

FINDING MINIMUM SPANNING TREES VIA LOCAL IMPROVEMENTS

LOUIGI ADDARIO-BERRY, JORDAN BARRETT, AND BENOÎT CORSINI

ABSTRACT. We consider a family of local search algorithms for the minimum-weight spanning tree, indexed by a parameter ρ . One step of the local search corresponds to replacing a connected induced subgraph of the current candidate graph whose total weight is at most ρ by the minimum spanning tree (MST) on the same vertex set. Fix a non-negative random variable X , and consider this local search problem on the complete graph K_n with independent X -distributed edge weights. Under rather weak conditions on the distribution of X , we determine a threshold value ρ^* such that the following holds. If the starting graph (the “initial candidate MST”) is independent of the edge weights, then if $\rho > \rho^*$ local search can construct the MST with high probability (tending to 1 as $n \rightarrow \infty$), whereas if $\rho < \rho^*$ it cannot with high probability.

1. INTRODUCTION

Local search is the name for an optimization paradigm in which optimal or near-optimal solutions are sought algorithmically, via sequential improvements which are “local” in that at each step, the search space consists only of neighbours (in some sense) of the current solution. Well-known algorithmic examples of this paradigm include simulated annealing, hill climbing, and the Metropolis-Hasting algorithm.

A recent line of research considers the behaviour of local search on *smoothed* optimization problems, in which the input is either fully random or is a random perturbation of a fixed input. The goal in this setting is to characterize the running time of local search and the quality of its output. Problems approached in this vein include *max-cut* [3, 4, 6], for which the allowed “local” improvements consist of moving a single vertex; *max-2CSP* and the *binary function optimization problem* [5], for which the allowed local improvements are bit flips; and Euclidean TSP [8], where the allowed local improvements consist of replacing edge pairs uv, wx with pairs uw, vx (when the result is still a tour).

In the current work, we analyze local search for the *random minimum spanning tree* problem, one of the first and foundational problems in combinatorial optimization. We now briefly describe our results (for more precise statements see Section 1.1, below). As input to the problem, we take the randomly-weighted complete graph $\mathbb{K}_n = (K_n, \mathbb{X})$, where $\mathbb{X} = (X_e, e \in E(K_n))$ are independent copies of a random variable X , and an arbitrary starting graph H_0 , which we aim to transform into the minimum-weight spanning tree MST. We fix a *threshold weight* $\rho > 0$; at step $k \geq 0$, a local improvement consists of choosing a connected induced subgraph of the current MST candidate H_k whose current total weight is at most ρ , and replacing it by the minimum weight spanning tree on the same vertex set.

Suppose that X is non-negative and has a density $f : [0, \infty) \rightarrow [0, \infty)$ which is continuous at 0 and satisfies $f(0) > 0$. Then writing $\rho^* = \sup\{x : \mathbb{P}(X > x) > 0\}$, we prove that if $\rho > \rho^*$ then there exist local search paths which output the MST, whereas if $\rho < \rho^*$ then local search cannot reach the MST (and, indeed, with high probability will only achieve an approximation ratio of order $\Theta(n)$).

1.1. Detailed statement of the results. Let $\mathbb{G} = (G, w) = (V, E, w)$ be a finite weighted connected graph, where $G = (V, E)$ is a graph and $w : E \rightarrow (0, \infty)$ are edge weights; set $V(\mathbb{G}) = V(G) = V$ and $E(\mathbb{G}) = E(G) = E$. For a subgraph H of G write $w(H) = \sum_{e \in E(H)} w(e)$ for its weight. A *minimum spanning tree* (MST) of \mathbb{G} is a spanning tree T of G which minimizes $w(T)$ among all spanning trees of

Date: May 10, 2022.

2020 Mathematics Subject Classification. 05C80, 60C05, 05C85.

Key words and phrases. Minimum spanning trees, local search, random graphs.

G . There is a unique MST provided all edge weights are distinct; we hereafter restrict our attention to weighted graphs \mathbb{G} where all edge weights are distinct (and more strongly where the edge weights are linearly independent over \mathbb{R}); we call such graphs *generic*. For a generic weighted graph \mathbb{G} , we write $\text{MST}(\mathbb{G})$ for the unique minimum spanning tree of \mathbb{G} .

For a weighted graph $\mathbb{G} = (V, E, w)$ and a set $S \subset V$, write $G[S]$ for the induced subgraph $G[S] = (S, E|_{S \times S})$ and $\mathbb{G}[S]$ for the induced weighted subgraph $\mathbb{G}[S] = (G[S], w|_{E(G[S])})$. Now, given a spanning subgraph H of G , define $\Phi(H, S) = \Phi_{\mathbb{G}}(H, S)$ as follows. If $H[S]$ is connected then let $\Phi(H, S)$ be the spanning subgraph with edge set $(E(H) \setminus E(H[S])) \cup E(\text{MST}(\mathbb{G}[S]))$; if $H[S]$ is not connected then let $\Phi(H, S) = H$. In words, to form $\Phi(H, S)$ from H , we replace $H[S]$ by the minimum-weight spanning tree of $\mathbb{G}[S]$, unless $H[S]$ is not connected.

Now suppose we are given a finite weighted connected graph $\mathbb{G} = (V, E, w)$, a spanning subgraph H of G , and a sequence $\mathbb{S} = (S_i, 1 \leq i \leq m)$ of subsets of V . Define a sequence of spanning subgraphs $(H_i, 0 \leq i \leq m)$ as follows. Set $H_0 = H$, and for $1 \leq i \leq m$ let $H_i = \Phi_{\mathbb{G}}(H_{i-1}, S_i)$. Using the previous definition of Φ , this simply corresponds to sequentially replacing the subgraph of H_{i-1} on S_i by its corresponding minimum spanning tree (assuming H_{i-1} is connected). We refer to \mathbb{S} as an *optimizing sequence* for the pair (\mathbb{G}, H) , and call $(H_i, 0 \leq i \leq m)$ the *subgraph sequence corresponding to* \mathbb{S} . We say \mathbb{S} is an *MST sequence* for (\mathbb{G}, H) if the final spanning subgraph H_m is the MST of \mathbb{G} .

The weight of step i of the sequence \mathbb{S} is defined as

$$\text{wt}(\mathbb{S}, i) = \text{wt}(\mathbb{G}, H, \mathbb{S}, i) := w(H_{i-1}[S_i]) = \sum_{e \in E(H_{i-1}[S_i])} w(e),$$

and the weight of the whole sequence is the maximal weight of a single step:

$$\text{wt}(\mathbb{S}) = \text{wt}(\mathbb{G}, H, \mathbb{S}) := \max \left\{ \text{wt}(\mathbb{S}, i) : 1 \leq i \leq m \right\}.$$

The *cost* of the pair (\mathbb{G}, H) is defined as

$$\text{cost}(\mathbb{G}, H) := \min \left\{ \text{wt}(\mathbb{S}) : \mathbb{S} \text{ is an MST sequence for } (\mathbb{G}, H) \right\}.$$

The following theorem is the main result of the current work. Write K_n for the complete graph with vertex set $[n] = \{1, \dots, n\}$, and $\mathbb{K}_n = (K_n, \mathbb{X})$ for the randomly weighted complete graph, where $\mathbb{X} = (X_e, e \in E(K_n))$ are independent $\text{Uniform}[0, 1]$ random variables. If $\mathbb{S} = (S_1, \dots, S_m)$ is an optimizing sequence for (\mathbb{K}_n, H_n) then we write $H_{n,0} = H_n$ and $H_{n,i} = \Phi_{\mathbb{K}_n}(H_{n,i-1}, S_i)$ for $1 \leq i \leq m$. Finally, we say a sequence $(E_n, n \geq 1)$ of events occurs with high probability if $\mathbb{P}(E_n) \rightarrow 1$ as $n \rightarrow \infty$.

Theorem 1.1. *Fix any sequence $(H_n, n \geq 1)$ of connected graphs with H_n being a spanning subgraph of K_n . Then for any $\varepsilon > 0$, as $n \rightarrow \infty$,*

- (a) *with high probability there exists an MST sequence \mathbb{S} for (\mathbb{K}_n, H_n) with $\text{wt}(\mathbb{S}) \leq 1 + \varepsilon$, and*
- (b) *there exists $\delta > 0$ such that with high probability, given any optimizing sequence $\mathbb{S} = (S_1, \dots, S_m)$ for (\mathbb{K}_n, H_n) with $\text{wt}(\mathbb{S}) \leq 1 - \varepsilon$, the final spanning subgraph $H_{n,m}$ has weight $w(H_{n,m}) \geq \delta n w(\text{MST}(\mathbb{K}_n))$.*

In particular, $\text{cost}(\mathbb{K}_n, H_n) \xrightarrow{\mathbb{P}} 1$ as $n \rightarrow \infty$.

We discuss possible refinements of and extensions to Theorem 1.1 in the conclusion, Section 4. We also explain in that section how to extend Theorem 1.1 to more general edge weight distributions than $\text{Uniform}[0, 1]$, as described just before Section 1.1.

1.2. Overview of the proof. In this section, we give an overview of the proof of Theorem 1.1, while postponing the proofs of the more technical aspects to Sections 2 and 3 and Appendix A. The lower bound of Theorem 1.1 is straightforward, so we provide it in full detail immediately.

Lower bound of Theorem 1.1. Fix $\varepsilon > 0$, and let $E_{n,\varepsilon} = \{e \in E(H_n) : X_e > 1 - \varepsilon\}$. The set $E_{n,\varepsilon}$ is a binomial random subset of $E(H_n)$ in which each edge is present with probability ε , so $\mathbb{P}(|E_{n,\varepsilon}| \geq \varepsilon n/2) \rightarrow 1$.

Note that, for any edge $e = uv \in E(H_n) \setminus E(\text{MST}(\mathbb{K}_n))$, and any optimizing sequence $\mathbb{S} = (S_1, \dots, S_m)$ for (\mathbb{K}_n, H_n) , if there is no set S_i with $u, v \in S_i$, then $e \in H_{n,m}$. It follows that for any optimizing sequence \mathbb{S} with $\text{wt}(\mathbb{S}) \leq 1 - \varepsilon$, the final spanning subgraph $H_{n,m}$ has $E_{n,\varepsilon} \subset E(H_{n,m})$ and so on the event that $|E_{n,\varepsilon}| \geq \varepsilon n/2$ we have

$$w(H_{n,m}) \geq n(1 - \varepsilon)\varepsilon/2.$$

To conclude, we use that $w(\text{MST}(\mathbb{K}_n)) \rightarrow \zeta(3)$ in probability [7]. It follows that with probability tending to 1, both $|E_{n,\varepsilon}| \geq \varepsilon n/2$ and $w(\text{MST}(\mathbb{K}_n)) \leq 2\zeta(3)$, and when both these events occur we have

$$w(H_{n,m}) \geq n(1 - \varepsilon)\varepsilon/2 \geq w(\text{MST}(\mathbb{K}_n)) \cdot n(1 - \varepsilon)\varepsilon/(4\zeta(3)).$$

Since this holds for any optimizing sequence with weight at most $1 - \varepsilon$, the result follows by taking $\delta = (1 - \varepsilon)\varepsilon/(4\zeta(3))$.

Upper bound of Theorem 1.1. We now turn to the key ideas underlying our proof of the upper bound. We begin with a deterministic fact.

Fact 1.2. *Any connected graph H with vertex set $[n]$ contains an induced subgraph with at least $\frac{1}{2}\sqrt{\log_2 n}$ vertices which is either a clique, a star, or a path.*

We prove the fact immediately since the proof is very short; but its proof can be skipped without consequence for the reader's understanding of what follows.

Proof of Fact 1.2. The result is trivial if $n \leq 16$ so assume $n > 16$. Let $m = n^{1/\sqrt{\log_2 n}} \geq 4$. If H has maximum degree less than m then it has diameter at least $\sqrt{\log_2 n} - 1 \geq \frac{1}{2}\sqrt{\log_2 n}$ so it contains a path of length at least $\frac{1}{2}\sqrt{\log_2 n}$. On the other hand, if H has maximum degree at least m then let v be a vertex of H with degree at least m and let N_v be the set of neighbours of v in H . By Ramsey's theorem, and more concretely the diagonal Ramsey upper bound $R(k, k) < 4^k$, the graph $H[N_v]$ contains a set S of size at least

$$\frac{1}{2} \log_2 m = \frac{1}{2} \frac{\log_2 n}{\sqrt{\log_2 n}} = \frac{1}{2} \sqrt{\log_2 n}$$

such that $H[S]$ is either a clique or an independent set. If $H[S]$ is a clique then we are done, and if $H[S]$ is an independent set then $H[S \cup \{v\}]$ is a star of size $|S| + 1$ so we are again done. \square

Fact 1.2 proves to be useful together with the following special case of the upper bound of Theorem 1.1, whose proof appears in Section 3.

Proposition 1.3. *Fix a sequence $(H_n, n \geq 1)$ of connected graphs such that, for all n , H_n is either a clique, a star, or a path with $V(H_n) = [n]$. Then for all $\varepsilon > 0$, with high probability $\text{cost}(\mathbb{K}_n, H_n) \leq 1 + \varepsilon$.*

We combine Proposition 1.3 with Fact 1.2 as follows. First, choose $V_n \subset [n]$ with $|V_n| \geq \frac{1}{2}\sqrt{\log n}$ such that $H_n[V_n]$ is a clique, a star or a path, and consider $\mathbb{K}_n[V_n]$, the restriction of the weighted complete graph \mathbb{K}_n to V_n . Let $\mathbb{S}'_n = (S'_0, \dots, S'_m)$ be an MST sequence for $(\mathbb{K}_n[V_n], H_n[V_n])$ of minimum cost. Now consider using the sequence \mathbb{S}'_n as an optimizing sequence for (\mathbb{K}_n, H_n) . In other words, we set $H_{n,i} = \Phi_{\mathbb{K}_n}(H_{n,i-1}, S'_i)$ for $1 \leq i \leq m$. Then $H_{n,m} = \Phi_{\mathbb{K}_n}(H_0, V_n)$, which is to say that $H_{n,m}$ consists of H_n with $H_n[V_n]$ replaced by $\text{MST}(\mathbb{K}_n[V_n])$. Moreover, by Proposition 1.3, $\text{wt}(\mathbb{S}'_n) = \text{wt}(\mathbb{K}_n, H_n, \mathbb{S}'_n) = \text{wt}(\mathbb{K}_n[V_n], H_n[V_n], \mathbb{S}'_n) \xrightarrow{\mathbb{P}} 1$; so with high probability we have transformed a "large" (i.e. whose size is $\geq \frac{1}{2}\sqrt{\log n}$) subgraph of H_n into its minimum spanning tree, using an optimizing sequence of cost at most $1 + o_{\mathbb{P}}(1)$.

The next step is to apply a procedure we call the *eating algorithm*, described in Section 2. This algorithm allows us to bound the minimum cost of an MST sequence in terms of the weighted diameters of minimum spanning trees of a growing sequence of induced subgraphs of the input graph, with each

graph in the sequence containing one more vertex than its predecessor. In the setting of Theorem 1.1, it allows us to find an MST sequence with weight at most $1 + o_{\mathbb{P}}(1)$ provided that the starting graph already contains a large subgraph on which it is equal to the MST. The key result of our analysis of the eating algorithm is summarized in the following proposition.

Proposition 1.4. *Fix a sequence $(H_n, n \geq 1)$ of connected graphs with $V(H_n) = [n]$. Fix any sequence of sets $(V_n, n \geq 1)$ such that $V_n \subset [n]$, $|V_n| \rightarrow \infty$ as $n \rightarrow \infty$, and $H_n[V_n]$ is connected for all $n \geq 1$. Let $H'_n = \Phi_{\mathbb{K}_n}(H_n, V_n)$, so that $H'_n[V_n] = \text{MST}(\mathbb{K}_n[V_n])$. Then for all $\varepsilon > 0$, with high probability $\text{cost}(\mathbb{K}_n, H'_n) \leq 1 + \varepsilon$.*

The proof of Proposition 1.4 appears in Section 2. We are now prepared to prove Theorem 1.1, modulo the proofs of Proposition 1.3 and Proposition 1.4.

Proof of Theorem 1.1. We already established the lower bound of the theorem, so it remains to show that for all $\varepsilon > 0$,

$$\mathbb{P}(\text{cost}(\mathbb{K}_n, H_n) \leq 1 + \varepsilon) \rightarrow 1$$

as $n \rightarrow \infty$. For the remainder of the proof we fix $\varepsilon > 0$.

Using Fact 1.2, let V_n be a subset of $[n]$ with size at least $\frac{1}{2}\sqrt{\log_2 n}$ such that $H_n[V_n]$ is a clique, a star or a path. Write $\mathbb{K}_n^- = \mathbb{K}_n[V_n]$ and $H_n^- = H_n[V_n]$, and let \mathbb{S}_n^- be an MST sequence for (\mathbb{K}_n^-, H_n^-) of minimum cost. By Proposition 1.3,

$$\mathbb{P}(\text{wt}(\mathbb{K}_n^-, H_n^-, \mathbb{S}_n^-) \leq 1 + \varepsilon) \rightarrow 1$$

as $n \rightarrow \infty$. Moreover, we have $\text{wt}(\mathbb{K}_n^-, H_n^-, \mathbb{S}_n^-) = \text{wt}(\mathbb{K}_n, H_n, \mathbb{S}_n^-)$: the weight of the sequence \mathbb{S}_n^- is the same with respect to $(\mathbb{K}_n^-, H_n^-) = (\mathbb{K}_n[V_n], H_n[V_n])$ as it is with respect to (\mathbb{K}_n, H_n) ; this is easily seen by induction. It follows that

$$\mathbb{P}(\text{wt}(\mathbb{K}_n, H_n, \mathbb{S}_n^-) \leq 1 + \varepsilon) \rightarrow 1.$$

Next let $H'_n = \Phi_{\mathbb{K}_n}(H_n, V_n)$, so $H'_n[V_n] = \text{MST}(\mathbb{K}_n[V_n])$. Since \mathbb{S}_n^- is an MST sequence for (\mathbb{K}_n^-, H_n^-) , this is also the graph resulting from using \mathbb{S}_n^- as an optimizing sequence for (\mathbb{K}_n, H_n) . Now let \mathbb{S}'_n be an MST sequence for H'_n of minimum cost. Since $|V_n| \rightarrow \infty$ and $H_n[V_n]$ is connected, it follows from Proposition 1.4 that

$$\mathbb{P}(\text{wt}(\mathbb{K}_n, H'_n, \mathbb{S}'_n) \leq 1 + \varepsilon) \rightarrow 1.$$

To conclude, note that the concatenation \mathbb{S}_n of \mathbb{S}_n^- and \mathbb{S}'_n is an MST sequence for (\mathbb{K}_n, H_n) , and

$$\text{wt}(\mathbb{K}_n, H_n, \mathbb{S}_n) = \max \left\{ \text{wt}(\mathbb{K}_n, H_n, \mathbb{S}_n^-), \text{wt}(\mathbb{K}_n, H'_n, \mathbb{S}'_n) \right\},$$

so $\mathbb{P}(\text{wt}(\mathbb{K}_n, H_n, \mathbb{S}_n) \leq 1 + \varepsilon) \rightarrow 1$ and thus $\mathbb{P}(\text{cost}(\mathbb{K}_n, H_n) \leq 1 + \varepsilon) \rightarrow 1$, as required. \square

The remainder of the paper proceeds as follows. In Section 2 we describe the eating algorithm and prove Proposition 1.4, modulo the proof of a key technical input (Theorem 2.3), an upper tail bound on the weighted diameter of $\text{MST}(\mathbb{K}_n)$, which is postponed to Appendix A. In Section 3 we prove Proposition 1.3 by using the details of the eating algorithm to generate a well bounded sequence of increasing MSTs that are each built from a clique, a star, or a path. We conclude in Section 4 by presenting the generalization of Theorem 1.1 to other edge weight distributions, and by discussing avenues for future research.

2. THE EATING ALGORITHM

In this section, we prove Proposition 1.4. Informally, we prove this proposition by showing that we can efficiently add vertices to an MST of a large subgraph of K_n , one at a time, via an optimizing sequence which has a low weight, with high probability. For a weighted graph $\mathbb{G} = (V, E, w)$, write $\text{wdiam}(\mathbb{G})$ for the weighted diameter of \mathbb{G} ,

$$\text{wdiam}(\mathbb{G}) := \max \left\{ \text{dist}_{\mathbb{G}}(u, v) : u, v \in V \right\},$$

where

$$\text{dist}_{\mathbb{G}}(u, v) := \min \left\{ w(P) : P \text{ is a path from } u \text{ to } v \text{ in } \mathbb{G} \right\}.$$

It is sometimes convenient to write $\text{wdiam}(G)$ for an unweighted graph G , where the appropriate choice of weights is clear from context. Finally, we also introduce the unweighted diameter

$$\text{diam}(G) := \max \left\{ \min \left\{ |E(P)| : P \text{ is a path from } u \text{ to } v \text{ in } G \right\} : u, v \in V \right\},$$

which will be used later in this work (in Section 3.2 and in Appendix A).

The key tool to prove Proposition 1.4 is the following proposition, which will be applied recursively.

Proposition 2.1. *Let $\mathbb{G} = (V, E, w)$ be a generic weighted graph with $V = [n]$ and $\max\{w(e) : e \in E\} \leq 1$. Suppose that H is a spanning subgraph of G and $H[n-1] = \text{MST}(\mathbb{G}[n-1])$. Then*

$$\text{cost}(\mathbb{G}, H) \leq 1 + \max \left\{ \text{wdiam}(\text{MST}(\mathbb{G}[n-1])), \text{wdiam}(\text{MST}(\mathbb{G})) \right\}.$$

The proof of Proposition 2.1 occupies the bulk of Section 2; it appears below in Sections 2.1 and 2.2.

Corollary 2.2 (The eating algorithm). *Let $\mathbb{G} = (V, E, w)$ be a weighted graph with $V = [n]$ and $\max\{w(e) : e \in E\} \leq 1$. Let H be a spanning subgraph of G and fix a non-empty set $U \subset [n]$ for which $H[U] = \text{MST}(\mathbb{G}[U])$. Let $U = U_0 \subset U_1 \subset \dots \subset U_k = V$ be any increasing sequence of subsets of V such that, for all $0 \leq i < k$, $U_{i+1} \setminus U_i$ is a singleton and $H[U_i]$ is connected. Then*

$$\text{cost}(\mathbb{G}, H) \leq 1 + \max \left\{ \text{wdiam}(\text{MST}(\mathbb{G}[U_i])) : 0 \leq i \leq k \right\}.$$

Proof. Set $F_0 = H$ and let $\mathbb{S}_1, \dots, \mathbb{S}_k$ and F_1, \dots, F_k be constructed inductively as follows. Given F_{i-1} , let \mathbb{S}_i be an MST sequence of minimal weight for the pair $(\mathbb{G}[U_i], F_{i-1}[U_i])$ and let $F_i = \Phi_{\mathbb{G}}(F_{i-1}, U_i)$. Note that $F_i[U_i]$ is the last graph of the subgraph sequence corresponding to \mathbb{S}_i .

By using that an optimizing sequence on $(\mathbb{G}[U_i], F_{i-1}[U_i])$ can also be seen as an optimizing sequence on (\mathbb{G}, F_{i-1}) of identical weight, we can bound the weight of the global optimizing sequence \mathbb{S} obtained by concatenating $\mathbb{S}_1, \dots, \mathbb{S}_k$ in that order. Indeed, we have that

$$\text{wt}(\mathbb{G}, H, \mathbb{S}) = \max \left\{ \text{wt}(\mathbb{G}[U_i], F_{i-1}[U_i], \mathbb{S}_i) : 1 \leq i \leq k \right\}.$$

Moreover, by the definition of F_{i-1} , we know that $F_{i-1}[U_{i-1}] = \text{MST}(\mathbb{G}[U_{i-1}])$ and by minimality of \mathbb{S}_i along with Proposition 2.1, it follows that for all $1 \leq i \leq k$,

$$\text{wt}(\mathbb{G}[U_i], F_{i-1}[U_i], \mathbb{S}_i) \leq 1 + \max \left\{ \text{wdiam}(\text{MST}(\mathbb{G}[U_{i-1}])), \text{wdiam}(\text{MST}(\mathbb{G}[U_i])) \right\}.$$

Since $\text{cost}(\mathbb{G}, H) \leq \text{wt}(\mathbb{G}, H, \mathbb{S})$, combining the last two results provides us with the desired upper bound for $\text{cost}(\mathbb{G}, H)$. \square

The importance of this corollary becomes clear in light of the next theorem, which provides strong tail bounds on the diameter of MSTs of randomly-weighted complete graphs.

Theorem 2.3. *Let $\mathbb{K}_n = (K_n, \mathbb{X})$ be the complete graph with vertex set $[n]$, endowed with independent, Uniform $[0, 1]$ edge weights $\mathbb{X} = (X_e, e \in E(K_n))$. Then for all n sufficiently large,*

$$\mathbb{P} \left(\text{wdiam}(\text{MST}(\mathbb{K}_n)) \geq \frac{7 \log^4 n}{n^{1/10}} \right) \leq \frac{4}{n^{\log n}}.$$

In particular, $\text{wdiam}(\text{MST}(\mathbb{K}_n)) \xrightarrow{\mathbb{P}} 0$ as $n \rightarrow \infty$.

The proof of Theorem 2.3 is postponed to Appendix A. We now use Corollary 2.2 and Theorem 2.3 to prove Proposition 1.4.

Proof of Proposition 1.4. Consider any sequence of sets $(V_n, n \geq 1)$ with $V_n \subset [n]$ and $|V_n| \rightarrow \infty$ as $n \rightarrow \infty$ and such that $H_n[V_n]$ is connected for all $n \geq 1$, and let $H'_n = \Phi_{\mathbb{K}_n}(H_n, V_n)$. Since H'_n is connected, we may list the vertices of $[n] \setminus V_n$ as v_1, \dots, v_k so that for all $1 \leq i \leq k$, vertex v_i is adjacent

to an element of $V_n \cup \{v_1, \dots, v_{i-1}\}$. Taking $U_0 = V_n$ and $U_i = V_n \cup \{v_1, \dots, v_i\}$ for $1 \leq i \leq k$, the sequence U_0, \dots, U_k satisfies the conditions of Corollary 2.2 with $\mathbb{G} = \mathbb{K}_n$. It follows that

$$(2.1) \quad \text{cost}(\mathbb{K}_n, H'_n) \leq 1 + \max \left\{ \text{wdiam}(\text{MST}(\mathbb{K}_n[U_i])) : 0 \leq i \leq k \right\}.$$

Moreover, since $|V_n| \rightarrow \infty$ as $n \rightarrow \infty$, for n sufficiently large we may apply Theorem 2.3 to $\mathbb{K}_n[U_i]$ for each $0 \leq i \leq k$ and obtain that

$$\mathbb{P}\left(\exists i : \text{wdiam}(\text{MST}(\mathbb{K}_n[U_i])) \geq |U_i|^{-\frac{1}{11}}\right) \leq \sum_{i=0}^k \mathbb{P}\left(\text{wdiam}(\text{MST}(\mathbb{K}_n[U_i])) \geq |U_i|^{-\frac{1}{11}}\right) \leq \sum_{i=0}^k \frac{1}{|U_i|^2};$$

where we have used that $\frac{7 \log^4 n}{n^{1/10}} \leq \frac{1}{n^{1/11}}$ and that $\frac{4}{n \log n} < \frac{1}{n^2}$ for n large. Since $|U_i| = |U_0| + i = |V_n| + i$, it follows that for all n sufficiently large,

$$\mathbb{P}\left(\exists i : \text{wdiam}(\text{MST}(\mathbb{K}_n[U_i])) \geq |U_i|^{-\frac{1}{11}}\right) \leq \sum_{s=|V_n|}^n \frac{1}{s^2} \leq \frac{1}{|V_n| - 1} \rightarrow 0.$$

In view of (2.1), this yields that

$$\mathbb{P}\left(\text{cost}(\mathbb{K}_n, H'_n) \geq 1 + \epsilon\right) \leq \mathbb{P}\left(\exists i : \text{wdiam}(\text{MST}(\mathbb{K}_n[U_i])) \geq \epsilon\right) \rightarrow 0,$$

as desired. \square

The remainder of Section 2 is devoted to proving Proposition 2.1.

2.1. A special case of Proposition 2.1. To prove Proposition 2.1, we need to bound $\text{cost}(\mathbb{G}, H)$ when H is a spanning subgraph of G with $H[n-1] = \text{MST}(\mathbb{G}[n-1])$. It is useful to first treat the special case that H only contains one edge which does not lie in $\text{MST}(\mathbb{G}[n-1])$, and more specifically that n is a leaf and H is a tree. We will later use this case as an input to the general argument.

Proposition 2.4. *In the setting of Proposition 2.1, if n is a leaf of H then*

$$\text{cost}(\mathbb{G}, H) \leq 1 + \text{wdiam}(\text{MST}(\mathbb{G}[n-1])).$$

The next lemma will be useful in the proof of both the special case and the general case; informally, it states that optimizing sequences never remove MST edges that are already present, and that optimizing sequences do not create cycles.

Lemma 2.5. *Let $(S_i, 1 \leq i \leq m)$ be an MST sequence for (\mathbb{G}, H) with corresponding spanning subgraph sequence $(H_i, 0 \leq i \leq m)$. Then*

- (1) *if $e \in \text{E}(\text{MST}(\mathbb{G}))$ and $e \in \text{E}(H_i)$, then $e \in \text{E}(H_j)$ for all $i \leq j \leq m$, and*
- (2) *if H_i is a tree, then H_j is a tree for all $i \leq j \leq m$.*

Proof. We use the standard fact that if $\mathbb{G} = (V, E, w)$ is a weighted graph with all edge weights distinct, then $e \in \text{E}(\text{MST}(\mathbb{G}))$ if and only if e is not the heaviest edge of any cycle in \mathbb{G} .

Fix $e \in \text{E}(\text{MST}(\mathbb{G}))$ and suppose that $e \in \text{E}(H_i)$. If the endpoints of e do not both lie in S_{i+1} then clearly $e \in \text{E}(H_{i+1})$ since H_i and H_{i+1} agree except on S_{i+1} . If the endpoints of e both lie in S_{i+1} then since e is not the heaviest edge of any cycle in \mathbb{G} , it is not the heaviest edge of any cycle in $\mathbb{G}[S_{i+1}]$. Thus $e \in \text{E}(\text{MST}(\mathbb{G}[S_{i+1}]))$, and so again $e \in \text{E}(H_{i+1})$. It follows by induction that $e \in \text{E}(H_j)$ for all $i \leq j \leq m$.

The second claim of the lemma is immediate from the fact that if T is any tree, S is a subset of $V(T)$ such that $T[S]$ is a tree, and T' is another tree with $V(T) = S$, then the graph with vertices $V(T)$ and edges $(\text{E}(T) \setminus \text{E}(T[S])) \cup \text{E}(T')$ is again a tree. \square

We now assume \mathbb{G} and H are as in Proposition 2.4. Define an optimizing sequence $\mathbb{S} = (S_i, 1 \leq i \leq n-1)$ for (\mathbb{G}, H) as follows. Let S_1 be the set of vertices on the path from n to 1 in $H_0 = H$, and let $H_1 = \Phi_{\mathbb{G}}(H_0, S_1)$. Then, inductively, for $1 < i \leq n-1$ let S_i be the set of vertices on the path from n to i in H_{i-1} and let $H_i = \Phi_{\mathbb{G}}(H_{i-1}, S_i)$. Since $H = H_0$ is a tree, by point 2 of Lemma 2.5 it follows that H_i is a tree for all i , so the paths S_i are uniquely determined and the sequence \mathbb{S} is well-defined.

Proposition 2.4 is now an immediate consequence of the following two lemmas.

Lemma 2.6. \mathbb{S} is an MST sequence for (\mathbb{G}, H) .

Proof. Since H_m is a tree, it suffices to show that $\text{MST}(\mathbb{G})$ is a subtree of H_m . Let $e \in \text{E}(\text{MST}(\mathbb{G}))$. Then either $e \in \text{E}(H_0[n-1])$ or $e = in$ for some $i \in [n-1]$. If $e \in \text{E}(H_0[n-1])$ then $e \in \text{E}(H_0)$ meaning that, by point 1 of Lemma 2.5, we have $e \in \text{E}(H_m)$. Otherwise, if $e = in$ for some $i \in [n-1]$, then $e \in \text{E}(\mathbb{G}[S_i])$ since S_i is the set of vertices on a path from n to i . Hence, $e \in \text{E}(\text{MST}(\mathbb{G}[S_i]))$, meaning that $e \in \text{E}(H_i)$. Once again, by point 1 of Lemma 2.5, this implies that $e \in \text{E}(H_m)$, proving that $\text{MST}(\mathbb{G})$ is a subtree of H_m . \square

Lemma 2.7. $\text{wt}(\mathbb{S}) \leq 1 + \text{wdiam}(\text{MST}(\mathbb{G}[n-1]))$

Proof. Let $i \in [n-1]$. Notice that the path from n to i in H_{i-1} contains a single edge from n to $[n-1]$. Hence, the weight of this path is bounded from above by $1 + \text{wdiam}(H_{i-1}[n-1])$. To prove the lemma it therefore suffices to show that $\text{E}(H_i[n-1]) \subseteq \text{E}(H_0[n-1]) = \text{E}(\text{MST}(\mathbb{G}[n-1]))$.

We prove this by induction on i , the base case $i = 0$ being automatic. For $i > 0$, suppose that $\text{E}(H_{i-1}[n-1]) \subseteq \text{E}(H_0[n-1])$. Fix any vertices $u, v \in S_i \cap [n-1]$ with $uv \notin \text{E}(H_{i-1})$ and let P be the path from u to v in H_{i-1} . Then P is a subpath of $H_{i-1}[S_i]$, and so by induction it is also a subpath of H_0 . Since $H_0[n-1] = \text{MST}(\mathbb{G}[n-1])$ it follows that P is a subpath of $\text{MST}(\mathbb{G}[n-1])$. This yields that uv is the edge with highest weight on the cycle created by closing P , and all the vertices of this cycle lie in S_i ; so $uv \notin \text{E}(\text{MST}(\mathbb{G}[S_i]))$ and thus $uv \notin \text{E}(H_i)$. This shows that $\text{E}(H_i[S_i]) \subseteq \text{E}(H_{i-1}[S_i]) \subseteq \text{E}(H_0[S_i])$. Since the rest of $H_{i-1}[n-1]$ and $H_i[n-1]$ are identical, it follows that $\text{E}(H_i[n-1]) \subseteq \text{E}(H_0[n-1])$, as required. \square

2.2. The general case of Proposition 2.1. We now lift the assumption that H is a tree; in this case, $\text{E}(H) \setminus \text{E}(H[n-1])$ could contain up to $n-1$ edges. As a result, the MST sequence previously defined in Section 2.1 does not provide us with the desired cost, since a path from n to $i \in [n-1]$ might contain additional edges with n as an endpoint, increasing the weight of the sequence. Thus, we require a more careful method. Informally, our approach is to first apply the method from the previous section to a sequence of subgraphs of $H[n-1]$, each of which is only joined to the vertex n by a single edge, but together which contain all the edges from n to $[n-1]$. We show that this yields a graph which contains the MST of G . We then prove that any cycles in the resulting graph can be removed at a low cost.

Let $\mathbb{G} = (V, E, w)$ be a generic weighted graph with $V = [n]$ and let H be a spanning subgraph of G with $H[n-1] = \text{MST}(\mathbb{G}[n-1])$. Let $\{v_1n, \dots, v_kn\} \subseteq \text{E}(H)$ be the set of edges in H with n as an endpoint, and for $1 \leq i \leq k$ let

$$V_i = \left\{ v \in [n-1] : \text{dist}_{(H[n-1], w)}(v_i, v) = \min \left\{ \text{dist}_{(H[n-1], w)}(v_i, v_j) : 1 \leq j \leq k \right\} \right\}.$$

That is to say, $(V_i, 1 \leq i \leq k)$ is the Voronoi partition of $[n-1]$ in $H[n-1]$ with respect to the vertices v_1, \dots, v_k ; it is indeed a partition since \mathbb{G} is generic.

Note that since $H[n-1] = \text{MST}(\mathbb{G}[n-1])$ it follows that $H[V_i] = \text{MST}(\mathbb{G}[V_i])$ for any $1 \leq i \leq k$. Moreover, vertex n has degree one in $H[V_i \cup \{n\}]$. Using Proposition 2.4, let $\mathbb{S}_i = (S_{i,j}, 1 \leq j \leq m_i)$ be an MST sequence for $(\mathbb{G}[V_i \cup \{n\}], H[V_i \cup \{n\}])$ with weight less than $1 + \text{wdiam}(\text{MST}(\mathbb{G}[V_i])) \leq 1 + \text{wdiam}(\text{MST}(\mathbb{G}[n-1]))$, and write $(H_{i,j}, 0 \leq j \leq m_i)$ for the corresponding subgraph sequence. Now set $m = m_1 + \dots + m_k$ and let $\mathbb{S}^* = (S_1^*, \dots, S_m^*)$ be formed by concatenating $\mathbb{S}_1, \dots, \mathbb{S}_m$, so

$$\mathbb{S}^* = (S_{1,1}, \dots, S_{1,m_1}, \dots, S_{k,1}, \dots, S_{k,m_k}),$$

and let (H_0^*, \dots, H_m^*) be the subgraph sequence corresponding to \mathbb{S}^* .

Lemma 2.8. We have $\text{MST}(\mathbb{G}) \subseteq H_m^*$, and $\text{wt}(\mathbb{S}^*) \leq 1 + \text{diam}(\text{MST}(\mathbb{G}[n-1]))$.

Proof. First, by assumption, $H_0[n-1] = \text{MST}(\mathbb{G}[n-1])$. Since $\text{MST}(\mathbb{G})[n-1]$ is a subgraph of $\text{MST}(\mathbb{G}[n-1])$, point 1 of Lemma 2.5 implies that $\text{MST}(\mathbb{G})[n-1]$ is a subgraph of H_i^* for all i , so in particular of H_m^* .

Next, since V_1, \dots, V_k are disjoint, we have $S_{i,j} \cap S_{i',j'} = \{n\}$ whenever $i \neq i'$, and it follows that $H_{m_1+\dots+m_{i-1}}^*[V_i \cup \{n\}] = H[V_i \cup \{n\}]$ for all $1 \leq i \leq k$. This implies that $H_{m_1+\dots+m_{i-1}+j}^*[V_i \cup \{n\}] = H_{i,j}$ for each $1 \leq j \leq m_i$, so in particular $H_{m_1+\dots+m_i}^*[V_i \cup \{n\}] = \text{MST}(\mathbb{G}[V_i \cup \{n\}])$.

Now fix any edge vn of $\text{MST}(\mathbb{G})$. Then $v \in V_i$ for some $1 \leq i \leq k$, so $vn \in E(\text{MST}(\mathbb{G}[V_i \cup \{n\}]))$. It follows that $vn \in H_{m_1+\dots+m_i}^*$, and thus by point 1 of Lemma 2.5 that vn is an edge of H_m^* . Therefore all edges of $\text{MST}(\mathbb{G})$ are edges of H_m^* , as required.

Finally, the bound on the weight of the sequence is immediate by the definition of \mathbb{S}^* and by using that $\text{wt}(\mathbb{G}[V_i \cup \{n\}], H[V_i \cup \{n\}], \mathbb{S}_i) = \text{wt}(\mathbb{G}, H, \mathbb{S}_i)$. \square

We are now left to deal with the edges $E(H_m^*) \setminus E(\text{MST}(\mathbb{G}))$. This is taken care of in the following lemma.

Lemma 2.9. *Let $\mathbb{G} = (V, E, w)$ be a generic weighted graph with $V = [n]$ and with all edge weights at most 1, and let H be a subgraph of G such that $\text{MST}(\mathbb{G})$ is a subgraph of H . Write $k = |E(H)| - (n-1)$. Then there exists an MST sequence $\mathbb{S}' = (S'_1, \dots, S'_k)$ with*

$$\text{wt}(\mathbb{S}') \leq 1 + \text{wdiam}(\text{MST}(\mathbb{G})).$$

Proof. If H is a tree then there is nothing to prove, so assume G contains at least one cycle (so $k \geq 1$). In this case there exist vertices u, v which are not adjacent in $\text{MST}(\mathbb{G})$ but are joined by an edge in H ; choose such u and v so that the length (number of edges) on the path P from u to v in $\text{MST}(\mathbb{G})$ is as small as possible. Let $S = V(P)$ be the set of vertices of the path P ; then $H[S]$ is a cycle (by the minimality of the length of P), and uv is the edge with largest weight on $H[S]$. It follows that $\text{MST}(\mathbb{G}[S]) = P$, so $\Phi_{\mathbb{G}}(H, S)$ has edge set $E = E(H) \setminus \{uv\}$. Moreover, since P is a path of $\text{MST}(\mathbb{G})$, it follows that

$$w(H[S]) = w(uv) + \text{wt}(P) \leq 1 + \text{wdiam}(\text{MST}(\mathbb{G})).$$

Since $\Phi_{\mathbb{G}}(H, S)$ contains $\text{MST}(G)$ but has one fewer edge than H , the result follows by induction. \square

We now combine Lemmas 2.8 and 2.9 to conclude the proof of Proposition 2.1.

Proof of Proposition 2.1. Let $\mathbb{S}^* = (S_1^*, \dots, S_m^*)$ be the optimization sequence defined above Lemma 2.8, and let (H_0^*, \dots, H_m^*) be the corresponding subgraph sequence. By that lemma, $\text{MST}(\mathbb{G})$ is a subgraph of H_m^* and $\text{wt}(\mathbb{S}^*) \leq 1 + \text{wdiam}(\text{MST}(\mathbb{G}[n-1]))$.

Next let $\mathbb{S}' = (S'_1, \dots, S'_k)$ be an MST sequence for (\mathbb{G}, H_m^*) of weight at most $1 + \text{wdiam}(\text{MST}(\mathbb{G}))$; the existence of such a sequence is guaranteed by Lemma 2.9. Then the concatenation

$$\mathbb{S} = (S_1^*, \dots, S_m^*, S'_1, \dots, S'_k)$$

of \mathbb{S}^* and \mathbb{S}' is an MST sequence for (\mathbb{G}, H) , of weight at most

$$\text{wt}(\mathbb{G}, H, \mathbb{S}) \leq 1 + \max \left\{ \text{wdiam}(\text{MST}(\mathbb{G}[n-1])), \text{wdiam}(\text{MST}(\mathbb{G})) \right\},$$

and the desired bound on $\text{cost}(\mathbb{G}, H)$ follows. \square

3. MST SEQUENCES FOR THE THE CLIQUE, THE STAR, AND THE PATH

This section is aimed at proving Proposition 1.3. We start by proving the result in the case of the clique, since it is straightforward using the result of Lemma 2.9. After that, the case of the star and the path are covered together; the proof in those cases uses the eating algorithm, Corollary 2.2, to find adequate sequences of increasing subsets on which to build increasing sequences of MSTs.

Proof of Proposition 1.3 (Case of the clique). Using Lemma 2.9, since $\text{MST}(\mathbb{K}_n)$ is a subgraph of $H_n = K_n$, it follows that

$$\text{cost}(\mathbb{K}_n, H_n) \leq 1 + \text{wdiam}(\text{MST}(\mathbb{K}_n)).$$

By Theorem 2.3 we have $\text{wdiam}(\text{MST}(\mathbb{K}_n)) \xrightarrow{\mathbb{P}} 0$, and the result follows. \square

3.1. MST sequences for the star and the path. In this section, we assume that H_n is either a star or a path. If H_n is a star, then by relabeling we may assume H_n has center n , so has edge set $\{e_1, \dots, e_{n-1}\}$ with $e_i = in$; call this star S_n . If H_n is a path, then by relabeling we may assume H_n is the path $P_n = 12 \dots n$, so has edge set $\{e_i, \dots, e_{n-1}\}$ with $e_i = i(i+1)$. In either case, with this edge labeling, for any $1 \leq i < j \leq n-1$, the set $V(i, j)$ defined as the endpoints in $\{e_i, \dots, e_{j-1}\}$ is connected in H_n . Note that $V(i, j) = \{i, \dots, j-1\} \cup \{n\}$ when H_n is a star and $V(i, j) = \{1, \dots, j\}$ when H_n is a path, and in both cases $|V(i, j)| = j - i + 1$. For the remainder of the section, it might be helpful to imagine that H_n is the path, $12 \dots n$.

Recall that $\mathbb{X} = (X_e, e \in E(K_n))$ is a set of independent Uniform $[0, 1]$ random variables. For $W \in (0, 1)$ and $2 \leq L < n - 1$, let

$$(3.1) \quad \mathbb{I} = \mathbb{I}(W, L) = (n - L) \wedge \min \left\{ i : \forall i \leq j < i + L, X_{e_j} \leq W \right\}.$$

Note that \mathbb{I} is a function of \mathbb{X} and more precisely that

$$\{\mathbb{I} \leq k\} \in \sigma \left(\{X_{e_i} \leq W\}, 1 \leq i < k + L \right),$$

where $\sigma(X)$ is the σ -algebra generated by X .

Next, let $\mathbb{U} = \mathbb{U}(\mathbb{I}) = (U_i, 0 \leq i < n - L)$ be the sequence of sets defined as follows.

$$(3.2) \quad (U_0, \dots, U_{n-L-1}) = \left(V(\mathbb{I}, \mathbb{I} + L), \dots, V(\mathbb{I}, n), V(\mathbb{I} - 1, n), \dots, V(1, n) \right).$$

In words, U_0 is the set of vertices that belong to the edges $e_{\mathbb{I}}, \dots, e_{\mathbb{I}+L-1}$ (that is $V(\mathbb{I}, \mathbb{I} + L)$); then we sequentially build U_1, \dots, U_{n-L-1} by first adding the vertices belonging to $e_{\mathbb{I}+L}, \dots, e_{n-1}$, then adding the vertices belonging to $e_{\mathbb{I}-1}, \dots, e_1$; see Figure 1 for a representation of \mathbb{I} and \mathbb{U} .

We now use the sequence \mathbb{U} to bound the cost of (\mathbb{K}_n, H_n) when H_n is a star or a path. The following lemma gives a first bound on the cost using \mathbb{U} .

Lemma 3.1. *Let H_n be the star S_n or path P_n . Then, conditionally given that $\mathbb{I}(W, L) < n - L$, we have*

$$\text{cost}(\mathbb{K}_n, H_n) \leq \max \left\{ WL, 1 + \max \left\{ \text{wdiam} \left(\text{MST}(\mathbb{K}_n[U_i]) \right) : 0 \leq i < n - L \right\} \right\}.$$

Proof. This result almost directly follows from Corollary 2.2. Indeed, let $H'_n = \Phi(H_n, U_0)$. Then the sets U_0, \dots, U_{n-L-1} satisfy the condition of Corollary 2.2 with $H = H'_n$, implying that

$$\text{cost}(\mathbb{K}_n, H'_n) \leq 1 + \max \left\{ \text{wdiam} \left(\text{MST}(\mathbb{K}_n[U_i]) \right) : 0 \leq i < n - L \right\}.$$

But now, by concatenating any minimal weight MST sequence for $(\mathbb{K}_n[U_0], H_n[U_0])$ and any minimal weight MST sequence for (\mathbb{K}_n, H'_n) , it follows that

$$\text{cost}(\mathbb{K}_n, H_n) \leq \max \left\{ \text{cost}(\mathbb{K}_n[U_0], H_n[U_0]), \text{cost}(\mathbb{K}_n, H'_n) \right\}.$$

In order to complete the proof of the lemma, note that, conditionally given $\mathbb{I} < n - L$,

$$w(H_n[U_0]) = \sum_{e \in E(H_n[U_0])} X_e \leq WL.$$

Taking $\mathbb{S} = (U_0)$, this yields

$$\text{cost}(\mathbb{K}_n[U_0], H_n[U_0]) \leq \text{wt}(\mathbb{K}_n[U_0], H_n[U_0], \mathbb{S}) = w(H_n[U_0]) \leq WL.$$


This proves the desired upper bound and concludes the proof of the lemma. \square

The next two results, combined with Lemma 3.1, will allow us to give the full proof of Proposition 1.3 when H_n is either a star or a path.

Proposition 3.2. *For any $\epsilon > 0$, for $W = \frac{1}{\log n}$ and $L = \lfloor \log \log n \rfloor$, as $n \rightarrow \infty$ we have*

$$\mathbb{P} \left(\exists U \in \mathbb{U}(\mathbb{I}(W, L)) : \text{wdiam} \left(\text{MST}(\mathbb{K}_n[U]) \right) > \epsilon \right) \rightarrow 0.$$

$$W = 0.2 \quad L = 3$$



the ordered line
with random
edge weights.

I is the first
index followed
by $L = 3$ edges
of weight less
than $W = 0.2$.

The sets in \mathbb{U}
are built using
 $L = 3$
and $I = 4$ to set
 $U_0 = \{4, 5, 6, 7\}$
before expanding
on both sides
(right then left).

FIGURE 1. An example of I and \mathbb{U} for an instance of the weighted ordered line (H_n, w) , with $W = 0.2$ and $L = 3$. First, I is set to be the first sequence of $L = 3$ consecutive edges with weights less than $W = 0.2$. In this example, $I = 4$. Then, given I , set $U_0 = V(I, I + L) = \{4, 5, 6, 7\}$ and expand first to the right and then to left to obtain U_1, \dots, U_{n-L-1} . In other words, in order to obtain U_1, U_2, U_3, U_4 , and U_5 , we sequentially add 8, 9, 3, 2, and 1 to U_0 .

Lemma 3.3. *Let $W = \frac{1}{\log n}$ and $L = \lfloor \log \log n \rfloor$. Then, for any $a > 0$, as $n \rightarrow \infty$ we have*

$$\mathbb{P}(I(W, L) \geq n^a) \rightarrow 0.$$

Lemma 3.3 is straightforward and we prove it immediately. On the other hand, Proposition 3.2 is quite technical and we dedicate Section 3.2 below to proving it.

Proof of Lemma 3.3. For any integer $k \geq 1$, by the definition of I ,

$$\begin{aligned} \mathbb{P}(I \geq kL + 1) &= \mathbb{P}\left(\forall i < kL + 1, \exists j \in \{i, \dots, i + L - 1\} : X_{e_j} > W\right) \\ &\leq \mathbb{P}\left(\forall i \in \{1, 1 + L, \dots, 1 + (k - 1)L\}, \exists j \in \{i, \dots, i + L - 1\} : X_{e_j} > W\right). \end{aligned}$$

But then, by independence of the weights of \mathbb{X} , we have

$$\mathbb{P}(I \geq kL + 1) = \prod_{i=0}^{k-1} \mathbb{P}\left(\exists j \in \{1 + iL, \dots, 1 + (i + 1)L - 1\} : X_{e_j} > W\right) = \prod_{i=0}^{k-1} (1 - W^L) \leq e^{-kW^L},$$

where the last inequality follows from the convexity of the exponential. Applying this result with $k = \lfloor \frac{n^a - 1}{L} \rfloor$, we obtain

$$\mathbb{P}(\mathbf{I} \geq n^a) \leq \mathbb{P}(\mathbf{I} \geq kL + 1) \leq \exp\left(-\left\lfloor \frac{n^a - 1}{L} \right\rfloor \cdot WL\right),$$

and the final expression tends to 0 as $n \rightarrow \infty$. \square

Proof of Proposition 1.3 (Case of the star and the path). Let $W = \frac{1}{\log n}$ and $L = \lfloor \log \log n \rfloor$. Fixing $\varepsilon > 0$, we have

$$\mathbb{P}\left(\text{cost}(\mathbb{K}_n, H_n) > 1 + \varepsilon\right) \leq \mathbb{P}\left(\text{cost}(\mathbb{K}_n, H_n) > 1 + \varepsilon \mid \mathbf{I} < n - L\right) + \mathbb{P}(\mathbf{I} = n - L).$$

Applying Lemma 3.3 with any $a < 1$, for large enough n we have

$$\mathbb{P}(\mathbf{I} = n - L) \leq \mathbb{P}(\mathbf{I} \geq n^a) \rightarrow 0.$$

Hence, we have

$$\mathbb{P}\left(\text{cost}(\mathbb{K}_n, H_n) > 1 + \varepsilon\right) = \mathbb{P}\left(\text{cost}(\mathbb{K}_n, H_n) > 1 + \varepsilon \mid \mathbf{I} < n - L\right) + o(1).$$

Since $WL \rightarrow 0$, combining the previous bound with Lemma 3.1 leads to

$$\begin{aligned} & \mathbb{P}\left(\text{cost}(\mathbb{K}_n, H_n) > 1 + \varepsilon\right) \\ & \leq \mathbb{P}\left(\max\left\{WL, 1 + \max\left\{\text{wdiam}(\text{MST}(\mathbb{K}_n[U_i]))\right\}\right\} > 1 + \varepsilon \mid \mathbf{I} < n - L\right) + o(1) \\ & = \mathbb{P}\left(\max\left\{\text{wdiam}(\text{MST}(\mathbb{K}_n[U])) : U \in \mathbb{U}\right\} > \varepsilon \mid \mathbf{I} < n - L\right) + o(1). \end{aligned}$$

The upper bound now follows from Proposition 3.2, once again since $\mathbb{P}(\mathbf{I} < n - L) \rightarrow 0$. \square

3.2. Proof of Proposition 3.2. In this section, we prove Proposition 3.2, which concludes the proof of Proposition 1.3. Before doing so, we state a proposition which is an important input to the proof.

Proposition 3.4. *Let $\mathbb{G} = (G, w)$ be a weighted graph. Let T be a subtree (not necessarily spanning) of \mathbb{G} and let $\mathbb{G}^* = (G, w^*)$ be a weighted graph such that $w^*(e) \leq w(e)$ for $e \in \mathbf{E}(T)$ and $w^*(e) = w(e)$ otherwise. Then*

$$\text{wdiam}(\text{MST}(\mathbb{G}^*)) \leq w^*(T) + |\mathbf{V}(T)| \times \text{wdiam}(\text{MST}(\mathbb{G})).$$

Moreover, if T is a subtree of $\text{MST}(\mathbb{G}^*)$, then

$$\text{wdiam}(\text{MST}(\mathbb{G}^*)) \leq w^*(T) + 2 \times \text{wdiam}(\text{MST}(\mathbb{G})).$$

Proof. Let us try to understand the relation between $\text{MST}(\mathbb{G})$ and $\text{MST}(\mathbb{G}^*)$. First note that

$$(3.3) \quad \mathbf{E}(\text{MST}(\mathbb{G}^*)) \subseteq \mathbf{E}(\text{MST}(\mathbb{G})) \cup \mathbf{E}(T).$$

Indeed, any edge $e \notin \mathbf{E}(T)$ has the same weight with respect to w and w^* . Then, for any $e \in \mathbf{E}(\text{MST}(\mathbb{G}^*)) \setminus \mathbf{E}(T)$, no cycle has e as the heaviest edge with respect to w^* , which implies that no cycle has e as the heaviest edge with respect to w , and thus $e \in \mathbf{E}(\text{MST}(\mathbb{G}))$.

Consider now a path P contained in $\text{MST}(\mathbb{G}^*)$. Using (3.3), we have

$$\mathbf{E}(P) \subseteq \mathbf{E}(\text{MST}(\mathbb{G})) \cup \mathbf{E}(T),$$

so we may uniquely decompose P into pairwise edge-disjoint paths P_0, \dots, P_{2k} , where $k \geq 1$, and P_i is a subpath of T for i odd and of $\text{MST}(\mathbb{G})$ for i even (it is possible that either or both of P_0, P_{2k} consists of a single vertex). Since $P_1, P_3, \dots, P_{2k-1}$ are disjoint subpaths of T , it follows that $k \leq |\mathbf{E}(T)|$

and that $\sum_{i \text{ odd}} w(P_i) \leq w(T)$. Moreover, each of the paths P_0, P_2, \dots, P_{2k} have weight at most $\text{wdiam}(\text{MST}(\mathbb{G}))$, so

$$(3.4) \quad \begin{aligned} \sum_{i \text{ even}} w(P_i) &\leq (k+1) \times \text{wdiam}(\text{MST}(\mathbb{G})) \\ &\leq (|\mathbb{E}(T)| + 1) \times \text{wdiam}(\text{MST}(\mathbb{G})) \\ &= |\mathbb{V}(T)| \times \text{wdiam}(\text{MST}(\mathbb{G})). \end{aligned}$$

The first bound of the proposition follows since

$$w(P) = \sum_{i \text{ even}} w(P_i) + \sum_{i \text{ odd}} w(P_i).$$

To establish the second bound, note that if T is a subtree of $\text{MST}(\mathbb{G}^*)$ then in the above decomposition of P we must have $k = 1$; a path in $\text{MST}(\mathbb{G}^*)$ may enter T and then leave it, after which it can never reenter T . In this case the first summation of (3.4) becomes

$$\sum_{i \text{ even}} w(P_i) \leq 2 \times \text{wdiam}(\text{MST}(\mathbb{G})),$$

so we obtain

$$w(P) = \sum_{i \text{ even}} w(P_i) + \sum_{i \text{ odd}} w(P_i) \leq w(T) + 2 \times \text{wdiam}(\text{MST}(\mathbb{G})),$$

as required. \square

For the remainder of this section we assume $W = \frac{1}{\log n}$ and $L = \lfloor \log \log n \rfloor$ and write $\mathbb{I} = \mathbb{I}(W, L)$. Consider the partition $\mathbb{U} = \mathbb{U}_r^- \cup \mathbb{U}_r^+ \cup \mathbb{U}_\ell$ where $\mathbb{U}_r^- = \mathbb{U}_r^-(\mathbb{I}) = (U_i, 0 \leq i \leq \min(L^{20}, n - \mathbb{I} - L))$, $\mathbb{U}_r^+ = \mathbb{U}_r^+(\mathbb{I}) = (U_i, \min(L^{20}, n - \mathbb{I} - L) < i \leq n - \mathbb{I} - L)$, and $\mathbb{U}_\ell = \mathbb{U}_\ell(\mathbb{I}) = (U_i, n - \mathbb{I} - L < i \leq n - L - 1)$. Then, in the case where $\mathbb{I} < n - L - L^{20}$, \mathbb{U}_r^- corresponds to adding the first L^{20} vertices on the right of U_0 , \mathbb{U}_r^+ corresponds to adding all remaining vertices on the right, and \mathbb{U}_ℓ corresponds to adding the vertices on the left of U_0 . We aim to prove tail bounds similar to that of Proposition 3.2 for each of the sets \mathbb{U}_r^- , \mathbb{U}_r^+ , and \mathbb{U}_ℓ , and we start with an important lemma regarding the distribution of \mathbb{G} conditioned on the value of \mathbb{I} .

Lemma 3.5. *Fix $k < n - L$ and let $\mathbb{K}_n^* = (K_n, \mathbb{X}^*)$ have the law of \mathbb{K}_n conditioned on the event that $\mathbb{I}(W, L) = k$. Then for any $e \in \{e_i, k \leq i < k + L\}$, X_e^* is a $\text{Uniform}[0, W]$; for any $e \notin \{e_i : 1 \leq i < k + L\}$, $X_{e_i}^*$ is a random $\text{Uniform}[0, 1]$, and the edge weights $(X_e^*, e \in \mathbb{E}(K_n) \setminus \{e_i, 1 \leq i < k\})$ are mutually independent and independent of $(X_e^*, e \in \{e_i, 1 \leq i < k\})$. It follows that there exists a coupling between $\mathbb{K}_n^* = (K_n, \mathbb{X}^*)$ and $\mathbb{K}_n' = (K_n, \mathbb{X}')$ where \mathbb{X}' is a set of independent $\text{Uniform}[0, 1]$, such that $X_e^* \leq X_e'$ if $e \in \{e_i : k \leq i < k + L\}$, and $X_e^* = X_e'$ if $e \in \mathbb{E}(K_n) \setminus \{e_i : 1 \leq i < k + L\}$.*

Proof. Using the definition of \mathbb{I} , we know that

$$\{\mathbb{I} = k\} \in \sigma\left(\{X_{e_i} \leq W : 1 \leq i < k + L\}\right),$$

from which it directly follows that the distribution of X_e is a $\text{Uniform}[0, 1]$ for any $e \notin \{e_{n,i} : 1 \leq i < k + L\}$. Furthermore, for any $e \in \{e_i : k \leq i < k + L\}$, X_e conditioned on $\mathbb{I} = k$ is the same as X_e conditioned on $X_e \leq W$. Since X_e is uniformly distributed, it follows that X_e conditioned on $\mathbb{I} = k$ is a $\text{Uniform}[0, W]$. Finally, note that

$$\{\mathbb{I} = k\} = \{X_{e_i} \leq W : k \leq i < k + L\} \cap \bigcap_{j=1}^{k-1} \left\{ \exists j \leq i < \min\{j + L, k\} : X_{e_i} > W \right\},$$

from which we see that the edges of $\mathbb{E}(K_n) \setminus \{e_i, 1 \leq i < k\}$ are conditionally independent of $\{e_i, 1 \leq i < k\}$ given that $\mathbb{I} = k$. It follows that all the edges in $\mathbb{E}(\mathbb{K}_n) \setminus \{e_i, 1 \leq i < k\}$ have independent weights in \mathbb{K}_n^* . The existence of the coupling asserted in the lemma is then an immediate consequence. \square

We now split the proof of Proposition 3.2 into proving analogous statements for the three different sets $\mathbb{U}_r^-, \mathbb{U}_r^+$, and \mathbb{U}_ℓ .

First right set \mathbb{U}_r^- .

Lemma 3.6. *For any $\epsilon > 0$, we have*

$$\mathbb{P}\left(\exists U \in \mathbb{U}_r^- : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon\right) \longrightarrow 0.$$

Proof. Fix $0 < a < 1$ and assume n is large enough so that $n^a < n - L - L^{20}$. Then, by Lemma 3.3, we have

$$\begin{aligned} & \mathbb{P}\left(\exists U \in \mathbb{U}_r^- : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon\right) \\ & \leq \mathbb{P}\left(\exists U \in \mathbb{U}_r^- : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon \mid \mathbf{I} < n - L - L^{20}\right) + \mathbb{P}(\mathbf{I} \geq n - L - L^{20}) \\ & = \mathbb{P}\left(\exists U \in \mathbb{U}_r^- : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon \mid \mathbf{I} < n - L - L^{20}\right) + o(1). \end{aligned}$$

Next fix $k < n - L - L^{20}$ and condition on the event $\mathbf{I} = k$. Under this conditioning, $\mathbb{U}_r^- = \mathbb{U}_r^-(\mathbf{I}) = \mathbb{U}_r^-(k)$ is a deterministic sequence of sets. Further recall from (3.2) that $U_0 = V(k, k + L)$ consists of the endpoints of the edges e_k, \dots, e_{k+L-1} , so equals $\{k, \dots, k + L\}$ if H_n is the path P_n and equals $\{k, \dots, k + L - 1, n\}$ if H_n is the star S_n . Let $T = H_n[U_0]$. Since $\mathbf{I} = k < n - L$, all edges in T have weight less than W . Now, suppose that all other edges of $\mathbb{K}_n[U_{L^{20}}]$ have weight larger than W . In this case, T is a subtree of $\text{MST}(\mathbb{K}_n[U_{L^{20}}])$, from which it follows that T is a subtree of $\text{MST}(\mathbb{K}_n[U_i])$ for any $0 \leq i \leq L^{20}$ (since $U_i \subset U_{L^{20}}$ for such U_i). Now, using that $\{\mathbf{I} = k\} \in \sigma(\{X_{e_i} : 1 \leq i < k + L\})$, we have

$$\mathbb{P}\left(\forall e \in \mathbf{E}(\mathbb{K}_n[U_{L^{20}}]) \setminus \mathbf{E}(T), X_e > W \mid \mathbf{I} = k\right) = (1 - W)^{\binom{L^{20}}{2} - L}.$$

Since $W = \frac{1}{\log n}$, we have $1 - W \geq \exp(-2W)$ for n large, so

$$\begin{aligned} \mathbb{P}\left(\mathbf{E}(T) \subset \mathbf{E}(\text{MST}(\mathbb{K}_n[U_{L^{20}}])) \mid \mathbf{I} = k\right) & \geq (1 - W)^{\binom{L^{20}}{2} - L} \\ & \geq \exp\left(-2W \left(\binom{L^{20}}{2} - L\right)\right) \\ & \geq \exp(-WL^{40}) \\ & \geq 1 - \frac{(\log \log n)^{40}}{\log n}, \end{aligned}$$

the last inequality holding since $W = \frac{1}{\log n}$, $L = \lfloor \log \log n \rfloor$, and $e^{-x} \geq 1 - x$ for $x \geq 0$. Hence,

(3.5)

$$\begin{aligned} & \mathbb{P}\left(\exists U \in \mathbb{U}_r^- : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon \mid \mathbf{I} = k\right) \\ & \leq \mathbb{P}\left(\exists U \in \mathbb{U}_r^- : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon, \mathbf{E}(T) \subset \mathbf{E}(\text{MST}(\mathbb{K}_n[U_{L^{20}}])) \mid \mathbf{I} = k\right) + \frac{(\log \log n)^{40}}{\log n}. \end{aligned}$$

Let $(\mathbb{K}_n^*, \mathbb{K}'_n)$ be as in Lemma 3.5. By the definition of \mathbb{K}_n^* and (3.5), we have that

$$\begin{aligned} & \mathbb{P}\left(\exists U \in \mathbb{U}_r^- : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon \mid \mathbf{I} = k\right) \\ & \leq \mathbb{P}\left(\exists U \in \mathbb{U}_r^-(k) : \text{wdiam}(\text{MST}(\mathbb{K}_n^*[U])) > \epsilon, \mathbf{E}(T) \subset \mathbf{E}(\text{MST}(\mathbb{K}_n^*[U_{L^{20}}]))\right) + \frac{(\log \log n)^{40}}{\log n}. \end{aligned}$$

Now, note that if $\mathbf{E}(T) \subset \mathbf{E}(\text{MST}(\mathbb{K}_n^*[U_{L^{20}}]))$, then for any $U \in \mathbb{U}_r^-(k)$, $\mathbf{E}(T) \subset \mathbf{E}(\text{MST}(\mathbb{K}_n^*[U]))$, since $\text{MST}(\mathbb{K}_n^*[U_{L^{20}}])[U]$ is a subgraph of $\text{MST}(\mathbb{K}_n^*[U])$. Applying Proposition 3.4 to $\text{MST}(\mathbb{K}_n^*[U])$ and

$\text{MST}(\mathbb{K}'_n[U])$, it follows that

$$\begin{aligned} & \mathbb{P}\left(\exists U \in \mathbb{U}_r^-(k) : \text{wdiam}(\text{MST}(\mathbb{K}'_n[U])) > \epsilon, \mathbf{E}(T) \subset \mathbf{E}(\text{MST}(\mathbb{K}'_n[U_{L^{20}}]))\right) \\ & \leq \mathbb{P}\left(\exists U \in \mathbb{U}_r^-(k) : w^*(T) + 2 \times \text{wdiam}(\text{MST}(\mathbb{K}'_n[U])) > \epsilon, \mathbf{E}(T) \subset \mathbf{E}(\text{MST}(\mathbb{K}'_n[U_{L^{20}}]))\right) \\ & \leq \mathbb{P}\left(\exists U \in \mathbb{U}_r^-(k) : w^*(T) + 2 \times \text{wdiam}(\text{MST}(\mathbb{K}'_n[U])) > \epsilon\right). \end{aligned}$$

Using that $w^*(T) \leq WL$ and combining the two previous inequalities yields the bound

$$(3.6) \quad \begin{aligned} & \mathbb{P}\left(\exists U \in \mathbb{U}_r^- : \text{wdiam}(\text{MST}(\mathbb{K}'_n[U])) > \epsilon \mid \mathbf{I} = k\right) \\ & \leq \mathbb{P}\left(\exists U \in \mathbb{U}_r^-(k) : \text{wdiam}(\text{MST}(\mathbb{K}'_n[U])) > (\epsilon - WL)/2\right) + \frac{(\log \log n)^{40}}{\log n}. \end{aligned}$$

We can now replace \mathbb{K}'_n by \mathbb{K}_n since they are identically distributed. Furthermore, recall that Theorem 2.3 states that, for n sufficiently large, we have

$$\mathbb{P}\left(\text{wdiam}(\text{MST}(\mathbb{K}_n)) \geq \frac{7 \log^4 n}{n^{1/10}}\right) \leq \frac{4}{n^{\log n}}.$$

Since $L \rightarrow \infty$ and $WL \rightarrow 0$ as $n \rightarrow \infty$, and since any set $U \in \mathbb{U}_r^-$ has size $|U| \geq |U_0| = L + 1$, we can choose n large enough so that, for any set $U \in \mathbb{U}_r^-$, we have $7 \log^4 |U| / |U|^{1/10} \leq (\epsilon - WL)/2$. It follows that

$$\begin{aligned} & \mathbb{P}\left(\exists U \in \mathbb{U}_r^-(k) : \text{wdiam}(\text{MST}(\mathbb{K}'_n[U])) > (\epsilon - WL)/2\right) \\ & \leq \mathbb{P}\left(\exists U \in \mathbb{U}_r^-(k) : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) \geq \frac{7 \log^4 |U|}{|U|^{1/10}}\right) \\ & \leq \sum_{U \in \mathbb{U}_r^-(k)} \frac{4}{|U|^{\log |U|}} \end{aligned}$$

The final step of the proof is to use that $\mathbb{U}_r^-(k) = (U_i, 0 \leq i \leq L^{20})$ where $|U_i| = |U_0| + i = L + i + 1$, along with the fact that $4/n^{\log n} \leq 1/n^2$ for n large enough, to obtain that

$$\mathbb{P}\left(\exists U \in \mathbb{U}_r^-(k) : \text{wdiam}(\text{MST}(\mathbb{K}'_n[U])) > (\epsilon - WL)/2\right) \leq \sum_{k=L+1}^{L+L^{20}+1} \frac{1}{k^2} \leq \frac{1}{L} \leq \frac{2}{\log \log n}.$$

Plugging this into (3.6), it follows that

$$\mathbb{P}\left(\exists U \in \mathbb{U}_r^- : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon \mid \mathbf{I} = k\right) \leq \frac{(\log \log n)^{40}}{\log n} + \frac{2}{\log \log n}.$$

Finally, since the previous inequality holds for any $k < n - L - L^{20}$, we have

$$\begin{aligned} & \mathbb{P}\left(\exists U \in \mathbb{U}_r^- : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon\right) \\ & = \mathbb{P}\left(\exists U \in \mathbb{U}_r^- : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon \mid \mathbf{I} < n - L - L^{20}\right) + o(1) \\ & \leq \frac{(\log \log n)^{40}}{\log n} + \frac{1}{(\log \log n)^{20}} + o(1) \rightarrow 0, \end{aligned}$$

which is the desired result. \square

Second right set \mathbb{U}_r^+ .

Lemma 3.7. *For any $\epsilon > 0$, we have*

$$\mathbb{P}\left(\exists U \in \mathbb{U}_r^+ : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon\right) \rightarrow 0.$$

Proof. Fix $0 < a < 1$ and assume n is large enough so that $n^a < n - L - L^{20}$. Then, by Lemma 3.3, we have

$$\begin{aligned} & \mathbb{P}\left(\exists U \in \mathbb{U}_r^+ : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon\right) \\ & \leq \mathbb{P}\left(\exists U \in \mathbb{U}_r^+ : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon \mid \mathbf{I} < n - L - L^{20}\right) + \mathbb{P}(\mathbf{I} \geq n - L - L^{20}) \\ & = \mathbb{P}\left(\exists U \in \mathbb{U}_r^+ : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon \mid \mathbf{I} < n - L - L^{20}\right) + o(1). \end{aligned}$$

Fix now $k < n - L - L^{20}$ and condition on the event $\mathbf{I} = k$. Let $T = H_n[U_0]$ and let $(\mathbb{K}'_n, \mathbb{K}^*_n)$ be given by the coupling in Lemma 3.5. Then, by Proposition 3.4,

$$\begin{aligned} & \mathbb{P}\left(\exists U \in \mathbb{U}_r^+ : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon \mid \mathbf{I} = k\right) \\ & \leq \mathbb{P}\left(\exists U \in \mathbb{U}_r^+(k) : \text{wdiam}(\text{MST}(\mathbb{K}^*_n[U])) > \epsilon\right) \\ & \leq \mathbb{P}\left(\exists U \in \mathbb{U}_r^+(k) : w^*(T) + |\mathbf{V}(T)| \times \text{wdiam}(\text{MST}(\mathbb{K}'_n[U])) > \epsilon\right) \\ & \leq \mathbb{P}\left(\exists U \in \mathbb{U}_r^+(k) : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > (\epsilon - WL)/(L + 1)\right), \end{aligned}$$

where the last step follows from the fact that $w^*(T) \leq WL$ conditionally given that $\mathbf{I} < n - L$, that $|\mathbf{V}(T)| = L + 1$, and that \mathbb{K}'_n is distributed as \mathbb{K}_n . Since $x \mapsto \frac{\log^3 x}{x^{1/10}}$ is a decreasing function for large enough x , since any set $U \in \mathbb{U}_r^+$ has size $|U| \geq |U_{L^{20}}| = L + L^{20} + 1$, and since $L = \lfloor \log \log n \rfloor \rightarrow \infty$ and $WL = \lfloor \log \log n \rfloor / \log n \rightarrow 0$, we can choose n large enough so that, for any $U \in \mathbb{U}_r^+$

$$\frac{7 \log^4 |U|}{|U|^{1/10}} \leq \frac{7 \log^4 (L^{20})}{(L^{20})^{1/10}} = \frac{7 \cdot 20^4 \log^4 \cdot (L)}{L^2} \leq \frac{\epsilon - WL}{L + 1}.$$

Then, recalling that $\mathbb{U}_r^+(k) = (U_i, L^{20} < i \leq n - k - L)$ where $|U_i| = |U_0| + i = L + i + 1$, Theorem 2.3 gives us

$$\begin{aligned} & \mathbb{P}\left(\exists U \in \mathbb{U}_r^+(k) : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > (\epsilon - WL)/(L + 1)\right) \\ & \leq \mathbb{P}\left(\exists U \in \mathbb{U}_r^+(k) : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \frac{7 \log^4 |U|}{|U|^{1/10}}\right) \\ & \leq \sum_{U \in \mathbb{U}_r^+(k)} \frac{4}{|U|^{\log |U|}} \\ & \leq \frac{1}{L + L^{20}}, \end{aligned}$$

where the last inequality uses that $x^{\log x} \geq 4x^2$ for x large enough, along with the fact that $|U_i| = L + i + 1$. Therefore,

$$\begin{aligned} & \mathbb{P}\left(\exists U \in \mathbb{U}_r^+ : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon\right) \\ & \leq \mathbb{P}\left(\exists U \in \mathbb{U}_r^+(k) : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > (\epsilon - WL)/(L + 1)\right) + o(1) \\ & \leq \frac{1}{L + L^{20}} + o(1) \rightarrow 0, \end{aligned}$$

concluding the proof of the lemma. \square

Left set \mathbb{U}_ℓ .

Lemma 3.8. *For any $\epsilon > 0$, we have*

$$\mathbb{P}\left(\exists U \in \mathbb{U}_\ell : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon\right) \rightarrow 0.$$

Proof. Fix $a < \frac{1}{4}$. Thanks to Lemma 3.3, we know that $\mathbb{P}(I \geq n^a) \rightarrow 0$. Moreover, note that under this event, any set $U \in \mathbb{U}_\ell$ has size $|U| \geq n - k \geq n - n^a$. Our strategy now is to prove that, due to the large size of these sets, conditioning on the event $\{I < n^a\}$ does not notably affect the structure of $\text{MST}(\mathbb{K}_n[U])$.

Let us try to understand how the edge weights $\{e_1, \dots, e_{n-1}\}$ behave given that $I < n^a$; call \mathbb{K}_n^a the random weighted graph corresponding to the distribution of \mathbb{K}_n conditionally given that $I < n^a$. Recall that $\{I < n^a\} \in \sigma(\{X_{e_i} : 1 \leq i < \lceil n^a \rceil + L\})$ and write $m = \lceil n^a \rceil + L - 1$ (note that e_1, \dots, e_m are the only edges affected when we condition on $I < n^a$). Let $\mathbf{A} = \{i \leq m : X_{e_i} \leq W\}$ and let \mathcal{A} be the collection of sets $A \subset [m]$ such that there exists $i < n^a$ with $\{i, \dots, i+L-1\} \subset A$. Then, by definition, $\{\mathbf{A} \in \mathcal{A}\} = \{I < n^a\}$. Now, for any $A \in \mathcal{A}$, conditionally given that $\mathbf{A} = A$, the weights of e_1, \dots, e_m are independent of each other and are distributed as $\text{Uniform}[0, W]$ or $\text{Uniform}[W, 1]$, according to whether or not the index i of the edge e_i lies in A . This means that for any $x_1, \dots, x_m \in [0, 1]$, and any $A \in \mathcal{A}$, we have

$$\begin{aligned} \mathbb{P}(\forall i \in [m] : X_{e_i} \leq x_i \mid \mathbf{A} = A, I < n^a) &= \mathbb{P}(\forall i \in [m] : X_{e_i} \leq x_i \mid \mathbf{A} = A) \\ &= \left(\prod_{i \in A} \frac{\min\{x_i, W\}}{W} \right) \left(\prod_{i \in [m] \setminus A} \frac{\max\{x_i, W\} - W}{1 - W} \right). \end{aligned}$$

Now, using that $\frac{\max\{x_i, W\} - W}{1 - W} \leq \frac{\min\{x_i, W\}}{W}$, it follows that

$$\mathbb{P}(\forall i \in [m] : X_{e_i} \leq x_i \mid \mathbf{A} = A, I < n^a) \leq \mathbb{P}(\forall i \in [m] : X'_{e_i} \leq x_i),$$

where $(X'_{e_1}, \dots, X'_{e_m})$ are independent $\text{Uniform}[0, W]$. This implies that there exists a generic weighted graph $\mathbb{K}'_n = (K_n, \mathbb{X}')$ with independent weights, where X'_e is a $\text{Uniform}[0, 1]$ if $e \notin \{e_1, \dots, e_m\}$ and a $\text{Uniform}[0, W]$ otherwise, and a coupling between \mathbb{K}'_n and \mathbb{K}_n^a such that $X'_e \leq X_e^a$ for any $e \in E(K_n)$. We now use this coupling to prove the lemma.

Consider the event

$$E' = \left\{ \forall k < n^a, \forall i \in [m], e_i \notin E(\text{MST}(\mathbb{K}'_n[V(k, n)])) \right\}$$

By using two union bounds, we have that

$$\mathbb{P}(E') \geq 1 - \sum_{k < n^a} \sum_{i \in [m]} \mathbb{P}(e_i \in E(\text{MST}(\mathbb{K}'_n[V(k, n)]))).$$

For k and i as in the above sum, if there exists $j \in V(k, n) \setminus e_i$ such that the weight of e_i is larger than the weight of the two other edges in the triangle $\Delta_{i,j}$ formed by e_i and j , then e_i is not in the MST of $\mathbb{K}'_n[V(k, n)]$. This means that

$$\begin{aligned} &\mathbb{P}(e_i \in E(\text{MST}(\mathbb{K}'_n[V(k, n)])) \mid X'_{e_i}) \\ &\leq \mathbb{P}(\forall j \in V(k, n) \setminus e_i, \max(X'_e : e \in \Delta_{i,j}) > X'_{e_i} \mid X'_{e_i}) \\ &= (1 - (X'_{e_i})^2)^{|V(k, n)| - 2} \end{aligned}$$

Using that X'_{e_i} is uniformly distributed over $[0, W]$ and that $|V(k, n)| = n - k + 1$, it follows that

$$\begin{aligned} \mathbb{P}(e_i \in E(\text{MST}(\mathbb{K}'_n[V(k, n)]))) &\leq \frac{1}{W} \int_0^W (1 - x^2)^{n-k+1} dx \\ &\leq \frac{1}{W} \int_0^\infty e^{-(n-k+1)x^2} dx \\ &= \frac{\sqrt{\pi}}{2W\sqrt{n-k+1}}, \end{aligned}$$

from which we obtain

$$\mathbb{P}(E') \geq 1 - \sum_{k < n^a} \sum_{i \in [m]} \frac{\sqrt{\pi}}{2W\sqrt{n-k-1}} \geq 1 - \frac{\sqrt{\pi}}{2} \frac{n^a m}{W\sqrt{n-n^a-1}} \rightarrow 1,$$

where the last convergence follows from $W = \frac{1}{\log n}$, $m = \lceil n^a \rceil + L - 1 = \lceil n^a \rceil + \lfloor \log \log n \rfloor - 1$, and $a < \frac{1}{4}$.

Combining the fact that $\mathbb{P}(I < n^a) \rightarrow 1$ with the definitions of \mathbb{K}_n^a and \mathbb{U}_ℓ , we now have that

$$\begin{aligned} (3.7) \quad & \mathbb{P}(\exists U \in \mathbb{U}_\ell : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon) \\ &= \mathbb{P}(\exists U \in \mathbb{U}_\ell : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon \mid I < n^a) + o(1) \\ &\leq \mathbb{P}(\exists k < n^a : \text{wdiam}(\text{MST}(\mathbb{K}_n^a[V(k, n)])) > \epsilon) + o(1), \end{aligned}$$

where the last inequality comes from the definition of \mathbb{K}_n^a , and is due to $\mathbb{U}_\ell = (V(I-1, n), \dots, V(1, n)) \subset (V(n^a-1, n), \dots, V(1, n))$ whenever $I < n^a$. Note that the coupling between \mathbb{K}_n^a and \mathbb{K}_n' only reduces the weight of the edges e_1, \dots, e_m in \mathbb{K}_n' relative to \mathbb{K}_n^a , from which it follows that, if $e_i \notin E(\text{MST}(\mathbb{K}_n^a[V(k, n)]))$ for some $i \in [m]$, then $e_i \notin E(\text{MST}(\mathbb{K}_n'[V(k, n)]))$. This implies that, conditionally given E' , the trees $\text{MST}(\mathbb{K}_n^a[V(k, n)])$ and $\text{MST}(\mathbb{K}_n'[V(k, n)])$ are equal. Using that $\mathbb{P}(E') \rightarrow 1$, we thus obtain

$$\begin{aligned} (3.8) \quad & \mathbb{P}(\exists k < n^a : \text{wdiam}(\text{MST}(\mathbb{K}_n^a[V(k, n)])) > \epsilon) \\ &= \mathbb{P}(\exists k < n^a : \text{wdiam}(\text{MST}(\mathbb{K}_n^a[V(k, n)])) > \epsilon \mid E') + o(1) \\ &= \mathbb{P}(\exists k < n^a : \text{wdiam}(\text{MST}(\mathbb{K}_n'[V(k, n)])) > \epsilon \mid E') + o(1). \end{aligned}$$

Finally, consider a coupling between \mathbb{K}_n' and \mathbb{K}_n where $X'_e \leq X_e$ for any $e \in E(K_n)$ and such that $X'_e = X_e$ whenever $e \notin \{e_1, \dots, e_m\}$. By using that $\text{MST}(\mathbb{K}_n') = \text{MST}(\mathbb{K}_n)$ whenever E' holds, it follows that

$$\begin{aligned} (3.9) \quad & \mathbb{P}(\exists k < n^a : \text{wdiam}(\text{MST}(\mathbb{K}_n'[V(k, n)])) > \epsilon \mid E') \\ &= \mathbb{P}(\exists k < n^a : \text{wdiam}(\text{MST}(\mathbb{K}_n[V(k, n)])) > \epsilon \mid E') \\ &= \mathbb{P}(\exists k < n^a : \text{wdiam}(\text{MST}(\mathbb{K}_n[V(k, n)])) > \epsilon) + o(1), \end{aligned}$$

where we used that $\mathbb{P}(E') \rightarrow 1$ for the last equality. Now, using Theorem 2.3 similarly as before, we obtain that

$$\mathbb{P}(\exists k < n^a : \text{wdiam}(\text{MST}(\mathbb{K}_n[V(k, n)])) > \epsilon) \rightarrow 0.$$

The proof of this lemma now follows by combining (3.7), (3.8), and (3.9). \square

With the above lemmas in hand, the proof of Proposition 3.2 is routine.

Proof of Proposition 3.2. Fix $\epsilon > 0$ and let $W = \frac{1}{\log n}$ and $L = \lfloor \log \log n \rfloor$. Then

$$\begin{aligned} \mathbb{P}(\exists U \in \mathbb{U} : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon) &= \mathbb{P}(\exists U \in \mathbb{U}_r^- \cup \mathbb{U}_r^+ \cup \mathbb{U}_\ell : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon) \\ &\leq \mathbb{P}(\exists U \in \mathbb{U}_r^- : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon) \\ &\quad + \mathbb{P}(\exists U \in \mathbb{U}_r^+ : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon) \\ &\quad + \mathbb{P}(\exists U \in \mathbb{U}_\ell : \text{wdiam}(\text{MST}(\mathbb{K}_n[U])) > \epsilon), \end{aligned}$$

and the right hand side converges to 0 by Lemma 3.6, 3.7, and 3.8, proving the proposition. \square

4. CONCLUSION

4.1. More general weight distributions. The extension of Theorem 1.1 from Uniform[0, 1] to more general weight distributions is quite straightforward. Fix a probability density function $f : [0, \infty) \rightarrow [0, \infty)$, and let $\rho^* = \sup\{x : \int_0^x f(y)dy < 1\}$. Let $\mathbb{X}' = (X'_e, e \in E(K_n))$ be independent random variables with density f , and let $\mathbb{K}'_n = (K_n, \mathbb{X}')$.

Theorem 4.1. *Suppose that $f(0) > 0$, that f is continuous at zero, and that $\rho^* < \infty$. Fix any sequence $(H_n, n \geq 1)$ of connected graphs with H_n being a spanning subgraph of K_n . Then for any $\varepsilon > 0$, as $n \rightarrow \infty$,*

- (a) *with high probability there exists an MST sequence \mathbb{S} for (\mathbb{K}'_n, H_n) with $\text{wt}(\mathbb{S}) \leq \rho^* + \varepsilon$, and*
- (b) *there exists $\delta > 0$ such that with high probability, given any optimizing sequence $\mathbb{S} = (S_1, \dots, S_m)$ for (\mathbb{K}'_n, H_n) with $\text{wt}(\mathbb{S}) \leq \rho^* - \varepsilon$, the final spanning subgraph $H_{n,m}$ has weight $w(H_{n,m}) \geq \delta n w(\text{MST}(\mathbb{K}_n))$.*

In particular, $\text{cost}(\mathbb{K}'_n, H_n) \xrightarrow{\mathbb{P}} \rho^*$ as $n \rightarrow \infty$.

The proof is very similar to that of Theorem 1.1, so we only describe the changes that are required to prove the more general version.

The proof of the lower bound, part (b), proceeds just as in the case of Uniform[0, 1] edge weights: for any $\varepsilon > 0$, any optimizing sequence $\mathbb{S} = (S_0, \dots, S_m)$ for (\mathbb{K}'_n, H_n) with $\text{wt}(\mathbb{S}) \leq \rho^* - \varepsilon$ leaves edges of weight greater than $\rho^* - \varepsilon$ untouched, so all such edges appear in the final subgraph $H_{n,m}$. The number of such edges is Binomial($|E(H_n)|, \int_{\rho^* - \varepsilon}^{\rho^*} f(x)dx$)-distributed, so with high probability there are a linear number of such edges. On the other hand, $w(\text{MST } \mathbb{K}_n) \rightarrow \zeta(3)/f(0)$ in probability [7], and the lower bound follows.

For the upper bound, note that the bounds on the total cost of the optimizing sequences we construct essentially all have the form $A+B$ where A is the greatest weight of a single edge, and B is the weighted diameter of the minimum spanning tree of some subgraph of K_n . In order to prove Theorem 1.1, we used that $A \leq 1$, and proved using Theorem 2.3 and Proposition 3.2 that we could take B as close to zero as we wished (by a careful choice of optimizing sequence). For the edge weights \mathbb{X}' , we can simply replace the bound $A \leq 1$ by the bound $A \leq \rho^*$. To show that we can make B as close to zero as we like, we can carry through the same proof as in the Uniform[0, 1] case, provided that versions of Theorem 2.3 and Proposition 3.2 are still available to us.

To see that Theorem 2.3 and Proposition 3.2 do essentially carry over to the setting of $\mathbb{K}'_n = (K_n, \mathbb{X}')$, we make use of the following coupling. For $t \in [0, \rho^*]$ let $g(t) = \mathbb{P}(X' \leq t)$, so that $g(X')$ is Uniform[0, 1]-distributed. We can thus couple the random weights \mathbb{X}' to independent Uniform[0, 1] weights $\mathbb{X} = (X_e, e \in E(K_n))$ by taking $X_e = g(X'_e)$, and thereby couple $\mathbb{K}'_n = (K_n, \mathbb{X}')$ to $\mathbb{K}_n = (K_n, \mathbb{X})$. The edge weights $\mathbb{X}' = (X'_e, e \in E(K_n))$ are almost surely pairwise distinct, and on this event, the ordering of $E(K_n)$ in increasing order of weight is the same for the weights \mathbb{X} and \mathbb{X}' and thus $\text{MST}(\mathbb{K}'_n) = \text{MST}(\mathbb{K}_n)$.

Since $f(0) > 0$ and f is continuous, for all u sufficiently small we have $f(u) > f(0)/2$ and $g(u) \geq uf(0)/2$. It follows in particular that if $X_e \leq uf(0)/2$ then $X'_e \leq 2X_e/f(0) \leq u$. This observation implies that, under the above coupling between \mathbb{K}_n and \mathbb{K}'_n , if $\text{wdiam}(\text{MST}(\mathbb{K}_n)) \leq uf(0)/2$ then $\text{wdiam}(\text{MST}(\mathbb{K}'_n)) \leq u$, and Theorem 2.3 thus yields that for all n sufficiently large,

$$(4.1) \quad \mathbb{P}\left(\text{wdiam}(\text{MST}(\mathbb{K}'_n)) \geq \frac{2}{f(0)} \frac{7 \log^4 n}{n^{1/10}}\right) \leq \mathbb{P}\left(\text{wdiam}(\text{MST}(\mathbb{K}_n)) \geq \frac{7 \log^4 n}{n^{1/10}}\right) \leq \frac{4}{n^{\log n}}.$$

Similarly, Proposition 3.2 implies that (in the notation of that proposition), for all $\varepsilon > 0$

$$\mathbb{P}\left(\exists U \in \mathbb{U} : \text{wdiam}(\text{MST}(\mathbb{K}'_n[U])) > 2\varepsilon/f(0)\right) \longrightarrow 0$$

as $n \rightarrow \infty$. But since $\varepsilon > 0$ was arbitrary, this implies that also

$$(4.2) \quad \mathbb{P}\left(\exists U \in \mathbb{U} : \text{wdiam}(\text{MST}(\mathbb{K}'_n[U])) > \varepsilon\right) \longrightarrow 0$$

for all $\varepsilon > 0$.

All the remaining ingredients of the proof of Theorem 1.1 use only information about the graph-theoretic structure of $\text{MST}(\mathbb{K}_n)$, not its weights, and so carry over to the setting of non-uniform weights (using the fact that $\text{MST}(\mathbb{K}'_n)$ and $\text{MST}(\mathbb{K}_n)$ have the same distributions as unweighted graphs – indeed, they are equal under the above coupling). By running the proof of Theorem 1.1 but replacing all expressions of the form $1 + \text{wdiam}(F)$ by $\rho^* + \text{wdiam}(F)$, and when needed invoking (4.1) and (4.2) in place of Theorem 2.3 and Proposition 3.2, respectively, we obtain Theorem 4.1.

Before concluding this subsection, we note that if $\rho^* = \infty$ then for any $r > 0$, the probability that at least one edge of H_n has weight at least r tends to 1, so $\mathbb{P}(\text{cost}(\mathbb{K}'_n, H_n) > r) \rightarrow 1$ as $n \rightarrow \infty$. Thus, in this case we also have $\text{cost}(\mathbb{K}'_n, H_n) \xrightarrow{\mathbb{P}} \rho^*$.

4.2. Open questions and future directions. This work introduces the notion of local minimum spanning tree searches and proves a weak law of large numbers for the cost of such local searches. The framework naturally suggests several directions for future research, some of which we now highlight.

- Our main results concern low-weight MST sequences \mathbb{S} for randomly weighted complete graphs, where $\text{wt}(\mathbb{S})$ is measured in the L_∞ sense: it is the maximum weight of any single step of the optimizing sequence. However, one may wish to vary the norm used to measure the weights of optimizing sequences. The other L_p norms are natural alternatives, and correspond to studying the values

$$\text{cost}_p(\mathbb{G}, H) = \min \left\{ \text{wt}_p(\mathbb{S}) : \mathbb{S} \text{ is an MST sequence for } (\mathbb{G}, H) \right\}$$

where

$$\text{wt}_p(\mathbb{S}) = \left(\sum_{i=1}^m (\text{wt}(\mathbb{S}, i))^p \right)^{\frac{1}{p}}$$

is the L^p norm of $(\text{wt}(\mathbb{S}, i), 1 \leq i \leq m)$.

At first sight, using L^1 weights may seem very natural, as it corresponds to the total weight of all the subgraphs modified by the sequence. Mathematically, however, in the setting considered in this paper the L^1 cost is quite easy to understand. Indeed, for \mathbb{K}_n and H_n as in Theorem 1.1, by considering the sequence $\mathbb{S} = ([n])$ which simply replaces H_n by $\text{MST}(\mathbb{K}_n)$ in one step, we obtain that

$$\text{cost}_1(\mathbb{K}_n, H_n) \leq w(H_n).$$

Conversely, since any edge of $e \in E(H_n) \setminus E(\text{MST}(\mathbb{K}_n))$ must be removed in order to form the MST, for any MST sequence $\mathbb{S} = (S_1, \dots, S_m)$, there must exist $i \in [m]$ such that $e \in E(H_{n,i-1}[S_i])$. This implies that

$$\text{wt}_1(\mathbb{S}) \geq \sum_{i \in [m]} w(H_{n,i-1}[S_i]) \geq \sum_{e \in E(H_n) \setminus E(\text{MST}(\mathbb{K}_n))} X_e = (1 + o_{\mathbb{P}}(1))w(H_n),$$

where the final asymptotic follows from the fact that H_n is chosen independently of \mathbb{X} and that any fixed edge belongs to the MST with probability $(n-1)/\binom{n}{2} = o_{\mathbb{P}}(1)$. Since the lower bound $\sum_{e \in E(H_n) \setminus E(\text{MST}(\mathbb{K}_n))} X_e$ does not depend on the choice of MST sequence \mathbb{S} , it is also a lower bound on $\text{cost}_1(\mathbb{K}_n, H_n)$, and thus

$$\frac{\text{cost}_1(\mathbb{K}_n, H_n)}{w(H_n)} \rightarrow 1$$

in probability. When $p < 1$, this argument can be adapted to prove the same convergence result for $\text{cost}_p(\mathbb{K}_n, H_n)$. However, when $p > 1$ it is less clear what behaviour to expect, and in particular it is unclear whether the dependence on the initial spanning subgraph H_n will play a more complicated role.

- Another natural modification of the setting is to measure the cost of a step by the *size*, rather than the weight, of the subgraph which is replaced by its MST. That is, we may define

$$\text{wt}'(\mathbb{G}, H, \mathbb{S}) := \max \left\{ |E(H_{i-1}[S_i])| : 1 \leq i \leq m \right\},$$

and study

$$\text{cost}'(\mathbb{G}, H) = \min \left\{ \text{wt}'(\mathbb{G}, H, \mathbb{S}) : \mathbb{S} \text{ is an MST sequence for } (\mathbb{G}, H) \right\}$$

For this notion of cost, even the behaviour of $\text{cost}'(\mathbb{K}_n, K_n)$ is unclear to us; how $\text{cost}'(\mathbb{K}_n, H_n)$ will depend on the starting graph H_n is likewise unclear. However, at a minimum we expect that $\text{cost}'(\mathbb{K}_n, H_n) \rightarrow \infty$ in probability, provided that the initial spanning subgraphs H_n are chosen independently of the weights.

- Our result proves the *existence* of MST sequences of weight at most $(\rho^* + \varepsilon)$, with high probability. However, our construction does not yield insight into the ubiquity of such sequences, and it would be interesting to know whether low-weight MST sequences can be found easily and without using “non-local” information. For example, suppose that at each step we choose a subgraph to optimize uniformly at random over all subgraphs of weight at most w . For which values of w will the resulting sequence be an MST sequence with high probability?
- What is the asymptotic behaviour of $\text{cost}(\mathbb{K}_n, H_n) - \rho^*$? In particular, is there a sequence a_n such that $a_n(\text{cost}(\mathbb{K}_n, H_n) - \rho^*)$ converges in distribution to a non-trivial random variable?
- What happens if (K_n, \mathbb{X}_n) is replaced by a different fixed connected, weighted graph $\mathbb{G}_n = (G_n, \mathbb{X}_n)$? How does the asymptotic behaviour of $\text{cost}(\mathbb{G}_n, H_n)$ depend on G_n ?
- What happens if the iid structure of the edge weights of \mathbb{K}_n is modified? For example, one might generate \mathbb{X}_n by first taking n independent, uniformly random points $P_1, \dots, P_n \in [0, 1]^d$, then letting $X_{ij} = |P_i - P_j|$ be the Euclidean distance between i and j .

APPENDIX A. BOUNDS ON THE WEIGHTED DIAMETER

In this section, we prove Theorem 2.3. The proof exploits Kruskal’s algorithm for constructing minimum spanning trees. We first recall a very useful connection between Kruskal’s algorithm run on the complete graph with independent Uniform[0, 1] edge weights $\mathbb{X} = (X_e, e \in E(K_n))$ and the Erdős-Rényi random graph process. In this setting, Kruskal’s algorithm may be phrased as follows. Write $N = \binom{n}{2}$

- Order the edges of $E(\mathbb{K}_n)$ in increasing order of weight as e_1, \dots, e_N .
- Let $F_0 = ([n], \emptyset)$ be the forest with vertex set n and no edges.
- For $1 \leq i \leq N$, if e_i joins distinct connected components of F_{i-1} then let $E(F_i) = E(F_{i-1}) \cup \{e_i\}$; otherwise let $F_i = F_{i-1}$.

The final forest F_N is $\text{MST}(\mathbb{K}_n)$.

The Erdős-Rényi random graph process can be described very similarly:

- Order the edges of $E(\mathbb{K}_n)$ in increasing order of weight as e_1, \dots, e_N .
- Let $G_0 = ([n], \emptyset)$ be the graph with vertex set n and no edges.
- For $1 \leq i \leq N$, let $E(G_i) = E(G_{i-1}) \cup \{e_i\}$.

It is straightforward to see by induction that F_i and G_i always have the same connected components and, more strongly, that F_i is the minimum spanning forest of G_i (in that each tree of F_i is the minimum spanning tree of the corresponding connected component of G_i).

We also take $G(n, p)$ to be the subgraph of \mathbb{K}_n with edge set $\{e \in E(\mathbb{K}_n) : X_e \leq p\}$. Since we ordered the edges in increasing order of weight as e_1, \dots, e_N , the edge set of $G(n, p)$ is thus $\{e_1, \dots, e_m\}$, where $m = m(p)$ is maximal so that $X_{e_m} \leq p$. We likewise let $F(n, p)$ be the subgraph of F_N consisting of all edges of F_N with weight at most p , and note that $F(n, p) = F_{m(p)}$.

With this coupling in hand, we next explain our approach to bounding the weighted diameter of $\text{MST}(\mathbb{K}_n)$. Our bound has two parts. Fix $p \in (0, 1)$, and let $T_{n,p}^{\max}$ be the largest connected component of $F(n, p)$, with ties broken lexicographically. Note that $T_{n,p}^{\max}$ is a subgraph of $\text{MST}(\mathbb{K}_n)$. Further

write $L_{n,p}$ for the greatest number of edges in any path of $\text{MST}(\mathbb{K}_n)$ which has exactly one vertex lying in $T_{n,p}^{\max}$. Finally, write W_n for the greatest weight of any edge of $\text{MST}(\mathbb{K}_n)$.

Proposition A.1. *For any $p \in (0, 1)$,*

$$\text{wdiam}(\text{MST}(\mathbb{K}_n)) \leq p(|T_{n,p}^{\max}| - 1) + 2W_n L_{n,p}.$$

Proof. Fix any path P in $\text{MST}(\mathbb{K}_n)$. Then the set of vertices of P contained in $T_{n,p}^{\max}$ form a subpath of P , since otherwise $\text{MST}(\mathbb{K}_n)$ would contain a cycle; call this subpath P_0 . Then P_0 contains at most $|T_{n,p}^{\max}|$ vertices, so at most $|T_{n,p}^{\max}| - 1$ edges, and each such edge has weight at most p . Moreover, the edges of P not lying in P_0 form at most two subpaths of P . Each of these subpaths has at most $L_{n,p}$ edges, so the number of edges of P which are not edges of P_0 is at most $2L_{n,p}$; and the edges of P which are not edges of P_0 all have weight at most W_n . \square

To exploit this bound and prove Theorem 2.3, we must bound $|T_{n,p}^{\max}|$ and $L_{n,p}$, for some well chosen value of p , and bound W_n . The latter bound is the easiest, and we take care of it first. We will need the following bound on the probability of connectedness of $G(n, p)$. We believe we have seen this bound in the literature, but were unable to find a reference, so we have included its short proof.

Lemma A.2. *Let $G \sim G(n, p)$. Then*

$$\mathbb{P}(G \text{ is not connected}) \leq e^{ne^{-\frac{np}{2}}} - 1$$

Proof. Let S be a subset of $[n]$ such that $S \neq \emptyset$ and $S \neq [n]$. Then

$$\mathbb{P}(S \text{ is not connected to } S^c \text{ in } G) = (1-p)^{|S|(n-|S|)}.$$

This implies that

$$\begin{aligned} \mathbb{P}(G \text{ is not connected}) &= \mathbb{P}(\exists S \subseteq [n] : 1 \leq |S| \leq n/2 \text{ and } S \text{ is not connected to } S^c \text{ in } G) \\ &\leq \sum_{S \subseteq [n] : 1 \leq |S| \leq n/2} \mathbb{P}(S \text{ is not connected to } S^c \text{ in } G). \end{aligned}$$

Combined with the previous result, this leads to

$$\mathbb{P}(G \text{ is not connected}) = \sum_{S \subseteq [n] : 1 \leq |S| \leq n/2} (1-p)^{|S|(n-|S|)} \leq \sum_{1 \leq k \leq n/2} \binom{n}{k} (1-p)^{k(n-k)}.$$

Use now that $(n-k) \geq n/2$ along with the fact that $1-p \geq 0$ to obtain that

$$\mathbb{P}(G \text{ is not connected}) \leq \sum_{1 \leq k \leq n} \binom{n}{k} (1-p)^{kn/2} = (1 + (1-p)^{n/2})^n - 1.$$

Finally, by using twice the convexity of exponential, we have

$$\mathbb{P}(G \text{ is not connected}) \leq \left(1 + e^{-\frac{pn}{2}}\right)^n - 1 \leq e^{ne^{-\frac{pn}{2}}} - 1,$$

which is the desired result. \square

Fact A.3. *For all n sufficiently large, it holds that $\mathbb{P}(W_n > 3 \log^2 n/n) \leq 1/n^{\log n}$.*

Proof. Under the above coupling, $F(n, p)$ and $G(n, p)$ have the same connected components, so

$$\mathbb{P}(W_n > 3 \log^2 n/n) = \mathbb{P}(F(n, 3 \log^2 n/n) \text{ is not connected}) = \mathbb{P}(G(n, 3 \log^2 n/n) \text{ is not connected}).$$

Use now the bound from Lemma A.2 to obtain that

$$\mathbb{P}(G(n, 3 \log^2 n/n) \text{ is not connected}) \leq \exp(ne^{-(3/2)\log^2 n}) - 1 = e^{n/(n^{\log n})^{3/2}} - 1 \leq 1/n^{\log n}$$

the final bound holding for all n sufficiently large. \square

Proof of Theorem 2.3. We prove the theorem by bounding $|T_{n,p}^{\max}|$ and $L_{n,p}$, for a carefully chosen value of p (spoiler: we will take $p = 1/n + 1/n^{11/10}$), then applying Proposition A.1. Our arguments lean heavily on results from [2], and we next introduce those results (and the terminology necessary to do so).

For $c > 0$, let $\alpha(c)$ be the largest real solution of $e^{-cx} = 1 - x$ (the quantity $\alpha(c)$ is the survival probability of a Poisson(c) branching process). The key to the proof is the fact that the size of the largest component of $G(n, p)$ is with high probability close to $n\alpha(np)$ when $p = (1 + o(1))/n$. We now provide a precise and quantitative version of this statement, with error bounds.

By [1, Exercise 21 (d)], for $\varepsilon \geq 0$ we have

$$2\varepsilon(1 - o(1)) \leq \alpha(1 + \varepsilon) \leq 2\varepsilon,$$

the first inequality holding as $\varepsilon \rightarrow 0$. In particular,

$$(A.1) \quad (3/2)\varepsilon \leq \alpha(1 + \varepsilon) \leq 2\varepsilon$$

for all $\varepsilon \geq 0$ sufficiently small.

For the remainder of the proof, fix $p = 1/n + 1/n^{11/10}$ and write $s^+ = n\alpha(n \log(1/(1-p))) + n^{3/4}$ and $s^- = n\alpha(n \log(1/(1-p))) - 2n^{3/4}$. (Aside: for the careful reader who is verifying the connections to the results from [1], note that $s^+ = t^+$ but $s^- \neq t^-$, where t^+, t^- are defined in [1, Proof of Theorem 4.4, Case 2]). By [1, Exercise 23 (a)], for all n sufficiently large we have

$$n\alpha(np) \leq n\alpha(n \log(1/(1-p))) \leq n\alpha(np) + \frac{2n^{1/2}}{1-p},$$

and using the above bounds on α , this yields

$$n^{9/10} \leq s^- \leq s^+ \leq 3n^{9/10},$$

for n sufficiently large.

Let \mathcal{C}^{\max} be the largest connected component of $G(n, p)$, and let $\mathcal{C}^{\text{runnerup}}$ be its second largest component. Using the previous inequality on s^- and s^+ , by [1, (4.7)] we have

$$(A.2) \quad \mathbb{P}(|\mathcal{C}^{\max}| \geq 3n^{9/10}) \leq \mathbb{P}(|\mathcal{C}^{\max}| \geq s^+) \leq ne^{-(25/2)n^{1/10}};$$

moreover, by [1, (4.10)], we have

$$(A.3) \quad \mathbb{P}(|\mathcal{C}^{\max}| \leq n^{9/10}) \leq \mathbb{P}(|\mathcal{C}^{\max}| \leq s^-) \leq 2ne^{-(25/2)n^{1/10}};$$

finally, by [1, (4.10) and (4.11)], we have

$$(A.4) \quad \mathbb{P}(|\mathcal{C}^{\text{runnerup}}| \geq n^{4/5}) \leq 5ne^{-(25/2)n^{1/10}}.$$

Furthermore, under the coupling between $G(n, p)$ and $F(n, p)$, we have $|\mathcal{C}^{\max}| = |T_{n,p}^{\max}|$, so (A.2) immediately gives us that for all n sufficiently large,

$$(A.5) \quad \mathbb{P}(|T_{n,p}^{\max}| \geq 3n^{9/10}) \leq ne^{-(25/2)n^{1/10}}.$$

It remains to bound $L_{n,p}$. For this, we use (A.3) and a Prim's-algorithm-type construction to control the greatest number of connected components of $G(n, p)$ that any path of $\text{MST}(\mathbb{K}_n)$ lying outside $T_{n,p}^{\max}$ passes through, and use (A.4) to bound the size of those components.

Condition on the graph $G(n, p)$, and fix a connected component \mathcal{C}_1 of $G(n, p)$ different from \mathcal{C}^{\max} . Let $f_1 = u_1v_1$ be the smallest-weight edge with exactly one endpoint in \mathcal{C}_1 , and let p_1 be its weight. Then $p_1 > p$, and f_1 is a cut-edge of $G(n, p_1)$. It follows that f_1 is an edge of $\text{MST}(\mathbb{K}_n)$. Moreover, by the exchangeability of the edge weights, the endpoint v_1 of f_1 not lying in \mathcal{C}_1 is uniformly distributed over the remainder of the vertices, so

$$\mathbb{P}\left(v_1 \notin \mathcal{C}^{\max} \mid G(n, p)\right) \leq 1 - \frac{|\mathcal{C}^{\max}|}{n - |\mathcal{C}_1|} < 1 - \frac{|\mathcal{C}^{\max}|}{n}.$$

If v_1 is not in \mathcal{C}^{\max} , then it lies in another connected component \mathcal{C}_2 . Let $f_2 = u_2v_2$ be the smallest-weight edge leaving $\mathcal{C}_1 \cup \mathcal{C}_2$, and let p_2 be its weight. Then f_2 is an edge of $\text{MST}(\mathbb{K}_n)$; to see this, note that any path γ connecting u_2 and v_2 which is not just the edge f_2 contains some edge e of

weight strictly greater than p_2 , meaning that f_2 is never the heaviest edge of any cycle. Moreover, the endpoint v_2 of f_2 not lying in $\mathcal{C}_1 \cup \mathcal{C}_2$ is uniformly distributed over the remainder of the graph, so once again

$$\mathbb{P}\left(v_2 \notin \mathcal{C}^{\max} \mid G(n, p), v_1 \notin \mathcal{C}^{\max}\right) < 1 - \frac{|\mathcal{C}^{\max}|}{n}.$$

Continuing this process, we construct a sequence $\mathcal{C}_1, \dots, \mathcal{C}_K$ of distinct connected components of $G(n, p)$ and a sequence f_1, \dots, f_K of edges of $\text{MST}(\mathbb{K}_n)$, where where $f_i = u_i v_i$ is the smallest-weight edge from $\mathcal{C}_1 \cup \dots \cup \mathcal{C}_i$ to the remainder of the graph, $\mathcal{C}_1, \dots, \mathcal{C}_K$ are all connected components of $G(n, p)$ different from \mathcal{C}^{\max} , and $v_K \in \mathcal{C}^{\max}$. To bound the length K of the sequences, we use that at each step of the construction, the conditional probability that $f_j = u_i v_j$ has an endpoint in \mathcal{C}^{\max} given $G(n, p)$ and given that e_1, \dots, e_{i-1} do not have an endpoint in \mathcal{C}^{\max} , is greater than $|\mathcal{C}^{\max}|/n$, and so

$$\begin{aligned} \mathbb{P}\left(K > k \mid G(n, p)\right) &= \mathbb{P}\left(v_k \notin \mathcal{C}^{\max} \mid G(n, p)\right) \\ &= \prod_{i=1}^k \mathbb{P}\left(v_i \notin \mathcal{C}^{\max} \mid G(n, p), v_1, \dots, v_{i-1} \notin \mathcal{C}^{\max}\right) \\ &\leq \left(1 - \frac{|\mathcal{C}^{\max}|}{n}\right)^k \end{aligned}$$

Note now that any path in $\text{MST}(\mathbb{K}_n)$ with one endpoint in \mathcal{C}_1 and the other endpoint in \mathcal{C}^{\max} passes through $\mathcal{C}_1, \dots, \mathcal{C}_K$ and edges f_1, \dots, f_K . Since each of the components $\mathcal{C}_1, \dots, \mathcal{C}_K$ has size at most that of $\mathcal{C}^{\text{runnerup}}$, it follows that the greatest number of edges in any path with one endpoint in \mathcal{C}_1 which only intersects \mathcal{C}^{\max} in one vertex is at most $K|\mathcal{C}^{\text{runnerup}}|$. Taking a union bound over the possible choices for \mathcal{C}_1 among all components of $G(n, p)$ different from \mathcal{C}^{\max} (there are less than n of them), it follows that

$$\begin{aligned} \mathbb{P}\left(L_{n,p} > k \mid |\mathcal{C}^{\text{runnerup}}| \mid G(n, p)\right) &\leq n \mathbb{P}\left(K > k \mid G(n, p)\right) \\ &\leq n \left(1 - \frac{|\mathcal{C}^{\max}|}{n}\right)^k. \end{aligned}$$

Recall now the tail bounds for $|\mathcal{C}^{\max}|$ and $|\mathcal{C}^{\text{runnerup}}|$ from (A.3) and (A.4) and use that

$$\begin{aligned} &\mathbb{P}\left(L_{n,p} > n^{9/10} \log^2 n \mid G(n, p), |\mathcal{C}^{\max}| > n^{9/10}, |\mathcal{C}^{\text{runnerup}}| < n^{4/5}\right) \\ &= \mathbb{P}\left(L_{n,p} > (n^{1/10} \log^2 n) \cdot n^{4/5} \mid G(n, p), |\mathcal{C}^{\max}| > n^{9/10}, |\mathcal{C}^{\text{runnerup}}| < n^{4/5}\right) \\ &\leq n \left(1 - \frac{n^{9/10}}{n}\right)^{n^{1/10} \log^2 n} \leq n e^{-\log^2 n} \end{aligned}$$

to obtain

$$(A.6) \quad \mathbb{P}\left(L_{n,p} > n^{9/10} \log^2 n\right) \leq 7n e^{-(25/2)n^{1/10}} + n e^{-\log^2 n} \leq \frac{2}{n^{\log n}},$$

the last bound holding for n large enough.

We can now conclude the proof of Theorem 2.3. By Fact A.3, for n sufficiently large, $\mathbb{P}(W_n > 3 \log^2 n/n) \leq 1/n^{\log n}$. Combined with (A.6), this implies that

$$\mathbb{P}\left(2W_n L_{n,p} > \frac{6 \log^4 n}{n^{1/10}}\right) \leq \frac{3}{n^{\log n}}.$$

Using the bound of Proposition A.1 and combining it with the previous inequality and (A.5), we obtain that

$$\begin{aligned} \mathbb{P}\left(\text{wdiam}(\text{MST}(\mathbb{K}_n)) > 3pn^{9/10} + \frac{6 \log^4 n}{n^{1/10}}\right) &\leq \mathbb{P}\left(|T_{n,p}^{\max}| \geq 3n^{9/10}\right) + \mathbb{P}\left(2W_n L_{n,p} > \frac{6 \log^4 n}{n^{1/10}}\right) \\ &\leq ne^{-(25/2)n^{1/10}} + \frac{3}{n^{\log n}} \leq \frac{4}{n^{\log n}}, \end{aligned}$$

the last inequality holding when n is large. Finally, since $p = 1/n + 1/n^{11/10}$, for n large we have $3pn^{9/10} + \frac{6 \log^4 n}{n^{1/10}} < \frac{7 \log^4 n}{n^{1/10}}$, so the bound of Theorem 2.3 follows. \square

ACKNOWLEDGEMENTS

This work was initiated at the Bellairs Research Institute of McGill University. The authors thank Gábor Lugosi, who initially proposed a version of the problem studied in this work, at a Bellairs workshop on probability and combinatorics.

REFERENCES

- [1] L. ADDARIO-BERRY, *Partition functions of discrete coalescents: from Cayley's formula to Frieze's $\zeta(3)$ limit theorem*, in XI Symposium on Probability and Stochastic Processes, Springer, 2015, pp. 1–45.
- [2] L. ADDARIO-BERRY, N. BROUTIN, AND B. REED, *The diameter of the minimum spanning tree of a complete graph*, in Discrete Mathematics and Theoretical Computer Science, Discrete Mathematics and Theoretical Computer Science, 2006, pp. 237–248.
- [3] O. ANGEL, S. BUBECK, Y. PERES, AND F. WEI, *Local max-cut in smoothed polynomial time*, in STOC'17—Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, ACM, New York, 2017, pp. 429–437.
- [4] A. BIBAK, C. CARLSON, AND K. CHANDRASEKARAN, *Improving the smoothed complexity of FLIP for max cut problems*, ACM Trans. Algorithms, 17 (2021), pp. Art. 19, 38.
- [5] X. CHEN, C. GUO, E. V. VLATAKIS-GKARAGKOUNIS, M. YANNAKAKIS, AND X. ZHANG, *Smoothed Complexity of Local Max-Cut and Binary Max-CSP*, Association for Computing Machinery, New York, NY, USA, 2020, p. 1052–1065.
- [6] M. ETSCHIED AND H. RÖGLIN, *Smoothed analysis of local search for the maximum-cut problem*, ACM Trans. Algorithms, 13 (2017).
- [7] A. M. FRIEZE, *On the value of a random minimum spanning tree problem*, Discrete Applied Mathematics, 10 (1985), pp. 47–56.
- [8] B. MANTHEY, *Smoothed analysis of local search algorithms*, in Algorithms and Data Structures, F. Dehne, J.-R. Sack, and U. Stege, eds., Cham, 2015, Springer International Publishing, pp. 518–527.

DEPARTMENT OF MATHEMATICS AND STATISTICS, MCGILL UNIVERSITY, MONTRÉAL, CANADA

Email address: `louigi.addario@mcgill.ca`

Email address: `jordan.barrett@mail.mcgill.ca`

Email address: `benoitcorsini@gmail.com`