The String of Diamonds Is Tight for Rumor Spreading*

Omer Angel[1], Abbas Mehrabian[2], and Yuval Peres[3]

1 **Department of Mathematics, University of British Columbia, Vancouver, BC, Canada**
`angel@math.ubc.ca`

2 **Department of Computer Science, University of British Columbia, Vancouver, BC, Canada**
`abbasmehrabian@gmail.com`

3 **Microsoft Research, Redmond, WA, USA**
`peres@microsoft.com`

Abstract

For a rumor spreading protocol, the spread time is defined as the first time that everyone learns the rumor. We compare the synchronous push&pull rumor spreading protocol with its asynchronous variant, and show that for any n-vertex graph and any starting vertex, the ratio between their expected spread times is bounded by $O\left(n^{1/3}\log^{2/3} n\right)$. This improves the $O(\sqrt{n})$ upper bound of Giakkoupis, Nazari, and Woelfel (in Proceedings of ACM Symposium on Principles of Distributed Computing, 2016). Our bound is tight up to a factor of $O(\log n)$, as illustrated by the string of diamonds graph.

1998 ACM Subject Classification C.2.1 Network Architecture and Design, G.2.2 Graph Theory, G.3 Probability and Statistics

Keywords and phrases randomized rumor spreading, push&pull protocol, asynchronous time model, string of diamonds

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2017.26

1 Introduction

Randomized rumor spreading is an important paradigm for information dissemination in networks with numerous applications in network science, ranging from spreading information in the WWW and Twitter to spreading viruses and diffusion of ideas in human communities. A well studied rumor spreading protocol is the *(synchronous) push&pull protocol*, introduced by Demers, Greene, Hauser, Irish, Larson, Shenker, Sturgis, Swinehart, and Terry [4] and popularized by Karp, Schindelhauer, Shenker, and Vöcking [11].

▶ **Definition 1** (Synchronous push&pull protocol). Suppose that one node s in a network G is aware of a piece of information, the 'rumor', and wants to spread it to all nodes

* This work started in the *Random Geometric Graphs and Their Applications to Complex Networks* workshop held in the Banff International Research Station in November 2016. The authors thank the workshop organizers and BIRS for making it happen. The first author is supported by NSERC. The second author was supported by an NSERC Postdoctoral Fellowship and a Simons-Berkeley Research Fellowship. Part of this work was done while he was visiting the Simons Institute for the Theory of Computing at UC Berkeley. This work was completed at Microsoft Research in Redmond. The authors thank Microsoft for their support.

© Omer Angel, Abbas Mehrabian, and Yuval Peres;
licensed under Creative Commons License CC-BY

quickly. The protocol proceeds in rounds; in each round $1, 2, \dots$, all vertices perform actions simultaneously. That is, each vertex x calls a random neighbor y, and the two share any information they may have: If x knows the rumor and y does not, then x tells y the rumor (a *push* operation), and if x does not know the rumor and y knows it, y tells x the rumor (a *pull* operation). Note that this is a synchronous protocol, e.g. a vertex that receives a rumor in a certain round cannot send it on in the same round. The synchronous spread time of G, denoted by $S(G, s)$, is the first time that everyone knows the rumor. Note that this is a discrete random variable.

A point to point communication network can be modeled as an undirected graph: the nodes represent the processors and the links represent communication channels between them. Studying rumor spreading has several applications to distributed computing in such networks, of which we mention just two (see [7] also). The first is in broadcasting algorithms: a single processor wants to broadcast a piece of information to all other processors in the network. The push&pull protocol has several advantages over other protocols: it puts less load on the edges than the naive flooding protocol; it is simple and naturally distributed (each node makes a simple local decision in each round; no knowledge of the global state or topology is needed; no internal states are maintained); it is scalable (the protocol is independent of the size of network: it does not grow more complex as the network grows) and it is robust (the protocol tolerates random node/link failures without the need for error recovery mechanisms).

A second application comes from the maintenance of databases replicated at many sites, e.g., yellow pages, name servers, or server directories. Updates to the database may be injected at various nodes, and these updates must propagate to all nodes in the network. In each round, a processor communicates with a random neighbor and they share any new information, so that eventually all copies of the database converge to the same contents. See [4] for details.

The above protocol assumed a synchronized model, i.e. all nodes take action simultaneously at discrete time steps. In many applications and certainly for modeling information diffusion in social networks, this assumption is not realistic. Boyd, Ghosh, Prabhakar, Shah [3] proposed an asynchronous time model with a continuous time line. This too is a randomized distributed algorithm for spreading a rumor in a graph, defined as follows. An *exponential clock* with rate λ is a clock that, once turned on, rings at times of a Poisson process with rate λ.

▶ **Definition 2** (Asynchronous push&pull protocol)**.** Given a graph G, independent exponential clocks of rate 1 are associated with the vertices of G, one to each vertex. Initially, one vertex s of G knows the rumor, and we turn on all clocks. Whenever the clock of a vertex x rings, it calls a random neighbor y: if x knows the rumor and y does not, then x tells y the rumor (a push operation); if x does not know the rumor and y knows it, y tells x the rumor (a pull operation). The asynchronous spread time of G, denoted by $A(G, s)$, is the first time that everyone knows the rumor.

Rumor spreading protocols in this model turn out to be closely related to Richardson's model for the spread of a disease [12, 6]. For a single rumor, the push&pull protocol is almost equivalent to the first passage percolation model introduced by Hammersley and Welsh [9] with edges having independent exponential weights (see also the survey [2]). The difference between the push&pull model and first passage percolation stems from the fact that in the rumor spreading models each vertex contacts one neighbor at a time, and so the rate at which x pushes the rumor to y is inversely proportional to the degree of x. A rumor

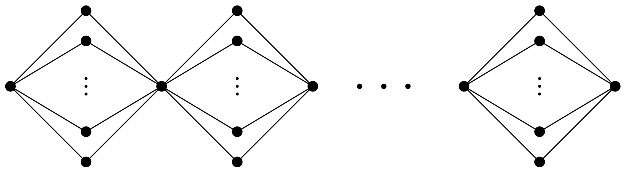

Figure 1 The string of diamonds graph.

can also be pulled from x to y at rate determined by the degree of y. On regular graphs, the asynchronous push&pull protocol, Richardson's model, and first passage percolation are essentially the same process, assuming appropriate parameters are chosen. For general graphs, the equivalence is to first passage percolation with exponential edge weights that are independent, but have different means. Hence, the degrees of vertices play a different role here than they do in Richardson's model or first passage percolation. A collection of known bounds for the average spread times of many graph classes is given in [1, Table 1].

Doerr, Fouz, and Friedrich [5] experimentally compared the spread time in the two time models. They state that 'Our experiments show that the asynchronous model is faster on all graph classes [considered here].' The first general relationship between the spread times of the two variants was given in [1], where it was proved using a coupling argument that

$$\frac{\mathbb{E}\left[S(G,s)\right]}{\mathbb{E}\left[A(G,s)\right]} = \widetilde{O}\left(n^{2/3}\right).$$

Here $\widetilde{O}$ (and $\widetilde{\Omega}$ below) allow for poly-logarithmic factors. Building on the ideas of [1] and using more involved couplings, Giakkoupis, Nazari and Woelfel [8] improved this bound to $O\left(n^{1/2}\right)$. In this note we improve the bound to $\widetilde{O}(n^{1/3})$. A graph was given in [1] with

$$\frac{\mathbb{E}\left[S(G,s)\right]}{\mathbb{E}\left[A(G,s)\right]} = \widetilde{\Omega}\left(n^{1/3}\right),$$

known as the *string of diamonds* (see Figure 1), which shows the exponent 1/3 is optimal.

We use a rather different coupling than previous ones. Our coupling is motivated by viewing rumor spreading as a special case of first passage percolation. This novel approach involves carefully intertwined Poisson processes. Our proof also yields a natural interpretation for the exponent 1/3: using non-trivial counting arguments, we prove that the longest distance the rumor can traverse during a unit time interval in the asynchronous protocol is $O(n^{1/3})$ (see the proof of Lemma 6), which is tight.

Regarding lower bounds, it is proved in [8] that $\mathbb{E}\left[A(G,s)\right] \le \mathbb{E}\left[S(G,s)\right] + O(\log n)$. We will use the following bounds that hold for all G and s (see [1, Theorem 1.3]):

$$\log n/5 \le \mathbb{E}\left[A(G,s)\right] \le 4n\,.$$

In this paper n always denotes the number of vertices of the graph, and all logarithms are in natural base.

1.1 Our results

For an n-vertex graph G and a starting vertex s, recall $A(G,s)$ and $S(G,s)$ denote the asynchronous and synchronous spread times, respectively. Our main theorem is the following.

▶ **Theorem 3.** *Given any $K > 0$, there is a $C > 0$ such that for any (G, s) and any $t \geq 1$ we have*

$$\mathbb{P}\Big[S(G,s) > C(t + t^{2/3}n^{1/3}\log n)\Big] \leq \mathbb{P}\Big[A(G,s) > t\Big] + Cn^{-K}.$$

▶ **Corollary 4.** *For any (G, s), we have $\mathbb{E}\left[S(G,s)\right] = O\left(\mathbb{E}\left[A(G,s)\right]^{2/3} n^{1/3}\log n\right)$.*

Proof. Apply Theorem 3 with $K = 1$ and $t = 3\mathbb{E}\left[A(G,s)\right] \leq 12n$. By Markov's inequality, $\mathbb{P}\left[S(G,s) > C(t + t^{2/3}n^{1/3}\log n)\right] \leq 1/3 + C/n \leq 1/2$ for n large enough. Since $t = O(n)$, this implies the median of $S(G,s)$, denoted by M, is $O(t^{2/3}n^{1/3}\log n)$. To complete the proof we need only show that $\mathbb{E}\left[S(G,s)\right] = O(M)$.

Consider the protocol which is the same as synchronous push&pull except that, if the rumor has not spread to all vertices by time M, then the new process reinitializes. Coupling the new process with push&pull, we obtain for any $i \in \{0, 1, 2, \dots\}$ that $\mathbb{P}\left[S(G,s) > iM\right] \leq 2^{-i}$. Thus,

$$\mathbb{E}\left[S(G,s)\right] = \sum_{i=0}^{\infty}\mathbb{P}\left[S(G,s) > i\right] \leq \sum_{i=0}^{\infty} M \times \mathbb{P}\left[S(G,s) > iM\right] \leq M \times \sum_{i=0}^{\infty} 2^{-i} = 2M. \qquad ◀$$

Since for all G and s, $\mathbb{E}\left[A(G,s)\right] = \Omega(\log n)$, we also obtain:

▶ **Corollary 5.** *For any (G, s) we have*

$$\frac{\mathbb{E}\left[S(G,s)\right]}{\mathbb{E}\left[A(G,s)\right]} = O\left(n^{1/3}\log^{2/3} n\right).$$

This corollary is tight up to logarithmic factors. Indeed, let G be the string of diamonds (see Figure 1) with m diamonds, each consisting of k paths of length 2, and let s be an end vertex of it. Then, $S(G,s) \geq 2m$ deterministically and $\mathbb{E}\left[A(G,s)\right] = O(\log n + m/\sqrt{k})$ (see [1] for the proof). If we let $m = \Theta(n^{1/3}(\log n)^{2/3})$ and $k = \Theta((n/\log n)^{2/3})$, we obtain a graph with

$$\frac{\mathbb{E}\left[S(G,s)\right]}{\mathbb{E}\left[A(G,s)\right]} = \Omega\left(n/\log n\right)^{1/3},$$

which means Corollary 5 is tight up to an $O(\log n)$ factor.

2 Proof of Theorem 3

For the rest of the paper we fix the graph G and the starting vertex s. Let $\Gamma(s,v)$ be the set of all simple paths in G from s to v. For a path γ, let $E(\gamma)$ be its set of edges and $|\gamma| \coloneqq |E(\gamma)|$ denote its length. Let $\deg(u)$ denote the degree of a vertex u.

For any ordered pair (u, v) of adjacent vertices, let $Y_{u,v}$ be an exponential random variable with rate $1/\deg(u)$. Assume these random variables are mutually independent. In the asynchronous protocol, since each vertex u calls any adjacent v at a rate of $1/\deg(u)$, we can write:

$$A \coloneqq A(G,s) = \max_{v\in V}\min_{\gamma\in\Gamma(s,v)}\sum_{xy\in E(\Gamma)}\min\{Y_{x,y}, Y_{y,x}\}. \tag{1}$$

Here $Y_{x,y}$ is the time it takes after one of x, y learns the rumor before x calls y.

For any positive integer L, consider the restriction to short paths

$$A_L \coloneqq \max_{v\in V} \min_{\substack{\gamma\in\Gamma(s,v)\\ |\gamma|\le L}} \sum_{xy\in E(\Gamma)} \min\{Y_{x,y}, Y_{y,x}\}.$$

For any L we have $A_L \ge A$. To bound A from below, we have the following "with high probability" stochastic domination result.

▶ **Lemma 6.** *There exists a $C>0$ such that for any $t\ge 1$ and $L \ge Ct^{2/3}n^{1/3}$ we have*

$$\mathbb{P}\left[A_L > t\right] \le \mathbb{P}\left[A > t\right] + e^{-L}\,.$$

Proof. We show that, in the asynchronous protocol, with probability $1-e^{-L}$, during the interval $[0,t]$, the rumor does not travel along any simple path of length L. We prove this by taking a union bound over all paths of length L. As there is no simple path of length n or more, we will assume $L<n$.

Consider a path γ with vertices $\gamma_0,\gamma_1,\dots,\gamma_L$. In order for the rumor to travel along γ, it is necessary that there are calls along the edges of γ in order, at some sequence of times $0\le t_1<\cdots<t_L\le t$. Since along each edge the rumor can travel via a push or a pull, the rate of calls along an edge xy is $1/\deg(x)+1/\deg(y)$. Since the volume of the L-dimensional simplex of possible sequences (t_i) is $t^L/L!$, the probability of such a sequence of calls along the path γ is at most

$$\frac{t^L}{L!}\prod_{i=1}^{L}\left(\frac{1}{\deg(\gamma_{i-1})}+\frac{1}{\deg(\gamma_i)}\right) \le \left(\frac{2et}{L}\right)^L \prod_{i=1}^{L}\frac{1}{\min(\deg(\gamma_{i-1}),\deg(\gamma_i))}. \tag{2}$$

In light of this, define

$$Q(\gamma)\coloneqq \prod_{i=1}^{|\gamma|}\frac{1}{\min(\deg(\gamma_{i-1}),\deg(\gamma_i))}.$$

Our objective is therefore a bound for $\sum_{|\gamma|=L} Q(\gamma)$.

For a path γ of length L, consider the sequence of degrees $(\deg(\gamma_i))_{i=0}^{L}$. We say the sequence has a *local minimum* at i if $\deg(\gamma_{i-1})>\deg(\gamma_i)\le\deg(\gamma_{i+1})$, and a *local maximum* at i if $\deg(\gamma_{i-1})\le\deg(\gamma_i)>\deg(\gamma_{i+1})$. In both of these definitions we use the convention that inequalities involving γ_{-1} or γ_{L+1} always hold. The edge set of γ can be partitioned into *segments* starting and ending at local maxima. For example, suppose $L=7$ and the degree sequence is

$$(\deg(\gamma_0),\deg(\gamma_1),\deg(\gamma_2),\deg(\gamma_3),\deg(\gamma_4),\deg(\gamma_5),\deg(\gamma_6),\deg(\gamma_7)) = (\mathbf{5},5,7,\mathbf{3},4,4,\mathbf{2},5).$$

Then the segments are $(\boldsymbol{\gamma_0},\gamma_1,\gamma_2)$, $(\gamma_2,\boldsymbol{\gamma_3},\gamma_4,\gamma_5)$, and $(\gamma_5,\boldsymbol{\gamma_6},\gamma_7)$. Thus, in each segment the degrees strictly decrease to a local minimum (bolded in the example), then weakly increase up to the local maximum at the end of the segment. (The first and last segments are special in that the local minimum could be at the beginning and end of the segment, respectively.) Henceforth, we use the term *segment* for a path with this property.

Each path gives rise to an ordered sequence of segments. Denote the segments of γ by $\sigma_1,\dots,\sigma_s$, and note that $s\le L/2+1$, since each segment except possibly the first and the last one contains at least two edges. The next observation is that we have $Q(\gamma)=\prod Q(\sigma_i)$; that is, the Q value of a path equals the product of Q values of its segments (this is true for any partition of a path into sub-paths). Note that not every sequence of segments can arise in this way: each segment must start at the last vertex of the previous segment. Since we are interested only in simple paths, the segments are otherwise disjoint. Thus for a collection of segments there is at most one order in which it could arise. Therefore,

$$\sum_{|\gamma|=L} Q(\gamma) \leq \sum_{s=1}^{L/2+1} \sum_{|\sigma_1|+\cdots+|\sigma_s|=L} \frac{1}{s!} \prod_{i=1}^{s} Q(\sigma_i)\,, \tag{3}$$

where the last sum is over s-tuples of segments whose lengths add up to L, but *without* the condition that they form a path (that is why we have an inequality rather than an equality).

We now bound the right-hand-side of (3). We say a segment has *type* $(x, \ell^-, \ell^+) \in V(G) \times \mathbb{Z} \times \mathbb{Z}$ if the local minimum is at a vertex x (called the *center* of the segment), and the segment has ℓ^- edges before x and ℓ^+ edges after x. (The example path above had $s = 3$ segments, of types $(\gamma_0, 0, 2)$, $(\gamma_3, 1, 2)$, and $(\gamma_6, 1, 1)$ respectively.) For a segment σ, let $\pi(\sigma)$ denote its type, and let $\mathcal{T}$ denote the set of all possible types.

For bounding the right-hand-side of (3), we first fix s and bound the number of options for the sequence $(\pi(\sigma_1), \ldots, \pi(\sigma_s))$. There are $n!/(n-s)!$ choices for the centers (the number of ways to choose s ordered distinct vertices), and at most 2^L choices for the lengths $\ell^\pm$ (the number of ways to write L as an ordered sum of natural numbers). Thus there are at most $2^L n!/(n-s)!$ options for $(\pi(\sigma_1), \ldots, \pi(\sigma_s))$. Enumerate these s-vectors of types by $T_1, \ldots, T_m \in \mathcal{T}^s$ with $m \leq 2^L n!/(n-s)!$, and let $T_{j,k}$ denote the kth component of T_j, i.e. the type specified for σ_k in T_j. Thus,

$$\begin{aligned} \sum_{|\sigma_1|+\cdots+|\sigma_s|=L} \prod_{i=1}^{s} Q(\sigma_i) &= \sum_{j=1}^{m} \sum_{(\pi(\sigma_1),\ldots,\pi(\sigma_s))=T_j} \prod_{i=1}^{s} Q(\sigma_i) \\ &\leq \sum_{j=1}^{m} \prod_{k=1}^{s} \left(\sum_{\pi(\sigma_k)=T_{j,k}} Q(\sigma_k) \right) \end{aligned}$$

Next, we claim that each of the brackets, which is the sum of Q values of segments of a given type can be bounded by 1. Fix some type (x, ℓ^-, ℓ^+), and let $\ell = \ell^- + \ell^+$. Then the constraints on the degrees along a segment $\sigma = v_0, v_1, \cdots, v_{\ell^-}, \cdots, v_\ell$ of this type imply $x = v_{\ell^-}$ and

$$Q(\sigma) = \prod_{i=1}^{\ell^-} \frac{1}{\deg(v_i)} \prod_{i=\ell^-}^{\ell-1} \frac{1}{\deg(v_i)}.$$

If we sum this up over all *walks* of length $\ell^- + \ell^+$ whose ℓ^-th vertex is x, we get 1 (since the number of choices for the neighbors cancel out the degree reciprocals). Restricting to simple paths with piecewise monotone degrees only decreases this. Thus we obtain

$$\sum_{|\sigma_1|+\cdots+|\sigma_s|=L} \prod_{i=1}^{s} Q(\sigma_i) \leq m \times 1 \leq 2^L n!/(n-s)!\,.$$

Plugging this back into (3) yields

$$\sum_{|\gamma|=L} Q(\gamma) \leq \sum_{s=1}^{L/2+1} 2^L n!/(n-s)!s! = 2^L \sum_{s=1}^{L/2+1} \binom{n}{s} \leq 2^L (2en/L)^{L/2+1},$$

where we used the standard bound $\sum_{i=0}^{k} \binom{n}{i} \leq (en/k)^k$ valid for all $0 \leq k \leq n$.

Therefore, by (2), the probability that the rumor travels along some path of length L is bounded by

$$\sum_{|\gamma|=L}\left(\frac{2et}{L}\right)^L Q(\gamma) \le \left(\frac{2et}{L}\right)^L 2^L(2en/L)^{L/2+1} \le c_1 n(c_2 nt^2/L^3)^{L/2},$$

which is at most e^{-L} for $L \ge Ct^{2/3}n^{1/3}$, completing the proof. ◀

In (1) we wrote $A(G,s)$ in a max-min form. We would like to write $S(G,s)$ in a similar way. To achieve this, let $q_{uv} = q_{vu}$ be the first (discrete) round at which one of u or v is informed. Suppose the first round *strictly after* q_{uv} that u calls v is F_{uv}. Then define $T_{u,v} = F_{uv} - q_{uv}$. Note that $T_{u,v}$ is a positive integer, and it is a geometric random variable: $\mathbb{P}[T_{u,v} \ge k] = (1 - 1/\deg(u))^{k-1}$ for any $k = 1, 2, \dots$. Moreover, observe that, both u and v are informed by round $q_{uv} + \min\{T_{u,v}, T_{v,u}\}$ hence, we have

$$S \coloneqq S(G,s) \le \max_{v\in V}\min_{\gamma\in\Gamma(s,v)}\sum_{xy\in E(\Gamma)}\min\{T_{x,y}, T_{y,x}\}. \tag{4}$$

Now we have a max-min expression for $S(G,s)$. However, the major trouble here is that the $\{T_{x,y}\}$ are not independent. We will stochastically dominate them by another collection $\{X_{x,y}\}$ of random variables, which are independent. To prove their independence, we first define the synchronous protocol in an equivalent but more convenient way.

Consider for each ordered pair $u \sim v$ a pair of exponential clocks $Z_{u,v}, Z'_{u,v}$, both with rate $1/\deg(u)$. All these clocks are independent. We say the clocks $Z_{u,v}, Z'_{u,v}$ are *located* at vertex u. Initially, the clocks $Z_{u,v}$ are turned on, and the clocks $Z'_{u,v}$ are off. For each round $1, 2, \dots$, we visit the vertices one by one. For each vertex u, we wait for the next clock at u to ring. If that ring comes from clock $Z_{u,v}$ or $Z'_{u,v}$, we say that u calls v in that round. Once the choice of calls at every vertex has been made, we use these to perform the push&pull operations in a round of the protocol. (Note that the time of the clocks is separate from the discrete rounds of the synchronous protocol: in each vertex, a different amount of time is elapsed on the clocks.) Moreover, whenever a vertex u gets informed of the rumor, for each adjacent v we turn off the clocks $Z_{u,v}$ and $Z_{v,u}$, and turn on $Z'_{u,v}$ and $Z'_{v,u}$. (If v was already informed, these status changes had already taken place.) Observe that, because of memorylessness of the exponential distribution, for each vertex u, this process generates a random sequence of independent uniform neighbors, so it is equivalent to the synchronous protocol.

Now let us see what are the random variables $T_{u,v}$ in this setup. For each ordered pair u, v, observe that the collection of ringing times of clocks $Z_{u,v}, Z'_{u,v}$ forms a Poisson process $P_{u,v}$ with rate $1/\deg(u)$. (It does not matter that the initial rings come from Z and subsequent rings from Z'.) Let

$$P_u \coloneqq \bigcup_{v\sim u} P_{u,v},$$

and note that P_u is a Poisson process with rate 1.

For a pair u, v, suppose the q_{uv}th point in P_u is at α, and suppose the first point of $P_{u,v}$ strictly larger than α is at β. Then, $T_{u,v}$ is precisely the number of points of P_u in the interval $(\alpha, \beta]$. Define $X_{u,v} = \beta - \alpha$. By construction, $X_{u,v}$ is the first time that clock $Z'_{u,v}$ rung from the time it was turned on, hence it is exponential with rate $1/\deg(u)$. Since these clocks are independent, the random variables $X_{u,v}$ are also independent. (Only the times at which the clocks are turned on depend on other clocks in a non-trivial manner.) Thus we have proven:

▶ **Lemma 7.** *The random variables $\{X_{x,y}\}$ defined above are mutually independent.*

On the other hand, we can use these to control the $T_{x,y}$:

▶ **Lemma 8.** *For any fixed K there is a fixed C such that with probability at least $1 - n^{-K}$, for all adjacent pairs x, y we have $T_{x,y} \le C \log n + CX_{x,y}$.*

Proof. We show that for any adjacent pair x, y, we have $T_{x,y} > C \log n + CX_{x,y}$ with probability at most n^{-K-2}, and then apply the union bound over all edges.

Observe that, conditioned on $X_{x,y} = t$, the random variable $T_{x,y} - 1$ is Poisson with rate $t \times (\deg(u) - 1)/\deg(u) \le t$. We will use a standard tail inequality for the Poisson distribution, which follows from Theorem 5.1(iii) in [10]. Let $\mathrm{Po}(t)$ denote a Poisson random variable with mean $t > 0$. Then for any $\alpha \ge 1$ we have $\mathbb{P}[\mathrm{Po}(t) \ge \alpha t] \le (e^{\alpha-1}\alpha^{-\alpha})^t$. This gives

$$\begin{aligned}\mathbb{P}[T_{x,y} - 1 > C \log n + CX_{x,y} | X_{x,y} = t] &\le \mathbb{P}[\mathrm{Po}(t) > C \log n + Ct] \\ &\le (e/C)^{C \log n} \le n^{-K-2}\end{aligned}$$

for $C \ge \max(e^2, K + 2)$. ◀

Our main result now follows easily from our lemmas.

Proof of Theorem 3. Given K, pick C sufficiently large so that Lemmas 6 and 8 hold. Fix $t \ge 1$ and let $L = Ct^{2/3}n^{1/3}$. We have

$$\begin{aligned}&\mathbb{P}[S > Ct + CL \log n] \\ &\le \mathbb{P}\left[\left(\max_{v\in V} \min_{\gamma\in\Gamma(s,v)} \sum_{xy\in E(\gamma)} \min\{T_{x,y}, T_{y,x}\}\right) > Ct + CL \log n\right] \\ &\le \mathbb{P}\left[\left(\max_{v\in V} \min_{\substack{\gamma\in\Gamma(s,v)\\|\gamma|\le L}} \sum_{xy\in E(\gamma)} \min\{T_{x,y}, T_{y,x}\}\right) > Ct + CL \log n\right] \\ &\le \mathbb{P}\left[\left(\max_{v\in V} \min_{\substack{\gamma\in\Gamma(s,v)\\|\gamma|\le L}} \sum_{xy\in E(\gamma)} C \log n + C\min\{X_{x,y}, X_{y,x}\}\right) > Ct + CL \log n\right] + n^{-K} \\ &\le \mathbb{P}\left[\left(\max_{v\in V} \min_{\substack{\gamma\in\Gamma(s,v)\\|\gamma|\le L}} \sum_{xy\in E(\gamma)} C\min\{X_{x,y}, X_{y,x}\}\right) > Ct\right] + n^{-K} \\ &= \mathbb{P}[A_L > t] + n^{-K} \\ &\le \mathbb{P}[A > t] + n^{-K} + e^{-Cn^{1/3}}.\end{aligned}$$

Here, the first inequality is copied from (4). The second inequality is because restricting the feasible region of a minimization problem can only increase its optimal value. The third inequality follows from Lemma 8. The fourth inequality is straightforward. The equality follows from the definition of A_L and noting that $\{X_{x,y}\}$ have the same joint distribution as $\{Y_{x,y}\}$, and the last inequality follows from Lemma 6. This completes the proof of Theorem 3. ◀

Acknowledgements. We would like to thank an anonymous referee for pointing out a mistake in the statement of Corollary 5 in an earlier version.

References

1 Hüseyin Acan, Andrea Collevecchio, Abbas Mehrabian, and Nick Wormald. On the push&pull protocol for rumor spreading. *SIAM Journal on Discrete Mathematics*, 31(2):647–668, 2017. Available in `https://arxiv.org/abs/1411.0948` (conference version in PODC'15). `doi:10.1137/15M1033113`.

2 A. Auffinger, M. Damron, and J. Hanson. 50 years of first passage percolation. *arXiv*, 1511.03262 [math.PR], 2016.

3 Stephen Boyd, Arpita Ghosh, Balaji Prabhakar, and Devavrat Shah. Randomized gossip algorithms. *IEEE Trans. Inform. Theory*, 52(6):2508–2530, 2006. `doi:10.1109/TIT.2006.874516`.

4 Alan Demers, Dan Greene, Carl Hauser, Wes Irish, John Larson, Scott Shenker, Howard Sturgis, Dan Swinehart, and Doug Terry. Epidemic algorithms for replicated database maintenance. In *Proceedings of the Sixth Annual ACM Symposium on Principles of Distributed Computing*, PODC'87, pages 1–12, New York, NY, USA, 1987. ACM. `doi:10.1145/41840.41841`.

5 B. Doerr, M. Fouz, and T. Friedrich. Experimental analysis of rumor spreading in social networks. In *Design and analysis of algorithms*, volume 7659 of *Lecture Notes in Comput. Sci.*, pages 159–173. Springer, Heidelberg, 2012. `doi:10.1007/978-3-642-34862-4_12`.

6 R. Durrett. Stochastic growth models: recent results and open problems. In *Mathematical approaches to problems in resource management and epidemiology (Ithaca, NY, 1987)*, volume 81 of *Lecture Notes in Biomath.*, pages 308–312. Springer, Berlin, 1989. `doi:10.1007/978-3-642-46693-9_21`.

7 Uriel Feige, David Peleg, Prabhakar Raghavan, and Eli Upfal. Randomized broadcast in networks. *Random Structures Algorithms*, 1(4):447–460, 1990. `doi:10.1002/rsa.3240010406`.

8 G. Giakkoupis, Y. Nazari, and P. Woelfel. How asynchrony affects rumor spreading time. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*, PODC'16, pages 185–194, New York, NY, USA, 2016. ACM. `doi:10.1145/2933057.2933117`.

9 J. M. Hammersley and D. J. A. Welsh. *First-Passage Percolation, Subadditive Processes, Stochastic Networks, and Generalized Renewal Theory*, pages 61–110. Springer Berlin Heidelberg, Berlin, Heidelberg, 1965.

10 S. Janson. Tail bounds for sums of geometric and exponential variables. Available in `http://www2.math.uu.se/~svante/papers/sjN14.pdf`.

11 R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized rumor spreading. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, pages 565–574, 2000. `doi:10.1109/SFCS.2000.892324`.

12 D. Richardson. Random growth in a tessellation. *Proc. Cambridge Philos. Soc.*, 74:515–528, 1973.