling; instead, only local link labeling is required to be able to select an outgoing port for a message to be sent and to identify the incoming point of a received message. Time is considered to be slotted in *synchronized steps*, also called *rounds*.

B. Buffers

Every node has a bounded size buffer where messages sent to the node are enqueued. The number of packets that buffers can hold is denoted by B and is a system parameter also referred to as buffer size given rise to what we henceforth refer to as *bounded buffer model* (or *bounded model*, for short). When B is very large, from the perspective of a finite-time execution, buffers are in practice "unbounded" and we refer to the model as the *unbounded buffer model* (or *unbounded model*, for short). We simply say *buffer model* when we want to refer to both the bounded and unbounded models.

C. Local Memory

Every node has a "small" (logarithmic in n) amount of memory. In particular, algorithms executed by nodes are allowed to store, in addition to the rumor, only a constant number of link identifiers (e.g., ids of connections to some of its neighbors).

D. Scheduling Policy

It specifies the order in which to process messages stored at a buffer. Decisions do not depend on the state of other nodes. Unless stated otherwise, we always assume that the scheduling policy is FIFO, with ties broken uniformly at random.

E. Local Steps: Sending, Delivering and Reading a Message

Each node runs a given algorithm in consecutive steps. At the start of a time step, every node can retrieve one message from its buffer (determined by the scheduling policy). Every node can send at most one message during each time step. This assumption, rather than for example allowing for a (potentially large) constant number of messages to be sent per time step is just for the sake of simplicity of exposition and in order to avoid unnecessarily cluttering notation. All our arguments can easily be adapted to yield analogous results but for nodes that can make $O(1)$ accesses to their buffers each time step. At the end of a step, messages arrive to their destination. If more messages than those that can be stored in a node's buffer arrive at any given step, then the excess messages are dropped. We always assume that the set of messages dropped are selected uniformly at random. Fig. 1 illustrates the order in which events take place during a step.

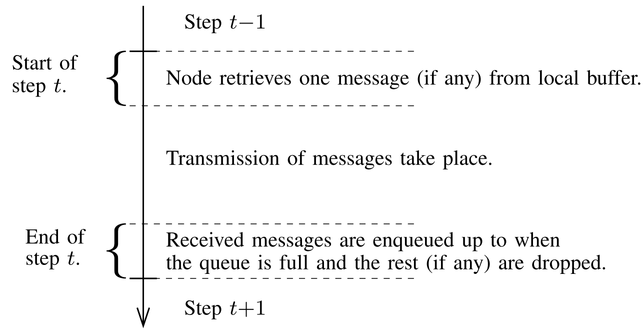

Fig. 1. Order of events that take place at a node during a step.

F. PUSH *and* PULL *Spreading Algorithms*

PUSH works in the buffer model as follows: if a node has the rumor stored in its memory, it sends a message with the rumor to one neighbor selected uniformly at random among all its neighbors (so called push action). If a node does not have the rumor stored in its memory, it waits until it receives the rumor (in the meantime, it keeps reading messages from its buffer, one reading attempt per step). PULL works analogously: If a node does not have the rumor stored in its memory, it sends a request to a neighbor selected uniformly at random among all its neighbors (so called pull request) and it reads one message from its buffer. If a node has the rumor stored in its memory, it retrieves messages from its buffer and answers a retrieved PULL request during the same step it is read. We recall that a node can read only one message from its buffer per time step — hence, the presence of a request in its buffer does not necessarily imply it will be answered immediately.

G. Rumor Spreading Problem

At the start of an execution, there is a single node, called *source node*, that has a rumor. The goal is to spread the rumor to every node in the network. In the execution of an algorithm, a node could be in one of two states: informed or uninformed. We say that a node is *informed* at a given step if it is the source node or it has retrieved from its buffer a message containing the rumor before or during the step. A node that is not yet informed is called *uninformed*. An uninformed node that has a message containing the rumor in its buffer is called *nearly-informed*. Note that a node can locally recognize and remember whether it is informed or not. However, an uninformed node cannot locally check if it is nearly-informed, as this would require reading from its buffer messages faster than one per step.

H. Time Complexity

Let $\mathcal{I}_t^{\mathcal{A}}(s)$ be the set of informed nodes at time t of an execution of the spreading algorithm $\mathcal{A}$ (for some specific occurrence of random events) when the rumor is initially at node s. We say that the *time complexity of* $\mathcal{A}$ is $O(T(n))$ w.h.p., if for every $c > 0$ and n sufficiently large, for every connected graph $G = (V, E)$ on n nodes and all possible sources s we have $\mathcal{I}_{T(n)}^{\mathcal{A}}(s) = V$ with probability at least $1 - n^{-c}$. We shall establish our w.h.p. results for $c = 1$.

In all cases considered, the adaptation to arbitrary $c > 0$ is straightforward.

Remark 1: The time complexity of PUSH in the buffer and classical models is the same. Intuitively this holds because if nodes use only PUSH, then uninformed nodes do not transmit any messages. Thus, no relevant queues are formed and the rate of propagation of the rumor throughout the network is as in the classical model. This explains why we henceforth only focus in studying PULL.

III. PULL IN THE BOUNDED BUFFER MODEL

When the underlying network is an arbitrary graph, the tighter known methods for upper bounding the time complexity of PULL in the classical model are based in graph conductance. In this section we derive analogous bounds for the time complexity of rumor spreading for PULL but in the buffer model. The bounds we establish are in terms of the so called weighted conductance, but for a particular set of weights that depends on the model and network parameters.

We start by recalling the notion of weighted conductance. Let $G = (V, E)$ be a graph with edge weights given by $\omega : E \to [0, 1]$. For $S, T \subseteq V$, the sum of the weights of edges with one endnode in S and the other in T is called the *weighted (S, T) edge cut in G* and is denoted $\mathrm{cut}^{\omega}_G(S, T)$, i.e., $\mathrm{cut}^{\omega}_G(S, T) = \sum_{st\in E: s\in S, t\in T} \omega(st)$. Moreover, the sum of the degrees of a set of nodes $S \subseteq V$ is called the *volume of S in G* and is denoted $\mathrm{vol}_G(S)$, i.e., $\mathrm{vol}_G(S) = \sum_{v\in S} \Delta_v$. The *weighted conductance of G* is defined as

$$\Phi_\omega(G) = \min_{S\subseteq V: 0<\mathrm{vol}_G(S)\leq |E|} \frac{\mathrm{cut}^{\omega}_G(S, V \setminus S)}{\mathrm{vol}_G(S)}.$$

When clear from context, we drop the subindex G from the notation introduced above. Also, if we omit ω, it is to be understood that it is identically 1 over E.

We now describe a natural weighted variant of PULL which takes into consideration random transmission failures. For every edge $e = uv \in E$, assuming v has the rumor, we consider a probability $\omega(e) \in [0, 1]$ such that a request sent by u to v is successfully replied to. The events associated to distinct edges are assumed independent. The probability that in any one round u chooses v as the transmission destination is as before, i.e., $1/\Delta_u$.

For the non-weighted case, i.e., when ω assigns weight 1 to every edge of G, the definition of $\Phi_\omega(G)$ coincides with the standard notion of graph conductance and the weighted versions of PUSH and PULL correspond to the classical cases, respectively. In this situation, Giakkoupis [16, Lemma 4] showed that the time complexity of PULL in the classical setting could be upper bounded in terms of the graph conductance. The proof argument in [16] readily generalizes to the weighted case and yields:

Theorem 1 (Giakkoupis [16]): For an initial set of informed vertices S_0 of $G = (V, E)$ and any fixed $\beta > 0$, all vertices get informed in at most $O(1)(\beta + 2)\big(1 + \Delta_{\max} / \mathrm{vol}_G(S_0)\big)\Phi^{-1}_\omega(G) \log |V|$ rounds of the PULL algorithm, with probability $1 - O(|V|^{-\beta})$.

We now informally discuss our key insight for analyzing the PULL protocol in the bounded buffer model. We view the rumor spreading process as a weighted version of the classical PULL protocol over time windows of length $2B$, for an appropriate choice of weights. The length of the window is motivated by the fact that a request sent at time t from node u to a neighbor node v, if successfully enqueued at node v, might end up spending B time steps at v's queue before being serviced, and when serviced, might end up spending another B time steps in u's queue if the response to the original request is successfully enqueued there. However, there is no guarantee that either a request or the response to a request will succeed in being enqueued at the adequate node destination. The probability of such a thing happening is what motivates us to consider weighted versions of rumor spreading, and justifies why we require the determination of the probability of occurrence of certain events such as the ones defined in the following results.

From now on, when not otherwise specified, we assume that the local rumor spreading protocol is PULL, that we are working in the buffer model, and that the underlying network is $G = (V, E)$. We denote by S_t (respectively, U_t) the set of informed (respectively, uninformed) nodes at time t. Clearly, we have $U_t = V \setminus S_t$.

Lemma 1: Consider $su \in \mathrm{cut}_G(S_t, U_t)$. Let $\mathcal{E}_{t,s,u}$ be the event "given that u sends a request to s during round $t + B$, the request sent by u is enqueued at s". Then,

$$\mathbb{P}(\mathcal{E}_{t,s,u}) \geq \Big(1 + \sum_{u'\in N_{U_t}(s): u'\neq u} \frac{1}{\Delta_{u'}}\Big)^{-1}.$$

Proof: Let $\mathcal{T}$ be the collection of uninformed neighbors of s that transmit a message to s during round $t+B$. Since $s \in S_t$, node s was informed at most by time step t. Thus, s does not send any rumor request during steps $t, \ldots, t+B-1$. Moreover, any such request enqueued at time t in one of the neighbors of s must have been dealt with (hence, purged from the queue) before time step $t + B$. It follows that during step $t + B$, messages sent to s can only come from nodes uninformed at time $t + B$. By definition of $\mathcal{E}_{t,s,u}$, we get

$$\mathbb{P}(\mathcal{E}_{t,s,u}) \geq \sum_{C\subseteq N_{U_{t+B}}(s): u\in C} \frac{1}{|C|} \cdot \mathbb{P}(\mathcal{T} = C).$$

Moreover, conditioned on u sending a request to s in round $t + B$,

$$\sum_{C\subseteq N_{U_{t+B}}(s): u\in C} \mathbb{P}(\mathcal{T} = C) = \mathbb{P}(u \in \mathcal{T}) = 1.$$

Jensen's inequality and convexity of $x \mapsto 1/x$ for $x > 0$, together with the fact that $U_{t+B} \subseteq U_t$ imply that

$$\begin{aligned}\mathbb{P}(\mathcal{E}_{t,s,u}) &\geq \Big(\sum_{C\subseteq N_{U_{t+B}}(s): u\in C} |C| \cdot \mathbb{P}(\mathcal{T} = C)\Big)^{-1} \\ &= \Big(\sum_{C\subseteq N_{U_{t+B}}(s)\setminus\{u\}} |C \cup \{u\}| \cdot \mathbb{P}(\mathcal{T} \setminus \{u\} = C)\Big)^{-1} \\ &\geq \Big(\sum_{C\subseteq N_{U_t}(s)\setminus\{u\}} |C \cup \{u\}| \cdot \mathbb{P}(\mathcal{T} \setminus \{u\} = C)\Big)^{-1}.\end{aligned}$$

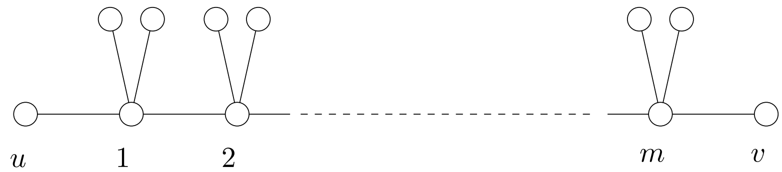

Fig. 2. Graph H.

The desired conclusion follows since the expression inside parenthesis in the last displayed term corresponds to 1 plus the expectation of the sum of Bernoulli random variables $X_{u'}$, $u' \in N_{U_t}(s) \setminus \{u\}$, each with expectation $1/\Delta_{u'}$. ■

Lemma 2: Consider $su \in \mathrm{cut}_G(S_t, U_t)$. Let $\mathcal{F}_{t,s,u}$ be the event "given that u sends a request to s during round $t + B$, the request is successfully enqueued at s and by round $t+2B$ node u becomes nearly informed". Then,

$$\mathbb{P}(\mathcal{F}_{t,s,u}) \geq \Big(1 + \sum_{\substack{u' \in N_{U_t}(s) \\ u' \neq u}} \frac{1}{\Delta_{u'}}\Big)^{-1} \Big(1 + \sum_{u' \in N_{U_t}(u)} \frac{1}{\Delta_{u'}}\Big)^{-1}.$$

Proof: Clearly $\mathcal{F}_{t,s,u}$ happens if $\mathcal{E}_{t,s,u}$ as defined in Lemma 1 occurs and the reply sent (at time $t + 2B$) by s to the request made by u (at time $t+B$) is successfully enqueued at u. Since $U_{t+2B} \subseteq U_t$ and given that the messages sent to u during time step $t + 2B$ either carry the rumor or originate in uninformed nodes, a similar proof argument as the one of Lemma 1 yields that

$$\mathbb{P}(\mathcal{F}_{t,s,u} | \mathcal{E}_{t,s,u}) \geq \Big(1 + \sum_{u' \in N_{U_t}(u)} \frac{1}{\Delta_{u'}}\Big)^{-1}.$$

To conclude, note that $\mathbb{P}(\mathcal{F}_{t,s,u}) = \mathbb{P}(\mathcal{F}_{t,s,u}|\mathcal{E}_{t,s,u})\mathbb{P}(\mathcal{E}_{t,s,u})$ and apply Lemma 1. ■

Henceforth, we denote by ω_G the map $\omega : E \to [0, 1]$ such that for $su \in E$

$$\omega(su) = \min_{\substack{\emptyset \subsetneq U \subsetneq V \\ su \in \mathrm{cut}_G(V \setminus U, U)}} \Big(1 + \sum_{\substack{u' \in N_U(s) \\ u' \neq u}} \frac{1}{\Delta_{u'}}\Big)^{-1} \Big(1 + \sum_{u' \in N_U(u)} \frac{1}{\Delta_{u'}}\Big)^{-1}.$$

Taking $U = V \setminus \{s\}$ in the left hand side above and given that $N_W(s) \subseteq N_V(s)$ for all $W \subseteq V$, it follows that

$$\omega(su) = \Big(1 + \sum_{\substack{u' \in N_V(s) \\ u' \neq u}} \frac{1}{\Delta_{u'}}\Big)^{-1} \Big(1 + \sum_{\substack{u' \in N_V(u) \\ u' \neq s}} \frac{1}{\Delta_{u'}}\Big)^{-1}.$$

The following (easily derived) claim relates Φ_ω with the classical notion of graph conductance Φ.

Proposition 1: Let $G = (V, E)$ be a connected graph where $|V| > 2$ and let $\omega = \omega_G$. If $\Delta_{\min}$ and $\Delta_{\max}$ are the smallest and largest degree in G, then $\big(\frac{\Delta_{\min}}{\Delta_{\max}+\Delta_{\min}-1}\big)^2 \Phi \leq \Phi_\omega \leq \big(\frac{\Delta_{\max}}{\Delta_{\max}+\Delta_{\min}-1}\big)^2 \Phi$. In particular, if G is Δ-regular, then $\Phi_\omega = \big(\frac{\Delta}{2\Delta-1}\big)^2 \Phi$. If the ratio of the degrees of neighboring nodes in G is at most δ, then $\big(\frac{1}{\delta+1}\big)^2 \Phi \leq \Phi_\omega \leq \big(\frac{2\delta}{\delta+2}\big)^2 \Phi$.

We now establish this section's main result.

Theorem 2: Let $G = (V, E)$ be a graph, $\omega = \omega_G$, and $\Phi_\omega = \Phi_\omega(G)$. For any initial set of informed vertices $S_0 \subseteq V$ and any fixed $\beta > 0$, with probability $1 - O(|V|^{-\beta})$, all vertices get informed in at most $O(1)B(\beta + 2)\big(1 + \Delta_{\max}/\operatorname{vol}(S_0)\big)\Phi_\omega^{-1} \log|V|$ rounds by PULL when buffers are of size B.

Proof: Follows directly from Theorem 1 and Lemma 2 by partitioning time into intervals of length $3B$, noting that $\mathbb{P}(\mathcal{F}_{t,s,u}) \geq \omega(su)$ and observing that if a node is nearly-informed by time step $t+2B$ it will be informed by time step $t+3B$. ■

We next describe a network for which Theorem 2 is essentially tight. The construction is motivated by the following footnote claim of Chierichetti *et al.* [4, Footnote 2] asserting that the conductance based upper bound on the time complexity of rumor spreading in the classical model given by Giakkoupis [16] is tight: *For any $\epsilon > 0$, positive integer n, and $\Phi = \Phi(G) \geq n^{-1+\epsilon}$, if G is a 3-regular graph of constant expansion on $O(n\Phi)$ nodes, then replacing each edge of G by a path of $O(\Phi^{-1})$ new nodes yields an $O(n)$ node graph with conductance* $\Omega(\Phi)$. We claim that if edges are replaced by the graph H of Fig. 2 instead of by paths, then a network for which Theorem 2 is tight is obtained. Intuitively, this happens because whenever a rumor arrives at an internal node in the u-v path of a copy of H whose buffer is full, it will henceforth encounter a full buffer in every internal node in the u-v path of every copy of H it moves into. Although Chierichetti et al. do not provide a detailed proof of their assertion, it is indeed not hard to reconstruct, although unfortunately it requires some space to corroborate carefully. The same is true, but even more so, in our case.

For a graph $G = (V, E)$ we let $G' = (V', E')$ be the graph obtained from G by replacing each of its edges, say with endpoints u and v, by the graph of H of Fig. 2, where $m = \lceil 1/\gamma \rceil$ and $0 < \gamma < 1$. The nodes of the resulting graph G' that arise as copies of nodes of degree 1 in H will be henceforth called leafs. As usual, denote by $D(G)$ the diameter of G. From now on, we assume each node's buffer is of size B. Let $s \in V$ be a node for which there is some $t \in V$ such that the number of edges in any path in G between s and t is $D(G)$ (such pair of nodes exists by definition of $D(G)$). We henceforth assume that the source of the rumor is s. Finally, for $\ell \in \mathbb{N}$, let V'_ℓ be the set of nodes $v' \in V' \setminus V$ such that for every path P in G' between s and v' it holds that at least ℓ internal nodes of P belong to $V' \setminus V$, i.e., $|(V(P) \setminus \{s, v'\}) \cap (V' \setminus V)| \geq \ell$.

Lemma 3: Before it becomes nearly-informed, non-leaf node $v' \in V'_\ell$ has at least $\min\{B, 2^\ell\} - 1$ messages queued in its buffer.

Proof: The statement is trivially true for $\ell = 0$. Assume it holds for $\ell > 0$. Consider a shortest path P in G' between s

and v'. Let u' be the internal node of P closest to v'. Clearly, u' is a non-leaf node and $u', v' \in V'_{\ell-1}$. By inductive hypothesis, just before u' becomes nearly-informed both u' and v' have at least $\min\{B, 2^{\ell-1}\} - 1$ messages queued in their buffers. Hence, the time step that u' becomes nearly-informed plus the time the rumor spends in u''s queue is at least $\min\{B, 2^{\ell-1}\}$. During such period of time, given that the leaves attached to v' in G' cannot become informed before v' and since v' can process at most 1 message per time step, its queue increases to at least $\min\{B, 2^{\ell}\} - 1$ messages. The claimed result follows. ∎

Corollary 1: The number of steps until node t is informed is at least

$$B\big(mD(G) - \lceil \log_2 B \rceil\big) \geq \frac{B}{\gamma}\big(D(G) - \log_2 B - O(1)\big).$$

Proof: Note that any path P in G' between s and t has at least $mD(G)$ of its internal nodes in $V' \setminus V$. However, all but at most $\ell^* := \lceil \log_2 B \rceil$ of such internal nodes can belong to V_{ℓ^*}. By Lemma 3, for every internal node v' of P in V_{ℓ^*}, the rumor takes at least one step to reach v' from an adjacent node and spends at least $\min\{B, 2^{\ell^*}\} - 1 \geq B-1$ steps enqueued at v'. Thus, the number of steps until t is informed is at least $B(mD(G) - \ell^*) = B(mD(G) - \lceil \log_2 B \rceil)$. The desired conclusion follows since $\lceil x \rceil \leq x + 1$ and $\gamma < 1$ ∎

The next two results are established as required to prove Chierichetti *et al.*'s [4] footnote claim.

Lemma 4: Let $G = (V, E)$ be a 3-regular graph and $\omega' = \omega_{G'}$. If $0 < \gamma < 1$, then $\Phi_{\omega'}(G') \leq \gamma$. Moreover, if G is a 3-regular expander graph, then $\Phi_{\omega'}(G') \geq \gamma\,\Omega(1)$.

Proof: Note that $|E'| = (3m + 1)|E|$ and $|E| = 3|V|/2$. Let S' be the set of new nodes that are created when a given edge uv of G is replaced by the graph H of Fig. 2. Clearly, $\mathrm{vol}_{G'}(S') = 6m \leq |E'|$ since $|E| \geq 3$, so $\mathrm{vol}_{G'}(S'^c) > \mathrm{vol}_{G'}(S')$. Also note that $|\mathrm{cut}_{G'}(S', V' \setminus S')| = 2$. It follows immediately from the definition of $\Phi_{\omega'}(\cdot)$, the 3-regularity of G, and some basic arithmetic, that

$$\Phi_{\omega'}(G') \leq \frac{\mathrm{cut}^{\omega'}_{G'}(S', V' \setminus S')}{\mathrm{vol}_{G'}(S')} = \frac{|\mathrm{cut}_{G'}(S', V' \setminus S')|}{(5/3) \cdot \mathrm{vol}_{G'}(S')} = \frac{1}{5m} \leq \frac{\gamma}{5} \leq \gamma.$$

Now, for the second part. Let S' be a non-trivial subset of V'. Because the maximum degree of any node of G' is 4 and each such node has at most two neighbors of degree 1, we get that $\omega'(e') \geq (2/7)^2$ for each $e' \in E'$, so

$$\Phi_{\omega'}(S') = \frac{\mathrm{cut}^{\omega'}_{G'}(S', V' \setminus S')}{\min\{\mathrm{vol}_{G'}(S'), \mathrm{vol}_{G'}(V' \setminus S')\}} \geq \frac{|\mathrm{cut}_{G'}(S', V' \setminus S')|/(2/7)^2}{\min\{\mathrm{vol}_{G'}(S'), \mathrm{vol}_{G'}(V' \setminus S')\}}.$$

Let $S = S' \cap V$. Note that $V \setminus S = (V' \setminus S') \cap V$. We can assume that $\mathrm{vol}_G(S) \leq |E| \leq \mathrm{vol}_G(V \setminus S)$, since otherwise we can replace S' by $V' \setminus S'$ in what follows.

Let F be the subgraph of G' induced by S'. Say a node $v \in S'$ is bad if the connected component of F to which v belongs does not contain a node that belongs to V. Denote by B' the set of bad nodes. Note that $B' \subseteq S'$, so $\mathrm{vol}_{G'}(S') = \mathrm{vol}_{G'}(B') + \mathrm{vol}_{G'}(S' \setminus B')$. Since every node of G' has degree at most 4, it follows that $\mathrm{vol}_{G'}(B') \leq 4|B'|$. Note that by the way in which badness is defined, $\mathrm{cut}_{G'}(B', S' \setminus B') = \emptyset$. Also, because the largest connected component induced by $V' \setminus V$ in G' is of size $3m$, it must hold that

$$\begin{aligned}\big|\mathrm{cut}_{G'}(S', V' \setminus S')\big| &= \big|\mathrm{cut}_{G'}(B', V' \setminus B')\big| + \big|\mathrm{cut}_{G'}(S' \setminus B', V' \setminus (S' \setminus B'))\big| \\ &\geq \frac{|B'|}{3m} + \big|\mathrm{cut}_{G'}(S' \setminus B', V' \setminus (S' \setminus B'))\big|.\end{aligned}$$

Putting everything together,

$$\begin{aligned}\Phi_{\omega'}(S') &\geq \frac{|\mathrm{cut}_{G'}(S', V' \setminus S')|}{(2/7)^2\,\mathrm{vol}_{G'}(S')} \\ &\geq \frac{1}{(2/7)^2} \cdot \frac{\frac{|B'|}{3m} + \big|\mathrm{cut}_{G'}(S' \setminus B', V' \setminus (S' \setminus B'))\big|}{4|B'| + \mathrm{vol}_{G'}(S' \setminus B')}.\end{aligned}$$

So, if $B' = S'$ (equivalently, $\mathrm{vol}_{G'}(S' \setminus B') = 0$), then $\Phi_{\omega'}(S') \geq \gamma\,\Omega(1)$ as claimed. Otherwise, recalling that for $a, c \geq 0$ and $b, d > 0$, it holds that $(a + c)/(b + d) \geq \min\{a/b, c/d\}$, we get that

$$\Phi_{\omega'}(S') \geq \frac{49}{4} \min\left\{\frac{|\mathrm{cut}_{G'}(S' \setminus B', V' \setminus (S' \setminus B'))|}{\mathrm{vol}_{G'}(S' \setminus B')}, \frac{1}{12m}\right\}.$$

Since for the graph H of Fig. 2 the volume of $H \setminus \{u, v\}$ is $6m$ and every node in $S' \setminus B'$ is in the same connected component of a node in S, if follows that $\mathrm{vol}_{G'}(S' \setminus B') \leq (6m + 1)\,\mathrm{vol}_G(S)$. Moreover, note that $|\mathrm{cut}_{G'}(S' \setminus B', V' \setminus (S' \setminus B'))| \geq |\mathrm{cut}_G(S, V \setminus S)|$. Thus,

$$\Phi_{\omega'}(S') \geq \frac{49}{4} \min\left\{\frac{|\mathrm{cut}_G(S, V \setminus S)|}{(6m + 1)\,\mathrm{vol}_G(S)}, \frac{1}{12m}\right\} = \gamma\,\Omega(1),$$

where the last inequality is due to the fact that $\mathrm{vol}_G(S) \leq |E|$ and G is an expander graph, so $|\mathrm{cut}_G(S, V \setminus S)|/\mathrm{vol}_G(S) = \Omega(1)$. ∎

Proposition 2: For every $\epsilon > 0$ and $B : \mathbb{N} \to \mathbb{N}$ such that $1 \leq B(n') < n'^{1-\epsilon}$, there is a graph $G' = (V', E')$ over n' vertices and $\omega' = \omega_{G'}$, such that: (1) $\Phi_{\omega'}(G') \geq B(n')/n'^{1-\epsilon}$ and, (2) the time complexity of PULL in the buffer model for G' is at least $\Omega(B(n')\log n'/\Phi_{\omega'}(G'))$.

Proof: Let $\gamma = \gamma(n') = B(n')/n'^{1-\epsilon} < 1$. Let $G = (V, E)$ be a 3-regular expander graph over n nodes such that $n' = (1 + 9m/2)n$ (note that n must be even, so n' is an integer). Also note that $n \geq (\frac{2}{9} - O(1))\gamma n'$. Construct $G' = (V', E')$ from G and choose s and t as discussed previously. Since G is 3-regular, $|E| = 3|V|/2$, and because of the way in which G' is constructed, $|V'| = |V| + 3m|E| = (1+9m/2)n = n'$. Moreover, since for any d-regular graph, $d \geq 3$, it holds that $D(G) \geq \log_{d-1}|V| - 1$, we get that $D(G) \geq \log_2 n - 1 \geq \log_2(\gamma n') - O(1)$. By Corollary 1 and our choice of γ, it follows that the time complexity for rumor spreading in G' is at least

$$\begin{aligned}&\frac{B(n')}{\gamma(n')}\Big(\log_2 \frac{\gamma(n')n'}{B(n')} - O(1)\Big) \\ &= \frac{B(n')}{\gamma(n')}\Big(\epsilon \log_2 n' - O(1)\Big) \\ &= \Omega\Big(\frac{B(n')}{\gamma(n')} \log_2 n'\Big).\end{aligned}$$

The conclusion follows because by Lemma 4, $\Phi_{\omega'}(G') = \Theta(\gamma(n'))$. ∎

IV. PULL IN THE UNBOUNDED BUFFER MODEL

In this section we study the behavior of PULL in the unbounded buffer model. We begin by highlighting the significant performance difference between PULL in the classical model and in the unbounded buffer model.

Proposition 3: There is a maximum degree 4 graph $G = (V, E)$ on n nodes and a source $s \in V$ for which PULL informs all nodes in $O(n \log n)$ rounds, w.h.p., in the classical model, and in $\Omega(1)2^{n/3}$ in the unbounded buffer model.

Proof: Let H, u, v and m be as depicted in Fig. 2. Consider as source of the rumor the node u. Note that H has $n = 3m + 2$ nodes. Since the time complexity of PUSH in any n-node connected graph is $O(n \log n)$ (see [13]), the bound also holds for H. To conclude, observe that by taking $B = +\infty$ in Lemma 3, we get that the number of rounds until v is informed is at least 2^m. Thus, H satisfies the conditions required of G in the statement. ■

We next establish a more positive counterpart to the preceding result, but restricted to regular graphs. The following two facts would be useful to establish separately.

Lemma 5: Let u and v be neighbors in G, and let $\Delta_{\max} = \Delta_{\max}(G)$. Suppose that v is uninformed when u becomes informed. Then, with probability at least $1 - 1/n^2$, from the time step when u becomes informed to the time step when v is nearly-informed, at most $2\Delta^2_{\max}(\Delta_{\max} - 1)\ln n$ messages get enqueued at v

Proof: Let $t = 2\Delta_{\max} \ln n$. The probability that v makes no request to u during the following $t = 2\Delta_{\max} \ln n$ steps after u is nearly-informed is at most $(1 - 1/\Delta_v)^t \le e^{-t/\Delta_v} \le e^{-t/\Delta_{\max}} = 1/n^2$. During an interval of time t, since no queue can increase by more than $(\Delta_{\max} - 1)$ per round, at most $t(\Delta_{\max} - 1)$ messages could have queued at u during the interval. Thus, with probability at least $1 - 1/n^2$ the node v becomes nearly-informed at most $t + t(\Delta_{\max} - 1) = t\Delta_{\max}$ steps after u. The desired conclusion follows recalling that at most $(\Delta_{\max} - 1)$ messages can be enqueued at any single node per time step. ■

Henceforth, given a source node in G, let V_d be the set of nodes of G at distance at most d from the source. Moreover, let τ_d denote the maximum, among all nodes $v \in V_d$, of the time v is informed. Since the only node at distance 0 from the source is the source itself, we clearly have $\tau_0 = 0$.

Lemma 6: Let $G = (V, E)$ be a connected Δ-regular graph. Then, for $d \in \mathbb{N}$, $d < D(G)$, with probability $1 - O(|V_{d+1} \setminus V_d|)/n^2$,

$$\tau_{d+1} \le \tau_d + \sqrt{\tau_d \Delta \ln n} + 2\Delta^2(\Delta - 1)\ln n.$$

Proof: Let $v \in V_{d+1} \setminus V_d$ and assume the first neighbor of v, say u, is informed at time $T \le \tau_d$. Note that such a neighbor need not belong to V_d. By Lemma 5, with probability at least $1-1/n^2$ node v enqueues at most $2\Delta^2(\Delta-1)\ln n$ messages before becoming nearly-informed but after u becomes informed.

We also need to bound the number of messages enqueued at v during the first T rounds. Now though, since none of v's neighbors is informed up to T, we know that each neighbor u of v sends a request to v with probability $1/\Delta$, independently of what other request v's neighbors send, and independent of what happened in previous rounds. We thus expect that roughly $T \sum_{u \in N_V(v)} \frac{1}{\Delta} = T$ requests would have arrived at v before time step T and that roughly T would have been serviced (given that v's queue could have remained empty during some rounds, not exactly T messages are necessarily serviced). We claim that the number of messages enqueued at v in round T is not significantly large. Let $X_{u,t}$ be the indicator of whether u sends a request to v during round t and let Q_t be the number of messages enqueued at v at the beginning of round t. We wish to bound the probability that Q_T is large. Let $A = \sqrt{T\Delta \ln n}$. Clearly, for $T < (\ln n)/\Delta$ we have $T\Delta < A$, so $\mathbb{P}(Q_T \ge A) = 0$, because in T steps at most $T(\Delta - 1)$ messages reach node v. Assume $T \ge (\ln n)/\Delta$. Partitioning according to the last round $0 \le t < T$ in which v's queue was empty, a union bound, the preceding discussion, a Chernoff [1, Th. A.14] bound, and Jensen's inequality, yield that

$$\begin{aligned}
\mathbb{P}(Q_T \ge \sqrt{T\Delta \ln n}) &= \mathbb{P}(Q_T \ge A)\\
&\le \sum_{0 \le t < T} \mathbb{P}(Q_T \ge A, Q_t = 0, \forall k \in [T-t], Q_{t+k} \ne 0)\\
&\le \sum_{0 \le t < T} \mathbb{P}\Big(\sum_{k \in [T-t]} \sum_{u \in N_V(v)} X_{u,t+k} \ge A + T - t\Big)\\
&\le \sum_{t=0}^{T-1} e^{-2T(\ln n)/(T-t)} \le (1 + o(1))Te^{-2\ln n \ln T}\\
&= O(1/n^2).
\end{aligned}$$

Since the time between v being nearly-informed and becoming informed is at most the queue size it had when becoming nearly-informed, it follows that with probability $1 - O(1/n^2)$ node v becomes informed by round

$$T + \sqrt{T\Delta \ln n} + 2\Delta^2(\Delta - 1)\ln n \le \tau_d + \sqrt{\tau_d \Delta \ln n} + 2\Delta^3 \ln n.$$

A union bound over the nodes in $V_{d+1} \setminus V_d$ yields the desired conclusion. ■

We can now establish this section's main result.

Proposition 4: The time complexity of PULL in the unbounded buffer model on an n-node Δ-regular network of diameter D is $4D\Delta \max\{\Delta^2, D\} \ln n$.

Proof: From Lemma 6, we get that with probability at least $1 - |V_D|/n^2 = 1 - 1/n$ the recurrence $\tau_d \le \tau_{d-1} + \sqrt{\tau_{d-1}\Delta \ln n} + 2\Delta^3 \ln n$ holds for all $d \in [D]$. By induction, one can show that $\tau_d \le 4d\Delta \max\{\Delta^2, D\} \ln n$ holds for all $d < D$. ■

V. CONCLUSIONS

In this paper we formally introduced a new model for randomized rumor spreading. The model encapsulates a more realistic scenario than classically considered in the literature. Specifically, we consider a restriction on the number of messages, per time step, that network nodes can act upon. We demonstrated differences between the buffer and classical model which are particularly striking in the case when the underlying network is not necessarily regular.

By definition of our buffer model, rumors spread slower than in the classical setting. For PULL, this is somewhat of an understatement, since as we have seen they can spread dramatically slower compared to the classical rumor spreading setting. The slowdown is somewhat less dramatic over regular networks. The slower propagation of rumors arises because FIFO queues prioritize messages that have been in the system longer. However, one actually would like to prioritize messages more recently generated, mainly for two reasons: (1) newer requests are more informative than older ones (a node requesting a rumor might have been informed since it sent the request), and (2) older messages arriving from a neighbor are less likely to carry a rumor than younger request responses from the same neighbor. The preceding discussion suggests that Last-In First-Out (LIFO) rather than FIFO should be used as scheduling policy. The work of Daum*et al.* [7] shows, stated in the language of this paper, that the slowdown of using PULL in the bounded buffer model with buffers of size 1 versus using the classic pull protocol can, w.h.p., be upper bounded by $O((\Delta_{\max}/\Delta_{\min})\log n)$ where $\Delta_{\max}$ and $\Delta_{\min}$ as used previously in this article. The same claim clearly holds for larger size buffers (even unbounded ones).

ACKNOWLEDGMENT

The authors would like to thank Prof. Dariusz R. Kowalski for many useful discussions during the early stages of this article. They are grateful to the anonymous reviewers for their careful reading of our manuscript and many observations.

REFERENCES

[1] N. Alon and J. Spencer, *The Probabilistic Method*. 3rd ed. New York, NY, USA: Wiley, 2008.
[2] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Randomized gossip algorithms," *IEEE Trans. Inf. Theory*, vol. 52, no. 6, pp. 2508–2530, Jun. 2006.
[3] F. Chierichetti, S. Lattanzi, and A. Panconesi, "Almost tight bounds for rumour spreading with conductance," in *Proc. 42nd ACm Symp. Theory Comput. (STOC)*, 2010, pp. 399–408.
[4] F. Chierichetti, S. Lattanzi, and A. Panconesi, "Rumour spreading and graph conductance," in *Proc. ACM-SIAM Symp. Discrete Algorithms (SODA)*, 2010, pp. 1657–1663.
[5] F. Chierichetti, S. Lattanzi, and A. Panconesi, "Rumor spreading in social networks," *Theor. Comput. Sci.*, vol. 412, no. 24, pp. 2602–2610, May 2011.
[6] B. S. Chlebus, K. Diks, and A. Pelc, "Fast gossiping with short unreliable messages," *Discrete Appl. Math.*, vol. 53, nos. 1–3, pp. 15–24, Sep. 1994.
[7] S. Daum, F. Kuhn, and Y. Maus, "Rumor spreading with bounded in-degree," in *Proc. 23rd Int. Colloq. Struct. Inf. Commun. Complexity (SIROCCO)*, Jul. 2016, pp. 323–339.
[8] A. Demers *et al.*, "Epidemic algorithms for replicated database maintenance," *Oper. Syst. Rev.*, vol. 22, no. 1, pp. 8–32, Jan. 1988.
[9] B. Doerr and M. Fouz, "Asymptotically optimal randomized rumor spreading," in *Proc. 38th Int. Colloq. Automata, Lang. Program. (ICALP)*, 2011, pp. 502–513.
[10] B. Doerr, M. Fouz, and T. Friedrich, "Social networks spread rumors in sublogarithmic time," in *Proc. 43rd ACM Symp. Theory Comput. (STOC)*, 2011, pp. 21–30.
[11] R. Elsässer, "On the communication complexity of randomized broadcasting in random-like graphs," in *Proc. 18th ACM Symp. Parallelism Algorithms Architectures (SPAA)*, 2006, pp. 148–157.
[12] R. Elsässer and T. Sauerwald, "On the runtime and robustness of randomized broadcasting," *Theor. Comput. Sci.*, vol. 410, no. 36, pp. 3414–3427, Aug. 2009.
[13] U. Feige, D. Peleg, P. Raghavan, and E. Upfal, "Randomized broadcast in networks," *Random Struct. Algorithms*, vol. 1, no. 4, pp. 447–460, 1990.
[14] N. Fountoulakis and K. Panagiotou, "Rumor spreading on random regular graphs and expanders," *Random Struct. Algorithms*, vol. 43, no. 2, pp. 201–220, Sep. 2013.
[15] A. M. Frieze and G. R. Grimmett, "The shortest-path problem for graphs with random arc-lengths," *Discrete Appl. Math.*, vol. 10, no. 1, pp. 57–77, Jan. 1985.
[16] G. Giakkoupis, "Tight bounds for rumor spreading in graphs of a given conductance," in *Proc. 28th Symp. Theor. Aspects Comput. Sci. (STACS)*, 2011, pp. 57–68.
[17] G. Giakkoupis and T. Sauerwald, "Rumor spreading and vertex expansion," in *Proc. ACM-SIAM Symp. Discrete Algorithms (SODA)*, 2012, pp. 1623–1641.
[18] M. Harchol-Balter, T. Leighton, D. Lewin, "Resource discovery in distributed networks," in *Proc. ACM Symp. Principles Distrib. Comput. (PODC)*, 1999, pp. 229–237.
[19] R. Karp, C. Schindelhauer, S. Shenker, and B. Vocking, "Randomized rumor spreading," in *Proc. 41st Annu. Symp. Found. Comput. Sci. (FOCS)*, Nov. 2000, pp. 565–574.
[20] D. W. Krumme, G. Cybenko, and K. N. Venkataraman, "Gossiping in minimal time," *SIAM J. Comput.*, vol. 21, no. 1, pp. 111–139, Feb. 1992.
[21] B. Pittel, "On spreading a rumor," *SIAM J. Appl. Math.*, vol. 47, no. 1, pp. 213–223, Feb. 1987.
[22] T. Sauerwald, "On mixing and edge expansion properties in randomized broadcasting," *Algorithmica*, vol. 56, no. 1, pp. 51–88, Jan. 2010.
[23] T. Sauerwald and A. Stauffer, "Rumor spreading and vertex expansion on regular graphs," in *Proc. ACM-SIAM Symp. Discrete Algorithms (SODA)*, 2011, pp. 462–475.
[24] R. van Renesse, Y. Minsky, and M. Hayden, "A gossip-style failure detection service," in *Proc. IFIP Int. Conf. Distrib. Syst. Platforms Open Distrib. Process. Middleware*, 1998, pp. 55–70.

Marcos Kiwi received the Engineering degree from U. de Chile in 1991 and the Ph.D. degree from MIT in 1996. He is Full Professor in the Department of Mathematical Engineering and an Associate Researcher at the Center for Mathematical Modeling (CNRS UMI 2807), both at U. de Chile, Santiago, Chile. His main research interests are theory of computation and random structures.

Christopher Thraves Caro is Assistant Professor at the Department of Mathematical Engineering, at the University of Concepción, Chile. Previously he has held different research positions such as: Research Engineer at LAAS-CNRS, Toulouse, France; Juan de la Cierva Researcher at the Rey Juan Carlos University, Madrid, Spain; Post-doc Researcher at the ASAP Team, INRIA Rennes-Bretagne Atlantique, France; and Post-doc Researcher at the CEPAGE Team, INRIA Bordeaux-Sud-Ouest, France. He received a PhD degree in Computer Science from the Rey Juan Carlos University and a PhD degree in Applied Mathematics from University of Chile in March 2008. The thesis was done under cotutelage between both University of Chile and Rey Juan Carlos University. His research interests include combinatorics, discrete optimization, and discrete mathematics.