UC Berkeley

UC Berkeley Previously Published Works

Title

On the Critical Value for 'Percolation' of Minimum-Weight Trees in the Mean-Field Distance Model

Permalink

<https://escholarship.org/uc/item/0qc791x2>

Journal

Combinatorics Probability and Computing, 7(1)

ISSN

0963-5483

Author

Aldous, D

Publication Date

1998

DOI

10.1017/S0963548397003155

Peer reviewed

On the Critical Value for “Percolation” of Minimum-Weight Trees in the Mean-Field Distance Model

David Aldous*
Department of Statistics
University of California
Berkeley CA 94720
aldous@stat.berkeley.edu
<http://www.stat.berkeley.edu/users/aldous>

October 8, 1996

Abstract

Consider the complete n -graph with independent exponential (mean n) edge-weights. Let $M(c, n)$ be the maximal size of subtree for which the average edge-weight is at most c . It is shown that $M(c, n)$ transitions from $o(n)$ to $\Omega(n)$ around some critical value $c(0)$, which can be specified in terms of a fixed point of a mapping on probability distributions.

*Research supported by N.S.F. Grant DMS96-22859

1 Introduction

To each edge e of the complete graph on $\{1, 2, \dots, n\}$ attach a weight w_e , where the (w_e) are independent exponential (mean n) r.v.'s. Call this randomly-weighted graph $\mathcal{W}_n$. For each subtree of $\mathcal{W}_n$, that is each tree t whose vertex-set is a subset of $\{1, 2, \dots, n\}$, write $|t|$ for the number of edges of t and $w(t) = \sum_{e \in t} w_e$ for the weight of the tree. So $w(t)/|t|$ is the average edge-weight of the tree t . Consider the maximum size of tree with average edge-weight at most c :

$$M(c, n) = \max\{|t| : w(t)/|t| \leq c\}. \quad (1)$$

It is natural to guess that there exists some critical value $c(0)$ such that (for large n) the random process $c \rightarrow M(c, n)$ makes the transition from $o(n)$ to $\Omega(n)$ as c increases through some neighborhood of $c(0)$. Theorem 1 verifies this guess.

Theorem 1 *There exists $c(0) \in [e^{-2}, e^{-1}]$ such that*

$$(c < c(0)) : \lim_n P(M(c, n) > \varepsilon n) = 0 \quad \forall \varepsilon > 0 \quad (2)$$

$$(c > c(0)) : \exists \varepsilon(c) > 0 \text{ such that } \lim_n P(M(c, n) > \varepsilon(c)n) = 1. \quad (3)$$

This is “natural” by analogy with the case where we replace the average edge-weight $w(t)/|t|$ by the maximum edge weight $\max_{e \in t} w_e$ and consider

$$\tilde{M}(c, n) = \max\{|t| : \max_{e \in t} w_e \leq c\}.$$

Then the analog of Theorem 1 holds with critical value 1. This is essentially the classical result that in the random graph process $G(n, P(\text{edge}) = \lambda/n)$ the time λ of first appearance of an $\Omega(n)$ -size connected component is asymptotic to 1. See [11] for a survey of variations of that classical theory.

The asymptotic bounds $\{e^{-2}, e^{-1}\}$ stated in Theorem 1 are rather easy: the lower bound emerges from a counting argument (section 3.4) and the upper bound is the exact asymptotic critical value when we restrict trees to be *paths* (section 3.5). The existence of the limit $c(0)$ is less easy, and we take an indirect approach, as follows. Write $\mathcal{M}$ for the set of probability distributions on $[0, \infty]$. Write δ_x for the probability measure concentrated at x . Write $0 < \xi_1 < \xi_2 < \dots$ for the times of a Poisson (rate 1) process of events. For fixed $c > 0$ define a map $\Gamma_c : \mathcal{M} \rightarrow \mathcal{M}$ as follows. Given $\mu \in \mathcal{M}$ let $(Y_i, i \geq 1)$ be independent with distribution μ , independent of (ξ_i) , and define $\Gamma_c \mu$ to be the distribution of $\sum_{i=1}^{\infty} (c - \xi_i + Y_i)^+$. Here $x^+ = \max(x, 0)$. Note that the distribution δ_∞ is trivially invariant under Γ_c .

Lemma 2 *There exists a critical value $0 \leq c(0) \leq 1$ with the following properties.*

(a) *For $c < c(0)$ the map Γ_c has an invariant measure μ_c such that $\mu_c[0, \infty) = 1$ and $\Gamma_c^k \delta_0 \rightarrow \mu_c$ as $k \rightarrow \infty$.*

(b) *For $c > c(0)$ the map Γ_c has no invariant measure except δ_∞ , and $\Gamma_c^k \delta_0 \rightarrow \delta_\infty$ as $k \rightarrow \infty$.*

In proving Theorem 1 we use Lemma 2 to provide the definition of $c(0)$. The key idea is the local approximation of the randomly-weighted complete graph $\mathcal{W}_n$ by a randomly-weighted infinite tree $\mathcal{T}^{(\infty)}$, described in section 2. This local approximation provides a systematic “probabilistic” approach to problems concerning $\mathcal{W}_n$: see [2] for this approach to Frieze’s result (4) and [3] for the random assignment problem. The proofs of Lemma 2 and Theorem 1 occupy section 3.

The feature that the solution is expressed in terms of a fixed point of a mapping on distributions occurs quite often in problems involving probability on trees or probabilistic analysis of recursive algorithms: see [4, 13] for algorithmic instances, [9] for queueing applications, and [12] for random walks on Galton-Watson trees.

Our setting may remind the reader of a well-known result of Frieze [7]:

$$EW_n \rightarrow \zeta(3) = \sum_{i=1}^{\infty} i^{-3} \text{ as } n \rightarrow \infty \quad (4)$$

where W_n is the average edge-weight in the minimum-weight spanning tree (MST). The tractability of (4) comes from the existence of the greedy algorithm for constructing the MST, which implies a simple criterion for whether a specified edge is in the MST. Theorem 1 can be rephrased in terms of minimum-weight trees spanning some subset of εn vertices, but the algorithmic problem of finding such trees is apparently hard¹. In this regard, attempting to estimate $c(0)$ by Monte Carlo simulation from its Theorem 1 interpretation would be difficult, but it is straightforward to use Lemma 2 as a basis for simulations, which indicate $c(0) \approx 0.263$.

Deeper study of the critical behavior of our minimum-average-weight subtrees (in the spirit of the deep known results [8] for critical behavior of large components in the usual random graph process) presents interesting challenges. In the usual random graph process, one can define an “incipient infinite component” $\mathcal{C}$ (to borrow terminology from percolation theory),

¹An expert said “it must be NP-complete” but could not provide a citation

which turns out to be the critical Poisson Galton-Watson branching process conditioned to be infinite [5]. In the setting of this paper one can presumably define an infinite random tree $\mathcal{C}^*$ as a weak limit of maximal subtrees of average edge-weight c_n , for some $c_n \downarrow c(0)$. Does $\mathcal{C}^*$ qualitatively resemble $\mathcal{C}$? More precisely, $\mathcal{C}$ has the *scaling property* that the number of vertices within distance d of a reference vertex grows as order d^2 ; does $\mathcal{C}^*$ have the same scaling property? This question is motivated by the *universality* paradigm from statistical physics, which asserts that scaling behavior at critical points should be independent of the details of a model.

2 The local approximation

Consider the infinite tree whose vertices are strings $\mathbf{b} = b_1 b_2 \dots b_i$, where $i \geq 0$ and $b_k \geq 1$, rooted at the empty string ϕ , with edges $(\mathbf{b}, \mathbf{b}j)$, where for $\mathbf{b} = b_1 b_2 \dots b_i$ the vertex $\mathbf{b}j = b_1 b_2 \dots b_i j$ is a *child* of $\mathbf{b}$. For each $\mathbf{b}$ attach weights $w_{\mathbf{b}, \mathbf{b}j}$ to the edges $(\mathbf{b}, \mathbf{b}j)$, where the weights are distributed as the times (ξ_j) of a Poisson (rate 1) process of events, independently as $\mathbf{b}$ varies. Call this weighted tree $\mathcal{T}^{(\infty)}$, and let $\mathcal{T}_m^{(\infty)}$ be the subtree consisting of the first m generations, i.e. strings $\mathbf{b} = b_1 b_2 \dots b_i$ with $i \leq m$, and let $\mathcal{T}_{m,L}^{(\infty)}$ be the further subtree where only the first L children of an individual are allowed, i.e. where each $b_i \leq L$. So the vertex-set of $\mathcal{T}_{m,L}^{(\infty)}$ is $N_L^{(m)} = \cup_{i=0}^m \{1, 2, \dots, L\}^i$.

Fix (m, L) and consider the weighted complete graph $\mathcal{W}_n$. For large enough n there is a natural way to define a random map $\iota : N_L^{(m)} \rightarrow \{1, 2, \dots, n\}$. Set $\iota(\phi) = 1$ and inductively, passing through $N_L^{(m)}$ in breadth-first order, let $\iota(\mathbf{b}j)$ be the vertex v for which $w_{\iota(\mathbf{b}), v}$ is minimized, over vertices v which are not the ι -value of any previous $\mathbf{b}'$. This construction yields a weighted tree, say $\mathcal{T}_{m,L}^{(n)}$, with edge-weights $w_{\iota(\mathbf{b}), \iota(\mathbf{b}j)}$, on the same vertex-set $N_L^{(m)}$ as $\mathcal{T}_{m,L}^{(\infty)}$. It is elementary that the order statistics $\eta_1^{(n)} < \eta_2^{(n)} < \dots < \eta_n^{(n)}$ of n independent exponential (mean n) r.v.'s satisfy

$$(\eta_1^{(n)}, \dots, \eta_L^{(n)}) \xrightarrow{d} (\xi_1, \dots, \xi_L) \text{ as } n \rightarrow \infty, L \text{ fixed}.$$

It follows that

$$\mathcal{T}_{m,L}^{(n)} \xrightarrow{d} \mathcal{T}_{m,L}^{(\infty)} \text{ as } n \rightarrow \infty, m, L \text{ fixed} \quad (5)$$

in the sense that the edge-weights converge in distribution. This *local approximation* is key to our methodology.

Recall the definition of Γ_c : $\Gamma_c \mu$ is the distribution of $\sum_{i=1}^{\infty} (c - \xi_i + Y_i)^+$. Write *dist* for “distribution” and write Γ_c^m for the m -fold iteration of Γ_c . The significance of the definition is that

$$\Gamma_c^m \delta_0 = \text{dist } \max(c|t| - w(t) : t \subset \mathcal{T}_m^{(\infty)}) \quad (6)$$

where t denotes a (possibly trivial) subtree containing the root ϕ . Identity (6) is established by induction on m : we get the maximum value over height- m subtrees by considering for each child $j = 1, 2, \dots$ of ϕ the maximal value Y_j over height- $(m-1)$ subtrees rooted at j , and including the branch through j if the contribution $c - \xi_j + Y_j$ made by that branch is positive. Similarly define $\Gamma_{c,L}$ by: $\Gamma_{c,L} \mu$ is the distribution of $\sum_{i=1}^L (c - \xi_i + Y_i)^+$. Then

$$\Gamma_{c,L}^m \delta_0 = \text{dist } \max(c|t| - w(t) : t \subset \mathcal{T}_{m,L}^{(\infty)}). \quad (7)$$

Remark. The right side of (6) clearly resembles a Lindley equation for workload in a tree-structured queueing process (cf. [9]), though we cannot give any precise queueing interpretation of our setup.

We record a simple lemma.

Lemma 3 *Let $q(k, a, L, n)$ be the probability that there exists a path $1 = v_0, v_1, \dots, v_j, j \leq k$ in $\mathcal{W}_n$ such that*

- (a) $\max_{i < j} w_{v_{i-1}, v_i} \leq a$
- (b) $|\{v : w_{v_{j-1}, v} < w_{v_{j-1}, v_j}\}| \geq L$.

Then for fixed k, a

$$\lim_{L \rightarrow \infty} \limsup_{n \rightarrow \infty} q(k, a, L, n) = 0.$$

Proof. The mean number of paths $1 = v_0, v_1, \dots, v_{j-1}, j \leq k$ satisfying (a) tends to $\sum_{i=0}^{k-1} a^i$ as $n \rightarrow \infty$. For each such path, the chance there exists v_j satisfying (a,b) tends to $P(\xi_{L+1} \leq a)$. So $\limsup_n q(k, a, L, n) \leq (\sum_{i=0}^{k-1} a^i) P(\xi_{L+1} \leq a)$. And clearly $\lim_{L \rightarrow \infty} P(\xi_{L+1} \leq a) = 0$.

3 Proof of Theorem 1

3.1 Proof of Lemma 2

The proof uses only simple monotonicity arguments. Let $\preceq$ be the usual “stochastically less than” partial order on $\mathcal{M}$:

$$\mu_1 \preceq \mu_2 \text{ iff } \mu_1[0, x] \geq \mu_2[0, x], 0 \leq x \leq \infty$$

or equivalently

$\mu_1 \preceq \mu_2$ iff there exist X_1, X_2 such that $\text{dist}(X_i) = \mu_i$ and $P(X_1 \leq X_2) = 1$.

It is easy to check from the definitions that Γ_c is an increasing map:

$$\mu_1 \preceq \mu_2 \text{ implies } \Gamma_c \mu_1 \preceq \Gamma_c \mu_2 \quad (8)$$

and that Γ_c increases with c :

$$c_1 \leq c_2 \text{ implies } \Gamma_{c_1} \mu \preceq \Gamma_{c_2} \mu. \quad (9)$$

By (8) and induction, $\Gamma_c^k \delta_0 \preceq \Gamma_c^{k+1} \delta_0$. It follows that there exists an increasing limit

$$\Gamma_c^k \delta_0 \uparrow \mu_c \text{ as } k \rightarrow \infty \quad (10)$$

where convergence is weak convergence on the compactified half-line $[0, \infty]$. It is easy to check that Γ_c is continuous w.r.t. increasing limits, and hence μ_c is invariant for Γ_c . But an invariant measure μ clearly has the property that $\mu(\infty) = 0$ or 1 . And by (9) μ_c is stochastically increasing with c . So if we define

$$c(0) = \inf\{c : \mu_c(\infty) = 1\}$$

then we have proved the assertions of the lemma, except for proving $c(0) \leq 1$.

Fix $c > 1$. Then $\Gamma_c \mu \succeq \text{dist}(Z + Y)$, where $\text{dist}(Y) = \mu$, $Z = c - \xi_1$ and Y, Z are independent. So inductively $\Gamma_c^k \delta_0 \succeq \text{dist}(Z_1 + \dots + Z_k)$ where the (Z_i) are independent copies of Z . But $EZ = c - 1 > 0$, so by the law of large numbers $\Gamma_c^k \delta_0 \rightarrow \delta_\infty$. So $c(0) \leq c$, establishing the result $c(0) \leq 1$.

3.2 The subcritical regime

The next lemma is the main ingredient for proving the subcritical behavior (2) in Theorem 1.

Lemma 4 *Let $c < c(0)$. For $x > 0, k \geq 1$ let $N(c, k, x, n)$ be the number of vertices v of $\mathcal{W}_n$ such that some tree t containing v has $|t| \leq 3k$ and $w(t) \leq c|t| - x$. Then*

$$\limsup_{n \rightarrow \infty} P(N(c, k, x, n) > \varepsilon n) \leq \varepsilon^{-1} \mu_c[x, \infty), \quad \varepsilon > 0.$$

Proof. Recall from section 2 the construction of $\mathcal{T}_{m,L}^{(n)}$. From (5,7), for fixed L

$$\max\{c|t| - w(t) : t \subseteq \mathcal{T}_{3k,L}^{(n)}\} \xrightarrow{d} \Gamma_{c,L}^{3k} \delta_0 \text{ as } n \rightarrow \infty$$

where “ $t \subseteq$ ” denotes a subtree containing the root ϕ . Restricting to subtrees of size at most $3k$ can only make the left side smaller, so

$$\max\{c|t| - w(t) : t \subseteq \mathcal{T}_{3k,L}^{(n)}, |t| \leq 3k\} \preceq_{n \rightarrow \infty} \Gamma_{c,L}^{3k} \delta_0$$

where $\preceq_{n \rightarrow \infty}$ means “asymptotically stochastically less than”. We now want to remove the restriction to “first L children” and say

$$\max\{c|t| - w(t) : t \subseteq \mathcal{W}_n, |t| \leq 3k, t \text{ contains vertex } 1\} \preceq_{n \rightarrow \infty} \Gamma_c^{3k} \delta_0. \quad (11)$$

Since $\Gamma_{c,L}^{3k} \delta_0 \preceq \Gamma_c^{3k} \delta_0$, the only way (11) could fail is if (with probability not $\rightarrow 0$) the maximum were attained by trees t_n containing a parent-child pair (v, v') for which $w_{v,v'}$ is the l_n 'th smallest edge-weight incident at v , for some $l_n \rightarrow \infty$. But this possibility is precluded by Lemma 3, which implies that such trees have $w(t_n) \rightarrow \infty$. So (11) is true. Combining (11) with the fact (10) that $\Gamma_c^{3k} \delta_0 \preceq \mu_c$, we see that the chance that vertex 1 has the property specified in the Lemma is asymptotically at most $\mu_c[x, \infty)$. So the result follows using Markov's inequality. $\square$

We also quote an elementary fact.

Lemma 5 *Fix $k \geq 1$. Any tree with at least k edges may be decomposed as a union of edge-disjoint subtrees, each subtree having between k and $3k$ (inclusive) edges.*

Proof of (2). Fix $k \geq 1$ and fix $c_1 < c_2 < c_3 < c(0)$. Suppose there exists a tree t^* with at least εn edges and with average edge-weight at most c_1 . Apply Lemma 5 to decompose into subtrees. Some subtrees may have average edge-weight $\geq c_2$, but the proportion of edges of t^* in such subtrees is at most c_1/c_2 , and so a proportion at least $1 - \frac{c_1}{c_2}$ of edges of t^* lie in subtrees with average edge-weight at most c_2 . Now any subset of e edges of a tree are incident to at least $e + 1$ distinct vertices. So at least $(1 - \frac{c_1}{c_2})\varepsilon n$ vertices lie in trees of size between k and $3k$ with average edge-weight at most c_2 . Defining x by

$$c_2 k = c_3 k - x,$$

so that $c_2 |t| \leq c_3 |t| - x$ whenever $k \leq |t| \leq 3k$, we have shown

$$P(M(c_1, n) \geq \varepsilon n) \leq P(N(c_3, k, x, n) \geq (1 - \frac{c_1}{c_2})\varepsilon n).$$

Applying Lemma 4,

$$\limsup_n P(M(c_1, n) \geq \varepsilon n) \leq (1 - \frac{c_1}{c_2})^{-1} \varepsilon^{-1} \mu_{c_3}[k(c_3 - c_2), \infty).$$

Since k is arbitrary, (2) follows.

3.3 The supercritical regime

The next lemma is the main ingredient for proving the supercritical behavior (3) in Theorem 1.

Lemma 6 *Fix $c > c(0)$ and an integer m . Then there exists an integer q , depending only on c and m , and an algorithm on $\mathcal{W}_n$ (for sufficiently large n) with the following properties. Given a uniform random initial vertex V_1 , the algorithm finds vertices $V_2, \dots, V_q$ by looking only at weights on edges for which one end-vertex is some $(V_i, i < q)$. With probability at least $1/2$, the algorithm “succeeds” in finding a tree whose vertices are a subset of $\{V_1, \dots, V_q\}$ containing both V_1 and V_q , such that the tree has at least m edges and has average edge-weight at most c .*

Proof. Since $\Gamma_c^k \delta_0 \rightarrow \delta_\infty$ as $k \rightarrow \infty$, we can choose k such that $\Gamma_c^k \delta_0(cm, \infty) > 1/2$, and then choose L such that $\Gamma_{c,L}^k \delta_0(cm, \infty) > 1/2$. From (7), this says that with probability $> 1/2$ there exists a subtree $t \subset \mathcal{T}_{k,L}^{(\infty)}$ such that $c|t| - w(t) > cm$. So by the local approximation (5), for sufficiently large n there exists, with probability $> 1/2$, a subtree $t \subset \mathcal{T}_{k,L}^{(n)}$ such that $c|t| - w(t) > cm$. Such a subtree certainly has at least m edges and has average edge-weight at most c . This existence result may be rephrased as an algorithm; at each step of the construction of $\mathcal{T}_{k,L}^{(n)}$ we check to see whether a subtree with the desired properties exists within the tree already constructed, and stop if it does. The last vertex examined has the property that the algorithm has not seen its edge-weights, except for those edges to previous vertices. At most $q = \sum_{i=0}^k L^i$ vertices are used, and by incorporating arbitrary extra vertices we may assume exactly q vertices are used. $\square$

Proof of (3). Take c, m, q as in Lemma 6. We shall describe how to construct a size $\Omega(n)$ tree in stages. In summary: at each stage we use Lemma 6 to examine q vertices (in the reduced graph from which vertices examined in previous stages were removed) and, if a suitable subtree is found, it is attached to the already-grown tree by a linking edge.

Here is the construction in detail. Take V_1 to be a uniform random vertex, apply Lemma 6, and suppose the algorithm succeeds. Let W_1 be the vertex $\notin \{V_1, \dots, V_q\}$ for which w_{V_q, W_1} is minimized. Remove vertices $\{V_1, \dots, V_q\}$ and incident edges from the graph. Then the reduced graph on $n - q$ vertices has edge-weights which are independent exponential (mean n) r.v.'s, because the algorithm choosing $\{V_1, \dots, V_q\}$ did not examine any of its edges, and W_1 is a uniform random vertex of the reduced graph. We want to apply Lemma 6, with n replaced by $n - q$, to the reduced graph. Because the edge-weights have mean n instead of $n - q$, we need to modify the criterion of “success” to say that the average edge-weight is at most $\frac{n}{n-q}c$. Apply the lemma to the reduced graph, to produce $\{W_1, \dots, W_q\}$. Suppose the algorithm succeeds. We then include the edge (V_q, W_1) as the “linking edge”, set X_1 to be the vertex $\notin \{W_1, \dots, W_q\}$ for which w_{W_q, X_1} is minimized, and remove vertices $\{W_1, \dots, W_q\}$ and incident edges from the graph. If alternatively the algorithm fails, remove vertices $\{W_1, \dots, W_q\}$ and incident edges from the graph, and let X_1 be the vertex in the remaining graph for which w_{V_q, X_1} is minimized.

In general, at the end of stage s we have constructed a subtree with a distinguished vertex V_* , and we have a reduced graph on $n - sq$ vertices whose internal edge-weights, and the weights on edges to V_* , are independent exponential (mean n). Stage $s + 1$ starts with the vertex Z_1 in the reduced graph for which w_{V_*, Z_1} is minimized, uses Lemma 6 to examine some $\{Z_1, Z_2, \dots, Z_q\}$ seeking a subtree with average weight at most $\frac{n}{n-sq}c$; if successful, the subtree is linked to the existing tree via the edge (V_*, Z_1) and vertex Z_q becomes the new distinguished vertex. Whether successful or not, vertices $\{Z_1, Z_2, \dots, Z_q\}$ are removed from the reduced graph.

Continue for $\delta n/q$ stages, and consider the resulting tree $\mathbf{T}$. With probability $\rightarrow 1$ as $n \rightarrow \infty$ (w.h.p.) at least $\delta n/(3q)$ stages were successes, and in this case $\mathbf{T}$ has at least $\frac{\delta n}{3q}(m+1) - 1$ edges, and its average edge-weight is at most

$$\frac{c}{1-\delta} + \frac{\gamma}{m+1}$$

where γ is the average weight of the “linking” edges. For large n , each linking edge-weight is approximately distributed as the time of first success in a Poisson (rate 1) process of events, where each event has chance $\geq 1/2$ to be a success. So the mean linking edge-weight is asymptotically at most 2, and so w.h.p. $\gamma \leq 3$. So w.h.p. $\mathbf{T}$ has at least $\frac{\delta m}{3q}n$ edges and average edge-weight at most $\frac{c}{1-\delta} + \frac{3}{m+1}$. By choosing δ small and m large, we establish (3).

3.4 The lower bound on $c(0)$

This is just the natural counting argument. Fix $c > 0$ and consider subtrees t of $\mathcal{W}_n$.

$$E|\{t : |t| = k, w(t) \leq ck\}| = \binom{n}{k+1} (k+1)^{k-1} P(\text{Poisson}(\frac{ck}{n}) \geq k)$$

using Cayley's formula and the representation of sums of exponential r.v.'s in terms of the Poisson distribution. Consider $n \rightarrow \infty$ and $k \rightarrow \infty$. The Poisson probability is asymptotic to the probability " $= k$ ", which is $e^{-ck/n} (ck/n)^k / k!$, and so

$$\begin{aligned} & k^{-1} \log E|\{t : |t| = k, w(t) \leq ck\}| \\ & \leq k^{-1} \log \binom{n}{k+1} + \log k + \log(ck/n) - k^{-1} \log k! + o(1). \end{aligned}$$

Since $\binom{n}{k+1} \leq n^{k+1} / (k+1)!$, the right side is at most

$$k^{-1} \log n + \log(ck^2) - \frac{2}{k} \log k! + o(1)$$

and using Stirling's formula this is

$$k^{-1} \log n + \log(ce^2) + o(1).$$

One can check that the $o(1)$ term is uniform over $\varepsilon n \leq k \leq n$, for fixed ε . So if $c < \rho e^{-2}$ for some $\rho < 1$, then ultimately

$$E|\{t : |t| = k, w(t) \leq ck\}| \leq \rho^k, \varepsilon n \leq k \leq n.$$

So

$$P(M(c, n) > \varepsilon n) \leq \sum_{\varepsilon n \leq k \leq n} E|\{t : |t| = k, w(t) \leq ck\}| \leq n\rho^{\varepsilon n} \rightarrow 0.$$

By (3) $c(0) \geq c$, implying $c(0) \geq e^{-2}$.

3.5 The critical value for paths

Specialize (1) to paths, i.e. define

$$M_*(c, n) = \max\{|t| : t \text{ is a path in } \mathcal{W}_n, w(t)/|t| \leq c\}.$$

The next result asserts that the critical value is now e^{-1} , implying the upper bound $c(0) \leq e^{-1}$.

Proposition 7

$$(c < e^{-1}) : \lim_n P(M_*(c, n) > \varepsilon n) = 0 \quad \forall \varepsilon > 0 \quad (12)$$

$$(c > e^{-1}) : \exists \varepsilon(c) > 0 \text{ such that } \lim_n P(M_*(c, n) > \varepsilon(c)n) = 1. \quad (13)$$

Proof. The subcritical result (12) is just the counting argument from section 3.4, counting paths instead of trees: we omit details. For the supercritical result (13), the key idea is that the weighted tree $\mathcal{T}^{(\infty)}$ in section 2 may be identified with the standard Yule process (continuous-time rate-1 branching process), by regarding the weight on edge $(\mathbf{b}, \mathbf{b}j)$ as the time between the birth of individual $\mathbf{b}$ and the birth of individual $\mathbf{b}j$. Now a standard fact about the Yule process is that the maximal generation G_s of the individuals born before time s satisfies

$$s^{-1}G_s \rightarrow e \text{ a.s. as } s \rightarrow \infty. \quad (14)$$

This is a simple special case of general results (see e.g. [10, 6]) about continuous-time branching random walk (regarding a generation- g individual as positioned at g); see [1] for a direct proof of the special case. Rephrasing (14) in terms of subtrees $t \subseteq \mathcal{T}^{(\infty)}$,

$$\begin{aligned} & \min\{w(t)/|t| : t \subseteq \mathcal{T}^{(\infty)} \text{ is a path of length } m \text{ from the root}\} \\ & \rightarrow e^{-1} \text{ a.s. as } m \rightarrow \infty. \end{aligned} \quad (15)$$

The supercritical result (13) can be deduced from (15) by constructing a $\Omega(n)$ -length path by linking path segments of fixed large size m , analogous to the construction of the $\Omega(n)$ -size tree in section 3.3. Again, we omit details.

4 Final remarks

(a) It is easy to see that at the critical value $c(0)$ the invariant distribution $\mu_{c(0)}$ of the map $\Gamma_{c(0)}$ is still supported on $[0, \infty)$. So we have the distributional identity

$$Y \stackrel{d}{=} \sum_{i \geq 1} (c(0) - \xi_i + Y_i)^+ \quad (16)$$

where $\text{dist } Y_i = \text{dist } Y = \mu_{c(0)}$. Now we can use (16) to define a branching Markov chain on $(0, \infty)$, in which an individual at position y has offspring

at positions (Y_i) specified as follows: condition the sum in (16) to equal y , and take the (Y_i) for which the summands are positive. It seems heuristically clear that the “incipient minimal-average-weight infinite tree” $\mathcal{C}^*$ mentioned at the end of the Introduction is just the family tree of this branching Markov chain, conditioned on non-extinction. But in the absence of explicit information about the distribution $\mu_{c(0)}$ it seems hard to study $\mathcal{C}^*$.

(b) In place of the complete graph, one could base a “mean-field” model on the infinite regular tree of degree r , putting independent exponential (mean r) weights on edges, and study the minimal value $c(r)$ of average edge-weight of infinite subtrees. Without checking details, we believe that the analogs of Theorem 1 and Lemma 2, with the sum defining Γ_c replaced by $\sum_{i=1}^{r-1} (c - \eta_i + Y_i)^+$, where $(\eta_1, \dots, \eta_{r-1})$ are independent exponential (mean r), are true (and much simpler to prove) in this alternate setting.

(c) The relation between our “percolation” result (Theorem 1) for trees and Frieze’s result (4) for spanning trees is analogous to the relation between the “percolation” result for *paths* (Proposition 7) and the TSP result for $\mathcal{W}_n$, i.e. the fact that the average edge-weight in the minimum-weight tour of all n vertices is asymptotic to a constant c_{TSP} . Ironically, the relative difficulty of the “percolation” and “all n ” results is reversed: the explicit value of c_{TSP} is not known.

References

- [1] David Aldous and Jim Pitman. The asymptotic speed and shape of a particle system. In J.F.C. Kingman and G.E.H. Reuter, editors, *Probability Statistics and Analysis*, pages 1–23. Cambridge University Press, 1983. London Math. Soc. Lecture Notes 79.
- [2] D.J. Aldous. A random tree model associated with random graphs. *Random Structures and Algorithms*, 1:383–402, 1990.
- [3] D.J. Aldous. Asymptotics in the random assignment problem. *Probab. Th. Rel. Fields*, 93:507–534, 1992.
- [4] D.J. Aldous. Greedy search on the binary tree with random edge-weights. *Combinatorics, Probability and Computing*, 1:281–293, 1992.
- [5] D.J. Aldous and J. Pitman. On the ascension of the giant component. In Preparation, 1996.
- [6] J.D. Biggins. The growth and spread of the general branching random walk. *Ann. Appl. Probab.*, 5:1008–1024, 1995.
- [7] A.M. Frieze. On the value of a random minimum spanning tree problem. *Discrete Appl. Math.*, 10:47–56, 1985.
- [8] S. Janson, D. E. Knuth, T. Łuczak, and B. Pittel. The birth of the giant component. *Random Structures and Algorithms*, 4:233–358, 1993.
- [9] F.I. Karpelevich, M. Ya. Kelbert, and Yu. M. Suhov. Higher-order Lindley equations. *Stochastic Process. Appl.*, 53:65–96, 1994.
- [10] J.F.C. Kingman. The first birth problem for an age-dependent branching process. *Ann. Probab.*, 3:790–801, 1975.
- [11] T. Łuczak. Phase transition phenomena in random discrete structures. *Discrete Math.*, 136:225–242, 1994.
- [12] R. Lyons, R. Pemantle, and Y. Peres. Ergodic theory on Galton-Watson trees: Speed of random walk and dimension of harmonic measure. *Ergodic Theory Dynamical Systems*, 15:593–619, 1995.
- [13] S.T. Rachev and L. Ruschendorf. Probability metrics and recursive algorithms. *Adv. in Appl. Probab.*, 27:770–799, 1995.