

UC Berkeley

UC Berkeley Previously Published Works

Title

Edge flows in the complete random-lengths network

Permalink

<https://escholarship.org/uc/item/8c95x2h0>

Journal

Random Structures and Algorithms, 37(3)

ISSN

1042-9832

Authors

Aldous, David J
Bhamidi, Shankar

Publication Date

2010-10-01

DOI

10.1002/rsa.20306

Peer reviewed

Edge Flows in the Complete Random-Lengths Network*

David J. Aldous,¹ Shankar Bhamidi²

¹Department of Statistics, University of California, Berkeley, California 94720-3860; e-mail: aldous@stat.berkeley.edu

²Department of Statistics, University of North Carolina, Chapel Hill, North Carolina 27599; e-mail: hamidi@email.unc.edu

Received 30 July 2007; accepted 28 November 2008; received in final form 18 February 2009

Published online 28 January 2010 in Wiley Online Library (wileyonlinelibrary.com).

DOI 10.1002/rsa.20306

ABSTRACT: Consider the complete n -vertex graph whose edge-lengths are independent exponentially distributed random variables. Simultaneously for each pair of vertices, put a constant flow between them along the shortest path. Each edge gets some random total flow. In the $n \rightarrow \infty$ limit we find explicitly the empirical distribution of these edge-flows, suitably normalized. © 2010 Wiley Periodicals, Inc. *Random Struct. Alg.*, 37, 271–311, 2010

Keywords: flow; percolation tree; random graph; random network

1. INTRODUCTION

Write network for an undirected graph whose edges e have positive real edge-lengths $\ell(e)$. In a n -vertex connected network, the distance $D(i, j)$ between vertices i and j is the length of the shortest route between them. Assuming generic edge-lengths, the shortest route is unique. For each ordered (source, destination) pair of vertices (i, j) , send flow of volume $1/n$ along the shortest route from i to j . (The normalization $1/n$ is arbitrary but convenient for (1) below). For each directed edge e (i.e., an edge e and a specified direction across e) of the network, let $f(e)$ be the total flow across the edge in that direction. Note

$$n^{-1} \sum_{\text{directed } e} f(e)\ell(e) = n^{-2} \sum_i \sum_j D(i, j) := \bar{D} \quad (1)$$

where \bar{D} is the average vertex-vertex distance.

Correspondence to: D. J. Aldous

*Supported by NSF (DMS0704159); PIMS, NSERC.

© 2010 Wiley Periodicals, Inc.

One can formulate a project to study the distribution of such edge-flows $f(e)$ in different models of random n -vertex networks. Such models include both deterministic graphs to which random edge-lengths are assigned, and random graphs of both the classical Erdős – Rényi or random regular type [9] and the more recent complex networks types [1, 11, 17, 22] again with real edge-lengths attached. As (1) implies, this project is a refinement of the project of studying $\mathbb{E}\bar{D}_n$, so we envisage a model sufficiently tractable that we know

$$\mathbb{E}\bar{D}_n = (1 + o(1))\bar{d}_n \quad (2)$$

for some explicit (\bar{d}_n) .

To set up some notation, return to the setting of a deterministic network. Because we are using shortest-path routing, we expect edges-flows to be correlated with edge-lengths, so let us study jointly edge-flows and edge-lengths by considering the empirical measure ψ^0 which puts weight $1/n$ on each point $(\ell(e), f(e))$:

$$\psi^0(\cdot, \cdot) := \frac{1}{n} \sum_{\text{directed } e} \mathbb{1}\{(\ell(e), f(e)) \in (\cdot, \cdot)\}.$$

So (1) becomes

$$\iint \ell u \psi^0(d\ell, du) = \bar{D}.$$

So when the lengths of short edges are order 1 we should normalize edge-flows by \bar{d}_n , that is consider the measure

$$\psi_n(\cdot, \cdot) := \frac{1}{n} \sum_{\text{directed } e} \mathbb{1}\{(\ell(e), f(e)/\bar{d}_n) \in (\cdot, \cdot)\}. \quad (3)$$

For a random network, \bar{D}_n is a random variable and $\psi_n(\cdot, \cdot)$ is a random measure, related by

$$\iint \ell y \psi_n(d\ell, dy) = \frac{\bar{D}_n}{\bar{d}_n}. \quad (4)$$

This notation is designed to suggest possible $n \rightarrow \infty$ limit behavior; that the random measures ψ_n converge to a nonrandom measure ψ which by (2) and under appropriate uniform integrability conditions must satisfy

$$\iint \ell y \psi(d\ell, dy) = 1. \quad (5)$$

The purpose of this article is to prove this result and identify ψ in one particular model, described in the next section. There is a fairly simple heuristic argument to identify ψ , shown in Section 1.3. The heuristic argument yields predictions for the limit ψ in many “locally tree-like” models, as discussed in Section 4.4. However for the proofs in this article we exploit special structure of our model, and it seems technically challenging to find rigorous proofs in the broader settings of Section 4.4.

1.1. The Complete Graph With Random Edge-Lengths

Our probability model for a random n -vertex network starts with the complete graph and assigns independent exponential(rate $1/n$) random lengths $L_{ij} = L_{ji} = L_e$ to the $\binom{n}{2}$ edges $e = (i, j)$. This model (which we denote by \mathcal{G}_n) and minor variants (uniform(0, 1) lengths; complete bipartite graph) have been studied in various contexts, for instance the length of minimum spanning tree [12], Steiner tree [10], minimum matching [3, 15, 19, 21], and traveling salesman tour [5, 18, 29]. Note that our scaling convention $\mathbb{E}L_e = n$ makes lengths a factor n larger than in most of the earlier literature. Most closely related to this article is the work of Janson [13] and van der Hofstad, Hooghiemstra, van Mieghem [27] who studied several aspects of the distances $D_n(i, j)$: see also Wästlund [28] for connections with minimum matching. In particular, it is known (16) that $\mathbb{E}\bar{D}_n = (1 + o(1)) \log n$ so that we use $\bar{d}_n := \log n$ to scale edge-flows.

1.2. The Main Result

To fix notation, the vertex-set is $[n] := \{1, 2, \dots, n\}$. All quantities in the n -vertex model \mathcal{G}_n depend on n ; our notation makes n explicit only where helpful. For each ordered pair (i, j) write $\pi(i, j)$ for the shortest path (considered here as a set of directed edges e) from i to j . Define

$$F_n(e) = \frac{1}{n} \sum_{i \in [n]} \sum_{j \in [n], j \neq i} \mathbb{1}\{e \in \pi(i, j)\} \tag{6}$$

so that $F_n(e)$ is the total flow across the directed edge e in the specified direction, when a flow of volume $1/n$ is put along the shortest path between each ordered vertex pair. Write $\#$ for cardinality.

Before stating the full result, Theorem 2, which requires some technical “modes of convergence” discussion, let us state an easier to interpret but less refined result.

Theorem 1. *As $n \rightarrow \infty$ for fixed $z > 0$,*

$$\frac{1}{n} \#\{e : F_n(e) > z \log n\} \rightarrow_{L^1} G(z) := \int_0^\infty \mathbb{P}(W_1 W_2 e^{-u} > z) du \tag{7}$$

where W_1 and W_2 are independent Exponential (1). In particular

$$\frac{1}{n} \mathbb{E} \#\{e : F_n(e) > z \log n\} \rightarrow G(z). \tag{8}$$

Note that the scaling of edge-lengths in \mathcal{G}_n does not affect the conclusions (7,8) which remain true if edge-lengths have Exponential(1) or Uniform(0, 1) distribution. See Section 4.1 for further discussion of the function $G(z)$. In particular its tail behavior is a stretched exponential

$$G(z) = \exp(-z^{\frac{1}{2}+o(1)}) \text{ as } z \rightarrow \infty \tag{9}$$

rather than an ordinary exponential as one might have guessed.

The full result deals with convergence of joint distributions of edge-lengths and edge-flows, but is more complicated to state because the relevant limit for edge-length

is a σ -finite measure. Let us write “distribution” (in quotation marks) in the σ -finite setting—it would be more precise to talk about push-forward of measures, but we find expressions like (10) easier to understand as stated. Take U_∞ independent of (W_1, W_2) above, with the “distribution” of U_∞ being Lebesgue measure on $[0, \infty)$. Then define a σ -finite measure ψ by

$$\psi \text{ is the “distribution” of } (U_\infty, W_1 W_2 e^{-U_\infty}). \tag{10}$$

By carefully declaring the underlying product space to be $[0, \infty) \times (0, \infty)$ we see that ψ is finite on compact subsets. Now recall the quantity ψ_n at (3), with $\bar{d}_n := \log n$, which records lengths and rescaled flows across all edges. This ψ_n is a random measure on $[0, \infty) \times (0, \infty)$ whose total mass is the nonrandom number $n^{-1} \times (\text{number of directed edges}) = n - 1$. Our main result is that ψ_n converges to ψ in the following sense.

Theorem 2.

$$\int h d\psi_n \rightarrow_{L^1} \int h d\psi$$

for all bounded continuous $h : [0, \infty) \times (0, \infty) \rightarrow \mathbb{R}$ with compact support. (11)

To explore what this means, recall that vague convergence $\nu_n \rightarrow \nu$ of σ -finite measures on $[0, \infty) \times (0, \infty)$ means $\int h d\nu_n \rightarrow \int h d\nu$ for bounded continuous test functions $h : [0, \infty) \times (0, \infty) \rightarrow \mathbb{R}$ with compact support. Consider the variant of Theorem 2 in which L^1 convergence is replaced by (the weaker notion of) convergence in probability; it is easy to check this variant is equivalent to the assertion

$\psi_n \rightarrow \psi$ in probability, when viewed as random variables taking values in the space of measures equipped with the vague topology.

Roughly speaking, Theorem 1 is derived from Theorem 2 by projecting $[0, \infty) \times (0, \infty)$ onto the the second (flow- volume) coordinate $(0, \infty)$. In doing this, a technical issue (uniform integrability over $(0, \infty)$) arises—see Section 2.9 for the detailed derivation.

Note the anticipated identity (5) holds because $\int_0^\infty ue^{-u} du = 1$. Section 4.3 states the analogous results for the distribution of flows through vertices instead of edges.

1.3. A Heuristic Argument

Here is a heuristic argument for why the limit is this particular function $G(z)$. The argument rests upon the elementary fact (cf. Lemma 5) that the number $N_n(\tau)$ of vertices within distance τ of a given vertex satisfies

$$e^{-\tau n} N_n(\tau_n) \xrightarrow{d} \text{Exponential}(1) \tag{12}$$

where (throughout this heuristic argument) $\tau_n \uparrow \infty$ slowly.

Consider a short edge e , that is an edge of length $O(1)$. Suppose there are $W'_e(\tau)$ vertices within a large distance τ of one end of e , and $W''_e(\tau)$ vertices within distance τ of the other end. A shortest-length path between distant vertices which passes through e must enter and exit the region above via some pair of vertices in the sets above (see Fig. 2), and there are $W'_e(\tau)W''_e(\tau)$ such pairs. The dependence on the length L_e is more subtle. We expect a longer edge to carry less flow because of the availability of alternate possible shortest paths.

It turns out that the flow through e depends on L_e as $\exp(-L_e)$, and in retrospect this can be viewed as a consequence of the exponential growth (12) of number of vertices in a ball and the alternate possible paths they provide.

Combining these two effects, the flow through edges e should be roughly proportional (as e varies) to $W'_e(\tau)W''_e(\tau)\exp(-L_e)$, that is roughly equal to $c_n W'_e(\tau) W''_e(\tau)\exp(-L_e)$ for some constant c_n . Taking $\tau = \tau_n \uparrow \infty$ slowly and using (12) we have $W'_e(\tau_n) \approx e^{\tau_n} W_1$ and $W''_e(\tau_n) \approx e^{\tau_n} W_2$ for independent Exponential(1) W_1, W_2 and so the flow across an edge e is approximately $c'_n W_1 W_2 \exp(-L_e)$. Finally, as $n \rightarrow \infty$ the normalized distribution $n^{-1}\#\{e : L_e \in \cdot\}$ over directed edges converges to the σ -finite distribution of U_∞ . This is heuristically how the limit joint distribution $(U_\infty, W_1 W_2 \exp(-U_\infty))$ arises.

1.4. Ingredients and Structure of the Proof

Here we give an overview of the structure of the proof, while mentioning (in the bullet points) the variety of tools used to handle the details.

Directly formalizing the heuristic argument gets us to Proposition 12 which gives the desired asymptotics for the conditional mean flow across an edge (conditioned on a large neighborhood). This result is based on

- analyzing the behavior of the percolation (flow from vertex 1) process in a given neighborhood (Sections 2.5 and 2.6)

but we also use

- the Yule process local approximation (section 2.3)
- a martingale property (Section 2.2).

At this point we realize that we need to show the actual flow is close to its conditional mean, which requires a conditional variance bound, deferred to Section 3. Granted the actual flow (a “global” quantity) is indeed close to its conditional mean (a “local” quantity), we can use

- a general weak law of large numbers for local functions on \mathcal{G}_n (Section 2.8).

to complete the proof of Theorems 1 and 2 in Section 2.9.

To prove (Section 3) the conditional variance bound, Proposition 19, we need to study the joint behavior of two shortest paths $\pi(1, 2), \pi(3, 4)$. This involves somewhat intricate conditioning arguments. The key ideas are

- finite- n bounds for mean intensities of short paths (Section 3.1)
- the size-biased Yule process (Section 3.4)
- conditional on existence of a given short path from vertex 1, the process of numbers of vertices within distance t from vertex 1 grows as a size-biased Yule process (Section 3.4).

We remark that while the overall structure of the argument could be adapted to more general random networks (in particular, those mentioned in Section 4.4), our proofs of the detailed lemmas rely on calculations with our particular, and simple, model, so it is not at all clear whether it is feasible to adapt to other models.

2. PROOFS

Section 2 provides the proofs of Theorems 1 and 2, except that one result (Proposition 19) is postponed and proved in Section 3.

2.1. Preliminaries

Exponential(λ) and Geometric(p) denote the exponential and geometric distributions, as usual parametrized by rate.

Here we collect without proof some standard properties of the random network model \mathcal{G}_n . For fixed n and for $t \geq 0$ define

$$N_n(t) := \text{number of vertices within distance } t \text{ from vertex } 1 \tag{13}$$

where we include vertex 1 itself;

$$S_{n,k} := \min\{t : N_n(t) = k\}, \quad 1 \leq k \leq n - 1$$

so that $S_{n,k+1}$ is the distance from vertex 1 to the k 'th nearest distinct vertex. Then

$$(S_{n,k+1} - S_{n,k}, 1 \leq k \leq n - 1) \text{ are independent Exponential}\left(\frac{k(n-k)}{n}\right). \tag{14}$$

Because the distance $D(1, 2)$ is distributed as $S_{n,V}$ for V uniform on $\{2, 3, \dots, n\}$, it is straightforward (see e.g. [14, 27] for similar calculations) to use (14) to write exact formulas for the mean, variance and generating function of $D(1, 2)$ and then deduce the $n \rightarrow \infty$ limit behavior

$$\mathbb{E}D(1, 2) - \log n \rightarrow c_1, \quad \text{var } D(1, 2) \rightarrow c_2, \quad D(1, 2) - \log n \xrightarrow{d} D_\infty \tag{15}$$

for finite constants c_1, c_2 and a distribution D_∞ discussed further in Section 4.6. Note that the average vertex–vertex distance \bar{D}_n at (1) has the same mean, but not the same distribution, as $D(1, 2)$, so

$$\mathbb{E}\bar{D}_n = \log n + O(1) \text{ as } n \rightarrow \infty. \tag{16}$$

There is a natural mental picture of (first passage) percolation, in which at time 0 there is water at vertex 1 only, and the water spreads along edges at speed 1. So $N_n(t)$ vertices have been wetted by time t . Each vertex $j \neq 1$ is first wetted via some edge $(i(j), j)$, and this collection of directed edges forms the percolation tree rooted at vertex 1. The “flow” in Theorem 2 from vertex 1 goes along the edges of this percolation tree, and no other edges.

Associated with the percolation process is a filtration (\mathcal{F}_t) , where \mathcal{F}_t is the information known at time t , illustrated informally as follows. Write T_i for the wetting time of vertex i . Then the values of T_i for which $T_i \leq t$ are in \mathcal{F}_t . On the event $\{T_i < t < T_j\}$, the information in \mathcal{F}_t about L_{ij} is that $L_{ij} > t - T_i$. By the memoryless property of the exponential distribution, on the event above the conditional distribution of $L_{ij} - (t - T_i)$ given \mathcal{F}_t is Exponential($1/n$), and the (L_{ij}) are conditionally independent given \mathcal{F}_t . More elaborate versions of this memoryless property appear in Lemmas 8 and 10.

An obvious consequence of the Exponential($1/n$) distribution of edge-lengths L_{ij} is that

$$n\mathbb{P}(L_{ij} \leq \tau) \rightarrow \tau \text{ as } n \rightarrow \infty.$$

In words, this says that the measure $n\mathbb{P}(L_{ij} \in \cdot)$ converges vaguely to Lebesgue measure on $(0, \infty)$.

2.2. A Martingale Property

In the percolation process above, write $\mathcal{W}_n(t)$ for the set of vertices wetted by time t . The following martingale property turns out to be useful.

Lemma 3. *Let W_n be a stopping time for the percolation process on \mathcal{G}_n . For each $v \in \mathcal{W}_n(W_n)$ let $Y(v)$ be the number of vertices $j \in [n]$ such that, in the shortest path from 1 to j , the last-visited vertex of $\mathcal{W}_n(W_n)$ is vertex v . Then*

$$\frac{1}{n} \mathbb{E}(Y(v) | \mathcal{F}_{W_n}) = \frac{1}{N_n(W_n)}.$$

Proof. Define $Y(v, t)$ as $Y(v)$ but counting only vertices j which are wetted by time t . As t increases, whenever a new vertex j is wetted via some edge $(i(j), j)$, the predecessor vertex $i(j)$ is a uniform random element of $\mathcal{W}(t-)$. It follows easily that the process $\frac{Y(v,t)}{N_n(t)}$, $t \geq W_n$ is a martingale. The optional sampling theorem shows

$$\mathbb{E} \left(\frac{Y(v, \infty)}{N_n(\infty)} \middle| \mathcal{F}_{W_n} \right) = \frac{Y(v, W_n)}{N_n(W_n)} = \frac{1}{N_n(W_n)}.$$

But $\frac{Y(v, \infty)}{N_n(\infty)} = \frac{Y(v)}{n}$. ■

2.3. The Yule Process Approximation

The Yule process $(N_\infty(t), 0 \leq t < \infty)$ is the population at time t in the continuous-time branching process started with one individual, in which individuals live forever and produce offspring at the times of a Poisson(1) process. Writing

$$S_{\infty,k} := \min\{t : N_\infty(t) = k\}$$

we have

$$(S_{\infty,k+1} - S_{\infty,k}, 1 \leq k \leq \infty) \text{ are independent Exponential}(k) \tag{17}$$

and so the Yule process is the natural $n \rightarrow \infty$ limit of the process (14) associated with the percolation tree. We quote some standard facts about the Yule process. For (a) see e.g. [24] p. 235.

Lemma 4.

- (a) $N_\infty(t)$ has Geometric(e^{-t}) distribution.
- (b) $e^{-t}N_\infty(t)$ is a martingale which is bounded in L^2 , and $e^{-t}N_\infty(t) \rightarrow W$ a.s. and in L^2 as $t \rightarrow \infty$, where W has Exponential(1) distribution.

It is intuitively clear that that the local structure of \mathcal{G}_n relative to one vertex converges to the Yule process. Abstractly [6], we call this notion of convergence of random graphs local weak convergence and the limit structure (the Yule process regarded as a “spatial” graph) is called the PWIT. But rather than work abstractly we will state only the more concrete consequences needed, such as the next lemma.

Lemma 5. *Fix $k \geq 1$ and $t < \infty$. For $1 \leq i \leq k$ let $N_n^{(i)}(t)$ be the number of vertices of \mathcal{G}_n within distance t from vertex i . Then as $n \rightarrow \infty$*

$$(N_n^{(1)}(t), \dots, N_n^{(k)}(t)) \xrightarrow{d} (N_\infty^{(1)}(t), \dots, N_\infty^{(k)}(t))$$

where the limits $N_\infty^{(i)}(t)$ are independent Geometric(e^{-t}).

Proof. For $k = 1$ this follows from (14,17) and Lemma 4(a). For general k , use the natural conditioning argument. ■

The following technical lemma shows one way in which the “exponential growth with rate 1” property of the Yule process (Lemma 4(b)) translates to the percolation process.

Lemma 6. *Let W_n be a randomized stopping time for the percolation process on \mathcal{G}_n . Fix $\varepsilon > 0, \sigma < \infty$ and a sequence (ω_n) such that $\omega_n \rightarrow \infty$ with $\omega_n \leq n^{1/2}$. Then as $n \rightarrow \infty$*

$$\mathbb{P}\left(1 - \varepsilon \leq \frac{N_n(W_n + \sigma)}{e^\sigma N_n(W_n)} \leq 1 + \varepsilon \mid \mathcal{F}_{W_n}\right) \rightarrow 1$$

uniformly on $\{\omega_n \leq N_n(W_n) \leq n/\omega_n\}$.

Proof. It is enough to show this holds conditionally on $N(W_n)$, that is to show

$$\mathbb{P}\left(1 - \varepsilon \leq \frac{N_n(t_n + \sigma)}{e^\sigma k_n} \leq 1 + \varepsilon \mid W_n = t_n, N_n(W_n) = k_n\right) \rightarrow 1 \tag{18}$$

whenever $k_n \rightarrow \infty, n/k_n \rightarrow \infty$. Note that the value of t_n does not affect the conditional probability.

First note that by (14,17) we can couple $N_n(\cdot)$ and $N_\infty(\cdot)$ by constructing each from the same i.i.d. Exponential(1) sequence (Y_i) via

$$\begin{aligned} S_{\infty,k} &= \sum_{i=1}^{k-1} \frac{1}{i} Y_i; & N_\infty(t) &= \max\{k : S_{\infty,k} \leq t\} \\ S_{n,k} &= \sum_{i=1}^{k-1} \frac{n}{i(n-i)} Y_i; & N_n(t) &= \max\{k : S_{n,k} \leq t\}. \end{aligned} \tag{19}$$

From this coupling we see that for $k_n < m_n$ with $k_n \rightarrow \infty, n/m_n \rightarrow \infty$ we have

$$\frac{S_{n,m_n} - S_{n,k_n}}{S_{\infty,m_n} - S_{\infty,k_n}} \rightarrow 1 \text{ a.s.} \tag{20}$$

By the homogeneous branching property of the Yule process and Lemma 4(a), conditional on $\{N_\infty(t_n) = k_n\}$, we can represent $N_\infty(t_n + \sigma)$ as the sum of k_n independent Geometric($e^{-\sigma}$) r.v.'s, and so (still conditionally) $\frac{N_\infty(W_n + \sigma)}{k_n e^\sigma} \rightarrow 1$ in probability. In terms of $(S_{\infty,k})$ this is equivalent to the (now unconditional) property that

$$\text{if } k_n \rightarrow \infty, m_n \sim k_n e^\sigma \text{ then } S_{\infty,m_n} - S_{\infty,k_n} \xrightarrow{p} \sigma.$$

Now by (20) and assumptions on (k_n) we see

$$\text{if } k_n \rightarrow \infty, m_n \sim k_n e^\sigma \text{ then } S_{n,m_n} - S_{n,k_n} \xrightarrow{p} \sigma.$$

This holds for each fixed $\sigma > 0$. Translating this back into an assertion about $(N_n(\cdot))$ establishes (18). ■

Recall that $D(i, j)$ denotes vertex-vertex distances in \mathcal{G}_n .

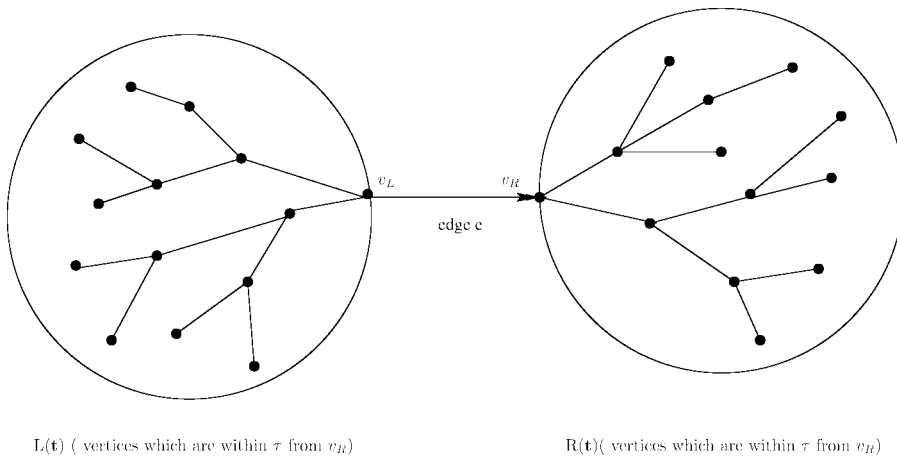


Fig. 1. Neighborhood about an edge.

Lemma 7. For disjoint subsets $B, C \subset [n]$ and for any $d > 0$,

$$\mathbb{E}\#\{(i, j) \in B \times C : D(i, j) \leq d\} \leq \#B\#C e^d / (n - 1).$$

Proof. By linearity and symmetry we reduce to the case where B and C are singletons. Then the left side equals $\mathbb{P}(D(1, 2) \leq d) = (n - 1)^{-1}(\mathbb{E}N_n(d) - 1)$. By the coupling (19) to the Yule process, and Lemma 4(a), $\mathbb{E}N_n(d) \leq \mathbb{E}N_\infty(d) = e^d$. ■

2.4. Local Structure in the n -Vertex Model

In this section, we give a result (Lemma 8) describing the global structure of \mathcal{G}_n conditional on a given local structure. The actual result is simple once stated, but requires some notational effort to set up. Note the notational convention that boldface \mathbf{t}, \mathbf{T} denote **tree**.

Fix a real $\tau > 0$. Let \mathbf{t} be a finite unlabelled tree with edge lengths, with the following properties (see Fig. 1).

- (i) There is a distinguished directed edge, whose end vertices can then be labelled as (v_L, v_R) , defining a partition of all the vertices of \mathbf{t} as $V(\mathbf{t}) = L(\mathbf{t}) \cup R(\mathbf{t})$. Here L and R are mnemonics for left and right, and the distinguished edge is directed left-to-right.
- (ii) Every vertex in $L(\mathbf{t})$ is within distance τ from v_L , and every vertex in $R(\mathbf{t})$ is within distance τ from v_R .

Write \mathbf{T}_τ for the set of such trees \mathbf{t} . For such \mathbf{t} , write $\ell(e)$ for the length of the distinguished edge e . And for each vertex $v \in L(\mathbf{t})$ write $b(v) = \tau - D(v, v_L)$ as the “distance to boundary” (as in Fig. 1, we envisage a boundary drawn at distance τ from v_L and from v_R). Similarly for $v \in R(\mathbf{t})$ write $b(v) = \tau - D(v, v_R)$. Finally, given \mathbf{t} and given a subset $A \subseteq [n]$ with $\#A = \#V(\mathbf{t})$, let \mathbf{t}_A denote some labelling of the vertices of \mathbf{t} by distinct labels from A .

Now consider the random network \mathcal{G}_n . We occasionally want to regard an edge (i, j) of \mathcal{G}_n as a point-set, so that a point on the edge is at some distance $0 < u < L_{ij}$ from vertex i and at distance $L_{ij} - u$ from vertex j , with this notion of distance extending in the natural way to

distances between a point on an edge and a distant vertex. Fix τ and vertices $v_L, v_R \in [n]$. Define the “neighborhood” $\mathcal{B}_\tau(v_L, v_R)$ as the subgraph of \mathcal{G}_n whose vertex-set consists of vertices v for which $\min(D(v, v_L), D(v, v_R)) \leq \tau$. Its edge-set is the subset of edges of \mathcal{G}_n such that, for every point along the edge, the distance to the closer of $\{v_L, v_R\}$ is at most τ . It also contains by fiat the distinguished directed edge (v_L, v_R) .

In general $\mathcal{B}_\tau(v_L, v_R)$ need not be a tree, but if it is a tree then clearly it is a tree of the form \mathbf{t}_A for some $\mathbf{t} \in \mathbf{T}_\tau$ and some $A \subseteq [n]$. In this case we, can define $b(i)$ (meaning distance to boundary of neighborhood) as above for vertices i of $\mathcal{B}_\tau(v_L, v_R)$, and we define $b(i) = 0$ for other vertices of \mathcal{G}_n .

Lemma 8. *Fix n, τ, v_L, v_R and \mathbf{t}_A . Then conditional on $\mathcal{B}_\tau(v_L, v_R) = \mathbf{t}_A$, the lengths (L_{ij}) of the edges of \mathcal{G}_n which are not edges of $\mathcal{B}_\tau(v_L, v_R)$ are independent r.v.’s for which $L_{ij} - b(i) - b(j)$ has Exponential($1/n$) distribution.*

Proof. We first check, by considering all the cases,

$$(i, j) \text{ is not an edge of } \mathcal{B}_\tau(v_L, v_R) \text{ if and only of } L_{ij} > b(i) + b(j). \tag{21}$$

■

Case 1. neither i nor j is in the vertex-set of $\mathcal{B}_\tau(v_L, v_R)$. Here $L_{ij} > 0 = b(i) + b(j)$ and (i, j) is not an edge of $\mathcal{B}_\tau(v_L, v_R)$.

Case 2. i but not j is in the vertex-set of $\mathcal{B}_\tau(v_L, v_R)$. Here the fact that j is not in the vertex-set implies $L_{ij} + \min(D(i, v_L), D(i, v_R)) > \tau$, which rearranges to $L_{ij} > b(i) = b(i) + b(j)$. And (i, j) is not an edge of $\mathcal{B}_\tau(v_L, v_R)$.

Case 3. both i and j are in the vertex-set of $\mathcal{B}_\tau(v_L, v_R)$. In this case there is a route from one of $\{v_L, v_R\}$, through edge (i, j) and back to one of $\{v_L, v_R\}$, whose length is

$$\min(D(i, v_L), D(i, v_R)) + L_{ij} + \min(D(j, v_L), D(j, v_R)) = 2\tau + L_{ij} - (b(i) + b(j)).$$

The definition for (i, j) to be an edge of $\mathcal{B}_\tau(v_L, v_R)$ is equivalent to the property that this route length is at most 2τ , in other words that $L_{ij} \leq b(i) + b(j)$.

This verifies (21). The edge-lengths (L_{ij}) are a priori independent Exponential($1/n$), and conditioning on all these inequalities leaves them independent with the stated distributions.

2.5. The Percolation Tree on a Neighborhood

We now come to the central idea of the proof, which is to study how the percolation tree behaves on a given neighborhood. Until further notice we adopt the setting of Lemma 8 and work conditionally on $\mathcal{B}_\tau(v_L, v_R) = \mathbf{t}_A$. Let \mathbb{E}_τ and \mathbb{P}_τ denote the conditional expectation and conditional probability operations, respectively; and we will use tilde notation \tilde{L}_{ij} to denote unconditioned quantities. Thus the conclusion of Lemma 8 can be rewritten as follows. Starting with independent Exponential($1/n$) r.v.’s (\tilde{L}_{ij}) we can construct the conditioned lengths (L_{ij}) as

$$L_{ij} = \tilde{L}_{ij} + b(i) + b(j), \quad (i, j) \notin \mathcal{B}_\tau(v_L, v_R). \tag{22}$$

This provides a coupling of the unconditioned and conditioned lengths.




Fig. 2. Flow from 1 passing through the neighborhood. Flow enters at time T_L along some edge (H', H) , and exits at time $T_L + 2\tau + \ell(e)$.

Now consider the percolation process started at vertex 1, and assume $1 \notin A$. Let T_L be the first time (in the conditioned model) that the percolating water gets to some point on an edge at distance τ from v_L via some path which does not include the distinguished edge. Define T_R similarly and then set $T = \min(T_L, T_R)$. So at time T the water is at distance $b(H)$ from some random vertex H of A . See Fig. 2.

Lemma 9.

- (a) *The (conditioned) distribution of T is the same as the (unconditioned) distribution of \tilde{T}_A , the first time in percolation on \mathcal{G}_n that some vertex in A is wetted. And the random vertex H is distributed uniformly on A , independent of T .*
- (b) *$N_n(T)$ is distributed as the smallest of $\#A$ uniform random samples without replacement from $\{2, 3, \dots, n\}$.*

Proof. Use (22) to construct the conditioned process from the unconditioned process. In the unconditioned process, it is clear by symmetry that H , the first vertex of A wetted, is uniform on A and independent of \tilde{T}_A . Obviously H is reached along some edge (H', H) with $H' \notin A$. In the conditioned process, at time \tilde{T}_A the percolating water has reached distance $b(H) = b(H) + b(H')$ from H , and hence is distance τ from either v_L or v_R (whichever is closer to H). So $T = \tilde{T}_A$ in the coupling. This gives (a).

To prove (b), let us first look at the unconditioned process. Fix a set of $\#A$ vertices other than vertex 1 and start the flow percolating from vertex 1 through the whole network at rate 1. Let T^* denote the first time that a vertex in the set A has been reached by the flow. We shall find the distribution of $N_n(T^*)$, the number of vertices wetted by time T^* . By part (a) this has the same distribution as $N_n(T)$ in the conditioned setup.

We recall the dynamics of the flow exploring the network. Suppose at some time t , there are k vertices, say $E_k = \{i_1 = 1, i_2, \dots, i_k\}$, which have been found by the flow. Then the time for the next vertex to be reached by the flow is exponentially distributed with rate $k(n - k)/n$ and, more importantly, this vertex is uniformly distributed over the set $[n] \setminus E_k$ of vertices available. Thus $N_n(T^*) \geq 2$ (since $1 \notin A$ already belongs to the set of vertices

explored by the flow). For any fixed $j \geq 2$ we have $N_n(T^*) = j$ if and only if the first $j - 1$ nodes reached by the flow do not belong to the set A but the j 'th node does belong to A . Using the “sequential uniform choices” dynamics of the flow process, we calculate

$$\mathbb{P}(N_n(T^*) = j) = \frac{\#A}{n - (j - 1)} \prod_{k=1}^{j-2} \left(1 - \frac{\#A + k}{n - k}\right).$$

But this is the distribution of the smallest of $\#A$ uniform random samples without replacement from the set $\{2, 3, \dots, n\}$, establishing (b). ■

The next lemma formalizes the idea “what do we know about edge-lengths at time T_L ?” As described earlier, on $\{T_L < T_R\}$ there is some vertex $H \in L(\mathbf{t}_A)$ such that $T_H - T_E = b(H)$, and vertex H gets wetted via some edge (H', H) where H' is in the set $\mathcal{W}(T_L)$ of vertices wetted by time T_L . Arguing as in Lemma 8 shows

Lemma 10. *Conditional on $\mathcal{B}_\tau(v_L, v_R) = \mathbf{t}_A$ and conditional on $\sigma(\mathcal{F}_{T_L}, H', H)$, on the event $\{T_L < T_R\}$, the collection of edge-lengths (L_{ij}) as (i, j) runs over all edges except*

- (i) (i, j) an edge of $\mathcal{B}_\tau(v_L, v_R)$
- (ii) $i, j \in \mathcal{W}(T_L)$
- (iii) $(i, j) = (H', H)$

are independent with distributions

$$\begin{aligned} L_{ij} &= \tilde{L}_{ij} + b(i) + b(j), & (i, j) \notin \mathcal{W}(T_L) \\ L_{ij} &= \tilde{L}_{ij} + (T_L - T_i) + b(j), & i \in \mathcal{W}(T_L). \end{aligned}$$

Recall we are conditioning on $\mathcal{B}_\tau(v_L, v_R) = \mathbf{t}_A$. Write

$$\sigma := 2\tau + \ell(e)$$

where e is the distinguished edge $e = (v_L, v_R)$ of \mathbf{t}_A . Estimating the mean flow through e is tantamount to estimating the mean of the random variable

$$M_1 := \#\{2 \leq j \leq n : e \in \pi(1, j)\} \tag{23}$$

counting the number of vertices $j \in [n]$ with $j \neq 1$ such that the shortest path from 1 to j passes through e . To analyze M_1 , we consider the set $\mathcal{R}^*(\mathbf{t}_A)$ of vertices in $R(\mathbf{t}_A)$ which are first reached via edge e :

$$\mathcal{R}^*(\mathbf{t}_A) := \{v \in R(\mathbf{t}_A) : e \in \pi(1, v)\}.$$

Note that from the definitions of T_L and T_R

$$\begin{aligned} \text{if } T_R \leq T_L & \quad \text{then} & \quad \#\mathcal{R}^*(\mathbf{t}_A) = 0 \text{ and } M_1 = 0 \\ \text{if } T_L < T_R \leq T_L + \sigma & \quad \text{then} & \quad 0 \leq \#\mathcal{R}^*(\mathbf{t}_A) \leq \#R(\mathbf{t}_A) \\ \text{if } T_R > T_L + \sigma & \quad \text{then} & \quad \#\mathcal{R}^*(\mathbf{t}_A) = \#R(\mathbf{t}_A). \end{aligned}$$

Lemma 11. $\mathbb{E}_\tau(M_1 | \mathcal{F}_{T_L + \sigma}) = \frac{n \#\mathcal{R}^*(\mathbf{t}_A)}{N(T_L + \sigma)}.$

Proof. In the notation of Lemma 3

$$M_1 = \sum_{v \in \mathcal{R}^*(\mathbf{t}_A)} Y(v). \tag{24}$$

Applying Lemma 3 with $W = T_L + \sigma$ (the fact we are working conditional on $\mathcal{B}_\tau(v_L, v_R) = \mathbf{t}_A$ doesn't affect the martingale property after time W) gives the second equality below:

$$\frac{1}{n} \mathbb{E}_\tau(M_1 | \mathcal{F}_W) = \sum_{v \in \mathcal{R}^*(\mathbf{t}_A)} \frac{1}{n} \mathbb{E}_\tau(Y(v) | \mathcal{F}_W) = \frac{\#\mathcal{R}^*(\mathbf{t}_A)}{N(W)}.$$

■

2.6. The Conditioned Mean Flow

We now start studying $n \rightarrow \infty$ asymptotics. Recall the definition (6) of the normalized flow $F_n(e)$ across an edge e of \mathcal{G}_n . The next result calculates the expected flow conditional on the neighborhood structure of \mathcal{G}_n around e .

Proposition 12. Fix τ and $\mathbf{t} \in \mathbf{T}_\tau$ and write $\sigma = 2\tau + \ell(e)$ where e is the distinguished edge e of \mathbf{t} . Let $v_L^n \neq v_R^n \in [n]$ and let $\{v_L^n, v_R^n\} \subseteq A_n \subset [n]$ satisfy $\#A_n = \#\mathbf{t}$. Then as $n \rightarrow \infty$, setting $e_n = (v_L^n, v_R^n)$,

$$\mathbb{E}_\tau(F_n(e_n)) = (1 + o(1))\#L(\mathbf{t})\#R(\mathbf{t})e^{-\sigma} \log n.$$

Here $\mathbb{E}_\tau(\cdot)$ denotes $\mathbb{E}(\cdot | \mathcal{B}_\tau(v_L^n, v_R^n) = \mathbf{t}_{A_n})$. All quantities except \mathbf{t}, τ, σ depend on n , though the dependence is often not made explicit in notation in the proof below.

Proof. Assume vertex $1 \notin A$. Recall the definition of M_1 from (23). We shall show that

$$\mathbb{E}_\tau(M_1) = (1 + o(1))\#L(\mathbf{t})\#R(\mathbf{t})e^{-\sigma} \log n. \tag{25}$$

In terms of flows of volume $1/n$ between each vertex-pair, the (conditional) mean contribution to the flow $F_n(e)$ through the distinguished edge arising from flow started at vertex $1 \notin A$ equals $n^{-1}\mathbb{E}(M_1)$. The same contribution arises from each of the $n - \#\mathbf{t}$ possible starting vertices $v \notin A$. For $v \in A$ the flow through e is trivially bounded by 1. So to prove Proposition 12 it is enough to prove (25).

Fix a sequence $\omega_n \rightarrow \infty$, with $\omega_n = o(\log n)$. Start the first passage percolation process from vertex 1. Recall that T_L denotes the first time the flow is within distance τ from v_L , and that by time $T_L + \sigma$ the flow has wetted every vertex in \mathbf{t}_A . We shall show that the dominant contribution to $\mathbb{E}_\tau(M_1)$ is from the “good” event

$$G^* := \{\omega_n \leq N_n(T_L) \leq n/\omega_n\} \cap \{T_R > T_L + \sigma\}. \tag{26}$$

For the details, observe that we can apply Lemma 6 to the percolation process on edges excluding the edges of \mathbf{t}_A and deduce

$$\begin{aligned} \mathbb{P}_\tau \left(1 - \varepsilon \leq \frac{N_{\text{avoid}}(T_L + \sigma)}{e^\sigma N(T_L)} \leq 1 + \varepsilon \mid \mathcal{F}_{T_L} \right) \\ \rightarrow 1 \text{ uniformly on } \{\omega_n \leq N(T_L) \leq n/\omega_n\} \end{aligned} \tag{27}$$

where $N_{\text{avoid}}(t)$ is the number of vertices wetted by time t via paths which use no edge of \mathbf{t}_A . For the rest of the argument, we work on the event $\{T_L < T_R\}$ (otherwise, $M_1 = 0$). Recall (Fig. 2) that flow enters the neighborhood at time T_L along some edge (H', H) . We claim: conditional on $\sigma(\mathcal{F}_{T_L}, H, H')$, on the event $\{T_L < T_R\}$, the expected number of vertices of $R(\mathbf{t}_A)$ wetted before time $T_L + \sigma$ by paths not using the distinguished edge is at most

$$(N(T_L) + \#L(\mathbf{t}))\#R(\mathbf{t})e^\sigma / (n - 1). \tag{28}$$

This follows from Lemmas 10 and 7 applied to $\mathcal{W}(T_L) \cup L(\mathbf{t}_A)$ and $R(\mathbf{t}_A)$, because the former lemma implies that the conditioned edge lengths can only be longer than the unconditioned edge-lengths in the latter lemma. Note that the expectation (28) tends to 0 on $\{N(T_L) \leq n/\omega_n\}$, so that

$$\mathbb{P}_\tau(T_R > T_L + \sigma | \mathcal{F}_{T_L}) \rightarrow 1 \text{ uniformly on } \{N(T_L) \leq n/\omega_n\} \cap \{T_L < T_R\}. \tag{29}$$

Next let $N_{\text{via}}(t)$ denote the number of vertices outside \mathbf{t}_A which have been wetted by time $T_L + \sigma$ using some path via \mathbf{t}_A . Again using Lemmas 10 and 7, applied now to $L(\mathbf{t}_A)$ and $[n] \setminus A$, we find

$$\mathbb{E}_\tau(\mathbb{1}\{T_L < T_R\}N_{\text{via}}(T_L + \sigma) | \mathcal{F}_{T_L}) \leq \#L(\mathbf{t})e^\sigma. \tag{30}$$

From the definitions,

$$0 \leq N(T_L + \sigma) - N_{\text{avoid}}(T_L + \sigma) \leq N_{\text{via}}(T_L + \sigma) + \#\mathbf{t}.$$

Combining this with (27, 29, 30) we deduce

$$\begin{aligned} \mathbb{P}_\tau\left(T_L + \sigma < T_R; 1 - \varepsilon \leq \frac{N(T_L + \sigma)}{e^\sigma N(T_L)} \leq 1 + \varepsilon \mid \mathcal{F}_{T_L}\right) \\ \rightarrow 1 \text{ uniformly on } \{\omega_n \leq N(T_L) \leq n/\omega_n\} \cap \{T_L < T_R\}. \end{aligned} \tag{31}$$

Recalling the definition (26) of G^* , Lemma 11 implies

$$\frac{\mathbb{E}_\tau[M_1 \mathbb{1}(G^*)]}{n \#R(\mathbf{t})} = \mathbb{E}_\tau\left[\frac{\mathbb{1}(G^*)}{N(T_L + \sigma)}\right]. \tag{32}$$

To obtain asymptotics for the right side, use the Lemma 9(b) description of the distribution of $N(T_L)$ to conclude that

$$\mathbb{P}_\tau(T_L < T_R, N_n(T_L) = m) = (1 + o(1)) \frac{\#L(\mathbf{t})}{n - 1} \text{ uniformly on } \{2 \leq m \leq n/\omega_n\}. \tag{33}$$

The harmonic sum estimate $\sum_{\omega_n}^{n/\omega_n} j^{-1} = (1 + o(1)) \log n$ leads to

$$\mathbb{E}_\tau\left[\frac{1}{e^\sigma N(T_L)} \mathbb{1}\{\omega_n \leq N(T_L) \leq n/\omega_n\} \mathbb{1}\{T_L + \sigma < T_R\}\right] = (1 + o(1))\#L(\mathbf{t})e^{-\sigma} n^{-1} \log n.$$

Combine this with (31) to get

$$\mathbb{E}_\tau \left[\frac{\mathbb{1}(G^*)}{N(T_L + \sigma)} \right] = (1 + o(1))\#L(\mathbf{t})e^{-\sigma}n^{-1} \log n$$

and thus by (32)

$$\mathbb{E}_\tau(M_1 \mathbb{1}(G^*)) = (1 + o(1))\#L(\mathbf{t})\#R(\mathbf{t})e^{-\sigma} \log n.$$

Recalling that $M_1 = 0$ on $\{T_R < T_L\}$, we can write

$$\mathbb{E}_\tau[M_1] = \mathbb{E}_\tau[M_1 \mathbb{1}(G^*)] + \mathbb{E}_\tau[M_1 \mathbb{1}(B_1)] + \mathbb{E}_\tau[M_1 \mathbb{1}(B_2)] + \mathbb{E}_\tau[M_1 \mathbb{1}(B_3)]$$

for the “bad” events

$$\begin{aligned} B_1 &:= \{T_L < T_R\} \cap \{N(T_L) > n/\omega_n\} \\ B_2 &:= \{T_L < T_R\} \cap \{N(T_L) < \omega_n\} \\ B_3 &:= \{T_L < T_R < T_L + \sigma\} \cap \{\omega_n \leq N(T_L) \leq n/\omega_n\} \end{aligned}$$

and we need to check for each B that $\mathbb{E}_\tau[M_1 \mathbb{1}(B)] = o(\log n)$. In each case we start by using Lemma 11.

$$\begin{aligned} \mathbb{E}_\tau[M_1 \mathbb{1}(B_1)] &= n\mathbb{E}_\tau \left(\frac{\#\mathcal{R}^*(\mathbf{t}_A)}{N(T_L + \sigma)} \mathbb{1}\{T_L < T_R\} \mathbb{1}\{N(T_L) > n/\omega_n\} \right) \\ &\leq \#R(\mathbf{t})\omega_n = o(\log n) \end{aligned}$$

because we chose $\omega_n = o(\log n)$.

$$\begin{aligned} \mathbb{E}_\tau[M_1 \mathbb{1}(B_2)] &= n\mathbb{E}_\tau \left(\frac{\#\mathcal{R}^*(\mathbf{t}_A)}{N(T_L + \sigma)} \mathbb{1}\{T_L < T_R\} \mathbb{1}\{N(T_L) \leq \omega_n\} \right) \\ &\leq n\#R(\mathbf{t})\mathbb{P}_\tau \left(T_L < T_R, N_n(T_L) \leq \omega_n \right) \text{ by (33)} \\ &= (1 + o(1))\#R(\mathbf{t})\#L(\mathbf{t})\omega_n \\ &= o(\log n). \end{aligned}$$

For the third event,

$$\mathbb{E}_\tau[M_1 \mathbb{1}(B_3)] = n\mathbb{E}_\tau \left(\frac{\#\mathcal{R}^*(\mathbf{t}_A)}{N(T_L + \sigma)} \mathbb{1}\{T_L < T_R < T_L + \sigma\} \mathbb{1}\{\omega_n \leq N_n(T_L) \leq n/\omega_n\} \right).$$

By (28) we have

$$\mathbb{P}_\tau \left(T_R < T_L + \sigma \mid \mathcal{F}_{T_L} \right) \leq (N(T_L) + \#L(\mathbf{t}))\#R(\mathbf{t})e^\sigma / (n - 1).$$

Because $N(T_L) \leq N(T_L + \sigma)$ and $\#\mathcal{R}^*(\mathbf{t}_A) \leq \#R(\mathbf{t})$, writing $C = \{T_L < T_R\} \cap \{\omega_n \leq N(T_L) \leq n/\omega_n\}$ we have

$$\begin{aligned} \mathbb{E}_\tau[M_1 \mathbb{1}(B_3)] &\leq n\#R(\mathbf{t})\mathbb{E}_\tau \left(\frac{(N(T_L) + \#L(\mathbf{t}))\#R(\mathbf{t})e^\sigma}{n - 1} \cdot \frac{1}{N(T_L)} \mathbb{1}(C) \right) \\ &\leq (1 + o(1))(\#R(\mathbf{t}) + 1)^2\#L(\mathbf{t})e^\sigma. \end{aligned}$$

This completes the proof. ■

We record a minor rephrasing of Proposition 12.

Corollary 13. *In the setting of Proposition 12, suppose vertices $1, 2, \notin A_n$. Then*

$$\mathbb{P}_\tau(e_n \in \pi(1, 2)) = (1 + o(1))\#L(\mathbf{t})\#R(\mathbf{t}) \left[\frac{\log n}{n} e^{-\sigma} \right].$$

Proof. By symmetry over vertices $j \notin A_n \cup \{1\}$ we have $\mathbb{P}_\tau(e_n \in \pi(1, 2)) = \frac{\mathbb{E}_\tau M_1^*}{n-1-\#\mathbf{t}}$ where M_1^* is defined as M_1 but excluding vertices $j \in A_n$. Since $M_1 - M_1^* \leq \#\mathbf{t}$, the corollary follows from (25). ■

2.7. Conditional Variance of the Flow

In the setting of Proposition 12, we want to show that the flow $F_n(e)$ is close to its conditional expectation, and the natural way to express this is via the conditional variance.

Proposition 14. *In the setting of Proposition 12.*

$$\frac{\text{var}_\tau(F_n(e_n)) | \mathcal{B}_\tau(v_L^n, v_R^n) = \mathbf{t}_{A_n}}{[\mathbb{E}_\tau(F_n(e_n)) | \mathcal{B}_\tau(v_L^n, v_R^n) = \mathbf{t}_{A_n}]^2} \leq \frac{1}{\#L(\mathbf{t})} + \frac{1}{\#R(\mathbf{t})} + \frac{1}{\#L(\mathbf{t})\#R(\mathbf{t})} + o(1) \text{ as } n \rightarrow \infty.$$

This formulation emphasizes that the relative variance of the conditional distribution gets smaller as the size of the neighborhood gets bigger.

The key step in proving Proposition 14 is the following Proposition. To set up notation, we may assume the edge e_n is $(n - 1, n)$ and the label set A_n is $\{n - \#V(\mathbf{t}) + 1, \dots, n\}$. Recall $\pi(1, 2)$ denotes the shortest path from 1 to 2. If this path uses e_n then there is some entrance-exit pair $(\alpha, \beta) \in L(\mathbf{t}) \times R(\mathbf{t})$ recording the first and last vertices of the neighborhood visited by the path (here we identify vertices of \mathbf{t} and \mathbf{t}_{A_n}).

Proposition 15. *Let (α, β) and (γ, δ) be pairs in $L(\mathbf{t}) \times R(\mathbf{t})$. As $n \rightarrow \infty$ $\mathbb{P}_\tau(\pi(1, 2)$ contains e_n with entrance-exit pair (α, β) ; $\pi(3, 4)$ contains e_n with entrance-exit pair (γ, δ))*

$$\leq (1 + o(1))\kappa_{\alpha,\beta,\gamma,\delta} \left[\frac{\log n}{n} e^{-\sigma} \right]^2$$

where

$$\kappa_{\alpha,\beta,\gamma,\delta} = 2^{1\{\gamma=\alpha\}+1\{\delta=\beta\}}. \tag{34}$$

Proposition 15 (more precisely, the variant Proposition 19 described in Section 2.10) will be proved in Section 3. Intuitively we have “ $= (1 + o(1))$ ” instead of “ $\leq (1 + o(1))$ ”, but the fact that we need only prove an inequality is technically helpful. Also intuitively, the constant κ arises as

$$\kappa_{\alpha,\beta,\gamma,\delta} = \mathbb{E}[W_\alpha W_\beta W_\gamma W_\delta]$$

where W_α is the limit Exponential(1) r.v. arising (cf. Lemma 5) in the growth of the percolation process from source α using flows avoiding any other vertex of the neighborhood.

Proof of Proposition 14. The sum of $\kappa_{\alpha,\beta,\gamma,\delta}$ over all choices of $\alpha, \beta, \gamma, \delta$ works out as $\#L(\mathbf{t})\#R(\mathbf{t})(\#L(\mathbf{t}) + 1)(\#R(\mathbf{t}) + 1)$. So

$$\begin{aligned} \mathbb{P}_\tau(\boldsymbol{\pi}(1, 2) \text{ contains } e_n; \boldsymbol{\pi}(3, 4) \text{ contains } e_n) \\ \leq (1 + o(1))\#L(\mathbf{t})\#R(\mathbf{t})(\#L(\mathbf{t}) + 1)(\#R(\mathbf{t}) + 1) \left[\frac{\log n}{n} e^{-\sigma} \right]^2. \end{aligned}$$

Using Corollary 13, we get a covariance bound

$$\begin{aligned} \mathbb{P}_\tau(\boldsymbol{\pi}(1, 2) \text{ contains } e_n; \boldsymbol{\pi}(3, 4) \text{ contains } e_n) - \mathbb{P}_\tau^2(\boldsymbol{\pi}(1, 2) \text{ contains } e_n) \\ \leq (1 + o(1))\#L(\mathbf{t})\#R(\mathbf{t})(\#L(\mathbf{t}) + \#R(\mathbf{t}) + 1) \left[\frac{\log n}{n} e^{-\sigma} \right]^2. \quad (35) \end{aligned}$$

The contribution to $F_n(e_n)$ from source-destination pairs (i, j) where i or j is in \mathbf{t} is negligible, so we may replace $F_n(e_n)$ by

$$G_n(e_n) := \frac{1}{n} \sum_{(i,j)} \mathbb{1}\{e_n \in \boldsymbol{\pi}(i, j)\}$$

where here and below the sum is over ordered pairs of vertices in $[n] \setminus A_n$. Writing $e = e_n$,

$$\begin{aligned} \text{var}_\tau \sum_{(i,j)} \mathbb{1}\{e \in \boldsymbol{\pi}(i, j)\} &\leq \mathbb{E}_\tau \sum_{(i,j)} \mathbb{1}\{e \in \boldsymbol{\pi}(i, j)\} \\ &+ \sum_{(i,j)} \sum_{(i',j') \neq (i,j)} [\mathbb{P}_\tau(\boldsymbol{\pi}(i, j) \text{ contains } e; \boldsymbol{\pi}(i', j') \text{ contains } e) - \mathbb{P}_\tau^2(\boldsymbol{\pi}(i, j) \text{ contains } e)]. \end{aligned}$$

Using symmetry and a compatibility condition (a directed edge e cannot be in the shortest path from i to j and also in the shortest path from j to k) we find

$$\begin{aligned} \text{var}_\tau G_n(e) &\leq n^{-1} \mathbb{E}_\tau G_n(e) + 2n \mathbb{P}_\tau(\boldsymbol{\pi}(1, 2) \text{ contains } e; \boldsymbol{\pi}(1, 3) \text{ contains } e) \\ &+ n^2 [\mathbb{P}_\tau(\boldsymbol{\pi}(1, 2) \text{ contains } e; \boldsymbol{\pi}(3, 4) \text{ contains } e) - \mathbb{P}_\tau^2(\boldsymbol{\pi}(1, 2) \text{ contains } e)]. \end{aligned}$$

The first term is $O(n^{-1} \log n)$ by Proposition 12. Bounding the second term crudely by $2n \mathbb{P}_\tau(\boldsymbol{\pi}(1, 2) \text{ contains } e)$ and using Corollary 13 shows the second term is $O(\log n)$. So the dominant term is the third term, which by (35) shows

$$\text{var}_\tau G_n(e) \leq (1 + o(1))\#L(\mathbf{t})\#R(\mathbf{t})(\#L(\mathbf{t}) + \#R(\mathbf{t}) + 1) [\log n e^{-\sigma}]^2.$$

Combining with Proposition 12, we have established Proposition 14. ■

2.8. WLLN for a Local Functional

The point of Propositions 12 and 14 is that the normalized flow $F_n(e)/\log n$, which a priori involves the global structure of \mathcal{G}_n , can be approximated by a certain functional, $\phi_n^e(e)$ below, which depends only on the “local” structure of \mathcal{G}_n near e . It is a general fact that empirical (random) distributions of such “local functionals” on \mathcal{G}_n converge to the limit nonrandom distribution associated with the Yule process/PWIT mentioned in Section 2.3. Rather than

prove a general result in this context (for the general result in a different context see [2] Proposition 7) we will just derive the specific result we need, Proposition 16.

Fix $\tau > 0$. Recall from Section 2.4 the definition of the neighborhood $\mathcal{B}_\tau(e)$ of a directed edge e of \mathcal{G}_n . For each directed edge e define

$$\phi_n^\tau(e) = \#\mathcal{L}(\mathcal{B}_\tau(e))\#\mathcal{R}(\mathcal{B}_\tau(e)) \exp(-2\tau - L_e) \mathbb{1}\{L_e \leq \tau\} \mathbb{1}\{\mathcal{B}_\tau(e) \text{ is a tree}\}. \tag{36}$$

Define Φ_n^τ as the empirical measure on $[0, \infty)^2$ obtained by putting weight $1/n$ on each point $(L_e, \phi_n^\tau(e))$ associated with the edges e of \mathcal{G}_n for which $L_e \leq \tau$:

$$\Phi_n^\tau(\cdot, \cdot) = \frac{1}{n} \sum_{\text{directed } e} \mathbb{1}\{(L_e, \phi_n^\tau(e)) \in (\cdot, \cdot)\} \tag{37}$$

and define the mean measure

$$\bar{\Phi}_n^\tau(\cdot, \cdot) = \mathbb{E}\Phi_n^\tau(\cdot, \cdot).$$

Define a limit measure

$$\bar{\Phi}_\infty^\tau(\cdot, \cdot) = \int_0^\tau \mathbb{P}\left(\left(u, \frac{W_1^\tau W_2^\tau}{e^{2\tau}} e^{-u}\right) \in \cdot \times \cdot\right) du$$

where W_1^τ and W_2^τ are independent Geometric($e^{-\tau}$). So $\bar{\Phi}_\infty^\tau$ has total mass τ .

Proposition 16. *For any continuous test function $h : [0, \infty)^2 \rightarrow \mathbb{R}$ with compact support,*

$$\int h d\Phi_n^\tau \rightarrow \int h d\bar{\Phi}_\infty^\tau \text{ in } L^1.$$

The proof rests upon the following straightforward lemma. Although superficially similar to Proposition 15 in using vertices 1, 2, 3, 4 as typical vertices, their role here is different. The precise statement is a bit fussy because the neighborhood must be a tree in order for left and right sides to be well-defined.

Lemma 17. *Fix $\tau > 0$ and the directed edges (1, 2) and (3, 4).*

- (a) $n\mathbb{P}(L_{12} \leq \tau, L_{12} \in \cdot)$ converges vaguely to Lebesgue measure on $[0, \tau]$.
- (b) Uniformly on $\{L_{12} \leq \tau\}$

$$\mathbb{P}(\mathcal{B}_\tau(1, 2) \text{ is a tree} | L_{12}) \rightarrow 1 \tag{38}$$

as $n \rightarrow \infty$. The same holds for edges (1, 2) and (3, 4); that is, uniformly on the set $\{L_{12} \leq \tau\} \cap \{L_{34} \leq \tau\}$

$$\mathbb{P}(\{\mathcal{B}_\tau(1, 2) \text{ is a tree}\} \cap \{\mathcal{B}_\tau(3, 4) \text{ is a tree}\} | L_{12}, L_{34}) \rightarrow 1. \tag{39}$$

- (c) Let $0 < \ell_{12}, \ell_{34} \leq \tau$. Write $(\tilde{N}_n^\tau(1), \tilde{N}_n^\tau(2))$ for the numbers of vertices in the left and right sides of the neighborhood $\mathcal{B}_\tau(1, 2)$; define $(\tilde{N}_n^\tau(3), \tilde{N}_n^\tau(4))$ similarly for the neighborhood $\mathcal{B}_\tau(3, 4)$ (these are well-defined when the neighborhoods are trees).

Conditional on the event $\{\mathcal{B}_\tau(1, 2) \text{ is a tree}\} \cap \{\mathcal{B}_\tau(3, 4) \text{ is a tree}\} \cap \{L_{12} = \ell_{12}, L_{34} = \ell_{34}\}$ we have

$$\left(\tilde{N}_n^\tau(1), \tilde{N}_n^\tau(2), \tilde{N}_n^\tau(3), \tilde{N}_n^\tau(4)\right) \xrightarrow{d} (W_1^\tau, W_2^\tau, W_3^\tau, W_4^\tau)$$

where the W 's are independent $\text{Geometric}(e^{-\tau})$.

Proof of Proposition 16. The mean measure $\bar{\Phi}_n^\tau$ equals

$$\frac{1}{n} \times n(n-1) \mathbb{P}\left(L_{12} \leq \tau, L_{12} \in \cdot, \frac{\tilde{N}_n^\tau(1)\tilde{N}_n^\tau(2)}{e^{2\tau}} e^{-L_{12}} \in \cdot\right).$$

Because $n\mathbb{P}(L_{12} \leq \tau, L_{12} \in \cdot)$ converges vaguely to Lebesgue measure on $[0, \tau]$, Lemma 17 (here only vertices 1 and 2 are relevant) implies vague convergence $\bar{\Phi}_n^\tau \rightarrow \bar{\Phi}_\infty^\tau$ of mean measures. To get L^2 convergence it is enough to show that for a generic test function h we have

$$\text{var}\left(\int h d\bar{\Phi}_n^\tau\right) = \text{var}\left(\frac{1}{n} \sum_{e: L_e \leq \tau} h(L_e, \phi_n(e))\right) \rightarrow 0.$$

Expanding the right side as the variance-covariance sum, the contribution to variance from terms (e, e') with 4 distinct end-vertices tends to 0 by Lemma 17 and the fact that $n^2\mathbb{P}(L_{12} \leq \tau, \mathcal{B}_\tau(1, 2) \text{ is a tree}, L_{12} \in \cdot, L_{34} \leq \tau, \mathcal{B}_\tau(3, 4) \text{ is a tree}, L_{34} \in \cdot)$ converges vaguely to Lebesgue measure on $[0, \tau]^2$. The contribution from terms with three distinct end-vertices is bounded by

$$\frac{1}{n} \|h\|_\infty^2 \mathbb{E}\Delta_1^2(\tau) \rightarrow 0$$

where $\Delta_1(\tau)$ is the number of edges at 1 with length less than τ . And the contribution from pairs (e, e) is bounded by $\|h\|_\infty^2 \mathbb{P}(L_{12} \leq \tau) \rightarrow 0$. ■

2.9. Completing the Proof of Theorems 1 and 2

The remainder of the proof uses only “soft” arguments.

Propositions 12 and 14 were stated for fixed $\mathbf{t} \in \mathbf{T}_\tau$, but it is clear that convergence is uniform over subsets of \mathbf{T}_τ on which the length of distinguished edge is bounded and the number of vertices is bounded. Rephrasing those Propositions gives, after some obvious manipulations:

Corollary 18. Fix $\tau > 0$ and $K < \infty$. As $n \rightarrow \infty$

$$\begin{aligned} \mathbb{E}\left(\frac{F_n(e_n)}{\log n} \middle| \mathcal{B}_\tau(e_n)\right) - \phi_n^\tau(e_n) &\rightarrow 0 \\ \text{var}\left(\frac{F_n(e_n)}{\log n} \middle| \mathcal{B}_\tau(e_n)\right) &\leq 3K^3 e^{-\tau} + o(1) \end{aligned}$$

uniformly over e_n satisfying

$$\mathcal{B}_\tau(e_n) \text{ is a tree ; } L_{e_n} \leq \tau; \quad \max(\#L(\mathcal{B}_\tau(e_n)), \#R(\mathcal{B}_\tau(e_n))) \leq Ke^\tau.$$

Now fix $\varepsilon > 0$. Applying Chebyshev’s inequality (conditional on $\mathcal{B}_\tau(e_n)$) and taking limits,

$$\begin{aligned} \limsup_n \mathbb{E} \frac{1}{n} \sum_{\text{directed } e} \mathbb{1}\{L_e \leq \tau, \mathcal{B}_\tau(e) \text{ is a tree}\} \mathbb{1}\{\max(\#L(\mathcal{B}_\tau(e)), \#R(\mathcal{B}_\tau(e))) \\ \leq Ke^\tau\} \mathbb{1}\left\{\left|\frac{F_n(e)}{\log n} - \phi_n^\tau(e)\right| > \varepsilon\right\} \leq 3\varepsilon^{-2}K^3e^{-\tau}. \end{aligned}$$

And using Lemma 17

$$\begin{aligned} \limsup_n \mathbb{E} \frac{1}{n} \sum_{\text{directed } e} \mathbb{1}\{L_e \leq \tau\} \mathbb{1}\{\mathcal{B}_\tau(e) \text{ is a tree}\} \mathbb{1}\{\max(\#L(\mathcal{B}_\tau(e)), \#R(\mathcal{B}_\tau(e))) > Ke^\tau\} \\ = \lim_n n \mathbb{E} \mathbb{1}\{L_{12} \leq \tau\} \mathbb{1}\{\mathcal{B}_\tau(1, 2) \text{ is a tree}\} \mathbb{1}\{\max(\#L(\mathcal{B}_\tau(1, 2)), \#R(\mathcal{B}_\tau(1, 2))) > Ke^\tau\} \\ = \tau \mathbb{P}(\max(W_1^\tau, W_2^\tau) > Ke^\tau) \\ \leq 2\tau \exp(-K + e^{-\tau}). \end{aligned}$$

Combining these bounds,

$$\begin{aligned} \limsup_n \mathbb{E} \frac{1}{n} \sum_{\text{directed } e} \mathbb{1}\{L_e \leq \tau\} \mathbb{1}\{\mathcal{B}_\tau(e) \text{ is a tree}\} \mathbb{1}\left\{\left|\frac{F_n(e)}{\log n} - \phi_n^\tau(e)\right| > \varepsilon\right\} \\ \leq 3\varepsilon^{-2}K^3e^{-\tau} + 2\tau \exp(-K + e^{-\tau}). \end{aligned}$$

Apply this with $K = \tau$ and then let $\tau \rightarrow \infty$:

$$\lim_\tau \limsup_n \mathbb{E} \frac{1}{n} \sum_{\text{directed } e} \mathbb{1}\{L_e \leq \tau\} \mathbb{1}\{\mathcal{B}_\tau(e) \text{ is a tree}\} \mathbb{1}\left\{\left|\frac{F_n(e)}{\log n} - \phi_n^\tau(e)\right| > \varepsilon\right\} = 0. \tag{40}$$

Also by Lemma 17 for each fixed τ

$$\limsup_n \mathbb{E} \frac{1}{n} \sum_{\text{directed } e} \mathbb{1}\{L_e \leq \tau, \mathcal{B}_\tau(e) \text{ not a tree}\} = 0 \tag{41}$$

Recall from Section 1.2 that Theorem 2 concerns measures ψ_n and ψ on the product space $[0, \infty) \times (0, \infty)$ (the first coordinate is length, the second is flow). Vague convergence $\nu_n \rightarrow \nu$ of measures on $[0, \infty) \times (0, \infty)$ is defined to mean $\int hd\nu_n \rightarrow \int hd\nu$ for bounded continuous test functions $h : [0, \infty) \times (0, \infty) \rightarrow \mathbb{R}$ with compact support, but in checking vague convergence we need consider only test functions which also have finite Lipschitz norm $\|h\|_{\text{Lip}}$. Fix such a test function h . We shall prove the assertion of Theorem 2, restated here as

$$\int hd\psi_n \rightarrow \int hd\psi \text{ in } L^1. \tag{42}$$

To prove this, recall from (3) and (37) the definitions

$$\psi_n(\cdot, \cdot) = \frac{1}{n} \sum_{\text{directed } e} \mathbb{1}\{(L_e, F_n(e)/\log n) \in (\cdot, \cdot)\}$$

$$\Phi_n^\tau(\cdot, \cdot) = \frac{1}{n} \sum_{\text{directed } e} \mathbb{1}\{(L_e, \phi_n^\tau(e)) \in (\cdot, \cdot)\}.$$

Fix h with support contained in $[0, \tau_0] \times [0, \infty)$. Then for $\tau > \tau_0$

$$\begin{aligned} \left| \int hd\psi_n - \int hd\Phi_n^\tau \right| &\leq 2\|h\|_\infty \frac{1}{n} \sum_{\text{directed } e} \mathbb{1}\{L_e \leq \tau_0, \mathcal{B}_\tau(e) \text{ is a tree}\} \mathbb{1}\left\{ \left| \frac{F_n(e)}{\log n} - \phi_n^\tau(e) \right| > \varepsilon \right\} \\ &\quad + 2\|h\|_\infty \frac{1}{n} \sum_{\text{directed } e} \mathbb{1}\{L_e \leq \tau_0, \mathcal{B}_\tau(e) \text{ not a tree}\} \\ &\quad + \varepsilon \|h\|_{\text{Lip}} \frac{1}{n} \sum_{\text{directed } e} \mathbb{1}\{L_e \leq \tau_0\}. \end{aligned}$$

Because $\mathbb{E} \frac{1}{n} \sum_{\text{directed } e} \mathbb{1}\{L_e \leq \tau_0\} \rightarrow \tau_0$ we can use (40,41) to deduce

$$\lim_{\tau} \limsup_n \mathbb{E} \left| \int hd\psi_n - \int hd\Phi_n^\tau \right| \leq \varepsilon \|h\|_{\text{Lip}} \tau_0.$$

Because ε is arbitrary, this shows

$$\lim_{\tau} \limsup_n \mathbb{E} \left| \int hd\psi_n - \int hd\Phi_n^\tau \right| = 0.$$

Proposition 16 allows us to replace the random measure Φ_n^τ by the limit mean measure $\bar{\Phi}_\infty^\tau$:

$$\lim_{\tau} \limsup_n \mathbb{E} \left| \int hd\psi_n - \int hd\bar{\Phi}_\infty^\tau \right| = 0.$$

But $\bar{\Phi}_\infty^\tau \rightarrow \psi$ vaguely as $\tau \rightarrow \infty$, and so we have proved (42), which is Theorem 2.

To prove Theorem 1, recall that the fact (16) $\mathbb{E}\bar{D}_n \sim \log n$ becomes, via (4),

$$\mathbb{E} \int \ell_y \psi_n(d\ell, dy) \rightarrow \int \ell_y \psi(d\ell, dy) = 1.$$

This enables us to extend the L^1 convergence (42) from continuous h with compact support to continuous $h \geq 0$ satisfying $\sup_{\ell,y} \frac{h(\ell,y)}{\ell_y} < \infty$. Using such functions to approximate the function $\mathbb{1}\{\ell > \varepsilon, y > z\}$ shows

$$\frac{1}{n} \#\{e : L_e > \varepsilon, F_n(e) > z \log n\} \rightarrow_{L^1} \psi((\varepsilon, \infty) \times (z, \infty)).$$

Because $\mathbb{E} \frac{1}{n} \#\{e : L_e \leq \varepsilon\} \rightarrow \varepsilon$ we can let $\varepsilon \rightarrow 0$ and deduce

$$\frac{1}{n} \#\{e : F_n(e) > z \log n\} \rightarrow_{L^1} \psi((0, \infty) \times (z, \infty))$$

which is the first assertion of Theorem 1.

2.10. Distance-Based Truncation of Flows

To avoid notational complications, the exposition above omitted one technical point. Recall that path-lengths $D(i, j)$ are $\log n \pm O(1)$ in probability. In seeking to prove Proposition 15 there are technical difficulties with unusually long paths, which we will handle by truncating them out. Precisely, instead of proving Proposition 15 we will prove

Proposition 19. *Fix $B < \infty$. Let (α, β) and (γ, δ) be pairs in $L(\mathbf{t}) \times R(\mathbf{t})$. Conditionally on $\mathcal{B}_\tau(n - 1, n) = \mathbf{t}_{A_n}$, as $n \rightarrow \infty$*

$$\mathbb{P}(\boldsymbol{\pi}(1, 2) \text{ contains } e_n \text{ with entrance-exit pair } (\alpha, \beta); \text{len}(\boldsymbol{\pi}(1, 2)) \leq \log n + B;$$

$$\boldsymbol{\pi}(3, 4) \text{ contains } e_n \text{ with entrance-exit pair } (\gamma, \delta), \text{len}(\boldsymbol{\pi}(3, 4)) \leq \log n + B)$$

$$\leq (1 + o(1))\kappa_{\alpha, \beta, \gamma, \delta} \left[\frac{\log n}{n} e^{-\sigma} \right]^2$$

where

$$\kappa_{\alpha, \beta, \gamma, \delta} = 2^{\mathbb{1}\{\gamma=\alpha\} + \mathbb{1}\{\delta=\beta\}}. \tag{43}$$

In this section, we explain (omitting some details at the end) why it is enough to prove Proposition 19 in place of Proposition 15. Consider the analog of flow $F_n(e)$ when contributions from paths of length $> \log n + B$ are ignored:

$$F_n^{[B]}(e) := \frac{1}{n} \sum_{i \in [n]} \sum_{j \in [n], j \neq i} \mathbb{1}\{e \in \boldsymbol{\pi}(i, j)\} \mathbb{1}\{\text{len}(\boldsymbol{\pi}(i, j)) \leq \log n + B\} \leq F_n(e).$$

An easy argument shows that for large B the global effect of truncation is negligible:

Lemma 20.

$$\lim_{B \rightarrow \infty} \limsup_n \mathbb{E} \frac{1}{n} \sum_e L_e \frac{F_n(e) - F_n^{[B]}(e)}{\log n} = 0.$$

Proof. Recall $D(i, j) = \text{len}(\boldsymbol{\pi}(i, j))$. Calculating the effect of truncation on edge-flows and on source-destination distances gives the identity

$$\sum_e L_e (F_n(e) - F_n^{[B]}(e)) = \frac{1}{n} \sum_{j \neq i} D(i, j) \mathbb{1}\{D(i, j) > \log n + B\}.$$

Taking expectations and using symmetry

$$\mathbb{E} \frac{1}{n} \sum_e L_e (F_n(e) - F_n^{[B]}(e)) = \frac{n(n-1)}{n^2} \mathbb{E} D(1, 2) \mathbb{1}\{D(1, 2) > \log n + B\}.$$

The result now follows from the mean and variance limits at (15). ■

Now we can choose $B_n \uparrow \infty$ sufficiently slowly that (by Lemma 20)

$$\mathbb{E} \frac{1}{n} \sum_e L_e \frac{F_n(e) - F_n^{[B_n]}(e)}{\log n} \rightarrow 0 \tag{44}$$

and such that

the conclusion of Proposition 19 holds for B_n in place of B .

The idea is now to repeat the arguments in Sections 2.7–2.9 using the truncated flow $F_n^{[B_n]}(e_n)$ in place of $F_n(e_n)$. This will establish Theorem 2 for the truncated flows, but then (44) establishes it for untruncated flows. The arguments would go through unchanged if we knew the truncated version of the conditional mean estimate of Proposition 12:

$$\mathbb{E}_\tau (F_n^{[B_n]}(e_n)) = (1 + o(1)) \#L(\mathbf{t}) \#R(\mathbf{t}) e^{-\sigma} \log n. \tag{45}$$

Of course the conditional upper bound “ $\leq (1 + o(1))$ ” in (45) follows from Proposition 12, but we need the lower bound “ $\geq (1 + o(1))$ ” in (45) in order to go from the upper bound on second moment to the upper bound on variance – cf. (35). However, from the conditional upper bound and because (44) implies a lower bound for unconditional expectation, Markov’s inequality implies that the conditional lower bound in (45) holds for all neighborhoods $\mathcal{B}_\tau(n - 1, n)$ excluding some occurring with probability $\rightarrow 0$ as $n \rightarrow \infty$. And this is enough to complete the proof of Theorem 2.

3. THE VARIANCE ESTIMATE

This section is devoted to the proof of Proposition 19, which will complete the proof of Theorem 2. The proof occupies Sections 3.1–3.3, with one result (Proposition 23) postponed and proved in Section 3.5.

Let us repeat the “conditioned” setting that we work in, throughout the section. There is a fixed tree \mathbf{t} with distinguished directed edge e . In the network \mathcal{G}_n , we fix edge $e_n = (n - 1, n)$ and label set $A_n = \{n - \#\mathbf{t} + 1, \dots, n\}$. Label \mathbf{t} as \mathbf{t}_{A_n} so that e is labeled $(n - 1, n)$. Then condition on $\mathcal{B}_\tau(n - 1, n) = \mathbf{t}_{A_n}$. Recall that Lemma 8 tells us the effect of this conditioning. In particular, for $i \notin A_n, j \in A_n$ the length of the edge-segment (i, j) from i to the boundary of the neighborhood has Exponential($1/n$) distribution, independently for different edges.

Roughly speaking, the issue in proving Proposition 19 is to estimate the dependence between the events

- (i) the shortest path $\pi(1, 2)$ between vertex 1 and vertex 2 uses edge e_n
- (ii) the shortest path $\pi(3, 4)$ between vertex 3 and vertex 4 uses edge e_n .

Corollary 13 tells us the asymptotic probabilities of these events, so a natural approach is to condition on (i) and seek to calculate the conditional probability of (ii). Now (i) breaks into two assertions:

- (ia) there is a short path (length $s = \log n \pm O(1)$) from 1 to 2 via e ;
- (ib) there is no shorter path from 1 to 2.

Now conditioning on (ia) can be implemented by conditioning on all edges in the path, and this does not affect lengths of other edges of \mathcal{G}_n . But event (ib) implicitly specifies that alternate short routes do not exist, and the effect of this conditioning on other edge-lengths of \mathcal{G}_n (while intuitively small) seems hard to handle rigorously. Instead, we shall carefully organize an argument to avoid ever conditioning on any “shortest path” event. In outline, the argument has three steps.

- Calculate chance of existence of paths π_{12}, π_{34} of specified lengths through e_n (Section 3.1).
- Conditional on existence of such paths, what is the chance they are the shortest paths? The percolation processes from vertices $i = 1, 2, 3, 4$ avoiding edge e_n become approximately size-biased Yule processes (Section 3.4) reaching $\tilde{W}_i e'$ vertices in distance t ; so the chance of a path from 1 to 2 of length t avoiding e_n is approximately $\exp(-\tilde{W}_1 \tilde{W}_2 e')$ (Section 3.5).
- These two estimates are combined in Section 3.3.

3.1. Joint Intensity for Two Short Paths Through e

For this section, we introduce some handy notation. We will describe the relationship (for an event B , a random variable T , and a function f)

$$\mathbb{P}(B, T \in (t, t + dt)) = f(t)dt$$

by the phrase

the event $[[B, T = t]]$ has intensity $f(t)$.

But we will describe events in words, rather than inventing ad hoc symbols, using the brackets $[[.]]$ to highlight the verbal description of the event.

For $\alpha \in L(\mathbf{t})$ and $\beta \in R(\mathbf{t})$ define $h_{\alpha,\beta}(s_1, s_2)$ to be the intensity (in s_1 and s_2) of the event:

[[there exists a path from 1 to v_L which crosses the neighborhood boundary at time s_1 and then first hits vertex α ;
and there exists a path from 2 to v_R which crosses the neighborhood boundary at time s_2 and then first hits vertex β .]]

Note that such paths, linked via the path from α to β in the neighborhood, specify a path π_{12} of length $s_1 + s_2 + \sigma$ from 1 to 2 via e . This path may or may not be the shortest path from 1 to 2, depending on lengths of other edges in \mathcal{G}_n .

Given also $\gamma \in L(\mathbf{t})$ and $\delta \in R(\mathbf{t})$, define an event which replicas the event above:

[[there exists a path from 3 to v_L which crosses the neighborhood boundary at time t_1 and then first hits vertex γ ;
and there exists a path from 4 to v_R which crosses the neighborhood boundary at time t_2 and then first hits vertex δ .]]

Again, such paths specify a path π_{34} from 3 to 4 via e . In order for it to be possible that both π_{12} and π_{34} are shortest paths, the following simple compatibility conditions must hold.

- (i) if the paths from 1 and from 3 meet at some vertex v_* outside the neighborhood, then they must coincide from v_* to the neighborhood.

- (ii) if the paths from 2 and from 4 meet at some vertex v^* outside the neighborhood, then they must coincide from v^* to the neighborhood.
- (iii) the set of vertices visited by the paths from 1 and 3 must be disjoint from the set of vertices visited by the paths from 2 and 4.

Define $H_{\alpha,\beta,\gamma,\delta}(s_1, s_2, t_1, t_2)$ to be the intensity of both events happening and the compatibility conditions holding.

Lemma 21.

$$h_{\alpha,\beta}(s_1, s_2) \leq \frac{1}{n^2} \exp(s_1 + s_2) \tag{46}$$

$$H_{\alpha,\beta,\gamma,\delta}(s_1, s_2, t_1, t_2) \leq \frac{\kappa}{n^4} \exp(s_1 + s_2 + t_1 + t_2). \tag{47}$$

Here $\kappa = \kappa_{\alpha,\beta,\gamma,\delta}$ as at (34).

Note these are inequalities for finite n . Heuristically they are asymptotic equalities in the ranges of interest to us.

Proof. The argument is based on exact formulas, starting with the following. The intensity of the event

[[there exists a path from 1 to v_L which crosses the neighborhood boundary at time s_1 and then first hits vertex α , taking exactly $k + 1$ steps]]

$$= \frac{(n - 1 - \#\mathbf{t})_k}{n^k} \frac{s_1^k}{k!} \frac{1}{n} \exp\left(\frac{-s_1 - b(\alpha)}{n}\right) \tag{48}$$

where we recall that $b(\alpha)$ is the distance from α to the neighborhood boundary. To prove this, take $0 < u_1 < u_2 < \dots < u_k < s_1$ and consider the probability that the j 'th step ends at distance $[u_j, u_j + du_j]$ from 1 and the boundary crossing is at distance $[s_1, s_1 + ds_1]$. This probability is

$$(n - 1 - \#\mathbf{t})_k \times \prod_{j=1}^k \frac{1}{n} \exp(-(u_j - u_{j-1})/n) du_j \times \frac{1}{n} \exp(-(s_1 + b(\alpha) - u_k)/n) ds_1 \tag{49}$$

where the first term indicates number of choices of k intermediate vertices, and the other terms are the Exponential(mean n) densities of edge-lengths. Because

$$\int \dots \int_{0 < u_1 < u_2 < \dots < u_k < s_1} du_1 \dots du_k = s_1^k / k! \tag{50}$$

we deduce (48).

The first and last terms of (48) are ≤ 1 . Summing over k shows that the intensity of

[[there exists a path from 1 to v_L which crosses the neighborhood boundary at time s_1 and then first hits vertex α]]

is $\leq \frac{1}{n} \exp(s_1)$. Combining this with the similar argument on the right side of the neighborhood gives (46).

To prove (47), because the left and right sides have analogous arguments, the issue is to show that the intensity of the event

[[there exists a path from 1 to v_L which crosses the neighborhood boundary at time s_1 and then first hits vertex α ;
there exists a path from 3 to v_L which crosses the neighborhood boundary at time t_1 and then first hits vertex γ]]

$$\leq 2^{1(\alpha=\gamma)} \frac{1}{n^2} \exp(s_1 + t_1). \tag{51}$$

Now in the case $\alpha \neq \gamma$, or for the contribution to the case $\alpha = \gamma$ from disjoint paths, we get density $\leq \frac{1}{n^2} \exp(s_1 + t_1)$ by arguments analogous to above. Let us show details of the more interesting case where $\alpha = \gamma$ and we consider the contribution from merging paths. Consider the intensity (in r, s_1, t_1) of the event:

[[there exists a path from 1 to v_L which crosses the neighborhood boundary at time s_1 and then first hits vertex α , taking exactly $k_1 + k_2 + 1$ steps;
there exists a path from 3 to v_L which crosses the neighborhood boundary at time t_1 and then first hits vertex α , taking exactly $k_3 + k_2 + 1$ steps;
these paths merge at some vertex v_* at distance r from the neighborhood boundary, the path from v_* to α using $k_2 + 1$ steps.]]

Analogous to (48) this intensity has an exact formula

$$\frac{(n - 1 - \#\mathbf{t})_{k_1+k_2+k_3-1}}{n^{k_1+k_2+k_3}} \frac{(s_1 - r)^{k_1}}{k_1!} \frac{(t_1 - r)^{k_3}}{k_3!} \frac{r^{k_2}}{k_2!} \frac{1}{n} \exp\left(\frac{-s_1 - t_1 + r - b(\alpha)}{n}\right).$$

This intensity is bounded by

$$\frac{1}{n^2} \frac{(s_1 - r)^{k_1} (t_1 - r)^{k_3} r^{k_2}}{k_1! k_3! k_2!}.$$

Summing over (k_1, k_2, k_3) shows that the intensity of

[[there exist paths from 1 (resp. 3) to v_L which cross the neighborhood boundary at time s_1 (resp. t_1) and then first hit vertex α , having merged at distance r before the boundary]]

is bounded by $\frac{1}{n^2} \exp(s_1 + t_1 - r)$. Integrating over r shows that the contribution to (51) from merging paths is $\leq \frac{1}{n^2} \exp(s_1 + t_1)$. This establishes (51). ■

3.2. A Cox Point Process

Here we introduce a process and a lemma; how the process arises will be seen in Section 3.5.

Take independent random variables \tilde{W}_1, \tilde{W}_2 with probability density $w e^{-w}$ on $0 < w < \infty$. Consider the Cox point process defined by: conditional on \tilde{W}_1, \tilde{W}_2 the points form a Poisson process of rate $(\tilde{W}_1 \tilde{W}_2 e^s, -\infty < s < \infty)$. (This is a particular case of the general class of Cox processes, which are Poisson processes built over a random intensity). Let \tilde{L} be the position of the leftmost point of this Cox process.

Lemma 22. $\int_{-\infty}^{\infty} e^s \mathbb{P}(\tilde{L} > s) ds = 1.$

Proof. Note $\mathbb{P}(\tilde{L} > s | \tilde{W}_1, \tilde{W}_2) = \exp(-\tilde{W}_1 \tilde{W}_2 e^s)$ and so

$$\begin{aligned} \int_{-\infty}^{\infty} e^s \mathbb{P}(\tilde{L} > s) ds &= \int_{-\infty}^{\infty} e^s \mathbb{E} \exp(-\tilde{W}_1 \tilde{W}_2 e^s) ds \\ &= \int_0^{\infty} \mathbb{E} \exp(-\tilde{W}_1 \tilde{W}_2 u) du \quad \text{setting } u = e^s \\ &= \mathbb{E} \left[\frac{1}{\tilde{W}_1 \tilde{W}_2} \right] = \left[\mathbb{E} \frac{1}{\tilde{W}_1} \right]^2 = 1. \end{aligned}$$

■

3.3. Conditional Distributions of Other Short Routes

Recall we are conditioning on $\mathcal{B}_\tau(n-1, n) = \mathbf{t}_{A_n}$, though this is not indicated in notation.

Fix times s_1, s_2, t_1, t_2 and vertices (maybe the same) $\alpha, \gamma \in L(\mathbf{t}_{A_n})$ and $\beta, \delta \in R(\mathbf{t}_{A_n})$. Recall from Section 3.1 that $H_{\alpha, \beta, \gamma, \delta}(s_1, s_2, t_1, t_2)$ denotes the intensity of the event

[[there exists a path π_{12} from 1 to 2 via e_n , where the path from 1 crosses the boundary of the neighborhood at time s_1 and then first hits vertex α , while the reverse path from 2 crosses the boundary of the neighborhood at time s_2 and then first hits vertex β ; similarly there exists a path π_{34} from 3 to 4 via e_n , where the path from 3 crosses the boundary of the neighborhood at time t_1 and then first hits vertex γ , while the reverse path from 4 crosses the boundary of the neighborhood at time t_2 and then first hits vertex δ]]

together with certain compatibility conditions. We will write $(\cdot | s_1, s_2, t_1, t_2)$ to denote conditioning on this event. For such paths we have

$$\text{len}(\pi_{12}) = s_1 + s_2 + \sigma, \quad \text{len}(\pi_{34}) = t_1 + t_2 + \sigma,$$

and we write $\text{len}(\pi_{12})$ and $\text{len}(\pi_{34})$ for these sums.

Write S_{12} (resp. S_{34}) for the length of the shortest path from 1 to 2 (resp. from 3 to 4) that does not use edge e . Let us first show that Proposition 19 reduces to the following proposition.

Proposition 23. *Fix $B < \infty$. Uniformly on $\{\max(\text{len}(\pi_{12}), \text{len}(\pi_{34})) \leq \log n + B\}$, as $n \rightarrow \infty$*

$$\begin{aligned} &\mathbb{P}(S_{12} > \text{len}(\pi_{12}), S_{34} > \text{len}(\pi_{34}) | s_1, s_2, t_1, t_2) \\ &\leq \mathbb{P}(\xi_1^{12} > \text{len}(\pi_{12}) - \log n, \xi_1^{34} > \text{len}(\pi_{34}) - \log n) + o(1) \end{aligned}$$

where $(\xi_1^{12}, \xi_2^{12}, \dots)$ and $(\xi_1^{34}, \xi_2^{34}, \dots)$ are independent Cox point processes as described in Section 3.2.

We will prove this in sections 3.4–3.5, but let us first show how to deduce Proposition 19 from Lemma 21 and Proposition 23.

Proof of Proposition 19. A path π_{12} via e_n using entrance-exit pair (α, β) with $\text{len}(\pi_{12}) = s_1 + \sigma + s_2$ is created as in the definition of $h_{\alpha, \beta}(s_1, s_2)$ from two paths with lengths-to-boundary s_1 and s_2 . Create π_{34} similarly, using paths of lengths t_1 and t_2 . The quantity in Proposition 19

$$\begin{aligned} &\mathbb{P}(\boldsymbol{\pi}(1, 2) \text{ contains } e_n \text{ with entrance-exit pair } (\alpha, \beta); \text{len}(\boldsymbol{\pi}(1, 2)) \leq \log n + B; \\ &\boldsymbol{\pi}(3, 4) \text{ contains } e_n \text{ with entrance-exit pair } (\gamma, \delta), \text{len}(\boldsymbol{\pi}(3, 4)) \leq \log n + B) \end{aligned}$$

can be calculated in terms of the intensity H at (47) as

$$\iiint \mathbb{1}\{s_1 + s_2 + \sigma \leq \log n + B\} \mathbb{1}\{t_1 + t_2 + \sigma \leq \log n + B\} \mathbb{P}(\boldsymbol{\pi}(1, 2) = \pi_{12}, \boldsymbol{\pi}(3, 4) = \pi_{34} | s_1, s_2, t_1, t_2) H_{\alpha, \beta, \gamma, \delta}(s_1, s_2, t_1, t_2) ds_1 ds_2 dt_1 dt_2. \tag{52}$$

Now

$$\mathbb{P}(\boldsymbol{\pi}(1, 2) = \pi_{12}, \boldsymbol{\pi}(3, 4) = \pi_{34} | s_1, s_2, t_1, t_2) \leq \mathbb{P}(S_{12} > \text{len}(\pi_{12}), S_{34} > \text{len}(\pi_{34}) | s_1, s_2, t_1, t_2),$$

this being an inequality because there might be shorter paths using e . Bounding the right side using Proposition 23 gives

$$\mathbb{P}(\boldsymbol{\pi}(1, 2) = \pi_{12}, \boldsymbol{\pi}(3, 4) = \pi_{34} | s_1, s_2, t_1, t_2) \leq \mathbb{P}(\xi_1^{12} > s^1) \mathbb{P}(\xi_1^{34} > t^1) + \varepsilon_{n,B}$$

where $\lim_n \varepsilon_{n,B} = 0$ and where

$$s^1 = s_1 + s_2 + \sigma - \log n; \quad t^1 = t_1 + t_2 + \sigma - \log n.$$

To upper bound (52), first fix s^1 and t^1 and calculate

$$\begin{aligned} & \int_{s_1+s_2+\sigma=\log n+s^1} \int_{t_1+t_2+\sigma=\log n+t^1} H_{\alpha, \beta, \gamma, \delta}(s_1, s_2, t_1, t_2) ds_1 dt_1 \\ & \leq \int_{s_1+s_2+\sigma=\log n+s^1} \int_{t_1+t_2+\sigma=\log n+t^1} \frac{\kappa}{n^4} \exp(s_1 + s_2 + t_1 + t_2) ds_1 dt_1 \text{ by Lemma 21} \\ & = \frac{\kappa}{n^2} \exp(s^1 + t^1 - 2\sigma) \times (\log n + s^1 - \sigma)(\log n + t^1 - \sigma) \\ & \leq \kappa e^{-2\sigma} \frac{(\log n + B)^2}{n^2} \exp(s^1 + t^1) \end{aligned}$$

where in the final line we assume $\max(s^1, t^1) \leq B$. Thus the quantity (52) is bounded by

$$\begin{aligned} & \kappa \frac{(\log n + B)^2}{n^2} e^{-2\sigma} \int_{-\infty}^B \int_{-\infty}^B (\mathbb{P}(\xi_1^{12} > s^1) \mathbb{P}(\xi_1^{34} > t^1) + \varepsilon_{n,B}) e^{s^1} ds^1 e^{t^1} dt^1 \\ & \leq \kappa \frac{(\log n + B)^2}{n^2} e^{-2\sigma} (1 + \varepsilon_{n,B} e^{2B}) \text{ by Lemma 22} \end{aligned}$$

establishing Proposition 19. ■

3.4. Size-Biasing the Percolation Process and Yule Process

What remains to be proved is Proposition 23. This section builds up to proving a result, Proposition 26, about the number of vertices seen by the percolation process from vertex 1 when we condition on existence of a short path of specified length from vertex 1. This will be used to prove Proposition 23 in the next section.

3.4.1. *Some Terminology.* Let us quickly revisit the structures (Section 2.1) within \mathcal{G}_n associated with percolation from vertex 1 and introduce more precise terminology. The percolation tree itself is the spanning tree consisting of all edges in the shortest paths $\pi(1, j)$, $2 \leq j \leq n$. The percolation tree process tells us at time (time = distance) t the subtree on vertices within distance t from vertex 1. And the percolation counting process $N_n(t)$ at (13) tells us at time t the number of vertices within distance t from vertex 1. We can use the same terminology for the Yule process of Section 2.3; the process $(N_\infty(t), 0 \leq t < \infty)$ is the Yule counting process. The underlying continuous-time branching process, run until time t and then regarded as a random tree with edge-lengths, is the Yule tree process at time t . This process run to time ∞ is the Yule tree or PWIT, a random infinite tree with edge-lengths.

3.4.2. *Heuristics for Size-Biasing.* Associated with the Yule counting process is the limit (Lemma 4(b)) random variable $W := \lim_t e^{-t} N_\infty(t)$ with probability density e^{-w} on $0 < w < \infty$. What can we say about the Yule tree conditioned on it having a vertex at some specified large distance t_0 ? The probability of this event given W is approximately proportional to W , so the posterior density of W given this event becomes approximately $w e^{-w}$. This is a basic instance of size-biasing. But instead of relying on Bayes calculations for single random variables, we describe next the more elegant approach to size-biasing the whole process based on a probabilistic construction (this type of construction is widely used in modern branching process theory [16]). In this method, the density $w e^{-w}$ arises as the density of the sum of two independent Exponential(1) random variables.

3.4.3. *The size-biased Yule process.* We are working toward a result of the type
 the $n \rightarrow \infty$ limit of the size-biased percolation process is the size-biased Yule process.

Before stating the result (Proposition 26) let us define and derive simple properties of the limit process, without justifying the “size-biased” name.

On the half-line \mathbb{R}^+ , put a “root” vertex at the origin and other vertices at the points $(P_i, i \geq 1)$ of a rate 1 Poisson process. Make each of these vertices the root of a Yule tree. Regarding the resulting structure as a random infinite tree with edge-lengths, call it the size-biased Yule tree, with root at the origin. Given a distance t , the size-biased Yule process at t is the subtree on vertices at distance less than t from the root, illustrated in Fig. 3. The associated counting process, giving the number of vertices at distance less than t from the root, is

$$\tilde{N}(t) = N_0(t) + \sum_{i:P_i \leq t} N_i(t - P_i) \tag{53}$$

where (N_0, N_1, N_2, \dots) are the Yule counting processes associated with the constituent Yule trees.

Call the original half-line the distinguished path to infinity. We will also use, for technical reasons, the variation where the distinguished path is cut at some at some large finite distance s from the origin, so its counting process is

$$\tilde{N}^s(t) = N_0(t) + \sum_{i:P_i \leq \min(s,t)} N_i(t - P_i).$$




Fig. 3. Size-biased Yule process.

We collect below some simple facts about these two processes. Our main aim is to understand the limiting behavior of $\tilde{N}(t)$ for large t , and similarly, the behavior of $\tilde{N}^s(t)$ for large s and t .

Lemma 24.

(a) There exists a random variable \tilde{W} with probability density we^{-w} , $w > 0$ such that

$$\lim_{t \rightarrow \infty} e^{-t} \tilde{N}(t) = \tilde{W}; \quad \lim_{s \rightarrow \infty} \lim_{t \rightarrow \infty} e^{-t} \tilde{N}^s(t) = \tilde{W} \tag{54}$$

where the convergence holds a.s. and in L^2 .

(b) For any $c, s, t > 0$

$$\mathbb{E} \tilde{N}^s(t) (c + \tilde{N}^s(t) + \#\{i : t \leq P_i \leq s\}) \leq 6e^{2t} + 2e^t (c + s).$$

Proof. The construction (53) implies that the size biased process $\tilde{N}(\cdot)$ can be represented in terms of the sum of two independent Yule processes

$$\tilde{N}(t) = N_0(t) + (N'(t) - 1) \tag{55}$$

because the contribution from the distinguished path to infinity behaves as another Yule process rooted at the origin, with the distinguished path representing the reproduction times of the initial ancestor. We subtract 1 to avoid double counting the root. Lemma 4(b) says we have independent Exponential(1) limits (a.s. and in L^2)

$$W_0 := \lim_t e^{-t} N_0(t); \quad W' = \lim_t e^{-t} N'(t)$$

and (54) follows easily. For (b), because $\tilde{N}^s(t)$ is independent of $\#\{i : t \leq P_i \leq s\}$, the quantity under consideration equals

$$(c + (s - t)^+) \mathbb{E} \tilde{N}^s(t) + \mathbb{E} (\tilde{N}^s(t))^2.$$

Use the inequality $\tilde{N}^s(t) \leq \tilde{N}(t)$ and the inequalities (from (55) and Lemma 4(a))

$$\mathbb{E} \tilde{N}(t) \leq 2e^t; \quad \text{var } \tilde{N}(t) \leq 2e^{2t}$$

to complete the proof of (b). ■

3.4.4. *The Percolation Counting Process Conditioned on Existence of a Path.* Proposition 26 will formalize the idea

Conditional on existence of a path from vertex 1 of specified length, the percolation process is approximately the size-biased Yule process.

In the following lemma, “number of vertices” excludes vertex 1, and we are conditioning on length-to-boundary being s .

Lemma 25. *Fix $\alpha \in L(\mathbf{t})$. Condition on the existence of a path of length $s + b(\alpha)$ from vertex 1 to vertex α . Let \tilde{P}_n^s denote the number of vertices on this path, and $U_{(1)} < U_{(2)} < \dots < U_{(\tilde{P}_n^s)}$ denote the distances of these vertices from vertex 1. Then*

(a) *The exact distribution of \tilde{P}_n^s is*

$$\mathbb{P}(\tilde{P}_n^s = k) = C(s, n) \frac{(n - 1 - \#\mathbf{t})_k s^k}{n^k k!}, \quad 0 \leq k \leq n - 1 - \#\mathbf{t} \tag{56}$$

where $C(s, n) = \left(\sum_{j=0}^{n-1-\#\mathbf{t}} \frac{(n-1-\#\mathbf{t})_j s^j}{n^j j!} \right)^{-1}$ is the normalizing constant.

(b) *Conditional on \tilde{P}_n^s the $(U_{(k)})$ are distributed as the order statistics of \tilde{P}_n^s independent Uniform(0, s) random variables.*

(c) *Suppose $s_n \rightarrow \infty$, $s_n = o(\sqrt{n})$. Then the variation distance between the distribution of $\tilde{P}_n^{s_n}$ and the Poisson(s_n) distribution tends to 0 as $n \rightarrow \infty$.*

Proof. Formula (49) says that the intensity of the event

[[there exists a path of length $s + b(\alpha)$ from 1 to α whose vertices are at distances $0 < u_1 < \dots < u_j < s$]]

is of the form $c(n, s, \alpha) (n - 1 - \#\mathbf{t})_j n^{-j}$. Now (a) follows from the integral identity (50), and (b) follows from the uniformity of the density in (u_1, \dots, u_j) . And (c) follows from (a) because the ratio $\frac{(n-1-\#\mathbf{t})_j}{n^j}$ tends to 1 for $j = o(\sqrt{n})$. ■

For the main result of this section, we study a certain pruned percolation process which we now define carefully. Recall that the percolation counting process $N_n(t)$ counts the number of vertices j such that there exists a path π from 1 to j of length $\leq t$. For the pruned percolation counting process we impose two extra restrictions on π . First, π must not use any vertex in the neighborhood $\mathcal{B}_r(n - 1, n) = \mathbf{t}_{A_n}$. Next, we will be conditioning on existence of a path, say $(1, \eta_1, \eta_2, \dots)$, of specified length from vertex 1. Say π contains a *short-cut* if the path π meets η_j for some $j \geq 1$. The second restriction on π is that π must not contain any short-cuts.

Proposition 26. *Consider a sequence s_n satisfying $\omega_n \leq s_n \leq \log n + B$, and a vertex $\alpha \in L(\mathbf{t})$. Condition on the existence of a path from 1 to α of length $s_n + b(\alpha)$. Let $\tilde{N}_n^{s_n}(t)$ be the pruned percolation counting process defined above. Then for each n there exists a random variable \tilde{W}_n having density we^{-w} on \mathbb{R}^+ such that*

$$\sup_{\omega_n \leq t \leq \log n - \omega_n} |e^{-t} \tilde{N}_n^{s_n}(t) - \tilde{W}_n| \longrightarrow 0 \text{ in probability}$$

as $n \rightarrow \infty$.

We give the proof in some detail; later (Proposition 28) we need the variant for percolation from several sources, and we will omit details of that variant.

3.4.5. *Proof of Proposition 26.* We start with some finite error bounds for the two size-biased Yule processes $\tilde{N}(\cdot)$ and $\tilde{N}^{s_n}(\cdot)$ that were introduced in Section 3.4.3.

Lemma 27. *Consider any sequences $\omega_n, s_n \rightarrow \infty$.*

(a) *Fix $\varepsilon > 0$. Recall the limiting random variable \tilde{W} from Lemma 24. Then there exists a constant C such that*

$$\mathbb{P}\left(\sup_{t \geq \omega_n} \left| e^{-t} \tilde{N}(t) - \tilde{W} \right| > \varepsilon\right) \leq C e^{-\omega_n} \cdot \varepsilon^{-2}.$$

(b) *Fix $\varepsilon > 0$ and consider the processes $\tilde{N}^{s_n}(\cdot)$. Then there exist random variables \tilde{W}_n having density $w e^{-w}$ on \mathbb{R}^+ such that*

$$\mathbb{P}\left(\sup_{t \geq \omega_n} \left| e^{-t} \tilde{N}^{s_n}(t) - \tilde{W}_n \right| > \varepsilon\right) \leq C \varepsilon^{-2} (e^{-2 \cdot s_n} + e^{-\omega_n}).$$

Proof. From (55) we see that $e^{-t} \cdot (\tilde{N}(t) + 1)$ is a martingale. Part (a) follows from Lemma 24 and the L^2 maximal inequality for martingales.

To prove part (b), let $Y(t)$ be a Yule process independent of $\tilde{N}^{s_n}(\cdot)$ and define the process

$$\begin{aligned} Z_n(t) &= \tilde{N}^{s_n}(t) && \text{for } t < s_n \\ &= \tilde{N}^{s_n}(t) + Y(t - s_n) - 1 && \text{for } t \geq s_n. \end{aligned}$$

Note that the process $Z_n(\cdot)$ has the same distribution as the untruncated size biased Yule process $\tilde{N}(\cdot)$. Thus there exists a limiting random variable \tilde{W}_n with density $x \cdot e^{-x}$ such that inequality (a) is satisfied with $Z_n(\cdot)$ in place of $\tilde{N}(\cdot)$. Now note that for any $t > s_n$

$$e^{-t} Z_n(t) = e^{-t} \tilde{N}^{s_n}(t) + e^{-s_n} \cdot [e^{-(t-s_n)} Y(t - s_n)] - e^{-t}.$$

For our desired asymptotics we can ignore the final $-e^{-t}$ term, and write

$$\mathbb{P}\left(\sup_{t \geq \omega_n} \left| \frac{\tilde{N}^{s_n}(t)}{e^t} - \tilde{W}_n \right| > \varepsilon\right) \leq \mathbb{P}\left(\sup_{t \geq \omega_n} \left| \frac{Z_n(t)}{e^t} - \tilde{W}_n \right| > \frac{\varepsilon}{2}\right) + \mathbb{P}\left(e^{-s_n} \cdot \sup_{t \geq 0} \frac{Y(t)}{e^t} > \frac{\varepsilon}{2}\right).$$

Applying the L^2 maximal inequality to the martingale $e^{-t} Y(t)$ gives $\mathbb{P}(e^{-s_n} \cdot \sup_{t \geq 0} \frac{Y(t)}{e^t} > \frac{\varepsilon}{2}) \leq C \varepsilon^{-2} e^{-2 \cdot s_n}$. Combine with part(a) of the Lemma applied to $Z_n(\cdot)$ to get the result. ■

We now give a construction of the pruned percolation counting process $\tilde{N}_n^{s_n}(t)$, designed for comparison with a similar construction of the size-biased Yule process. Recall we are conditioning on existence of a “distinguished” path from 1 to α of length $s_n + b(\alpha)$. Write

$0 = U_0 < U_1 < U_2 < \dots < U_{P_n} < s_n$ for the distances from 1 to the vertices within this path. Define

$$A_n(t) := \#\{i \geq 0 : U_i \leq t\}.$$

We can write

$$\tilde{N}_n^{s_n}(t) = A_n(t) + G_n(t) \tag{57}$$

where $G_n(t)$ is the number of vertices in the pruned percolation counting process which are not on the distinguished path. By definition the process $G_n(\cdot)$ evolves as the counting process satisfying

$$\mathbb{P}(G_n(t + dt) - G_n(t) = 1 | \mathcal{G}_n(t)) = n^{-1} (A_n(t) + G_n(t)) (n - P_n - \#\mathbf{t} - G_n(t)) dt \tag{58}$$

because the number of vertices wetted at t equals $A_n(t) + G_n(t)$ and the number of available vertices to be wetted equals $n - P_n - \#\mathbf{t} - G_n(t)$, the terms P_n and $\#\mathbf{t}$ arising from the two restrictions in the definition of pruned. The filtration used here has $\mathcal{G}_n(0)$ as the σ -field generated by the (U_i) and then $\mathcal{G}_n(t) = \sigma(\mathcal{G}_n(0), G_n(s), 0 \leq s \leq t)$.

To relate this construction to the size-biased Yule process, it does no harm to assume (by variation distance convergence, Lemma 25(c)) that P_n has exactly $\text{Poisson}(s_n)$ distribution, so that (U_1, \dots, U_{P_n}) are the points of a rate-1 Poisson point process on $(0, s_n)$. Now we could construction the size-biased Yule process, cut at s_n , via

$$\tilde{N}^{s_n}(t) = A_n(t) + C_n(t)$$

where $C_n(\cdot)$ evolves as the counting process satisfying

$$\mathbb{P}(C_n(t + dt) - C_n(t) = 1 | \mathcal{G}(t)) = (A_n(t) + C_n(t)) dt \tag{59}$$

for appropriate filtration $(\mathcal{G}(t))$. But it is more useful to couple the two processes by first defining $\tilde{N}_n^{s_n}(\cdot)$ via (57) and then defining $\tilde{N}^{s_n}(\cdot)$ via

$$\tilde{N}^{s_n}(t) = A_n(t) + G_n(t) + B_n(t) \tag{60}$$

where $B_n(\cdot)$ evolves as the counting process with $B_n(0) = 0$ and

$$\mathbb{P}(B_n(t + dt) - B_n(t) = 1 | \mathcal{G}_n(t)) = (b_n(t) + B_n(t)) dt \tag{61}$$

where (subtracting (58) from (59)) $b_n(t) + B_n(t) = (A_n(t) + G_n(t) + B_n(t)) - n^{-1} (A_n(t) + G_n(t)) (n - P_n - \#\mathbf{t} - G_n(t))$, which works out as

$$0 \leq b_n(t) = (A_n(t) + G_n(t)) \frac{P_n + \#\mathbf{t} + G_n(t)}{n}.$$

In particular, $\tilde{N}_n^{s_n}(t) \leq \tilde{N}^{s_n}(t)$, and $B_n(\cdot)$ is the number of extra vertices in the size-biased Yule process but not in the pruned percolation process. In view of Lemma 27(b), to prove Proposition 26 it is sufficient to prove

$$\sup_{\omega_n \leq t \leq \log n - \omega_n} e^{-t} B_n(t) \longrightarrow 0 \text{ in probability.} \tag{62}$$

Note that we can write

$$\begin{aligned} b_n(t) &= n^{-1} \tilde{N}_n^{s_n}(t) (\tilde{N}_n^{s_n}(t) + \#\mathbf{t} + \#\{i : t \leq U_i \leq s_n\}) \\ &\leq n^{-1} \tilde{N}_n^{s_n}(t) (\tilde{N}_n^{s_n}(t) + \#\mathbf{t} + \#\{i : t \leq U_i \leq s_n\}) \end{aligned}$$

and then Lemma 24(c) implies

$$Eb_n(t) \leq n^{-1}(6e^{2t} + 2e^t(\#\mathbf{t} + s_n)) := a_n(t), \text{ say.} \tag{63}$$

Consider the event

$$\Omega_n := \left\{ B_n \left(\frac{1}{3} \log n \right) = 0 \right\}$$

that the two processes coincide up to time $\frac{1}{3} \log n$. Using (61)

$$1 - \mathbb{P}(\Omega_n) \leq \int_0^{\frac{1}{3} \log n} \mathbb{E}b_n(t) dt \leq \int_0^{\frac{1}{3} \log n} a_n(t) dt \rightarrow 0$$

and so $\mathbb{P}(\Omega_n) \rightarrow 1$. Next observe that $e^{-t}B_n(t)$ is a submartingale, because

$$d(e^{-t}B_n(t)) = e^{-t}dB_n(t) - e^{-t}B_n(t)dt$$

and so

$$\mathbb{E}(d(e^{-t}B_n(t)) | \mathcal{G}_n(t)) = e^{-t}b_n(t)dt \geq 0.$$

Appealing to the L^1 maximal inequality for submartingales, to prove (62) it is now enough to prove

$$\mathbb{E}e^{-t_n}B_n(t_n)\mathbb{1}(\Omega_n) \rightarrow 0 \text{ for } t_n = \log n - \omega_n. \tag{64}$$

Write $f_n(t) = \mathbb{E}B_n(t)\mathbb{1}(\Omega_n)$ for $t \geq \frac{1}{3} \log n$, so that $f_n(\frac{1}{3} \log n) = 0$. Using (61), $f'_n(t) \leq f_n(t) + a_n(t)$ and so

$$(e^{-t}f_n(t))' \leq e^{-t}a_n(t).$$

Using (63), for $t \geq \frac{1}{3} \log n$ we have $a_n(t) \leq Ce^{2t}/n$ for some constant C . So $(e^{-t}f_n(t))' \leq Ce^t/n$ and then (64) holds because

$$e^{-t_n}f_n(t_n) = \int_{\frac{1}{3} \log n}^{t_n} Ce^t/n dt \leq Cn^{-1} \exp(t_n) \rightarrow 0.$$

3.5. Proof of Proposition 23

Proposition 26 studied the pruned percolation counting process starting from vertex 1. We now want to consider four such processes running concurrently, starting from vertices 1, 2, 3, 4 (sources). In this setting, if one of the flows reaches a vertex j which was previously reached by a different flow, we say a collision occurs, and vertex j is only counted in the counting process $(\tilde{N}_n^{(i)}(t)$ below) for the source i whose flow first reaches j .

Recall the setting of Proposition 23: we say “conditional on s_1, s_2, t_1, t_2 ” to mean conditional on the event described at the beginning of Section 3.3. Note that we condition only on the lengths and not on the internal structure of the four distinguished path segments. The values of s_1, s_2, t_1, t_2 (which depend on n) are assumed to satisfy

$$\omega_n \leq s^1 := s_1 + s_2 + \sigma \leq \log n + B; \quad \omega_n \leq t^1 := t_1 + t_2 + \sigma \leq \log n + B. \quad (65)$$

For $i = 1, 2, 3, 4$ write $\tilde{N}_n^{(i)}(t)$ for the number of vertices reached by the flow started at source i before time t , in the concurrent flow process. This process differs from the four separate processes in two ways. First, the elimination of collisions, as described earlier. Second, we extend the notion of (forbidden) short-cuts to say that a path of the percolation process may not meet any of the four distinguished paths. Of course for each i , the number $\tilde{N}_n^{(i)}(t)$ is bounded by the corresponding number in the percolation flow from i when the other flows are not present, which was the context of Proposition 26.

Proposition 28. *Conditional on s_1, s_2, t_1, t_2 , there exist random variables $\tilde{W}_n^{(i)}$ such that*

$$\left(\tilde{W}_n^{(1)}, \tilde{W}_n^{(2)}, \tilde{W}_n^{(3)}, \tilde{W}_n^{(4)}\right) \xrightarrow{d} \left(\tilde{W}^{(1)}, \tilde{W}^{(2)}, \tilde{W}^{(3)}, \tilde{W}^{(4)}\right)$$

where the limit r.v.'s are independent with density we^{-w} ; and such that, for any $\omega_n \leq t_n \leq \frac{1}{2}(\log n + B)$,

$$e^{-t_n \tilde{N}_n^{(i)}(t_n)} - \tilde{W}_n^{(i)} \rightarrow 0 \text{ in probability} \quad (66)$$

for each $1 \leq i \leq 4$.

Proof. The proof involves only minor modifications of the proof of Proposition 26—the essential issue is to show that the two changes (collisions; short cuts) in going from separate to concurrent processes has negligible effect. We omit details. ■

Recall the definition of S_{12}, S_{34} and the definition of the Cox point processes from Section 3.3. Let S_{12}^* , (resp. S_{34}^*) be the times of the first collision (within the concurrent flow process) between the flow processes starting from 1 and 2 (resp. from 3 and 4). Note that if a collision occurs between the flow processes started at 1 and 2 at time t , then there is a path of length $2t$ from 1 to 2, and (because we do not allow flow through the neighborhood $\mathcal{B}_t(n-1, n)$) this path does not use the distinguished edge of the neighborhood. So $S_{12} \leq 2S_{12}^*$ and $S_{34} \leq 2S_{34}^*$. This is an inequality because there might be shorter paths that were “pruned away” in the processes we have studied. So to prove Proposition 23 it is enough to prove the following

Proposition 29. *Conditional on s_1, s_2, t_1, t_2 ,*

$$(2S_{12}^* - \log n, 2S_{34}^* - \log n) \xrightarrow{d} (\xi_1^{12}, \xi_1^{34}) \quad (67)$$

as $n \rightarrow \infty$.

Proof. Condition on the concurrent flow process until time t , and suppose the flows from source 1 and source 2 have not collided before time t . Then the instantaneous conditional probability-per-unit-time of a collision (hazard rate) equals

$$\lambda_n(t) = \frac{2\tilde{N}_n^{(1)}(t) \tilde{N}_n^{(2)}(t)}{n}$$

because the unseen length of each possible edge has Exponential(1/n) distribution. We are interested in the recentered process $2S_{12}^* - \log n$. This process has hazard rate

$$\tilde{\lambda}_n(s) = \frac{1}{2}\lambda_n \left(\frac{1}{2}s + \frac{1}{2}\log n \right) = e^s \frac{\tilde{N}_n^{(1)}(\frac{1}{2}\log n + \frac{1}{2}s)}{e^{\frac{1}{2}\log n + \frac{1}{2}s}} \cdot \frac{\tilde{N}_n^{(2)}(\frac{1}{2}\log n + \frac{1}{2}s)}{e^{\frac{1}{2}\log n + \frac{1}{2}s}} \tag{68}$$

Now use Proposition 28 to conclude

$$e^{-s}\tilde{\lambda}_n(s) \xrightarrow{P} \tilde{W}^{(1)}\tilde{W}^{(2)} \text{ uniformly on } 2\omega_n - \log n \leq s \leq B.$$

This easily implies $2S_{12}^* - \log n \xrightarrow{d} \xi_1^{12}$, because ξ_1^{12} is defined to have hazard rate $e^s \tilde{W}^{(1)}\tilde{W}^{(2)}$ on $-\infty < s < \infty$. The joint convergence (67) follows by the same argument, the independence of the limits ($\tilde{W}^{(i)}$) in Proposition 28 implying independence of the limits (ξ_1^{12}, ξ_1^{34}) here. ■

4. FURTHER DISCUSSION

4.1. Analysis of the Limit Function $G(z)$

Recall

$$G(z) = \int_0^\infty P(W_1 W_2 e^{-u} > z) du$$

where W_1 and W_2 are independent Exponential(1). The Mellin transform of G is

$$\begin{aligned} \Phi(y) &= \int_0^\infty z^{y-1} G(z) dz \\ &= \int_0^\infty \int_0^\infty z^{y-1} P\left(W_1 > \frac{ze^u}{W_2}\right) dudz \\ &= \int_0^\infty \int_0^\infty \int_0^\infty \exp(-w) z^{y-1} \exp\left(-\frac{ze^u}{w}\right) dzdudw \\ &= \int_0^\infty \int_0^\infty e^{-w} \Gamma(y) \left(\frac{e^u}{w}\right)^{-y} dwdu \\ &= \Gamma(y) \int_0^\infty e^{-uy} \int_0^\infty e^{-w} w^y dwdu \\ &= \Gamma(y)\Gamma(y+1) \int_0^\infty e^{-uy} du \\ &= \frac{\Gamma(y)\Gamma(y+1)}{y} \\ &= (\Gamma(y))^2. \end{aligned}$$

Checking a table of Mellin transforms ([23] II.5.34) we see

$$G(z) = 2K_0(2z^{1/2})$$




Fig. 4. The function $G(z)$, drawn using Mathematica’s numerical integration toolbox.

where K_0 is the modified Bessel function of the second kind. The standard asymptotics of K_0 ([7] 4.12.6) say

$$K_0(x) \sim \sqrt{\frac{\pi}{2x}} \exp(-x) \quad \text{as } x \rightarrow \infty$$

and so

$$G(z) \sim \pi^{1/2} z^{-1/4} \exp(-2\sqrt{z}) \quad \text{as } z \rightarrow \infty.$$

(See Figure 4).

4.2. Methodology of Relating Local and Global Structure

As illustrated by the heuristic argument in Section 1.3, the conceptual point of Theorem 2 is that a quantity depending on the “global” structure of the network can be studied statistically via a “local” (i.e., large fixed distance) calculation. This reduction to local structure is, to our understanding, the central point in the powerful nonrigorous cavity method of statistical physics [20]. In our attempted mathematical reformulations of the cavity method as applied to combinatorial optimization problems such as TSP [4–6] in this random network model \mathcal{G}_n , we made explicit use of the $n \rightarrow \infty$ limit structure (the PWIT of Section 2.3) of this model as viewed from a random vertex. In these harder problems one needs rather abstract, often as yet not rigorously justified, arguments to connect local and global structure. The problem in this article seems conceptually easier in that we can use concrete calculations instead.

4.3. Flows Through Vertices

In the setting of Theorem 2 one could alternatively consider flows through vertices instead of edges. Let us state this alternative result and indicate the derivation of the limit distribution without giving details of proof.

Write $F_n^*(v)$ for the flow through vertex $v \in [n]$. Let $(W_i, i \geq 1)$ be independent Exponential(1) r.v.'s and let $0 < \xi_1 < \xi_2 < \dots$ be the points of a Poisson (rate 1) process on $(0, \infty)$. Define

$$\Xi := \sum_i \sum_{j \neq i} W_i W_j \exp(-\xi_i - \xi_j).$$

Corollary 30. *In the setting of Theorem 2, as $n \rightarrow \infty$ for fixed $z > 0$,*

$$\frac{1}{n} \#\{v : F_n^*(v) \leq z \log n\} \rightarrow_{L_1} \mathbb{P}(\Xi \leq z).$$

The formula is most neatly derived using the $n \rightarrow \infty$ limit PWIT structure of \mathcal{G}_n [6]. Relative to a typical vertex v of the PWIT, the edge-lengths (ξ_i) to adjacent vertices (v_i) are distributed as the points of a Poisson (rate 1) process on $(0, \infty)$. For each v_i let $N_i(t)$ be the number of vertices within distance t from v_i using paths not via v . Then $e^{-t}N_i(t) \rightarrow W_i$ for i.i.d. Exponential(1) r.v.'s (W_i) . The relative volume of flow through the two edges $v_i \rightarrow v \rightarrow v_j$ will then be $W_i W_j \exp(-\xi_i - \xi_j)$ by the argument for Theorem 2.

4.4. Different Models of Random Networks

The heuristic argument of Section 1.3 can be carried over to a variety of random networks models. For example, fix a degree distribution $P(\Delta = i), i \geq 1$ with finite $2 + \varepsilon$ moment. There are several ways (e.g., the ‘‘configuration model’’) to formalize the idea of a n -vertex graph which is random subject to the constraint that the $n \rightarrow \infty$ asymptotic degree distribution is Δ . Such models have local weak limits which are simple branching processes; looking outward from a typical edge e , each end-vertex is the founder of a Galton-Watson branching process with offspring distribution

$$P(\Delta^* = i) = (i + 1)P(\Delta = i + 1)/E\Delta.$$

Thus we expect the average vertex-vertex distance D_n in such a random graph to behave as

$$\mathbb{E}D_n \sim \frac{\log n}{\log \mathbb{E}\Delta^*}.$$

See [26] for proofs for several models. Now make a random network by assigning independent random lengths η_e to edges e (note that here we do not scale edge-lengths with n). Then the Galton-Watson process above becomes a general Markov branching process in which individuals have Δ^* offspring at independent ages (η_i) ; the population size process $N_\infty(t)$ has some Malthusian growth constant θ and some a.s. limit $\exp(-\theta t)N_\infty(t) \rightarrow Z$. The heuristic argument from Section 1.3 now suggests that the limit joint distribution of edge-lengths and relative edge-flows will be

$$(\eta, Z_1 Z_2 \exp(-\theta \eta))$$

where (η, Z_1, Z_2) are independent. But giving a rigorous proof for the models in [26] may be technically challenging.

Finally, one might consider models on the two-dimensional lattice with i.i.d. random edge-lengths. Here, studying lengths of shortest routes is tantamount to studying (unoriented) first passage percolation [25]. However, if i is close to i' and j is close to j' then we expect the routes from i to j and from i' to j' to coincide except near the endpoints. This suggests a quite different distribution of edge-flows, more specifically that $F(e)$ should have a power-law tail.

4.5. Random Demands

A small variation of our model is to assume that the total flow to be routed from vertex i to vertex j is a random variable D_{ij}/n instead of $1/n$; the flow is still routed along the same shortest path as in the uniform demand case. So the flow across edge e is

$$F_n(e) = \frac{1}{n} \sum_{i \in [n]} \sum_{j \in [n], j \neq i} D_{ij} 1\{e \in \pi(i, j)\}.$$

Because the flow across an edge e is made up from many different source-destination pairs, it is straightforward to add a “law of large numbers” step to the proof of Theorem 2 and obtain the following corollary.

Corollary 31.

(a) *Suppose $D_{ij} \geq 0$ are independent with common mean $0 < \mu < \infty$ and with uniformly bounded second moments. Then*

$$\frac{1}{n} \#\{e : F_n(e) > z\mu \log n\} \rightarrow_{L^1} G(z), \quad z > 0.$$

(b) *Suppose instead the gravitational model $D_{ij} = D_i D_j$ where $D_i \geq 0$ are independent random variables with common mean $0 < \mu < \infty$ and with uniformly bounded second moments. Then*

$$\frac{1}{n} \#\{e : F_n(e) > z\mu^2 \log n\} \rightarrow_{L^1} G(z), \quad z > 0.$$

4.6. Joint Distributions for Shortest Paths

As described in Section 2.1, various aspects of shortest paths in the model \mathcal{G}_n have been studied. The following ideas are developed in [8], inspired by this article. There is a known (implicitly, at least) limit distribution

$$D_n(1, 2) - \log n \xrightarrow{d} D(1, 2)$$

for distance between a typical pair of vertices. Now fix $k \geq 3$. We expect a joint limit

$$(D_n(1, 2) - \log n, \dots, D_n(1, k) - \log n) \xrightarrow{d} (D(1, 2), \dots, D(1, k)) \tag{69}$$

and it turns out the limit distribution is

$$(D(1, 2), \dots, D(1, k)) \stackrel{d}{=} (\xi_1 + \eta_{12}, \dots, \xi_1 + \eta_{1k})$$

where ξ_1 has the double exponential distribution

$$\mathbb{P}(\xi \leq x) = \exp(-e^{-x}), \quad -\infty < x < \infty$$

the η_{ij} have logistic distribution

$$\mathbb{P}(\eta \leq x) = \frac{e^x}{1 + e^x}, \quad -\infty < x < \infty$$

and (here and below) the r.v.'s in the limits are independent. Now we can go one step further: we expect a joint limit for the array

$$(D_n(i, j) - \log n, 1 \leq i < j \leq k) \xrightarrow{d} (D(i, j), 1 \leq i < j \leq k)$$

and the joint distribution of the limit is

$$(D(i, j), 1 \leq i < j \leq k) \stackrel{d}{=} (\xi_i + \xi_j - \xi_{ij}, 1 \leq i < j \leq k)$$

where the limit r.v.'s all have the double exponential distribution. And this is proved in [8], Theorem 3. This implies two representations for the original limit distribution:

$$D(1, 2) \stackrel{d}{=} \xi_1 + \eta_{12} \stackrel{d}{=} \xi_1 + \xi_2 - \xi_{12}.$$

ACKNOWLEDGMENTS

The authors thank a referee for a careful reading and prompting expository improvements.

REFERENCES

- [1] R. Albert and A.-L. Barabási, Statistical mechanics of complex networks, *Rev Mod Phys* 74 (2002), 47–97.
- [2] D. J. Aldous, Asymptotic fringe distributions for general families of random trees, *Ann Appl Probab* 1 (1991), 228–266.
- [3] D. J. Aldous, The $\zeta(2)$ limit in the random assignment problem, *Random Struct Alg* 18 (2001), 381–418.
- [4] D. J. Aldous, Percolation-like scaling exponents for minimal paths and trees in the stochastic mean-field model, *Proc R Soc Lond Ser A Math Phys Eng Sci* 461 (2005), 825–838.
- [5] D. J. Aldous and A. G. Percus, Scaling and universality in continuous length combinatorial optimization, *Proc Natl Acad Sci USA* 100 (2003), 11211–11215.
- [6] D. J. Aldous and J. M. Steele, The objective method: Probabilistic combinatorial optimization and local weak convergence, In H. Kesten, editor, *Probability on discrete structures*, *Encyclopaedia of Mathematical Sciences* 110, Springer-Verlag, Berlin, 2003, pp. 1–72.
- [7] G. E. Andrews, R. Askey, and R. Roy, *Special functions*, *Encyclopedia of Mathematics and its Applications*, Vol. 71, Cambridge University Press, Cambridge, 1999.
- [8] S. Bhamidi, First passage percolation on locally treelike networks. I. Dense random graphs, *J Math Phys* 49 (2008), 125218, 27 pp.
- [9] B. Bollobás, *Random graphs*, Academic Press, London, 1985.

- [10] B. Bollobás, D. Gamarnik, O. Riordan, and B. Sudakov, On the value of a random minimum weight Steiner tree, *Combinatorica* 24 (2004), 187–207.
- [11] R. Durrett, *Random graph dynamics*, Cambridge University Press, Cambridge, 2006.
- [12] A. M. Frieze, On the value of a random minimum spanning tree problem, *Discrete Appl Math* 10 (1985), 47–56.
- [13] S. Janson, One, two and three times $\log n/n$ for paths in a complete graph with random weights, *Combin Probab Comput* 8 (1999), 347–361.
- [14] D. G. Kendall, La propagation d'une épidémie ou d'un bruit dans une population limitée, *Publ Inst Stat Univ Paris* 6 (1957), 307–311.
- [15] S. Linusson and J. Wästlund, A proof of Parisi's conjecture on the random assignment problem, *Probab Th Rel Fields* 128 (2004), 419–440.
- [16] R. Lyons, R. Pemantle, and Y. Peres, Conceptual proof of $L \log L$ criteria for mean behavior of branching processes, *Ann Probab* 23 (1995), 1125–1138.
- [17] J. F. F. Mendes and S. N. Dorogovtsev, *Evolution of networks: From biological nets to the internet and WWW*, Oxford University Press, Oxford, 2003.
- [18] M. Mézard and G. Parisi, A replica analysis of the travelling salesman problem, *J Physique* 47 (1986), 1285–1296.
- [19] M. Mézard and G. Parisi, On the solution of the random link matching problem, *J. Physique* 48 (1987), 1451–1459.
- [20] M. Mézard and G. Parisi, The cavity method at zero temperature, *J Stat Phys* 111 (2003), 1–34.
- [21] C. Nair, B. Prabhakar, and M. Sharma, Proofs of the Parisi and Coppersmith-Sorkin random assignment conjectures, *Random Struct Alg* 27 (2005), 413–444.
- [22] M. E. J. Newman, The structure and function of complex networks, *SIAM Rev* 45 (2003), 167–256.
- [23] F. Oberhettinger, *Tables of Mellin transforms*, Springer-Verlag, New York, 1974.
- [24] S. M. Ross, *Stochastic processes*, 2nd edition, Wiley, New York, 1996.
- [25] R. T. Smythe and J. C. Wierman, First-passage percolation on the square lattice, *Lecture Notes in Mathematics*, Vol. 671, Springer, Berlin, 1978.
- [26] H. van den Esker, R. van der Hofstad, and G. Hooghiemstra, Universality for the distance in finite variance random graphs: Extended version, arXiv:math.PR/0605414, Nancy, France.
- [27] R. van der Hofstad, G. Hooghiemstra, and P. van Mieghem, The flooding time in random graphs, *Extremes* 5 (2002), 111–129.
- [28] J. Wästlund, Random assignment and shortest path problems, In *Proceedings of the fourth colloquium on mathematics and computer science: Algorithms, trees, combinatorics and probabilities*, Nancy, France, 2006.
- [29] J. Wästlund, The travelling salesman problem in the stochastic mean field model, Available at <http://www.math.chalmers.se/~wastlund/>, 2006. Accessed 1 February 2009.