



Successive shortest paths in complete graphs with random edge weights

LSE Research Online URL for this paper: <http://eprints.lse.ac.uk/106487/>

Version: Published Version

Article:

Gerke, Stefanie, Mezei, Balazs F. and Sorkin, Gregory ORCID: 0000-0003-4935-7820 (2020) Successive shortest paths in complete graphs with random edge weights. *Random Structures and Algorithms*, 57 (4). 1205 - 1247. ISSN 1042-9832

<https://doi.org/10.1002/rsa.20962>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>



RESEARCH ARTICLE

WILEY

Successive shortest paths in complete graphs with random edge weights

Stefanie Gerke¹ | Balázs F. Mezei^{1,2} | Gregory B. Sorkin³

¹Department of Mathematics, Royal Holloway University of London, Egham, United Kingdom

²Department of Computer Science, University of Oxford, Wolfson Building, Oxford, United Kingdom

³Department of Mathematics, The London School of Economics and Political Science, London, United Kingdom

Correspondence

Gregory B. Sorkin, Department of Mathematics, The London School of Economics and Political Science, Houghton Street, London WC2A 2AE, United Kingdom.
Email: g.b.sorkin@lse.ac.uk

Abstract

Consider a complete graph K_n with edge weights drawn independently from a uniform distribution $U(0, 1)$. The weight of the shortest (minimum-weight) path P_1 between two given vertices is known to be $\ln n/n$, asymptotically. Define a second-shortest path P_2 to be the shortest path edge-disjoint from P_1 , and consider more generally the shortest path P_k edge-disjoint from all earlier paths. We show that the cost X_k of P_k converges in probability to $2k/n + \ln n/n$ uniformly for all $k \leq n - 1$. We show analogous results when the edge weights are drawn from an exponential distribution. The same results characterize the collectively cheapest k edge-disjoint paths, that is, a minimum-cost k -flow. We also obtain the expectation of X_k conditioned on the existence of P_k .

KEYWORDS

Dijkstra's algorithm, minimum-cost flow, optimization in random structures, robust optimization, second-cheapest structure, shortest path

1 | INTRODUCTION

It is a standard problem to find the shortest s - t path in a graph, that is, the cheapest path P_1 between specified vertices s and t , and its cost X_1 , where the cost of a path is the sum of the costs of its edges. We will use the terms “cost” and “weight” interchangeably, and reserve “length” for the number of edges in a path.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2020 The Authors. *Random Structures & Algorithms* published by Wiley Periodicals LLC.

Consider the complete graph $G = K_n$ with each edge $\{u, v\}$ having weight $w(u, v)$, where the $w(u, v)$ are i.i.d. random variables with exponential distribution $\text{Exp}(1)$ or uniform distribution $U(0, 1)$ (we consider both versions). In this random setting, a well-known result of Janson [14] is that as $n \rightarrow \infty$,

$$\frac{X_1}{\ln n/n} \xrightarrow{p} 1. \quad (1.1)$$

We define the second cheapest path P_2 , with cost X_2 , to be the cheapest s - t path edge-disjoint from P_1 , and in general define P_k , with cost X_k , to be the cheapest s - t path edge-disjoint from $P_1 \cup \dots \cup P_{k-1}$, provided such a path exists. We also think of this as finding path P_k after the preceding paths' edges have been removed. Our question is how the costs X_k behave in the limit as $n \rightarrow \infty$ (this limit is implicit throughout). Our main result is the following.

Theorem 1.1. *In the complete graph K_n with i.i.d. uniform $U(0, 1)$ edge weights, with X_k the cost of the k th cheapest path,*

$$\frac{X_k}{2k/n + \ln n/n} \xrightarrow{p} 1 \quad (1.2)$$

uniformly for all $k \leq n - 1$. That is, for any $\varepsilon > 0$, asymptotically almost surely, for every $k = 1, \dots, n - 1$,

$$1 - \varepsilon \leq \frac{X_k}{2k/n + \ln n/n} \leq 1 + \varepsilon. \quad (1.3)$$

Naturally, with $k = 1$, (1.2) recovers Janson's result (1.1), since $2/n = o(\ln n/n)$.

As discussed shortly, in contrast to many cases, the result for the uniform distribution does not extend immediately to all distributions with positive density at 0. However, we have a corresponding result for exponentially distributed edge weights. Given an edge-weight distribution, let $W_{(k)}$ be the (random) weight of the k th cheapest edge out of a vertex (the k th order statistic of $n - 1$ edge weights).

Theorem 1.2. *In the complete graph K_n with i.i.d. exponential edge weights with mean 1,*

$$\frac{X_k}{2\mathbb{E}W_{(k)} + \ln n/n} \xrightarrow{p} 1 \quad (1.4)$$

uniformly for all $k \leq n - 1$.

We give the guiding intuition behind the formula (1.4) in Section 1.1. Note that $\mathbb{E}W_{(k)} = \sum_{i=1}^k \frac{1}{n-i}$ in the exponential case (see e.g., Lemma 4.2). In the uniform case, $\mathbb{E}W_{(k)} = k/n$, so (1.2) in Theorem 1.1 can also be written as (1.4).

Rather than finding the k successive cheapest paths, we may alternatively wish to find the k edge-disjoint paths of *collective* minimum cost. Equivalently, where every edge of G has capacity 1, we may be interested in the minimum-cost k -flow from s to t in G . The following remark shows that this problem leads to essentially the same costs. (The analogous "collective" problem for minimum spanning trees [MSTs] is solved in [8], and [15] shows that for MSTs, the "successive" version leads to strictly larger costs.)

Remark 1.3. In the complete graph K_n with i.i.d. edge weights with distribution $U(0, 1)$ or exponential with mean 1, the minimum-cost k -flow has cost F_k satisfying

$$\frac{F_k}{\sum_{i=1}^k (2\mathbb{E}W_{(i)} + \ln n/n)} \xrightarrow{p} 1 \quad (1.5)$$

uniformly for all $k \leq n - 1$.

As in (1.3), the statement consists of high-probability upper and lower bounds. The upper bounds here, for the two models, follow immediately from the upper bounds of (1.2) and (1.4). The lower bounds follow from the lower bound on $S_k := \sum_{i=1}^k X_i$ (see (7.1)) in (7.4) and its analogue for the exponential case, as those bounds hold for any set of k edge-disjoint paths. (The main work in Section 7, not needed here, is to extract lower bounds on X_k from the lower bounds on S_k .)

Remark 1.4. P_k is always defined for all $k \leq n/2$, but, at least for n even, may be undefined for all $k > n/2$.

Proof. There are $n - 2$ length-2 s - t paths. Any path P_k can destroy (share an edge with) at most two such paths (since P_k uses just one edge incident to each of s and t). Also, the single-edge path $\{s, t\}$ is destroyed only by the path P_k consisting of just this edge. So, for P_1, \dots, P_k to destroy all length-1 and length-2 paths requires $k \geq (n - 2)/2 + 1 = n/2$, so for $k \leq n/2$, certainly path P_k exists.

Conversely, a construction described in 1892 by Lucas [17, pp. 162-164], which he attributes to Walecki, shows that a complete graph K_{2r} can be decomposed into r edge-disjoint Hamilton paths whose $2r$ terminals are all distinct. For n even, decompose $G = K_n \setminus \{s, t\}$ in this way, then link s to one “start” terminal of each such path and t to the other “end” terminal, giving $(n - 2)/2$ edge-disjoint s - t paths. The edge $\{s, t\}$ gives another path, for $n/2$ paths in all. The only edges not used by these paths are a star from s to the Hamilton paths’ end terminals, and another star from t to their start terminals, and as there are no other unused edges to connect these two stars, there is no further s - t path. With nonzero probability, the edge weights are such that $P_1, \dots, P_{n/2}$ are these $n/2$ paths, so that $P_{n/2+1}$ does not exist. ■

Remark 1.4 implies that, at least for n even, $\mathbb{E}[X_k]$ is undefined for $k > n/2$. The following theorem establishes $\mathbb{E}X_k$ for $k \leq n/2$, and for all $k \leq n - 1$, gives the expectation conditioned on the (high-probability) event that P_k exists.

Theorem 1.5. *In both the uniform and exponential models, for $k \leq n - 1$, a.a.s. P_k exists, and*

$$\mathbb{E}[X_k \mid P_k \text{ exists}] = (1 + o(1))(2\mathbb{E}W_{(k)} + \ln n/n), \tag{1.6}$$

uniformly in k .

For $k \leq n/2$, by Remark 1.4 the conditioning is null, so it is immediate from Theorem 1.5 that $E[X_k] = (1 + o(1))(2\mathbb{E}W_{(k)} + \ln n/n)$.

1.1 | Intuition

The intuitive picture is that path P_k should use the k th cheapest edges out of s and t , whose costs are denoted $W_{(k)}^s$ and $W_{(k)}^t$ respectively. Then, if we ignore previous paths’ use of other edges in $G \setminus \{s, t\}$, by (1.1) the opposite endpoints of these two edges should be connected by a path of cost about $\ln n/n$. This suggests that $X_k \leq W_{(k)}^s + W_{(k)}^t + \ln n/n$, and this is our guiding intuition. Obviously, the path P_k does not have to use the k th cheapest edge, its middle section may cost more or less than $\ln n/n$, and as earlier paths use up edges, the costs of these middle sections may rise. It is true, though, that $\sum_{i=1}^k X_i \geq \sum_{i=1}^{k-1} (W_{(i)}^s + W_{(i)}^t)$ (summing only to $k - 1$ on the right-hand side to avoid doubly counting

edge $\{s, t\}$, and we use this in proving the lower bounds on X_k (in Section 7 for uniform and Section 8.7 for exponential) and, more surprisingly, in proving the upper bounds on X_k for large k (in Section 5 generically, the details treated in Section 6 and 8).

Our upper bounds are obtained by reasoning as follows. Janson [14] analyzes the shortest s - t path, and shortest-path tree (SP tree or SPT) on s , in the randomly edge-weighted graph $G = K_n$, showing that the cost of P_1 is asymptotically almost surely, almost exactly $\ln n/n$. When the path P_1 is deleted, this prunes away a root-level branch of the SP tree. The SP tree is a uniform random tree, and using known properties of such trees (see e.g., [19]) it is not hard to show that what remains of the SP tree is likely to be large; capitalizing on this we can find an almost equally cheap path P'_2 . This line of argument also shows that there remains a cheap path after deleting P'_2 , but we need to know what happens when we delete the true second-shortest path P_2 , and at this point the argument fails because it gives no characterization of P_2 , only of P'_2 . We do know, however, that P_2 is cheap (no more expensive than P'_2), and of course uses just one edge incident to each of s and t , and we will show that deleting *any* edge set with these properties (including P_2 as a possibility) must still leave a cheap path P'_3 , and so forth. This “adversarial” deletion argument is developed in Section 3.2 to prove Theorem 1.1.

1.2 | Context

The question fits with a broad research theme on optimization (and satisfiability) problems on random structures. The novel element here is the “robustness” aspect of finding cheap structures even after the cheapest has been removed, and in this we were motivated by a recent study by Janson and Sorkin [15] of the same question for successive MSTs, again for K_n with uniform or exponential random edge weights. The results for shortest paths and MSTs are dramatically different. For MSTs, it is a celebrated result of Frieze [10] that as $n \rightarrow \infty$ the cost of the MST T_1 satisfies $w(T_1) \xrightarrow{p} \zeta(3) \stackrel{\text{def}}{=} \sum_{k=1}^{\infty} 1/k^3$, and [15] shows that each subsequent tree’s cost has $w(T_k) \xrightarrow{p} \gamma_k$ with the γ_k strictly increasing (and $2k - 2\sqrt{k} < \gamma_k < 2k + 2\sqrt{k}$). That is very different from the case here, for paths, where for $k = o(\ln n)$ we have X_k asymptotically equal to X_1 .

Further context is given in the discussion of open problems in Section 2.2.

1.3 | Edge weight distributions

As remarked earlier, in many contexts (including for the length X_1 of a shortest path) the result for any distribution with positive density at 0 follows immediately from that for the uniform distribution $U(0, 1)$, but that is not the case for the successive paths considered here.

Remark 1.6. Janson proves the X_1 case in the exponential model but provides standard “black-box” reasoning that it holds also for the uniform distribution, for any distribution with density 1 at 0 (i.e., with cumulative distribution function (CDF) $\mathbb{P}(X \leq x) = x + o(x)$ for $x \searrow 0$), and, after simple rescaling, for any distribution with positive density at 0. Simply, if there is a path of cost $o(1)$ in some such model, each edge w must also cost $o(1)$, and, coupling with the uniform distribution by replacing w with $w' = F(w)$, with F the CDF, $w' \leq (1 + o(1))w$, and thus the same path is similarly cheap in the uniform model. By the same token, if a path is cheap in any model, the same path has asymptotically the same cost in any other model, and thus the cheapest paths have asymptotically the same cost.

Remark 1.7. In our setting this argument does not apply: to find path P_k we must know the nature of the $k - 1$ previous paths; their costs are not enough. For $k = o(n)$, however, the standard argument

applies within our proofs, since the proofs rely only on edges of cost $o(1)$. However, for larger k there are genuine difficulties. Our argument for the exponential case, in Section 8, largely parallels that for uniform but requires new calculations for the upper bound, and one new idea for the lower bound (in Section 8.7). It is not clear for what other edge-weight distributions (even those with density 1 at 0) (1.4) will hold.

2 | OPEN PROBLEMS

2.1 | Poisson multigraph model

The issue of possible nonexistence of paths P_k for $k > n/2$ (see Remark 1.4) is obviated if, as in [15], we work in a Poisson multigraph model. Here, each pair of vertices $\{u, v\}$ of K_n is joined by infinitely many edges, whose weights are drawn from a Poisson process of rate 1 (so that the cheapest $\{u, v\}$ edge has exponentially distributed cost of mean 1). By construction, in this model every s - t path is always available (possibly at a higher cost).

Conjecture 2.1. *In the Poisson multigraph model, $\frac{X_k}{2k/n + \ln n/n} \xrightarrow{P} 1$ uniformly for all $k \leq n - 1$, and $\frac{\mathbb{E}X_k}{2k/n + \ln n/n} \rightarrow 1$ for all $k \leq n - 1$.*

Actually, in this model there is no need to stop at $k = n - 1$, but it is not clear how far out we can go (especially preserving uniform convergence).

2.2 | Other models

Most narrowly, it would be interesting to characterize successive shortest paths that are vertex-disjoint rather than edge-disjoint, and (in the style of Remark 1.3 for edge-disjoint paths) the k vertex-disjoint paths of collective minimum cost. In this model, guessing that path lengths stay around $\ln n$, we would expect P_k to be defined up to k about $n/\ln n$.

More broadly, it would be interesting to explore different edge-weight distributions, different structures, and different graphs.

As noted earlier, we have results for uniformly and exponentially distributed edge weights, but not for arbitrary distributions. As mentioned, results for the single shortest path follow by standard arguments for any distribution with positive density near 0. For a distribution with density tending to 0 or ∞ at 0, shortest paths were studied in [4]. In particular, they consider the case when edge weights are i.i.d. and have the same distribution as Z^p , where $Z \sim \text{Exp}(1)$ and $p > 0$ is a fixed parameter; in this setting, the shortest path has length $p \ln n$ and its cost is $\ln n/n^p$ times a p -dependent constant. A variant where the edge-weight distribution may depend on n is studied in [7].

To what distributions does Theorem 1.2 extend? Restricting to distributions with positive density near 0, the arguments in Section 8 should immediately extend for all $k = o(n)$. For larger k , the “middle” of each path should remain short, so the issue is the edges incident on s and t in P_k . Certainly (1.4) will fail if the order statistics of edges incident to s are not concentrated, for example if the edge distribution is a mixture of $U(0, 1)$ and an atom at 2 or (for a continuous example) a mixture of $U(0, 1)$ and the Pareto distribution with CDF $1 - 1/x$ for $x \geq 1$. It might be true that (1.4) holds more generally if the expectation $2\mathbb{E}W_{(k)}$ is replaced by $W_{(k)}^s + W_{(k)}^t$. However, to obtain the needed lower bound for the exponential model (see Section 8), we had to address the fact that the k th path does not necessarily use the edges of cost $W_{(k)}^s$ and $W_{(k)}^t$; we also needed exponential-specific calculations for the upper bound.




FIGURE 1 Cartoon of a robust subgraph R of G , showing the vertices s and t , their respective structures R_s and R_t including SPTs represented by triangles (some “failed” and thus not shown), and the cheap edges connecting triangles in R_s and R_t . Vertices s and t have down-degree (number of children) r_0 , and vertices at levels 1 and 2 (in R_s and R_t) have down-degrees r_1 and r_2 respectively

One could explore other structural models. MSTs have already been explored in [15] for the successive version and in [8] for the collective version. But for many other models the single cheapest structure is well studied but the successive and collective extensions have not been explored: this includes perfect matchings in complete bipartite graphs $K_{n,n}$ [2, 21], perfect matchings in complete graphs K_n [20], and Hamilton cycles (i.e., the traveling salesman problem) in K_n [22].

One could also consider graphs other than complete graphs, in the style of studies of the MST in a random regular graph [3], and of first-passage percolation in Erdős–Rényi random graphs [5] and hypercubes [18].

3 | UPPER BOUND FOR SMALL k

In this section we prove the upper bound of Theorem 1.1 for all $k = o(\sqrt{n})$; larger values are treated in the next section.

As discussed in the introduction, we can characterize the cheapest path P_1 and subsequent paths that are *cheap* but not necessarily *cheapest*, putting us at a loss to characterize what remains on deletion of a subsequent *cheapest* path. We address this in this section. Given k , we show a construction of a subgraph $R = R^{(k)}$ of G designed so that, as we will show in turn, its s – t paths are all cheap, and no deletion of edges from R subject to certain constraints can destroy all these paths. We show that the union of the k shortest paths satisfies these constraints, so that there remains a cheap s – t path in R and thus in G , and use this to prove Theorem 1.1.

Specifically, we will define a structure R , sketched in Figure 1, that has many cheap and spread-out paths between s and t , within which we will always find a cheap path. A crucial point is that each

step of the construction occurs in a complete induced subgraph of G of size $n - o(n)$ with all edges unconditioned.

We will show, assuming that

$$X_i \leq (1 + \varepsilon) \left(\frac{2i}{n} + \frac{\ln n}{n} \right) \quad (3.1)$$

for all $i \leq k$, that the same holds for $i = k + 1$. We will do so by showing that after deleting k paths, each of cost $\leq (1 + \varepsilon)(2k/n + \ln n/n)$ from G , some or all of whose edges may lie in R , there remains a path in R satisfying the same cost bound, and so this must also be true of P_{k+1} .

Consistent with this approach, and because to prove convergence in probability it suffices to consider an arbitrarily small, fixed ε (see around (1.3)), throughout this section we assume that $\varepsilon > 0$ is fixed. Thus, in the $n \rightarrow \infty$ limit implicit throughout,

$$\varepsilon = \Theta(1), \quad (3.2)$$

and ε (and functions of ε) may be absorbed into the constants implicit in any Landau-notation expression.

Remark 3.1. Most of the calculations below hold for any $\varepsilon > 0$, but a few (e.g., (3.7) and (3.8)) hold only for ε sufficiently small. This is not restrictive here, in proving convergence in probability, but to characterize expectation, Section 9.1 requires ε to be a large constant (to assure sufficiently small failure probabilities). The proof of Lemma 9.1 addresses the changes needed.

Before going into detail let us sketch the construction of R . We first build up a tree R_s on s , starting from s at level 0, the opposite endpoints of edges out of level i forming level $i + 1$. We will always choose “cheap” edges, but not always the cheapest ones, as explained later. From s we will choose $k + r_0$ cheap edges; from each of these $k + r_0$ level-1 vertices we choose r_1 cheap edges; from each of the $(k + r_0)r_1$ level-2 vertices we choose r_2 cheap edges; and on each of the $(k + r_0)r_1r_2$ level-3 vertices we construct a shortest-path tree comprising d vertices. We do a similar construction on t to form R_t . Finally, we link R_s and R_t using cheap edges between their SPTs. The values of the parameters r_0 , r_1 , r_2 , and d are given in (3.11), (3.13), (3.17), and (3.19), and it is confirmed in Section 3.9 that the construction uses only a small fraction of G 's vertices,

$$|V(R)| = O((k + r_0)r_1r_2d) = o(n), \quad (3.3)$$

a fact we rely on in the construction.

We will repeatedly use the following Chernoff bound, which in fact holds under more general conditions; see for example [12, Theorem 1, eq. (4)].

Lemma 3.2. *Let $X \sim \text{Bi}(n, p)$ be a binomial random variable with mean $\lambda = np$. Then for any $\varepsilon > 0$, $\mathbb{P}(X < (1 - \varepsilon)\lambda) \leq \exp(-\varepsilon^2\lambda/2)$.*

3.1 | Cheap paths are short

We show that, w.h.p., every cheap path in G is also short. The following lemma asserts the contrapositive. The result is used in (3.9) to restrict the number of edges the adversary can delete.

Lemma 3.3. *In both the uniform and exponential models, with probability $1 - O(n^{-1.9})$, simultaneously for all l with $\ln n \leq l < n$, every s - t path of length l has cost $\geq l/(19n)$.*

Proof. We start with the uniform distribution. Here, with $X = \sum_{i=1}^l X_i$, $X_i \sim U(0, 1)$ i.i.d., X has the Irwin-Hall distribution and it is a standard result that $\mathbb{P}(X \leq a) \leq a^l/l!$ (see e.g., [9, eq. 8]). Recall that Stirling's approximation is also a lower bound. Thus,

$$\mathbb{P}\left(X \leq \frac{l}{19n}\right) \leq \frac{(l/19n)^l}{l!} \leq \frac{(l/19n)^l}{\sqrt{2\pi l}(l/e)^l} < \left(\frac{e}{19n}\right)^l.$$

The cost of a fixed path of length l has the same law as X . Over the $\leq n^l$ choices for such a path, the number M_l of “cheap paths” (of cost $< l/(19n)$) satisfies (by Markov's inequality)

$$\mathbb{P}(M_l > 0) \leq \mathbb{E}M_l \leq n^l \mathbb{P}\left(X \leq \frac{l}{19n}\right) \leq n^l \left(\frac{e}{19n}\right)^l = \left(\frac{e}{19}\right)^l.$$

Summing over $l \geq \ln n$, the probability that there is a cheap path of any such length is $O((e/19)^{\ln n}) = O(n^{-1.9})$.

Since an $\text{Exp}(1)$ random weight X' can be obtained from a $U(0, 1)$ r.v. X by setting $X' = -\ln(1-X) > X$, the exponential weight stochastically dominates the uniform, so the result for uniform immediately implies that for exponential. ■

3.2 | Adversarial edge deletions

As noted in the introduction, we introduce an edge-deleting adversary whose powers allow it to delete the paths P_1, \dots, P_k , but which is more easily characterized than those paths are. We now specify what the adversary is permitted to do.

Let

$$s = s(k) := 2k + \ln n. \quad (3.4)$$

(From context it should be easy to distinguish this use of s from that as the source of an s - t path.) Let w_0 be the “target cost” of a path, namely

$$w_0 = w_0(k) := \frac{s}{n} = \frac{2k}{n} + \frac{\ln n}{n}. \quad (3.5)$$

Define a “heavy” edge to be one of cost

$$\geq \frac{1}{11} \varepsilon w_0. \quad (3.6)$$

Assuming that each of P_1, \dots, P_k has weight $\leq (1+\varepsilon)w_0$, the number of heavy edges in $P_1 \cup \dots \cup P_k$ is at most

$$\frac{k(1+\varepsilon)w_0}{\frac{1}{11} \varepsilon w_0} < \frac{12k}{\varepsilon} < \frac{12s}{\varepsilon}. \quad (3.7)$$

Also, modulo the one-time failure probability $O(n^{-1.9})$ from Lemma 3.3, by that lemma each path has length at most

$$(1 + \epsilon)w_0 \cdot 19n < 20s. \quad (3.8)$$

Thus, the length of all k paths taken together (i.e., the number of edges in $P_1 \cup \dots \cup P_k$) is at most

$$20ks < 10s^2. \quad (3.9)$$

And of course the k paths include

$$\text{exactly } k \text{ edges incident on each of } s \text{ and } t. \quad (3.10)$$

Subject to these assumptions—that each of P_1, \dots, P_k has weight $\leq (1 + \epsilon)w_0$ and that the high-probability conclusion of Lemma 3.3 holds— $P_1 \cup \dots \cup P_k$ satisfies all three of the constraints (3.7), (3.9), and (3.10) on heavy edges, all edges, and “incident” edges. An adversary who can delete any edge set subject to these constraints is able to delete $P_1 \cup \dots \cup P_k$, which is all we require. However, to simplify analysis we will give the adversary even more power.

At the root of R we will allow the adversary to delete edges subject only to (3.10); at level 1, additional edges subject only to the “heavy-edge budget” (3.7); and at levels 2 and 3 and for “middle” edges, additional edges subject only to the “edge-count budget” (3.9).

We will show how to choose the parameters of R so that every s – t path in R has cost $\leq (1 + \epsilon)w_0$, and so that R is “robust”: after the adversarial deletions, at least one path remains. Specifically, we will arrange that there remains a path in which the “root” edge incident to s costs $\leq \frac{k}{n} + \frac{1}{9}\epsilon w_0$, the edge out of level 1 is heavy but has cost $\leq \frac{1}{9}\epsilon w_0$, the edge out of level 2 may be light or heavy and also has cost $\leq \frac{1}{9}\epsilon w_0$, the path through the SP tree has total cost $\leq \frac{1}{2} \frac{\ln n}{n} + \frac{1}{9}\epsilon w_0$, the central edge joining this to the opposite SP tree adds cost $\leq \frac{1}{9}\epsilon w_0$, and the continuation of this path to t has the symmetrical properties. It is immediate that such a path has total cost $\leq (2k + \ln n)/n + 9 \cdot \frac{1}{9}\epsilon w_0 = (1 + \epsilon)w_0$. (But see (3.28) for confirmation, after the construction is detailed.)

3.3 | Level 0, cheapest edges

On s , add to R the $k + r_0$ edges of lowest cost, excluding $\{s, t\}$ from consideration, with

$$r_0 = \left\lceil \frac{1}{10}\epsilon s \right\rceil = \Theta(s). \quad (3.11)$$

Consider this step a *failure* if any selected edge has cost greater than $\frac{k}{n} + \frac{1}{9}\epsilon w_0$. There are $n' = n - 2 = (1 - o(1))n$ edges under consideration, with weights i.i.d. $U(0, 1)$, and failure occurs iff the number X of edges with weights in the interval $[0, \frac{k}{n} + \frac{1}{9}\epsilon w_0]$ is smaller than $k + r_0$. Note that $X \sim \text{Bi}(n', \frac{k}{n} + \frac{1}{9}\epsilon w_0)$, thus $\mathbb{E}X = (1 - o(1))(k + \frac{1}{9}\epsilon s)$, and failure means that $X < k + r_0$, that is, that

$$\frac{X}{\mathbb{E}X} < (1 + o(1)) \frac{k + r_0}{k + \frac{1}{9}\epsilon s} = (1 + o(1)) \frac{k + \frac{1}{10}\epsilon s}{k + \frac{1}{9}\epsilon s},$$

which by $s > 2k$ is

$$< (1 + o(1)) \frac{k + \frac{1}{10}\varepsilon \cdot 2k}{k + \frac{1}{9}\varepsilon \cdot 2k} = (1 + o(1)) \frac{(1 + \frac{2}{10}\varepsilon)k}{(1 + \frac{2}{9}\varepsilon)k} < 1 - \frac{1}{50}\varepsilon = 1 - \Omega(\varepsilon).$$

By Lemma 3.2, then, the probability of failure is

$$\mathbb{P}(X < (1 - \Omega(\varepsilon))\mathbb{E}X) \leq \exp(-\Omega(\varepsilon^2)\mathbb{E}X/2) \leq \exp(-\Omega(\varepsilon^2 \cdot \varepsilon s)) \leq \exp(-\Theta(s)), \quad (3.12)$$

the final expression using that ε is constant (see (3.2)).

So, modulo the given failure probability, every selected edge incident on s has cost $\leq \frac{k}{n} + \frac{1}{9}\varepsilon w_0$, and after the adversarial deletion of k of these edges, r_0 remain. The selection of edges conditions the costs of the other edges incident on s , but none will play any role in the analysis.

The purpose of the next two levels is to expand the number of edges to the point where the adversary cannot delete all of them, because of the heavy-edge budget (3.7) for edges out of level 1, and the edge-count budget (3.9) for edges out of level 2 and beyond. At the same time, we try to minimize the number of vertices introduced into the construction so that it will remain $o(n)$ for as large a k as possible.

3.4 | Level 1, cheapest heavy edges

From each neighbor v of s along the edges just added, add to R the

$$r_1 := \lceil 10,000/\varepsilon^2 \rceil = \Theta(1) \quad (3.13)$$

cheapest *heavy* edges from v to any of the $n' = n(1 - o(1))$ vertices not yet added (see (3.3)), as before also excluding vertex t . Consider this step a *failure* if any added edge has cost greater than $\frac{1}{9}\varepsilon w_0$. For each neighbor v there are n' edges under consideration, with weights i.i.d. $U(0, 1)$, and failure occurs iff the number X of edges with weights in the interval $[\frac{1}{11}\varepsilon w_0, \frac{1}{9}\varepsilon w_0]$ is smaller than r_1 . Note that $X \sim \text{Bi}(n', (\frac{1}{9} - \frac{1}{11})\varepsilon w_0)$, thus $\mathbb{E}X = (1 - o(1))(\frac{1}{9} - \frac{1}{11})\varepsilon s = \Theta(\varepsilon s)$. Failure means that $X < r_1 < \mathbb{E}X/2$, so by Lemma 3.2 the probability of failure for a given v is $\leq \exp(-\Theta(\varepsilon s))$. The number of level-1 vertices v is $k + r_0 = O(s)$, so by the union bound the probability of any failure is

$$\leq O(s) \exp(-\Theta(\varepsilon s)) \leq \exp(-\Theta(s)), \quad (3.14)$$

by suitable adjustment of the constants implicit in Θ .

This edge selection conditions the costs of the other edges incident on each v , but none will play any role in the analysis. The adversary must leave r_0 edges out of the root, expanding to

$$r_0 r_1 \geq \frac{1}{10}\varepsilon s \cdot 10,000/\varepsilon^2 = 1,000s/\varepsilon$$

(heavy) edges out of level 1, of which (by (3.7)) he can delete at most $12s/\varepsilon$, leaving (very generously calculated) at least

$$\rho_1 := 120s/\varepsilon = \Theta(r_0 r_1) \quad (3.15)$$

edges out of level 1. The vertices at the opposite endpoints of these edges constitute level 2.

3.5 | Level 2, cheapest edges

From each level 2 vertex v in turn, add to R the cheapest r_2 edges to any of the $n' = n(1 - o(1))$ vertices not yet added, again also excluding vertex t from consideration. Here choose r_2 so as to make

$$\rho_2 := \rho_1 r_2 = 12s^2, \tag{3.16}$$

namely taking

$$r_2 = \frac{12s^2}{\rho_1} = \frac{12s^2}{120s/\varepsilon} = \frac{1}{10}\varepsilon s = \Theta(\varepsilon s). \tag{3.17}$$

Consider this step a *failure* if any added edge has cost greater than $\frac{1}{9}\varepsilon w_0$. For each neighbor v there are $n' = (1 - o(1))n$ edges under consideration, with weights i.i.d. $U(0, 1)$, and failure occurs iff the number X of edges with weights in the interval $[0, \frac{1}{9}\varepsilon w_0]$ is smaller than r_2 . Note that $X \sim \text{Bi}(n', \frac{1}{9}\varepsilon w_0)$, thus $\mathbb{E}X = (1 - o(1))\frac{1}{9}\varepsilon s = \Theta(\varepsilon s)$. Failure means that $X < r_2 < 0.99\mathbb{E}X$, so by Lemma 3.2 the probability of failure for a given v is $\leq \exp(-\Theta(\varepsilon s))$. The number of level-2 vertices v is $(k + r_0)r_1 = O(s)$ so by the union bound the probability of any failure is

$$\leq \exp(-\Theta(\varepsilon s)). \tag{3.18}$$

This edge selection conditions the costs of the other edges incident on each v , but none will play any role in the analysis. The adversary had to leave at least ρ_1 edges out of level 1, expanding to $\rho_1 r_2 = \rho_2 = 12s^2$ edges out of level 2, of which by (3.9) he can delete at most $10s^2$, leaving at least $2s^2$ edges out of level 2. The vertices at the opposite endpoints of these edges constitute level 3.

3.6 | Level 3, shortest-path trees

We now grow each level-3 vertex v to a tree T_v with d vertices, including v , choosing

$$d := \left\lceil \sqrt{\frac{n \ln n}{2s^3}} \right\rceil < \sqrt{n}. \tag{3.19}$$

We grow these trees one after another, always working within the $n' = n(1 - o(1))$ vertices not yet added, and again always excluding vertex t from consideration.

Controlling the lengths of the paths in T_v would allow a choice of d as large as \sqrt{n} , but we make it smaller to keep the number of vertices in R as small as possible (and thus keep it to $o(n)$ for s as large as possible).

Here it will be convenient to work with exponentially rather than uniformly distributed edge weights. There are various easy ways to arrange this. We do so by temporarily replacing each uniform weight w with a weight $w' = -\ln(1 - w)$; it is standard that these transformed weights are exponentially distributed, and that $w' \geq w$. We construct a SPT of order d using the transformed weights; it will not be an SPT for the original weights, but its paths will be short under the original weights, which is all that we care about.

Define the distance $\text{dist}(u, v)$ between two vertices to be the cost of a minimum-weight path between them, and define the radius $\text{rad}(T_v)$ of an SPT T_v to be the maximum distance from v to any vertex in T_v . The radius is described by the following claim, which we phrase in a generic setting with n vertices and a root vertex s .

Claim 3.4. In a complete graph K_n with i.i.d. exponential edge weights with mean 1, the radius $X = \text{rad}(T_s)$ of a SPT T_s of order d is

$$X = \sum_{i=1}^{d-1} X_i, \tag{3.20}$$

where the X_i are independent random variables with $X_i \sim \text{Exp}(i(n - i))$.

Proof. Following [14], think of the process of finding shortest paths from s to other vertices as first-passage percolation or “infection spreading” starting from s . Let $L := L(r)$ be the set of vertices within radius (distance) r of s ; we think of gradually increasing r , starting with $r = 0$ where $L = \{s\}$. It is well known that each edge $(v, u) \in L(r) \times (V \setminus L(r))$ has exponentially distributed weight W conditioned by $W + \text{dist}(s, v) \geq r$, and that these random weights are independent. This can be seen by imagining that infection has spread to radius r from s , including to the vertex v and additionally a length $r - \text{dist}(s, v)$ further along the edge (v, u) , and appealing to the memoryless property of the exponential distribution; it can also be verified by analyzing Dijkstra’s algorithm in this randomized setting.

It follows that the distance X_1 to the vertex nearest s is distributed as $X_1 \sim \text{Exp}(n - 1)$; the additional distance to the next vertex is X_2 with $X_2 \sim \text{Exp}(2(n - 2))$ and independent of X_1 (for total distance $X_1 + X_2$); and when there are i vertices in the tree, the additional distance to the next is $X_i \sim \text{Exp}(i(n - i))$, with all the X_i independent, for total distance as claimed. ■

We will only use trees T_v whose radius is $X \leq (1 + \frac{2}{9}\epsilon) \frac{1}{2} \ln n/n < \frac{1}{2} \ln n/n + \frac{1}{9}\epsilon w_0$. Call a tree a failure (and do not include it in the structure R) if $X > (1 + \frac{2}{9}\epsilon) \frac{1}{2} \ln n/n$. Declare the construction of level 3 a failure if more than $0.01s^2$ trees fail.

Since (3.20) is monotone increasing in d , the larger the d , the greater the probability of failure, so in the next paragraphs we will pessimistically take d to be \sqrt{n} (ignoring integrality since \sqrt{n} is large). In this case, applying Claim 3.4 to T_v , constructed in a complete graph of order $n' = n(1 - o(1))$, the expectation of X is

$$\mu := \mathbb{E}X = \sum_{i=1}^{d-1} \frac{1}{i(n' - i)} = \frac{1 + o(1)}{n} \sum_{i=1}^{d-1} \frac{1}{i} = (1 + o(1)) \frac{\ln d}{n} = (1 + o(1)) \frac{1}{2} \frac{\ln n}{n}. \tag{3.21}$$

Thus, failure of T_v implies that

$$\frac{X}{\mu} > 1 + \frac{1}{5}\epsilon. \tag{3.22}$$

To bound the probability of this event we require one more lemma (also used later in proving Lemma 4.2).

Lemma 3.5 ([13, Theorem 5.1]). *Let $X = \sum_{i=1}^n X_i$ with $X_i \sim \text{Exp}(a_i)$ independent rate- a_i random variables, where $a_i \geq 0$. Write $a_* := \min_i a_i$ and $\mu := \mathbb{E}X = \sum_{i=1}^n \frac{1}{a_i}$. Then:*

for any $\lambda = 1 + \epsilon > 1$,

$$\mathbb{P}(X \geq \lambda\mu) \leq \lambda^{-1} e^{-a_*\mu(\lambda^{-1} - \ln \lambda)} \leq \exp(-\Omega(\alpha_*\mu)) \tag{3.23}$$

for any $\lambda = 1 - \epsilon < 1$,

$$\mathbb{P}(X \leq \lambda\mu) \leq e^{-a_*\mu(\lambda - 1 - \ln \lambda)} \leq \exp(-\Omega(\alpha_*\mu)), \tag{3.24}$$

and for any $\epsilon > 0$,

$$\mathbb{P}(|X - \mu| \geq \epsilon\mu) \leq 2 \exp(-\Omega(\alpha_*\mu)). \tag{3.25}$$

The constants implicit in the $\Omega(\cdot)$ expressions are positive and only depend on ϵ .

Proof. The inequalities in (3.23) and (3.24) in terms of λ are directly from [13, Theorem 5.1]. The remaining expressions, including (3.25), follow immediately. ■

From (3.23) of Lemma 3.5, the probability of the event in (3.22) (and thus that of T_v failing) is at most

$$\mathbb{P}\left(X - \mu > \frac{1}{5}\epsilon\mu\right) \leq \exp(-\Omega(\alpha_*\mu)) = \exp(-\Omega(\ln n)), \tag{3.26}$$

using that $\alpha_* = n' = (1 - o(1))n$, μ is given by (3.21), and $\epsilon = \Theta(1)$.

The total number of trees built is $N = (k+r_0)r_1r_2$, which, with reference to (3.4), (3.11), (3.13), and (3.17), is $\Theta(s^2)$. By (3.26), each tree independently fails with at most some probability $p = o(1)$. Thus, the number of trees surviving dominates $\text{Bi}(N, 1 - p)$, with expectation $\lambda = N(1 - p) = N(1 - o(1))$. Failure at level 3 means that at least $0.01s^2 = \Theta(N)$ trees fail, equivalently the number surviving is at most some $\lambda(1 - \Theta(1))$, which by Lemma 3.2 has probability

$$\exp(-\Omega(s^2)). \tag{3.27}$$

Remark 3.6. When construction of a tree T_v rooted at a level-3 vertex v is finished, the edge between any vertex a of T_v and any vertex b in $V' \setminus V(T_v)$ has weight $w(a, b)$ that—still in the uniform model with edge weights temporarily transformed to be exponentially distributed—is exponentially distributed conditional upon being $\geq \text{rad}(T_v) - d(v, a)$. Equivalently, the edge $\{a, b\}$ gives a v -to- b path (through a) with cost $\text{rad}(T_v) + X_{a,b}$, where the “excess” $X_{a,b}$ has simple exponential distribution $X_{a,b} \sim \text{Exp}(1)$ (with no conditioning). Furthermore, the $X_{a,b}$ are independent, over all choices of a and b .

Call R_s the now-complete construction on s . Note that there is no conditioning on edges between the remaining vertices; in particular, the SPT infection process (or equivalently Dijkstra’s algorithm) as described in Claim 3.4 never looked at edges between uninfected vertices.

3.7 | Symmetric construction on vertex t

Just as we have constructed R_s , we now make a similar construction R_t for vertex t , with the same branching factors out of levels 0, 1, and 2 and similar SPTs on level-3 vertices. Since the number n' of vertices available after constructing R_s still satisfies $n' = (1 - o(1))n$, and because the construction on s did not look at nor condition any edge between these vertices, the construction on t enjoys the same properties as that on s .

3.8 | Edges between the trees on s and t

It remains only to complete paths between s and t , which we do by adding cheap edges (where present) between the SPTs in R_s and those in R_t .

Let T_u be an SPT rooted at a level-3 vertex u of R_s , and T_v one rooted at a level-3 vertex v of R_t . Let a and b be any vertices in T_u and T_v respectively. By Remark 3.6, edge $\{a, b\}$ gives a u -to- b path with cost $\text{rad}(T_u) + X_{a,b}$, the collection of all the excesses $X_{a,b}$ being i.i.d. each with distribution $X_{a,b} \sim \text{Exp}(1)$. Thus, $\{a, b\}$ gives a u -to- v path with cost $\leq \text{rad}(T_u) + X_{a,b} + \text{rad}(T_v)$.

Select, and add to the full construction R , any such “middle edge” $\{a, b\}$ having $X_{a,b} \leq \frac{1}{9}\epsilon w_0$. This completes the construction of R .

3.9 | Order of R , failure probability, and path costs

It is worth first confirming that the construction uses, as claimed, $o(n)$ vertices. The number of vertices used is of order $(k+r_0)r_1r_2d$, which by (3.11), (3.13), (3.17), and (3.19) is $O(s^2d)$. Recalling from (3.19) that $d = \left\lceil \sqrt{n \ln n / 2s^3} \right\rceil$, as long as the ceiling function does not affect the order of d , the total number of vertices is $O(s^2d) = O(\sqrt{ns \ln n})$, which is $o(n)$ for $s = o(n/\ln n)$. However, the ceiling function does affect the order of d when $n \ln n / 2s^3 < 1$, that is, when $s > (\frac{1}{2}n \ln n)^{1/3}$; in this case, $d = 1$, the total number of vertices used is $O(s^2)$, and this is still $o(n)$ if $s = o(\sqrt{n})$. Taking the two cases together, the construction is valid up to any $s = o(\sqrt{n})$, or equivalently for any $k = o(\sqrt{n})$.

Failures at levels 0, 1, and 2 each occur w.p. $\leq \exp(-\Theta(s))$ (by (3.12), (3.14), and (3.18)), and at level 3 w.p. $\leq \exp(-\Omega(s^2))$ (by (3.27)), so by the union bound the probability of any failure is $\leq \exp(-\Theta(s))$.

We now confirm that, assuming that the construction was successful, any s - t path in R through successful SPTs has cost $\leq (1 + \epsilon)w_0$. (Remember that there may be some unsuccessful SPTs.) By assumption of success, any level-0 edge on s or t has cost $\leq \frac{k}{n} + \frac{1}{9}\epsilon w_0$, any level-1 edge has cost $\leq \frac{1}{9}\epsilon w_0$, and any level-2 edge also has cost $\leq \frac{1}{9}\epsilon w_0$. Each successful level-3 tree T in R_s or R_t has radius $\text{rad}(T) \leq (1 + \frac{2}{9}\epsilon)\frac{1}{2} \ln n/n \leq \frac{1}{2} \ln n/n + \frac{1}{9}\epsilon w_0$, and each selected middle edge $\{a, b\}$ connects the roots of two trees at an excess cost (above the sum of the two radii) of $X_{a,b} \leq \frac{1}{9}\epsilon w_0$. The total of the 9 upper bounds in question is

$$2 \cdot \frac{k}{n} + 2 \cdot \frac{1}{2} \ln n/n + 9 \cdot \frac{1}{9}\epsilon w_0 = (1 + \epsilon)w_0. \quad (3.28)$$

3.10 | Robustness of R

We now show that, after the deletion of the k cheapest paths in G , there remains at least one path in R (that uses successful SPTs). Recall from Section 3.2 that deletion of the k cheapest paths in G is conservatively modeled as an adversarial deletion subject to: (3.10), the deletion of exactly k edges incident on each of s and t ; (3.7), the number of heavy edges deleted at level 1; and (3.9), the total number of edges deleted elsewhere in R (at levels 2 and 3, and joining R_s and R_t).

Without loss of generality we may assume that the adversary does not delete an edge within an SPT T , nor a middle edge from such a tree to a facing one, since deleting the level-2 edge into the level-3 root of T destroys more paths in R at the same budgetary cost.

By the assumption of success, there are at most $0.01s^2$ failed SPTs on each of s and t , and for simplicity we will deal with them by imagining all trees to be successful but allowing the adversary his choice of this many SPTs to delete; by the argument above we can model this as deletion of edges into the roots of these trees, and simply add $0.02s^2$ to this budget.

Let us now allow the adversary to delete k edges from each of s and t , $12s/\epsilon$ edges out of level 1 for each (double-counting the heavy-edge budget), and $10.02s^2$ edges out of level 2 for each (again

double-counting). Can he destroy all s - t paths? We have not yet made any high-probability structural assertion about the middle edges, so this is a probabilistic question: what is the probability, over the randomness still present in the middle edges, that there is an adversarial deletion destroying all paths?

Of the $k + r_0 = O(s)$ edges on s , the adversary chooses k to delete; there are at most $2^{O(s)}$ ways to do so. Any choice leaves $\Theta(r_0 r_1) = \Theta(s)$ edges out of level 1, of which the adversary is able to delete a positive fraction, again in at most $2^{O(s)}$ ways. Any choice leaves $\Theta(s^2)$ edges out of level 2, of which the adversary is able to delete a positive fraction, in at most $2^{O(s^2)}$ ways. The adversary makes a similar set of choices on t , but still this comes to just $2^{O(s^2)}$ possible outcomes in all.

A given deletion choice destroys all paths precisely if it leaves no middle edge of excess $\leq \frac{1}{9}\epsilon w_0$. (Remember that, w.l.o.g., we have excluded deletions in and between the SPTs at level 3.) By construction, any deletion choice leaves $\Theta(s^2)$ edges out of level 2 and thus, by (3.19), $\Theta(s^2 d) = \Omega(\sqrt{ns \ln n})$ vertices in SPTs in each of R_s and R_t , for $\Omega(ns \ln n)$ potential middle edges. A middle edge is *selected* if its excess cost (in the exponential model) is $w' = -\ln(1 - w) \leq \frac{1}{9}\epsilon w_0$, that is, if $1 - w \geq \exp(-\frac{1}{9}\epsilon w_0)$, thus is *rejected* with probability $\exp(-\frac{1}{9}\epsilon w_0)$. There is no path only if every potential edge is rejected, which happens w.p. $\leq \exp(-\frac{1}{9}\epsilon w_0 \cdot ns \ln n) = \exp(-\Omega(s^2 \ln n))$. Taking the union bound over all adversarial choices, the probability that any choice leaves no paths is

$$2^{O(s^2)} \exp(-\Omega(s^2 \ln n)) = \exp(-\Omega(s^2 \ln n)). \tag{3.29}$$

This is dominated by the failure probabilities $\exp(-\Theta(s))$ for other steps.

3.11 | Success for each k , and for all k

We have shown that, for any $k = o(\sqrt{n})$, subject to an absence of failures, we can generate a robust structure $R^{(k)}$ in which, after adversarial deletions, there remains an s - t path of cost $\leq (1 + \epsilon)w_0(k)$. (Remember that w_0 and s are simple functions of k , per (3.4) and (3.5). Here we retain the argument k we usually suppress.) There are two types of failures possible. The first is that the graph fails Lemma 3.3's conclusion that "cheap paths are short"; this occurs w.p. $O(n^{-1.9})$. The second is that $R^{(k)}$ is not robust; this occurs w.p. $O(\exp(-\Omega(s(k))))$.

Assume success in generating $R^{(k)}$. We claim that P_1, \dots, P_{k+1} all have cost $\leq (1 + \epsilon)w_0(k)$ (call this "cheap"). Suppose not. Then there is some $i \leq k$ for which P_1, \dots, P_i are cheap but P_{i+1} is not. Our adversary's budget allows it to delete P_1, \dots, P_i , and by assumption of success this leaves a cheap path P in $R^{(k)}$. Thus there is a cheap $i + 1$ st path in G , a contradiction.

It follows that for each k , $X_{k+1} \leq (1 + \epsilon)w_0(k)$ with probability

$$1 - O(n^{-1.9}) - O(\exp(-\Omega(s(k)))) \tag{3.30}$$

A simple calculation shows that w.h.p. $X_{k+1} \leq (1 + \epsilon)w_0(k)$ simultaneously for all k in this range, proving the upper bound in (1.3). By the union bound, the probability of failure to build a robust structure for any k is at most

$$\begin{aligned} \sum_{k=0}^{\infty} \exp(-\Omega(s(k))) &\leq \ln n \exp(-\Omega(\ln n)) + \sum_{k=\ln n}^{\infty} \exp(-\Omega(k)) \\ &= \exp(-\Omega(\ln n)) = n^{-\Omega(1)}. \end{aligned} \tag{3.31}$$

Including the probability of failure in applying Lemma 3.3, the total failure probability is $O(n^{-1.9} + n^{-\Omega(1)}) = o(1)$.

3.12 | Limitation to small k

We have established Theorem 1.1 up to any $k = o(\sqrt{n})$, and the construction of $R^{(k)}$ was tailored to such values. For levels using heavy edges, fanout is limited to $O(s)$. On the other hand, the meet-in-the-middle argument requires that each side grow large, to $\Omega(\sqrt{n/s})$. Thus, for small k , a more-than-constant number of levels is needed. Summing heavy edges over this many levels would exceed the target weight $(1 + \varepsilon)w_0$, so light edges are needed. The adversary may delete $\Theta(s^2)$ light edges, so the construction must contain at least this many. The construction explicitly required each light edge to lead to a new vertex, and we do not readily see how to do otherwise as long as we are using SPTs, thus intrinsically limiting s (thus k) to $O(\sqrt{n})$. For larger k , however, we can obtain sufficient heavy-edge fanout in constant depth, permitting a simpler construction described in Section 5.

4 | EDGE ORDER STATISTICS

In this section we establish results on order statistics needed in later sections. Let $\{W_{(k)}\}_{k=1}^{n-1}$ be the order statistics of $n - 1$ i.i.d. random variables, variously uniform $U(0, 1)$ or exponential $\text{Exp}(1)$. We choose $n - 1$ rather than n as the parameter both because many expressions are more natural in this parametrization, and because this way $W_{(k)}$ is the cost of the k th cheapest edge incident to a fixed vertex $v \in K_n$.

The following lemma is used in Section 6.3.

Lemma 4.1. *Let $l = n^{-0.99}$. Consider the unit interval $[0, 1]$ with n points placed uniformly and independently at random. Then w.h.p. every interval of length at least $l' \geq l$ contains at least $0.99l'n$ points.*

Proof. Partition the unit interval into contiguous intervals I_i each of length $L := l/1000$, using $\lfloor 1/L \rfloor$ such intervals (possibly leaving a small interval near 1 not covered). Any interval I of length $l' \geq l$ has at least a $998/1000$ fraction of its length covered by intervals $I_i \subset I$, and we will show that w.h.p. every interval I_i contains at least $0.999Ln$ points (that is, at least a 0.999 fraction of the expectation). If so, it follows that I has at least $0.999 \cdot 0.998l'n \geq 0.99l'n$ points.

The distribution of the number of points in each interval I_i of length L follows the binomial distribution $\text{Bi}(n, L)$. By Lemma 3.2,

$$\mathbb{P}(\text{Bi}(n, L) \leq 0.999Ln) \leq \exp(-\Omega(Ln)),$$

where the sign in the Ω is taken as positive. The probability that any interval I_i contains less than 0.999 points is, by the union bound, at most,

$$\lfloor 1/L \rfloor \cdot \exp(-\Omega(Ln)) = \exp(-\Omega(n^{0.01})) = o(1) \quad (4.1)$$

as desired. ■

The following lemma is used in (6.10) and (7.5).

Lemma 4.2. *Let $\{W_{(k)}\}_{k=1}^{n-1}$ be the order statistics of $n - 1$ i.i.d. random variables, either all uniform $U(0, 1)$ or all exponential $\text{Exp}(1)$. For any $\varepsilon > 0$ and $a = a(n) = \omega(1)$, w.h.p.*

$$1 - \varepsilon \leq \frac{W_{(k)}}{\mathbb{E}W_{(k)}} \leq 1 + \varepsilon$$

simultaneously for all k in the range $a \leq k \leq n - 1$.

Proof. Without loss of generality, we may assume that $a \leq n/10$.

Exponential case. It is standard that, where $Z_i \sim \text{Exp}(i)$ are independent exponential r.v.s., we may generate the $W_{(k)}$ as

$$W_{(k)} = \sum_{i=1}^k Z_{n-i}. \tag{4.2}$$

Using a superscripted E to highlight the exponential model, $W_{(k)}$ has mean

$$\mu_k = \mu_k^{(E)} := \mathbb{E}W_{(k)} = \sum_{i=1}^k \frac{1}{n-i} = H(n-1) - H(n-k-1) \sim \ln(n) - \ln(n-k); \tag{4.3}$$

the change by 1 in the logarithms' arguments avoids $\ln(0)$ when $k = n - 1$ and remains asymptotically correct.

By (3.25),

$$\mathbb{P}(|W_{(k)} - \mu_k| \geq \varepsilon \mu) \leq 2 \exp(-\Omega((n-k)\mu_k)). \tag{4.4}$$

By the union bound, it suffices to show that the sum over k from a to $n - 1$ of the RHS of (4.4) is $o(1)$. We treat the sum in two ranges. For $k \leq \frac{n}{2}$, $(n-k)\mu_k \geq \frac{n}{2} \cdot \frac{k}{n} = \frac{k}{2}$. Thus,

$$\sum_{k=a}^{n/2} \exp(-\Omega((n-k)\mu_k)) \leq \sum_{k=a}^{n/2} \exp(-\Omega(k)) \leq O(ae^{-\Omega(a)}) \rightarrow 0, \tag{4.5}$$

since $a = \omega(1)$. For $k > \frac{n}{2}$, for brevity let $\bar{k} = n - k$. Then $\mu_k \sim \ln n - \ln(\bar{k})$ by (4.3) and

$$\begin{aligned} \sum_{\bar{k}=1}^{n/2} \exp(-\Omega((n-k)\mu_k)) &= \sum_{\bar{k}=1}^{n/2} \exp(-\bar{k} \Omega(\ln n - \ln(\bar{k}))) = \sum_{\bar{k}=1}^{n/2} \left(\frac{\bar{k}}{n}\right)^{\Omega(\bar{k})} \\ &\leq \left(\frac{1}{n}\right)^{\Omega(1)} \sum_{\bar{k}=1}^{n/2} \bar{k}^{\Omega(1)} \left(\frac{\bar{k}}{n}\right)^{\Omega(\bar{k}-1)} = n^{-\Omega(1)} = o(1), \end{aligned} \tag{4.6}$$

where the explicit inequality factors out the $\bar{k} = 1$ term, from which, since $\bar{k}/n \leq 1/2$, the later terms decrease geometrically. This concludes the exponential case.

Uniform case: Let $U_i \sim U(0, 1)$ be i.i.d. uniform random variables and $W_i \sim \text{Exp}(1)$ i.i.d. exponential random variables. Because the exponential distribution has CDF $F(x) = 1 - \exp(-x)$, we may couple the two sets of variables as $U_i = F(W_i)$ or equivalently $W_i = f(U_i)$ with $f(x) = F^{-1}(x) = -\ln(1-x)$. Because f is increasing, $W_{(k)} = f(U_{(k)})$. Now using superscript U to distinguish the uniform model, the mean is well known to be

$$\mu_k = \mu_k^{(U)} := \mathbb{E}U_{(k)} = \frac{k}{n}. \tag{4.7}$$

We want to show that with high probability, for all k in the range $a \leq k \leq n - 1$,

$$(1 - \varepsilon)\mu_k^{(U)} \leq U_{(k)} \leq (1 + \varepsilon)\mu_k^{(U)}$$

or equivalently,

$$f((1 - \varepsilon)\mu_k^{(U)}) \leq W_{(k)} \leq f((1 + \varepsilon)\mu_k^{(U)}).$$

From the exponential case already proved, taking error bound $\varepsilon/2$, we know that w.h.p., for all k ,

$$(1 - \varepsilon/2)\mu_k^{(E)} \leq W_{(k)} \leq (1 + \varepsilon/2)\mu_k^{(E)},$$

so it suffices to show that, for all k (deterministically),

$$f((1 - \varepsilon)\mu_k^{(U)}) \leq (1 - \varepsilon/2)\mu_k^{(E)} \quad \text{and} \quad f((1 + \varepsilon)\mu_k^{(U)}) \geq (1 + \varepsilon/2)\mu_k^{(E)}.$$

This is so. Using (4.7), (4.3), and convexity of f ,

$$f((1 - \varepsilon)\mu_k^{(U)}) = f((1 - \varepsilon)k/n) \leq (1 - \varepsilon)f(k/n) = (1 - \varepsilon) \ln\left(\frac{n}{n-k}\right) \leq (1 - \varepsilon/2)\mu_k^{(E)};$$

$$f((1 + \varepsilon)\mu_k^{(U)}) = f((1 + \varepsilon)k/n) \geq (1 + \varepsilon)f(k/n) = (1 + \varepsilon) \ln\left(\frac{n}{n-k}\right) \geq (1 + \varepsilon/2)\mu_k^{(E)}. \quad \blacksquare$$

5 | UPPER BOUND FOR LARGE k , SKETCH

5.1 | Introduction

To address larger values of k we use a different construction, generating s - t paths of length 4. A straightforward extension of the previous argument to this construction would let us get up to $k = n - f(n)$ for an arbitrarily slowly growing function f , but not to $k = n - 1$ because it requires $k + 1/\varepsilon^2$ edges incident on each of s and t (thus requires that $k + 1/\varepsilon^2 \leq n - 1$).

Getting all the way to $k = n - 1$ requires a couple of additional ideas. Again, we will introduce an adversary with a cost budget that with high probability exceeds the cost of the first k cheapest paths. First, we observe that much of the adversary's cost budget must be spent on edges incident to s and t , leaving less to delete other edges, thus allowing a smaller structure R to be sufficiently robust. In particular, the k cheapest paths from s to t must use edges incident on s of total weight at least $\sum_{i=1}^k W_{(i)}^s$ where

$$W_{(i)}^v \tag{5.1}$$

is the cost of the i th cheapest edge incident to v . (We may omit the superscript when it is either generic or clear from context.) One technical detail is that, where R includes the $k + r_0$ cheapest edges incident to s , we will control $W_{(k+r_0)} - W_{(k)}$ directly, using results on order statistics from Section 4, rather than through a high-probability upper bound on $W_{(k+r_0)}$ and a high-probability lower bound on $W_{(k)}$. Finally, it is no longer adequate to allow path costs to exceed their nominal values by an $\varepsilon = \Theta(1)$ factor, as such large excesses would swell the adversary's budget too quickly, so we more tightly control the excess cost of each path P_k as a function of k (and n , implicitly).

The details later will be clearer if we sketch the argument now, with most details but without the calculations. We will argue for k from $n^{4/10}$ to $n - 1$. (We must start with some $k = o(n^{1/2})$ since that is as far as the "small k " argument extended, and we need $k = \omega(n^{1/3})$ since below this the new construction's path costs would exceed the $2k/n$ target.)




FIGURE 2 The robust structure $R = R^{(k)}$ after adversarial deletion of k edges on s , leaving r_0 edges to some vertices $V'_s \subseteq V_s$, and likewise for t and V'_t . The middle vertices are pruned to $M' = M \setminus (V'_s \cup V'_t)$, and edges from M' to V'_s and V'_t are in R if they have weight between ϵ_k and $2\epsilon_k$. Here, edges from just one representative vertex $v \in M'$ are illustrated

5.2 | Structure R

Figure 2 illustrates the robust structure $R = R^{(k)}$ after adversarial deletion of root edges, as discussed in Section 5.6. The construction is based on parameters $r_0 = r_0(k)$ and ϵ_k to be defined later. Start with R consisting of just the vertices s and t . Add to R the $k + r_0$ edges incident on s of lowest cost, and let V_s be the set of opposite endpoints of these edges. Do the same for t , generating vertex set V_t . Take $M := V(G) \setminus \{s, t\}$ as a collection of “middle vertices.”

Note that V_s , V_t , and M may well have vertices in common, but our analysis will use a subgraph of R where the relevant subsets of these three sets are disjoint, and it is easier to understand the construction imagining them to be disjoint. Add to R each edge e in $M \times V_s$ and $M \times V_t$ that is “heavy but not too heavy,” with cost $W(e) \in (\epsilon_k, 2\epsilon_k)$. This concludes the construction of the structure R .

5.3 | Path weights

It is immediate that every s - t path in R has cost at most

$$W_{(k+r_0)}^s + 2\epsilon_k + 2\epsilon_k + W_{(k+r_0)}^t. \tag{5.2}$$

We will show (in (6.11) for uniform and (8.15) for exponential) that, subject to the nonoccurrence of certain unlikely failure events, (5.2) is at most

$$W_{(k+1)}^s + W_{(k+1)}^t + 7\epsilon_k. \tag{5.3}$$

We will show in Section 5.6 that, after deletion of the first k paths, there remains an s - t path in R (again subject to nonoccurrence of unlikely failure events), whereupon it follows that

$$X_{k+1} \leq W_{(k+1)}^s + W_{(k+1)}^t + 7\epsilon_k. \tag{5.4}$$

5.4 | Adversary

We define an adversary who is “sufficiently strong” to delete the first k paths. For $k \leq n^{4/10}$, taking $\varepsilon = 0.1$, (3.30) implies that w.p.

$$1 - O(n^{-1.9}) - O(\exp(-\Omega(n^{4/10}))) = 1 - O(n^{-1.9}) = 1 - o(1), \quad (5.5)$$

we have that

$$X_k \leq X_{n^{4/10}} \leq 3n^{4/10}/n. \quad (5.6)$$

For $k > n^{4/10}$, further assume the absence of the failure events alluded to just above, so that (5.4) holds. Then, hiding a sum of the $\ln n/n$ terms of (1.3) in the $o(\cdot)$ term below,

$$\begin{aligned} \sum_{i=1}^k X_i &= \sum_{i=1}^{n^{4/10}} X_i + \sum_{i=n^{4/10}+1}^k X_i \\ &\leq \frac{3(n^{4/10})^2}{n} + \sum_{i=n^{4/10}+1}^k X_i \\ &\leq 3n^{-2/10} + \sum_{i=n^{4/10}+1}^k \left(W_{(i)}^s + W_{(i)}^t + 7\varepsilon_{i-1} \right) \\ &=: U_k. \end{aligned} \quad (5.7)$$

Thus, the first k paths’ edges have total weight at most U_k .

Furthermore, the first k paths’ edges incident on s and t are all distinct except, possibly, for the edge $\{s, t\}$. Therefore, not counting edge $s-t$ at all, the cost of these “incident” edges is at least

$$I_k := \sum_{i=1}^{k-1} \left(W_{(i)}^s + W_{(i)}^t \right). \quad (5.8)$$

(In proving Claim 8.2 we will use a slightly different lower bound I_k on the weight of the incident edges.)

It follows that the first k paths’ “middle edges” (edges other than the incident edges) cost at most $U_k - I_k$. We will explicitly define a budget B_k satisfying

$$B_k \geq U_k - I_k. \quad (5.9)$$

We will allow the adversary to delete any k edges in G incident on each of s and t , possibly including the edge $s-t$ (enough to let it delete the incident edges of the first k paths), and to delete any other edges in G of total cost at most B_k (enough to let it delete the middle edges of the first k paths). Thus, the adversary is sufficiently strong to delete the first k paths.

The adversary’s allowable deletions in G mean that also in R it deletes at most k edges incident on each of s and t , and middle edges of total cost at most B_k .

5.5 | Budgets B_k

The budgets B_k will be defined explicitly in the details. For the model with uniformly distributed edge weights we will do so in two ranges of k , corresponding to Claims 6.1 and 6.2, and likewise in the

model with exponentially distributed edge weights, corresponding to Claims 8.1 and 8.2. For Claims 6.2 and 8.2 we will establish (5.9) directly.

For Claims 6.1 and 8.1 we will establish (5.9) by the following reasoning; we will only need to check (5.10), (5.11), and (5.12) below. We will show that the budgets satisfy

$$B_{k+1} \geq B_k + 8\varepsilon_k. \quad (5.10)$$

(Roughly speaking, given B_k we will set ε_k as small as possible while keeping $R^{(k)}$ robust to the adversary with budget B_k . Then, we will set B_{k+1} as small as possible, namely by taking equality in (5.10). Behind the scenes, we derive B_k by solving the differential-equation equivalent of (5.10) satisfied with equality.)

We will show that (5.9) is satisfied in the base case, by showing that

$$B_k \geq U_k \quad \text{for } k = n^{4/10}. \quad (5.11)$$

Then, (5.9) is established for all k by induction on k :

$$U_{k+1} - I_{k+1} = (U_{k+1} - U_k) - (I_{k+1} - I_k) + (U_k - I_k)$$

which by (5.7), (5.8), and the inductive hypothesis (5.9) is

$$\begin{aligned} &\leq (W_{(k+1)}^s + W_{(k+1)}^t + 7\varepsilon_k) - (W_{(k)}^s + W_{(k)}^t) + B_k \\ &\leq B_k + 8\varepsilon_k \quad (\text{see below}) \end{aligned} \quad (5.12)$$

$$\leq B_{k+1} \quad (\text{by (5.10)}). \quad (5.13)$$

To justify (5.12) it suffices to show that $W_{(k+1)} - W_{(k)}$ is at most $0.1\varepsilon_k$, and we do so in (6.12) for the uniform case and in (8.17) for the exponential case. In both cases, $r_0 = \omega(1)$, and $W_{(k+r_0)} - W_{(k+1)} = O(\varepsilon_k)$ (as used in going from (5.2) to (5.3)), making this conclusion unsurprising.¹

5.6 | Robustness of R

We wish to make R robust against the adversary, so that after the deletions just described, R should retain an s - t path w.h.p., so that (5.4) holds and X_{k+1} is small. It will suffice to show that, to delete all s - t paths in R ,

$$\begin{aligned} &\text{after deletion of } k \text{ edges incident on each of } s \text{ and } t, \text{ an adversary would still} \\ &\text{have to delete middle edges of total cost more than } B_k, \end{aligned} \quad (5.14)$$

and thus it is powerless to do so.

Obtaining this robustness requires choosing ε_k sufficiently large in the construction. With reference to Figure 2, on deletion of any k edges on each of s and t , the level-1 sets are in effect pruned to V'_s and V'_t , each of cardinality r_0 . Should V'_s and V'_t have vertices in common, or if $t \in V'_s$ or $s \in V'_t$, then there

¹In proving Claims 6.2 and 8.2 we will set $r_0 = 1$, so this reasoning does not apply. Indeed, in Claim 8.2 (the large- k exponential case) (5.12) would be false— $W_{(k+1)} - W_{(k)}$ can be much larger than ε_k —but (to reiterate) it is not needed there, as we establish (5.9) directly.

is an s - t path. So, assume that V'_s and V'_t are disjoint and do not contain s nor t . Consider only middle vertices $M' \subseteq M$ not appearing in V'_s nor V'_t , that is, $M' = M \setminus \{V'_s \cup V'_t\}$. We will have $r_0 = o(n)$, so $|M'| = n - 2 - 2r_0 > 0.99n$. Note that edges in $M' \times V'_s, M' \times V'_t, \{s\} \times V'_s$, and $\{t\} \times V'_t$ are all distinct.

Consider a choice of the k deletions on s and t to be fixed in advance. (We will eventually take a union bound over all such choices.) The weights of edges in $M' \times V'_s$ and $M' \times V'_t$ have not even been observed yet, so each has (unconditioned) $U(0, 1)$ distribution, all are independent (by distinctness of the edges), and thus each such edge is included in R with probability ε_k , independently.

A vertex $v \in M'$ is connected to V'_s by

$$Z_v^s \sim \text{Bi}(r_0, \varepsilon_k) \quad (5.15)$$

edges, with mean

$$\lambda := \mathbb{E}Z_v^s = r_0 \varepsilon_k. \quad (5.16)$$

Define Z_v^t symmetrically, and note that Z_v^s and Z_v^t are i.i.d. Intuitively, if λ is small, Z_v^s is usually 0, is 1 with probability about λ , and rarely any larger value. So, the probability that v is connected to both V'_s and V'_t is about λ^2 , in which case to destroy s - t paths through v the adversary must delete an edge of cost at least ε_k . So, to delete all s - t paths, over the nearly n vertices in M' the adversary would have to delete edges of expected total weight at least

$$\varepsilon_k n \lambda^2. \quad (5.17)$$

We will choose ε_k so that

$$\varepsilon_k n \lambda^2 > B_k, \quad (5.18)$$

which hopefully will ensure (see Remark 5.1) that a path must remain (i.e., that R is robust).

Let us give a back-of-the-envelope calculation. In the uniform case we expect $W_{(k)}$ to be about k/n , so letting $r_0 = \varepsilon_k n$ means that $W_{(k+r_0)} - W_{(k)}$ will be about ε_k , justifying (5.3). Then (5.16) gives $\lambda = \varepsilon_k^2 n$, so (5.18) indicates that we need to take $\varepsilon_k^5 n^3 > B_k$. As noted after (5.10), roughly speaking, we obtain B_k and ε_k by solving this and (5.10) with equality as a system of differential equations.

Remark 5.1. This intuitive argument proves to be essentially sound, but to make it rigorous will take some work. Chiefly, $\mathbb{P}(Z_v^s > 0)$ is of course not exactly $\mathbb{E}Z_v^s = \lambda$ even when λ is small, and we will also have to consider the case when λ is large. Also, where the intuition is based on expectations, we must calculate the probability of the “failure” event that all paths can be deleted at a cost less than B_k . Finally, we must take the union bound of this failure event over all choices of root edges at s and t (but, as in the small- k case, this turns out to change nothing).

6 | UPPER BOUND FOR LARGE k , UNIFORM MODEL

In this section we fill in the details of the steps from Section 5 and show that they conclude the proof of the upper bound in Theorem 1.1. Specifically, to control the *path weights* (these emphasized keywords match section titles) we must show that (5.2) is at most (5.3). For the *adversary* we need only show (5.9); as noted earlier, for large k (Claim 6.2) we will do this directly, while for medium k (Claim 6.1) we will argue that the *budgets* B_k satisfy (5.11) and (5.12). And for *robustness* we will prove that the probability of failure is small (i.e., it is unlikely that the adversary can destroy all s - t paths in $R^{(k)}$).

6.1 | Claims, and implications for Theorem 1.1

We first state the two precise claims we make for large k , in two ranges. We use symbolic constants C_B, C_ε, C'_B , and C'_ε in the claims and the proofs, as it makes the calculations clearer. Whenever we encounter an inequality that the constants must satisfy, we will highlight with a parenthetical “check” that they do so.

Claim 6.1. For $k \in [n^{4/10}, n - 14\sqrt{n}]$, let $B_k = \left(C_B k n^{-3/5} + (2n^{-1/5})^{4/5} \right)^{5/4}$ and $\varepsilon_k = C_\varepsilon n^{-3/5} B_k^{1/5}$, with $C_B = 32$ and $C_\varepsilon = 5$. Then, asymptotically almost surely, simultaneously for all k in this range,

$$X_{k+1} \leq W_{(k+1)}^s + W_{(k+1)}^t + 8\varepsilon_k. \tag{6.1}$$

Remark: In proving Claim 6.1 we will set

$$r_0 := \varepsilon_k n. \tag{6.2}$$

From the definitions of B_k and ε_k in Claim 6.1, both are increasing in k , and we will make frequent use of the following inequalities. For n sufficiently large,

$$B_k = \Theta \left(k^{5/4} n^{-3/4} + n^{-1/5} \right) \tag{6.3}$$

$$B_k \leq B_n \leq 1.01 C_B^{5/4} n^{1/2} \tag{6.4}$$

$$B_k \geq B_{n^{4/10}} \geq 2n^{-1/5} \tag{6.5}$$

$$\varepsilon_k = \Theta \left(k^{1/4} n^{-3/4} + n^{-16/25} \right) \tag{6.6}$$

$$\varepsilon_k \leq \varepsilon_n \leq 1.01 C_\varepsilon C_B^{1/4} n^{-1/2} \tag{6.7}$$

$$\varepsilon_k \geq \varepsilon_{n^{4/10}} \geq 1.14 C_\varepsilon n^{-16/25}. \tag{6.8}$$

Claim 6.2. For $k \in (n - 14\sqrt{n}, n - 2]$, let

$$B_k = C'_B \sqrt{n} \quad \text{and} \quad \varepsilon_k = C'_\varepsilon n^{-1/6} \tag{6.9}$$

with $C'_B = 78$ and $C'_\varepsilon = 5$. Then, asymptotically almost surely, simultaneously for all k in this range,

$$X_{k+1} \leq W_{(k+1)}^s + W_{(k+1)}^t + 8\varepsilon_k.$$

Remark: In proving Claim 6.2 we will set $r_0 := 1$. Note that here B_k and ε_k are constants independent of k , but we retain the subscript for consistency with the notation of Section 5.1.

We will prove the two claims shortly.

Proof of the upper bound of (1.3) in Theorem 1.1. Given $\varepsilon > 0$ from Theorem 1.1, apply Lemma 4.2 to the order statistics $W_{(k)}^s$ and $W_{(k)}^t$ with ε in the lemma as our $\varepsilon/2$ and $a = n^{4/10}$. Then by Claim 6.1 w.h.p., simultaneously for all $k \in [n^{4/10} + 1, n - 14\sqrt{n}]$,

$$X_k \leq W_{(k)}^s + W_{(k)}^t + 8\varepsilon_{k-1} \leq (1 + \varepsilon/2)2k/n + 8\varepsilon_{k-1} \leq (1 + \varepsilon)(2k/n + \ln n/n); \tag{6.10}$$

the key point is that $\varepsilon_{k-1} \leq \varepsilon_k = o(k/n)$, which follows from (6.6). Specifically, by (6.6), $\varepsilon_k/(k/n) = O(k^{-3/4} n^{1/4} + k^{-1} n^{9/25})$, which by $k \geq n^{4/10}$ is $O(n^{-0.3} n^{0.25} + n^{-4/10} n^{0.36}) = o(1)$.

Likewise, by Claim 6.2, inequality (6.10) holds w.h.p. simultaneously for all $k \in [n - 14\sqrt{n}, n - 2]$. Again, we need only show that $\varepsilon_k = o(k/n)$, which holds because here $k/n = \Theta(1)$ while by definition $\varepsilon_k = o(1)$. ■

We now prove the two claims, by filling in the details for Sections 5.2 and 5.6.

6.2 | Structure R

With reference to Section 5.2, all that we need to confirm is that $k + r_0 \leq n - 1$. For Claim 6.1, by hypothesis $k \leq n - 14\sqrt{n}$, and provided that $1.01C_\varepsilon C_B^{1/4} \leq 13$ (check), by (6.7) $\varepsilon_k \leq 13n^{-1/2}$, whereupon $r_0 = \varepsilon_k n \leq 13\sqrt{n}$. For Claim 6.2, with $r_0 = 1$, $k + r_0 \leq n - 1$ is immediate.

6.3 | Path weights

With reference to Section 5.3, we establish that the bound (5.3) holds w.h.p. simultaneously for all $k \geq n^{4/10}$. With $W_{(k)}$ representing the cost of the k th cheapest edge incident on some fixed vertex (which we will take to be s and then t in turn), it suffices to show that

$$W_{(k+r_0)} \leq W_{(k+1)} + 1.1\varepsilon_k \quad (6.11)$$

holds with high probability for all $k \geq n^{4/10}$.

For Claim 6.2, with $r_0 = 1$, (6.11) is immediate. For Claim 6.1, with $r_0 = \varepsilon_k n$, generate the variables $W_{(k)}$ by placing $n - 1$ points uniformly at random on the unit interval I , associating $W_{(k)}$ with the k th smallest point. It suffices to show that, w.h.p., each interval $(W_{(k+1)}, W_{(k+1)} + 1.1\varepsilon_k)$ contains at least r_0 points. For all $k \in [n^{4/10}, n - 14\sqrt{n} - 1]$, Claim 6.1 has $\varepsilon_k \geq n^{-0.99}$ by (6.8), so Lemma 4.1 shows that w.p. $1 - \exp(-\Omega(n^{0.01}))$, every interval of length $\geq 1.1\varepsilon_k$ in $[0, 1]$ contains at least $r_0 \stackrel{=}{=} \varepsilon_k n$ points, and in particular this holds for all the intervals $(W_{(k+1)}, W_{(k+1)} + 1.1\varepsilon_k)$.

We assume henceforth that the graph G is “good” in the sense that (6.11) holds for all $k \geq n^{4/10}$ for vertices s and t , and that for all $k \leq n^{4/10}$ we have the upper bounds on X_k from (1.3), as proved to hold w.h.p. in Section 3.

6.4 | Adversary

With reference to Section 5.4, we need only verify (5.9), and this will be done in the next subsection.

6.5 | Budgets B_k

With reference to Section 5.5, we first establish (5.12). This follows from

$$W_{(k+1)}^s - W_{(k)}^s \leq 0.1\varepsilon_k. \quad (6.12)$$

The reasoning for this is the same as for (6.11): each interval of length $0.1\varepsilon_k$ contains at least one point. The parameters are trivial to check.

Next, we show that the parameters of Claim 6.1 satisfy (5.9), for which as argued in Section 5.5 it suffices to show that they satisfy (5.10) and (5.11). We start with (5.11), the base case. Here $k = n^{4/10}$, $B_k \geq 3n^{-2/10}$ from (6.5), and $U_k = 3n^{-2/10}$ from (5.7), establishing (5.11).

To establish (5.10), first note that $\frac{\partial}{\partial k} B_k = \frac{5}{4} C_B B_k^{1/5} n^{-3/5}$ is an increasing function. Then,

$$B_{k+1} - B_k \geq \frac{\partial}{\partial k} B_k = \frac{5}{4} C_B B_k^{1/5} n^{-3/5} = \frac{5}{4} \frac{C_B}{C_\epsilon} \epsilon_k \geq 8\epsilon_k,$$

since $C_B \geq \frac{8.4}{5} C_\epsilon$ (check).

We now establish (5.9) for the parameters of Claim 6.2. With $k^* = \lfloor n - 14\sqrt{n} \rfloor$, the point where Claim 6.1 ends and just before Claim 6.2 begins, the previous case showed that $B_{k^*} \geq U_{k^*} - I_{k^*}$, and by (6.4) $B_{k^*} \leq 77\sqrt{n}$. Then, for k from $k^* + 1$ to $n - 2$,

$$\begin{aligned} U_k - I_k &= (U_{k^*} - I_{k^*}) + [(U_k - U_{k^*}) - (I_k - I_{k^*})] \\ &\leq B_{k^*} + \left[\sum_{i=k^*+1}^k (W_{(i)}^s + W_{(i)}^t + 7\epsilon_{i-1}) - \sum_{i=k^*}^{k-1} (W_{(i)}^s + W_{(i)}^t) \right] \quad (\text{see (5.7) and (5.8)}) \\ &\leq B_{k^*} + \sum_{i=k^*+1}^{n-2} 7\epsilon_{i-1} + (W_{(k)}^s + W_{(k)}^t - W_{(k^*)}^s - W_{(k^*)}^t) \\ &\leq 77\sqrt{n} + (14\sqrt{n}) \cdot 7C'_\epsilon n^{-1/6} + 2 \quad (\text{see (6.9)}) \\ &\leq 78\sqrt{n} \\ &\leq B_k \quad (\text{see (6.9)}), \end{aligned} \tag{6.13}$$

using that $C'_B \geq 78$ (check).

6.6 | Minimum of two binomial variables

Before addressing robustness of the structure R , we require a lemma (Lemma 6.4) on the minimum Z of two i.i.d. binomial $\text{Bi}(n, p)$ random variables. There is a genuine difference in the cases when the common mean $\lambda = np$ is large or small: if λ is large then Z is likely to be close to λ , making $\mathbb{E}Z = \Theta(\lambda)$; if λ is small then Z will most often be 0, occasionally 1 (with probability about λ^2), and rarely anything larger, making $\mathbb{E}Z = \Theta(\lambda^2)$. The lemma relies on the following property of the median of a binomial random variable. (A weaker form of (6.14) and thus of Lemma 6.4 can be obtained from Lemma 3.2 in lieu of using the median.)

Theorem 6.3 (Hamza [11, Theorem 2]). *A binomial random variable X has median satisfying $|\text{Med}(X) - \mathbb{E}X| \leq \ln 2$.*

In this discrete setting $\text{Med}(X)$ is not unique: it can be any value m for which $\mathbb{P}(X \leq m) \geq 1/2$ and $\mathbb{P}(X \geq m) \geq 1/2$. [11] defines it uniquely as the smallest integer m such that $\mathbb{P}(X \leq m) > 1/2$; as desired, this gives $\mathbb{P}(X \geq \text{Med}(X)) = 1 - \mathbb{P}(X \leq \text{Med}(X) - 1) \geq 1 - 1/2 = 1/2$. (For other results on the binomial median see Kaas and Buhrman [16], in particular, Corollary 1. Stronger results for the Poisson distribution are given by Choi [6], proving a conjecture of Chen and Rubin, and by Adell and Jodrá [1].)

Lemma 6.4. *Let Z_1, Z_2 be i.i.d. $\text{Bi}(n, p)$ random variables, $Z := \min(Z_1, Z_2)$ and $\lambda := \mathbb{E}Z_1 = np$.*

(1) If $\lambda \geq 2$, then

$$\mathbb{P}(Z \geq 0.65\lambda) > 1/4. \tag{6.14}$$

(2) If $\lambda \leq 2$, then

$$\mathbb{P}(Z \geq 1) > 0.18\lambda^2. \tag{6.15}$$

Proof. In the first case,

$$\text{Med}(Z_1) \geq \lambda - \ln 2 = \frac{\lambda - \ln 2}{\lambda} \lambda \geq \frac{2 - \ln 2}{2} \lambda \geq 0.65\lambda,$$

so $\mathbb{P}(Z_1 \geq 0.65\lambda) \geq \mathbb{P}(Z_1 \geq \text{Med}(Z_1)) \geq 1/2$. The same holds of course for Z_2 , and the result follows by independence.

In the second case we again use independence, and here

$$\mathbb{P}(Z_1 \geq 1) = 1 - (1 - p)^n \geq 1 - \exp(-\lambda) = \frac{1 - \exp(-\lambda)}{\lambda} \cdot \lambda \geq 0.43\lambda.$$

The last inequality comes from minimizing $\frac{1 - \exp(-x)}{x}$ over $0 \leq x \leq 2$; the function is decreasing so the minimum is at $x = 2$. ■

6.7 | Robustness in Claim 6.1

With reference to Section 5.6, let us complete the robustness argument for Claim 6.1, showing that (5.14) holds with high probability. Here we have taken $r_0 = \epsilon_k n$, so that the number of edges from a middle vertex to V'_s (see (5.15)) is $Z_v^s \sim \text{Bi}(\epsilon_k n, \epsilon_k)$, with mean $\lambda = r_0 \epsilon_k = \epsilon_k^2 n$ (see (5.16)).

Recall that if λ is small we expect (see (5.17)) that to destroy all paths the adversary will have to delete edges of total weight at least $\epsilon_k n \lambda^2 = \epsilon_k^5 n^3$, which will exceed B_k . And, if λ is large, then each Z_v will have expectation close to $\lambda = \epsilon_k^2 n$, for a total cost $\epsilon_k n$ times larger, namely $\epsilon_k^3 n^2$, and again this exceeds B_k . We now replace these rough calculations with detailed probabilistic ones, applying Lemma 6.4 to Z_v in the two cases of λ small and large.

For the adversary to delete all s - t paths via v , he must delete at least

$$Z_v := \min(Z_v^s, Z_v^t)$$

edges, and to destroy all paths he must delete at least

$$N := \sum_{v \in M'} Z_v$$

edges. As described in Section 5.6, we imagine a fixed deletion of k edges on each of s and t giving neighbor sets V'_s and V'_t and a set M' of middle vertices; we will eventually take a union bound over all such choices.

If $\lambda \geq 2$, then by Lemma 6.4, for each $v \in M'$, $\mathbb{P}(Z_v^s \geq 0.65\lambda) \geq 1/4$. Thus, N stochastically dominates $0.65\lambda \cdot \text{Bi}(0.99n, 1/4)$, with expectation $> 0.1608\lambda n$. We shall consider it a *failure* if $N \leq 0.16\lambda n$.

Assuming success, since each edge costs at least ε_k to delete, it costs at least $0.16\varepsilon_k\lambda n = 0.16\varepsilon_k^3 n^2$ to delete them all. This exceeds B_k :

$$\begin{aligned} \frac{0.16 \cdot \varepsilon_k^3 n^2}{B_k} &= 0.16 \cdot C_\varepsilon^3 n^{-9/5} B_k^{-2/5} n^2 \quad (\text{by definition of } \varepsilon_k) \\ &\geq 0.15 \cdot C_\varepsilon^3 C_B^{-1/2} n^{1/5} n^{-1/5} \quad (\text{by (6.4)}) \\ &> 1, \end{aligned}$$

using that $0.15 \cdot C_\varepsilon^3 C_B^{-1/2} > 1$ (check).

Failure means that $N/(0.65\lambda) \sim \text{Bi}(0.99n, 1/4) \leq (0.16/0.65)n$. Noting that $0.99 \cdot 1/4 > 0.16/0.65$, by Lemma 3.2, the probability of failure is $\exp(-\Omega(n))$. By the union bound, the total of the failure probabilities, over all rounds (values of k) and all adversary choices of the k root edges at s and t , is small:

$$\begin{aligned} \sum_k \binom{k+r_0}{r_0}^2 \cdot \exp(-\Omega(n)) & \tag{6.16} \\ & \leq \sum_k (n^{r_0})^2 \exp(-\Omega(n)) \\ & = \sum_k \exp(2\varepsilon_k n \ln n - \Omega(n)) \quad (\text{by } r_0 = \varepsilon_k n) \\ & \leq n \exp(-\Omega(n)) \quad (\text{using } \varepsilon_k n = O(n^{1/2}) \text{ from (6.7)}) \\ & = o(1). \end{aligned}$$

If $\lambda < 2$, then by Lemma 6.4 N stochastically dominates $\text{Bi}(0.99n, 0.18\lambda^2)$, with expectation $> 0.175\lambda^2 n$. We shall consider it a *failure* if $N \leq 0.17\lambda^2 n = 0.17\varepsilon_k^4 n^3$. Each edge costs at least ε_k to delete. Assuming success, it thus costs at least $0.17\varepsilon_k^5 n^3$ to delete them all, which exceeds B_k :

$$\begin{aligned} \frac{0.17\varepsilon_k^5 n^3}{B_k} &= 0.17C_\varepsilon^5 \quad (\text{by definition of } \varepsilon_k) \\ &> 1, \end{aligned}$$

using that $0.17C_\varepsilon^5 > 1$ (check).

By Lemma 3.2, the probability of failure is

$$\mathbb{P}(N \leq 0.17\varepsilon_k^4 n^3) = \exp(-\Omega(\varepsilon_k^4 n^3)). \tag{6.17}$$

Over all rounds (values of k) and adversary choices of edges incident to s and t , the total failure probability is at most

$$\begin{aligned} \sum_k \binom{k+r_0}{r_0}^2 \cdot \mathbb{P}(N < 0.17\varepsilon_k^4 n^3) & \\ & \leq \sum_k \exp(2\varepsilon_k n \ln n - \exp(-\Omega(\varepsilon_k^4 n^3))) \\ & \leq n \exp(-\Omega(\varepsilon_k^4 n^3)), \end{aligned}$$

because $\varepsilon_k n \ln n$ is dominated by $\varepsilon_k^4 n^3$: the latter is larger by a factor $\varepsilon_k^3 n^2 / \ln n$, which by (6.8) is $\Omega(n^{-48/25} n^2 / \ln n) = \Omega(n^{2/25} / \ln n) = \omega(1)$. Continuing, this is

$$\begin{aligned} &\leq n \exp(-\Omega(n^{11/25})) \quad (\text{invoking (6.8) again}) \\ &= o(1). \end{aligned} \tag{6.18}$$

6.8 | Robustness in Claim 6.2

Again, our aim is to establish robustness of R by showing that (5.14) holds with high probability, and the argument is similar to but simpler than that of Section 6.7.

Since $r_0 = 1$, both V'_s and V'_t have size 1. For a vertex $v \in M'$, let Z_v be the number of paths from V'_s to V'_t via v . There is only one such possible path, hence

$$Z_v \sim \text{Bernoulli}(\varepsilon_k^2).$$

To destroy all s - t paths the adversary must delete at least

$$N := \sum_{v \in M'} Z_v$$

edges. N stochastically dominates $\text{Bi}(0.99n, \varepsilon_k^2)$, which has expectation $0.99\varepsilon_k^2 n$. We declare the event $N \leq 0.98\varepsilon_k^2 n$ a *failure*. Assuming success, destroying all s - t paths would cost at least $\varepsilon_k N \geq 0.98\varepsilon_k^3 n$. This exceeds B_k , since

$$\frac{0.98\varepsilon_k^3 n}{B_k} = \frac{0.98C'_\varepsilon{}^3}{C'_B},$$

and $0.98C'_\varepsilon{}^3 > C'_B$ (check).

The probability of failure is

$$\mathbb{P}(N \leq 0.98\varepsilon_k^2 n) = \exp(-\Omega(\varepsilon_k^2 n)) = \exp(-\Omega(n^{2/3})). \tag{6.19}$$

Over all rounds and adversary choices, using that $\binom{k+r_0}{r_0} = \binom{k+1}{1} \leq n$, the total failure probability is at most

$$\begin{aligned} &\sum_k \binom{k+r_0}{r_0}^2 \cdot \mathbb{P}(N \leq 0.98\varepsilon^2 n) \\ &\leq (14\sqrt{n}) n^2 \exp(-\Omega(n^{2/3})) \quad (\text{by (6.19)}) \\ &= o(1). \end{aligned} \tag{6.20}$$

7 | LOWER BOUND

In this section, we establish the lower bound in (1.3) of Theorem 1.1. Section 7.1 establishes the lower bound on X_k directly for $k \leq \sqrt{\ln n}$. Values $k \geq \sqrt{\ln n}$ are treated in the subsequent parts. In Section 7.2, Lemma 7.1 establishes a lower bound on the running totals S_k ,

$$S_k := \sum_{i=1}^k X_i. \tag{7.1}$$

In Section 7.3, Lemma 7.2 obtains a lower bound on X_k using Lemma 7.1's lower bound on S_k , the previously established upper bound on X_k from Theorem 1.1, and the monotonicity of X_k .

7.1 | Lower bound for small k

We begin with $k \leq \sqrt{\ln n}$. For any fixed $\varepsilon > 0$, we know from [14] that w.h.p.

$$X_1 > (1 - \varepsilon/2) \frac{\ln n}{n}. \tag{7.2}$$

Assuming that (7.2) holds, it follows immediately, and deterministically, that for all $k \leq \sqrt{\ln n}$,

$$X_k \geq X_1 \geq (1 - \varepsilon/2) \frac{\ln n}{n} \geq (1 - \varepsilon) \frac{2k + \ln n}{n}. \tag{7.3}$$

The first inequality holds because the sequence X_k is monotone increasing, the next by assumption on X_1 , the next by $k = o(\ln n)$.

7.2 | Lower bound on the running totals

Lemma 7.1. *For any $\varepsilon > 0$, w.h.p., simultaneously for every $k \leq n - 1$,*

$$S_k \geq (1 - \varepsilon) \sum_{i=1}^k \left(\frac{2i + \ln n}{n} \right). \tag{7.4}$$

Proof. Write $W_{(i)}^s$ and $W_{(i)}^t$ for the order statistics of edge weights out of s and t , respectively. By Lemma 4.2, w.h.p.,

$$W_{(k)}^s, W_{(k)}^t \in \left[(1 - \varepsilon/2) \frac{k}{n}, (1 + \varepsilon/2) \frac{k}{n} \right] \quad \text{for all } k \geq \sqrt[3]{\ln n}, \tag{7.5}$$

and we will assume throughout the proof that (7.5) holds.

We prove the assertion in two ranges of k .

For $\ln^{11/10} n \leq k \leq n - 1$, the k paths must use at least $k - 1$ edges on each of s and t , all distinct (k edges each, ignoring the edge $\{s, t\}$ if it is used). Then, using (7.5), we get that w.h.p., for all k in the range,

$$\begin{aligned} S_k &\geq \sum_{i=1}^{k-1} \left(W_{(i)}^s + W_{(i)}^t \right) \geq \sum_{i=\sqrt[3]{\ln n}}^{k-1} (1 - \varepsilon/2) \frac{2i}{n} \\ &= (1 - \varepsilon/2) \left(\sum_{i=1}^k \frac{2i + \ln n}{n} - \sum_{i=1}^k \frac{\ln n}{n} - \sum_{i=1}^{\sqrt[3]{\ln n}-1} \frac{2i}{n} - \frac{2k}{n} \right) \end{aligned} \tag{7.6}$$

$$\geq (1 - o(1))(1 - \varepsilon/2) \sum_{i=1}^k \frac{2i + \ln n}{n} \quad (\text{see below}) \tag{7.7}$$

$$\geq (1 - \varepsilon) \sum_{i=1}^k \left(\frac{2i + \ln n}{n} \right). \tag{7.8}$$

To justify (7.7) it suffices to show that the first sum in (7.6) is of strictly larger order than the other terms. The first sum is at least $\sum_{i=k/2}^k 2i/n = \Omega(k^2/n)$, which since $k \geq \ln^{11/10} n$ is also $\Omega(k \ln^{11/10} n/n)$ and $\Omega(\ln^{22/10} n/n)$; we will use all three formulations. The second term is of order $O(k \ln n/n)$, negligible compared with the middle formulation. The third term is $O(\ln^{2/3} n/n)$, negligible compared with the last formulation. And the fourth term, of order $O(k/n)$, is negligible compared with the first formulation.

For $1 \leq k \leq \ln^{11/10} n$, let $\delta = \varepsilon/3$ and let $G' = G - s - t$. Let N_s and N_t be the endpoints of the cheapest $\ln^3 n$ edges out of s and t respectively. Note that these sets are independent of the edge weights of G' .

If any path P_i , $i \leq k$, uses a root edge (edge incident on s or t) *not* among the $\ln^3 n$ cheapest edges of s or t , then by (7.5) this edge costs at least $(1 - \varepsilon) \ln^3 n/n$, thus $S_k \geq (1 - \varepsilon) \ln^3 n/n$. Then (7.4) follows because this is larger than the RHS of (7.4), namely $\Theta((k^2 + k \ln n)/n) = O(\ln^{11/5} n/n)$ for this range of k . Thus we may assume that for all $i \leq k$, each path P_i goes via some $s' \in N_s$, $t' \in N_t$.

For $s' \in N_s$, $t' \in N_t$, define $A(s', t')$ to be the event that t' is one of the $(n - 2)^{1-\delta}$ nearest vertices of s' , by cost, in G' . Clearly, for each pair s', t' , $\mathbb{P}(A(s', t')) = (n - 2)^{-\delta}$. Let A be the union of these events, that is, the event that any such pair has this property. By the union bound,

$$\mathbb{P}(A) \leq (\ln^3 n)^2 (n - 2)^{-\delta} = o(1).$$

We assume henceforth that A does not hold: the $\ln^3 n$ cheapest root edges at s and t do not happen to sample any “nearest” pairs in G' .

By assumption that A does not hold, in the *exponential* model (where each edge is i.i.d. $\text{Exp}(1)$) for G' , for each $s' \in N_s$, $t' \in N_t$, the distance $d(s', t')$ stochastically dominates $Y \sim \sum_{i=1}^{n-2} \text{Exp}(i(n-2-i))$ by (3.20). We have $\mathbb{E}Y = (1 + o(1))(1 - \delta) \ln n/n$ by (3.21) (just adjusting its last equation where the value of d is substituted in). Applying Lemma 3.5’s (3.25) with $\mu = \mathbb{E}Y$ as above, $a^* = n - 3$, and $\lambda = 1 - \delta$, that in the exponential model G' ,

$$\mathbb{P}\left(d_{G'}(s', t') \leq (1 - \delta) \frac{(1 + o(1))(1 - \delta) \ln n}{n}\right) \leq \exp(-\Theta(n \cdot \ln n/n \cdot \delta^2)) = n^{-\Theta(1)}.$$

Since $(1 + o(1))(1 - \delta)^2 \geq (1 - \frac{3}{4}\varepsilon)$, by the union bound this implies, still in the exponential model G' ,

$$\mathbb{P}\left(\exists s' \in N_s, t' \in N_t : d_{G'}(s', t') \leq (1 - \frac{3}{4}\varepsilon) \ln n/n\right) \leq (\ln^3 n)^2 n^{-\Theta(1)} = o(1). \tag{7.9}$$

By standard coupling arguments (see Remark 1.6), this also implies that (7.9) holds in the *uniform* model G in which we are working.

Thus w.h.p., for all $s' \in N_s, t' \in N_t$, we have $d_{G'}(s', t') \geq (1 - \frac{3}{4}\varepsilon) \ln n/n$; assume this holds. We already assumed that each path P_i , $i \leq k$, goes via some $s' \in N_s, t' \in N_t$, so its nonroot edges contribute at least $d_{G'}(s', t') \geq (1 - \frac{3}{4}\varepsilon) \ln n/n$ to S_k . Then, for all k in this range,

$$\begin{aligned} S_k &\geq \sum_{i=1}^k \left(1 - \frac{3}{4}\varepsilon\right) \frac{\ln n}{n} + \sum_{i=1}^{k-1} \left(W_{(i)}^s + W_{(i)}^t\right) \\ &\geq \sum_{i=1}^k \left(1 - \frac{3}{4}\varepsilon\right) \frac{\ln n}{n} + \left(1 - \frac{1}{2}\varepsilon\right) \sum_{i=\frac{3}{\sqrt{\ln n}}}^{k-1} \frac{2i}{n} \quad (\text{by (7.5)}) \\ &\geq (1 - \varepsilon) \sum_{i=1}^k \left(\frac{2i + \ln n}{n}\right). \end{aligned} \tag{7.10}$$

To justify the final inequality, rewrite the second sum in (7.10) as $\sum_{i=1}^k \frac{2i}{n} - \frac{2k}{n} - \sum_{i=1}^{\sqrt[3]{\ln n}-1} \frac{2i}{n}$ and observe that both its second term, $2k/n$, and its final term, which is of order $O(\sqrt[3]{\ln n}^2/n)$, are negligible compared with the first sum in (7.10), which is of order $\Omega(k \ln n/n)$. ■

7.3 | Lower bound for large k

Lemma 7.2. For any $\varepsilon > 0$, w.h.p., simultaneously for every $k \in [\sqrt{\ln n}, n-1]$,

$$X_k \geq (1 - \varepsilon) \left(\frac{2k + \ln n}{n} \right).$$

Proof. Let $\delta = \varepsilon^2/9$ and define

$$c_k = \frac{2k + \ln n}{n}, \quad L_k = (1 - \delta) \sum_{i=1}^k c_i, \quad U_k = (1 + \delta) \sum_{i=1}^k c_i. \quad (7.11)$$

W.h.p., simultaneously for all k , $S_k \geq L_k$ (by Lemma 7.1) and $S_k \leq U_k$ (by the upper bound of Theorem 1.1, already proved). Henceforth, assume that both hold, so $L_k \leq S_k \leq U_k$. The rest of the argument is deterministic. For any positive integer $t < k$, using that X_k is monotone increasing, we have

$$\begin{aligned} tX_k &\geq X_k + \dots + X_{k-t+1} \\ &= S_k - S_{k-t} \\ &\geq L_k - U_{k-t}. \end{aligned} \quad (7.12)$$

Thus

$$\begin{aligned} X_k &\geq \frac{1}{t} (L_k - U_{k-t}) = \frac{1}{t} \left((1 - \delta) \sum_{i=1}^k c_i - (1 + \delta) \sum_{i=1}^{k-t} c_i \right) \\ &\geq \frac{1}{t} \left(\sum_{i=k-t+1}^k c_i - 2\delta \sum_{i=1}^k c_i \right) \geq \frac{1}{t} (tc_{k-t} - 2\delta kc_k) = c_{k-t} - \frac{2\delta kc_k}{t} \\ &= c_k - \frac{2t}{n} - \frac{2\delta kc_k}{t} \\ &\geq c_k - \frac{tc_k}{k} - \frac{2\delta k}{t} c_k \quad (\text{using that } c_k/k > 2/n) \\ &= c_k \left(1 - \frac{t}{k} - \frac{2\delta k}{t} \right). \end{aligned}$$

Ignoring integrality for a moment, setting $t = k\sqrt{2\delta}$ would make the last expression $c_k(1 - 2\sqrt{2\delta})$. Since this $t = \Theta(k) = \omega(1)$, rounding it can be seen to change the expression by a factor $1 + o(1)$, so we may safely write

$$X_k \geq c_k(1 - 3\sqrt{\delta}) = (1 - \varepsilon) \frac{2k + \ln n}{n}. \quad \blacksquare$$

8 | EXPONENTIAL MODEL

In this section we prove Theorem 1.2, the analogue of Theorem 1.1 for exponentially distributed edge weights.

For small k , results for the exponential case follow from those for the uniform. We first argue that the upper bound of Theorem 1.1 also holds in the exponential case for any $k = o(n)$. Couple the two models, so that any edge of weight $w = o(1)$ in one model has cost $w' = w(1 + o(1))$ in the other. The uniform-model upper-bound constructions in Section 3 (for $k = o(n^{1/2})$) and Sections 5 and 6 (for larger k) only use edges of weight $o(1)$ (when $k = o(n)$), and therefore the same upper bounds hold for the exponential model; the multiplicative difference of $1 + o(1)$ can be subsumed into the factor $1 + \varepsilon$ already present. (In the construction of Sections 5 and 6, the “middle edges” are of cost $o(1)$ for *all* k , but the “incident edges” have larger cost for k large. In particular, for large k , (6.11) will no longer hold in the exponential case until we adjust r_0 and ε_k appropriately.)

For the lower bound too, the argument in Section 7 carries over for all $k = o(n)$. The lower bounds L_k on the prefix sums S_k derived in Sections 7.1 and 7.2 carry over to the exponential case because the edge costs are equal to within $1 + o(1)$ factors in the two models. The upper bounds U_k on the prefix sums are simply the sums of the individual upper bounds on X_k , and we have just argued that these change only by a $1 + o(1)$ factor. Section 7.3 only uses L_k and U_k to derive lower bounds on X_k , so with these both changed only by $1 + o(1)$ factors, its results carry over verbatim.

Our task, then, is to prove the upper and lower bounds in Theorem 1.2 for larger k . For the upper bound, arguing for $k > n^{0.4}$ (there is no advantage to a larger starting value), we use the same approach as for the uniform model in Section 5.

For the lower bound, we argue for $k \geq n^{9/10}$. Unfortunately, the method used in Section 7 for the uniform distribution does not extend; let us explain why. The lower bound there came from (7.12), $tX_k \geq L_k - U_{k-t}$, valid for any functions L and U with $L_k \leq S_k \leq U_k$. Here, we would take L_k as the sum I_k of incident edges as in (5.8) and U_k as the sum of the X_k upper bounds as in (5.4). Recall that we defined B_k so that $B_k \geq U_k - L_k$, as in (5.9). Then we can rewrite the previous lower bound approach as $X_k \geq \frac{1}{t}(L_k - U_{k-t}) \geq \frac{1}{t}(L_k - L_{k-t}) + \frac{1}{t}(L_{k-t} - U_{k-t}) \geq \frac{1}{t} \sum_{i=k-t}^{k-1} W_{(i)} - \frac{1}{t} B_{k-t}$. For large k , $W_{(k)}$ and therefore X_k are $\Theta(\ln n)$. Since the B_k grow to size $\Theta(n^{1/2})$ (in the exponential case as well as the uniform case), we are thus limited by the second term to $t = \Omega(n^{1/2+o(1)})$. However, from (4.3), such a large value of t would mean that the average given by the first term is significantly different from $W_{(k)}$.

The desired lower bound would be immediate if we could claim that P_k necessarily used the k th cheapest edge on s (of cost $W_{(k)}^s$) or a later one, and likewise for t . We will prove something close to this. We argue in Section 8.7 that every pair of vertices (excluding both s and t) is joined by a path of cost at most δ (for some small δ to be specified) that is edge-disjoint from *all* P_i , $i = 1, \dots, n-1$. We will show that this implies that path P_k uses an edge on s that is at most δ cheaper than $W_{(k)}^s$, and likewise for t , yielding a sufficient lower bound.

8.1 | Claims, and implications for Theorem 1.2

In order to establish upper bounds on X_k in the exponential model, we use the same structure $R^{(k)}$ as described in Section 5.2. Then (5.4) follows as before, and we can continue to define U_k as in (5.7). For convenience define

$$\bar{k} = n - k. \quad (8.1)$$

As before we will treat k in two ranges, and we start now with the smaller range.

Claim 8.1. For $k \in [n^{4/10}, n - \sqrt{n}]$, let

$$B_k := \left(\frac{2n^{1/25} + C_B(n^{3/5} - \bar{k}^{3/5})}{n^{1/5}} \right)^{5/4} \quad \text{and} \quad \varepsilon_k := C_\varepsilon B_k^{1/5} n^{-1/5} \bar{k}^{-2/5}, \quad (8.2)$$

with $C_B = 44$ and $C_\varepsilon = 4$. Then, asymptotically almost surely,

$$X_{k+1} \leq W_{(k+1)}^s + W_{(k+1)}^t + 8\varepsilon_k. \tag{8.3}$$

Remark: In proving Claim 8.1 we will set

$$r_0 := \varepsilon_k \bar{k}. \tag{8.4}$$

because it roughly equates $W_{(k+r_0)} - W_{(k)}$ and ε_k ; see (4.3). In this regime integrality is not an issue: r_0 is large, per (8.9).

It is clear that both B_k and ε_k in (8.2) are increasing in k , even over the larger range $k \in [0, n]$. We will make use of the following bounds, holding for n sufficiently large. Here, (8.5) uses that at $k = n - \Theta(\sqrt{n})$, $\bar{k}^{3/5}$ dominates $2n^{1/25}$, while (8.6) takes $k = 0$.

$$B_k \leq B_{n-\sqrt{n}} \leq C_B^{5/4} n^{1/2} \tag{8.5}$$

$$B_k \geq B_{n^{4/10}} \geq 2n^{-1/5} \tag{8.6}$$

$$\varepsilon_k \leq C_\varepsilon B_k^{1/5} n^{-1/5} n^{-1/2 \cdot 2/5} \leq C_\varepsilon C_B^{1/4} n^{-3/10} \tag{8.7}$$

$$\varepsilon_k \geq C_\varepsilon (B_{n^{4/10}})^{1/5} n^{-1/5} \bar{k}^{-2/5} \geq C_\varepsilon n^{-6/25} \bar{k}^{-2/5} \tag{8.8}$$

$$r_0 = \bar{k} \varepsilon_k \stackrel{(8.8)}{\geq} C_\varepsilon n^{-6/25} \bar{k}^{3/5} \geq C_\varepsilon n^{3/50}. \tag{8.9}$$

Claim 8.2. For $k \in (n - \sqrt{n}, n - 2]$, let

$$B_k := C'_B \sqrt{n} \quad \text{and} \quad \varepsilon_k := C'_\varepsilon n^{-1/6}, \tag{8.10}$$

with $C_B = 115$ and $C_\varepsilon = 5$. Then, asymptotically almost surely, simultaneously for all k in this range,

$$X_{k+1} \leq W_{(k+1)}^s + W_{(k+1)}^t + 8\varepsilon_k. \tag{8.11}$$

Remark: In proving Claim 8.2 we will set

$$r_0 := 1. \tag{8.12}$$

As in Claim 6.2, B_k and ε_k are constants independent of k , but we retain the subscript for consistency with the notation of Section 5.1.

Proof of the upper bounds in Theorem 1.2. Analogous to the argument in Section 6.1, it is sufficient to check that $\varepsilon_k = o(\mathbb{E}W_{(k)})$. Since $\mathbb{E}W_{(k)} \sim \ln\left(\frac{n}{n-k}\right) \geq \frac{k}{n}$ (see (4.3)), it is enough to show that $\varepsilon_k = o(k/n)$.

For $k \leq n^{0.99} = o(n)$, by first-order approximation,

$$n^{3/5} - \bar{k}^{3/5} \stackrel{\text{def}}{=} n^{3/5} - (n-k)^{3/5} \sim \frac{3}{5} n^{-2/5} k, \tag{8.13}$$

so $B_k = \Theta\left(n^{-1/5} + n^{-3/4} k^{5/4}\right)$. Hence, from Claim 8.1, specifically 8.2,

$$\varepsilon_k = \Theta\left((n^{-1/25} + n^{-3/20} k^{1/4}) n^{-1/5} n^{-2/5}\right) = \Theta\left(n^{-16/25} + n^{-3/4} k^{1/4}\right) = o(k/n) \tag{8.14}$$

as $k \geq n^{4/10}$.

For $k > n^{0.99}$, we have in Claim 8.1 that $\varepsilon_k = O(n^{-3/10})$ by (8.5), and so $\varepsilon_k = o(k/n)$, while in Claim 8.2, $\varepsilon_k = \Theta(n^{-1/6}) = o(k/n)$. ■

8.2 | Path weights

To show inequality (5.3) it suffices to show that

$$W_{(k+r_0)} - W_{(k+1)} \leq 1.1\epsilon_k. \quad (8.15)$$

In Claim 8.2, we have defined $r_0 := 1$, so (8.15) is trivial. For Claim 8.1, $\Delta := W_{(k+r_0)} - W_{(k+1)}$ has the same distribution as $\sum_{i=k+2}^{k+r_0} X(n-i)$, where $X(a) \sim \text{Exp}(a)$ and these variables are all independent. Thus Δ is stochastically dominated by the sum of $r_0 - 1$ independent random variables $X(\bar{k} - r_0)$. Since $r_0 = \bar{k}\epsilon_k$, we have that $\mathbb{E}\Delta \leq r_0/(\bar{k} - r_0) = \epsilon_k/(1 - \epsilon_k)$, and from Lemma 3.5 it follows that $\mathbb{P}(\Delta > 1.1\epsilon_k) = O(\exp(-\Theta(r_0)))$. From (8.9), by the union bound, there is a negligible chance that (5.3) fails in any round.

8.3 | Budgets in Claim 8.1

As before, we need to define a B_k satisfying (5.9) and, as before, ϵ_k can be guessed from (5.18), then checked to satisfy yield robustness as in Section 6.7 and 6.8. The base case, confirming (5.11), is given by $k = n^{4/10}$, where by (8.6)

$$B_k \geq 2n^{-1/5} \stackrel{\text{def}}{=} U_k. \quad (8.16)$$

To verify (5.10), it is straightforward to check that $\frac{\partial^2}{\partial k^2} B_k$ is positive, so $\frac{\partial}{\partial k} B_k$ is increasing, and

$$B_{k+1} - B_k \geq \frac{\partial}{\partial k} B_k = \frac{5}{4} \frac{3}{5} \frac{C_B}{C_\epsilon} = \frac{3}{4} \frac{C_B}{C_\epsilon} \geq 8\epsilon_k,$$

since $\frac{3}{4} \frac{C_B}{C_\epsilon} \geq 8$ (check). Finally, we establish (5.12). We show that w.h.p. for all k in the range,

$$\Delta := W_{(k+1)}^s - W_{(k)}^s \leq 0.1\epsilon_k. \quad (8.17)$$

Note that $\Delta \sim \text{Exp}(\bar{k} - 1)$, so

$$\mathbb{P}(\Delta > 0.1\epsilon_k) = \exp(-0.1\epsilon_k \cdot (\bar{k} - 1)) = \exp(-\Omega(r_0)) = \exp(-n^{\Omega(1)})$$

by (8.9). Then, by the union bound there is a negligible chance that (8.17) fails for any k .

8.4 | Robustness in Claim 8.1

With reference to Section 5.6, we complete the robustness argument for Claim 8.1, showing that (5.14) holds with high probability. Here we have taken $r_0 = \epsilon_k \bar{k}$, so the number of edges from a middle vertex to V'_S (see (5.15)) is $Z'_v \sim \text{Bi}(\epsilon_k \bar{k}, \epsilon_k)$, with mean

$$\lambda = r_0 \epsilon_k = \epsilon_k^2 \bar{k} \quad (8.18)$$

(see (5.16)). Recall that if λ is small we expect (see (5.17)) that to destroy all paths the adversary will have to delete edges of total weight at least $\epsilon_k n \lambda^2 = \epsilon_k^5 n \bar{k}^2$, which will exceed B_k . And, if λ is large, then each Z_v will have expectation close to $\lambda = \epsilon_k^2 \bar{k}$, for a total cost $\epsilon_k n$ times larger, namely $\epsilon_k^3 n \bar{k}$, and again this exceeds B_k .

We now show the details of these rough calculations, including the probabilistic details, applying Lemma 6.4 to Z_v in the two cases of λ small and large.

For the adversary to delete all s - t paths via v , he must delete at least

$$Z_v := \min(Z_v^s, Z_v^t)$$

edges, and to destroy all paths he must delete at least

$$N := \sum_{v \in M'} Z_v$$

edges. As described in Section 5.6, we imagine a fixed deletion of k edges on each of s and t , giving neighbor sets V'_s and V'_t and a set M' of middle vertices, eventually taking a union bound over all such choices.

If $\lambda \geq 2$, then by Lemma 6.4, for each $v \in M'$, $\mathbb{P}(Z_v^s \geq 0.65\lambda) \geq 1/4$. Thus, N stochastically dominates $0.65\lambda \cdot \text{Bi}(0.99n, 1/4)$, with expectation $> 0.1608\lambda n$. We shall consider it a *failure* if $N \leq 0.16\lambda n$. Assuming success, since each edge costs at least ϵ_k to delete, it costs at least $0.16\epsilon_k \lambda n = 0.16\epsilon_k^3 n \bar{k}$ to delete them all. This exceeds B_k :

$$\begin{aligned} \frac{0.16 \epsilon_k^3 n \bar{k}}{B_k} &= 0.16 C_\epsilon^3 n^{-3/5} \bar{k}^{-6/5} B_k^{-2/5} n \bar{k} \quad (\text{by definition of } \epsilon_k) \\ &= 0.16 C_\epsilon^3 B_k^{-2/5} n^{2/5} \bar{k}^{-1/5} \\ &\geq 0.15 C_\epsilon^3 C_B^{-1/2} n^{1/5} n^{-1/5} \quad (\text{by (8.5)}) \\ &> 1, \end{aligned}$$

using that $0.15 \cdot C_\epsilon^3 C_B^{-1/2} > 1$ (check).

Failure means that $N/(0.65\lambda) \sim \text{Bi}(0.99n, 1/4) \leq (0.16\lambda n)/(0.65n) = (0.16/0.65)n$. Noting that $0.99 \cdot 1/4 > 0.16/0.65$, by Lemma 3.2, the probability of failure is $\exp(-\Omega(n))$. By the union bound, the total of the failure probabilities, over all rounds and all adversary choices of the k root edges at s and t , is small:

$$\begin{aligned} \sum_k \binom{k+r_0}{r_0}^2 \cdot \exp(-\Omega(n)) & \tag{8.19} \\ & \leq \sum_k (n^{r_0})^2 \exp(-\Omega(n)) \\ & = \sum_k \exp(2\epsilon_k \bar{k} \ln n - \Omega(n)) \quad (\text{by } r_0 = \epsilon_k \bar{k}) \\ & \leq n \exp(-\Omega(n)) = o(1), \end{aligned}$$

the penultimate inequality using $\epsilon_k \bar{k} = O(n^{7/10})$ by (8.7).

If $\lambda < 2$, then by Lemma 6.4 N stochastically dominates $\text{Bi}(0.99n, 0.18\lambda^2)$, with expectation $> 0.175\lambda^2 n$. We shall consider it a *failure* if $N \leq 0.17\lambda^2 n = 0.17\epsilon_k^4 n \bar{k}^2$. Each edge costs at least ϵ_k to delete. Assuming success, it thus costs at least $0.17\epsilon_k^5 n \bar{k}^2$ to delete them all, which exceeds B_k :

$$\begin{aligned} \frac{0.17\epsilon_k^5 n \bar{k}^2}{B_k} &= 0.17C_\epsilon^5 \quad (\text{by definition of } \epsilon_k) \\ &> 1, \end{aligned}$$

using that $0.17C_\epsilon^5 > 1$ (check).

By Lemma 3.2, the probability of failure is

$$\mathbb{P}(N \leq 0.17\epsilon_k^4 n \bar{k}^2) = \exp(-\Omega(\epsilon_k^4 n \bar{k}^2)). \quad (8.20)$$

Over all rounds and adversary choices of edges incident to s and t , the total failure probability is at most

$$\begin{aligned} & \sum_k \binom{k+r_0}{r_0}^2 \cdot \mathbb{P}(N < 0.17\epsilon_k^4 n \bar{k}^2) \\ & \leq \sum_k \exp(2\epsilon_k \bar{k} \ln n - \exp(-\Omega(\epsilon_k^4 n \bar{k}^2))) \\ & \leq n \exp(-\Omega(\epsilon_k^4 n \bar{k}^2)), \end{aligned}$$

because $\epsilon_k^4 n \bar{k}^2$ is larger than $\epsilon_k \bar{k}$ by a factor $\epsilon_k^3 n \bar{k}$, which by (8.8) is $\Omega(n^{-18/25} \bar{k}^{-6/5} n \bar{k}) = \Omega(n^{7/25} \bar{k}^{-1/5}) = \Omega(n^{2/25})$. Continuing, this is

$$\begin{aligned} & \leq n \exp(-\Omega(n^{1/25} \bar{k}^{2/5})) \quad (\text{invoking (8.8) again}) \\ & = o(1). \end{aligned} \quad (8.21)$$

8.5 | Budgets in Claim 8.2

We now establish (5.9) for the parameters of Claim 8.2. Section 8.3 showed that (5.9) holds for k up to $k^* := \lfloor n - \sqrt{n} \rfloor$, the point where Claim 8.1 ends and just before Claim 8.2 begins, so in particular $B_{k^*} \geq U_{k^*} - I_{k^*}$. For the regime of Claim 8.2, we redefine I_k from (5.8). Recall that I_k is a lower bound on the edges incident to s and t used by the first k paths. Previously, the sum defining I_k in (5.8) went to $k-1$ to avoid double counting the $\{s, t\}$ edge. In this regime, however, we need the sum to go k , as the $W_{(i)}$ increase rapidly. The weight of the $\{s, t\}$ edge is distributed as $\text{Exp}(1)$, thus w.h.p. it costs at most $n^{0.01}$. For $k > k^*$, define

$$I_k := \sum_{i=1}^k \left(W_{(k)}^s + W_{(k)}^t \right) - n^{0.01}, \quad (8.22)$$

so that w.h.p. I_k is a lower bound on the incident edges: the $n^{0.01}$ term resolves the potential double-counting of $\{s, t\}$. We are now ready to check that (5.9) holds. Following the derivation of (6.13), for k from $k^* + 1$ to $n-2$,

$$\begin{aligned} U_k - I_k &= (U_{k^*} - I_{k^*}) + [(U_k - U_{k^*}) - (I_k - I_{k^*})] \\ &\leq B_{k^*} + \sum_{k=k^*+1}^{n-2} 7\epsilon_k - (W_{(k^*)}^s + W_{(k^*)}^t - n^{0.01}) \quad (\text{see (5.7), (5.8), and (8.22)}) \\ &\leq 114\sqrt{n} + \sqrt{n} \cdot 7C'_\epsilon n^{-1/6} + n^{0.01} \quad (\text{see (8.5) and (8.10)}) \\ &\leq 115\sqrt{n} \\ &\leq B_k \quad (\text{see (8.10)}), \end{aligned} \quad (8.23)$$

using that $C'_B \geq 115$ (check).

8.6 | Robustness in Claim 8.2

Again, our aim is to establish robustness of R by showing that (5.14) holds with high probability, and the argument is similar to but simpler than that for robustness in Claim 8.1.

Since $r_0 = 1$, both V'_s and V'_t have size 1. For a vertex $v \in M'$, let Z_v be the number of paths from V'_s to V'_t via v . There is only one such possible path, hence

$$Z_v \sim \text{Bernoulli}(\epsilon_k^2).$$

To destroy all s - t paths the adversary must delete at least

$$N := \sum_{v \in M'} Z_v$$

edges. N stochastically dominates $\text{Bi}(0.99n, \epsilon_k^2)$, with expectation at least $0.99\epsilon_k^2 n$. We declare the event $N \leq 0.98\epsilon_k^2 n$ a *failure*. Assuming success, destroying all s - t paths would cost at least $\epsilon_k N \geq 0.98\epsilon_k^3 n$. This exceeds B_k , since by (8.10) and $0.98C'_\epsilon{}^3 > C'_B$ (check),

$$\frac{0.98\epsilon_k^3 n}{B_k} = \frac{0.98C'_\epsilon{}^3}{C'_B} > 1.$$

The probability of failure is

$$\mathbb{P}(N \leq 0.98\epsilon_k^2 n) = \exp(-\Omega(\epsilon_k^2 n)) = \exp(-\Omega(n^{2/3})). \tag{8.24}$$

Over all rounds and adversary choices, using that $\binom{k+r_0}{r_0} = \binom{k+1}{1} \leq n$, the total failure probability is at most

$$\begin{aligned} \sum_k \binom{k+r_0}{r_0}^2 \cdot \mathbb{P}(N \leq 0.98\epsilon^2 n) &\leq \sqrt{n} n^2 \exp(-\Omega(n^{2/3})) \quad (\text{by (8.24)}) \\ &= o(1). \end{aligned} \tag{8.25}$$

8.7 | Lower bound

As argued in the introduction of this section, for any $k = o(n)$, the lower bound follows from the uniform case. Thus it is sufficient if we show the lower bound for $k \geq n^{9/10}$, which we do now.

Remark 8.3. With high probability, for every pair of vertices u and v in $G' = G - s - t$, there is a u - v path in G' of cost at most $\delta = 20n^{-1/6}$ that is edge-disjoint from P_1, \dots, P_{n-1} .

Proof. The proof of Claim 8.2 showed that w.h.p., for all k in the claim's range (up to $k = n - 2$), there is a cheap s - t path (of cost given by (8.11)) disjoint from P_1, \dots, P_k , because for a given pair of neighbors u, v of s and t , there is a u - v path in G' that is edge-disjoint from these k paths and has cost at most $4\epsilon_k = 20n^{-1/6} \stackrel{\text{def}}{=} \delta$ (see (8.10)). The existence of a $k + 1$ st s - t path limits k to $n - 2$ since after that there are no new neighbors u and v of s and t , but the rest of the argument extends to $k = n - 1$.

In particular, extending the definition (8.10) of B_k and ϵ_k to $k = n - 1$, the derivation of (8.23) extends without change and shows that the budget B_{n-1} covers the middle edges of all paths P_1, \dots, P_{n-1} , and the robustness argument also extends and shows (8.24) to hold for $k = n - 1$. Since

the failure probability in (8.24) is exponentially small, and there are fewer than n^2 pairs $\{u, v\}$ in G' , w.h.p. there is a cheap path (of cost $\leq \delta$) for every pair. ■

For the remainder of this section we assume that the high-probability conclusion of Remark 8.3 holds.

Let H_k^s be the weight of the heaviest edge incident to s used by the first k paths, and let L_k^s be the weight of the lightest edge incident to s not used by the first k paths. Define H_k^t and L_k^t likewise.

We claim that for all k from 1 to $n - 1$, with $\delta = 20n^{-1/6}$ as in Remark 8.3,

$$H_k^s - L_k^s \leq \delta. \quad (8.26)$$

We argue by contradiction. Given k , let P_i , $i \leq k$, be the path using the edge of weight H_k^s . By Remark 8.3, we can construct an s - t path Q whose s -incident edge is the one of weight L_k^s , whose t -incident edge is the same as that of P_i , and whose middle edges cost at most δ and are not used in P_1, \dots, P_{n-1} . This path Q is cheaper than P_i : its s -incident edge is cheaper by $H_k^s - L_k^s > \delta$, its t -incident edge has the same cost, and its middle edges (costing at most δ) cost at most δ more than those of P_i . Also, Q is edge-disjoint from the first $i - 1$ paths: its s -incident edge L_k^s is not used even by the first k paths, the middle edges are disjoint from those of all $n - 1$ paths, and its t -incident edge is that used by P_i (so not used by a previous path). Thus, Q should have been chosen in preference to P_i , a contradiction, establishing (8.26).

Trivially, $H_k^s \geq W_{(k)}^s$. Thus, from (8.26),

$$L_k^s \geq H_k^s - \delta \geq W_{(k)}^s - \delta. \quad (8.27)$$

For $k \leq n - 2$, the edge of P_{k+1} incident to s costs at least L_k^s and the edge incident to t at least L_k^t . If P_{k+1} is not the single-edge path $\{s, t\}$ these two edges are distinct, so that $X_{k+1} \geq L_k^s + L_k^t$. If P_{k+1} is the single-edge path $\{s, t\}$ then P_k is not, and $X_{k+1} \geq X_k \geq L_{k-1}^s + L_{k-1}^t$. Either way, by (8.27),

$$\begin{aligned} X_{k+1} &\geq L_{k-1}^s + L_{k-1}^t \\ &\geq W_{(k-1)}^s + W_{(k-1)}^t - 2\delta. \end{aligned} \quad (8.28)$$

Recall that we are concerned here with $k \geq n^{9/10}$. By Lemma 4.2, for all such k , and for any $\gamma > 0$, w.h.p. $W_{(k)} \geq (1 - \gamma)\mathbb{E}W_{(k)}$. Since the exponential random variable is stochastically greater than the uniform, $\mathbb{E}W_{(k)} > k/n = \Omega(n^{-1/10})$, while $\delta = 20n^{-1/6} = o(\mathbb{E}W_{(k)})$. From (4.3) it is clear that $\mathbb{E}W_{(k-1)} \sim \mathbb{E}W_{(k+1)}$ (for any $k = \omega(1)$), and we subsume the asymptotic error into the constant γ . Thus, from (8.28), for any $\gamma > 0$, w.h.p., for all $k \geq n^{9/10}$,

$$X_k \geq (1 - \gamma)2\mathbb{E}W_{(k)},$$

completing the proof of the lower bound in Theorem 1.2.

9 | EXPECTATION

In this section we prove Theorem 1.5. We treat the uniform and exponential models at the same time. Let \mathcal{P}_k be the event that P_k exists. Clearly $\mathbb{P}(\mathcal{P}_k) \geq \mathbb{P}(\mathcal{P}_{n-1})$. By Theorem 1.1 (for the uniformly random model) and Theorem 1.2 (for the exponential model), $\mathbb{P}(\mathcal{P}_{n-1}) = 1 - o(1)$. This establishes the first

part of the theorem. Then, let $\mu_k = 2\mathbb{E}W_{(k)} + \ln n/n$ (so for the uniform model, $\mu_k = w_0(k)$). It suffices to show that

$$\mathbb{E}[X_k \mid \mathcal{P}_k] = (1 + o(1))\mu_k \tag{9.1}$$

uniformly in k .

First, we show the lower bound implicit in (9.1). Fix $\varepsilon > 0$. Let \mathcal{L}_k be the event that (jointly) \mathcal{P}_k exists and $X_k \geq (1 - \varepsilon)\mu_k$. By Theorem 1.1 (for the uniform model) and Theorem 1.2 (for the exponential model), \mathcal{L}_k holds with probability $1 - o(1)$ uniformly in k . Thus,

$$E[X_k \mid \mathcal{P}_k] \geq \mathbb{P}(\mathcal{L}_k)\mathbb{E}[X_k \mid \mathcal{P}_k \wedge \mathcal{L}_k] \geq (1 - o(1))(1 - \varepsilon)\mu_k.$$

Since this holds for any ε , we have that

$$\mathbb{E}[X_k \mid \mathcal{P}_k] \geq (1 - o(1))\mu_k.$$

We now establish the corresponding upper bound.

9.1 | Small k

First, we consider the range $k \leq n^{4/10}$. We will need the following lemma in (9.8).

Lemma 9.1. *There exists an absolute constant $C > 0$ such that, for all $\varepsilon > C$, in both the exponential and uniform models, for all $k = o(\sqrt{n})$ the probability of the event*

$$X_k > (1 + \varepsilon)\mu_k \tag{9.2}$$

is $O(n^{-1.9})$.

Proof. By the reasoning given in the introduction of Section 8, it is sufficient to show the result in the uniform case, where $\mu_k = \frac{2k + \ln n}{n}$. We use the same argument as developed in Section 3, where we prove Theorem 1.1 up to $k = o(\sqrt{n})$. Our argument in Section 3 (see (3.1)) was that for any sufficiently small $\varepsilon > 0$,

$$\text{if } X_i \leq (1 + \varepsilon)\left(\frac{2i}{n} + \frac{\ln n}{n}\right) \text{ for all } i \leq k, \text{ then w.h.p. the same holds for } i = k + 1. \tag{9.3}$$

We proved this by constructing a structure $R = R^{(k)}$ in G , in which after deleting k paths, each of cost $\leq (1 + \varepsilon)(2k/n + \ln n/n)$ from G , w.h.p. there remains a path in R satisfying the same cost bound. By (3.30), the probability of failure was $O(n^{-1.9}) + \exp(-\Theta(s(k)))$. This does not suffice since for k small the second term may exceed $O(n^{-1.9})$ (recall $s = 2k + \ln n$).

To prove the lemma, we will show that for some sufficiently large constant ε , the failure probability in (9.3) is $O(n^{-1.9})$. As noted in Remark 3.1, a few parts of the argument developed in Section 3 rely on ε being sufficiently small, and here we will detail the changes needed. Principally, we will make one modification (a simplification) to Section 3's construction of R . We will also track the dependence of key Landau-notation expressions on ε .

Recall from (3.4) and (3.5) that $s = 2k + \ln n$ and $w_0 = s/n$.

Paralleling the structure of Section 3, we start by reviewing the adversary's edge-count budget. This was given by (3.9) which, through its dependence on (3.8), held only for sufficiently small ε . For sufficiently large ε , modulo the one-time failure probability $O(n^{-1.9})$ from Lemma 3.3, each of the first k paths has length $\leq (1 + \varepsilon)w_0 \cdot 19n < 20s\varepsilon$, and the total length of the first k paths is at most

$$20ks\varepsilon < 10s^2\varepsilon, \quad (9.4)$$

so we now take this to be the adversary's budget.

We build level-0 edges of R exactly as in Section 3.3, and using the same parameter r_0 . That is, we add the cheapest $k + r_0$ edges incident on s , with $r_0 = \left\lceil \frac{1}{10}\varepsilon s \right\rceil$ as in (3.11); the opposite endpoints of these edges are the level-1 vertices. Recall that we declared this step a failure if the number X of edges with weights in the interval $[0, \frac{k}{n} + \frac{1}{9}\varepsilon w_0]$ is smaller than $k + r_0$. Note that $X \sim \text{Bi}(n', \frac{k}{n} + \frac{1}{9}\varepsilon w_0)$, thus $\mathbb{E}X = (1 - o(1))(k + \frac{1}{9}\varepsilon s)$, and failure means that $X < k + r_0$, that is, that

$$\frac{X}{\mathbb{E}X} = (1 + o(1)) \frac{k + \frac{1}{10}\varepsilon s}{k + \frac{1}{9}\varepsilon s} \leq \frac{10}{11}$$

for ε sufficiently large. Then, analogously to (3.12), the failure probability by Lemma 3.2 is at most

$$\mathbb{P}(X < \frac{10}{11}\mathbb{E}X) \leq \exp(-\Omega(\mathbb{E}X)) \leq \exp(-\Omega(\varepsilon s)). \quad (9.5)$$

We skip constructing level-1 edges as in Section 3.4, instead setting the level-2 vertices identical to level-1 vertices. (There are no edges between these levels; we have "level 2" only to keep the level numbering the same as before.)

We build level-2 edges exactly as before, with the same parameter r_2 , linking to each level-2 vertex its cheapest $r_2 = \frac{1}{10}\varepsilon s$ neighbors (which become the level-3 vertices). The calculations in (3.5) hold for any $\varepsilon > 0$, and from (3.18) the probability of any failure on this level is

$$\leq \exp -\Theta(\varepsilon s). \quad (9.6)$$

The adversary's deletions of edges incident on s must leave r_0 vertices at level 1 (a.k.a. level 2), thus $r_0 r_2 = \varepsilon^2 s^2 / 100$ edges leading to level 3. By (9.4) the adversary is allowed to delete at most $10s^2\varepsilon$ edges, so for ε sufficiently large, at least $2s^2$ level-3 vertices remain; this is the same as before, and will continue to suffice.

From level 3 we construct SPTs just as in (3.6), whose calculations hold for any $\varepsilon > 0$. To recapitulate, these trees are built to a size (3.19) independent of ε , the calculations made are valid for all ε , and the result (here as in Section 3) is that each tree fails with some probability $o(1)$, but the level as a whole fails only if at least $0.01s^2$ trees fail, which occurs with probability only $\exp(-\Omega(s^2))$ (see (3.27)).

This concludes the modified construction of R . The remainder of the argument is unchanged from Section 3. In the absence of failures, the maximum weight of any s - t path in R remains at most $(1 + \varepsilon)w_0$ per (3.28) (indeed, a little less as we have skipped the level-1 edges). The number of successful level-3 trees is $\Omega(s^2)$ as before, and the calculations leading to the probability that an adversary can destroy all cheap paths in R are unaffected: this probability remains $\exp(-\Omega(s^2 \ln n))$ as in (3.29), which is dominated by other failure probabilities.

Tallying up, as in Section 3.11, we have a one-time failure probability of $O(n^{-1.9})$ from Lemma 3.3. Out of levels 0, 2 and 3 we have failure probabilities given respectively by (9.5), (9.6), and (3.27), namely $\exp(-\Omega(\varepsilon s))$, $\exp(-\Omega(\varepsilon s))$ and $\exp(-\Omega(s^2))$. Since $s > \ln n$, for some ε sufficiently large, the net failure probability is $O(n^{-1.9})$, as claimed. ■

Let C be the constant in Lemma 9.1. Separately, fix any sufficiently small $\varepsilon > 0$. Let

$$\begin{aligned} U_1 &= [0, (1 + \varepsilon)\mu_k), \\ U_2 &= [(1 + \varepsilon)\mu_k, C\mu_k), \\ U_3 &= [C\mu_k, \infty). \end{aligned}$$

Let \mathcal{A}_i be the event that $X_k \in U_i$. By Theorem 1.1, $\mathbb{P}(\mathcal{A}_1) = 1 - o(1)$ and $\mathbb{P}(\mathcal{A}_2) = o(1)$, and by Lemma 9.1, $\mathbb{P}(\mathcal{A}_3) = O(n^{-1.9})$.

Since here we are considering $k \leq n^{4/10} \leq n/2$, with reference to the proof of Remark 1.4, one possible choice for P_k is some path of length 2 (there must remain at least one such), and thus, deterministically,

$$X_k \leq W_s + W_t, \tag{9.7}$$

where W_v denotes most expensive edge out of v ($W_v = W_{(n-1)}^v$ in the notation of (5.1)).

In the uniform model, (9.7) means that, deterministically, $X_k \leq 2$. Then,

$$\begin{aligned} \mathbb{E}[X_k] &= \mathbb{P}(\mathcal{A}_1)\mathbb{E}[X_k \mid \mathcal{A}_1] + \mathbb{P}(\mathcal{A}_2)\mathbb{E}[X_k \mid \mathcal{A}_2] + \mathbb{P}(\mathcal{A}_3)\mathbb{E}[X_k \mid \mathcal{A}_3] \\ &\leq (1 - o(1)) \cdot (1 + \varepsilon)\mu_k + o(1) \cdot (1 + C)\mu_k + O(n^{-1.9}) \cdot 2 \\ &\leq (1 + \varepsilon + o(1))\mu_k, \end{aligned} \tag{9.8}$$

since $\mu_k > \ln n/n$. As this holds for arbitrarily small $\varepsilon > 0$,

$$\mathbb{E}[X_k] \leq (1 + o(1))\mu_k. \tag{9.9}$$

For the exponential model the same argument applies, once we control $\mathbb{E}[X_k \mid \mathcal{A}_3]$. We make use of the following inequality. Let Z be a random variable with CDF F , and \mathcal{A} be an event with $\mathbb{P}(\mathcal{A}) = \alpha$. Then,

$$\mathbb{E}[Z \mid \mathcal{A}] \leq \mathbb{E}[Z \mid Z > F^{-1}(1 - \alpha)]. \tag{9.10}$$

In the case that Z is an exponential random variable with rate λ , $F(z) = 1 - \exp(-\lambda z)$, so $F^{-1}(1 - \alpha) = -\ln(\alpha)/\lambda$. By the memoryless property of the exponential, the RHS of (9.10) is $\mathbb{E}[Z] + F^{-1}(1 - \alpha)$, giving

$$\mathbb{E}[Z \mid \mathcal{A}] \leq \frac{1 - \ln(\alpha)}{\lambda}. \tag{9.11}$$

Recall from (4.2) that $W_v = \sum_{i=1}^{n-1} Z_i$ where $Z_i \sim \text{Exp}(i)$. Condition on the event \mathcal{A}_3 , taking $\alpha := \mathbb{P}(\mathcal{A}_3) = O(n^{-1.9})$. By (9.11),

$$\mathbb{E}[W_k \mid \mathcal{A}_3] = \sum_{i=1}^{n-1} \mathbb{E}[Z_i \mid \mathcal{A}_3] \leq \sum_{i=1}^{n-1} \frac{1 - \ln(\alpha)}{i} \sim (1 - \ln(\alpha)) \ln n = O(\ln^2 n). \tag{9.12}$$

By (9.7), (9.12) and linearity of expectation,

$$\mathbb{P}(\mathcal{A}_3)\mathbb{E}[X_k \mid \mathcal{A}_3] \leq \alpha\mathbb{E}[W_s + W_t \mid \mathcal{A}_3] = 2\alpha O(\ln^2 n) = O(n^{-1.9} \ln^2 n), \quad (9.13)$$

which is $o(\mu_k)$ since $\mu_k > \ln n/n$. Thus (9.8) holds also for the exponential model (the change to the middle line of the calculation affects nothing), whereupon so does (9.9).

9.2 | Large k

For $k \geq n^{4/10}$, we gather the failure events in Section 5. First, we have $X_{n^{4/10}} \leq 3n^{4/10}/n$ with failure probability $O(n^{-1.9})$, from (5.6) and (5.5). Then, we have to check two types of failures: failure of (5.3) to be an upper bound on (5.2) (because the edge order statistics are not as expected), and violation of (5.4) (because R fails to be robust against the adversary).

Failure of (5.3) as an upper bound is, in the uniform model, checked through violation of (6.11), the paragraph after (6.11) showing failure to occur w.p. at most $\exp(-\Omega(n^{0.01}))$. Likewise, in the exponential model it is checked in and following (8.15), with a failure probability of $O(\exp(-\Omega(n^{3/50})))$.

The failure probability of (5.4) in the uniform model is calculated for three cases: near (6.16) as $n \exp(-\Omega(n))$, near (6.18) as $n \exp(-\Omega(n^{11/25}))$, and near (6.20) as $14n^{5/2} \exp(-\Omega(n^{2/3}))$. The failure probability in the exponential model is also calculated for three cases: near (8.19) as $n \exp(-\Omega(n))$, near (8.21) as $n \exp(-\Omega(n^{1/25}))$, and near (8.25) $n^{5/2} \exp(-\Omega(n^{2/3}))$.

Thus, the failure probabilities for (5.3) and (5.4) are all $O(\exp(-n^{0.01}))$, so the probability of any failure affecting any $k > n^{4/10}$ is $O(n^{-1.9})$.

Let

$$\begin{aligned} U_1 &= [0, (1 + \varepsilon)\mu_k) \\ U_2 &= [(1 + \varepsilon)\mu_k, \infty), \end{aligned}$$

and let \mathcal{A}_i be the event that P_k exists and $X_k \in U_i$. Thus $\mathbb{P}(\mathcal{A}_1) = 1 - o(1)$ and $\mathbb{P}(\mathcal{A}_2) = O(n^{-1.9})$.

Conditioning on the event \mathcal{P}_k that P_k exists, this path clearly has cost

$$X_k \leq Z := \sum_{v \in V(G)} W_v$$

(analogous to (9.7)). In the uniform model, deterministically, $Z \leq n$. In the exponential model, the event \mathcal{A}_2 here has the same probability as event \mathcal{A}_3 in Section 9.1, so we may reuse (9.12), obtaining

$$\mathbb{E}[Z \mid \mathcal{A}_2] = \sum_{v \in V(G)} \mathbb{E}[W_v \mid \mathcal{A}_2] = n O(\ln^2 n) = o(n^{1.1}).$$

Thus, in both the uniform and exponential cases,

$$\begin{aligned} \mathbb{E}[X_k \mid \mathcal{P}_k] &= \mathbb{P}(\mathcal{A}_1)\mathbb{E}[X_k \mid \mathcal{A}_1] + \mathbb{P}(\mathcal{A}_2)\mathbb{E}[X_k \mid \mathcal{A}_2] \\ &\leq (1 - o(1)) \cdot (1 + \varepsilon)\mu_k + O(n^{-1.9}) \cdot o(n^{1.1}) \\ &= (1 - o(1))(1 + \varepsilon)\mu_k, \end{aligned} \quad (9.14)$$

since $\mu_k > 2k/n > n^{-6/10} = \omega(n^{-0.8})$. As this holds for arbitrarily small $\varepsilon > 0$, for all $k \geq n^{4/10}$,

$$\mathbb{E}[X_k \mid \mathcal{P}_k] \leq (1 + o(1))\mu_k, \quad (9.15)$$

completing the proof.

ACKNOWLEDGMENTS

We thank Alan Frieze and Wes Pegden for an initial discussion of the second-shortest path, and Alan for noticing that minimum-cost k -flow (Remark 1.3) was not an open problem but immediately implied by our other results. We also thank two anonymous referees for helpful suggestions.

REFERENCES

1. J. A. Adell and P. Jodrá, *The median of the Poisson distribution*, *Metrika*. **61** (2005), 337–346.
2. D. J. Aldous, *The $\zeta(2)$ limit in the random assignment problem*, *Random Structures Algorithms*. **18** (2001), 381–418.
3. A. Beveridge, A. Frieze, and C. McDiarmid, *Random minimum length spanning trees in regular graphs*, *Combinatorica*. **18** (1998), 311–333.
4. S. Bhamidi and R. van der Hofstad, *Weak disorder asymptotics in the stochastic mean-field model of distance*, *Ann. Appl. Probab.* **22** (2012), 29–69.
5. S. Bhamidi, G. Hooghiemstra, and R. van der Hofstad, *First passage percolation on the Erdős-Rényi random graph*, *Combin. Probab. Comput.* **20** (2011), 683–707.
6. K. P. Choi, *On the medians of gamma distributions and an equation of Ramanujan*, *Proc. Amer. Math. Soc.* **121** (1994), 245–251.
7. M. Eckho, J. Goodman, R. van der Hofstad, and F. R. Nardi, *Short paths for first passage percolation on the complete graph*, *J. Stat. Phys.* **151** (2013), 1056–1088.
8. A. Frieze and T. Johansson, *On edge-disjoint spanning trees in a randomly weighted complete graph*, *Combin. Probab. Comput.* **27** (2018), 228–244.
9. A. Frieze, W. Pegden, and G. B. Sorkin, *The distribution of minimum-weight cliques and other subgraphs in graphs with random edge weights*, *SIAM J. Discrete Math.* **32** (2018), 2115–2133.
10. A. Frieze, *On the value of a minimum spanning tree problem*, *Discrete Appl. Math.* **10** (1985), 47–56.
11. K. Hamza, *The smallest uniform upper bound on the distance between the mean and the median of the binomial and Poisson distributions*, *Statist. Probab. Lett.* **23** (1995), 21–25.
12. S. Janson, *On concentration of probability*. *Contemporary Combinatorics*, Bolyai Soc. Math. Stud. János Bolyai Math. Soc., Budapest, Vol. 10, 2002, pp. 289–301.
13. S. Janson, *Tail bounds for sums of geometric and exponential variables*, *Statist. Probab. Lett.* **135** (2017).
14. S. Janson, *One, two and three times $\log n/n$ for paths in a complete graph with random weights*, *Combin. Probab. Comput.* **8** (1999), 347–361.
15. S. Janson and G. B. Sorkin, *Successive minimum spanning trees*, 2019, arXiv: 1906.01533.
16. R. Kaas and J. M. Buhrman, *Mean, median and mode in binomial distributions*, *Stat. Neerl.* **34** (1980), 13–18.
17. É. Lucas, *Récréations Mathématiques*. A facsimile of the 1896 Second Printing, Gauthier-Villars, Paris, Vol. 2, available at <https://archive.org/details/recreationmatedou02lucarich>, 1892.
18. A. Martinsson, *Unoriented first-passage percolation on the n -cube*, *Ann. Appl. Probab.* **26** (2016), 2597–2625.
19. C. Su, Q. Feng, and Z. Hu, *Uniform recursive trees: Branching structure and simple random downward walk*, *J. Math. Anal. Appl.* **315** (2006), 225–243.
20. J. Wästlund, *Random matching problems on the complete graph*, *Electron. Commun. Probab.* **13** (2008), 258–265.
21. J. Wästlund, *An easy proof of the $\zeta(2)$ limit in the random assignment problem*, *Electron. Commun. Probab.* **14** (2009), 261–269.
22. J. Wästlund, *The mean field traveling salesman and related problems*, *Acta Math.* **204** (2010), 91–150.

How to cite this article: Gerke S, Mezei BF, Sorkin GB. Successive shortest paths in complete graphs with random edge weights. *Random Struct Alg.* 2020;1–43. <https://doi.org/10.1002/rsa.20962>