

The longest minimum-weight path in a complete graph

Louigi Addario-Berry Nicolas Broutin Gábor Lugosi *

February 15, 2009

Abstract

We consider the minimum-weight path between any pair of nodes of the n -vertex complete graph in which the weights of the edges are i.i.d. exponentially distributed random variables. We show that the longest of these minimum-weight paths has about $\alpha^* \log n$ edges where $\alpha^* \approx 3.5911$ is the unique solution of the equation $\alpha \log \alpha - \alpha = 1$. This answers a question posed by Janson [8].

1 Introduction

We consider the complete graph K_n on n vertices $\{1, \dots, n\} = [n]$, augmented with independent exponential mean n (or $\text{Exp}(n)$ for short) edge weights $\{X_e : e \in E(K_n)\}$. For any subgraph $H = (V(H), E(H))$ of K_n we write $|H|$ for $|E(H)|$ and let

$$w(H) = \sum_{e \in E(H)} X_e.$$

For $i, j \in [n]$ we let P_{ij} be the minimum-weight path from i to j and write $W_{ij} = w(P_{ij})$ (adopting the convention $P_{ii} = \emptyset$, $w_{ii} = 0$). For any fixed vertex i , $\cup_{j \neq i} P_{ij}$ is a tree, the *shortest path tree* SPT_i rooted at i . For $t \geq 0$ and $k \in [n]$, we let $\text{SPT}_i(t)$ be the subtree of SPT_i induced by nodes j with $W_{ij} \leq t$.

Let t_k^i be the first time for which $\text{SPT}_i(t)$ contains at least $k+1$ vertices, and write $t_0^i = 0$. Due to the memoryless property of the exponential distribution, for each $k \in [n]$, the location of attachment of the vertex added at time t_k^i is uniform among the vertices of $\text{SPT}_i(t_{k-1}^i)$; in other words, $\text{SPT}_i(t_k^i)$ is distributed like a *random recursive tree* with k vertices. Random recursive trees have been well-studied [13]; in particular, it is known that the depth (number of edges on the path from a uniformly random node to the root) in a recursive tree with n vertices is asymptotic to $\log n$ in probability, see Devroye [3], and the height (greatest number of edges on any path starting from the root) is asymptotic to $e \log n$ in probability, see Devroye [3] and Pittel [9].

*Email: louigi@gmail.com, nicolas.broutin@m4x.org, lugosi@upf.es .

A variety of authors have studied the *weighted* structure of a single tree SPT_1 (or SPT_i for any i), and of the family of shortest paths $\{P_{ij} : i, j \in [n]\}$. Van der Hofstad, Hooghiemstra, and van Mieghem [14, 15] prove that, letting W be the total weight of SPT_1 ,

$$\frac{(W - \zeta(2))}{\sqrt{n}} \xrightarrow[n \rightarrow \infty]{} \mathcal{N}(0, 4\zeta(3))$$

in distribution. Janson [8] proved the beautiful result that

$$\frac{W_{12}}{\log n} \xrightarrow[n \rightarrow \infty]{} 1, \quad \frac{\max_{1 \leq j \leq n} W_{1j}}{\log n} \xrightarrow[n \rightarrow \infty]{} 2, \quad \text{and} \quad \frac{\max_{1 \leq i, j \leq n} W_{ij}}{\log n} \xrightarrow[n \rightarrow \infty]{} 3$$

in probability, and provided more detailed information about the distribution of W_{12} and of $\max_{1 \leq j \leq n} W_{1j}$. (Janson's results in fact hold for some edge weight distributions other than exponential, see Section 5.) On the other hand, the unweighted structure of shortest path trees has been studied by Hooghiemstra and Van Mieghem [7] who derived the limiting distribution of the pair $(W_{12}, |P_{12}|)$ suitably rescaled.

It is immediate from the aforementioned results for the depth of nodes and the height in random recursive trees that

$$\frac{|P_{12}|}{\log n} \xrightarrow[n \rightarrow \infty]{} 1 \quad \text{and} \quad \frac{\max_{1 \leq j \leq n} |P_{1j}|}{\log n} \xrightarrow[n \rightarrow \infty]{} e$$

in probability. Furthermore, as noted by Janson [8], the tail bounds for the height of random recursive trees established by Devroye [3] imply that

$$\limsup_{n \rightarrow \infty} \frac{\max_{1 \leq i, j \leq n} |P_{ij}|}{\log n} \leq \alpha^* \quad \text{in probability,} \quad (1)$$

where $\alpha^* \approx 3.5911$ is the unique solution of $\alpha \log \alpha - \alpha = 1$. Janson [8] asked whether there is a constant $c \in [e, \alpha^*]$ such that $\max_{1 \leq i, j \leq n} |P_{ij}| / \log n \rightarrow c$ in probability. It is the purpose of this note to answer Janson's question in the affirmative:

Theorem 1. *For any function $\omega(n)$ tending to $+\infty$ with n , for some constant $L > 0$, and for all $\delta \in (0, 1)$, for all n sufficiently large, with probability greater than $1 - \delta$,*

$$\alpha^* \log n - L \log \log n \leq \max_{1 \leq i, j \leq n} |P_{ij}| \leq \alpha^* \log n + \omega(n).$$

Furthermore,

$$\alpha^* \log n - L \log \log n \leq \mathbf{E} \left[\max_{1 \leq i, j \leq n} |P_{ij}| \right] \leq \alpha^* \log n + L.$$

Of course, it follows immediately from the first part of Theorem 1 that $\max_{1 \leq i, j \leq n} |P_{ij}| / \log n \rightarrow \alpha^*$ in probability. It is a stroke of luck that the limit

in Theorem 1 is $\alpha^* \log n$, as (1) and other results from Devroye [3] will make proving the upper bounds of the theorem easy. We will prove the lower bounds using the second moment method, applied to a suitably defined set of shortest paths with special properties that make them amenable to analysis. Before we get going, a brief remark is in order: we have to work with exponential edge weights in order to exploit the memoryless property of the exponential distribution. In the conclusion (Section 5) we explain how to extend our results to a range of other edge weights, including uniform edge weights.

2 A useful lemma and the upper bound

For $i \in [n]$ and $k \in [n-1]$ let $\tau_k^i = t_{k+1}^i - t_k^i$ denote the k -th interarrival time. By the memoryless property of the exponential distribution, $\{\tau_k^i, 1 \leq k < n\}$ are independent. Furthermore, τ_k^i is the minimum of $k(n-k)$ independent exponential random variables with mean n , hence τ_k^i has an exponential distribution with rate parameter $k(n-k)/n$. Using this fact, the following lemma provides upper bounds on the size of $\text{SPT}_i(t)$.

Lemma 2. *For all $t \geq 0$ and integers $m \in [n-1]$ with $m \geq e^t$,*

$$\mathbf{P}\{|\text{SPT}_i(t)| \geq m\} \leq 3 \cdot \sqrt{m/e^t} \cdot e^{-m/e^t}.$$

Proof. Observe that, for any $c > 0$,

$$\mathbf{P}\{|\text{SPT}_i(t)| \geq m\} = \mathbf{P}\left\{\sum_{k=1}^m \tau_k^i \leq t\right\} = \mathbf{P}\left\{\prod_{k=1}^m \exp(-c\tau_k^i) \geq e^{-ct}\right\}.$$

By Markov's inequality and the independence of the interarrival times $\tau_1^i, \dots, \tau_{n-1}^i$, we have

$$\mathbf{P}\{|\text{SPT}_i(t)| \geq m\} \leq e^{ct} \prod_{k=1}^m \mathbf{E}[\exp(-c\tau_k^i)] \leq e^{ct} \prod_{k=1}^m \frac{k}{k+c} \leq e^{ct} \cdot \frac{\Gamma(c+1)}{m^c},$$

where Γ is the Gamma function. Using the strengthening of Stirling's formula $\Gamma(c+1) \leq \sqrt{2\pi c}(c/e)^c e^{1/12}$ (see, e.g., [10]), and setting $c = m/e^t$, we obtain

$$\mathbf{P}\{|\text{SPT}_i(t)| \geq m\} \leq \sqrt{2\pi} e^{1/12} \cdot \sqrt{c} \cdot e^{-c} \leq 3 \cdot \sqrt{c} \cdot e^{-c},$$

as claimed. \square

We will also use the following result from Devroye [3], providing an upper bound on the heights of random recursive trees. In the following, $h(T)$ denotes the height of a rooted tree T .

Theorem 3 ([3], Theorem 9). *Let T_m be a random recursive tree with $m \geq 2$ nodes. Then, for $x > 1$,*

$$\mathbf{P}\{h(T_m) \geq x \log m\} \leq e^{x-1} m^{x-x \log x}.$$

The upper bounds in Theorem 1 follow immediately from Theorem 3.

Proof of Theorem 1, upper bounds. For any $t \geq 0$, by a union bound we have, for $n > 2$,

$$\mathbf{P} \left\{ \max_{1 \leq i, j \leq n} |P_{ij}| \geq \alpha^* \log n + t \right\} \leq n \mathbf{P} \left\{ \max_{1 \leq j \leq n} |P_{1j}| \geq \alpha^* \log n + t \right\}.$$

The event in the right-hand side above holds if the SPT_i has height at least $\alpha^* \log n + t$. Observe also that, by the definition of $\alpha^* \geq e$, for $t \geq 0$,

$$(\alpha^* + t) \log(\alpha^* + t) - (\alpha^* + t) \geq 1 + t.$$

Since SPT_i is distributed like a random recursive tree on n nodes, Theorem 3 then implies that

$$\mathbf{P} \left\{ \max_{1 \leq i, j \leq n} |P_{ij}| \geq \alpha^* \log n + t \right\} \leq e^{\alpha^* + t / \log n} e^{-t}. \quad (2)$$

Choosing $t = t(n) \rightarrow \infty$ proves the upper bound in probability. Now, using the bound (2), we have, for $n \geq 8$,

$$\begin{aligned} \mathbf{E} \left[\max_{1 \leq i, j \leq n} |P_{ij}| \right] &\leq \alpha^* \log n + 1 + \sum_{\ell=1}^{\infty} 2^\ell \mathbf{P} \left\{ \max_{1 \leq i, j \leq n} |P_{ij}| \geq \alpha^* \log n + 2^{\ell-1} \right\} \\ &\leq \alpha^* \log n + 1 + e^{\alpha^*} \cdot \sum_{\ell=1}^{\infty} 2^\ell \exp(-2^{\ell-2}) \\ &= \alpha^* \log n + O(1), \end{aligned}$$

proving the upper bound in expectation. \square

3 Towards the lower bound

Fix a path $P = v_1, \dots, v_{k+1}$, $k \leq 12 \log n$ (where 12 is somewhat arbitrary, the key point being that it's bigger than $3\alpha^*$). It is easily seen that for $w > 0$, letting $\text{Po}(w)$ denote a Poisson mean w random variable, we have

$$\mathbf{P} \{w(P) \leq wn\} = \mathbf{P} \{\text{Po}(w) \geq k\} = \sum_{i=k}^{\infty} \frac{w^i e^{-w}}{i!}.$$

When $w = o(1)$, the above sum is dominated by its first term and $e^{-w} = 1 - o(1)$, so we have

$$\mathbf{P} \{w(P) \leq wn\} = (1 + o(1)) \frac{w^k}{k!}. \quad (3)$$

We next show that *given* that P has small weight, it is very likely to be the minimum-weight path between its endpoints. More precisely, let $P^*(v, w)$ be the minimum-weight path in K_n between two vertices v and w . Then we have the following.

Lemma 4. Fix $c > 2$. For all n sufficiently large, for all $\epsilon > c \log \log n / \log n$ and any path $P = v_1, \dots, v_{k+1}$ in K_n with $1 \leq k \leq 12 \log n$, we have

$$\mathbf{P} \{P \neq P^*(v_1, v_{k+1}) \mid w(P) \leq (1 - \epsilon) \log n\} \leq \frac{13k^2}{n^\epsilon}.$$

The restriction $c > 2$ is not necessary in the above lemma, but the upper bound becomes trivial when $\epsilon \leq 2 \log \log n / \log n$.

Proof. Let $\text{SPT}_v(t)$ denote the shortest path tree started from a vertex v and stopped at time t . For $k \geq 2$, if $P \neq P^*(v_1, v_{k+1})$ then it must be the case that for some $i = 1, \dots, k$, in $K_n \setminus E(P)$, $\text{SPT}_{v_i}((1 - \epsilon) \log n)$ contains one of v_{i+1}, \dots, v_{k+1} .

Allowing connections using the edges along the path P only increases the probability of this event. Let K_n^P be K_n where the edge weights X_e , $e \in E(P)$, along the path P have been replaced by independent copies. Let A_i be the event that, in K_n^P , $\text{SPT}_{v_i}((1 - \epsilon) \log n)$ contains one of v_{i+1}, \dots, v_{k+1} for some $i = 1, \dots, k$. The remark above then implies by the union bound that

$$\mathbf{P} \{P \neq P^*(v_1, v_{k+1}) \mid w(P) \leq (1 - \epsilon) \log n\} \leq \sum_{i=1}^k \mathbf{P} \{A_i\}. \quad (4)$$

Let N_i denote the number of nodes in $\text{SPT}_{v_i}((1 - \epsilon) \log n)$ in K_n^P . Observe that, conditioning on N_i , the probability that A_i does *not* occur is

$$\begin{aligned} 1 - \mathbf{P} \{A_i \mid N_i\} &= \left(\prod_{j=1}^{N_i} \left(1 - \frac{k-i+1}{n-j} \right) \right) \cdot \mathbf{1}[n - N_i > k - i + 1] \\ &\geq \max \left(1 - \frac{(k-i+1)N_i}{n - N_i}, 0 \right), \end{aligned} \quad (5)$$

where $\mathbf{1}[\cdot]$ denotes the indicator function. Let $E_i(0)$ be the event that $N_i \leq 2n^{1-\epsilon}$, for $j = 1, \dots, \lfloor \log_2(n^\epsilon/2k) \rfloor - 1$, let $E_i(j)$ be the event that $2^j n^{1-\epsilon} < N_i \leq 2^{j+1} n^{1-\epsilon}$, and let $E_i(\lfloor \log_2(n^\epsilon/2k) \rfloor)$ be the event that $N_i > 2^{\lfloor \log_2(n^\epsilon/2k) \rfloor} n^{1-\epsilon}$. Expressing $\mathbf{P} \{A_i\}$ as a sum of conditional probabilities, we have

$$\begin{aligned} \mathbf{P} \{A_i\} &= \sum_{j=0}^{\lfloor \log_2(n^\epsilon/2k) \rfloor} \mathbf{P} \{A_i \mid E_i(j)\} \cdot \mathbf{P} \{E_i(j)\} \\ &\leq \sum_{j=0}^{\lfloor \log_2(n^\epsilon/2k) \rfloor} \mathbf{P} \{A_i \mid E_i(j)\} \cdot 3\sqrt{2^j} e^{-2^j}, \end{aligned} \quad (6)$$

by Lemma 2. Using (5) to bound $\mathbf{P}\{A_i \mid E_i(j)\}$,

$$\begin{aligned} \mathbf{P}\{A_i\} &\leq \sum_{j=0}^{\lfloor \log_2(n^\epsilon/2k) \rfloor - 1} \frac{(k-i+1)2^{j+1}n^{1-\epsilon}}{n-2^{j+1}n^{1-\epsilon}} \cdot \frac{3\sqrt{2^j}}{e^{2^j}} + \frac{3\sqrt{2^{\lfloor \log_2(n^\epsilon/2k) \rfloor}}}{e^{2^{\lfloor \log_2(n^\epsilon/2k) \rfloor}}} \\ &\leq \frac{6(k-i+1)}{n^\epsilon(1-1/2k)} \sum_{j=0}^{\infty} \frac{2^{3j/2}}{e^{2^j}} + \frac{3\sqrt{2^{\lfloor \log_2(n^\epsilon/2k) \rfloor}}}{e^{2^{\lfloor \log_2(n^\epsilon/2k) \rfloor}}} \\ &\leq \frac{12(k-i+1)}{n^\epsilon} + n^{-\epsilon}, \end{aligned} \quad (7)$$

where we have used the fact that $\sum_{j \geq 0} 2^{3j/2}e^{-2^j} \leq 1$ and the last inequality holds for all $k \leq 12 \log n$ and n sufficiently large. It follows that, since $k \geq 1$,

$$\sum_{i=1}^k \mathbf{P}\{A_i\} \leq \frac{1}{n^\epsilon} \left(12 \binom{k+1}{2} + k \right) \leq \frac{13k^2}{n^\epsilon},$$

which, together with (4), completes the proof. \square

The following fact follows immediately from the proof of Lemma 4 (there are just fewer terms in the sum (6)); it will be useful later.

Corollary 5. *Fix $c > 2$. For all n sufficiently large, for all $\epsilon > c \log \log n / \log n$ and any path $P = v_1, \dots, v_{k+1}$ in K_n with $1 \leq k \leq 12 \log n$, we have*

$$\mathbf{P}\{P \neq P^*(v_1, v_{k+1}) \mid (1-2\epsilon) \log n \leq w(P) \leq (1-\epsilon) \log n\} \leq \frac{13k^2}{n^\epsilon}.$$

It follows from Lemma 4, at least intuitively, that to get a lower bound of k on the length of the longest minimum-weight path, we can instead bound from below the probability that there is *some* path with k edges of weight $(1-\epsilon) \log n$. Given a positive integer k and real $\epsilon > 0$, let $\mathcal{P}_{k,\epsilon} = \mathcal{P}_{k,\epsilon}(n)$ be the set of paths with k edges and weight at most $(1-\epsilon) \log n$ in K_n . As a first step, we remark that by (3),

$$\mathbf{E}|\mathcal{P}_{k,\epsilon}| \sim n^{k+1} \left(\frac{(1-\epsilon) \log n}{n} \right)^k \frac{1}{k!} \sim \frac{n}{\sqrt{2\pi k}} \left(\frac{e \log n}{k} \right)^k (1-\epsilon)^k. \quad (8)$$

Let α_ϵ be the solution of

$$\alpha \log \alpha - \alpha(1 + \log(1-\epsilon)) = 1, \quad (9)$$

so $\alpha_\epsilon < \alpha^*$ and $\alpha_\epsilon \rightarrow \alpha^* > e$ as $\epsilon \rightarrow 0$. When $t = k - \alpha_\epsilon \log n = o(\sqrt{\log n})$, we then have

$$\mathbf{E}|\mathcal{P}_{k,\epsilon}| = \frac{(1+o(1))}{\sqrt{2\pi\alpha_\epsilon \log n}} \left(\frac{e}{\alpha_\epsilon} \right)^t (1-\epsilon)^t,$$

and letting

$$k_\epsilon = \alpha_\epsilon \log n - \beta_\epsilon \log \log n \quad \text{where} \quad \beta_\epsilon = \frac{1}{2(1 + \log \alpha_\epsilon - \log(1-\epsilon))} \quad (10)$$

we have the following.

Lemma 6. *Let $\alpha_\epsilon, \beta_\epsilon$ and k_ϵ be defined as in (9) and (10) above. Given ϵ with $0 < \epsilon < 1/2$, there is a positive constant c_ϵ such that $\mathbf{E}|\mathcal{P}_{k_\epsilon, \epsilon}| = (1 + o(1))c_\epsilon$ and, for $|t| = o(\sqrt{\log n})$,*

$$\mathbf{E}|\mathcal{P}_{k_\epsilon+t, \epsilon}| = (1 + o(1))c_\epsilon \cdot \left(\frac{e}{\alpha_\epsilon}\right)^t (1 - \epsilon)^t.$$

In particular, it follows that if $t = t(n) \rightarrow -\infty$ then $\mathbf{E}|\mathcal{P}_{k_\epsilon+t, \epsilon}(n)| \rightarrow \infty$. To derive a lower bound on the probability that there exists such a path, we use the second moment method, and now introduce the version of it we require. Given any random set \mathcal{S} of paths in K_n and two paths P and Q , let $q_{\mathcal{S}}(P, Q)$ be the probability that P and Q are both in \mathcal{S} , and let

$$\Delta(\mathcal{S}) = \sum_{P, Q} q_{\mathcal{S}}(P, Q), \quad (11)$$

where the sum is over pairs P, Q of distinct but intersecting paths in K_n . By Corollary 4.3.4 of Alon et al. [1], we then have

$$\mathbf{P}\{|\mathcal{S}| = 0\} \leq \frac{1}{\mathbf{E}|\mathcal{S}|} + \frac{\Delta(\mathcal{S})}{(\mathbf{E}|\mathcal{S}|)^2}. \quad (12)$$

Given the preceding discussion, a natural choice for the set \mathcal{S} would be $\mathcal{P}_{k_\epsilon+t, \epsilon}$, for some $t = t(n)$ tending to $-\infty$ with n . Unfortunately, for this choice of \mathcal{S} and for the values of t we wish to consider, the quantity $\Delta(\mathcal{S})$ is too large for (12) to yield a useful bound. However, it is both useful and instructive to proceed as though this was our choice of \mathcal{S} , and see how far we can get.

Given paths P and Q in K_n , let

$$q_\epsilon(P, Q) = \mathbf{P}\{w(P) \leq (1 - \epsilon) \log n, w(Q) \leq (1 - \epsilon) \log n\}.$$

Also, for $t \in \mathbb{R}$, let

$$\Delta_t = \Delta_t(n, \epsilon) = \sum_{P, Q} q_\epsilon(P, Q) \quad (13)$$

where the sum is over pairs P, Q of distinct but intersecting paths with $\lceil k_\epsilon + t \rceil$ edges in K_n . (So, Δ_t is just $\Delta(\mathcal{S})$ when \mathcal{S} is the set of paths in $\mathcal{P}_{\lceil k_\epsilon + t \rceil, \epsilon}$.)

3.1 Light intersecting paths

In order to bound Δ_t , we first decompose the sum in (13). For a given real number t and integers i, j with $1 \leq j \leq i < \lceil k_\epsilon + t \rceil$, let $\Delta_{t, i, j}$ be the sum of $q_\epsilon(P, Q)$ over paths P, Q with $\lceil k_\epsilon + t \rceil$ edges, such that P and Q share i edges, and these i edges form precisely j connected components. Then, we have

$$\Delta_t = \sum_{1 \leq j \leq i < \lceil k_\epsilon + t \rceil} \Delta_{t, i, j}. \quad (14)$$

We now need to consider (a) the probability that two paths of k edges (for k near to k_ϵ) sharing i edges both have weight at most $(1 - \epsilon) \log n$, and (b) the number of such pairs of paths. The following lemma bounds the former probability. Counting the number of terms of a given sum $\Delta_{t,i,j}$ is the subject of Section 3.2.

Lemma 7. *Given two paths P, Q in K_n , each consisting of k edges, i of which are common to P and Q , and any $s \geq 0$, we have,*

$$\mathbf{P}\{w(P) \leq sn, w(Q) \leq sn\} \leq 4^{k-i} \cdot \frac{s^{2k-i}}{(2k-i)!}.$$

Proof. If we consider exponential edge weights with mean 1 instead of the exponential mean n for the individual edge weights, we are lead to evaluate $\mathbf{P}\{w(P) \leq s, w(Q) \leq s\}$. Observe that the sum of ℓ exponential random variables is has a Gamma($1, \ell$) distribution, with density function $f_\ell(t) = t^{\ell-1}e^{-t}/(\ell-1)!$ and distribution function $F_\ell(t)$. So, conditioning on the aggregated weight of the i shared edges, we see that

$$\mathbf{P}\{w(P) \leq s, w(Q) \leq s\} = \int_0^s f_i(t)F_{k-i}(t)^2 dt.$$

However, for $t \geq 0$, $f_\ell(t) \leq t^{\ell-1}/(\ell-1)!$ and $F_\ell(t) \leq t^\ell/\ell!$, which implies that

$$\mathbf{P}\{w(P) \leq s, w(Q) \leq s\} \leq \int_0^s \frac{t^{i-1}}{(i-1)!} \frac{(s-t)^{2(k-i)}}{(k-i)!^2} dt = \frac{s^{2k-i}}{(2k-i)!} \binom{2(k-i)}{k-i}.$$

Using the classical bound for the central binomial coefficients $\binom{2n}{n} \leq 4^n$ completes the proof. \square

3.2 The number of intersecting pairs

What is actually need is to count the number of *pairs* (P, Q) , where P and Q are two paths containing k edges such that $P \cap Q$ has i edges in j connected components. Our aim is to analyze shortest paths in the graph, and hence it suffices to consider *self-avoiding* paths (that do not intersect themselves even at vertices). Let $N_{k,i,j}$ denote the number of such pairs.

Lemma 8. *The number of pairs of self-avoiding paths (P, Q) of length k such that $P \cap Q$ contains i edges in j connected components satisfies*

$$N_{k,i,j} \leq n^{2k+2-i-j} (2k^3)^j.$$

Proof. Observe first that if $k < i + 2j - 2$ then $N_{k,i,j} = 0$, since for a given path P , if Q shares i edges in j connected components with P then Q has at least $i + 2j - 2$ edges. We now assume that $k \geq i + 2j - 2$.

We first focus on the choice of P with the edges of $P \cap Q$ distinguished. First fix a path P , self-avoiding, as an ordered sequence of $k + 1$ vertices. There

are $\binom{n}{k+1} \cdot (k+1)!$ such choices. (In fact, we are double counting here, as the reversed sequence would yield the same path P ; we can afford to ignore this fact in obtaining our upper bound.) We next choose the sequence of sizes of the j parts of $P \cap Q$ as they appear along P . Since $P \cap Q$ contains i edges, it suffices to split i into j ordered parts; there are $\binom{i+j}{j}$ possibilities for this partition. Now that we have the sequence of sizes of the parts of $P \cap Q$, it remains to choose the positions in P where these edges appear. We can count the number of such choices as follows: once we have removed $P \cap Q$, there remains a sequence of $k-i$ edges of P . The portions of $P \cap Q$ can be inserted at any of the $k-i+1$ separating positions (the extremities are included), and hence there are $\binom{k-i+1}{j}$ choices. So in the end, fixing the path P together with the edges of $P \cap Q$ can be done in

$$\binom{n}{k+1} (k+1)! \cdot \binom{i+j}{j} \cdot \binom{k-i+1}{j} \quad (15)$$

distinct ways.

It now remains to choose the second path Q , also self-avoiding, so that it intersects P at the distinguished edges. We first choose the order in which the j parts of $P \cap Q$ appear in Q : there are $j!$ possible choices. $P \cap Q$ contains $i+j$ vertices and we need $k+1-i+j$ other vertices to complete Q . It is possible that some vertices of $Q \setminus (P \cap Q)$ are in P , so we can choose an ordered sequence of these vertices in $\binom{n-i-j}{k+1-i-j} (k+1-i-j)!$ different ways. (In fact, not all such sequences yield valid choices for $Q \setminus P$. For example, $Q \setminus P$ should not contain two consecutive vertices from P in the same part, or $Q \cap P$ will not be what we claimed. However, we only seek an upper bound, and so can afford to ignore this issue.) Finally, we need to choose how the j (now ordered) parts of $P \cap Q$ interlace with these $(k+1-i-j)$ extra vertices. The extra vertices define $k-i-j+1$ intervals (with the extremities): there are $\binom{k-i-j+1}{j}$ ways to choose j of them (and insert the parts of $P \cap Q$ at these spots). It follows by this argument that for a fixed path P with distinguished edges forming j parts, there are *at most*

$$j! \cdot \binom{n-i-j}{k+1-i-j} (k+1-i-j)! \cdot \binom{k-i-j+1}{j} \quad (16)$$

possible choices for the path Q .

Now, the desired number $N_{k,i,j}$ of pairs (P, Q) of paths of length k such that $P \cap Q$ contains i edges in j connected parts is *at most* the product of the numbers appearing in (15) and (16). The desired bound follows by routine bounding. \square

3.3 Paths intersecting at least twice

Lemmas 7 and 8 yield more than sufficient control over all $\Delta_{t,i,j}$ with $j \geq 2$, which we quantify in the following lemma.

Lemma 9. *There is an $\epsilon_0 > 0$ such that for all $0 < \epsilon \leq \epsilon_0$, if $k = \alpha_\epsilon \log n + o(\sqrt{\log n})$, then for all n sufficiently large*

$$\frac{\sum_{2 \leq j \leq i < k} \Delta_{t,i,j}}{\mathbf{E}[|\mathcal{P}_{k,\epsilon}|]^2} \leq n^{-0.95}.$$

Proof. Combining the bound on the number of intersecting paths (Lemma 8) with Lemma 7 with $s = (1 - \epsilon) \log n/n$, we obtain

$$\begin{aligned} \Delta_{t,i,j} &:= N_{k,i,j} \cdot \mathbf{P}\{w(P) \leq sn, w(Q) \leq sn\} \\ &\leq n^{2-j} (2k^3)^j ((1 - \epsilon) \log n)^{2k-i} \frac{4^{k-i}}{(2k-i)!} \\ &\leq n^{2-j} (2k^3)^j \left(\frac{(1 - \epsilon)e \log n}{2k-i} \right)^{2k-i} 4^{k-i}, \end{aligned}$$

since, for $\ell \geq 1$, we have $\ell! \geq (\ell/e)^\ell$. We have assumed that $j \geq 2$, so for n large enough, $(2k^3)^j n^{2-j} \leq 4k^6$. It follows that

$$\sum_{2 \leq j \leq i < k} \Delta_{t,i,j} \leq 4k^8 \max_{2 \leq i < k} \left\{ 4^{k-i} \left(\frac{e(1 - \epsilon) \log n}{(2k-i)} \right)^{2k-i} \right\}.$$

Using the estimate (8) for $\mathbf{E}|\mathcal{P}_{k,\epsilon}|$, it follows that, for n large enough,

$$\frac{\sum_{2 \leq j \leq i < k} \Delta_{t,i,j}}{\mathbf{E}[|\mathcal{P}_{k,\epsilon}|]^2} \leq \frac{7\pi k^9}{n^2} \max_{2 \leq i < k} \left\{ \left(\frac{4e(1 - \epsilon) \log n}{2k-i} \right)^{-i} \left(\frac{2k}{2k-i} \right)^{2k} \right\}. \quad (17)$$

We bound the right-hand side of (17) by allowing i to take real values. We write $k = \gamma \log n$ and $x = \beta \log n$ (so $0 \leq \beta \leq \gamma$); we then see that

$$\left(\frac{4e(1 - \epsilon) \log n}{2k - x} \right)^{-x} \left(\frac{2k}{2k - x} \right)^{2k} = e^{g(\beta) \log n},$$

where

$$g(\beta) := \beta \log \left(\frac{2\gamma - \beta}{4e(1 - \epsilon)} \right) + 2\gamma \log \left(\frac{2\gamma}{2\gamma - \beta} \right).$$

For γ close to α^* , the function g is suitably approximated by g^* defined by

$$g^*(\beta) := \beta \log \left(\frac{2\alpha^* - \beta}{4e} \right) + 2\alpha^* \log \left(\frac{2\alpha^*}{2\alpha^* - \beta} \right).$$

By differentiation, we see that $g^*(\beta)$ is maximized on $[0, \alpha^*]$ by taking $\beta = 2\alpha^* - 4$, at which point $g^*(\beta) \approx 1.02$. For any δ , we may ensure that $|\gamma - \alpha^*| \leq \delta$ for n large by choosing ϵ sufficiently small. Since $g(\beta) = g_\gamma(\beta)$ is bounded and continuous in both γ and β away from $\beta \in [2\gamma, \infty)$, for all ϵ sufficiently small and all n sufficiently large we have

$$|g(\beta) - g^*(\beta)| \leq 0.01$$

for all $k = \alpha_\epsilon \log n + o(\sqrt{\log n})$ and all $\beta \in [0, \gamma]$. It follows that $g(\beta) \leq 1.04$ for all $\beta \in [0, \gamma]$, which, combined with (17), yields that

$$\frac{\sum_{2 \leq j \leq i < k} \Delta_{t,i,j}}{\mathbf{E} [|\mathcal{P}_{k,\epsilon}|]^2} \leq \frac{7\pi k^9}{n^{0.96}} = O\left(\frac{\log^9 n}{n^{0.96}}\right).$$

This proves the lemma. \square

4 Dealing with paths intersecting only once

Unfortunately, for paths intersecting only once (i.e., when $j = 1$), for some values of i the quantity $\Delta_{t,i,1}$ is too large for us to apply the straightforward approach used above. (Note that this is not an artefact of our upper bounds: Lemmas 7 and 8 are essentially tight up to logarithmic factors.) The most natural and naive way to deal with this complication is to simply throw away all pairs of shortest paths P and Q whose intersection is connected, and try to bound the probability that one of the remaining minimum-weight paths is “long” (in the same sense as above). This is essentially our approach. However, in order to keep a handle on the conditioning imposed in doing so, it is useful to proceed “from the other direction”: building a set of shortest paths with special properties that guarantee that (a) no pair of such paths has an intersection which is connected, yet (b) the set contains a minimum-weight path with about $\alpha^* \log n$ edges.

In order to describe this set, we first need to introduce a few concepts. We say that paths P and P' *intersect once* if $P \cap P'$ has only one connected component containing at least one edge. (There may be other components of $P \cap P'$ which are isolated vertices.) We say that a path P is a *local optimum* if, for all paths P' with $|P| = |P'|$ that intersect P once, we have $w(P') > w(P)$. From now on, let \mathcal{O}_k denote the set of paths with k edges that are local optima, let $\mathcal{O}_{k,\epsilon}$ denote the set of elements of \mathcal{O}_k with weight at most $(1 - \epsilon) \log n$. Also, let $\mathcal{P}^* = \mathcal{P}(n) = \{P_{ij} : 1 \leq i < j \leq n\}$ denote the set of shortest paths in K_n . Note that $\mathcal{O}_{k,\epsilon} \subseteq \mathcal{P}_{k,\epsilon}$.

AN INSTRUCTIVE EXAMPLE. To motivate our next definitions, consider the unlikely (impossible) but instructive event that we find a path $P = (v_1, \dots, v_{k+1})$, all of whose edges have weight exactly 1, and consider a path Q of length k intersecting P in exactly $i < k$ consecutive edges —say $Q \cap P = (v_1, \dots, v_{i+1})$, for example. Let Q_1 (respectively Q_2) be the component of $Q \setminus P$ containing v_1 (respectively v_{i+1}). In order that $w(Q) \leq k$, then certainly we must have both $w(Q_1) \leq k - i$ and $w(Q_2) \leq k - i$. On the other hand, one of Q_1 and Q_2 has at least $(k - i)/2$ edges. It follows that if $w(Q) \leq k$ then in K_n^P either $\text{SPT}_{v_1}(k - i)$ or $\text{SPT}_{v_{i+1}}(k - i)$ has height at least $(k - i)/2$.

This observation is key to our approach. Roughly speaking, we wish to consider long minimum-weight paths with the special property that none of the shortest path trees leaving them are “too tall”. This property will guarantee that the paths are local optima (Lemma 14), and we can then use the second

moment method to prove concentration for the number of such paths. However, the idea suggested by the above sketch gives a little too much away in terms of what “too tall” means; in particular, it is not careful enough about the interplay between the contributions of Q_1 and Q_2 to the length of Q . When we formalize our idea, we will have to deal with this interplay to make the details work out.

4.1 The fluctuations of conditioned partial sums

In the above example, there was a second simplification: the edge weights have the extremely desirable property that for every subpath P' of P , $w(P')$ is *exactly* $\mathbf{E}[w(P') \mid w(P) = k]$, which makes the conditions we need to impose on the shortest path trees leaving P very easy to state. It is, of course, too much to ask for paths whose edge weights are as well behaved as in the above example. In general, the weights of subpaths will fluctuate considerably from their conditional expected value given $w(P)$. Our tool for controlling the size of these fluctuations is the *law of the iterated logarithm*; we will use the version found in Rogers and Williams [11]. (These bounds are in fact stronger than we require; however, we did not see a substantially simpler proof of a simpler but sufficient result; furthermore, Corollary 12 is perhaps of interest independent of its role in the current work.)

Theorem 10 ([11], Corollary 16.5). *Let Y_1, Y_2, \dots be independent and identically distributed random variables with finite variance, and let $S_n = \sum_{j=1}^n Y_j$. Then*

$$\mathbf{P} \left\{ \limsup_{n \rightarrow \infty} \frac{S_n - n\mathbf{E}Y_1}{\sqrt{\mathbf{Var}[Y_1] 2n \log \log n}} = 1 \right\} = 1.$$

We then immediately have the following finite statement, more useful for our purposes.

Corollary 11. *Let Y_1, Y_2, \dots be independent and identically distributed random variables with finite variance, and let $S_n = \sum_{j=1}^n Y_j$. For all $\delta > 0$ there is a constant $C > 0$ depending on δ and the distribution of Y_1 (we write $C(\delta, Y_1)$ for short) such that*

$$\mathbf{P} \left\{ \sup_{n \geq 1} \left| \frac{S_n - n\mathbf{E}Y_1}{\sqrt{\mathbf{Var}[Y_1] 2n \log \log n}} \right| \leq C \right\} \geq 1 - \delta.$$

(To avoid annoying technicalities in our formulae, in Corollary 11 and hereafter when we write $\log \log n$ we mean $\max(\log \log n, 1)$.) At this point, the fact that we are considering exponential random variables comes in very handy, as it allows us to apply Corollary 11 conditional upon the value of S_n . More precisely, we have

Corollary 12. *Suppose Y_1 is an exponential random variable. Then for all $\delta > 0$, there is $C' = C'(\delta)$ such that for all n sufficiently large,*

$$\mathbf{P} \left\{ \forall k \in [n], \left| \frac{S_k - S_n(k/n)}{\sqrt{\mathbf{Var}[Y_1 \mid S_n] 2k \log \log k}} \right| \leq C' \mid S_n \right\} \geq 1 - \delta. \quad (18)$$

We emphasize that the probability in (18) is a *random variable*, measurable with respect to S_n ; the content of the lemma is that this random variable is *deterministically* at least $1 - \delta$. In particular, if we let \mathcal{E} be the event in (18) (so (18) bounds $\mathbf{P}\{\mathcal{E} \mid S_n\}$), then $\mathbf{P}\{\mathcal{E}\} = \mathbf{E}_{S_n}[\mathbf{P}\{\mathcal{E} \mid S_n\}] \geq 1 - \delta$ and more strongly, for any event \mathcal{E}' which is measurable with respect to S_n ,

$$\mathbf{P}\{\mathcal{E} \mid \mathcal{E}'\} = \mathbf{E}_{S_n}[\mathbf{P}\{\mathcal{E} \mid S_n\} \mid \mathcal{E}'] \geq 1 - \delta.$$

Proof of Corollary 12. Fix $\delta > 0$ and n . For $i = 1, \dots, n$, let $E_i = S_i/S_n$. Then the set $\{E_1, \dots, E_n\}$ is independent of S_n and distributed like n independent $[0, 1]$ -uniform random variables (for a proof, see Chapter 8 of [12]). It follows in particular that E_1 , the minimum of n independent uniforms, has a $\text{Beta}(1, n-1)$ distribution and $\mathbf{Var}[Y_1 \mid S_n] = (S_n/n)^2(n-1)/(n+1) \geq (S_n/n)^2/3$, for $n \geq 2$. Furthermore, the independence ensures that proving (18) reduces to finding C' such that the following *unconditioned* statement holds

$$\mathbf{P}\left\{\sup_{1 \leq k \leq n} \left| \frac{(n/S_n)S_k - k}{\sqrt{2k \log \log k}} \right| \leq C'\right\} \geq 1 - \delta. \quad (19)$$

By choosing C' large enough we can certainly ensure that this holds for small n ; shortly we will take advantage of this fact. By Corollary 11 we have that

$$\mathbf{P}\left\{\sup_{1 \leq k \leq n} \left| \frac{S_k - k}{\sqrt{2k \log \log k}} \right| \leq C_1\right\} \geq 1 - \frac{\delta}{2}, \quad (20)$$

where $C_1 = C(\delta/2, Y_1)$ as in the statement of Corollary 11. If, on the other hand, there is $k \in \{1, \dots, n\}$ such that

$$\left| \frac{(n/S_n)S_k - k}{\sqrt{2k \log \log k}} \right| \geq C_1 + c,$$

for some $c > 0$, then either the event whose probability is bounded in (20) must fail to hold, or by the triangle inequality we must have

$$\left| \frac{(n/S_n - 1)S_k}{\sqrt{2k \log \log k}} \right| \geq c,$$

for some k , or more simply $|S_n - n| \geq c\sqrt{2k \log \log k} \cdot S_n/S_k$. For this to occur, we must either have $|S_n - n| \geq \sqrt{cn}$ or $S_k \geq \sqrt{c(2k \log \log k)} \cdot (\sqrt{n} - \sqrt{c})$. It follows that, for n large enough that $\sqrt{n} - \sqrt{c} \geq \sqrt{n/2}$ (and recalling that $\log \log k \geq 1$ for all k by convention), we have

$$\begin{aligned} \mathbf{P}\left\{\sup_{1 \leq k \leq n} \left| \frac{(n/S_n - 1)S_k}{\sqrt{2k \log \log k}} \right| \geq c\right\} &\leq \sum_{k=1}^n \mathbf{P}\{S_k \geq k\sqrt{c}\} + \mathbf{P}\{|S_n - n| \geq \sqrt{cn}\} \\ &\leq \sum_{k=1}^n e^{-kH(\sqrt{c})} + \frac{1}{c}, \end{aligned} \quad (21)$$

where we have bounded the first term using a large deviations bound for sums of exponential random variables [2] with $H(x) := x - 1 - \log x$, and the second term by Chebyshev's inequality. We now choose c large enough that the right-hand side of (21) above is at most $\delta/2$ for all n . Combining (21) with (20) and taking $C' = C_1 + c$ (or slightly larger if necessary, to deal with small n) then establishes (19) and completes the proof. \square

Motivated by Corollary 12, we say that a path $P = (v_1, \dots, v_{k+1})$ in K_n is C -legal (for a given $C > 0$), and write $P \in \mathcal{L}_C = \mathcal{L}_C(n)$, if for all $i = 1, \dots, k$,

$$\left| \frac{k}{w(P)} \cdot \sum_{j=1}^i X_{v_j v_{j+1}} - i \right| \leq C \sqrt{2i \log \log i}, \quad (22)$$

and

$$\left| \frac{k}{w(P)} \cdot \sum_{j=k-i+1}^k X_{v_j v_{j+1}} - i \right| \leq C \sqrt{2i \log \log i}. \quad (23)$$

We need both conditions as we wish to ensure the path is well-behaved “from both ends”. Though these conditions only restrict the weight of subpaths of P starting from an end, they can be combined to control the behavior of other subpaths. More precisely, let

$$m(i, k) = \min\{i - 1, k + 1 - i\} \quad \text{and} \quad s(i, j, k) = \max\{m(i, k), m(j, k)\}. \quad (24)$$

Intuitively, $s(i, j, k)$ is the furthest distance from either i or j to one of the ends of the path. We then have

Lemma 13. *If $P = (v_1, \dots, v_{k+1})$ is C -legal then for all $1 \leq i < j \leq k + 1$,*

$$\left| \frac{k}{w(P)} \cdot \sum_{m=i}^{j-1} X_{v_m v_{m+1}} - (j - i) \right| \leq 2C \cdot \sqrt{2s(i, j, k) \log \log(s(i, j, k))}. \quad (25)$$

Proof. Write $P(i)$ (respectively $P(i, j)$) for the path (v_1, \dots, v_i) (respectively (v_i, \dots, v_j)), then apply (22) and (23) to $P(i)$ or $P \setminus P(i)$ according as $i \leq (k + 1)/2$ or $i > (k + 1)/2$ and to $P(j)$ or $P \setminus P(j)$, similarly. Finally, use one of the equalities

$$P(i, j) = P(j) \setminus P(i-1) = (P \setminus P(i-1)) \setminus (P \setminus P(j)) = P \setminus (P(i-1) \cup (P \setminus P(j))),$$

together with the triangle inequality to prove the result. \square

4.2 Bounding the heights of shortest path trees

We have now introduced all the technical apparatus we will need to control the path P itself. Returning to our example from page 11, we recall that our main goal is to bound the *heights of the shortest path trees leaving P* ; Lemma 13 tells us just how tall we can allow them to be. We say that P is C -bonsai if:

(\star) for all $i = 1, \dots, k+1$, and all integers ℓ with $\ell \geq m(i, k)/40$, in K_n^P

$$\text{SPT}_{v_i}(\ell + 2C\sqrt{500\ell \log \log \ell}) \text{ has height less than } \left(\frac{9k}{10w(P)} \right) \cdot \ell.$$

We denote the set of C -bonsai paths in K_n by $\mathcal{B}_C = \mathcal{B}_C(n)$.

Lemma 14. *For any $C \geq 0$, if $w(P) \leq k \leq 4w(P)$ and P is both C -legal and C -bonsai, i.e., $P \in \mathcal{L}_C \cap \mathcal{B}_C$, then P is a local optimum.*

Proof of Lemma 14. Let $P = (v_1, \dots, v_{k+1}) \in \mathcal{L}_C \cap \mathcal{B}_C$. Our aim is to show that for any integers i, j with $1 \leq i < j \leq k+1$ and any path Q of length k whose intersection with P is (v_i, \dots, v_j) , we have $w(Q) > w(P)$. By symmetry we may assume that $i < (k+1)/2$ (so $m(i, k) = i-1$) and that $m(i, k) \leq m(j, k)$ (so $s(i, j, k) = m(j, k)$).

Let Q_i (respectively Q_j) be the component of $Q \setminus P$ containing v_i (respectively v_j). We must have $|Q_i| + |Q_j| = k - (j - i) \geq m(i, k) + m(j, k)$. Let $k_i = (10/9)(w(P)/k)|Q_i|$ and define k_j similarly. If $|Q_i| \geq km(i, k)/(40w(P))$ then since P is C -bonsai and $w(P)/k \geq 1/4$,

$$w(Q_i) \geq k_i + 2C\sqrt{500k_i \log \log k_i} \geq k_i + 2C\sqrt{m(i, k) \log \log m(i, k)}.$$

(The constant 500 ensures that the second inequality holds.) Similarly, if $|Q_j| \geq (m(j, k)/40)k/w(P)$, then

$$w(Q_j) \geq k_j + 2C\sqrt{500k_j \log \log k_j} \geq k_j + 2C\sqrt{m(j, k) \log \log m(j, k)}.$$

If both these conditions occur then by the definitions of k_i and k_j we have

$$w(Q_i) + w(Q_j) \geq \frac{w(P)}{k}(|Q_i| + |Q_j|) + 2C\sqrt{s(i, j, k) \log \log s(i, j, k)}.$$

Since $w(P) \leq k$, combining the preceding equation with the lower bound on $w(P \cap Q)$ given by Lemma 13 completes the proof in this case.

On the other hand, if $|Q_j| \leq km(j, k)/(40w(P)) \leq m(j, k)/10$ then $|Q_i| \geq m(i, k)$ and

$$|Q_i| \geq (k - (j - i)) - \frac{m(j, k)}{10} \geq \frac{9(k - (j - i))}{10} \geq \frac{9m(j, k)}{10}.$$

Since P is C -bonsai we have $w(Q_i) \geq k_i + 2C\sqrt{500k_i \log \log k_i}$. Moreover, in this case, $k_i \geq (w(P)/k)(k - (j - i))$ and $k_i \geq m(j, k)/4 = s(i, j, k)/4$, the preceding equation implies that

$$w(Q_i) \geq \frac{w(P)}{k}(k - (j - i)) + 2C\sqrt{s(i, j, k) \log \log s(i, j, k)},$$

which, combined with Lemma 13 and the fact that $w(P) \leq k$, completes the proof in this case. The case $|Q_i| \leq (m(i, k)/40)(k/w(P))$ is similar to but easier than the case $|Q_j| \leq (m(i, k)/40)(k/w(P))$. \square

We are now in a position to fully define the special set of paths we wish to consider. Given $\epsilon > 0$ and a positive integer k , let $\mathcal{U}_{\epsilon,k,C} = \mathcal{U}_{\epsilon,k,C}(n)$ be the set of paths P of length k in K_n for which the following events hold

$$\begin{aligned} A_1 &:= \{(1 - 2\epsilon) \log n \leq w(P) \leq (1 - \epsilon) \log n\}, \\ A_2 &:= \{P \in \mathcal{P}^*\} \\ A_3 &:= \{P \in \mathcal{L}_C \cap \mathcal{B}_C\}. \end{aligned} \tag{26}$$

We remark that if $\epsilon > 0$ is small enough (in particular, any $\epsilon = o(1)$ suffices), $C > 0$ is fixed, $t = t(n) = o(\sqrt{\log n})$, and $P \in \mathcal{U}_{\epsilon, \lceil k_\epsilon + t \rceil, C}$ (where k_ϵ was defined in (10)) then for n large enough, $w(P) \leq k \leq 4w(P)$, so, by Lemma 14, P is a local optimum. The following lemma bounds the expected size of $\mathcal{U}_{\epsilon,k,C}$.

Lemma 15. *There exist $C > 0$ and $\delta > 0$ such that for any fixed $c > 2$ and any ϵ with $\delta > \epsilon > c \log \log n / \log n$, for any function $t(n) = o(\sqrt{\log n})$, and for all n sufficiently large*

$$\mathbf{E}|\mathcal{U}_{\epsilon, \lceil k_\epsilon + t \rceil, C}| \geq \delta \mathbf{E}|\mathcal{P}_{\lceil k_\epsilon + t \rceil, \epsilon}|.$$

Proof. Recall that $h(T)$ denotes the (unweighted) height of a rooted tree T . Fix $\epsilon > 0$ as above, and a path P of length $k = \lceil k_\epsilon + t \rceil$. Define the following conditional probability

$$\mathbf{P}_\epsilon \{ \cdot \} := \mathbf{P} \{ \cdot \mid (1 - 2\epsilon) \log n \leq w(P) \leq (1 - \epsilon) \log n \}.$$

Then, by the definition (26), we have

$$\begin{aligned} \mathbf{P} \{ A_1, A_2, A_3 \} &= \mathbf{P}_\epsilon \{ A_2, A_3 \} \cdot \mathbf{P} \{ (1 - 2\epsilon) \log n \leq w(P) \leq (1 - \epsilon) \log n \} \\ &= \mathbf{P}_\epsilon \{ A_2, A_3 \} \cdot (1 + o(1)) \mathbf{P} \{ w(P) \leq (1 - \epsilon) \log n \}, \end{aligned}$$

by (3). So by linearity of expectation, to prove the lemma it suffices to show that $\mathbf{P}_\epsilon \{ A_2, A_3 \}$ is bounded away from zero. To do so we first write

$$\begin{aligned} \mathbf{P}_\epsilon \{ A_2, A_3 \} &= \mathbf{P}_\epsilon \{ P \in \mathcal{P}^*, P \in \mathcal{L}_C, P \in \mathcal{B}_C \} \\ &\geq \mathbf{P}_\epsilon \{ P \in \mathcal{B}_C, P \in \mathcal{L}_C \} - \mathbf{P}_\epsilon \{ P \notin \mathcal{P}^* \}. \end{aligned}$$

For ϵ as above, $\mathbf{P}_\epsilon \{ P \notin \mathcal{P}^* \} \rightarrow 0$ as $n \rightarrow \infty$ by Corollary 5, so we need only prove that $\mathbf{P}_\epsilon \{ P \in \mathcal{B}_C, P \in \mathcal{L}_C \}$ is bounded away from zero. To do so, rather than considering the set \mathcal{B}_C we consider a slightly different collection, which we denote \mathcal{B}_C^ϵ . We let $P \in \mathcal{B}_C^\epsilon$ if (\star) page 14 holds with the term $w(P)$ replaced by $(1 - \epsilon) \log n$. Given that $w(P) \leq (1 - \epsilon) \log n$, we have $\mathcal{B}_C^\epsilon \subseteq \mathcal{B}_C$, so $\mathbf{P}_\epsilon \{ P \in \mathcal{B}_C, P \in \mathcal{L}_C \} \geq \mathbf{P}_\epsilon \{ P \in \mathcal{B}_C^\epsilon, P \in \mathcal{L}_C \}$. But the event $P \in \mathcal{B}_C^\epsilon$ is independent from the events $P \in \mathcal{L}_C$ and

$$A_1 = \{(1 - 2\epsilon) \log n \leq w(P) \leq (1 - \epsilon) \log n\},$$

since they are determined by disjoint sets of edges. Thus, we have

$$\mathbf{P}_\epsilon \{ P \in \mathcal{B}_C, P \in \mathcal{L}_C \} \geq \mathbf{P} \{ P \in \mathcal{B}_C^\epsilon \} \mathbf{P}_\epsilon \{ P \in \mathcal{L}_C \}.$$

Now choose $C = C'(1/2)$, where $C'(1/2)$ is as in Corollary 12. By that corollary we then have $\mathbf{P}_\epsilon \{P \in \mathcal{L}_C\} \geq 1/2$, so

$$\mathbf{P}_\epsilon \{P \in \mathcal{B}_C, P \in \mathcal{L}_C\} \geq \frac{1}{2} \mathbf{P} \{P \in \mathcal{B}_C^\epsilon\}.$$

To complete the proof, it thus suffices to prove that $\mathbf{P} \{P \in \mathcal{B}_C^\epsilon\}$ is bounded below by a constant.

We assume ϵ is small enough and n is large enough (say $n \geq n_0$ for some n_0) that

$$\frac{(1-\epsilon) \log n}{\lceil k_\epsilon + t \rceil} \leq \frac{1}{3.5}. \quad (27)$$

(Note that this is possible by the value of α^* . We also remark that the preceding equation is the only place we use the hypothesis that ϵ is “small”.) To further simplify our lives, we now consider a subset of \mathcal{B}_C^ϵ that we denote $\bigcap_{i=1}^{k+1} \mathcal{B}_{C,i}$, where $\mathcal{B}_{C,i}$ is the set of all paths $Q = (v_1, \dots, v_{k+1})$ in K_n for which, for all integers ℓ with $\ell \geq m(i, k)/40$ (recall that $m(i, k)$ is defined in (24)), in K_n^Q

$$h \left(\text{SPT}_{v_i} \left(\ell + 2C \sqrt{500\ell \log \log \ell} \right) \right) < 3.15\ell. \quad (28)$$

The set $\bigcap_{i=1}^{k+1} \mathcal{B}_{C,i}$ is indeed a subset of \mathcal{B}_C^ϵ , since $9k/(10(1-\epsilon) \log n) \geq 3.15$ by (27). The events $P \in \mathcal{B}_{C,i}$, $i = 1, \dots, k+1$ are all increasing in the weight values, so by the FKG inequality [1, 4, 5, 6], we have

$$\mathbf{P} \{P \in \mathcal{B}_C^\epsilon\} \geq \prod_{i=1}^{k+1} \mathbf{P} \{P \in \mathcal{B}_{C,i}\} \geq \prod_{i=1}^{\lfloor (k+1)/2 \rfloor} \mathbf{P} \{P \in \mathcal{B}_{C,i}\}^2, \quad (29)$$

where the second inequality takes advantage of the symmetry. For small ℓ , the probability that (28) occurs is bounded away from zero: it is at least the probability that all $n-1$ edges leaving v_i have weight at least $\ell + 2C\sqrt{500\ell \log \log \ell}$, and every edge weight has an exponential distribution with mean n . Thus, choose an integer ℓ_0 large enough that the following four conditions are satisfied for $\ell \geq \ell_0$,

$$\begin{cases} \ell + 2C\sqrt{240\ell \log \log \ell} \leq 1.05\ell, \\ 1 - e^{-\ell/10}/2 \geq \exp(-e^{-\ell/10}) \end{cases} \quad \text{and} \quad \begin{cases} e^{3-\ell/8} \leq e^{-\ell/10}/4 \\ 3e^{\ell/40} \exp(-e^{-\ell/20}) \leq e^{-\ell/10}/4. \end{cases}$$

Then there is $\alpha_0 > 0$ such that for $\ell \leq \ell_0$, the event in (28) occurs with probability at least α_0 . Furthermore, for each ℓ the event in (28) is increasing in the weight values, so the FKG inequality yields

$$\begin{aligned} \mathbf{P} \{P \in \mathcal{B}_{C,i}\} &\geq \prod_{\ell=\lceil m(i,k)/40 \rceil}^{\infty} \mathbf{P} \left\{ h \left(\text{SPT}_i \left(\ell + 2C\sqrt{240\ell \log \log \ell} \right) \right) < 3.15\ell \right\} \\ &\geq \alpha_0^{\xi_i} \cdot \prod_{\ell \geq \ell_0 \vee m(i,k)/40} \mathbf{P} \{h(\text{SPT}_i(1.05\ell)) < 3.15\ell\}, \end{aligned} \quad (30)$$

where we have written $\xi_i = \max\{\ell_0 - m(i, k)/40, 0\}$. We now turn to bounding these remaining probabilities for $\ell \geq \ell_0$. By Lemma 2, and our choice for ℓ_0 ,

$$\mathbf{P} \{ |\text{SPT}_i(1.05\ell)| \geq e^{1.1\ell} \} \leq 3e^{\ell/40} \exp(-e^{\ell/40}) \leq \frac{e^{-\ell/10}}{4}.$$

Furthermore, with the following values and our choice for ℓ_0 , the bound of Theorem 3 yields

$$\mathbf{P} \{ h(\text{SPT}_i(1.05\ell)) \geq 3.15\ell \mid |\text{SPT}_i(1.05\ell)| < e^{1.1\ell} \} \leq \frac{e^{-\ell/10}}{4}.$$

It follows from the two preceding equations that for $\ell \geq \ell_0$

$$\mathbf{P} \{ h(\text{SPT}_i(1.05\ell)) < 3.15\ell \} \geq \left(1 - \frac{e^{-\ell/10}}{4}\right)^2 \geq 1 - \frac{e^{-\ell/10}}{2} \geq e^{-e^{-\ell/10}}.$$

Combining this inequality with (30) yields

$$\mathbf{P} \{ P \in \mathcal{B}_{C,i} \} \geq \alpha_0^{\xi_i} \exp\left(-\sum_{\ell=\lceil m(i,k)/40 \rceil}^{\infty} e^{-\ell/10}\right) \geq \alpha_0^{\xi_i} \exp(-26e^{-m(i,k)/400}).$$

Observe that the multiplicative factor $\alpha_0^{\xi_i} = 1$ but for the values of i such that $m(i, k) \leq 40\ell_0$. Observe also that for $i \leq (k+1)/2$, we have by definition (24) $m(i, k) = i - 1$. So, together with (29) the preceding equation yields

$$\mathbf{P} \{ P \in \mathcal{B}_C^c \} \geq \alpha_0^{80\ell_0^2} \exp\left(-52 \sum_{i=1}^{\lfloor (k+1)/2 \rfloor} e^{-(i-1)/400}\right),$$

which completes the proof, since the series in the right-hand side converges. \square

With Lemma 15 under our belt, we can now complete our proof of Theorem 1

Proof of Theorem 1, lower bounds. Let C be as in Lemma 15, choose $\epsilon = \epsilon(n) = 3 \log \log n / \log n$, and choose a function $t = t(n) = o(\sqrt{\log n})$ such that $t(n) \rightarrow -\infty$ with n ; we will make a more precise choice for $t(n)$ shortly. As noted just before the statement of Lemma 15, if $\epsilon = o(1)$, $C > 0$ is fixed, $t = o(\sqrt{\log n})$, and $P \in \mathcal{U}_{\epsilon, \lceil k_\epsilon + t \rceil, C}$ then for n large enough, $w(P) \leq k \leq 4w(P)$, so, by Lemma 14, P is a local optimum. Our choices of ϵ, C , and t satisfy these conditions, so for n large enough we have

$$\mathcal{U}_{\lceil k_\epsilon + t \rceil, \epsilon, C} \subseteq \mathcal{O}_{\lceil k_\epsilon + t \rceil, \epsilon} \subseteq \mathcal{P}_{\lceil k_\epsilon + t \rceil, \epsilon}.$$

By Lemmas 6 and 15, and the definition of k_ϵ in (10), there is $K > 0$ such that choosing $t(n) = K \log \log n$, for n large enough we have

$$\mathbf{E}|\mathcal{U}_{\lceil k_\epsilon + t \rceil, \epsilon}| \geq \log n.$$

Recalling the definition (11) of $\Delta(\mathcal{U}_{\lceil k_\epsilon + t \rceil, \epsilon})$, we remark that since $\mathcal{U}_{\lceil k_\epsilon + t \rceil, \epsilon, C} \subseteq \mathcal{O}_{\lceil k_\epsilon + t \rceil, \epsilon}$, the contribution of pairs of paths who intersect once to $\Delta(\mathcal{U}_{\lceil k_\epsilon + t \rceil, \epsilon})$, is zero. Also, since $\mathcal{U}_{\lceil k_\epsilon + t \rceil, \epsilon, C} \subseteq \mathcal{P}_{\lceil k_\epsilon + t \rceil, \epsilon}$, we have

$$\Delta(\mathcal{U}_{\lceil k_\epsilon + t \rceil, \epsilon}) \leq \sum_{2 \leq j \leq i < \lceil k_\epsilon + t \rceil} \Delta_{t,i,j},$$

so by Lemma 9, for n large enough, we have $\Delta(\mathcal{U}_{\lceil k_\epsilon + t \rceil, \epsilon}) \leq n^{-0.95} \mathbf{E} [|\mathcal{P}_{\lceil k_\epsilon + t \rceil, \epsilon}|]^2$. By our choice of ϵ there is $K' > 0$ such that $k_\epsilon + t \geq \alpha^* \log n - K' \log \log n$. It follows by the second moment method (12) and Lemma 15 that

$$\begin{aligned} \mathbf{P} \left\{ \max_{1 \leq i, j \leq n} |P_{ij}| \leq \alpha^* \log n - K' \log \log n \right\} &\leq \mathbf{P} \{ |\mathcal{U}_{\lceil k_\epsilon + t \rceil, \epsilon}| = 0 \} \\ &< \frac{1}{\mathbf{E} |\mathcal{U}_{\lceil k_\epsilon + t \rceil, \epsilon}|} + \frac{\Delta(\mathcal{U}_{\lceil k_\epsilon + t \rceil, \epsilon})}{\mathbf{E} [|\mathcal{U}_{\lceil k_\epsilon + t \rceil, \epsilon}|]^2} \\ &\leq \frac{1}{\log n} + \frac{n^{-0.95}}{\delta^2} \xrightarrow{n \rightarrow \infty} 0, \end{aligned}$$

proving the lower bound in probability. We then have from the previous bound

$$\mathbf{E} \left[\max_{1 \leq i, j \leq n} |P_{ij}| \right] \geq (\alpha^* \log n - K' \log \log n) \left(1 - \frac{1}{\log n} - \frac{n^{-0.95}}{\delta^2} \right),$$

which proves the lower bound in expectation. \square

5 Conclusion

We have shown that the longest minimum-weight path in a complete graph with i.i.d. exponentially distributed edge weights is asymptotically $\alpha^* \log n$ in probability, where $\alpha^* \approx 3.5911$ satisfies $\alpha^* \log \alpha^* = \alpha^* + 1$. The difficult part of the proof is the lower bound as the upper bound follows easily from Devroye's [3] results on the height of a random recursive tree and a simple union bound. It may come as a surprise that such a simple union bound leads to correct bounds, since it means that the n shortest path trees SPT_i essentially behave as if they were independent.

The lower bound is proved by a second-moment argument for a carefully chosen random set of minimum-weight paths, all of length about $\alpha^* \log n$, that satisfy certain regularity conditions. This careful choice allows us to control the variance of the size of the set of such paths. The proof shows that, with high probability, there exist minimum-weight paths of length about $\alpha^* \log n$ whose total weight is close to $\log n$. Thus, these extra-long paths have a total weight just like a typical minimum-weight path.

Finally, just as Janson [8] did in proving his results discussed in the introduction, we can use a standard coupling argument to extend Theorem 1 to distributions other than exponential. To explain the coupling, it is useful to divide the edge weights by n , so to consider the $\text{Exp}(1)$ edge weights $X'_e = X_e/n$.

Let Z be any non-negative random variable, with distribution function $F(t) = \mathbf{P}\{Z \leq t\}$, and let $F^{-1}(t) = \sup\{x : F(x) \leq t\}$. Let U be a random variable uniform on $[0, 1]$. Then $F^{-1}(U)$ is distributed as Z . We may thus couple the exponential edge weights X'_e to uniform edge weights U_e by setting $X'_e = F_E^{-1}(U_e)$, where $F_E(t)$ is the distribution function of an $\text{Exp}(1)$ random variable. We have $F_E(t) = 1 - e^{-t} = t - t^2/2 + O(t^3)$ as $t \rightarrow 0$, and it follows that for all edges e of K_n , $|X'_e - U_e| = O((X'_e)^2)$, uniformly when $X'_e \leq 1/2$.

Of particular use for us is the following consequence: for all edges e with $X'_e \leq 12 \log n/n$ (say), we have $|X'_e - U_e| = O(\log^2 n/n^2)$, so for any path P with $O(\log n)$ edges and with $w'(P) := \sum_{e \in E(P)} X'_e \leq 12 \log n/n$, we have

$$\left| w'(P) - \sum_{e \in E(P)} U_e \right| = O\left(\frac{\log^3}{n^2}\right). \quad (31)$$

It then follows easily that the conclusions of Theorem 1 hold for uniform $[0, 1]$, as well as exponential, edge weights. (Technically, one must redo all the arguments of the paper keeping (31) in mind, and accordingly adjust some of the constants very slightly; we omit this step.)

In fact all we really need is a bound of $o(1/n)$ in (31) in order to extend Theorem 1 to uniform $[0, 1]$ edge weights. Thus, by another coupling as above, we may extend Theorem 1 from uniforms to any other edge weights Y whose distribution function F_Y satisfies $F_Y(t) = ct + o(t \log^2 t)$ for some constant $c > 0$, as $t \rightarrow 0$.

References

- [1] N. Alon, J. Spencer, and P. Erdős. *The Probabilistic Method*. Wiley, New York, NY, second edition, 2000.
- [2] A. Dembo and O. Zeitouni. *Large Deviation Techniques and Applications*. Springer Verlag, second edition, 1998.
- [3] L. Devroye. Branching processes in the analysis of the heights of trees. *Acta Informatica*, 24:277–298, 1987.
- [4] C.M. Fortuin, P.W. Kasteleyn, and J. Ginibre. Correlational inequalities for partially ordered sets. *Communications of Mathematical Physics*, 22: 89–103, 1971.
- [5] G. R. Grimmett. *Percolation*, volume 321 of *A Series of Comprehensive Studies in Mathematics*. Springer Verlag, second edition, 1999.
- [6] T.E. Harris. A lower bound for the critical probability in a certain percolation process. *Proceedings of the Cambridge Philosophical Society*, 56:13–20, 1960.
- [7] G. Hooghiemstra and P. Van Mieghem. The weight and hopcount of the shortest path in the complete graph with exponential weights. *Combinatorics, Probability and Computing*, 2008.

- [8] S. Janson. One, two and three times $\log n/n$ for paths in a complete graph with random weights. *Combinatorics, Probability and Computing*, 8(1999): 347–361, 1999.
- [9] B. Pittel. Note on the height of recursive trees and m -ary search trees. *Random Structures and Algorithms*, 5:337–347, 1994.
- [10] H. Robbins. A remark on Stirling’s formula. *American Mathematical Monthly*, 62:26–29, 1955.
- [11] L.C.G. Rogers and D. Williams. *Diffusions, Markov processes, and martingales*, volume 1. Cambridge University Press, 2000.
- [12] Galen R. Shorack and Jon A. Wellner. *Empirical Processes with Applications to Statistics*. John Wiley & Sons, Inc, 1986.
- [13] R. T. Smythe and H. M. Mahmoud. A survey of recursive trees. *Theoretical Probability and Mathematical Statistics*, 51:1–27, 1995.
- [14] R. van der Hofstad, G. Hooghiemstra, and P. van Mieghem. Size and weight of shortest path trees with exponential link weights. *Combinatorics, Probability and Computing*, 15(06):903–926, 2006.
- [15] R. van der Hofstad, G. Hooghiemstra, and P. Van Mieghem. The weight of the shortest path tree. *Random Structures & Algorithms*, 30(3):359–379, 2007.