Combinatorics, Probability and Computing (2020), **29**, pp. 190–199
doi:10.1017/S0963548319000385

CAMBRIDGE UNIVERSITY PRESS

ARTICLE

The string of diamonds is nearly tight for rumour spreading†

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

[1]Department of Mathematics, University of British Columbia, Vancouver, BC, V6T 1Z2, Canada, [2]Department of Computer Science, University of British Columbia, Vancouver, BC, V6T 1Z4, Canada and [3]Microsoft Research, Redmond, WA 98052, USA

*Corresponding author. Email: abbas.mehrabian@gmail.com

(Received 18 July 2017; revised 2 August 2019; first published online 4 November 2019)

Abstract

For a rumour spreading protocol, the spread time is defined as the first time everyone learns the rumour. We compare the synchronous push&pull rumour 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 $O(n^{1/3}\log^{2/3} n)$. This improves the $O(\sqrt{n})$ upper bound of Giakkoupis, Nazari and Woelfel (2016). Our bound is tight up to a factor of $O(\log n)$, as illustrated by the string of diamonds graph. We also show that if, for a pair α, β of real numbers, there exist infinitely many graphs for which the two spread times are n^α and n^β in expectation, then $0 \leqslant \alpha \leqslant 1$ and $\alpha \leqslant \beta \leqslant \frac{1}{3} + \frac{2}{3}\alpha$; and we show each such pair α, β is achievable.

2010 MSC Codes: Primary 68Q87; Secondary 05C81, 60C05, 68W15

1. Introduction

Randomized rumour spreading is an important paradigm for information dissemination in networks with numerous applications in network science, ranging from spreading of information on the Web or Twitter to diffusion of ideas and spreading of viruses in human communities. A well-studied rumour spreading protocol is the *(synchronous) push&pull protocol*, introduced by Demers *et al.* [4] and popularized by Karp *et al.* [10].

Definition 1.1 (synchronous push&pull protocol). Suppose that one node s in a network G is aware of a piece of information, the 'rumour', and wants to spread it to all nodes quickly. The synchronous protocol proceeds in rounds; in each round $1, 2, \ldots,$ all vertices perform their random actions simultaneously. Each vertex x calls a random neighbour y, and the two share any information they may have: if x knows the rumour and y does not, then x tells y the rumour (a *push*

†A preliminary version of this paper appeared in *Proceedings of the 21st International Workshop on Randomization and Computation (RANDOM'2017)*. This full version includes a new result, namely Theorem 2.5.

§Supported by NSERC.

¶Supported by an NSERC Postdoctoral Fellowship and a Simons–Berkeley Research Fellowship. Part of this work was done while this author was at the Simons Institute for the Theory of Computing at UC Berkeley. Currently at McGill University.

operation); if x does not know the rumour and y knows it, y tells x the rumour (a *pull* operation). Note that this is a synchronous protocol, that is, a vertex that receives a rumour in a certain round cannot also send it on in the same round, even though the vertex may be involved in multiple simultaneous calls initiated by other vertices. The *synchronous spread time* of G, denoted by $S(G, s)$, is the first time that everyone knows the rumour. This is a discrete random variable.

A point-to-point communication network can be modelled as an undirected graph: the nodes represent the processors and the links represent communication channels between them. The study of rumour spreading has several applications to distributed computing in such networks, of which we mention just two (see also [7]). 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, since 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 and does not grow more complex as the network grows); it is robust, in that 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 neighbour 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 assumes a synchronized computation and communication model, that is, all nodes take action simultaneously at discrete time steps. In many applications, and certainly for modelling information diffusion in social networks, this assumption is not realistic. In light of this, Boyd, Ghosh, Prabhakar and Shah [3] proposed an asynchronous model with a continuous timeline. This too is a randomized distributed algorithm for spreading a rumour in a graph, defined below. An *exponential clock* with rate λ is a clock that, once turned on, rings at times of a Poisson process with rate λ.

Definition 1.2 (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 rumour, and all clocks are turned on. Whenever the clock of a vertex x rings, it calls a random neighbour y. If x knows the rumour and y does not, then x tells y the rumour (a push operation); if x does not know the rumour and y knows it, y tells x the rumour (a pull operation). The *asynchronous spread time* of G, denoted by $A(G, s)$, is the first time that everyone knows the rumour.

Rumour spreading protocols in this model turn out to be closely related to Richardson's model for the spread of a disease [6, 11]. Moreover, the push&pull protocol is also quite similar to the first passage percolation model introduced by Hammersley and Welsh [9] with edges having independent exponential weights (see the survey [2]). The difference between the push&pull model and first passage percolation stems from the fact that, in rumour spreading models, each vertex contacts one neighbour at a time, so the rate at which x pushes the rumour to y is inversely proportional to the degree of x. A rumour can also be pulled from x to y; this happens at a rate determined by the degree of y. On regular graphs, the asynchronous push&pull protocol, Richardson's model and first passage percolation are fundamentally equivalent, assuming appropriate parameters are chosen. For general graphs, the push&pull model is equivalent 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 in first passage percolation.

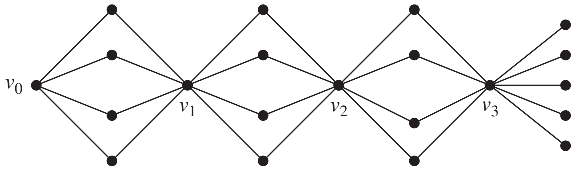

Figure 1. The string of diamonds graph $\mathcal{S}_{3,4,5}$.

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 times in the two time models. They write '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

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

Here and below $\widetilde{O}$ (and $\widetilde{\Omega}$) allows for polylogarithmic factors. Building on the ideas of [1] and using more involved couplings, Giakkoupis, Nazari and Woelfel [8] improved this bound to $O(n^{1/2})$. Our main contribution is to further improve this bound to $\widetilde{O}(n^{1/3})$.

Theorem 1.3. *For any* (G,s)*, we have*

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

An explicit graph was given in [1] with

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

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

While we also use a coupling argument, our argument is rather different from previous ones. Our coupling is motivated by viewing rumour 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 (discrete) distance that the rumour can traverse during a unit time interval in the asynchronous protocol is $O(n^{1/3})$ (see the proof of Lemma 3.1). The string of diamonds shows that this is best possible.

We shall make use of the following general bounds. It is proved in [8] that $\mathbb{E}[A(G,s)] \leqslant \mathbb{E}[S(G,s)] + O(\log n)$. Moreover, for all G and s (see [1, Theorem 1.3]), we have

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

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

2. Proof of Theorem 1.3.

For an n-vertex graph G and a starting vertex s, recall that $A(G,s)$ and $S(G,s)$ denote the asynchronous and synchronous spread times, respectively. Our main technical result is the following theorem (proved in Section 3), which has several corollaries.

Theorem 2.1. *Given any* $K > 0$*, there is a* $C > 0$ *such that, for any* (G,s) *and any* $t \geqslant 1$*, we have*

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

Corollary 2.2. *For any* (G, s), *we have* $\mathbb{E}[S(G,s)] = O(\mathbb{E}[A(G,s)]^{2/3} n^{1/3} \log n)$.

Proof. Apply Theorem 2.1 with $K = 1$ and $t = 3\mathbb{E}[A(G,s)] \leqslant 12n$. By Markov's inequality,

$$\mathbb{P}[S(G,s) > C(t + t^{2/3} n^{1/3} \log n)] \leqslant 1/3 + C/n \leqslant 1/2 \quad \text{for } n \text{ large enough.}$$

Since $t = O(n)$, this implies that 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 $\mathbb{E}[S(G,s)] = O(M)$. Consider the protocol which is the same as synchronous push&pull, except that if the rumour has not spread to all vertices by time M, then the process reinitializes. Clearly the spread time for this model is larger than the spread time for the synchronous model. Coupling the new process with push&pull, we obtain for any $i \in \{0, 1, 2, \ldots\}$ that $\mathbb{P}[S(G,s) > iM] \leqslant 2^{-i}$. Thus,

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

We obtain Theorem 1.3 from Corollary 2.2 by recalling that, for all G and s, $\mathbb{E}[A(G,s)] = \Omega(\log n)$.

Theorem 1.3 is tight up to an $O(\log n)$ factor: consider the following construction.

Definition 2.3 ($\mathcal{S}_{m,k,l}$). Let $m \geqslant 1$, $k \geqslant 2$ and $l \geqslant 0$ be integers. The graph $\mathcal{S}_{m,k,l}$ is built as follows. Start with $m+1$ vertices $v_0, v_1, \ldots, v_m$. For each $0 \leqslant i \leqslant m-1$, add k edge-disjoint paths of length 2 between v_i and v_{i+1}. Finally, add l new vertices and join them to v_m (see Figure 1 for an example). The graph $\mathcal{S}_{m,k,l}$ has $m(k+1) + l + 1$ vertices and $2km + l$ edges. If $l = 0$, this is called a 'string of diamonds' in [1].

The spread times of this graph are given by the following lemma, whose proof can be found in Section 4.

Lemma 2.4. *We have* $\mathbb{E}[S(\mathcal{S}_{m,k,l}, v_0)] = \Theta(m)$ *and* $\mathbb{E}[A(\mathcal{S}_{m,k,l}, v_0)] = \Theta(\log n + m/\sqrt{k})$.

If we let $m = \Theta(n^{1/3}(\log n)^{2/3})$ and $k = \Theta((n/\log n)^{2/3})$ such that $km + m + 1 = n$, we obtain a graph $\mathcal{S}_{m,k,0}$ with

$$\frac{\mathbb{E}[S(\mathcal{S}_{m,k,0}, v_0)]}{\mathbb{E}[A(\mathcal{S}_{m,k,0}, v_0)]} = \Omega(n/\log n)^{1/3},$$

which means Theorem 1.3 is tight up to an $O(\log n)$ factor.

It turns out that using our results and the above construction, we can get a more refined picture of what values the pair $(A(G,s), S(G,s))$ can take. More precisely, for α, β, we say the pair of exponents (α, β) is *attainable* if there exist infinitely many graphs (G, s) for which $\mathbb{E}[A(G,s)] = \widetilde{\Theta}(n^{\alpha})$ and $\mathbb{E}[S(G,s)] = \widetilde{\Theta}(n^{\beta})$. One may wonder which pairs (α, β) are attainable? The following theorem answers this question.

Theorem 2.5. *The pair* (α, β) *is attainable if and only if* $0 \leqslant \alpha \leqslant 1$ *and* $\alpha \leqslant \beta \leqslant \frac{1}{3} + \frac{2}{3}\alpha$.

Proof. The necessity of $0 \leqslant \alpha \leqslant \beta \leqslant 1$ follows from results in [1] mentioned above. Corollary 2.2 gives $\beta \leqslant \frac{1}{3} + \frac{2}{3}\alpha$.

To see that these conditions are sufficient, assume (α, β) satisfy $0 \leqslant \alpha \leqslant 1$ and $\alpha \leqslant \beta \leqslant \frac{1}{3} + \frac{2}{3}\alpha$. If $\beta > 0$, let $m = [n^{\beta}/2]$, $k = [n^{2\beta - 2\alpha}]$, and $l = n - 1 - m(k+1)$ so that $l \geqslant 0$ for n large enough. Lemma 2.4 gives

$$\mathbb{E}[S(\mathcal{S}_{m,k,l}, v_0)] = \Theta(m) = \Theta(n^\beta),$$
$$\mathbb{E}[A(\mathcal{S}_{m,k,l}, v_0)] = \Theta(\log n + m/\sqrt{k}) = \Theta(\log n + n^\alpha) = \widetilde{\Theta}(n^\alpha).$$

If $\beta = 0$, then $\alpha = 0$. In this case the star graph on n vertices has $\mathbb{E}[S(G,s)] = \Theta(1)$ and $\mathbb{E}[A(G,s)] = \Theta(\log n) = \widetilde{\Theta}(1)$ for any vertex s, as required (this is because the expected value of the maximum of n independent exponential random variables of mean 1 is $\Theta(\log n)$: see [1, Section 2.2] for details). $\square$

3. Proof of Theorem 2.1

In this section we fix the graph G and the starting vertex s. We first introduce several notations. For any vertex $v \in G$, 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 $|\gamma| := |E(\gamma)|$ 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 $1/\deg(u)$, so that these random variables are all independent. In the asynchronous protocol, since each vertex u calls any adjacent v at a rate of $1/\deg(u)$, we can write

$$A := A(G,s) = \max_{v\in V} \min_{\gamma\in\Gamma(s,v)} \sum_{xy\in E(\gamma)} \min\{Y_{x,y}, Y_{y,x}\}. \tag{3.1}$$

To see this, simply interpret $Y_{x,y}$ as the time it takes after one of x, y learns the rumour before x calls y.

For any positive integer L, we introduce the restriction to short paths:

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

For any L we trivially have $A_L \geqslant A$. To bound A from below, we have the following result, giving stochastic domination 'with high probability'.

Lemma 3.1. *There exists a C_0 such that, for any $C > C_0$, $t \geqslant 1$ and $L \geqslant Ct^{2/3}n^{1/3}$, we have*

$$\mathbb{P}[A_L > t] \leqslant \mathbb{P}[A > t] + e^{-L}.$$

Proof. We show that, in the asynchronous protocol, with probability $1 - e^{-L}$, during the interval $[0,t]$, the rumour does not travel along any simple path of length L. This automatically implies that the rumour does not travel along any longer path either. 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 may assume $L < n$.

Consider a path γ with vertices $\gamma_0, \gamma_1, \ldots, \gamma_L$. In order for the rumour to travel along γ, it is necessary that calls are made along the edges of γ in the order given by γ, at some sequence of times $0 \leqslant t_1 < \cdots < t_L \leqslant t$. Since along each edge the rumour 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

$$\frac{t^L}{L!}\prod_{i=1}^{L}\left(\frac{1}{\deg(\gamma_{i-1})} + \frac{1}{\deg(\gamma_i)}\right) \leqslant \left(\frac{2et}{L}\right)^L Q(\gamma), \tag{3.2}$$

where we define

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

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

https://doi.org/10.1017/S0963548319000385 Published online by Cambridge University Press

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 $\deg(\gamma_{i-1}) > \deg(\gamma_i) \leqslant \deg(\gamma_{i+1})$, and a *local maximum* at i if $\deg(\gamma_{i-1}) \leqslant \deg(\gamma_i) > \deg(\gamma_{i+1})$. 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

$$(\deg(\gamma_0), \ldots, \deg(\gamma_7)) = (\mathbf{5}, 5, 7, \mathbf{3}, 4, 4, \mathbf{2}, 5).$$

The local minima are shown in bold. 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 (again, in bold), then weakly increase up to the local maximum at the end of the segment. Henceforth, we use the term *segment* for a path with degrees having this property. (The first and last segments are special in that the local minimum could be at the beginning and end of the segment, respectively.)

Each path provides an ordered sequence of segments. Denote the segments of γ by $\sigma_1, \ldots, \sigma_s$, and note that $s \leqslant L/2 + 1$, since each segment (except possibly the first and the last ones) 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 subpaths). Note also 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,

$$\sum_{|\gamma|=L} Q(\gamma) \leqslant \sum_{s=1}^{L/2+1} \sum_{|\sigma_1|+\cdots+|\sigma_s|=L} \frac{1}{s!} \prod_{i=1}^{s} Q(\sigma_i), \tag{3.3}$$

where the second sum is over ordered 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). The $1/s!$ factor comes from the aforementioned fact that at most one ordering of each s-tuple is possible and needs to be counted.

We now bound the right-hand side of (3.3). 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 *centre* 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 (3.3), we first fix s and bound the number of options for the sequence $(\pi(\sigma_1), \ldots, \pi(\sigma_s))$. There are n^s choices for the centres, 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 $2^L n^s$ 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 $m \leqslant 2^L n^s$, and let $T_{j,k}$ denote the kth component of T_j, *i.e.* the type specified for σ_k in T_j. Thus,

$$\sum_{|\sigma_1|+\cdots+|\sigma_s|=L} \prod_{i=1}^{s} Q(\sigma_i) = \sum_{j=1}^{m} \sum_{(\pi(\sigma_1),\ldots,\pi(\sigma_s))=T_j} \prod_{i=1}^{s} Q(\sigma_i) \leqslant \sum_{j=1}^{m} \prod_{k=1}^{s} \Bigg(\sum_{\pi(\sigma_k)=T_{j,k}} Q(\sigma_k) \Bigg).$$

Next, we claim that each term in the last product, 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 $\ell = \ell^- + \ell^+$. The constraints on the degrees along a segment $\sigma = v_0, v_1, \ldots, v_{\ell^-}, \ldots, v_\ell$ 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 over all *walks* of length $\ell^- + \ell^+$ whose ℓ^-th vertex is x, but waiving the degree monotonicity constraint, then we get 1 (since the number of choices for the neighbours cancels

https://doi.org/10.1017/S0963548319000385 Published online by Cambridge University Press

out the degree reciprocals). Restricting to simple paths with piecewise monotone degrees only decreases this. Thus we obtain

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

Plugging this back into (3.3) yields

$$\sum_{|\gamma|=L} Q(\gamma) \leqslant \sum_{s=1}^{L/2+1} 2^L n^s / s! \leqslant \left(\frac{8en}{L}\right)^{L/2+1}.$$

(We use here that $L < n$, hence each term is less than half the next and the sum is at most twice its last term.)

Therefore, by (3.2), the probability that the rumour travels along some path of length L is bounded by

$$\sum_{|\gamma|=L} \left(\frac{2et}{L}\right)^L Q(\gamma) \leqslant \left(\frac{2et}{L}\right)^L \left(\frac{8en}{L}\right)^{L/2+1} \leqslant C_1 n (C_2 n t^2 / L^3)^{L/2},$$

which is at most e^{-L} for $L \geqslant C t^{2/3} n^{1/3}$, completing the proof. □

In (3.1) 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 learns the rumour. Suppose the first round *strictly after* q_{uv} at which u calls v is F_{uv}, and define $T_{u,v} = F_{uv} - q_{uv}$. Note that $T_{u,v}$ is a positive integer, and is a geometric random variable: $\mathbb{P}[T_{u,v} \geqslant k] = (1 - 1/\deg(u))^{k-1}$ for any $k = 1, 2, \ldots$. Moreover, observe that both u and v are informed by round $q_{uv} + \min\{T_{u,v}, T_{v,u}\}$, hence we have

$$S := S(G, s) \leqslant \max_{v \in V} \min_{\gamma \in \Gamma(s,v)} \sum_{xy \in E(\gamma)} \min\{T_{x,y}, T_{y,x}\}. \tag{3.4}$$

We now have a max-min expression for $S(G, s)$. However, a major difficulty in the synchronous model 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. Initially, the clocks $Z_{u,v}$ are turned on, and the clocks $Z'_{u,v}$ are off. At later times we may turn off $Z_{u,v}$ and turn on $Z'_{u,v}$. We say the clocks $Z_{u,v}, Z'_{u,v}$ are *located* at vertex u. Continuous time at each vertex will advance separately, though there will be synchronized rounds as defined below.

For each round $1, 2, \ldots$, we visit the vertices one by one. For each vertex u, we let all active clocks located at u advance, until one of the clocks rings. 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: at each vertex, a different amount of time has elapsed on the clocks.) Having determined the spread of the rumour at the present round, whenever a vertex u gets informed of the rumour, 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 will have already taken place at an earlier round.) Observe that because of memorylessness of the exponential distribution, and since all clocks at u have the same rate, this process generates a random sequence of independent uniform neighbours, so it is equivalent to the synchronous protocol.

https://doi.org/10.1017/S0963548319000385 Published online by Cambridge University Press

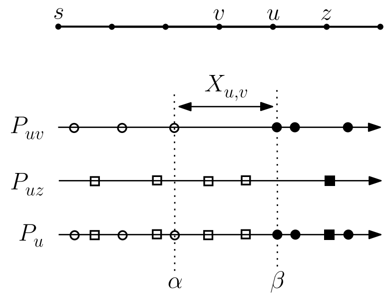

Figure 2. Illustration of the proof of Lemma 3.2: the graph is shown at the top, together with the three processes P_{uv}, P_{uz} and P_u. The rumour starts from vertex s. Suppose that at discrete round $q_{uv} = 5$ (continuous time α at vertex u), vertex v is informed from the left; at this moment clock Z_{uv} (corresponding to empty circles) is turned off and Z'_{uv} (black circles) is turned on. After $T_{u,v} = 3$ discrete rounds (continuous time $X_{u,v} = \beta - \alpha$ has passed), vertex u calls v and gets informed at discrete round $F_{uv} = 8$ (continuous time β); at this moment clock Z_{uz} (white squares) is turned off and Z'_{uz} (black squares) is turned on.

Now let us see what the random variables $T_{u,v}$ are in this set-up. For each ordered pair u, v, observe that the combined 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 several initial rings come from Z and subsequent rings from Z'.) Let

$$P_u := \bigcup_{v \sim u} P_{u,v},$$

and note that P_u is a Poisson process with rate 1. See Figure 2 for an illustration.

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 clock $Z'_{u,v}$ has rung since the time it was turned on, hence it is exponential with rate $1/\deg(u)$. Since the clocks are independent, the random variables $X_{u,v}$ are also independent. The times at which the Z' clocks are turned on depend on other clocks in a non-trivial manner, but do not affect the $X_{u,v}$ variables. Thus we have proved the following.

Lemma 3.2. *The random variables $\{X_{u,v}\}$ defined above are mutually independent.*

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

Lemma 3.3. *For every K and large enough $C \geqslant C_0(K)$, with probability at least $1 - n^{-K}$, for all adjacent pairs u, v we have $T_{u,v} \leqslant C \log n + CX_{u,v}$.*

Proof. We show that for any adjacent pair x, y, we have $\mathbb{P}(T_{u,v} > C \log n + CX_{u,v}) \leqslant n^{-K-2}$, and then apply the union bound over all edges.

Observe that, conditioned on $X_{u,v} = t$, the random variable $T_{u,v} - 1$ is Poisson with rate $t \times (\deg(u) - 1)/\deg(u) \leqslant t$. Indeed, this is the number of rings over time t of the $\deg(u) - 1$ active clocks on edges (u, w) with $w \neq v$. Let $\mathrm{Poi}(t)$ denote a Poisson random variable with mean $t > 0$. For $m \geqslant et$, we have $\mathbb{P}(\mathrm{Poi}(t) = m) \leqslant e^{-1}\mathbb{P}(\mathrm{Poi}(t) = m - 1)$, hence $\mathbb{P}(\mathrm{Poi}(t) > et + m) \leqslant e^{-m}$. This gives

$$\mathbb{P}[T_{u,v} - 1 > (K+2)\log n + eX_{u,v} | X_{u,v} = t] \leqslant \mathbb{P}[\mathrm{Poi}(t) > (K+2)\log n + et] \leqslant n^{-K-2}.$$

The claim follows with $C = \max(e, K+2)$. $\square$

Theorem 2.1 now follows easily from our lemmas.

Proof of Theorem 2.1. Given K, pick C sufficiently large so that Lemmas 3.1 and 3.3 hold. Fix $t \geqslant 1$ and let $L = Ct^{2/3}n^{1/3}$. We have

$$\begin{aligned} &\mathbb{P}[S > Ct + CL\log n] \\ &\leqslant \mathbb{P}\left[\left(\max_{v \in V} \min_{\gamma \in \Gamma(s,v)} \sum_{xy \in E(\gamma)} \min\{T_{x,y}, T_{y,x}\}\right) > Ct + CL\log n\right] \end{aligned}$$

https://doi.org/10.1017/S0963548319000385 Published online by Cambridge University Press

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

Here, the first inequality is copied from (3.4). 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 3.3. 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 3.1. This completes the proof of Theorem 2.1. □

4. Proof of Lemma 2.4

In this section we show that

$$2m \leqslant \mathbb{E}[S(\mathcal{S}_{m,k,l},v_0)] \leqslant 4m+1 \quad \text{and} \quad \mathbb{E}[A(\mathcal{S}_{m,k,l},v_0)] = \Theta(\log n + m/\sqrt{k}).$$

Fix $m\geqslant 1$, $k\geqslant 1$ and $l\geqslant 0$, and let $G=\mathcal{S}_{m,k,l}$. Recall that $v_0,\dots,v_m$ are the vertices connecting the diamonds in $\mathcal{S}_{m,k,l}$.

Since the graph distance between v_0 and v_m is $2m$, we have $S(G,v_0)\geqslant 2m$ deterministically. Fix $0\leqslant i\leqslant m-1$ and suppose that at some time v_i is informed and v_{i+1} is uninformed. We claim that the expected time to inform v_{i+1} is at most 4. Let u be some common neighbour of v_i and v_{i+1}. It takes two rounds in expectation for u to pull the rumour from v_i, and another two rounds for it to push the rumour to v_{i+1}, so the claim follows. Once all the v_i are informed, every other vertex will be informed in the next round. Therefore, $\mathbb{E}[S(G,v_0)]\leqslant 4m+1$.

Next we show $\mathbb{E}[A(G,v_0)] = O(\log n + m/\sqrt{k})$. Let Y_i denote the communication time between v_i and v_{i+1} (the first time that v_{i+1} learns the rumour, assuming initially only v_i knows the rumour). Between v_i and v_{i+1} there are k disjoint paths of length 2, so Y_i is stochastically dominated by $Z:=\min\{Z_1,\dots,Z_k\}$, where the Z_i are independent random variables equal in distribution to the sum of two independent exponential random variables with rate 1/2. (The difference between Y and Z stems from calls initiated at v_i, v_{i+1}.) Since each Z_i has density $(t/4)e^{-t/2}$ on $\mathbb{R}_+$, we have

$$\mathbb{P}[Z>t] = \left(1+\frac{t}{2}\right)^k e^{-kt/2}.$$

The change of variable $u=k(t/2+1)$ gives

$$\mathbb{E}[Z] = \int_0^\infty \mathbb{P}[Z>t]\,\mathrm{d}t = \frac{2e^k}{k^{k+1}}\int_k^\infty u^k e^{-u}\,\mathrm{d}u.$$

The integral of $u^k e^{-u}$ from 0 to ∞ is $k!$, so

$$\mathbb{E}[Z] \leqslant \frac{2e^k k!}{k^{k+1}} = O(1/\sqrt{k}).$$

Hence, the expected time for all the v_i to learn the rumour is at most $O(mk^{-1/2})$. After this has happened, any other vertex pulls the rumour in Exp (1) time. The expected value of the maximum

https://doi.org/10.1017/S0963548319000385 Published online by Cambridge University Press

of at most n independent Exp (1) variables is the harmonic sum $H_n \leqslant 1 + \log n$, so $\mathbb{E}[A(G, v_0)] = O(\log n + mk^{-1/2})$.

Finally, we show $\mathbb{E}[A(G, v_0)] = \Omega(\log n + mk^{-1/2})$. The bound $\mathbb{E}[A(G, v_0)] = \Omega(\log n)$ holds for any n-vertex graph G (see [1, Theorem 1.3]), so we need only show that $\mathbb{E}[A(G, v_0)] = \Omega(mk^{-1/2})$. In fact, since each of the intermediate v_i is a cut-vertex, we need only show that $\mathbb{E}[Y_i] = \Omega(k^{-1/2})$ for each i.

Suppose that at time s only v_i is informed. For any $t > 0$, if v_{i+1} is informed by time $s + t$, then during the time interval $[s, s + t]$, either the clock of v_i has rung at least once, or the clock of v_{i+1} has rung at least once, or the clock of one of their k common neighbours has rung at least twice. Since the ringing times at each vertex are a Poisson process, we find

$$\mathbb{P}[Y_i \leqslant t] \leqslant 2(1 - e^{-t}) + k(1 - e^{-t} - te^{-t}) \leqslant 2t + kt^2/2.$$

Hence, with $t = 1/3\sqrt{k} \leqslant 1/3$,

$$\mathbb{E}[Y_i] \geqslant \frac{1}{3\sqrt{k}}\mathbb{P}\left[Y_i \geqslant \frac{1}{3\sqrt{k}}\right] \geqslant \frac{1}{3\sqrt{k}}(1 - 2/3 - 1/18) = \Omega\left(\frac{1}{\sqrt{k}}\right),$$

completing the proof of Lemma 2.4.

Acknowledgements

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. We thank the workshop organizers and BIRS for making it happen. The work was completed at Microsoft Research in Redmond. We are grateful to Microsoft for their support.

We would like to thank an anonymous referee of the RANDOM'2017 workshop for pointing out a mistake in the statement of Theorem 1.3 in an earlier version.

References

[1] Acan, H., Collevecchio, A., Mehrabian, A. and Wormald, N. (2017) On the push&pull protocol for rumor spreading. *SIAM J. Discrete Math.* **31** 647–668.

[2] Auffinger, A., Damron, M. and Hanson, J. (2017) *50 Years of First-Passage Percolation*, Vol. 68 of University Lecture Series, American Mathematical Society.

[3] Boyd, S., Ghosh, A., Prabhakar, B. and Shah, D. (2006) Randomized gossip algorithms. *IEEE Trans. Inform. Theory* **52** 2508–2530.

[4] Demers, A., Greene, D., Hauser, C., Irish, W., Larson, J., Shenker, S., Sturgis, H., Swinehart, D. and Terry, D. (1987) Epidemic algorithms for replicated database maintenance. In *Proceedings of the Sixth Annual ACM Symposium on Principles of Distributed Computing (PODC '87)*, ACM, pp. 1–12.

[5] Doerr, B., Fouz, M. and Friedrich, T. (2012) Experimental analysis of rumor spreading in social networks. In *Design and Analysis of Algorithms*, Vol. 7659 of Lecture Notes in Computer Science, Springer, pp. 159–173.

[6] Durrett, R. (1989) Stochastic growth models: Recent results and open problems. In *Mathematical Approaches to Problems in Resource Management and Epidemiology (Ithaca, NY, 1987)*, Vol. 81 of Lecture Notes in Biomathematics, Springer, pp. 308–312.

[7] Feige, U., Peleg, D., Raghavan, P. and Upfal, E. (1990) Randomized broadcast in networks. *Random Struct. Alg.* **1** 447–460.

[8] Giakkoupis, G., Nazari, Y. and Woelfel, P. (2016) How asynchrony affects rumor spreading time. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing (PODC '16)*, ACM, pp. 185–194.

[9] Hammersley, J. M. and Welsh, D. J. A. (1965) *First-Passage Percolation, Subadditive Processes, Stochastic Networks, and Generalized Renewal Theory*, Springer, pp. 61–110.

[10] Karp, R., Schindelhauer, C., Shenker, S. and Vöcking, B. (2000) Randomized rumor spreading. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, IEEE, pp. 565–574.

[11] Richardson, D. (1973) Random growth in a tessellation. *Proc. Cambridge Philos. Soc.* **74** 515–528.

Cite this article: Angel O, Mehrabian A and Peres Y (2020). The string of diamonds is nearly tight for rumour spreading. *Combinatorics, Probability and Computing* **29**, 190–199. https://doi.org/10.1017/S0963548319000385