

Asynchronous Rumor Spreading in Preferential Attachment Graphs

Benjamin Doerr¹, Mahmoud Fouz², and Tobias Friedrich^{1,2}

¹ Max-Planck-Institut für Informatik, Saarbrücken, Germany

² Universität des Saarlandes, Saarbrücken, Germany

Abstract. We show that the asynchronous push-pull protocol spreads rumors in preferential attachment graphs (as defined by Barabási and Albert) in time $O(\sqrt{\log n})$ to all but a lower order fraction of the nodes with high probability. This is significantly faster than what synchronized protocols can achieve; an obvious lower bound for these is the average distance, which is known to be $\Theta(\log n / \log \log n)$.

1 Introduction

Online social networks like *Facebook* and *Twitter* are changing the way people communicate, organize and act collectively. They are starting to take the lead over traditional news media in their ability to spread news at a remarkable speed. One striking example was the first picture of US Airways Flight 1549's crash landing on the Hudson River, which became known to a broad audience through Twitter even before TV channels started to report on the accident.

The theoretical model most widely used for social networks is the so-called *preferential attachment* (PA) model, which was introduced in a seminal paper by Barabási and Albert [1]. It builds on the paradigm that new vertices attach to already present vertices with a probability proportional to their degree. Several papers prove that this model indeed enjoys many properties observed in social networks and many other real world networks, e.g., a power law distribution of the vertex degrees, a small diameter and a small average degree [2, 4]. The precise definition of the PA model can be found in Section 2. Note that later extended definitions for PA graphs were given (with the preference not anymore proportional to the degree); in this paper, we shall always refer to the original one.

In this paper, we revisit the *rumor spreading* problem in PA graphs, i.e., the spread of one piece of information in a graph. The classical rumor spreading process is modeled on a discrete time line. A simple protocol assumes that in each time step (or *round*) every node that knows the

rumor forwards it to a randomly chosen neighbor. This is known as the *push strategy*. For many network topologies, this strategy is a very efficient way to spread a rumor. Let n denote the number of vertices of a graph. Then the push model with high probability (i.e., with probability $1 - o(1)$) sends the rumor to all vertices in time $\Theta(\log n)$, if the graph is a complete graph [19], a hypercube [15], an Erdős-Rényi random graph $G_{n,p}$ with $p \geq (1 + \varepsilon) \log(n)/n$ [15], or a random regular graph [17]. In contrast to this, Chierichetti, Lattanzi, and Panconesi [7] showed that the push model with non-vanishing probability needs $\Omega(n^\alpha)$ rounds on PA graphs for some $\alpha > 0$.

Opposite to the push strategy is the *pull strategy*: each vertex in each round contacts a random neighbor and learns the rumor if its contact knows the rumor. There is a symmetry between the two models [6, 10], hence these results also hold for the pull model.

Karp, Schindelhauer, Shenker, and Vöcking [22] pointed out that for complete graphs, the pull strategy is inferior to the push strategy until roughly $n/2$ vertices are informed, and then the pull strategy becomes more effective. This motivates to combine both approaches. In this so-called *push-pull strategy* each vertex contacts another vertex chosen uniformly at random among its neighbors. It *pushes* the rumor in case it has the rumor, and *pulls* the rumor in case the neighbor has the rumor. For complete graphs and many Erdős-Rényi random graphs, this protocol also needs $\Theta(\log n)$ rounds, though with better implicit constants [9, 13, 22]. Its main advantage here is that it allows to define protocols using fewer messages. Chierichetti et al. [6] relate the broadcast time of the push-pull strategy to the conductance of graphs; graphs with conductance Φ have a broadcast time of $O(\log^2(\Phi^{-1}) \Phi^{-1} \log n)$ with high probability. Giakkoupis [20] recently improved this bound to $O(\Phi^{-1} \log n)$ which is tight.

For preferential attachment graphs, however, the push-pull strategy is much better than push or pull alone. Chierichetti et al. [7] showed that with this strategy, $O(\log^2 n)$ rounds suffice with high probability. Recently, we showed that in fact the push-pull strategy succeeds to inform all nodes in $\Theta(\log n)$ rounds [11]. Surprisingly, if the push-pull strategy is slightly modified to prevent that a node contacts the same neighbor twice in a row, then with high probability already $\Theta(\log n / \log \log n)$ rounds suffice [11], which is the diameter of the PA graph.

All these results assume a *synchronized* model, in which all nodes take action simultaneously at discrete time steps. In many applications and certainly in real-world social networks, this assumption is not very

plausible. One can also argue (see, e.g., [5]) that time-synchronization contradicts the idea of a self-organized broadcasting protocol. Boyd et al. [5] therefore proposed an *asynchronous time model* with a *continuous* time line. Each node has its own clock that ticks at the times of a rate 1 Poisson process independent from the clocks of other nodes. The protocol now specifies for every node what to do when its own clock ticks.

The rumor spreading problem in the asynchronous time model has so far received less attention. The push-pull protocol in this model, however, turns out to be closely related to Richardson’s model for the spread of a disease and to first-passage percolation. In this sense, for the hypercube, Fill and Pemantle [16] and Bollobás and Thomason [3] showed that the asynchronous push-pull protocol spreads a rumor to all nodes in time $\Theta(\log n)$. Similarly, for the complete graph, Janson [21] showed a bound of $\Theta(\log n)$. Note that these bounds match the same asymptotics as in the synchronized case. We also suspect that the same bounds hold in case all but $o(n)$ nodes are to be informed.

Fountoulakis, Panagiotou, and Sauerwald [18] have recently studied the push-pull protocol in the asynchronous time model for random graphs in the Ching-Lu model [8] with a given expected degree distribution that follows a power law with exponent in $(2, 3)$. For these graphs, they show a constant runtime to inform $n - o(n)$ nodes. Note that these graphs are quite different from our PA graphs, e.g., their average diameter is known to be $\Theta(\log \log n)$ (see [8]), whereas for PA graphs the average diameter is also $\Theta(\log n / \log \log n)$ (see [12]).

Our results: We study the push-pull protocol in the asynchronous time model on PA graphs and prove that it spreads a rumor in time $O(\sqrt{\log n})$ to $n - o(n)$ nodes in the PA model with high probability. The protocol thus beats the average distance of $\Theta(\log n / \log \log n)$, which is a natural lower bound for the synchronized protocol achieving this aim. To inform all nodes, however, our protocol is shown to need $\Theta(\log n)$ time. This is mainly due to few nodes that require $\Omega(\log n)$ time to contact or be contacted by a neighbor for the first time.

These results show that the asynchronous push-pull protocol behaves quite differently than the synchronized one, despite the fact that each node still contacts one neighbor per time unit on average. The discrepancy between informing all nodes and almost all nodes reflects an often observed ‘long tail’ behavior in real world networks. Such effects are less visible in the synchronized case [11].

2 Precise Model and Preliminaries

Preferential attachment graphs were first introduced by Barabási and Albert [1]. In this work, we follow the formal definition of Bollobás et al. [2, 4]. Let G be an undirected graph. We denote by $\deg_G(v)$ the degree of a vertex v in G .

Definition 1 (Preferential attachment graph). *Let $m \geq 2$ be a fixed parameter. The random graph G_m^n is an undirected graph on the vertex set $V := \{1, \dots, n\}$ inductively defined as follows.*

G_m^1 consists of a single vertex with m self-loops. For all $n > 1$, G_m^n is built from G_m^{n-1} by adding the new node n together with m edges $e_n^1 = \{n, v_1\}, \dots, e_n^m = \{n, v_m\}$ inserted one after the other in this order. Let $G_{m,i-1}^n$ denote the graph right before the edge e_n^i is added. Let $M_i = \sum_{v \in V} \deg_{G_{m,i-1}^n}(v)$ be the sum of the degrees of all the nodes in $G_{m,i-1}^n$. The endpoint v_i is selected randomly such that $v_i = u$ with probability $\deg_{G_{m,i-1}^n}(u)/(M_i + 1)$, except for n that is selected with probability $(\deg_{G_{m,i-1}^n}(n) + 1)/(M_i + 1)$.

This definition implies that when e_n^i is inserted, the vertex v_i is chosen with probability proportional to its degree (except for $v_i = n$). Since many real-world social networks are conjectured to evolve using similar principles, the PA model can serve as a model for social networks. Another property observed in many real-world networks has been formally proven for preferential attachment graphs, namely that the degree distribution follows a power-law [4].

For $m = 1$ the graph is disconnected with high probability; so we focus on the case $m \geq 2$. Under this assumption, Bollobás and Riordan [2] showed that the diameter is only $\Theta(\log(n)/\log \log n)$ with high probability.

With a slight abuse of notation we write $(u, v) \in E$ or $(v, u) \in E$ both to denote $\{u, v\} \in E$. The definition of G_m^n can lead to multiple edges and self-loops, though they typically make up only a vanishing fraction of the edges.

We examine the following broadcasting protocol.

Definition 2 (Asynchronous push-pull strategy). *Each node has a clock that ticks at the times of a rate 1 Poisson process. Whenever the clock of a vertex u ticks, it chooses uniformly at random a neighbor v . If u knows the rumor, it sends the rumor to v (“push”). If v knows the rumor, it sends the rumor to u (“pull”).*

We say that an edge (u, v) *fires*, whenever the clock of node u ticks and u contacts v . We call the time span between two ticks of a clock a *round*. The length of a round is *exponentially* distributed with mean 1. Since the exponential distribution is memoryless, the length of a round is independent over time. The following elementary lemma shows that also the time when a node contacts a *specific* neighbor is exponentially distributed.

Lemma 1. *Let u be a node of degree d that is connected to a node v . Let T denote the time span until u contacts v . Then, $\mathbb{P}[T > x] = e^{-x/d}$.*

3 Statement of Results

Theorem 1. *With probability $1 - o(1)$, the asynchronous push-pull protocol broadcasts a rumor from any node of G_m^n to (i) all but $o(n)$ nodes in time $O(\sqrt{\log n})$, (ii) and to all nodes in time $\Theta(\log n)$.*

The proofs of the upper bounds in Theorem 1 consist of three main steps. In Section 4.3, we analyze the time needed until the rumor reaches a so-called *useful node*. Roughly speaking, a node is useful if its degree is at least polylogarithmic (see Section 4.2 for details). We prove that a useful node is reached in time $O((\log \log n)^2)$ with probability $1 - o(1)$ and in time $O(\log n)$ with probability $1 - o(n^{-2})$. The later bound is used for the case when all nodes are to be informed.

The core of the proof (see Section 4.4) consists of showing that once a useful node u has been informed, within $O(\sqrt{\log n})$ time the rumor is propagated to node 1. To this aim, we show that there is a short path from u to 1 such that every second node has degree exactly m that is traversed in time $O(\sqrt{\log n})$. To prove such a fast traversal we exploit edges that fire fast. In particular, we use the fact that the minimum of k i.i.d. exponential random variables with mean 1 is also exponentially distributed with mean $1/k$.

The result then follows from the following symmetry property.

Lemma 2. *Assume that if the rumor starts in node u , it reaches node v in time t with probability p . This implies the reverse statement: if the rumor is initiated by v , then it reaches u in time t with probability p .*

4 Analysis of the Asynchronous Push-Pull Model

4.1 Alternative model

In the random process generating G_m^n the random decisions made at each step depend heavily on the previous random decisions. Bollobás and Riordan [2] therefore suggested an alternative way of generating G_m^n that is

more accessible. We first describe the model for $m = 1$ and then generalize it to arbitrary m .

Let (x_i, y_i) for $i \in [n] := \{1, 2, \dots, n\}$ be n independently and uniformly chosen pairs from $[0, 1] \times [0, 1]$. With probability one, all these numbers are distinct. By reordering each pair if necessary, we assume that $x_i < y_i$ for every $i \in [n]$. Suppose that after relabeling, $y_1 < y_2 < \dots < y_n$. We set $W_0 := 0$ and $W_i := y_i$ for $i \in [n]$. The graph G_1^n is now defined by having an edge (i, j) if and only if $W_{j-1} < x_i < W_j$. Define $w_j := W_j - W_{j-1}$.

Similarly, for G_m^n , we sample mn pairs $(x_{i,j}, y_{i,j})$ independently and uniformly from $[0, 1] \times [0, 1]$ with $x_{i,j} < y_{i,j}$ for $i \in [n]$ and $j \in [m]$. We relabel the variables such that $y_{i,j}$ is increasing in lexicographic order: $y_{1,1} < y_{1,2} < \dots < y_{1,m} < y_{2,1} < \dots < y_{n,1} < \dots < y_{n,m}$. We set $W_0 := 0$ and $W_i := y_{i,m}$ for $i \in [n]$. The graph is now defined by having an edge (i, j) for each $k \in [m]$ such that $W_{j-1} < x_{i,k} < W_j$. As before, define $w_j = W_j - W_{j-1}$. We write $\ell_{i,k}$ for the node j such that $W_{j-1} < x_{i,k} < W_j$. Note that given $y_{1,1}, \dots, y_{n,m}$, the random variables $x_{1,1}, \dots, x_{n,m}$ are independent with $x_{i,k}$ being chosen uniformly from $[0, y_{i,k}]$. We instead assume that if $y_{1,1}, \dots, y_{n,m}$ are given, then each $x_{i,k}$ is chosen independently and uniformly from $[0, W_i]$. By this slight modification, we can work with the values of the W_i 's and ignore the values of the $y_{i,j}$'s. This modification only increases the probability of a loop at i . It is straightforward to check that each step of our proof remains valid if the probability of a loop is not increased. Thus, the validity of our proof is not affected.

We give a few properties of the alternative model, that are useful in the analysis. Let $s = 2^a$ be the smallest power of 2 larger than $\log^{10} n$, and let 2^b be the largest power of 2 smaller than $2n/3$. Let $I_t = [2^t + 1, 2^{t+1}]$.

Lemma 3 (Bollobás and Riordan [2]). *Let $m \geq 2$ be fixed. Using the definitions above, each of the following five events holds with probability $1 - o(1)$.*

- $E_1 := \{|W_i - \sqrt{i/n}| \leq \frac{1}{10} \sqrt{i/n} \text{ for all } i \in [s, n]\}$
- $E_2 := \{|\{i \in I_t \mid w_i \geq 1/(10\sqrt{in})\}| \geq 2^{t-1} \text{ for all } t \in [a, b]\}$
- $E_3 := \{w_1 \geq \frac{4}{\log(n)\sqrt{n}}\}$
- $E_4 := \{w_i \geq \log^2(n)/n \text{ for all } i < n^{1/5}\}$
- $E_5 := \{w_i < \log^2(n)/n \text{ for all } i \geq n/2\}$.

Note that the event E_5 is slightly adjusted for our purposes. In the original paper, the authors show that for $i \geq n/\log^5 n$; we have $w_i < n^{-4/5}$. It is easy to check that (essentially) the same proof holds for the above version.

Instead of working directly with the alternative model where the W_i 's are random variables, we use the following *typical social network* model where the W_i 's are fixed numbers that satisfy the properties E_1, \dots, E_5 . Since by Lemma 3, these properties hold with high probability, all results proven for a typical social network model carry over to G_m^n with high probability. Let $0 < W_1 < \dots < W_n < 1$ be distinct real numbers and let $w_i = W_i - W_{i-1}$. Assume that W_1, \dots, W_n satisfy the properties E_1, \dots, E_5 . A typical social network $G_m(W_1, \dots, W_n)$ is obtained by connecting each node i with the nodes $\ell_{i,1}, \dots, \ell_{i,m}$, where each $\ell_{i,k}$ is a node chosen randomly with $\mathbb{P}[\ell_{i,k} = j] = w_j/W_i$ for all $j \leq i$.

We always assume to have a typical social network $G := G_m(W_1, \dots, W_n)$.

4.2 Useful nodes

We use the notion of a *useful* node that was introduced by Bollobás and Riordan [2]. A node i is useful if $w_i \geq \log^2(n)/n$. Note that we are slightly relaxing the original definition in [2] where the authors also assumed that $i \leq n/\log^5(n)$. We have by E_5 that $i < n/2$ for all useful nodes. Furthermore by E_4 , every $i < n^{1/5}$ is useful. The following properties of non-useful nodes were proved in [11].

Lemma 4. *With probability $1 - o(1)$, the following event holds*

- $E_6 := \{\deg_G(v) \leq 5m \log^2 n \text{ for all non-useful } v\}$.

Lemma 5. *Assume that E_6 holds. With prob. $1 - n^{-1/5+o(1)}$, we have*

- $E_7 := \{\text{for all non-useful } v, \text{ there exists at most one cycle whose nodes are all connected to } v \text{ via non-useful paths of length at most } \frac{\log n}{(\log \log n)^2}\}$.

Lemma 6 (Bollobás and Riordan [2]). *Let v be a fixed non-useful node. Then for all $k \in [m]$, the prob. that $\ell_{v,k}$ is a useful node is at least $\log^{-3} n$. This event is independent from all other random decisions $\ell_{v',k'}$ with $(v', k') \neq (v, k)$.*

Note that in the original lemma, the authors only state a bound on the probability that $\ell_{v,1}$ is a useful node. However, the same proof yields the above version. Also, Lemma 6 remains valid if we condition on E_6 .

4.3 Informing the first useful node

Let $G = G_m(W_1, \dots, W_n)$ be a typical social network. Assume that also E_6 and E_7 hold. In this section, all probabilities are taken over the product space of the random graph G and the random decisions of the rumor spreading process.

Lemma 7. *Let u be a fixed node. The rumor initiated by u reaches a useful node in time $O((\log \log n)^2)$ with probability $1 - o(1)$, and in time $O(\log n)$ with probability $1 - o(n^{-2})$.*

4.4 Informing node 1

Similar to the synchronized case, we use constant degree nodes to establish *fast links* between large degree nodes. More precisely, once a neighbor of a constant degree node is informed, the time until it has pulled the rumor from this neighbor and pushed it to one specific neighbor is (essentially) exponentially distributed. Thus, independent of their own degrees, two nodes that are connected via a third node of constant degree exchange information in time exponentially distributed.

Starting from one informed useful node, we study how fast the rumor spreads to the surrounding ‘neighborhoods’ of nodes. We consider neighborhoods alternating between small nodes and *good* nodes i of relatively large weight w_i . The small nodes act as fast links between the levels of good nodes that ensure a large expansion. In particular, we make use of the fact that a good node i has a high degree and since every small neighbor of i independently pulls the rumor in time exponentially distributed, we can argue that a considerable fraction of the small neighbors of i will be informed very fast. The more neighbors of informed nodes there are, the faster the rumor will spread to sufficiently many neighbors that form the next level of informed node. In contrast, in the synchronized case, it would always take at least one time step for a neighbor to pull the rumor.

We consider informed neighborhoods at suitably chosen time steps on the continuous time line. The smaller these steps are chosen, the smaller the achieved expansion factor is at each step. On the other hand, smaller time steps allow us to progress faster through the different neighborhood levels. By carefully choosing each step size, we can balance out these opposing effects in order to achieve the following runtime.

Theorem 2. *Let W_1, \dots, W_n be such that E_1, \dots, E_5 are satisfied. Let G be a random graph from $G_m(W_1, \dots, W_n)$. Let $v \in [n]$ be a useful node. With probability $1 - o(n^{-1})$, using the asynchronous push-pull protocol, a rumor present at v reaches node 1 in $O(\sqrt{\log n})$ steps.*

For our argument using fast links, we will need many nodes of constant degree. The following simple lemma proves that there is a linear number of nodes $i \in [\frac{2}{3}n, n]$ that have a degree equal to m . We call such nodes *small*. If not explicitly stated, all probabilities in this section are taken over the random graph $G_m(W_1, \dots, W_n)$, where W_1, \dots, W_n are given numbers that satisfy properties E_1, \dots, E_5 .

Lemma 8. *Let $\varepsilon_m := \frac{1}{8}e^{-3m}$. With probability $1 - e^{-\Omega(n)}$, there are at least $\varepsilon_m n$ small nodes in $[\frac{2}{3}n, n]$.*

Crucial for a large expansion in each step are *good* nodes of large weight. We say a node i is *good* if

$$i \in [s + 1, 2^b] \text{ and } w_i \geq 1/(10\sqrt{in}), \quad (1)$$

where, as before, $s = 2^a$ is the smallest power of 2 larger than $\log^{10} n$ and 2^b is the largest power of 2 smaller than $\frac{2}{3}n$. Let u be a useful node. Let $t_0 < t'_0 < t_1 < t'_1 < \dots$ denote discrete time steps to be specified later. We consider neighborhoods of u that are informed in these time intervals. In particular, we define sets Γ_k and Γ'_k recursively as follows. We set $\Gamma_0 = \{u\}$. Given the set Γ_k , Γ'_k consists of all *small* nodes $i \geq \frac{2}{3}n$ that contact a neighbor in Γ_k in time $[t_k, t'_k]$ and have not been included in any Γ'_ℓ with $\ell \leq k - 1$. Similarly, Γ_k is defined as the set of all *good* nodes that are contacted by a neighbor in Γ'_{k-1} in time $[t'_{k-1}, t_k]$ and have not been included in any Γ_ℓ with $\ell \leq k - 1$. Note that for all $k \geq 0$, Γ_k only contains nodes $i < \frac{2}{3}n$, while Γ'_k only contains nodes $i \geq \frac{2}{3}n$. This is true for Γ_0 since u is useful and by E_5 , all useful nodes are smaller than $n/2$. We define the *weight* of a set Γ_k by

$$f_k := \begin{cases} w_u & \text{if } k = 0 \\ \sum_{i \in \Gamma_k} \frac{1}{\sqrt{in}} & \text{if } k \geq 1. \end{cases} \quad (2)$$

Since for $k \geq 1$, Γ_k only contains good nodes, and by definition, $w_u = f_0$, we have for $k \geq 0$,

$$\sum_{i \in \Gamma_k} w_i \geq f_k/10. \quad (3)$$

We denote by $N_k = \Gamma_0 \cup \Gamma_1 \cup \dots \cup \Gamma_k$ (note that the Γ'_i are not included). Let $C_0 \subseteq [\frac{2}{3}n, n]$ be the set of small nodes and for $k \geq 1$, let $C_k = C_0 \setminus \{\Gamma'_0, \dots, \Gamma'_{k-1}\}$ be the set of small nodes excluding nodes in $\Gamma'_0, \Gamma'_1, \dots, \Gamma'_{k-1}$. By Lemma 8, we have $C_0 \geq \varepsilon_m n$ with probability $1 - e^{-\Omega(n)}$.

The next lemma shows that we achieve an exponential expansion in terms of f_k in each level as long as there is still a linear number of small nodes in C_k and similarly, as long as for each interval $I_t := [2^t + 1, 2^{t+1}]$, where $t \in [a, b)$, there are still 2^{t-2} good nodes that are not N_k .

Lemma 9. *Let $c > 0$ be a sufficiently large constant, $k \geq 0$ be such that $\log^4(n)/\sqrt{n} \geq f_k \geq \log^2(n)/n$ and $|C_k| \geq \varepsilon_m n/2$. Let $\Delta_k = c \log_2^{-1/2}(\varepsilon_m f_k n / \log^2 n)$. Set $t'_k := t_k + \Delta_k$ and $t_{k+1} := t'_k + \Delta_k$. Then given C_k and $\Gamma_0, \Gamma'_0, \Gamma_1, \Gamma'_1, \dots, \Gamma_k$, with prob. $1 - O(n^{-6/5})$, one of the following is satisfied. (i) $|N_{k+1} \cap I_t| \geq 2^{t-2}$, for some $t \in [a, b)$, or (ii) $f_{k+1} \geq 2f_k$.*

5 Conclusion

We have shown that for PA graphs the asynchronous push-pull protocol informs almost all nodes in $O(\sqrt{\log n})$ time. This shows, in an even stronger way than the previous $\tilde{O}(\log n)$ bounds for the synchronized protocol [11], that randomized rumor spreading is very effective in network topologies resembling real-world networks.

From a broader perspective, our result also indicates that in naturally asynchronous settings, it might be a misleading oversimplification to assume a synchronized protocol.

Bibliography

- [1] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.
- [2] B. Bollobás and O. Riordan. The diameter of a scale-free random graph. *Combinatorica*, 24:5–34, 2004.
- [3] B. Bollobás and A. Thomason. *On Richardson’s Model on the Hypercube*. Cambridge University Press, 1997.
- [4] B. Bollobás, O. Riordan, J. Spencer, and G. Tusnády. The degree sequence of a scale-free random graph process. *Random Structures & Algorithms*, 18:279–290, 2001.
- [5] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 52:2508–2530, 2006.
- [6] F. Chierichetti, S. Lattanzi, and A. Panconesi. Almost tight bounds for rumour spreading with conductance. In *42nd ACM Symposium on Theory of Computing (STOC)*, pages 399–408, 2010.

- [7] F. Chierichetti, S. Lattanzi, and A. Panconesi. Rumor spreading in social networks. *Theoretical Computer Science*, 412:2602–2610, 2011.
- [8] F. R. K. Chung and L. Lu. The average distance in a random graph with given expected degrees. *Internet Mathematics*, 1:91–113, 2003.
- [9] A. J. Demers, D. H. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. E. Sturgis, D. C. Swinehart, and D. B. Terry. Epidemic algorithms for replicated database maintenance. *Operating Systems Review*, 22:8–32, 1988.
- [10] B. Doerr, T. Friedrich, and T. Sauerwald. Quasirandom rumor spreading: Expanders, push vs. pull, and robustness. In *36th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 366–377, 2009.
- [11] B. Doerr, M. Fouz, and T. Friedrich. Social networks spread rumors in sublogarithmic time. In *43rd ACM Symposium on Theory of Computing (STOC)*, pages 21–30, 2011.
- [12] S. Dommers, R. van der Hofstad, and G. Hooghiemstray. Diameters in preferential attachment models. *Journal of Statistical Physics*, 139:72–107, 2010.
- [13] R. Elsässer. On the communication complexity of randomized broadcasting in random-like graphs. In *18th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 148–157, 2006.
- [14] M. Evans, N. Hastings, and B. Peacock. *Statistical Distributions*. John Wiley & Sons, Inc., 3rd edition, 2000.
- [15] U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized broadcast in networks. *Random Structures & Algorithms*, 1:447–460, 1990.
- [16] J. A. Fill and R. Pemantle. Percolation, first-passage percolation, and covering times for Richardson’s model on the n -cube. *Annals of Applied Probability*, 3:593–629, 1993.
- [17] N. Fountoulakis and K. Panagiotou. Rumor spreading on random regular graphs and expanders. In *14th International Workshop on Randomization and Computation (RANDOM)*, pages 560–573, 2010.
- [18] N. Fountoulakis, K. Panagiotou, and T. Sauerwald. Ultra-fast rumor spreading in social networks. In *23rd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1642–1660, 2012.
- [19] A. M. Frieze and G. R. Grimmett. The shortest-path problem for graphs with random arc-lengths. *Discrete Applied Mathematics*, 10: 57–77, 1985.
- [20] G. Giakkoupis. Tight bounds for rumor spreading in graphs of a given conductance. In *28th International Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 57–68, 2011.

- [21] S. Janson. One, two and three times $\log n/n$ for paths in a complete graph with random weights. *Combinatorics, Probability & Computing*, 8:347–361, 1999.
- [22] R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized rumor spreading. In *41st IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 565–574, 2000.