

RESEARCH OUTPUTS / RÉSULTATS DE RECHERCHE

Speeding up non-Markovian First Passage Percolation with a few extra edges

Medvedev, Alexey; Pete, Gábor

Published in:
Advances in Applied Probability

DOI:
[10.1017/apr.2018.39](https://doi.org/10.1017/apr.2018.39)

Publication date:
2018

[Link to publication](#)

Citation for published version (HARVARD):
Medvedev, A & Pete, G 2018, 'Speeding up non-Markovian First Passage Percolation with a few extra edges', *Advances in Applied Probability*, vol. 50, no. 3, pp. 858-886. <https://doi.org/10.1017/apr.2018.39>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Speeding up non-Markovian First Passage Percolation with a few extra edges

Alexey Medvedev ^{*} Gábor Pete [†]

Abstract

One model of real-life spreading processes is First Passage Percolation (also called SI model) on random graphs. Social interactions often follow bursty patterns, which are usually modelled with i.i.d. heavy-tailed passage times on edges. On the other hand, random graphs are often locally tree-like, and spreading on trees with leaves might be very slow, because of bottleneck edges with huge passage times. Here we consider the SI model with passage times following a power law distribution $\mathbb{P}(\xi > t) \sim t^{-\alpha}$, with infinite mean. For any finite connected graph G with a root s , we find the largest number of vertices $\kappa(G, s)$ that are infected in finite expected time, and prove that for every $k \leq \kappa(G, s)$, the expected time to infect k vertices is at most $O(k^{1/\alpha})$. Then, we show that adding a single edge from s to a random vertex in a random tree \mathcal{T} typically increases $\kappa(\mathcal{T}, s)$ from a bounded variable to a fraction of the size of \mathcal{T} , thus severely accelerating the process. We examine this acceleration effect on some natural models of random graphs: critical Galton-Watson trees conditioned to be large, uniform spanning trees of the complete graph, and on the largest cluster of near-critical Erdős-Rényi graphs. In particular, at the upper end of the critical window, the process is already much faster than exactly at criticality.

Keywords: temporal networks; near-critical random graphs; Galton-Watson tree; Erdős-Rényi graph; Pólya urn process; spreading phenomena; SI model; First Passage Percolation; bursty time series; non-Markovian processes. **MSC2010 classification:** 60K35; 60K37; 05C80; 82C99; 90B18

1 Introduction

1.1 Motivation and background

Spreading is one of the most important dynamic processes on complex networks as it is the basis of a broad range of phenomena from epidemic contagion to diffusion of innovations [Ves12, PSCVMM15]. One of the original, and still primary, reasons for studying networks is to understand the mechanisms by which diseases, information, computer viruses, rumors, innovations spread over them [New03]. The spreading problem initially came from epidemiology, and the terminology still remembers these roots: in the simplest spreading model, the two-state *susceptible-infected (SI) model*, any particular person can be either in susceptible (S) or in infected (I) state, with the transition rule $S \rightarrow I$, meaning that once an infection is obtained, the person stays infected forever. It is usually assumed that initially there is only one infected person.

^{*}Namur Institute for Complex Networks (naXys), Université de Namur, Rempart de la Vierge, 8, Namur 5000 Belgium and ICTEAM, Université Catholique de Louvain, Av. Georges Lemaitre 4, 1348 Ottignies-Louvain-la-Neuve, Belgium. Email: `an_medvedev@yaho dot com`

[†]Alfréd Rényi Institute of Mathematics, Reáltanoda u. 13-15., Budapest 1053 Hungary, and Budapest University of Technology and Economics, Egrý J. u. 1, Budapest 1111 Hungary. Email: `robagetep@gmail dot com`

In the usual mathematical model of SI spreading, a rooted connected simple graph $G = (V, E, s)$ is given, with vertices representing people, and edges representing the connections between them, through which the infection can pass. Initially, only the root $s \in V$ is infected. The edges $e \in E$ have i.i.d. random positive weights $\xi(e)$, with common distribution ξ , representing the passage time of the infection. Infection is then transmitted along the edges from infected vertices to susceptible ones with passage times according to the edge weights (measured from the moment of infecting the first endpoint of the edge). The process runs until all vertices get infected. Alternatively, and this is how the model was first defined in the mathematical community under the name *First Passage Percolation (FPP)* [HW65], one may consider $\xi(e)$ as the length of the edge $e \in E$, and then we can think of some fluid flowing through this random medium at speed 1, starting from the source s . In the two models, SI and FPP, the moment when a vertex is reached by the infection/fluid will be the same. Then, we can measure spreading via two different, but of course closely related stochastic processes: either by $(T_k)_{k=1}^{|V|}$, the sequence of times when exactly k vertices are infected in the process, or by $(N_t)_{t \geq 0}$, the number of vertices reached by time t . In FPP one is typically more concerned with the process $(N_t)_{t \geq 0}$ (see, e.g., [Kes87]), while in the current paper we consider the process $(T_k)_{k=1}^{|V|}$, therefore we will prefer the SI model notation.

The SI spreading model has attracted much attention in the network science community, giving rise to a solid number of papers (e.g., [IM09, KKP⁺11, MH13, JPKK14, HK14]). In the mathematical literature, there has been an enormous number of papers studying this model on various graphs and using various distributions for ξ ; see, e.g., [HW65, Kes87, FP93, BvdHH10, BvdHH11, BvdHK14, BvdHH17]; two surveys are [AHD15] on the first 50 years of the model on Euclidean lattices and [vdH17, Chapter 8] on other graphs. Although originally ξ was typically assumed to have Exponential or Bernoulli distribution, more recently, motivated by the bursty nature of communications in many real networks, ξ with *power law distributions* $\text{pow}(\alpha)$ have also been considered: $\mathbb{P}(\xi > t) = 1 \wedge (t/t_0)^{-\alpha}$, where $t_0 > 0$ and $\alpha > 0$. In this case the model is non-Markovian, which creates many obstacles in studying it; this will be the situation in the present work, as well.

Most sparse (e.g., bounded average degree) random graph models produce graphs that are *locally tree-like*: a large neighbourhood of a random root is a tree with high probability, maybe with some extra edges decorating the tree. Typically, *supercritical* random graph models have been studied, where the unique giant connected component looks locally like a fast-growing supercritical Galton-Watson tree, maybe with decorations. (In more precise terms, the Benjamini-Schramm limit of the giant clusters is an infinite unimodular random tree with average degree larger than 2, or more generally, an invariantly non-amenable unimodular random graph [BS01, AL07]). This local fast-growing tree structure has been used in several works very successfully to understand the behavior of SI and FPP on the entire graph [BvdHK14, BvdHH17, BvdHK17], even with a general absolutely continuous distribution for ξ .

However, it was observed in [IM09] that the SI spreading with heavy-tailed ξ on *finite trees*, when started from a typical site, may be very slow. Real-life examples with such ξ include, for example, information cascades in Twitter [GCSdA⁺17, BGD⁺17]. A striking feature of slow spreading when ξ has infinite mean was noticed in [HK14] via computer simulations, which is most apparent when not the curve $t \mapsto \langle N_t \rangle$ is considered, but the “inverse” curve $k \mapsto \langle T_k \rangle$. Namely, running the simulation of the process $T = (T_k)_{k=1}^n$ on a tree for M times, and defining $\langle T_k \rangle = \frac{1}{M} \sum_{i=1}^M T_k^{(i)}$, the *spreading curve* $(\langle T_k \rangle, k/n)$ exhibits “uncontrolled large plateaux”, which do not decrease and do not converge as we increase the number of runs (see the dark blue curves on Figure 1.1 (a)). Plateaux of similar type were also empirically found in [BGD⁺17]. On the other hand, adding just one extra edge to the tree (*which does not change the local statistics of the graph!*) makes the

spreading curve quite smooth; see the lighter green curves on Figure 1.1 (a), and also Figure 1.1 (b) that shows SI spreading on the cycle, with power law exponent $\alpha \in (1/2, 1)$.

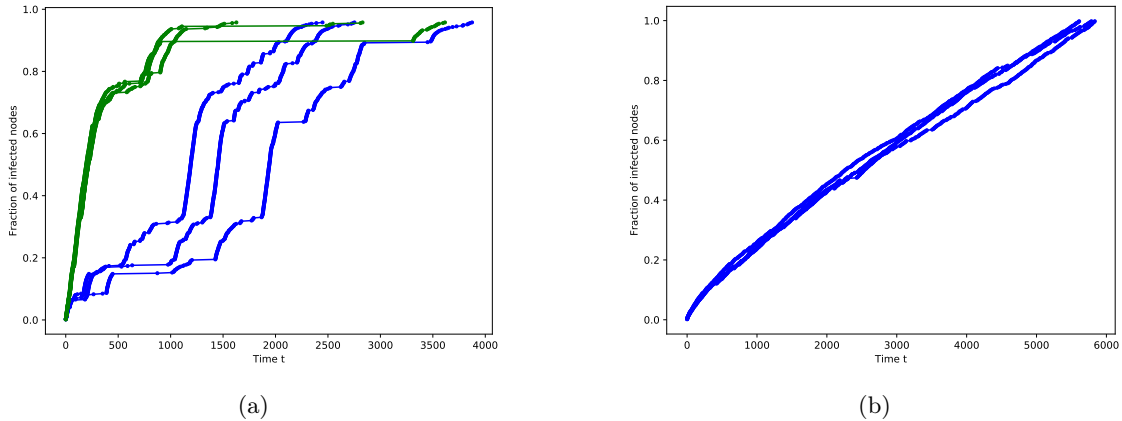


Figure 1.1: Spreading curves of the SI model simulation with power law weights ξ with $\alpha = 0.8$. (a) For the three dark blue curves, the underlying graph is a tree with 472 vertices; for the three lighter green curves, one fixed edge is added between the root and a random vertex; b) the underlying graph is a cycle C_n with 472 vertices.

The first motivation for our paper was to understand the striking effect of adding a few edges to large finite trees. After describing this phenomenon in detail for arbitrary deterministic graphs, we address some natural models of random trees, using a variety of techniques. Since supercritical trees are typically not finite (or, if a supercritical GW tree is conditioned to be finite, it becomes subcritical [LP16, Section 5.7]), the most natural and interesting case seems to be adding a random edge to a critical random tree, conditioned to be large in one way or another. Or, in the most natural model that we will consider, the largest cluster in a *near-critical Erdős-Rényi graph* $G(n, \frac{1}{n} + \frac{\lambda}{n^{4/3}})$, for $\lambda \ll 0$ is very likely to be a large critical tree, while for $\lambda \gg 0$ is very likely to be a critical tree with several extra edges (see Section 6 for references). In particular, the SI process with power law weights for $\lambda \ll 0$ is slow and has uncontrolled large plateaux, while most of the graph is infected in a fast and smooth way for $\lambda \gg 0$ (see Figure 6.1).

The plateaux are explained by the presence of so-called temporal “bottlenecks”, which are large passage times that occur on some particular edges. Indeed, if there is a k for which $\mathbb{E}(T_k) < \infty$ but $\mathbb{E}(T_{k+1}) = \infty$, then the Law of Large Numbers tells us that

$$\frac{1}{M} \left(\sum_{i=1}^M T_{k+1}^{(i)} - \sum_{i=1}^M T_k^{(i)} \right) \rightarrow \infty \text{ a.s., as } M \rightarrow \infty,$$

and hence the first uncontrolled plateau in the spreading curve is between $\langle T_k \rangle$ and $\langle T_{k+1} \rangle$. The next plateau occurs as if we started the process anew after the first bottleneck, without the already infected part of the graph. Hence all the later plateaux can be investigated with the method we describe below, and here we will focus on the time when the first plateau appears.

It is easy to obtain some intuition how the temporal bottlenecks disappear for the SI process on the cycle with $\alpha \in (1/2, 1)$: there are always at least *two* edges through which the infection might proceed, and the minimum of two independent $\text{pow}(\alpha)$ variables already has a finite expectation. In fact, we will show in Subsection 2.1 that $\mathbb{E}(T_k) \asymp k^{1/\alpha} < \infty$ for every $1 \leq k \leq n$. However, because of the non-Markovian nature of the process (with already old edges getting occupied even later), making such an argument rigorous for a general graph with cycles is not at all trivial.

In any case, we see that for cycles $\alpha = 1/2$ must be a threshold, and in the current paper we will concentrate on the case of $\alpha \in (1/2, 1)$. Nevertheless, our results can be extended to the case

when the infection has $d > 2$ ways to proceed and $\alpha > 1/d$; simplistically, the main reason for this bound is that then the minimum of d independent $\text{pow}(\alpha)$ variables has a finite expectation. See Subsection 1.3 for more details.

1.2 Results and outline

Summarizing the above assumptions, in this paper we consider the SI spreading model with i.i.d. power law transmission times and we study the role of cycles in speeding up of the process when $\alpha \in (1/2, 1)$ and thus eliminating large plateaux on (or ‘smooth’) the average spreading curve up to a certain level. We rigorously identify when the first possible temporal bottleneck appears and its relation to the graph structure. We show that for each finite connected graph G on n vertices there exists a specific threshold $\kappa(G, s)$, where s is an initially infected vertex, such that the average time to infect $k \leq \kappa(G, s)$ vertices is finite, and is even bounded by a polynomial function in k , and when $k > \kappa(G, s)$ the average time is infinite. The result, proved in Section 3, can be stated as follows.

Theorem 1.1. *Consider a graph G on n vertices with a root s and the SI spreading process $T = (T_j)_{j=1}^n$ with power law weights with $\alpha \in (1/2, 1)$. Then there exists the number $\kappa(G, s)$ such that for each k , where $1 \leq k \leq \kappa(G, s)$, the expected time to infect k vertices is bounded by*

$$\mathbb{E}(T_k) \leq Ck^{1/\alpha},$$

and for $k > \kappa(G, s)$, the expectation $\mathbb{E}(T_k) = \infty$.

The number $\kappa(G, s)$ identifies the position of the first temporal bottleneck for the process, and has a simple combinatorial description, which is presented in the following Lemma.

Lemma 1.2. *Let G be a finite rooted graph with root s and let T be the SI spreading process on G with weights having absolutely continuous distribution ξ , such that $\mathbb{E}(\xi) = \infty$. Then,*

$$\kappa(G, s) = \min_{e \in E(G)} |\mathcal{C}(s, G \setminus e)|,$$

where $|\mathcal{C}(s, G \setminus e)|$ is the size of the connected component of vertex s in the graph G without edge e .

In Section 2, we will describe two examples, the cycle and the star graph, for both of which we have $\kappa = n - 1$ and $\mathbb{E}(T_\kappa) \asymp n^{1/\alpha}$, but for quite different reasons. We have been unable to identify either of these examples as the ‘hardest graph to infect’, which can be considered as an explanation for the non-triviality of Theorem 1.1.

Next, we will study the value of $\kappa(G, s)$ in some near-critical random graph models, and show that by adding one or a few edges, it increases from a bounded value to a positive proportion of all the vertices. In our first example, we add an extra edge to a *critical Galton-Watson tree* conditioned to have depth at least N , between the starting vertex s and a uniform random vertex. Here our main tool is comparison with the critical GW tree conditioned to be infinite [Kes86], and our goal is to be as elementary as possible, to develop arguments that may be applicable to other critical random tree models that one may be interested in (see Section 4). In our second example, we add one random edge to the *uniform spanning tree* of the complete graph; this model makes it possible to use elegant methods relying on Loop-Erased Random Walks (via Wilson’s algorithm [Wil96]) and Pólya urns with time-dependent increments [Pem90], very different from the first example (see Section 5). Finally, our third example is the ‘physically’ most relevant one: the *near-critical Erdős-Rényi graph* $G(n, \frac{1}{n} + \frac{\lambda}{n^{4/3}})$, $\lambda \in \mathbb{R}$. Here the analysis is based on fine understanding of the metric

structure of the graph [ABBG10b], encoded typically as a Gromov-Hausdorff scaling limit using Aldous' Continuum Random Tree [Ald91b] as a basic building block (see Section 6).

In the first two of these models, in order to raise $\kappa(G, s)$ significantly, it is important that the extra edge is incident to s , or in other words, that s lies in the cycle that we have created. Instead of adding the extra edge, we could have moved s onto the "spine" of the tree, so that there are two disjoint long paths emanating from s , and this would have a very similar accelerating and smoothing effect. So, in these cases, adding the extra edge could just be considered as a simple way to make the starting vertex special. This also makes sense from a practical point of view: if one's goal is to speed up the infection process, it is easier to create a new link, or to start the infection at two random places (which is almost the same), than finding a special vertex in the graph to start the process. In the third case, near-critical Erdős-Rényi graphs, the starting vertex will not be special, and therefore $\kappa(G, s)$ will remain of constant order. However, after infecting a tiny portion of the graph, progressing with k a little bit (although typically with a huge T_k), we get to a vertex with a large κ , and from there we infect most of the graph in a fast and smooth way.

We now state our exact results for the three near-critical random graph models. In our first example, we add a single extra edge e between the initially infected vertex s and a uniform random vertex in a critical Galton-Watson tree \mathcal{T} . This graph will be denoted by \mathcal{T}_{+e} , while $|\mathcal{T}|$ denotes the number of vertices in \mathcal{T} .

Theorem 1.3. *Let \mathcal{T}^N be a critical Galton-Watson tree, conditioned on surviving to at least N generations, $Z_N > 0$, and consider the SI spreading process with power law weights with $\alpha \in (1/2, 1)$. Then as $N \rightarrow \infty$,*

- (1) *the sequence of r.v. $\kappa(\mathcal{T}^N, s)$ is tight;*
- (2) *for any $\varepsilon > 0$ there exists $\delta > 0$, such that*

$$\mathbb{P}\left(\frac{\kappa(\mathcal{T}_{+e}^N, s)}{|\mathcal{T}_{+e}^N|} > \delta\right) > 1 - \varepsilon.$$

Secondly, a similar result, but with a different approach, is provided for the model of uniform spanning tree of the complete graph.

Theorem 1.4. *Let $\text{UST}(n)$ be the uniform spanning tree of the complete graph on n vertices, and consider the SI spreading process with power law weights with $\alpha \in (1/2, 1)$, with starting vertex s . Let e be a uniform random extra edge with one endpoint being s . Then, $\forall \varepsilon > 0$ there is $\delta > 0$ such that, for all n large enough,*

$$\mathbb{P}\left(\frac{\kappa(\text{UST}(n)_{+e}, s)}{n} > \delta\right) > 1 - \varepsilon.$$

Thirdly:

Theorem 1.5. *Let $\mathcal{C} = \mathcal{C}_1^n(\lambda)$ be the largest cluster of the Erdős-Rényi graph $G(n, \frac{1}{n} + \frac{\lambda}{n^{4/3}})$.*

- (1) *Let σ be a uniform random vertex in \mathcal{C} . Then, as $n \rightarrow \infty$, the sequence of variables $\kappa(\mathcal{C}, \sigma)$ is tight for any fixed $\lambda \in \mathbb{R}$: for every $\varepsilon > 0$, if $K > 0$ is large enough, then*

$$\mathbb{P}(\kappa(\mathcal{C}, \sigma) < K) > 1 - \varepsilon,$$

for all $n > n_0(\lambda, \varepsilon)$ large enough.

- (2) For every $\lambda \in \mathbb{R}$ and $\epsilon > 0$, there exists a $\delta_0(\epsilon) > 0$ that is independent of λ , and a $\delta_1(\lambda, \epsilon) > 0$ such that

$$\mathbb{P}\left(\delta_0|\mathcal{C}| < \max_{s \in V(\mathcal{C})} \kappa(\mathcal{C}, s) < (1 - \delta_1)|\mathcal{C}|\right) > 1 - \epsilon.$$

In particular, the spreading curve may have fast smooth parts for a positive proportion of the volume, but the maximum proportion is bounded away from 1.

- (3) For every $\epsilon > 0$, if $\lambda > 0$ is large enough, then

$$\mathbb{P}\left(\forall x \in V(\mathcal{C}) \exists s \in V(\mathcal{C}) \text{ with } k_x(s) < \epsilon|\mathcal{C}| \text{ and } \kappa(\mathcal{C}'_x, s)/|\mathcal{C}| > 1 - \epsilon\right) > 1 - \epsilon,$$

where $k_x(s)$ is the random number of steps in which the infection from x reaches s , and \mathcal{C}'_x is the graph we obtain from \mathcal{C} by removing the vertices infected before s .

1.3 Extensions and open questions

According to Lemma 1.2, the first bottleneck appears in the place where the graph G has a bridge. It is a simple corollary that if the graph is 2-edge-connected, then it contains no temporal bottlenecks for the infection, which is specifically true for a cycle.

It is possible to extend the result of Theorem 1.1 to the case of $\alpha \in (1/d, 1)$, when $d > 2$, by defining for each finite connected graph G with a root s the number

$$\kappa_d(G, s) = \min_{e_1, \dots, e_{d-1} \in E(G)} |\mathcal{C}(s, G \setminus \{e_1, \dots, e_{d-1}\})|.$$

Then, for every $k \leq \kappa_d(G, s)$, the expected time to infect k vertices satisfies $\mathbb{E}T_k \leq Ck^{1/\alpha}$. In order to obtain this result, one would first extend Lemma 3.2 to

$$\mathbb{E}\left(\min\{X, Y_1 - t, \dots, Y_{d-1} - t\} \mid Y_1, \dots, Y_{d-1} > t\right) \asymp t^{1-\alpha},$$

which has the same order of magnitude as before. Then, in Section 3.2, one should use the delayed process with always having d active edges in the front, and the Q process with $d - 1$ always old edges, getting the bound $\mathbb{E}T_k \leq Ck^{1/\alpha}$ at the end. However, it is less clear how to extend our speed-up results. In this case, the bottleneck-free subgraphs are d -edge-connected, hence adding an extra edge and creating a cycle may not always have a speed up effect, and more sophisticated constructions would need to be introduced.

We have explained above why it is natural to consider *critical* random trees. Nevertheless, one could also condition a *subcritical* GW tree to be large, and then, depending on the offspring distribution, one would get a tree either very similar to a long path (with tiny trees hanging from it), or to a large star-graph (many spikes, all of bounded depth) [Jan12], so in fact our two introductory examples basically cover subcritical trees, already. For *supercritical* GW trees, our questions are already not interesting: on the one hand, adding a single edge to an infinite tree has certainly an insignificant effect; on the other hand, as discussed above, for supercritical graphs there is a quite universal and fast spreading behaviour [BvdHH17], where a large tail for ξ does not produce new phenomena.

About the models we have considered, one could also ask more detailed questions. Is there a *distributional scaling limit* for the process on a critical Galton-Watson tree conditioned to have depth at least N ? Time should be scaled by $N^{1/\alpha}$, as for an iid sum of these variables (FPP on a half-line). And what is the scaling limit on the line segment $[-N/2, N/2]$, started from the origin, time scaled by $N^{1/\alpha}$?

In the proof of Theorem 1.5, part (1), we are basically showing that the largest cluster of $G(n, \frac{1}{n} + \frac{\lambda}{n^{4/3}})$, viewed from a random vertex, resembles a critical GW tree conditioned to have large depth, and then we can use part (1) of Theorem 1.3. With more work, one should be able to show that the *Benjamini-Schramm limit of the largest cluster* is in fact a GW tree with Poisson(1) offspring, conditioned to be infinite. Surprisingly, we have not found a proof of this in the literature. Beyond the upper end of the critical window, the results of [DKLP11] do imply this claim.

1.4 Notation

Given two random variables X_1 and X_2 , taking values in some spaces \mathcal{X}_1 and \mathcal{X}_2 , the random variable (\hat{X}_1, \hat{X}_2) with values in $\mathcal{X}_1 \times \mathcal{X}_2$ is a *coupling* of them if \hat{X}_i has the same marginal distribution as X_i for $i = 1, 2$, i.e., if $\mathbb{P}(\hat{X}_i \in E) = \mathbb{P}(X_i \in E)$ for all measurable subsets E of \mathcal{X}_i .

If X and Y are two discrete random variables on the same space \mathcal{X} , with distributions

$$\mathbb{P}(X = x) = p_x, \quad \mathbb{P}(Y = y) = q_y, \quad x, y \in \mathcal{X},$$

then the *total variation distance* between the measures p and q is

$$d_{TV}(p, q) = \frac{1}{2} \sum_x |p_x - q_x| = \inf \left\{ \mathbb{P}(\hat{X} \neq \hat{Y}) : (\hat{X}, \hat{Y}) \text{ is a coupling of } X \text{ and } Y \right\}.$$

The second equality is usually referred to as Strassen's theorem (see e.g. [Str65] or [vdH16], p.59).

If X and Y are two real-valued random variables, we say that Y *stochastically dominates* X , denoted as $X \preceq Y$, if for every $x \in \mathbb{R}$, the following inequality holds:

$$\mathbb{P}(X \geq x) \leq \mathbb{P}(Y \geq x).$$

Stochastic domination $X \preceq Y$ implies $\mathbb{E}(X) \leq \mathbb{E}(Y)$. We use the fact [Lin99] that X stochastically dominates Y if and only if there is a coupling (\hat{X}, \hat{Y}) of X and Y such that $\mathbb{P}(\hat{X} \geq \hat{Y}) = 1$.

If f and g are real-valued functions defined on \mathbb{N} or \mathbb{R} , then $f \asymp g$ and $f = \Theta(g)$ mean that there exist constants $0 < c < C < \infty$ such that $cf(x) < g(x) < Cf(x)$ for all x . We write $f \sim g$, as $x \rightarrow x_0$, to mean that $\lim_{x \rightarrow x_0} f(x)/g(x) = 1$, and write $f = o(g)$ to mean that $\lim_{x \rightarrow x_0} f(x)/g(x) = 0$.

2 Two basic examples

2.1 Spreading on a cycle

Consider the example of SI spreading with power law weights on the cycle C_n with n vertices, with $\alpha \in (1/2, 1)$. The process on the cycle can be understood via the process on the integer line \mathbb{Z} : the time T_k of the infection of the k 'th vertex has the same distribution on \mathbb{Z} as on C_n , for every $k \leq n$. On \mathbb{Z} , we will denote by X_i the power law distributed random weight on the edge $(i-1, i)$ if $i > 0$, and on the edge $(i, i+1)$ if $i < 0$.

Proposition 2.1. *For the SI spreading process $(T_k)_{k=1}^\infty$ on \mathbb{Z} with power law weights $\alpha \in (1/2, 1)$, the expected time to infect k vertices satisfies*

$$\mathbb{E}(T_k) \asymp k^{1/\alpha},$$

where the constant factors depend only on α .

Proof. Let $S_k = \sum_{i=1}^k X_i$ and $S_k^* = \sum_{i=-1}^{-k} X_i$. Note that

$$\min\{S_{\lfloor k/2 \rfloor}, S_{\lfloor k/2 \rfloor}^*\} \leq T_{k+1} \leq \min\{S_k, S_k^*\}.$$

Then it is enough to prove that $\mathbb{E}(\min\{S_k, S_k^*\}) \asymp k^{1/\alpha}$. It is well-known (see [Dur10, Theorem 3.7.2]) that the sum S_k as $k \rightarrow \infty$ is in the domain of attraction of the stable law Y with the same parameter α :

$$\mathbb{P}(S_k/k^{1/\alpha} > t) \xrightarrow[k \rightarrow \infty]{} \mathbb{P}(Y > t).$$

Denote $\bar{S}_k = S_k/k^{1/\alpha}$. The convergence is given via the convergence of characteristic functions, where the limit characteristic function is given by [Dur10]:

$$\phi_Y(t) = \lim_{k \rightarrow \infty} \phi_{\bar{S}_k}(t) = C_{\text{sgn}(t)} \exp(-b|t|^\alpha), \quad (2.1)$$

where the constants, $C_{-1} = \bar{C}_1$ and $b > 0$, depend on α . Hence, in the bounded interval $|t| < 1$ the convergence in (2.1) is uniform in t , thus we can write

$$\phi_{\bar{S}_k}(t) = C_{\text{sgn}(t)} \exp(-b|t|^\alpha)(1 + o(1)),$$

where $o(1) \rightarrow 0$ as $k \rightarrow \infty$ uniformly in $|t| < 1$. Using the relation between the tail distribution and the characteristic function, given by the following inequality [Dur10, Eq. (3.3.1)]:

$$\mathbb{P}(|X| > 2/u) \leq \frac{1}{u} \int_{-u}^u (1 - \phi_X(t)) dt,$$

where X is a random variable with characteristic function $\phi_X(t)$, we derive that when t is sufficiently large then for all k ,

$$\mathbb{P}(\bar{S}_k > t) \leq t \int_{-2/t}^{2/t} 1 - (1 + o(1)) C_{\text{sgn}(t)} \exp(-b|x|^\alpha) dx < t \int_{-2/t}^{2/t} C_2 |x|^\alpha dx = C_3 t^{-\alpha}, \quad (2.2)$$

where $C_3 > 0$ is constant that depends on α . Thus we have for sufficiently large t :

$$\mathbb{P}(\min\{S_k, S_k^*\}/k^{1/\alpha} > t) \leq C_4 t^{-2\alpha},$$

where $C_4 > 0$ is constant that depends on α . Since S_k is positive then we can find a random variable Z with power law tail with exponent 2α such that $|\min\{S_k, S_k^*\}/k^{1/\alpha}| < Z$ a.s. for all $k > 0$, and thus by Dominated Convergence Theorem for $\alpha > 1/2$ we have convergence of expectations

$$\mathbb{E}(\min\{S_k, S_k^*\}/k^{1/\alpha}) \xrightarrow[k \rightarrow \infty]{} \mathbb{E}(\min\{Y, Y^*\}).$$

where Y, Y^* are stable with parameter α . The minimum of Y, Y^* has power law tail with exponent 2α thus has finite expectation and we have:

$$\mathbb{E}(\min\{S_k, S_k^*\}/k^{1/\alpha}) \asymp 1,$$

for all $k > 0$, which implies the statement of the proposition. \square

2.2 Spreading on a star

The star graph St_n consists of a distinguished root vertex 0 and vertices $\{1, 2, \dots, n-1\}$ attached to it. On St_n , we consider the SI spreading process $T = (T_k)_{k=1}^n$ started from the root, with power law distributed random weights with $\alpha \in (1/2, 1)$, denoted as X_1, X_2, \dots, X_{n-1} .

Proposition 2.2. *On the graph St_n with $n \geq 2$, the expected time to infect k vertices, for $k \leq n-1$, is bounded by*

$$\mathbb{E}(T_k) \leq C_\alpha k^{1/\alpha},$$

where $C_\alpha > 0$ is a constant that depends only on α . For $k = n-1$, in order to infect all but one vertices, we have

$$\mathbb{E}(T_{n-1}) \asymp n^{1/\alpha},$$

with the implicit constant factors depending on α .

Proof. Denote by $X_{(1)}^{n-1} < \dots < X_{(n-1)}^{n-1}$ the order statistics of X_1, \dots, X_{n-1} . Then we have $T_{k+1} = X_{(k)}^{n-1}$, and it is obvious that

$$X_{(k)}^{n-1} \leq X_{(k)}^{k+1},$$

for $k \leq n-2$, hence the second statement of the theorem implies the first one.

It is straightforward to calculate the tail distribution of $X_{(k)}^{k+1}$:

$$\begin{aligned} \mathbb{P}(X_{(k)}^{k+1} > t) &= 1 - \mathbb{P}(X_{(k)}^{k+1} < t) \\ &= 1 - (k+1) \mathbb{P}(X_1, \dots, X_k < t, X_{k+1} > t) - \mathbb{P}(X_1, \dots, X_{k+1} < t) \\ &= 1 - (k+1)(1-t^{-\alpha})^k t^{-\alpha} - (1-t^{-\alpha})^{k+1}. \end{aligned} \quad (2.3)$$

To get an upper bound on $\mathbb{E}(X_{(k)}^{k+1})$, we integrate (2.3) over $t > 0$, using the bound $(1-t^{-\alpha})^k > 1-kt^{-\alpha}$ for $t > k^{1/\alpha}$, and the bound $\mathbb{P}(X_{(k)}^{k+1} > t) \leq 1$ for $t < k^{1/\alpha}$. After cancellations,

$$\begin{aligned} \mathbb{E}(X_{(k)}^{k+1}) &\leq \int_0^{k^{1/\alpha}} 1 dt + (k+1)k \int_{k^{1/\alpha}}^\infty t^{-2\alpha} dt \\ &= k^{1/\alpha} + (k+1)k \frac{k^{1/\alpha-2}}{2\alpha-1} \leq C_\alpha k^{1/\alpha}. \end{aligned}$$

In order to get a lower bound, we use the bound $(1-t^{-\alpha})^k < 1-kt^{-\alpha} + \frac{1}{2}k^2t^{-2\alpha}$ for $t > k^{1/\alpha}$, and the bound $\mathbb{P}(X_{(k)}^{k+1} > t) \geq 0$ for $t < k^{1/\alpha}$:

$$\begin{aligned} \mathbb{E}(X_{(k)}^{k+1}) &\geq \left(k(k+1) - \frac{(k+1)^2}{2} \right) \int_{k^{1/\alpha}}^\infty t^{-2\alpha} dt - \frac{(k+1)k^2}{2} \int_{k^{1/\alpha}}^\infty t^{-3\alpha} dt \\ &\sim \frac{1}{2} \left(\frac{1}{2\alpha-1} - \frac{1}{3\alpha-1} \right) k^{1/\alpha} \geq c_\alpha k^{1/\alpha}, \end{aligned}$$

with some $c_\alpha > 0$. Thus,

$$\mathbb{E}(X_{(k)}^{k+1}) \asymp k^{1/\alpha},$$

which finishes proof of the theorem. \square

3 General deterministic graphs

3.1 Preliminary lemmas

We present here two lemmas that will be used in tandem in the proof of Lemma 3.5.

Lemma 3.1. *Let $(b_n)_{n=1}^\infty$ be a positive sequence that satisfies the following recursive inequality for some $C > 0$ and $0 < \alpha < 1$:*

$$b_{n+1} \leq b_n + Cb_n^{1-\alpha}. \quad (3.1)$$

Then

$$b_n \leq dn^{1/\alpha},$$

with $d = \max\{b_1, (\alpha C)^{1/\alpha}\}$.

Proof. We prove the statement by induction. By definition, the statement holds for $b_1 \leq d$. Suppose the statement holds for some $n > 1$: for any k with $1 \leq k \leq n$, we have $b_k \leq dk^{1/\alpha}$. Now rewrite (3.1) as

$$b_{k+1} - b_k \leq Cb_k^{1-\alpha}.$$

Making a telescopic sum, then using the induction hypothesis and bounding the sum with an integral, we obtain

$$\begin{aligned} b_{n+1} - b_1 &\leq \sum_{k=1}^n Cb_k^{1-\alpha} \leq \sum_{k=1}^n Cd^{1-\alpha}k^{1/\alpha-1} \leq \int_1^{n+1} Cd^{1-\alpha}x^{1/\alpha-1}dx \\ &= \alpha Cd^{1-\alpha} \left((n+1)^{1/\alpha} - 1 \right) \leq d \left((n+1)^{1/\alpha} - 1 \right). \end{aligned} \quad (3.2)$$

Adding $b_1 \leq d$ to this inequality, we arrive at $b_{n+1} \leq d(n+1)^{1/\alpha}$, as desired. \square

Lemma 3.2. *Let X and Y be i.i.d. power law distributed random variables with $\alpha \in (1/2, 1)$. Then, for any $t > 1$:*

$$\mathbb{E}(\min\{X, Y - t\} | Y > t) \asymp t^{1-\alpha},$$

with the constant factors depending on α .

Proof. The conditional tail distribution of the minimum is the following:

$$\mathbb{P}(\min\{X, Y - t\} > s | Y > t) = \frac{\mathbb{P}(X > s, Y > t + s)}{\mathbb{P}(Y > t)} = \begin{cases} s^{-\alpha} \left(1 + \frac{s}{t}\right)^{-\alpha}, & \text{if } s > 1; \\ \left(1 + \frac{s}{t}\right)^{-\alpha}, & \text{if } 0 < s < 1. \end{cases}$$

Using the substitution $u = \frac{s}{t}$ we write the expected value as follows:

$$\begin{aligned} \mathbb{E}(\min\{X, Y - t\} | Y > t) &= \int_0^\infty \mathbb{P}(\min\{X, Y - t\} > s | Y > t) ds \\ &= \int_0^1 \left(1 + \frac{s}{t}\right)^{-\alpha} ds + t^{-\alpha} \int_1^\infty \left(\frac{s}{t} \left(1 + \frac{s}{t}\right)\right)^{-\alpha} ds \\ &= t \int_0^{1/t} (1 + u)^{-\alpha} du + t^{1-\alpha} \int_{1/t}^\infty (u(1 + u))^{-\alpha} du. \end{aligned}$$

In the first integral, $1 \leq 1 + u \leq 2$, hence that integral is $\asymp 1/t$ and the term is altogether $\asymp 1$. To calculate the second term, we split the interval of integration into two parts again:

$$t^{1-\alpha} \int_{1/t}^{\infty} (u(1+u))^{-\alpha} du = t^{1-\alpha} \left[\int_{1/t}^1 (u(1+u))^{-\alpha} du + \int_1^{\infty} (u(1+u))^{-\alpha} du \right]. \quad (3.3)$$

The first integral on the RHS of (3.3) is less than $\int_0^1 u^{-\alpha} du = \frac{1}{1-\alpha}$. The second integral can be bounded in the following way:

$$\int_1^{\infty} (1+u)^{-2\alpha} du \leq \int_1^{\infty} (u(1+u))^{-\alpha} du \leq \int_1^{\infty} u^{-2\alpha} du,$$

which is clearly $\asymp \frac{1}{2\alpha-1}$. Summing up the three terms we have calculated, we have proved the lemma. \square

3.2 Model and main result

We consider finite simple connected rooted graphs $G = (V, E, s)$ on n vertices with m edges. We denote by $V(G)$ the vertex set of G and by $|G|$ the total size of it, or $|G| = |V|$. We call a path between s and t an (s, t) -path. The *graph distance* $d(s, t)$ between vertices s and t as the length of the shortest (s, t) -path.

Denote the weight, or the infection passage time, attached to the edge $e \in E$ as X_e , where X_e is defined on the probability space $(\mathbb{R}_+, \mathcal{F}, \mathbb{P})$. The probability space $\Omega = (\mathbb{R}_+^m, \mathcal{F}^m, \mathbb{P}^m)$ is the space of all possible random assignments of weights to edges of the graph G equipped with the product measure. Denote by $P(s, t)$ the shortest weighted (s, t) -path and by $|P(s, t)|$ the total weight of this path. The *SI* spreading process on the graph G is a stochastic process $T = \{T_k : k \in \{1, \dots, n\}\}$, where T_k is defined as the minimum over $H_k \in \mathcal{G}_k$ of the maximum over vertices $t \in V(H_k)$ of the total weight of the shortest weighted (s, t) -path. In symbols,

$$T_k = \min_{H_k \in \mathcal{G}_k} \max_{t \in V(H_k)} |P(s, t)|,$$

where \mathcal{G}_k is the set of subtrees of G on k vertices with the same root s . We call an edge as *occupied* at time t , if both end vertices are in the infected state I , *unoccupied*, if both are in the susceptible state S , and *active*, if one is in state I and the other is in state S .

The process T is defined on the space Ω , equipped with natural filtration $\mathcal{F} = \{\mathcal{F}_k : k \in \{1, \dots, n\}\}$. Denote the sample sequence $T(\omega) = \{T_k(\omega) : k \in \{1, \dots, n\}\}$, where $\omega \in \Omega$. Each sample sequence $T(\omega)$, $\omega \in \Omega$, defines an order $\epsilon_{T(\omega)} = (e_1^\epsilon, e_2^\epsilon, \dots, e_m^\epsilon)$ on the edge set, in which the edges are occupied by the process. Given $T_k(\omega)$ we define for each k the *infection trail tree* $G(T_k(\omega))$ as the subtree from \mathcal{G}_k that consists of those edges that successfully passed the infection, i.e., the occupied edges on the shortest paths between the root s and all $k-1$ other infected vertices. It may happen that at some time $T_k(\omega)$ two or more edges, incident to a newly infected vertex, become occupied at the same time. In this case we assume that first the edge on the shortest (weighted) path to the root is occupied, and the rest are occupied with respect to some fixed generic order $\bar{\epsilon}$ on the edge set to eliminate ambiguity. Also it may happen that some vertex has two or more shortest paths with equal total weight, but since the system is finite, this event happens with zero probability.

In the current framework, for each $\omega \in \Omega$ the spreading curve is the set of pairs $\{(\mathbb{E}(T_k), k/n) : 1 \leq k \leq n\}$. Our goal is to mathematically define the position of the first temporal bottleneck, responsible for the plateau on the spreading curve. Further we restrict ourselves to consideration of power law distributed weights with $\alpha \in (1/2, 1)$, however some lemmas consider broad class of weights.

First temporal bottleneck. Remember we call an edge active at time t if one of its incident vertices is in state S and the other is in state I . In other words, an active edge is an edge that currently transmits an infection. Let $\omega \in \Omega$ and the *front of the epidemic* $F(T_k(\omega))$ be the set of edges, that are active at time $T_k(\omega)$, where $k \in \{1, \dots, n\}$, in the sample sequence $T(\omega)$. Define $\kappa(G, s)$ to be the *maximal* number of vertices k such that for each sample sequence $T(\omega)$ and for each $j < k$, the front $F(T_j(\omega))$ has at least two active edges. In other words, it is the *minimal* k such that there exists $\omega \in \Omega$ with $F(T_k(\omega))$ having one or zero active edges. We say the active edge e is *old* and has age $\tau > 0$ at time t , if the edge has become active at time $t - \tau$. If $\tau = 0$, then an edge is called *new*. We now prove that if there is a sample sequence $T(\omega)$ and a number i , for which the front of the epidemics $F(T_i(\omega))$ has one active edge, then there is a large plateau at this point on the spreading curve.

Lemma 3.3. *Let G be a finite rooted graph and let T be the SI spreading process with absolutely continuous passage times ξ , such that $\mathbb{E}(\xi) = \infty$. Let there exists an $\omega_0 \in \Omega$ and a number $i \in \mathbb{N}$, such that the front $F(T_i(\omega_0))$ has one active edge. Then for each j , where $i+1 \leq j \leq n$, the expected infection time is*

$$\mathbb{E}(T_j) = \infty.$$

Proof. The sample sequence $T(\omega_0)$ defines the order of occupation of the edge set $\epsilon_{T(\omega)} = (e_1^\epsilon, e_2^\epsilon, \dots, e_m^\epsilon)$. Since all edge weights have absolutely continuous distribution and the number of edges is finite, there exists a subset $\mathcal{A}(\omega_0)$ of sample sequences of positive measure with the same order of occupation of edges as in $\epsilon_{T(\omega)}$. More precisely, there exists a small $\varepsilon > 0$ such that the set:

$$\mathcal{A}(\omega_0) = \{\omega : |X_{e_j^\epsilon}(\omega) - X_{e_j^\epsilon}(\omega_0)| < \varepsilon\},$$

which has positive measure for any $\varepsilon > 0$. For this, one can take ε to be smaller than half of the minimum of all the absolute differences between the edge weights of different edges (which is almost surely positive).

Then, since the front $F(T_i(\omega_0))$ has one active edge, then for each $\omega \in \mathcal{A}(\omega_0)$, the front $F(T_i(\omega))$ also has one active edge, say e . Therefore, we have

$$\mathbb{E}(T_{i+1} - T_i \mid \mathcal{A}(\omega_0)) = \mathbb{E}(X_e) = \infty,$$

and by the law of total expectation

$$\mathbb{E}(T_{i+1}) = \mathbb{E}(T_{i+1} - T_i) + \mathbb{E}(T_i) > \mathbb{E}(T_{i+1} - T_i \mid \mathcal{A}(\omega_0))\mathbb{P}(\mathcal{A}(\omega_0)) = \infty.$$

Since $T_i \leq T_{i+1}$ for all i , then we have $\mathbb{E}(T_j) = \infty$, for all $j > i + 1$, which finishes the proof of the lemma. \square

By Lemma 3.3, $\kappa(G, s)$ captures the exact number of nodes that are infected before the possible emergence of the first plateau. It has a simple combinatorial description, as we stated in Lemma 1.2.

Proof of Lemma 1.2. Suppose that there exists $\omega \in \Omega$ and the number k , where $0 < k < n$, such that $F(T_k(\omega))$ has one active edge $e \in E$. Then at time $T_k(\omega)$ we can divide vertices of G into two classes: infected (in state I) and susceptible (in state S). In the induced subgraph on infected vertices all edges are occupied, and in the subgraph on susceptible vertices all edges are unoccupied, and since the only edge is active, then it is the only one edge between these two subgraphs. Hence, the active edge e is a cut edge, and the size of the infected subgraph equals to $k_{inf} = |\mathcal{C}(s, G \setminus e)|$. Since $\kappa(G, s)$ is defined as the minimum over $\omega \in \Omega$, hence can be a bottleneck edge in the process, thus

$$\kappa(G, s) = \min_{e \in E(G)} |\mathcal{C}(s, G \setminus e)|.$$

This finishes the proof of the Lemma. \square

Delayed process \bar{T} . Fix an arbitrary order on the edge set $\epsilon = (e_1^\epsilon, e_2^\epsilon, \dots, e_m^\epsilon)$. Define the process $\bar{T} = \{\bar{T}_k : k \in \{1, \dots, n\}\}$, coupled with the original process T as follows. Start with the infected root s . Then choose the two active edges incident to s with smallest indices in the order ϵ and spread the infection through them. At the time when one of these edges becomes occupied, choose the next active edge from E with the smallest index in ϵ and repeat the procedure. If the two active edges share one susceptible vertex, then, when one edge gets occupied, choose two new active edges, given by the smallest indices in ϵ . The process runs until there are no more new active edges to take, and the remaining times are assumed to be infinite.

Obviously, for each $k < \kappa(G, s)$ and each $\omega \in \Omega$ the front $F(\bar{T}_k(\omega))$ has two active edges, since the delayed process can be turned into a particular realization of an original process with order of edge occupation ϵ .

The process \bar{T} stochastically dominates the process T , which is proved in the following lemma.

Lemma 3.4. *Let $G = (V, E, s)$ be a finite connected rooted graph and let T be the SI spreading process on G with absolutely continuous passage times ξ . Then the delayed process \bar{T} stochastically dominates the process T .*

Proof. Consider an $\omega \in \Omega$. Then for any k , where $1 \leq k \leq n$, the sample sequence of the delayed process $\bar{T}_k(\omega)$ induces the infection trail tree $G(\bar{T}_k(\omega))$ and we have

$$\bar{T}_k(\omega) = \max_{t \in V(G(\bar{T}_k(\omega)))} |P(s, t)|.$$

On the other hand, the original process is given by the minimum over all possible subtrees on k vertices:

$$T_k(\omega) = \min_{H_k \in \mathcal{G}_k} \max_{t \in V(H_k)} |P(s, t)|.$$

Therefore, we have $T_k(\omega) \leq \bar{T}_k(\omega)$ for all k , and, hence, $T \preceq \bar{T}$. \square

Despite the fact that the delayed process runs slower than the original, the next we define the process Q , which is even slower, but is necessary to achieve our final statement.

Process Q . Define the $Q = \{Q_k : k \in \{1, \dots, n\}\}$ to be the stochastic process, in which at each time Q_k there are two active edges with weights X and Y in the front: one of them is always old, with the age of the process, and an other is new. In symbols, let

$$\begin{aligned} Q_1 &= 0, \\ Q_2 &= \min\{X, Y\}, \\ Q_{k+1} &= Q_k + \min\{X, Y - Q_k | Q_k, Y > Q_k\}, \end{aligned} \tag{3.4}$$

where the unconditional X, Y are i.i.d. edge weights. The process Q qualitatively constitutes the particular scenario the infection can spread on \mathbb{Z} , having an ever old edge Y and spreading only along new edges X in one direction. The following lemma provides a bound on the expected time to infect k vertices in the process Q .

Lemma 3.5. *Consider the process Q_k defined above in (3.4) with $X, Y \sim \text{pow}(\alpha)$. Then, for $\alpha \in (1/2, 1)$ and for each k , where $k \geq 1$, we have*

$$\mathbb{E}(Q_k) \leq dk^{1/\alpha}.$$

where $d > 0$ is a constant that depends on α .

Proof. Using the law of total expectation, Lemma 3.2 and Jensen's inequality we have that

$$\begin{aligned} \mathbb{E}(Q_{k+1}) &= \mathbb{E}(Q_k + \min\{X, Y - Q_k\}) = \mathbb{E}(Q_k) + \mathbb{E}(\mathbb{E}(\min\{X, Y - Q_k\} | Q_k)) \\ &\leq \mathbb{E}(Q_k) + C\mathbb{E}(Q_k^{1-\alpha}) \\ &\leq \mathbb{E}(Q_k) + C\mathbb{E}(Q_k)^{1-\alpha}, \end{aligned}$$

where $C > 0$ is a large enough constant. Then immediately we have $\mathbb{E}(Q_k) \leq b_k$, where b_k is defined with a recursion

$$\begin{aligned} b_{k+1} &= b_k + Cb_k^{1-\alpha}, \\ b_1 &= \mathbb{E}(Q_2) = \frac{2\alpha}{2\alpha - 1} \leq (\alpha C)^{1/\alpha} := d. \end{aligned}$$

By Lemma 3.1, this sequence is bounded and we have for any $k \geq 1$:

$$\mathbb{E}(Q_k) \leq dk^{1/\alpha}.$$

This finishes the proof of the lemma. □

Define the random variable X_s with the following probability measure:

$$\mathbb{P}(X_s > t) := \mathbb{P}(X - s > t | X > s),$$

where $s > 0$. We call the random variable X to have *shifted power law* distribution $\text{shiftpow}(\alpha)$ with $\alpha > 0$, if $\mathbb{P}(X > t) = (t + 1)^{-\alpha}$, for all $t \geq 0$. Notionally, it represents a power law random variable with a shift of a size of the minimum cutoff. Note that for any $s > 0$, if the random variable $X \sim \text{shiftpow}(\alpha)$ with $\alpha > 0$, then $X_s \stackrel{d}{=} (s + 1)X$ and therefore, for any $s_1 < s_2$:

$$X_{s_1} \preceq X_{s_2}. \tag{3.5}$$

In other words, if we consider the SI spreading with shifted power law passage times, then the older edges dominate the newer ones. We can now prove the main theorem of this section:

Proof of Theorem 1.1. Let a random variable $X \sim \text{pow}(\alpha)$, then $(X - 1) \sim \text{shiftpow}(\alpha)$. Define $T^{(X-1)}$ to be the SI spreading process coupled to T , with shifted power law passage times with the same parameter α as in T . The shift in times is deterministic, hence by time T_k , the process with shifted times is faster than the original process by the cumulative shift of not more than $k - 1$:

$$T_k - T_k^{(X-1)} \leq k - 1,$$

and equality holds only if the infection trail tree is a path of length $k - 1$. The cumulative shift depends on the shape of the infection trail tree given by the process at time T_k and can therefore be non-deterministic.

Since for $X, Y \sim \text{shiftpow}(\alpha)$, then by Lemma 3.4 the delayed process $\bar{T}^{(X-1)}$ with the same weights dominates the process $T^{(X-1)}$, and we have for any $k \leq \kappa(G, s)$:

$$T_k \preceq \bar{T}_k^{(X-1)} + (k - 1).$$

Now consider the process Q with shifted power law weights denoted as $Q^{(X-1)}$. Since in this case the old edges dominate newer ones, the process $Q^{(X-1)}$ dominates the process $\bar{T}^{(X-1)}$, hence we have

$$T_k \preceq Q_k^{(X-1)} + (k - 1). \quad (3.6)$$

Since the shift is negative, then we have $Q^{(X-1)} \preceq Q$. Hence, we have for each k , where $1 \leq k \leq \kappa(G, s)$,

$$T_k \preceq Q_k + (k - 1).$$

By Lemma 3.5 we have

$$\mathbb{E}(T_k) \leq dk^{1/\alpha} + (k - 1) \leq (d + 1)k^{1/\alpha},$$

where $d > 0$ depends on α , which finishes the proof of the theorem. \square

From the statement of Lemma 1.2 we have the following simple corollary.

Corollary 3.6. *Let G be a finite 2-edge-connected rooted graph with root s and let T be the SI spreading process on G with power law weights, where $1/2 < \alpha < 1$. Then for each $k \leq n$,*

$$\mathbb{E}(T_k) \leq Ck^{1/\alpha}.$$

Thus, we established the existence of specific number of nodes $\kappa(G, s)$ for each finite connected rooted graph G , up to which a spreading curve of the SI process with power law passage times with $\alpha \in (1/2, 1)$ has no large plateaux. In case of trees, there is an analytic expression for κ . Consider a rooted tree \mathcal{T} with root s on n vertices. Let the degree of the root be equal to d_s . Denote subtrees hanging from the root as $\mathcal{T}_1, \dots, \mathcal{T}_{d_s}$ (see Figure 3.1, a). Then, by Lemma 1.2,

$$\kappa(\mathcal{T}, s) = |\mathcal{T}| - \max\{|\mathcal{T}_i| : 1 \leq i \leq d_s\}. \quad (3.7)$$

Now add an extra edge e between s and a randomly chosen vertex in the tree. Then, we obtain a graph, denoted as \mathcal{T}_{+e} , which consists of a cycle C_k of length k with N rooted trees $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_N$, where $N \geq d_s$ a.s., attached to it by edges e_1, e_2, \dots, e_N (see Figure 3.1, b). Then, by Lemma 1.2, we have

$$\kappa(\mathcal{T}_{+e}, s) = |\mathcal{T}_{+e}| - \max\{|\mathcal{T}_i| : 1 \leq i \leq N\}. \quad (3.8)$$

In the following sections we consider models of random trees, and we use the above formulas to compare how the addition of an extra edge changes κ .

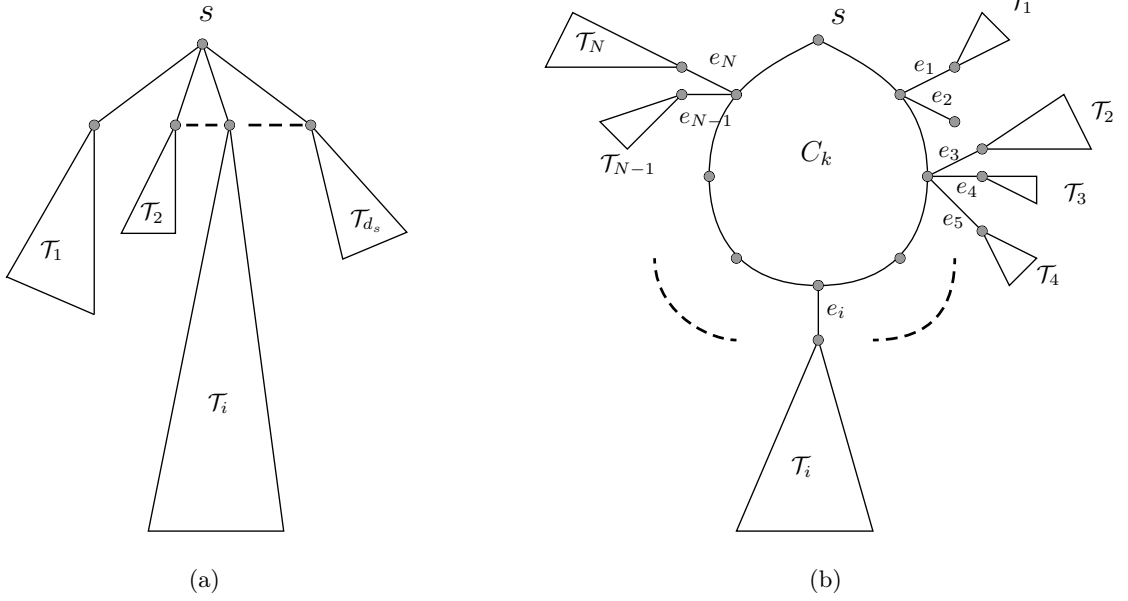


Figure 3.1: Schematic structure: a) the tree \mathcal{T} with the root s having degree d_s ; b) the graph \mathcal{T}_{+e} of a tree \mathcal{T} with an extra edge e : a cycle C_k with hanging trees $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_N$. The sizes of triangles represent the relative size of hanging trees.

4 Critical Galton-Watson trees

4.1 Preliminaries

A Galton-Watson (GW) tree \mathcal{T} is defined as a geneological tree of a *Galton-Watson process* $\langle Z_n : n \geq 0 \rangle$ of evolution of a particle system and we assume that $Z_0 = 1$ (see e.g. [Jan12] for the introduction). The process is characterized by the offspring distribution ξ , which is assumed in context of the paper to be non-degenerate, integer and have finite variance $\text{Var}(\xi) < \infty$. The process (and the tree) is called *critical*, if $\mathbb{E}(\xi) = 1$.

The GW trees can be viewed as rooted labeled trees. The root of the tree \mathcal{T} corresponds to the particle in the generation Z_0 , and it is denoted by $\langle 0 \rangle$. A generic particle in the generation Z_k is indexed as $\langle 0, l_1, \dots, l_k \rangle$, where $l_r \geq 1$, $1 \leq r \leq k$. The particles $\langle 0, l_1, \dots, l_{k-1}, j \rangle$, where $j = 1, 2, \dots$, denote the children of the particle $\langle 0, l_1, \dots, l_{k-1} \rangle$ in generation $k-1$. Of course, not for all j does $\langle 0, l_1, \dots, l_{k-1}, j \rangle$ correspond to an actual vertex of \mathcal{T} . Let $N(0, l_1, \dots, l_{k-1})$ be the number of children of $\langle 0, l_1, \dots, l_{k-1} \rangle$ in the process. Then $\langle 0, l_1, \dots, l_{k-1}, j \rangle$ is a vertex of \mathcal{T} for $1 \leq j \leq N(0, l_1, \dots, l_{k-1})$.

Denote the set of all GW trees as $\langle GW \rangle$ and a randomly chosen GW tree as \mathcal{T} . The size of \mathcal{T} is defined as the number of vertices it contains and is denoted as $|\mathcal{T}|$. It is well known that a CGW tree is almost surely finite (e.g. Theorem 3.1, p. 84, [vdH16]) and the following theorem provides a bound on the size of a CGW tree [Kol86].

Lemma 4.1. *Let \mathcal{T} be a CGW tree with integer offspring distribution ξ , such that $\text{Var}(\xi) := \sigma^2 < \infty$. Then for $n \rightarrow \infty$,*

$$\mathbb{P}(|\mathcal{T}| = n) = \frac{1}{\sqrt{2\pi\sigma}} n^{-3/2} (1 + o(1)).$$

The *height* $H(\mathcal{T})$ of a GW tree \mathcal{T} is the length of a longest path from the root or the maximum N , such that $Z_N > 0$. The following limit theorem about the height of the tree \mathcal{T} holds [KNS66].

Lemma 4.2. *Let \mathcal{T} be a CGW tree with offspring distribution ξ , such that $\text{Var}(\xi) := \sigma^2 < \infty$. Then we have,*

$$\lim_{N \rightarrow \infty} N \mathbb{P}(H(\mathcal{T}) > N) = \lim_{N \rightarrow \infty} N \mathbb{P}(Z_N > 0) = \frac{2}{\sigma^2}.$$

The following theorem provides an upper bound on the probability of having a tree of height at least N conditioned on the exact size of this tree [ABDJ13], [Kol86].

Lemma 4.3. *Let \mathcal{T} be a CGW tree with offspring distribution ξ , such that $\text{Var}(\xi) := \sigma^2 < \infty$. Then there exist positive constants C and c , such that*

$$\mathbb{P}(H(\mathcal{T}) \geq x \mid |\mathcal{T}| = n) \leq C e^{-cx^2/n}.$$

We consider the set of CGW trees *conditioned on* $Z_N > 0$, where $N > 0$, as the subset of trees $\langle GW \rangle$ with height at least N . Denote this set of conditioned CGW trees as $\langle GW \mid Z_N > 0 \rangle$. The expected limit size of the k 'th generation in such trees is given in the following Theorem [MM78].

Lemma 4.4. *Let \mathcal{T} be a CGW tree with offspring distribution ξ , such that $\text{Var}(\xi) := \sigma^2 < \infty$. Then we have:*

$$\lim_{N \rightarrow \infty} \mathbb{E}(Z_k \mid Z_N > 0) = 1 + k\sigma^2.$$

Kesten in [Kes86] proved that CGW trees, conditioned on $Z_N > 0$, converge in distribution to an infinite CGW tree \mathcal{T}^∞ , which is the geneological tree of the critical Galton-Watson process conditioned on *non-extinction*. The infinite tree has the following construction. The tree \mathcal{T}^∞ has two types of vertices: *normal* and *special*, with root being special. Normal vertices have offsprings according to independent copies of ξ , while special nodes have a number of offsprings according to the size-biased distribution $\hat{\xi}$, where

$$\mathbb{P}(\hat{\xi} = k) := kp_k,$$

and $k = 0, 1, 2, \dots$. Every offspring of a normal vertex is normal. When a special vertex produces a number of offsprings, one of its children is selected uniformly at random and becomes special, while all other children are normal.

An alternative construction of the tree \mathcal{T}^∞ is to start by taking an infinite path γ of special vertices from the root, which is called a *spine*, and then attach $\nu = \hat{\xi} - 1$ independent CGW trees at each node of the spine. Since each CGW tree is a.s. finite, it follows that \mathcal{T}^∞ a.s. has exactly one infinite path from the root, viz. the spine.

4.2 Speed of convergence to the infinite tree

We write as $(\mathcal{T}[k] = T)$ and $(\mathcal{T}^\infty[k] = T)$ the event that the first k generations of the tree \mathcal{T} and \mathcal{T}^∞ respectively match the first k generations of a given tree T . Denote by $\#T_k$ the size of k 'th generation in the tree T . The following lemma holds for the trees \mathcal{T} and \mathcal{T}^∞ [Kes86].

Lemma 4.5. *Let \mathcal{T} be a CGW tree with offspring distribution ξ . Then, for any rooted vertex-labeled tree T of at least k generations:*

$$\lim_{N \rightarrow \infty} \mathbb{P}(\mathcal{T}[k] = T \mid Z_N > 0) = \#T_k \cdot \mathbb{P}(\mathcal{T}[k] = T).$$

Then

$$\mathbb{P}(\mathcal{T}^\infty[k] = T) = \#T_k \cdot \mathbb{P}(\mathcal{T}[k] = T).$$

It is natural that as $N \rightarrow \infty$ the conditioned tree $\mathcal{T}^N := (\mathcal{T} | Z_N > 0)$ and \mathcal{T}^∞ w.h.p. start looking similar. The question now is how large (as a function of N) that similar part is. The following useful proposition answers this question. Versions of this proposition, with the tree conditioned to have exactly N vertices, have appeared in [Ker98, Theorem 5] and [Stu15, Theorem 6.5], but we have not found our exact statement in the literature, hence include a proof.

Proposition 4.6. *Let \mathcal{T}^N be a CGW tree conditioned on $Z_N > 0$ and \mathcal{T}^∞ be an infinite CGW tree. Then, as $N \rightarrow \infty$, for any $\varepsilon > 0$ there exist $\delta > 0$ and a coupling between \mathcal{T}^N and \mathcal{T}^∞ , such that*

$$\mathbb{P}(\mathcal{T}^N[\delta N] \neq \mathcal{T}^\infty[\delta N]) < \varepsilon.$$

Proof. In order to prove the statement of the lemma we show that the conditioned measure and the infinite measure are close in total variation distance. First we establish bounds on the conditioned measure. Consider a rooted tree T with height k , where $k \leq \delta N$ and $\delta > 0$ is small. Then by Bayes' formula,

$$\begin{aligned} \mathbb{P}(\mathcal{T}[k] = T | Z_N > 0) &= \frac{\mathbb{P}(Z_N > 0 | \mathcal{T}[k] = T)}{\mathbb{P}(Z_N > 0)} \mathbb{P}(\mathcal{T}[k] = T) \\ &= \frac{\mathbb{P}(Z_{N-k}^{(1)} > 0 \cup \dots \cup Z_{N-k}^{(\#T_k)} > 0)}{\mathbb{P}(Z_N > 0)} \mathbb{P}(\mathcal{T}[k] = T), \end{aligned} \quad (4.1)$$

where $Z_{N-k}^{(i)}$ denotes the $(N-k)$ 'th generation in the copy of the CGW process $Z^{(i)}$, started from a vertex at level k . By Lemma 4.2, for a large enough N there exists $\varepsilon_0 > 0$ such that,

$$\frac{2}{\sigma^2 N} (1 - \varepsilon_0) < \mathbb{P}(Z_N > 0) < \frac{2}{\sigma^2 N} (1 + \varepsilon_0). \quad (4.2)$$

Also, when $N-k$ is large enough, there exists $\varepsilon_1 > 0$, such that

$$\frac{2}{\sigma^2(N-k)} (1 - \varepsilon_1) < \mathbb{P}(Z_{N-k}^{(i)} > 0) < \frac{2}{\sigma^2(N-k)} (1 + \varepsilon_1), \quad (4.3)$$

where $1 \leq i \leq \#T_k$. In order to simplify the further calculations we take common $\varepsilon_2 := \max(\varepsilon_0, \varepsilon_1)$ instead of ε_0 and ε_1 in (4.2) and (4.3), and the conditioned measure is denoted as $\mathbb{P}_N(\cdot) := \mathbb{P}(\cdot | Z_N > 0)$.

Upper bound. We use the union bound on the right-hand side of (4.1) and together with (4.2) and (4.3), we obtain

$$\begin{aligned} \mathbb{P}_N(\mathcal{T}[k] = T) &\leq \frac{\#T_k \mathbb{P}(Z_{N-k} > 0)}{\mathbb{P}(Z_N > 0)} \mathbb{P}(\mathcal{T}[k] = T) \\ &< \frac{N}{(N-k)} \#T_k \mathbb{P}(\mathcal{T}[k] = T) \frac{1 + \varepsilon_2}{1 - \varepsilon_2}. \end{aligned} \quad (4.4)$$

Therefore, we can write that for small enough k there exists $\varepsilon_3 > 0$, such that

$$\mathbb{P}_N(\mathcal{T}[k] = T) < \frac{N}{(N-k)} \#T_k \mathbb{P}(\mathcal{T}[k] = T) (1 + \varepsilon_3). \quad (4.5)$$

Lower bound. We rewrite (4.1) using (4.5) as follows:

$$\begin{aligned} \mathbb{P}_N(\mathcal{T}[k] = T) &= \frac{1}{\mathbb{P}(Z_N > 0)} \left(1 - \mathbb{P}(Z_{N-k}^{(1)} = 0 \cap \dots \cap Z_{N-k}^{(\#T_k)} = 0) \right) \mathbb{P}(\mathcal{T}[k] = T) \\ &= \frac{1}{\mathbb{P}(Z_N > 0)} \left(1 - (1 - \mathbb{P}(Z_{N-k} > 0))^{\#T_k} \right) \mathbb{P}(\mathcal{T}[k] = T) \\ &> \frac{1}{\mathbb{P}(Z_N > 0)} \left(1 - \left(1 - \frac{2(1 - \varepsilon_1)}{\sigma^2(N-k)} \right)^{\#T_k} \right) \mathbb{P}(\mathcal{T}[k] = T). \end{aligned} \quad (4.6)$$

Since for any x , where $x > 0$, we have $1 - x < \exp(-x) < 1 - x + x^2/2$, then for any $n \geq 1$

$$1 - (1 - x)^n > 1 - \exp(-nx) > nx - \frac{(nx)^2}{2}. \quad (4.7)$$

We rewrite (4.6) using (4.7) for $x = \mathbb{P}(Z_{N-k} > 0)$ and $n = \#T_k$ as follows:

$$\mathbb{P}_N(\mathcal{T}[k] = T) > \frac{\mathbb{P}(\mathcal{T}[k] = T)}{\mathbb{P}(Z_N > 0)} \left(\#T_k \mathbb{P}(Z_{N-k} > 0) - \frac{1}{2} (\#T_k \mathbb{P}(Z_{N-k} > 0))^2 \right). \quad (4.8)$$

Now use (4.3) and we obtain:

$$\begin{aligned} \mathbb{P}_N(\mathcal{T}[k] = T) &> \frac{\mathbb{P}(\mathcal{T}[k] = T)}{\mathbb{P}(Z_N > 0)} \left(\frac{2\#T_k(1 - \varepsilon_2)}{\sigma^2(N - k)} - \frac{1}{2} \left(\frac{2\#T_k(1 + \varepsilon_2)}{\sigma^2(N - k)} \right)^2 \right) \\ &> \mathbb{P}(\mathcal{T}[k] = T) \#T_k \left(\frac{N}{N - k} \frac{1 - \varepsilon_2}{1 + \varepsilon_2} - \#T_k \frac{C_1 N}{(N - k)^2} \right), \end{aligned} \quad (4.9)$$

where $C_1 = \frac{(1 + \varepsilon_2)}{2\sigma^2} < 1/\sigma^2$. Therefore, we can write that for small enough k there exists $\varepsilon_4 > 0$ and a bounded $C_2 > 0$, that depends on σ and ε_2 , such that

$$\mathbb{P}_N(\mathcal{T}[k] = T) > \mathbb{P}(\mathcal{T}[k] = T) \#T_k \left(\frac{N}{N - k} - \#T_k \frac{C_2 N}{(N - k)^2} \right) (1 - \varepsilon_4). \quad (4.10)$$

Combining the (4.5) and (4.10), and choosing $\varepsilon_5 := \max\{\varepsilon_3, \varepsilon_4\}$, we obtain the following bounds on the probability $\mathbb{P}_N(\mathcal{T}[k] = T)$:

$$\left(\frac{N}{N - k} - \#T_k \frac{C_2 N}{(N - k)^2} \right) (1 - \varepsilon_5) \leq \frac{\mathbb{P}_N(\mathcal{T}[k] = T)}{\#T_k \mathbb{P}(\mathcal{T}[k] = T)} \leq \frac{N}{N - k} (1 + \varepsilon_5). \quad (4.11)$$

Total variation distance. Now we bound the total variation distance between conditioned and infinite measures. From the upper bound in (4.11) we obtain that when k is small enough, the following inequality holds:

$$\begin{aligned} \mathbb{P}_N(\mathcal{T}[k] = T) - \mathbb{P}(\mathcal{T}^\infty[k] = T) &\leq \left(\frac{N}{N - k} (1 + \varepsilon_5) - 1 \right) \#T_k \mathbb{P}(\mathcal{T}[k] = T) \\ &= \left(\left(\frac{N}{N - k} - 1 \right) + \frac{N}{N - k} \varepsilon_5 \right) \#T_k \mathbb{P}(\mathcal{T}[k] = T), \end{aligned}$$

and, on the other hand, from the lower bound in (4.11) we obtain

$$\begin{aligned} \mathbb{P}(\mathcal{T}^\infty[k] = T) - \mathbb{P}_N(\mathcal{T}[k] = T) &\leq \left(1 - \frac{N}{N - k} (1 - \varepsilon_5) + \#T_k \frac{C_\varepsilon N}{(N - k)^2} (1 - \varepsilon_5) \right) \\ &\quad \cdot \#T_k \mathbb{P}(\mathcal{T}[k] = T) \\ &= \left(\left(1 - \frac{N}{N - k} \right) + \frac{N}{N - k} \varepsilon_5 + \#T_k \frac{C_\varepsilon N}{(N - k)^2} (1 - \varepsilon_5) \right) \\ &\quad \cdot \#T_k \mathbb{P}(\mathcal{T}[k] = T). \end{aligned}$$

Comparing both bounds we see that all summands are positive, except of $\left(1 - \frac{N}{N - k} \right)$, thus we can inverse the sign it and derive the bound for an absolute value. Summing those bounds over all possible trees of height k , we obtain

$$\begin{aligned} \sum_T \left| \mathbb{P}_N(\mathcal{T}[k] = T) - \mathbb{P}(\mathcal{T}^\infty[k] = T) \right| &\leq \sum_T \left(\left(\frac{N}{N - k} - 1 \right) + \varepsilon_5 \frac{N}{N - k} + \#T_k \frac{C_\varepsilon N}{(N - k)^2} (1 - \varepsilon_5) \right) \\ &\quad \cdot \#T_k \mathbb{P}(\mathcal{T}[k] = T). \end{aligned} \quad (4.12)$$

From the fact that we have a measure on the set of infinite trees and from Lemma 4.4 we have

$$\begin{aligned}\sum_T \mathbb{P}(\mathcal{T}^\infty[k] = T) &= \sum_T \#T_k \mathbb{P}(\mathcal{T}[k] = T) = 1, \\ \mathbb{E}(Z_k | \mathcal{T}^\infty) &= \sum_T (\#T_k)^2 \mathbb{P}(\mathcal{T}[k] = T) = 1 + k\sigma^2.\end{aligned}$$

Therefore we can rewrite (4.12) for $k = \delta N$, when $\delta > 0$ is small, and obtain

$$\sum_T \left| \mathbb{P}_N(\mathcal{T}[\delta N] = T) - \mathbb{P}(\mathcal{T}^\infty[\delta N] = T) \right| \leq \frac{\delta}{1-\delta} + \frac{\varepsilon_5}{1-\delta} + C'_2 \delta \frac{1-\varepsilon_5}{(1-\delta)^2} + \frac{C_2}{N} \frac{1-\varepsilon_5}{(1-\delta)^2},$$

where $C'_2 = C_2\sigma^2 < 1$. Hence, for any $\varepsilon_6 > 0$ we can find large N and small $\delta > 0$, such that

$$\sum_T \left| \mathbb{P}_N(\mathcal{T}[\delta N] = T) - \mathbb{P}(\mathcal{T}^\infty[\delta N] = T) \right| \leq \varepsilon_6. \quad (4.13)$$

Denote the projection of measures \mathbb{P}_N and \mathbb{P}_∞ onto the trees with common first δN layers $\mathcal{T}[\delta N]$ as $\mathbb{P}_N|_{\delta N}$ and $\mathbb{P}_\infty|_{\delta N}$ respectively. Then, by (4.13) and the definition of the total variation distance we have

$$d_{TV}(\mathbb{P}_N|_{\delta N}, \mathbb{P}_\infty|_{\delta N}) \leq \frac{1}{2}\varepsilon_6.$$

Hence by Strassen's Theorem there exists a coupling of random variables $\mathcal{T}[\delta N]$ and $\mathcal{T}^\infty[\delta N]$, with the same d_{TV} . This finishes the proof of Proposition 4.6. \square

4.3 Proof of Theorem 1.3

Proof. (1) Consider first the infinite CGW tree \mathcal{T}^∞ . Following the construction of \mathcal{T}^∞ we denote the spine as γ and label the unconditioned CGW trees attached to the root as t_1, t_2, \dots, t_ν , and the rest of the tree as t_0 . The number of the unconditioned trees ν is represented by shifted size-biased i.i.d. random variables $\nu = \hat{\xi} - 1$. Let $\mu := \mathbb{E}\nu$, then it is straightforward to show that $\mu = \sigma^2 < \infty$. By Markov's inequality for any given $C_1 > 0$,

$$\mathbb{P}(\nu > C_1) < \frac{\sigma^2}{C_1}.$$

Hence, for any $\varepsilon_1 > 0$ there exists $C_1 > 0$, such that

$$\mathbb{P}(\nu < C_1) > 1 - \varepsilon_1.$$

Then, using the law of total probability we bound the total size of these trees in the following way:

$$\begin{aligned}\mathbb{P}\left(\sum_{i=1}^{\nu} |t_i| > C_1 K\right) &= \mathbb{P}\left(\sum_{i=1}^{\nu} |t_i| > C_1 K \mid \nu < C_1\right) \mathbb{P}(\nu < C_1) \\ &\quad + \mathbb{P}\left(\sum_{i=1}^{\nu} |t_i| > C_1 K \mid \nu > C_1\right) \mathbb{P}(\nu > C_1) \\ &< \mathbb{P}\left(\sum_{i=1}^{C_1} |t_i| > C_1 K\right) + \varepsilon_1.\end{aligned} \quad (4.14)$$

It remains to show that the total size of C_1 trees is bounded. Using the union bound we can write:

$$\mathbb{P}\left(\sum_{i=1}^{C_1} |t_i| > C_1 K\right) < \mathbb{P}(\text{at least one of } |t_i| > K) < C_1 \mathbb{P}(|t_i| > K),$$

and by Lemma 4.1 and bounding the sum with an integral, we obtain that for large K ,

$$\mathbb{P}(|t_i| > K) \leq C_2 \sum_{k=K+1}^{\infty} k^{-3/2} \leq C_2 \int_K^{\infty} x^{-3/2} dx = \frac{C_3}{\sqrt{K}}, \quad (4.15)$$

where $C_3 > 0$ is constant that depends on ξ . Hence, for any $\varepsilon_2 > 0$ and given $C_1 > 0$ there exists large K , such that

$$\mathbb{P}\left(\sum_{i=1}^{\nu} |t_i| < C_1 K\right) > 1 - \varepsilon_2. \quad (4.16)$$

Since the total size of trees t_i , where $i \geq 1$, is bounded w.h.p., then their height is bounded too. By Proposition 4.6, for large N and any $\varepsilon_3 > 0$ we can find $\delta > 0$ and a coupling of \mathcal{T}^N and \mathcal{T}^∞ , such that the tree $\mathcal{T}^N[\delta N]$ is same as $\mathcal{T}^\infty[\delta N]$ with probability at least $(1 - \varepsilon_3)$, hence by (4.16) for any $\varepsilon_4 > 0$ and large N there exists $K > 0$ such that in the conditioned tree \mathcal{T}^N :

$$\mathbb{P}_N\left(\sum_{i=1}^{\nu} |t_i| < C_1 K\right) > 1 - \varepsilon_4. \quad (4.17)$$

That is, the total size of trees t_i , where $1 \leq i \leq \nu$, is absolutely bounded. On the other hand, as $N \rightarrow \infty$, the conditioning $Z_N > 0$ ensures that the total size of \mathcal{T}^N grows to infinity, which then must be due to the tree t_0 . Thus, by formula (3.7), where $\mathcal{T}_i = t_i$ for $i \geq 0$, we have that $\kappa(\mathcal{T}^N, s)$ is exactly equal to $\sum_{i=1}^{\nu} |t_i|$. Therefore, (4.17) implies the tightness $\kappa(\mathcal{T}^N, s)$, as claimed.

(2) In order to prove the second statement, we make use of formula (3.8). We consider the fraction of the conditioned tree that w.h.p. resembles the infinite tree up to the depth δN , where $\delta > 0$, then show that the total size of the trees that hang off the first $\delta^2 N$ vertices on the spine constitutes a non-zero but small fraction of the total size of the conditioned tree w.h.p. Thus, the probability that the extra edge is attached to any of these trees is small, and hence all these trees will hang off from the cycle created, ensuring that κ is positive.

We will first consider the infinite tree \mathcal{T}^∞ , then will translate our results using Proposition 4.6. Denote as γ_k , where $k \geq 1$, the initial part of γ between the root and the vertex at depth k inclusive. Following the construction of \mathcal{T}^∞ , the number of unconditioned trees attached to γ_k is $S_k = \sum_{i=1}^k \nu_i$, where $\nu_i = \hat{\xi}_i - 1$, and $\hat{\xi}_i$ is the shifted size-biased version of the offspring distribution ξ . Label these unconditioned trees in the breadth-first order as $t_1, t_2, \dots, t_{\nu_1}, \dots, t_{S_k}$.

Let $n := \delta N$ and $n' := \delta^2 N$. Consider the $S_{n'}$ unconditioned trees that hang off $\gamma_{n'}$, and let $t_0 := \mathcal{T}^\infty \setminus (\gamma_{n'} \cup t_1 \cup \dots \cup t_{S_{n'}})$. We show that the t_i 's, where $1 \leq i \leq S_{n'}$, do not go deeper than generation n in \mathcal{T}^∞ w.h.p., or in other words, they have height at most $n - n'$. To start with, by Lemma 4.2, for any $\varepsilon'_1 > 0$, small $\delta > 0$ and i , for all large enough N we have

$$\mathbb{P}(H(t_i) > n - n') = \mathbb{P}(H(t_i) > (1 - \delta)\delta N) < \frac{2}{\sigma^2(1 - \delta)\delta N} (1 + \varepsilon'_1).$$

We will now prove that the number of unconditioned trees $S_{n'}$ is linear in n' w.h.p. Remember that $\mu = \mathbb{E}(\nu) < \infty$. Hence, by the SLLN, for any $\varepsilon'_2 > 0$ and any $0 < k'_1 < 1 < K'_1 < \infty$, if n' is large enough, then

$$\mathbb{P}(k'_1 \mu n' < S_{n'} < K'_1 \mu n') > 1 - \varepsilon'_2. \quad (4.18)$$

Hence, using the law of total probability and a union bound, we have

$$\begin{aligned} \mathbb{P}\left(\exists i \in \{1, \dots, \nu_{n'}\} : H(t_i) > n - n'\right) &< \mathbb{P}\left(\exists i \in \{1, \dots, K'_1 \mu n'\} : H(t_i) > n - n'\right) + \varepsilon'_2 \\ &< \frac{2K'_1 \mu \delta}{\sigma^2(1 - \delta)}(1 + \varepsilon'_1) + \varepsilon'_2. \end{aligned}$$

Thus, for any $\varepsilon'_3 > 0$ there exists $\delta > 0$ small enough, such that

$$\mathbb{P}(H(t_i) < n - n' \forall i \in \{1, \dots, \nu_{n'}\}) > 1 - \varepsilon'_3.$$

We will now show that the total size of $S_{n'}$ unconditioned trees is w.h.p. of order n'^2 . Indeed, for any $0 < k'_2 < K'_2 < \infty$, using (4.18) again,

$$\mathbb{P}\left(\sum_{i=1}^{S_{n'}} |t_i| > K'_2 n'^2\right) < \mathbb{P}\left(\sum_{i=1}^{K'_1 \mu n'} |t_i| > K'_2 n'^2\right) + \varepsilon'_2, \quad (4.19)$$

and similarly,

$$\mathbb{P}\left(\sum_{i=1}^{S_{n'}} |t_i| < k'_2 n'^2\right) < \mathbb{P}\left(\sum_{i=1}^{k'_1 \mu n'} |t_i| < k'_2 n'^2\right) + \varepsilon'_2. \quad (4.20)$$

Lemma 4.1 implies that $\sum |t_i|$ belongs to the domain of attraction of Lévy's stable law with exponent $1/2$; see [Dur10, Theorem 3.7.2]. Therefore, $\sum_{i=1}^{k'_1 \mu n'} |t_i| / (k'_1 \mu n')^2$ has a non-degenerate limit law that is independent of k' . It follows that if K'_2 is chosen large enough in (4.19), and k'_2 is chosen small enough in (4.20), then the first terms on the right hand sides become small. Altogether,

$$\mathbb{P}\left(\frac{1}{n'^2} \sum_{i=1}^{S_{n'}} |t_i| \in [k'_2, K'_2]\right) > 1 - \varepsilon'_4. \quad (4.21)$$

Now, by Proposition 4.6, for large N and any $\varepsilon' > 0$ we can find $\delta > 0$ and a coupling of \mathcal{T}^N and \mathcal{T}^∞ , such that the tree $\mathcal{T}^N[n]$ is same to $\mathcal{T}^\infty[n]$ with probability at least $1 - \varepsilon'$. There exists an image of γ_n in \mathcal{T}^N , and all trees that hang off $\gamma_{n'}$ do not go deeper than generation Z_n w.h.p. This and (4.21) imply that for any $\varepsilon'_5 > 0$ and large N there exist $k'_2, K'_2 > 0$ such that in the conditioned tree \mathcal{T}^N :

$$\mathbb{P}_N\left(\frac{1}{n'^2} \sum_{i=1}^{S_{n'}} |t_i| \in [k'_2, K'_2]\right) > 1 - \varepsilon'_5. \quad (4.22)$$

Now we are ready to prove that the extra edge is attached to the subtree t_0 in \mathcal{T}^N w.h.p. The probability of this event can be bounded as follows, for any $\delta_0 > 0$:

$$\begin{aligned} \mathbb{P}(e \in t_0) &> \mathbb{P}\left(e \in t_0 \mid \frac{|\mathcal{T}_{+e}^N \setminus t_0|}{|\mathcal{T}_{+e}^N|} < \delta_0\right) \mathbb{P}\left(\frac{|\mathcal{T}_{+e}^N \setminus t_0|}{|\mathcal{T}_{+e}^N|} < \delta_0\right) \\ &> (1 - \delta_0) \mathbb{P}\left(\frac{|\mathcal{T}_{+e}^N \setminus t_0|}{|\mathcal{T}_{+e}^N|} < \delta_0\right), \end{aligned} \quad (4.23)$$

where we used that the attachment is uniform.

Now notice that (4.22) implies, in particular, that for any $\varepsilon_5 > 0$ there is a small enough $\beta > 0$ such that

$$\mathbb{P}_N(|\mathcal{T}^N| > \beta N^2) > 1 - \varepsilon_5. \quad (4.24)$$

Using (4.24) and (4.21), we have

$$\mathbb{P}\left(\frac{|\mathcal{T}_{+e}^N \setminus t_0|}{|\mathcal{T}_{+e}^N|} < \frac{K'_2 \delta^4}{\beta}\right) > 1 - \varepsilon_5 - \varepsilon'_5.$$

Since δ can be chosen to be small enough, combining this with (4.23) we get that for any $\varepsilon'_6 > 0$ there exists $\delta' > 0$ such that for all large enough N ,

$$\mathbb{P}(e \in t_0) > 1 - \varepsilon'_6. \quad (4.25)$$

Now since the statement of the theorem about the fraction of $\kappa(\mathcal{T}_{+e}^N, s)$ and $|\mathcal{T}_{+e}^N|$, the large volume of the tree can decrease this fraction, thus we need to prove that the probability to have a very large tree \mathcal{T}^N on the N^2 scale is also small. Let us first write Bayes' formula:

$$\mathbb{P}_N(|\mathcal{T}^N| = M) = \mathbb{P}(|\mathcal{T}| = M \mid Z_N > 0) = \frac{\mathbb{P}(Z_N > 0 \mid |\mathcal{T}| = M)}{\mathbb{P}(Z_N > 0)} \mathbb{P}(|\mathcal{T}| = M).$$

Since the condition $Z_N > 0$ implies that $|\mathcal{T}^N| > N$, we may assume that $M > N$. Then, by Lemmas 4.1, 4.2 and 4.3, it follows that

$$\mathbb{P}_N(|\mathcal{T}^N| = M) < C_1 e^{-c_1 \frac{N^2}{M}} N M^{-3/2}. \quad (4.26)$$

Taking a large $B > 0$, we thus have:

$$\begin{aligned} \mathbb{P}(|\mathcal{T}_{+e}^N| > B N^2) &< C_1 \sum_{m > B N^2} e^{-c_1 \frac{N^2}{m}} N m^{-3/2} \\ &< C_1 \int_{B N^2}^{\infty} e^{-c_1 \frac{N^2}{m}} N m^{-3/2} dm = C_1 \int_B^{\infty} e^{-\frac{c_1}{x}} x^{-3/2} dx < C_2 B^{-1/2}. \end{aligned}$$

Hence, for any $\varepsilon_6 > 0$ we can find $B > 0$ such that

$$\mathbb{P}(|\mathcal{T}_{+e}^N| < B N^2) > 1 - \varepsilon_6. \quad (4.27)$$

Now use formula (3.8) for $\kappa(\mathcal{T}_{+e}^N, s)$. If the other endpoint of the extra edge e is in t_0 , then $\kappa(\mathcal{T}_{+e}^N, s) > \min\{|t_0|, \sum_{i=1}^{S'_n} |t_i|\}$, and if δ is chosen small enough so that $\beta > 2K'_2 \delta^4$, and the events of (4.24) and (4.21) are satisfied, then this minimum is $\sum_{i=1}^{S'_n} |t_i|$. That is, combining everything, we have

$$\mathbb{P}\left(\frac{\kappa(\mathcal{T}_{+e}^N, s)}{|\mathcal{T}_{+e}^N|} > \frac{k'_2 \delta^4}{B}\right) > 1 - \varepsilon_5 - \varepsilon'_5 - \varepsilon_6 - \varepsilon'_6.$$

This finishes proof of the second part of Theorem 1.3. \square

5 Uniform spanning tree of the complete graph

The uniform spanning tree $\text{UST}(n)$ is a uniform random element of the set of spanning trees of the complete graph K_n on n vertices (of which there are n^{n-2} , as Cayley's formula tells us). Marking a vertex s of K_n as the root, we have a uniform rooted spanning tree. This random tree model can be considered as a variation of the first one: instead of producing a "large critical tree" by conditioning a critical Galton-Watson tree to have depth larger than a given value, one could condition on having

a given volume. In particular, a critical Poisson Galton-Watson tree, conditioned to have volume n , turns out to be exactly the rooted $\text{UST}(n)$; see, e.g., [Jan12] for this equivalence.

In this section we derive the analogous result about speeding up of the SI spreading with power law weights on the rooted $\text{UST}(n)$ with an extra edge added between the root and a random vertex in the tree. This model could probably be studied using variations of the methods of the previous section, but we use a rather different approach, based on Loop-Erased Random Walks (LERW) and time-inhomogeneous Pólya urns. The use of LERW is called for by Wilson’s algorithm to construct $\text{UST}(n)$ (see [Wil96] or [LP16, Section 4.1]). The rooted $\text{UST}(n, s)$ with an extra edge is obtained using a slight variation of Wilson’s algorithm: we first add the extra edge $e = (x_0, x_1)$, then the first LERW path is run from x_0 to x_1 . This procedure does not spoil the uniformity of measure because of the symmetries of Wilson’s algorithm.

Colored Wilson’s algorithm for the construction of $\text{UST}(n)_{+e}$:

1. Enumerate the vertices as $\{x_0, x_1, \dots, x_{n-1}\}$, where $x_0 = s$, and add the edge (x_0, x_1) .
2. Run a LERW from x_1 until hitting x_0 ; the union of this walk with the existing edge will be denoted by C_{-1} .
3. Run a LERW from x_2 until hitting C_{-1} . This walk, including its endpoint, will be denoted by R_0 , while C_{-1} minus the endpoint of R_0 will be B_0 . We also set $C_0 = R_0 \cup B_0$.
4. Now, recursively, given $C_i = R_i \cup B_i$ for some $i \geq 0$, run a LERW from x_{i+3} until hitting C_i , denoting the walk, excluding its endpoint, by Δ_{i+1} . If the endpoint is in R_i , then we let $R_{i+1} = R_i \cup \Delta_{i+1}$ and $B_{i+1} = B_i$; otherwise, we let $B_{i+1} = B_i \cup \Delta_{i+1}$ and $R_{i+1} = R_i$.
5. Iterate this procedure until reaching $C_I = \{x_0, x_1, \dots, x_{n-1}\}$. In retrospect, set also $\Delta_{-1} = C_{-1} \setminus \{x_0\}$ and $\Delta_0 = C_0 \setminus C_{-1}$.

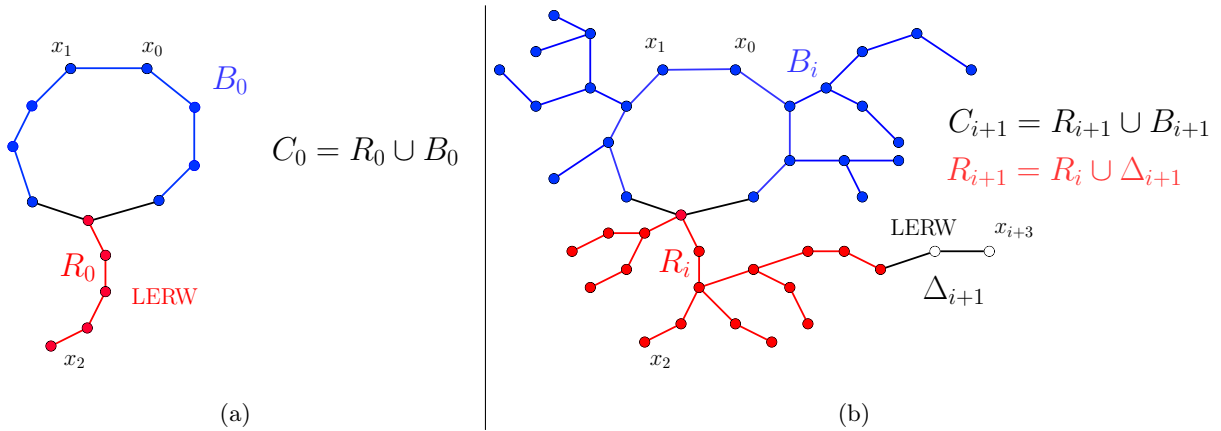


Figure 5.1: Illustration of the colored Wilson’s algorithm for constructing the rooted $\text{UST}(n, s)$. (a) Step 3, with the emergence of two colored classes. (b) Step 4, the LERW from x_{i+3} hitting R_i .

We will refer to C_i as the set of colored vertices, red in R_i , blue in B_i . We are keeping track of the colors in order to understand equation (3.8) for $\kappa(\text{UST}(n), s)$ — if the largest subtree is red, then $\kappa(\text{UST}(n), s) = |B_I|$, otherwise we have $\kappa(\text{UST}(n), s) \geq |R_I|$. Thus, our goal can be formulated as follows:

Theorem 5.1. *In the above framework of constructing a rooted $\text{UST}(n)$, the final partition (R_I, B_I) of the set $\{x_0, x_1, \dots, x_{n-1}\}$, as $n \rightarrow \infty$, is asymptotically almost surely non-trivial: for every $\varepsilon > 0$ there exists $\delta > 0$ such that*

$$\mathbb{P}\left(\frac{\kappa(\text{UST}(n), s)}{n} > \delta\right) \geq \mathbb{P}\left(|R_I|/|C_I| \in (\delta, 1 - \delta)\right) > 1 - \varepsilon,$$

for all $n > n_\varepsilon$.

Note that Δ_{i+1} will be added to R_i or B_i with probability $|R_i|/(|R_i| + |B_i|)$ and $|B_i|/(|R_i| + |B_i|)$, respectively; moreover, the entire set Δ_{i+1} is independent of whether we add it to R_i or B_i . Therefore, we can think of the process $(|R_i|, |B_i|)_{i \geq 0}$ as a Pólya urn with random, symmetric, time-inhomogeneous increments: in each round i , we add $|\Delta_i|$ balls either to the red or to the blue balls, with probabilities proportional to the current color counts. In such a Pólya urn process, when $(|\Delta_i|)_{i \geq 1}$ is uniformly bounded, the following result holds [Pem90].

Theorem 5.2. *Let the Pólya urn $(R_i, B_i)_{i \geq 0}$ has increments Δ_i bounded by some constant M for all $i \geq 1$. Let $C_i = R_i + B_i$. Then the limiting distribution of $\lim_{i \rightarrow \infty} (R_i/C_i, B_i/C_i)$ has no atoms in either coordinate.*

Corollary 5.3. *In the setting of the previous theorem, for any stopping time I for the process $(R_i, B_i)_{i \geq 0}$, the distribution of $(R_I/C_I, B_I/C_I)$ has no atoms at 0 or 1.*

Proof of the corollary. By the strong Markov property of $(|R_i|, |B_i|)_{i \geq 0}$, we can continue the urn process from (R_I, B_I) , and if $|R_I|/|C_I|$ had an atom at 0 or 1, then it would survive in the limit law of $|R_i|/|C_i|$. \square

Of course, our increments Δ_i are not at all bounded; however, it will turn out that the sequence $(|\Delta_i|/\sqrt{n})_{i \geq 1}^I$ is asymptotically almost surely bounded, i.e., for every $\epsilon > 0$ there exists $M < \infty$ such that

$$\mathbb{P}\left[|\Delta_i|/\sqrt{n} < M \text{ for all } i = 1, 2, \dots, I\right] > 1 - \epsilon, \quad (5.1)$$

for all $n > n_\epsilon$ large enough. Hence, with probability $1 - \epsilon$, we can consider $(|R_i|, |B_i|)/\sqrt{n}$ as a symmetric time-inhomogeneous Pólya urn process with bounded increments. Since the normalizations by \sqrt{n} cancel each other out, Corollary 5.3 implies that $|R_I|/|C_I|$ cannot have an atom at 0 or 1 of a uniformly positive size, as $n \rightarrow \infty$. Finally, we can take the mixture of these laws $|R_I|/|C_I|$ over the law of the initial configuration $(|R_0|, |B_0|)/\sqrt{n}$ and over the sequence $(\Delta_i/\sqrt{n})_{i \geq 1}^I$, to obtain the final result.

Proof of (5.1). First, we need to prove that the initial configuration $(|R_0|, |B_0|)/\sqrt{n}$ is non-trivial. It is a well-known result that the length of the first LERW path C_{-1} has Rayleigh distribution in the limit [CP00]. We present here the derivation taken from [PR04], since we will use a generalization of this argument. Pick two vertices x and y in the complete graph K_n , let $L = d_{\text{UST}(n)}(x, y)$ be the length of the LERW path from x to y , and denote by γ_k the first k vertices of the path itself. Because of the symmetries of K_n , we can write the following, with any specific sequence $x = x_0, x_1, \dots, x_{k-1}$ of distinct vertices that does not contain y :

$$\begin{aligned} \mathbb{P}(L = k \mid L \geq k) &= \mathbb{P}(L = k \mid \gamma_k = \{x_0, \dots, x_{k-1}\}) \\ &= \mathbb{P}\left(\text{Step}(x_{k-1}, y) \mid \text{Past}(x_0, \dots, x_{k-1}), \text{Future}(x_0, \dots, x_{k-1})\right); \end{aligned} \quad (5.2)$$

here, $\text{Past}(x_0, \dots, x_{k-1})$ is the event that the SRW from x to y is currently at x_{k-1} and its loop-erasure so far is $\{x_0, \dots, x_{k-1}\}$; the event $\text{Future}(x_0, \dots, x_{k-1})$ is that the SRW continued from x_{k-1} will not return to $\{x_0, \dots, x_{k-1}\}$ before hitting y ; the event $\text{Step}(x_{k-1}, y)$ is that the first step of the SRW continued from x_{k-1} will be to y ; and finally, the equality between the two lines is by considering the last visit of the SRW to x_{k-1} . We can then continue (5.2) as follows:

$$\frac{\mathbb{P}(\text{Step}(x_{k-1}, y) \mid \text{Past}(x_0, \dots, x_{k-1}))}{\mathbb{P}(\text{Future}(x_0, \dots, x_{k-1}) \mid \text{Past}(x_0, \dots, x_{k-1}))} = \frac{\frac{1}{n-1}}{\frac{1}{n-1} + \frac{n-k-1}{n-1} \frac{1}{k+1}} = \frac{k+1}{n}, \quad (5.3)$$

where the second term in the denominator corresponds to first stepping to some vertex not in $\{x_0, \dots, x_{k-1}, y\}$, then hitting y before $\{x_0, \dots, x_{k-1}\}$. Now we can write a telescopic product:

$$\begin{aligned} \mathbb{P}(L = k) &= \mathbb{P}(L = k \mid L \geq k) \cdot \mathbb{P}(L > k-1 \mid L \geq k-1) \cdots \mathbb{P}(L > 1 \mid L \geq 1) \\ &= \frac{k+1}{n} \prod_{j=1}^{k-1} \left(1 - \frac{j+1}{n}\right). \end{aligned} \quad (5.4)$$

In order to understand the size of (5.4), note that for all $0 < x < 1/2$ we have $\exp(-x - x^2) < 1 - x < \exp(-x)$, and thus, for any $t > 0$ and n large enough,

$$\exp\left(-\frac{\lfloor t\sqrt{n} \rfloor^2}{2n} - \frac{\lfloor t\sqrt{n} \rfloor^3}{3n^2}\right) < \prod_{j=1}^{\lfloor t\sqrt{n} \rfloor} \left(1 - \frac{j+1}{n}\right) < \exp\left(-\frac{\lfloor t\sqrt{n} \rfloor^2}{2n}\right).$$

Therefore, (5.4) gives, as $n \rightarrow \infty$,

$$\mathbb{P}(L = \lfloor t\sqrt{n} \rfloor) \sim \exp\left(-\frac{t^2}{2}\right) \frac{t}{\sqrt{n}},$$

which proves the Rayleigh limit law for L/\sqrt{n} .

In the further steps of the algorithm, the LERW hits the set C_i instead of a single vertex, thus using similar arguments we can obtain the conditional distribution of the size of Δ_i , for each $i \geq -1$:

$$\mathbb{P}(|\Delta_{i+1}| = k \mid |C_i|) = \frac{k + |C_i|}{n} \prod_{j=1}^{k-1} \left(1 - \frac{j + |C_i|}{n}\right). \quad (5.5)$$

In particular, using the same calculus estimate as before,

$$\mathbb{P}\left(R_0 = \lfloor r\sqrt{n} \rfloor \mid B_0 = \lfloor b\sqrt{n} \rfloor\right) \sim \frac{r+b}{\sqrt{n}} \exp\left(-\frac{(r+b)^2 - b^2}{2}\right).$$

That is, $(R_0, B_0)/\sqrt{n}$ indeed has an absolutely continuous limit law on $[0, \infty)^2$.

Continuing in the same way, assuming $|C_i| < n/4$ and $t\sqrt{n} < n/4$ in order for the calculus estimate to work, we have

$$\mathbb{P}\left(|\Delta_{i+1}| = \lfloor t\sqrt{n} \rfloor \mid |C_i|\right) \sim \left(\frac{t}{\sqrt{n}} + \frac{|C_i|}{n}\right) \exp\left(-\frac{t^2}{2} - \frac{t|C_i|}{\sqrt{n}}\right).$$

Summing up, this implies for $|C_i| < n/4$ that

$$\mathbb{P}\left(|\Delta_{i+1}| > t\sqrt{n} \mid |C_i|\right) \asymp \exp\left(-\frac{t^2}{2} - \frac{t|C_i|}{\sqrt{n}}\right). \quad (5.6)$$

For $|C_i| \geq n/4$, equation (5.5) implies an exponential tail for $|\Delta_{i+1}|$ directly. We will also use a slightly different version of (5.6):

$$\mathbb{P}\left(|\Delta_{i+1}| \geq t \frac{n}{|C_i|} \mid |C_i|\right) \asymp \exp\left(-t - t^2 \frac{n}{2|C_i|^2}\right). \quad (5.7)$$

Now, in order to prove boundedness of $(|\Delta_i|/\sqrt{n})_{i \geq 1}$, for all the steps till I , we need the probability $\mathbb{P}(|\Delta_{i+1}| > t\sqrt{n})$ to decay fast as i increases. In light of (5.6), this is equivalent to $|C_i|$ increasing steadily. In fact, we will prove that there exist constants $a, A > 0$ such that

$$\mathbb{P}(|C_i|^2 < ani) < \exp(-Ai), \quad (5.8)$$

where, moreover, $A \rightarrow \infty$ as $a \rightarrow 0$. Assuming this, using (5.6),

$$\begin{aligned} \mathbb{P}(|\Delta_{i+1}| > t\sqrt{n}) &= \mathbb{P}(|\Delta_{i+1}| > t\sqrt{n}, |C_i|^2 \geq ani) + \mathbb{P}(|\Delta_{i+1}| > t\sqrt{n}, |C_i|^2 < ani) \\ &< \mathbb{P}\left(|\Delta_{i+1}| > t\sqrt{n} \mid |C_i|^2 \geq ani\right) + \mathbb{P}(|C_i|^2 < ani) \\ &< C \exp(-t\sqrt{ai}) + \exp(-Ai). \end{aligned} \quad (5.9)$$

Given any $\epsilon > 0$, if we take first A large (by taking a small a), then t large, we get

$$\sum_{i=1}^{\infty} \left(C \exp(-t\sqrt{ai}) + \exp(-Ai) \right) < \epsilon,$$

and (5.1) is proved.

Now, in order to prove (5.8), consider the increments

$$X_{i+1} := |C_{i+1}|^2 - |C_i|^2 = 2|C_i||\Delta_{i+1}| + |\Delta_{i+1}|^2, \quad (5.10)$$

of the process $(|C_i|^2)_{i \geq 1}$. If we prove that, conditioned on an arbitrary C_i , there exist constants $b_1, b_2 > 0$ such that

$$\mathbb{P}(X_{i+1} < b_1 n \mid C_i) < b_2, \quad (5.11)$$

with $b_2 \rightarrow 0$ as $b_1 \rightarrow 0$, then (5.8) will follow by a simple argument: (5.11) implies that the random variable $I := \#\{1 \leq j \leq i : X_{j+1} > b_1 n\}$ stochastically dominates a Binom($i, 1 - b_2$) variable, thus for any $A < \infty$, if $b_2 > 0$ is small enough, then $\mathbb{P}(I > i/2) \geq 1 - \exp(-Ai)$, and on the event $\{I > i/2\}$ we have $|C_i|^2 > b_1 ni/2$, implying (5.8).

To start the proof of (5.11), we can assume that $C_i < n/4$, since otherwise $C_{i+1} \geq C_i + 1$ implies $X_{i+1} \geq n/2$ deterministically. Now, if $|C_i| < \sqrt{n}$, then (5.6) implies that $\mathbb{P}(|\Delta_{i+1}| > b_1 \sqrt{n} \mid C_i) \rightarrow 1$ as $b_1 \rightarrow 0$, and hence the term $|\Delta_{i+1}|^2$ in (5.10) ensures that (5.11) holds. On the other hand, if $\sqrt{n} \leq |C_i| < n/4$, then (5.7) implies that $\mathbb{P}(|C_i||\Delta_{i+1}| > b_1 \sqrt{n} \mid C_i) \rightarrow 1$ as $b_1 \rightarrow 0$, and the term $|C_i||\Delta_{i+1}|$ in (5.10) ensures that (5.11) holds. That is, in all cases we are done. \square

6 Critical versus near-critical Erdős-Rényi graphs

We will consider the Erdős-Rényi graph $G(n, p)$ in its critical window for the emergence of a giant cluster, at $p = p_n(\lambda) = 1/n + \lambda/n^{4/3}$, with $\lambda \in (-\infty, \infty)$. It is always helpful to use the standard monotone coupling of these random graphs: the edges of the complete graph K_n have i.i.d. $U_e \sim \text{Unif}[0, 1]$ labels, and then the graph with edge set $\{e \in E(K_n) : U_e \leq p\}$ is distributed like $G(n, p)$. This way, we can think of raising p or λ as adding edges to the graph.

It is well-known that the cluster of a typical vertex in $G(n, p_n(\lambda))$ in the critical case $\lambda = 0$ is locally a GW tree with Poisson offspring distribution with mean 1, while more globally, the largest clusters have size of order $n^{2/3}$ (see [Pet17, Section 12.3] for an outline of the proof and references). As we are raising λ , extra edges appear in the standard coupling; since the number of possible edges in a cluster of size $n^{2/3}$ is of order $n^{4/3}$, the number of extra edges in each large component is approximately Poisson ($f(\lambda)$) for some function f , while, of course, the extra edges also merge some of these components. That is, the large scale structure of large critical versus near-critical clusters resembles but does not exactly coincide with our first example: a critical random tree conditioned to be large, plus a constant number of random edges.

One result making this picture more precise is that the probability that the largest cluster $\mathcal{C}_1^n(\lambda)$ of $G(n, p_n(\lambda))$ is a tree converges to some $r_0(\lambda) > 0$, which decays rapidly as $\lambda \rightarrow \infty$ (see [JKLP93]). Moreover, conditioned on the largest cluster being a tree on N vertices, it is clearly a uniform random tree $\text{UST}(N)$. This means that, for $\lambda = 0$, with a decent positive probability the SI spreading on G will encounter bottlenecks everywhere during the process, and, because of these bottlenecks at random locations, the averaged spreading curve will not converge and will produce jumps. However, at large $\lambda > 0$, a typical largest cluster will have one or more extra edges, hence, on a typical realization of the cluster, we expect to see the smoothing effect. The following results from our numerical simulations show that, even at a moderately off-critical value $\lambda = 5$, the smoothing effect is typically quite severe (see Figure 6.1).

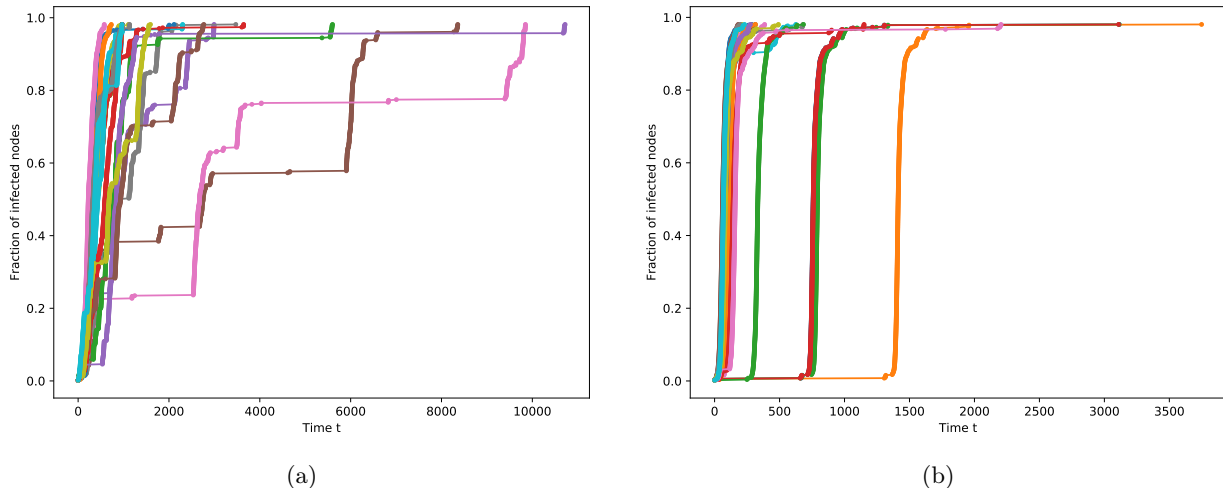


Figure 6.1: Simulation of SI spreading with power law inter-event times with $\alpha = 0.8$ on the largest component of a near-critical Erdős-Rényi graph with n vertices. The edge weights are fixed, 20 runs are shown with random starting vertices. (a) $\lambda = 0$, $n = 6000$ and the component has 541 vertices, no surplus edges. (b) $\lambda = 5$, $n = 1000$ the component has 515 vertices, 27 surplus edges.

Of course, in order to apply our Theorem 1.1, we need to know the structure of G with its extra edges, to prove that $\kappa(G, s)$ is typically positive when λ is large. A seminal result of Aldous [Ald97] says that the decreasingly ordered vector $\mathbf{C}^n(\lambda) := (|\mathcal{C}_1^n(\lambda)|, |\mathcal{C}_2^n(\lambda)|, \dots)$ of the sizes of the clusters of $G(n, p_n(\lambda))$, together with the vector $\mathbf{S}^n(\lambda) := (S_1^n(\lambda), S_2^n(\lambda), \dots)$ of the number of surplus edges in each $\mathcal{C}_i^n(\lambda)$, meaning the number of edges additional to the minimum possible value $|\mathcal{C}_i^n(\lambda)| - 1$, has a joint scaling limit:

$$\left(\frac{\mathbf{C}^n(\lambda)}{n^{2/3}}, \mathbf{S}^n(\lambda) \right) \xrightarrow{d} \left(\mathbf{C}^\infty(\lambda), \mathbf{S}^\infty(\lambda) \right), \quad (6.1)$$

where the distribution of the limiting vector $\mathbf{C}^\infty(\lambda)$ is given by the excursion lengths (in decreasing

order) of a one-dimensional Brownian motion with a parabolic drift, $B_t + \lambda t - t^2/2$, above its running minimum, and each surplus is given by the arrivals of a rate one planar Poisson point process inside the excursion. This implies, in particular, that there is a limiting distribution for the number of surplus edges in the largest cluster, $\mathbb{P}(S_1^n(\lambda) = k) \rightarrow \mathbb{P}(S_1^\infty(\lambda) = k) = r_k(\lambda)$, and it can be described using Brownian computations. For instance, for large $\lambda > 0$, it is not hard to show that the Brownian perturbation around the parabola is unlikely to ruin completely the excursion of area $\Theta(\lambda^3)$ given by the positive part of the parabola, and therefore

$$\mathbb{E}S_1^\infty(\lambda) \asymp \lambda^3 \quad \text{and} \quad S_1^\infty(\lambda) \xrightarrow{\mathbb{P}} \infty, \quad \text{as } \lambda \rightarrow \infty. \quad (6.2)$$

For a proof via counting arguments, see [JKLP93].

Building on the above [Ald97], the metric structure of large near-critical clusters was described in [ABBG10a, ABBG10b]. To start with, the reader should recall the notion of Aldous' *Brownian Continuum Random Tree* [Ald91a, Ald91b, Ald93], abbreviated as CRT from now on, which is a random real metric tree built from the Brownian Excursion measure, together with a mass measure that is supported on the leaves of the tree. It is the *Gromov-Hausdorff-Prokhorov limit* of critical GW trees with finite variance offspring distribution, conditioned to have volume N , edge lengths scaled by \sqrt{N} , with the mass measure coming from the uniform distribution on the vertices. (See, e.g., [Ald91b, LG05, ADH13] for the definitions and overviews.) Then it was proved in [ABBG10a] that the vector of metric spaces $(\mathcal{C}_1^n(\lambda), \mathcal{C}_2^n(\lambda), \dots)$ converges to some $(\mathcal{C}_1^\infty(\lambda), \mathcal{C}_2^\infty(\lambda), \dots)$ in the Gromov-Hausdorff topology extended to vectors in a suitable way. For the limit, the following construction was given in [ABBG10b, Procedure 1].

Construction. Sample the sizes (C_1, C_2, \dots) and surpluses (S_1, S_2, \dots) for $(\mathcal{C}_1^\infty(\lambda), \mathcal{C}_2^\infty(\lambda), \dots)$ according to the limit distribution in (6.1). In particular, S_1 is given by the distribution $r_k(\lambda)$ mentioned before. Conditionally on these data, the coordinates $\mathcal{C}_i^\infty(\lambda)$ will be independent from each other, and the distributions depend on the size C_i only through a scaling. So, it is enough to describe $\mathcal{C}_i^\infty(\lambda)$ conditionally on the scaled volume being 1 and the surplus being some $k \in \mathbb{N}$.

- If $k = 0$ then let the component simply be a Brownian CRT of total mass 1.
- If $k = 1$ then let (X_1, X_2) be a Dirichlet($\frac{1}{2}, \frac{1}{2}$) random vector, let $\mathcal{T}_1, \mathcal{T}_2$ be independent Brownian CRT's of sizes X_1 and X_2 , and identify the root of \mathcal{T}_1 with a uniform leaf of \mathcal{T}_1 and with the root of \mathcal{T}_2 , to make a "lollipop" shape. (For the definition of the Dirichlet distributions, see [ABBG10b, Section 3.1] or Wikipedia.)
- If $k \geq 2$, let K be a random 3-regular multigraph with $2(k-1)$ vertices chosen from a finite list, according to the following probability measure:

$$\mu_k(K) := \frac{1}{Z_k} \left(2^{t(K)} \prod_{e \in E(K)} \text{mult}(e)! \right)^{-1}, \quad (6.3)$$

where $t(K)$ is the number of loops in the multigraph K , $\text{mult}(e)$ is the multiplicity of the edge e , and Z_k is the normalization factor to get a probability measure. This K will be called the *kernel* of the cluster. Then:

1. Order the edges of K arbitrarily as $e_1, \dots, e_{3(k-1)}$, with $e_i = u_i v_i$.
2. Let $(X_1, \dots, X_{3(k-1)})$ be a Dirichlet($\frac{1}{2}, \dots, \frac{1}{2}$) random vector.

3. Let $\mathcal{T}_1, \dots, \mathcal{T}_{3(k-1)}$ be independent Brownian CRT's, with tree \mathcal{T}_i having mass X_i , and for each i let r_i and s_i be the root and a uniform leaf of \mathcal{T}_i .
4. Form the component by replacing edge $u_i v_i$ with tree \mathcal{T}_i , identifying r_i with u_i and s_i with v_i , for $i = 1, \dots, 3(k-1)$.

We will now pull back this information on the scaling limit to the discrete graphs:

Proof of Theorem 1.5. (1) As a warm-up, note that the weaker result

$$\kappa(\mathcal{C}_1^n, \sigma) / |\mathcal{C}_1^n| \xrightarrow{\mathbb{P}} 0 \quad (6.4)$$

follows easily from the scaling limit description. It is well-known that the mass measure of the CRT is fully supported on the leaves; that is, a uniform random vertex in a discrete random tree on N vertices that converges to the CRT will have, with probability tending to 1, only a single macroscopic path emanating from it, i.e., going to distance of order \sqrt{N} . By the mass measure having no atoms in the scaling limit, pieces of vanishing diameter have vanishing volume in the discrete tree, and hence by cutting the one macroscopic path, we can separate most of the tree from the uniform vertex. Together with the construction above that glues finitely many CRT's, we obtain (6.4).

In order to prove the stronger statement, knowledge about the scaling limit will not suffice, since that cannot distinguish between different distances of size $o(\sqrt{|\mathcal{C}_1^n|})$ and between different volumes of size $o(|\mathcal{C}_1^n|)$. That is, we need to go back to the discrete model, and the breadth-first search exploration process used in [Ald97].

Let \mathcal{B} be the event that the largest cluster is the cluster \mathcal{C} of vertex 1. By symmetry, it is enough to prove that $\mathbb{P}(\kappa(\mathcal{C}, 1) > K \mid \mathcal{B}) < \epsilon$ if $K > K_0(\epsilon)$ is large enough, for all $n > n_0(\lambda, \epsilon)$. Furthermore, let \mathcal{A}_δ be the event that the exploration process of the cluster of 1 finds that \mathcal{C} has size at least $\delta n^{2/3}$. Now observe the following two facts:

- For any $\delta > 0$, there exists $\tilde{\delta} > 0$, tending to 0 with δ , such that, for all large enough n ,

$$\mathbb{P}(\mathcal{B} \mid \mathcal{A}_\delta) > \tilde{\delta} \quad \text{and} \quad \mathbb{P}(\mathcal{A}_\delta \mid \mathcal{B}) > 1 - \tilde{\delta}.$$

These follow easily from [Ald97]. Indeed, for the first inequality, the event {all clusters other than \mathcal{C} are smaller than $\delta n^{2/3}$ } has probability at least some $\tilde{\delta}$, it is positively correlated with \mathcal{A}_δ , and their intersection implies \mathcal{B} . The second inequality follows from the size of the largest cluster having a scaling limit with no atom at 0.

- For the exploration process, given \mathcal{A}_δ with any $\delta > 0$, the proof of Theorem 1.3 (1) applies, hence, for any $\eta > 0$, if $K > K_0(\delta, \eta)$ is large enough, then

$$\mathbb{P}(\kappa(\mathcal{C}, 1) > K \mid \mathcal{A}_\delta) < \eta.$$

Now, combining these two facts, with $\eta = \tilde{\delta}^2$ and writing \mathcal{R}_K for the event $\{\kappa(\mathcal{C}, 1) > K\}$,

$$\begin{aligned} \mathbb{P}(\mathcal{R}_K \mid \mathcal{B}) &= \frac{\mathbb{P}(\mathcal{R}_K \cap \mathcal{B})}{\mathbb{P}(\mathcal{B})} \leq \frac{\mathbb{P}(\mathcal{R}_K \cap \mathcal{A}_\delta) + \mathbb{P}(\mathcal{B} \setminus \mathcal{A}_\delta)}{\mathbb{P}(\mathcal{B})} \\ &\leq \frac{\mathbb{P}(\mathcal{R}_K \mid \mathcal{A}_\delta) \mathbb{P}(\mathcal{A}_\delta)}{\mathbb{P}(\mathcal{B} \cap \mathcal{A}_\delta)} + \frac{\mathbb{P}(\mathcal{B} \setminus \mathcal{A}_\delta)}{\mathbb{P}(\mathcal{B})} \\ &= \frac{\mathbb{P}(\mathcal{R}_K \mid \mathcal{A}_\delta)}{\mathbb{P}(\mathcal{B} \mid \mathcal{A}_\delta)} + \mathbb{P}(\mathcal{A}_\delta^c \mid \mathcal{B}) < \frac{\tilde{\delta}^2}{\tilde{\delta}} + \tilde{\delta} = 2\tilde{\delta}. \end{aligned}$$

Since $\delta > 0$ can be chosen so small that $2\tilde{\delta} < \epsilon$, this finishes the proof of part (1).

(2) For the lower bound, \mathcal{C} always contains a spanning tree with the structure of the $k = 0$ case, converging to Aldous' CRT. In this limit, removing any vertex whose distance from the set of leaves is positive breaks the tree into pieces of positive mass measure (by the self-similar nature of the CRT). Translating this to the discrete tree, of volume N , with probability close to 1, there exist vertices that have at least two disjoint paths of length $\delta\sqrt{N}$ if δ is small enough, and the resulting subtrees have volume $\delta_0 N$.

For the upper bound, just notice that the surplus is at most some $k = k(\lambda, \epsilon)$ with probability at least $1 - \epsilon$, then we break the total volume N into at most $3k - 3$ pieces, hence the largest of these pieces has volume at least $N/(3k - 3)$. Regardless of how this piece fits in the kernel, by the structure of the CRT, with probability close to 1 if δ is small enough, it has a subtree of volume at least $\delta N/(3k - 3)$ that can be separated from \mathcal{C} by a single edge. This implies the claimed bound $(1 - \delta_1(\lambda, \epsilon))N$.

(3) The intuition for the proof is that the macroscopic structure of the largest cluster is more and more determined by the kernel as $\lambda \rightarrow \infty$, and this kernel more and more becomes an expander graph. This is similar to the strategy used in [BKW14] to understand the mixing time of random walk on the giant cluster in a supercritical Erdős-Rényi graph. We will need two lemmas:

Lemma 6.1. *If (X_1, \dots, X_n) is a Dirichlet($\frac{1}{2}, \dots, \frac{1}{2}$) random vector, then $\max_{1 \leq i \leq n} X_i$ converges to 0 in probability, as $n \rightarrow \infty$.*

Proof. The marginal distribution of each X_i is Beta($\frac{1}{2}, \frac{n-1}{2}$). It is well-known and easy to calculate that $\mathbb{E}(\text{Beta}(\frac{1}{2}, \frac{n-1}{2})^2) \sim \frac{3}{n^2}$. Thus, Markov's inequality for X_i^2 and a union bound give

$$\mathbb{P}\left(\max_{1 \leq i \leq n} X_i > \frac{1}{n^{1/4}}\right) \leq n \mathbb{P}(X_i^2 > \frac{1}{n^{1/2}}) \leq n \frac{(3 + o(1))\sqrt{n}}{n^2} \rightarrow 0,$$

as desired. □

Lemma 6.2. *For any $\epsilon > 0$ fixed, if K is sampled according to μ_k , then the probability that K has a cut-edge that has at least ϵk vertices on either side goes to 0 as $k \rightarrow \infty$.*

Proof. As explained, e.g., in [Wor99], the weights in (6.3) mean that K can be generated by the model of taking $3(k - 1)$ independent random pairs on the set of $2(k - 1)$ vertices, conditioned on 3-regularity. Then the results of [BKW14, Section 5] apply, meaning that K is an expander with probability tending to 1, as $k \rightarrow \infty$. □

The combination of (6.2) with Lemma 6.2 gives us that the kernel has no “large-scale cut-edges” with high probability as $\lambda \rightarrow \infty$. Combining also with Lemma 6.1, we get that there is a 2-connected core of \mathcal{C} from which all subgraphs hanging off have small volume. This immediately implies the statement. □

Acknowledgements

We are indebted to János Kertész for drawing our attention to the phenomenon studied in the paper, for many useful discussions, and constant support; to Júlia Komjáthy for suggesting that we look at near-critical Erdős-Rényi graphs; to Balázs Ráth for many comments and discussions regarding the manuscript; and to Louigi Addario-Berry for some references.

Most of the work was done while AM was a PhD student at the Central European University, Budapest. AM also gratefully acknowledges financial support from the European Research Council grant “Limits of discrete structures”, 617747, ARC (Federation Wallonia-Brussels) project "Big Data Models" and from Grant 16-01-00499 of the Russian Foundation of Basic Research. GP acknowledges support from the Hungarian National Research, Development, and Innovation Office, NKFIH grant K109684, and from an MTA Rényi Institute “Lendület” Research Group.

References

- [ABBG10a] L. Addario-Berry, N. Broutin, and C. Goldschmidt. The continuum limit of critical random graphs. *Probability Theory and Related Fields*, 152(3):367–406, 2010.
- [ABBG10b] Louigi Addario-Berry, Nicolas Broutin, and Christina Goldschmidt. Critical random graphs: Limiting constructions and distributional properties. *Electron. J. Probab.*, 15:no. 25, 741–775, 2010.
- [ABDJ13] Louigi Addario-Berry, Luc Devroye, and Svante Janson. Sub-Gaussian tail bounds for the width and height of conditioned Galton-Watson trees. *Ann. Probab.*, 41(2):1072–1087, 2013.
- [ADH13] Romain Abraham, Jean-François Delmas, and Patrick Hoscheit. A note on the Gromov-Hausdorff-Prokhorov distance between (locally) compact metric measure spaces. *Electronic Journal of Probability*, 18, 2013.
- [AHD15] Antonio Auffinger, Jack Hanson, and Michael Damron. 50 years of first passage percolation. *arXiv preprint arXiv:1511.03262*, 2015.
- [AL07] David Aldous and Russell Lyons. Processes on unimodular random networks. *Electron. J. Probab.*, 12(54):1454–1508, 2007.
- [Ald91a] David Aldous. The continuum random tree I. *The Annals of Probability*, pages 1–28, 1991.
- [Ald91b] David Aldous. The continuum random tree II: an overview. *Stochastic analysis*, 167:23–70, 1991.
- [Ald93] David Aldous. The continuum random tree III. *The Annals of Probability*, pages 248–289, 1993.
- [Ald97] David Aldous. Brownian excursions, critical random graphs and the multiplicative coalescent. *The Annals of Probability*, 25(2):812–854, 1997.
- [BGD⁺17] Ayan Kumar Bhowmick, Martin Gueuning, Jean-Charles Delvenne, Renaud Lambiotte, and Bivas Mitra. Temporal pattern of (re)tweets reveal cascade migration. In *Proceedings of IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining ASONAM 2017, Sydney, Australia*, 2017. to appear.
- [BKW14] Itai Benjamini, Gady Kozma, and Nicholas Wormald. The mixing time of the giant component of a random graph. *Random Structures & Algorithms*, 45(3):383–407, 2014.

- [BS01] Itai Benjamini and Oded Schramm. Recurrence of distributional limits of finite planar graphs. *Electron. J. Probab.*, 6:no. 23, 1–13, 2001.
- [BvdHH10] Shankar Bhamidi, Remco van der Hofstad, and Gerard Hooghiemstra. First passage percolation on random graphs with finite mean degrees. *Ann. Appl. Probab.*, 20(5):1907–1965, 10 2010.
- [BvdHH11] Shankar Bhamidi, Remco van der Hofstad, and Gerard Hooghiemstra. First passage percolation on the Erdős–Rényi random graph. *Combinatorics, Probability and Computing*, 20(5):683–707, 2011.
- [BvdHH17] S Bhamidi, R van der Hofstad, and G Hooghiemstra. Universality for first passage percolation on sparse random graphs. *Ann. Probab.*, to appear, 2017+. arXiv:1210.6839.
- [BvdHK14] Shankar Bhamidi, Remco van der Hofstad, and Júlia Komjáthy. The front of the epidemic spread and first passage percolation. *Journal of Applied Probability*, 51A:101–121, 12 2014.
- [BvdHK17] Enrico Baroni, Remco van der Hofstad, and Júlia Komjáthy. Nonuniversality of weighted random graphs with infinite variance degree. *Journal of Applied Probability*, 54(1):146–164, 2017.
- [CP00] Michael Camarri and Jim Pitman. Limit distributions and random trees derived from the birthday problem with unequal probabilities. *Electron. J. Probab.*, 5(2):1–18, 2000.
- [DKLP11] Jian Ding, Jeong Han Kim, Eyal Lubetzky, and Yuval Peres. Anatomy of a young giant component in the random graph. *Random Struct. Alg.*, 39:139–178, 2011.
- [Dur10] Rick Durrett. *Probability: Theory and examples*. Cambridge University Press, 4th edition, 2010.
- [FP93] James Allen Fill and Robin Pemantle. Percolation, first-passage percolation and covering times for Richardson’s model on the n-cube. *Ann. Appl. Probab.*, 3(2):593–629, 05 1993.
- [GCSdA⁺17] Yerali Gandica, Joao Carvalho, Fernando Sampaio dos Aidos, Renaud Lambiotte, and Timoteo Carletti. Stationarity of the inter-event power-law distributions. *PLOS ONE*, 12(3):1–10, 03 2017.
- [HK14] Dávid X Horváth and János Kertész. Spreading dynamics on networks: the role of burstiness, topology and non-stationarity. *New Journal of Physics*, 16(7):073037, 2014.
- [HW65] J. Hammersley and D. Welsh. First-passage percolation, subadditive processes, stochastic networks, and generalized renewal theory. In *Proc. Internat. Res. Semin., Statist. Lab., Univ. California, Berkeley*, pages 61–110. Springer-Verlag, New York, 1965.
- [IM09] José Luis Iribarren and Esteban Moro. Impact of human activity patterns on the dynamics of information diffusion. *Phys. Rev. Lett.*, 103:038702, Jul 2009.

- [Jan12] Svante Janson. Simply generated trees, conditioned Galton-Watson trees, random allocations and condensation. *Probab. Surveys*, 9:103–252, 2012.
- [JKLP93] Svante Janson, Donald E. Knuth, Tomasz Łuczak, and Boris Pittel. The birth of the giant component. *Random Structures & Algorithms*, 4(3):233–358, 1993.
- [JPKK14] Hang-Hyun Jo, Juan I. Perotti, Kimmo Kaski, and János Kertész. Analytically solvable model of spreading dynamics with non-Poissonian processes. *Phys. Rev. X*, 4:011041, Mar 2014.
- [Ker98] Götz Kersting. On the height profile of a conditioned Galton-Watson tree. *arXiv preprint arXiv:1101.3656*, 1998.
- [Kes86] Harry Kesten. Subdiffusive behavior of random walk on a random cluster. *Annales de l’Institut Henri Poincaré. Probabilités et Statistiques*, 22(4):425–487, 1986.
- [Kes87] Harry Kesten. Percolation theory and first-passage percolation. *Annals of Probability*, 15(4):1231–1271, 1987.
- [KKP⁺11] M. Karsai, M. Kivelä, R. K. Pan, K. Kaski, J. Kertész, A.-L. Barabási, and J. Saramäki. Small but slow world: How network topology and burstiness slow down spreading. *Phys. Rev. E*, 83:025102, Feb 2011.
- [KNS66] H. Kesten, P. Ney, and F. Spitzer. The Galton-Watson process with mean one and finite variance. *Theory of Probability & Its Applications*, 11(4):513–540, 1966.
- [Kol86] Valentin F. Kolchin. *Random mappings*. Optimization Software Inc., 1986.
- [LG05] Jean-François Le Gall. Random trees and applications. *Probab. Surv.*, 2(245):10–1214, 2005.
- [Lin99] Torgny Lindvall. On Strassen’s theorem on stochastic domination. *Electron. Commun. Probab.*, 4:no. 7, 51–59, 1999.
- [LP16] Russell Lyons and Yuval Peres. *Probability on Trees and Networks*, volume 42 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, New York, 2016. Available at <http://pages.iu.edu/~rdlyons/>.
- [MH13] Naoki Masuda and Petter Holme. Predicting and controlling infectious disease epidemics using temporal networks. *F1000 prime reports*, 5:6, 2013.
- [MM78] A. Meir and J. W. Moon. On the altitude of nodes in random trees. *Canadian Journal of Mathematics*, 30:997–1015, 1978.
- [New03] M. E. J. Newman. The structure and function of complex networks. *SIAM Review*, 45(2):167–256, 2003.
- [Pem90] Robin Pemantle. A time-dependent version of Pólya’s urn. *Journal of Theoretical Probability*, 3(4):627–637, 1990.
- [Pet17] Gábor Pete. *Probability and Geometry on Groups*. In preparation, 2017. Available at <http://www.math.bme.hu/~gabor/PGG.pdf>.

- [PR04] Yuval Peres and David Revelle. Scaling limits of the uniform spanning tree and loop-erased random walk on finite graphs. October 2004. arXiv:math/0410430.
- [PSCVMV15] Romualdo Pastor-Satorras, Claudio Castellano, Piet Van Mieghem, and Alessandro Vespignani. Epidemic processes in complex networks. *Rev. Mod. Phys.*, 87:925–979, Aug 2015.
- [Str65] V. Strassen. The existence of probability measures with given marginals. *The Annals of Mathematical Statistics*, 36(2):423–439, 1965.
- [Stu15] Benedikt Stuffer. Random enriched trees with applications to random graphs. *arXiv preprint arXiv:1504.02006*, 2015.
- [vdH16] R. van der Hofstad. *Random Graphs and Complex Networks*, volume 1. Cambridge University Press, 2016. <http://www.win.tue.nl/~rhofstad/NotesRGCN.pdf>.
- [vdH17] R. van der Hofstad. *Random Graphs and Complex Networks*, volume 2. In preparation, 2017. <http://www.win.tue.nl/~rhofstad/NotesRGCNII.pdf>.
- [Ves12] Alessandro Vespignani. Modelling dynamical processes in complex socio-technical systems. *Nature Physics*, 8:32–39, 2012.
- [Wil96] David Bruce Wilson. Generating random spanning trees more quickly than the cover time. In *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing*, STOC '96, pages 296–303, New York, NY, USA, 1996. ACM.
- [Wor99] Nicholas C Wormald. Models of random regular graphs. *London Mathematical Society Lecture Note Series*, pages 239–298, 1999.