

Near-minimal spanning trees: A scaling exponent in probability models

David J. Aldous^{a,1}, Charles Bordenave^{b,2} and Marc Lelarge^{c,3}

^aUniversity of California, Department of Statistics, 367 Evans Hall # 3860, Berkeley CA 94720-3860, USA. E-mail: aldous@stat.berkeley.edu

^bÉcole Normale Supérieure, Département d'Informatique, 45 rue d'Ulm, 75230 Paris Cedex 5, France. E-mail: charles.bordenave@ens.fr

^cINRIA-ENS, 45 rue d'Ulm, 75230 Paris Cedex 5, France. E-mail: marc.lelarge@ens.fr

Received 22 October 2006; revised 3 April 2007; accepted 1 June 2007

Abstract. We study the relation between the minimal spanning tree (MST) on many random points and the “near-minimal” tree which is optimal subject to the constraint that a proportion δ of its edges must be different from those of the MST. Heuristics suggest that, regardless of details of the probability model, the ratio of lengths should scale as $1 + \Theta(\delta^2)$. We prove this *scaling result* in the model of the lattice with random edge-lengths and in the Euclidean model.

Résumé. Nous étudions la relation entre l'arbre couvrant minimal (ACM) sur des points aléatoires et l'arbre “quasi” optimal sous la contrainte qu'une proportion δ de ses arêtes soit différente de celles de l'ACM. Un raisonnement heuristique suggère que quelque soit le modèle probabiliste sous-jacent, le ratio des longueurs des deux arbres doit varier en $1 + \Theta(\delta^2)$. Nous montrons ce *résultat d'échelle* pour le modèle de la grille avec des longueurs d'arêtes aléatoires et pour le modèle Euclidien.

MSC: 05C80; 60K35; 68W40

Keywords: Combinatorial optimization; Continuum percolation; Disordered lattice; Local weak convergence; Minimal spanning tree; Poisson point process; Probabilistic analysis of algorithms; Random geometric graph

1. Introduction

This paper gives details of one aspect of the following broad project [1]. Freshman calculus tells us how to find a minimum x_* of a smooth function $f(x)$: set the derivative $f'(x_*) = 0$ and check $f''(x_*) > 0$. The related series expansion tells us, for points x near to x_* , how the distance $\delta = |x - x_*|$ relates to the difference $\varepsilon = f(x) - f(x_*)$ in f -values: ε scales as δ^2 . This *scaling exponent 2* persists for functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$: if x_* is a local minimum and $\varepsilon(\delta) := \min\{f(x) - f(x_*): |x - x_*| = \delta\}$, then $\varepsilon(\delta)$ scales as δ^2 for a generic smooth function f .

Combinatorial optimization, exemplified by the *traveling salesman problem* (TSP), is traditionally viewed as a quite distinct subject, with theoretical analysis focussed on the number of steps that algorithms require to find the optimal solution. To make a connection with calculus, compare an arbitrary tour \mathbf{x} through n points with the optimal (minimum-length) tour \mathbf{x}_* by considering the two quantities

$$\delta_n(\mathbf{x}) = \frac{\{\text{number of edges in } \mathbf{x} \text{ but not in } \mathbf{x}_*\}}{n},$$

¹Research supported by N.S.F. Grant DMS0203062.

²Research supported in part by the EuroNGI network.

³Research supported by Science Foundation Ireland Grant SFI04/RP1/I512.

$$\varepsilon_n(\mathbf{x}) = \frac{\{\text{length difference between } \mathbf{x} \text{ and } \mathbf{x}_*\}}{s(n)},$$

where $s(n)$ is the length of the minimum length tour. Now define $\varepsilon_n(\delta)$ to be the minimum value of $\varepsilon_n(\mathbf{x})$ over all tours \mathbf{x} for which $\delta_n(\mathbf{x}) \geq \delta$. Although the function $\varepsilon_n(\delta)$ will depend on n and the problem instance, we anticipate that for typical instances drawn from a suitable probability model it will converge in the $n \rightarrow \infty$ limit to some deterministic function $\varepsilon(\delta)$. The *universality* paradigm from statistical physics [8] suggests there might be a scaling exponent α defined by

$$\varepsilon(\delta) \sim \delta^\alpha \quad \text{as } \delta \rightarrow 0$$

and that the exponent should be robust under model details.

There is fairly strong evidence [1] that for TSP the scaling exponent is 3. This is based on analytic methods in a *mean-field* model of interpoint distances (distances between pairs of points are random, independent for different pairs, thus ignoring geometric constraints) and on Monte Carlo simulations for random points in 2, 3 and 4 dimensional space. The analytic results build upon a recent probabilistic reinterpretation [2] of the work of Krauth and Mézard [9] establishing the average length of mean-field TSP tours. But neither part of these TSP assertions is rigorous, and indeed rigorous proofs in d dimensions seem far out of reach of current methodology. In contrast, for the *minimum spanning tree* (MST) problem, a standard algorithmically easy problem, a simple heuristic argument (Section 1.2) strongly suggests that the scaling exponent is 2 for any reasonable probability model. The goal of this paper is to work through the details of a rigorous proof.

Why study such scaling exponents? For a combinatorial optimization problem, a larger exponent means that there are more near-optimal solutions, suggesting that the algorithmic problem of finding the optimal solution is intrinsically harder. So scaling exponents may serve to separate combinatorial optimization problems of an appropriate type into a small set of classes of increasing difficulty. For instance, the *minimum matching* and *minimum Steiner tree* problems are expected to have scaling exponent 3, and thus be in the same class as TSP in a *quantitative* way, as distinct from their qualitative similarity as NP-complete problems under worst-case inputs. In contrast, algorithmically easy problems are expected to have scaling exponent 2, analogously to the ‘‘calculus’’ scaling exponent. One plausible explanation is that the near-optimal solutions in such problems differ from the optimal solution via only ‘‘local changes,’’ each local change affecting only a number of edges which remains $O(1)$ as $\delta \rightarrow 0$.

1.1. Background

Steele [11] and Yukich [13] give general background concerning combinatorial optimization over random points.

A *network* is a graph whose edges e have positive real *lengths* $\text{len}(e)$. Let G be a finite connected network. Recall the notion of a *spanning tree* (ST) T in G . Identifying T as a set of edges, write $\text{len}(T) = \sum_{e \in T} \text{len}(e)$. A *minimal spanning tree* (MST) is a ST of minimal length; such a tree always exists but may not be unique. The classical greedy algorithm (*Kruskal’s algorithm* [7]) for constructing a MST yields two fundamental properties which we record without proof in Lemma 1.

Let G_t be the subnetwork consisting of those edges e of G with $\text{len}(e) < t$. For arbitrary vertices v, w define

$$\text{perc}(v, w) = \inf\{t: v \text{ and } w \text{ in same component of } G_t\}. \quad (1)$$

For an edge $e = (v, w)$ of G write $\text{perc}(e) = \text{perc}(v, w) \leq \text{len}(e)$ and also define the *excess*

$$\text{exc}(e) = \text{len}(e) - \text{perc}(e) \geq 0.$$

Lemma 1. *Suppose all the edge-lengths in G are distinct.*

(a) *There is a unique MST, say T , and it is specified by the criterion*

$$e \in T \quad \text{if and only if } \text{exc}(e) = 0.$$

(b) *For any vertices v, w*

$$\text{perc}(v, w) = \max\{\text{len}(e): e \text{ on path from } v \text{ to } w \text{ in } T\}.$$

1.2. The heuristic argument

Given a probability model for n random points and their interpoint lengths, define a measure $\mu_n(\cdot)$ on $(0, \infty)$ in terms of the expectation

$$\mu_n(0, x) = \frac{1}{n} \mathbb{E} \left| \left\{ \text{edges } e: 0 < \text{len}(e) - \text{perc}(e) < x \right\} \right|.$$

For any reasonable model with suitable scaling of edge-lengths we expect an $n \rightarrow \infty$ limit measure $\mu(\cdot)$, with a density $f_\mu(x) = d\mu/dx$ having a non-zero limit $f_\mu(0^+)$ as $x \downarrow 0$.

Now modify the MST by adding an edge e with $\text{len}(e) - \text{perc}(e) = b$, for some small b , to create a cycle; then delete the longest edge $e' \neq e$ of that cycle, which necessarily has $\text{len}(e') = \text{perc}(e)$. This gives a spanning tree containing exactly one edge not in the MST and having length greater by b . Repeat this procedure with every edge e for which $0 < \text{len}(e) - \text{perc}(e) < \beta$, for some small β . For large n , the number of such edges should be $n\mu_n(0, \beta) \approx n f_\mu(0^+) \beta$ to first order in β , and assuming there is negligible overlap between cycles, each of the new edges will increase the tree length by $\sim \beta/2$ on average. So we expect (Lemma 6)

$$\delta(\beta) \sim f_\mu(0^+) \beta, \quad \varepsilon(\beta) \sim \frac{f_\mu(0^+) \beta^2}{2}.$$

This construction should yield essentially the minimum value of ε for given δ , so we expect

$$\varepsilon(\delta) \sim \frac{\delta^2}{2 f_\mu(0^+)} \tag{2}$$

and in particular we expect the scaling exponent to be 2.

1.3. Results

Our goal is to formalize the argument above in the context of the following two probability models for n random points. Fix dimension $d \geq 2$ (the case $d = 1$ is of course rather special).

Model 1 (The disordered lattice). Start with the discrete d -dimensional cube $\mathbb{C}_m^d = [1, 2, \dots, m]^d$, so there are $n = m^d$ vertices and there are $2d$ edges at each non-boundary vertex. Then take the edge-lengths to be i.i.d. random variables ξ_e , whose common distribution ξ has finite mean and some bounded continuous density function $f_\xi(\cdot)$.

Model 2 (Random Euclidean). Take the continuum d -dimensional cube $[0, n^{1/d}]^d$ of volume n . Put down n independent uniformly distributed random points in this cube. Take the complete graph on these n vertices, with Euclidean distance as edge-lengths.

The results of this paper will remain valid in a slightly more general framework than Model 2 in which points are put down independently at random in the cube $[0, n^{1/d}]^d$ with common density $f(n^{1/d}x)$ on \mathbb{R}^d , with f having support on $[0, 1]^d$ and being bounded away from zero. To avoid technicalities, we restrict ourselves to the case f constant.

Each model is set up so that nearest-neighbor distances are order 1 and the MST T_n has mean length of order n . To formalize the ideas in the introduction we define the random variable

$$\varepsilon_n(\delta) := \min \left\{ \frac{\text{len}(T'_n) - \text{len}(T_n)}{n} : |T'_n \setminus T_n| \geq \delta n \right\}, \tag{3}$$

where the minimum is over spanning trees T'_n and where $T'_n \setminus T_n$ is the set of edges in T'_n but not in T_n .

Theorem 2. *In either model, we have*

$$(a) \quad \limsup_{\delta \downarrow 0} \delta^{-2} \limsup_n \mathbb{E} \varepsilon_n(\delta) < \infty,$$

and,

$$(b) \quad \liminf_{\delta \downarrow 0} \delta^{-2} \liminf_n \mathbb{E} \varepsilon_n(\delta) > 0.$$

Structure of the paper

In Section 2, we do calculations in the finite models: we prove Theorem 2 for Model 1 and part (a) of the theorem for Model 2. In Section 3, we introduce the limit infinite random network (limit in the sense of *local weak convergence* [4]) and its associated minimal spanning forest. We show how results from continuum percolation theory allow us to show part (b) of Theorem 2 for Model 2.

2. Proofs for the finite network

2.1. The upper bound: Model 1 with $d = 2$

We first consider Model 1 with $d = 2$ and then consider the other cases.

The upper bound rests upon a simple construction of near-minimal spanning trees, illustrated in Fig. 1.

The figure illustrates a particular kind of configuration. There is a 4-cycle of edges $abcd$ where, for some x ,

$$\text{len}(a) = x, \quad \text{len}(b) \in (x, x + \delta), \quad \text{len}(c) < x, \quad \text{len}(d) < x$$

and where the eight other edges touching the cycle have lengths $> x + \delta$. With such a configuration (within a larger configuration on \mathbb{C}_m^2), edges adc are in the MST, and edge b is not. We can modify the minimal spanning tree by removing edge a and adding edge b ; this creates a new spanning tree whose extra length equals $\text{len}(b) - x$.

Thus given a realization of the edge-lengths on the $m \times m$ discrete square, partition the square into adjacent 3×3 regions; on each region where the configuration is as in Fig. 1, make the modification above. This changes the MST T_n into a certain near-minimal spanning tree T'_n . On each 3×3 square, the probability of seeing the Fig. 1 configuration equals

$$q(\delta) := \int_0^\infty f(x) (F(x + \delta) - F(x)) F^2(x) (1 - F(x + \delta))^8 dx.$$

Here f and F are the density and distribution functions of edge-lengths. And the (unconditioned) increase in edge-length of spanning tree caused by the possible modification equals

$$r(\delta) := \int_0^\infty f(x) \left(\int_x^{x+\delta} (y - x) f(y) dy \right) F^2(x) (1 - F(x + \delta))^8 dx.$$




Fig. 1. A special configuration on the 3×3 grid.

Letting $n \rightarrow \infty$ with fixed δ , and using the weak law of large numbers,

$$n^{-1}|T'_n \setminus T_n| \xrightarrow{P} \frac{1}{9}q(\delta), \tag{4}$$

$$n^{-1}(\text{len}(T'_n) - \text{len}(T_n)) \xrightarrow{P} \frac{1}{9}r(\delta). \tag{5}$$

Because we defined $\varepsilon_n(\cdot)$ in terms of spanning trees which differ from the MST by a *non-random* proportion of edges, we need a detour to handle expectations over events of asymptotically zero probability. We defer the proof.

Lemma 3. (a) For any sequence T_n^* of spanning trees, the sequence $n^{-1} \text{len}(T_n^*)$ is uniformly integrable.

(b) There exist spanning trees T_n'' such that

$$|T_n'' \setminus T_n| \geq a_n,$$

where $a_n/n \rightarrow 1/2$.

Now consider the spanning tree T_n^* defined to be T'_n if $n^{-1}|T'_n \setminus T_n| \geq \frac{1}{10}q(\delta)$ and to be T_n'' if not. It follows from (4,5) and Lemma 3 that

$$n^{-1}|T_n^* \setminus T_n| \geq \frac{1}{10}q(\delta), \quad \text{for large } n,$$

$$\limsup_n n^{-1} \mathbb{E}(\text{len}(T_n^*) - \text{len}(T_n)) \leq \frac{1}{9}r(\delta).$$

Then from the definitions of $q(\delta), r(\delta)$ and the assumption that $f(\cdot)$ is bounded it is easy to check

$$q(\delta) \sim c\delta, \quad r(\delta) \sim \frac{1}{2}\delta q(\delta) \quad \text{as } \delta \downarrow 0 \tag{6}$$

for a certain $0 < c < \infty$. This establishes the upper bound (a) in Theorem 2.

Proof of Lemma 3. Part (a) is automatic because, writing \sum_e for the sum over all edges of \mathbb{C}_m^2 , the sequence $n^{-1} \sum_e \xi_e$ is uniformly integrable. For (b), note that the cube \mathbb{C}_m^2 with $2m(m-1)$ edges can be regarded as a subgraph of the discrete torus \mathbb{Z}_m^2 with $2m^2$ edges. Take a uniform random spanning tree \tilde{T}_n on \mathbb{Z}_m^2 , delete edges not in \mathbb{C}_m^2 and add back boundary edges to make some (non-uniform) random spanning tree \mathcal{T}_n on \mathbb{C}_m^2 . By symmetry of the torus we have $\mathbb{P}(e \in \tilde{T}_n) = \frac{m^2-1}{2m^2}$ for each edge e of the torus, and it follows that $\mathbb{P}(e \in \mathcal{T}_n) = \frac{m^2-1}{2m^2}$ for each non-boundary edge of the cube. Since there are $4(m-1)$ boundary edges and $2(m-1)(m-2)$ non-boundary edges, for any spanning tree \mathbf{t} we have

$$\mathbb{E}|\mathcal{T}_n \cap \mathbf{t}| \leq 4(m-1) + (n-1)\frac{m^2-1}{2m^2} = 4(n^{1/2}-1) + \frac{(n-1)^2}{2n}.$$

So

$$\begin{aligned} \mathbb{E}|\mathcal{T}_n \setminus \mathbf{t}| &= (n-1) - \mathbb{E}|\mathcal{T}_n \cap \mathbf{t}| \\ &\geq a_n := (n-1) - 4(n^{1/2}-1) - \frac{(n-1)^2}{2n}. \end{aligned}$$

So for any spanning tree \mathbf{t} there exists some spanning tree \mathbf{t}^* such that $|\mathbf{t}^* \setminus \mathbf{t}| \geq a_n$. Applying this fact to the MST gives (b). □




Fig. 2. A special configuration on the 3×3 square.

2.2. Upper bound: Other cases

The argument for Model 1 in the case $d \geq 3$ involves only very minor modifications of the proof above, so we turn to Model 2 with $d = 2$ (the case $d \geq 3$ is similar). Here it is natural to consider a different notion of special configuration, see Figure 2.

Here there is a 3×3 square containing a concentric 1×1 square. There are three points within the larger square, all being inside the smaller square. In the triangle abc formed by the three points, writing x for the length of the second longest edge length, the length of the longest edge is in the interval $(x, x + \delta)$, and $x + \delta < 1$. For such a configuration (within a configuration on a $m \times m$ square containing the 3×3 square), edges ac are in the MST, and edge b is not. We can modify the minimal spanning tree by removing edge a and adding edge b ; this creates a new spanning tree whose extra length equals $\text{len}(b) - x$.

We now repeat the argument from the previous section, and the overall logic is the same. One gets different formulas for $q(\delta), r(\delta)$ but they have the same relationship (6). The weak law (4), (5) is easily established. The only non-trivial difference is that we need to replace the technical Lemma 3 by the following technical lemma.

Lemma 4. (a) *There exists c_1 such that for any n and any configuration on n points in the square of area n , the MST \hat{T}_n has $\text{len}(\hat{T}_n) \leq c_1 n$.*

(b) *For sufficiently large n , there exist spanning trees T_n'' such that $\text{len}(T_n'') \leq 12c_1 n$ and*

$$n^{-1} |T_n'' \setminus T_n| \geq \frac{1}{2}.$$

Proof. Part (a) follows from the analogous result for TSP – see [11] inequality (2.14). For (b), let ξ_1, \dots, ξ_n be the positions of the n random points and recall that T_n is their MST. Classify these points ξ_i as “odd” or “even” according to whether the number of edges in the path inside T_n from ξ_i to ξ_1 is odd or even. Let $(\hat{\xi}_i)$ be a configuration obtained from (ξ_i) by moving each “odd” point a distance $11c_1$ in some arbitrary direction. Let \hat{T}_n be the MST on $(\hat{\xi}_i)$. Let T_n'' be the spanning tree on (ξ_i) defined by

$$(\xi_i, \xi_j) \in T_n'' \quad \text{iff} \quad (\hat{\xi}_i, \hat{\xi}_j) \in \hat{T}_n.$$

Suppose (ξ_i, ξ_j) is an edge of both T_n and T_n'' . Since one end-vertex is odd and the other is even, it is easy to see: either (i) $\text{len}(\xi_i, \xi_j) \geq 5c_1$; or (ii) $\text{len}(\hat{\xi}_i, \hat{\xi}_j) \geq 5c_1$. But by part (a) there are at most $n/5$ edges satisfying (i), and similarly for (ii). So $|T_n \cap T_n''| \leq 2n/5$. Noting that

$$\text{len}(T_n'') \leq 11c_1(n - 1) + \text{len}(\hat{T}_n) \leq 12c_1 n$$

using (a), we have established (b). □

2.3. The lower bound: A discrete lemma

The lower bound argument rests upon the following simple lemma.

Lemma 5. *Consider a finite connected network with distinct edge-lengths. If T is the MST and T' is any ST then*

$$\text{len}(T') - \text{len}(T) \geq \sum_{e' \in T' \setminus T} \text{exc}(e').$$

Proof. Suppose $|T' \setminus T| = k \geq 1$. It is enough to show that there exist $e' \in T' \setminus T$ and $e \in T \setminus T'$ such that

- (i) $T^* = T' \setminus \{e'\} \cup \{e\}$ is a ST;
- (ii) $|T^* \setminus T| = k - 1$;
- (iii) $\text{len}(e') - \text{len}(e) \geq \text{exc}(e')$

for then we can continue inductively.

To prove this we first choose an arbitrary $e' \in T' \setminus T$. Consider $T' \setminus \{e'\}$. This is a two-component forest; so the path in T linking the end-vertices of e' must contain some edge $e \in T \setminus T'$ which links these two components. So choose some such edge e . Properties (i) and (ii) are clear. Apply Lemma 1(b) to the end-vertices of e' to see that $\text{perc}(e') \geq \text{len}(e)$. So

$$\text{len}(e') - \text{len}(e) \geq \text{len}(e') - \text{perc}(e') = \text{exc}(e')$$

which is (iii).

We will also need the following integration lemma; part (a) will be used for Model 1 and part (b) for Model 2. \square

Lemma 6. (a) *Let ξ and W be independent real-valued random variables such that ξ has a density function bounded by a constant b . Then for any event $A \subseteq \{\xi > W\}$ we have*

$$\mathbb{E}(\xi - W)\mathbb{1}_A \geq \frac{\mathbb{P}^2(A)}{2b}.$$

(b) *Let $(V_i, i \geq 1)$ be real-valued r.v.'s such that $\mu(0, x) := \mathbb{E} \sum_i \mathbb{1}_{(0 < V_i < x)}$ satisfies $\limsup_{x \downarrow 0} \frac{\mu(0, x)}{x} < \infty$. Then there exists a function $g(s) \sim \beta s^2$ as $s \downarrow 0$, for some $\beta > 0$, such that for any sequence of events $A_i \subseteq \{V_i > 0\}$,*

$$\mathbb{E} \sum_i V_i \mathbb{1}_{A_i} \geq g\left(\sum_i \mathbb{P}(A_i)\right).$$

Proof. (a) It is sufficient to prove

$$\mathbb{E}[(\xi - W)\mathbb{1}_A | W] \geq \frac{\mathbb{P}^2(A|W)}{2b} \quad \text{a.s.}, \tag{7}$$

then by Jensen's inequality, we get

$$\mathbb{E}[\mathbb{E}[(\xi - W)\mathbb{1}_A | W]] \geq \frac{\mathbb{E}[\mathbb{P}^2(A|W)]}{2b} \geq \frac{\mathbb{P}^2(A)}{2b}.$$

Since ξ and W are independent of each other, Eq. (7) reduces to

$$\mathbb{E}[\xi \mathbb{1}_A] \geq \frac{\mathbb{P}^2(A)}{2b} \quad \text{for } A \subseteq \{\xi > 0\}.$$

We can couple ξ to a r.v. U such that

- (i) $U = 0$ on $\{\xi \leq 0\}$.
- (ii) U has constant density b on $(0, \mathbb{P}(\xi > 0)/b)$.
- (iii) $U \leq \xi$.

Now it suffices to prove

$$\mathbb{E}[U\mathbb{1}_A] \geq \frac{\mathbb{P}^2(A)}{2b} \quad \text{for } A \subseteq \{U > 0\}. \tag{8}$$

But it is clear that, for a given value of $\mathbb{P}(A)$, the choice of $A \subseteq \{U > 0\}$ that minimizes $\mathbb{E}[U\mathbb{1}_A]$ is of the form $A_c := \{0 < U < c\}$ for some $c < \mathbb{P}(\xi > 0)/b$. A brief calculation gives

$$\mathbb{E}[U\mathbb{1}_{A_c}] = \frac{\mathbb{P}^2(A_c)}{2b}$$

establishing (8).

(b) For small $s > 0$ define $g(s)$ by

$$g(s) = \int_0^{c(s)} x\mu(dx), \quad \text{where } \mu(0, c(s)) = s.$$

By hypothesis there exists $\gamma > 0$ such that $c(s) \geq \gamma s$ for small s . So

$$g(s) \geq \int_{c(s)/2}^{c(s)} x\mu(dx) \geq c\left(\frac{s}{2}\right) \times \frac{s}{2} = \frac{\gamma s^2}{4}.$$

Taking $A_i^s := \{0 < V_i \leq c(s)\}$ we have

$$\sum_i \mathbb{P}(A_i^s) = \mu(0, c(s)) = s; \quad \mathbb{E} \sum_i V_i \mathbb{1}_{A_i^s} = \int_0^{c(s)} x\mu(dx) = g(s).$$

This is clearly the choice of (A_i) which minimizes the left side subject to $\sum_i \mathbb{P}(A_i) = s$, and so for arbitrary (A_i) we have

$$\mathbb{E} \sum_i V_i \mathbb{1}_{A_i} \geq g\left(\sum_i \mathbb{P}(A_i)\right).$$

□

2.4. The lower bound in Model 1

We treat the case $d = 2$, but $d \geq 3$ involves only minor changes. Recall $\mathbb{C}_m^2 = G^{(n)}$ has $c_n := 2(n - n^{1/2})$ edges. Fix $\delta > 0$. Consider a pair (T'_n, T_n) attaining the minimum in the definition (3) of $\varepsilon_n(\delta)$. For a uniform random edge e_n of \mathbb{C}_m^2 ,

$$\mathbb{P}(e_n \in T'_n \setminus T_n) = \frac{\mathbb{E}|T'_n \setminus T_n|}{c_n} \geq \frac{\delta}{2} \tag{9}$$

and

$$\begin{aligned} \mathbb{E}\varepsilon_n(\delta) &= \frac{1}{n} \mathbb{E}(\text{len}(T'_n) - \text{len}(T_n)) \\ &\geq \frac{1}{n} \mathbb{E} \sum_{e \in T'_n \setminus T_n} \text{exc}(e) \quad \text{by Lemma 5} \\ &= \frac{c_n}{n} \mathbb{E} \text{exc}(e_n) \mathbb{1}_{(e_n \in T'_n \setminus T_n)}. \end{aligned} \tag{10}$$

For a fixed edge e of \mathbb{C}_m^2 we can write

$$\text{exc}(e) = (\xi^{(n)}(e) - W^{(n)}(e))^+$$

where $\xi^{(n)}(e)$ is the edge-length of $e = (v, v^*)$ and where

$$W^{(n)}(e) = \inf\{t: v \text{ and } v^* \text{ in the same component of } G_t^{(n)} \setminus \{e\}\}.$$

Note (and this is the key special feature that makes Model 1 easy to study) that $\xi^{(n)}(e)$ and $W^{(n)}(e)$ are independent. Since $\text{exc}(e) > 0$ on $\{e \in T'_n \setminus T_n\}$ we see that the quantity at (10) is of the form appearing in Lemma 6(a). So

$$\begin{aligned} \frac{n}{c_n} \mathbb{E} \varepsilon_n(\delta) &\geq \mathbb{E}(\xi^{(n)}(e_n) - W^{(n)}(e_n)) \mathbb{1}_{(e_n \in T'_n \setminus T_n)} \quad \text{by (10)} \\ &\geq \frac{\mathbb{P}^2(e_n \in T'_n \setminus T_n)}{2\bar{f}} \quad \text{by Lemma 6(a)} \\ &\geq \frac{\delta^2}{8\bar{f}} \quad \text{by (9),} \end{aligned}$$

where \bar{f} is the bound on the density of ξ . Because $c_n \sim 2n$ we have established part (b) of Theorem 2 in this case.

3. The minimum spanning forest and continuum percolation

It remains to prove the lower bound in Model 2. Rather than doing calculations with the finite model, we consider the limit Poisson process on the plane, and exploit the well known connection between the *minimum spanning forest* (MSF) and continuum percolation. We then relate the finite models to the infinite limits in Section 3.3, as an instance of *local weak convergence* [4] of random graphical structures.

3.1. Minimum spanning forests

Here is a general definition, in the context of a countable-vertex network G with distinct edge-lengths (see [5] for more detailed treatment). As in Section 1.1 let G_t be the subnetwork consisting of those edges e of G with $\text{len}(e) < t$. Define the MSF by:

an edge (v, w) is in the MSF if and only if, for $t = \text{len}(v, w)$, vertices v and w are in different components of G_t and at least one of these components is finite.

Consider a Poisson point process $\Phi = \sum_i \delta_{\eta_i}$ of rate 1 in \mathbb{R}^d . Add an extra point O at the origin. Consider $\Phi^O = \sum_i \delta_{\eta_i} + \delta_O$ as the vertices of a network G (the complete graph with Euclidean edge-lengths). With probability one, Φ^O has only finitely many points in any bounded subset of \mathbb{R}^d and all of the interpoint distances are distinct. As in Section 1.1, we define for arbitrary points η_i and η_j of Φ^O ,

$$\text{perc}(\eta_i, \eta_j) = \inf\{t: \eta_i \text{ and } \eta_j \text{ are in the same component of } G_t\}.$$

We now give some properties of the MSF denoted \mathcal{F}_∞ on this network and show how Lemma 1 extends to this setting.

Lemma 7. (a) *We have $e \in \mathcal{F}_\infty$ if and only if $\text{len}(e) = \text{perc}(e)$.*

(b) *For any vertex-pair u, v write $u \rightarrow v$ for the set of paths π from u to v . Then, a.s.*

$$\text{perc}(u, v) = \min_{\pi: u \rightarrow v} \max\{\text{len}(e): e \in \pi\}. \tag{11}$$

Proof. Let us say that G has the uniqueness property if for every vertex-pair $u, v \in G$, the graph $G_{\text{len}(u, v)}$ has at most one infinite component (note that this notion was used in the proof of Lemma 2.1 in [12]). Part (a) will follow from

the fact that Φ^O has the uniqueness property, which implies:

$$\begin{aligned} e = (v, w) \in \mathcal{F}_\infty &\iff v \text{ and } w \text{ are in different components of } G_{\text{len}(e)} \\ &\iff \text{perc}(e) \geq \text{len}(e) \\ &\iff \text{perc}(e) = \text{len}(e). \end{aligned}$$

To show that Φ^O has the uniqueness property almost surely it is enough to show

$$\mathbb{P}(\forall u \in \Phi^O, G_{\text{len}(O,u)} \text{ has at most one infinite component}) = 1. \quad (12)$$

This last fact follows from Theorem 1.8 (and Remark 1.10) of [6], which implies (see also [5]),

$$\mathbb{P}(G_t \text{ includes at most one infinite component for each } t \in \mathbb{R}) = 1. \quad (13)$$

Note that (12) can be proved without appealing to the simultaneous uniqueness result as follows:

$$\begin{aligned} &\mathbb{P}(\forall u \in \Phi^O, G_{\text{len}(O,u)} \text{ has at most one infinite component}) \\ &= \lim_{n \rightarrow \infty} \mathbb{P}(\forall u \in \Phi^O \cap B(n), G_{\text{len}(O,u)} \text{ has at most one infinite component}) \\ &\geq \lim_{n \rightarrow \infty} \mathbb{P}(\forall u \in \Phi^O \cap B(n), G_{\text{len}(O,u)} \setminus B(n) \text{ has at most one infinite component}), \end{aligned}$$

where $B(n)$ is the ball of center the origin and radius n and for any network G on \mathbb{R}^d , $G \setminus B(n)$ is the subnetwork with edges and vertices in $\mathbb{R}^d \setminus B(n)$. By independence and the fact that there can be at most one infinite component in continuum percolation (see Theorem 3.6 in [10]), we have

$$\mathbb{P}(\forall u \in \Phi^O \cap B(n), G_{\text{len}(O,u)} \setminus B(n) \text{ has at most one infinite component} \mid \Phi^O \cap B(n)) = 1 \quad \text{a.s.}$$

which proves (12).

We now prove (b). Let $t = \text{perc}(u, v)$, the definition of $\text{perc}(u, v)$ may be restated easily as:

$$t = \text{perc}(u, v) = \inf_{\pi: u \rightarrow v} \max\{\text{len}(e): e \in \pi\}.$$

Hence (b) amounts to prove that with probability one, this infimum is indeed a minimum. Note that $G_t(u) \cap G_t(v) = \emptyset$ and by (13) a.s. at least one of these two clusters, say $G_t(u)$, is finite. Let E be the set of edges with exactly one of its end vertices in $G_t(u)$ and the other one in $G_t(u)^c$, and with edge length less than $t + 1$. The set E is a.s. finite and then we easily see that $\min\{\text{len}(e), e \in E\} = t = \text{perc}(u, v)$ since u and v are in the same component of $G_{t+\epsilon}$ for any $\epsilon > 0$. Let $e^* = \arg \min\{\text{len}(e), e \in E\}$ and write $e^* = (a, b)$ with $a \in G_t(u)$ and $b \in G_t(u)^c$. Since a.s. we have $\text{len}(e) \neq \text{len}(e^*) = t$ for any $e \neq e^*$, a.s. we have $G_{t+}(u) := \bigcap_{\epsilon > 0} G_{t+\epsilon}(u) = G_t(u) \cup \{e^*\} \cup G_t(b)$ and $G_{t+}(v) = G_t(v)$. But the definition of $t = \text{perc}(u, v)$ implies that $G_{t+}(u) = G_{t+}(v)$, and hence $b \in G_t(v)$. It follows that

$$\inf_{\pi: u \rightarrow v} \max\{\text{len}(e): e \in \pi\} = \text{len}(e^*) = \min_{\pi: u \rightarrow v} \max\{\text{len}(e): e \in \pi\},$$

and (b) follows. □

3.2. Finite density

Define the measure μ on $(0, +\infty)$ by

$$\begin{aligned} \mu(0, x) &= \mathbb{E} \sum_i \mathbb{1}(0 < \text{len}(O, \eta_i) - \text{perc}(O, \eta_i) < x) \\ &= \mathbb{E} \sum_i \mathbb{1}(0 < \text{len}(O, \eta_i) - \text{perc}(O, \eta_i) \leq x). \end{aligned}$$

The next lemma formalizes the heuristic idea $f_\mu(0^+) < \infty$ from Section 1.2.

Proposition 8. *In Model 2, we have,*

$$\limsup_{x \downarrow 0} \frac{\mu(0, x)}{x} < \infty.$$

For $(X_1, \dots, X_n) \in (\mathbb{R}^d)^n$ we define $\Phi^{X_1, \dots, X_n} = \Phi + \sum_{i=1}^n \delta_{X_i}$, and write $\mathbb{P}^{X_1, \dots, X_n}$ for the probability measure associated with the random variable Φ^{X_1, \dots, X_n} . Using Campbell’s formula, we have

$$\mu(0, x) = \omega_d \int_0^\infty \mathbb{P}^{O, \underline{t}}(\text{perc}(O, \underline{t}) \in [t - x, t]) t^{d-1} dt,$$

where \underline{t} is the point $(t, 0, \dots, 0)$ and $\omega_d = \frac{2\pi^{d/2}}{\Gamma(d/2)}$ is the surface of the unit sphere.

We need to introduce some continuum percolation terminology. For any r and λ , we define the probability measure $\mathbb{P}_r^{O, \underline{t}}$ under which Φ is a Poisson point process of intensity 1 and an edge e from the complete graph $\Phi^{O, \underline{t}}$ is said to be open (resp. closed) if $\text{len}(e) < r$ (resp. $\text{len}(e) \geq r$). We denote by G^O the open cluster containing the origin: $G^O = G_r^O$. Let r_c be the critical radius for the Poisson continuum percolation model of density 1 and deterministic radius, i.e. for $r < r_c$ the number of vertices in any open cluster is finite whereas for $r > r_c$ there exists an unique unbounded open cluster.

Write C, C_1, C_2 for positive constants not depending on the parameters of the problem.

Lemma 9. *For any $\epsilon > 0$, we have*

$$\begin{aligned} \text{for } 0 < t < r_c - \epsilon, \quad & \mathbb{P}^{O, \underline{t}}(\text{perc}(O, \underline{t}) \in [t - x, t]) \leq C_1 x, \\ \text{for } t > r_c + \epsilon, \quad & \mathbb{P}^{O, \underline{t}}(\text{perc}(O, \underline{t}) \in [t - x, t]) \leq C_1 x e^{-C_2 t}. \end{aligned}$$

We first introduce some notations. The edge-length is the Euclidean distance denoted $\text{len}(u, v) = |u - v|$. For a set $S \subset \mathbb{R}^d$, we denote by $d(S) = \sup\{|x - y|, x, y \in S\}$ its diameter. For $x \in \mathbb{R}^d$ and $r > 0$, $B(x, r)$ denotes the open ball of radius r centered at x . For $t > 0$ we denote $S(t) = [-t, t]^d$. Under the probability measure $\mathbb{P}_r^{O, \underline{t}}$, the occupied region is $\bigcup_{X \in \Phi^{O, \underline{t}}} B(X, r/2)$ and the vacant region is the complement of the occupied region. The occupied component of the origin W is defined by $\mathbb{P}_r^{O, \underline{t}}(W = \bigcup_{X \in G^O} B(X, r/2)) = 1$. The vacant component containing the point $\underline{t}/2$ is denoted by V . More generally, for $r > 0$ the occupied region at level r is $\bigcup_{X \in \Phi^{O, \underline{t}}} B(X, r/2)$ and we denote $W_r = \bigcup_{X \in G_r^O} B(X, r/2)$ the occupied component of the origin at level r and V_r the vacant component containing the point $\underline{t}/2$ at level r .

Since we may assume that all interdistances are different, there exists an unique pair (X, Y) in the support of $\Phi^{O, \underline{t}}$ such that $\text{perc}(O, \underline{t}) = |X - Y|$ (see Lemma 7).

First consider the case $t < r_c - \epsilon$. Let $S_z = zt/2 + S(t/2)$ where $z \in \mathbb{Z}^d$. If the event $\{\text{perc}(O, \underline{t}) \in [t - x, t]\}$ occurs, there is some $z \in \mathbb{Z}^d$ such that $W \cap S_z \neq \emptyset$ and there exists $X, Y \in \Phi \cap S_z$ such that $|X - Y| \in [t - x, t]$. Note that we have for any $z \in \mathbb{Z}^d$,

$$\mathbb{P}_t^{O, \underline{t}}(\exists X, Y \in \Phi \cap S_z, |X - Y| \in [t - x, t]) \leq C(1 + t^{2d-1})x.$$

Hence we have

$$\begin{aligned} & \mathbb{P}^{O, \underline{t}}(\text{perc}(O, \underline{t}) \in [t - x, t]) \\ & \leq \sum_{z \in \mathbb{Z}^d} \mathbb{P}_t^{O, \underline{t}}(W \cap S_z \neq \emptyset, \exists X, Y \in \Phi(S_z), |X - Y| \in [t - x, t]) \\ & \leq \left(K + \sum_{z, \|z\| \geq 2} \mathbb{P}_t^{O, \underline{t}}(W \cap S_z \neq \emptyset) \right) C(1 + t^{2d-1})x, \end{aligned} \tag{14}$$

where $\|z\| := \max(|z_1|, |z_2|)$ and K is a constant depending on d . Lemma 3.3 of [10] ensures that the sum of (14) is finite for $t < r_c - \epsilon$.

The case $t > r_c + \epsilon$ is quite similar. If the event $\{\text{perc}(O, \underline{t}) \in [t - x, t)\}$ occurs, there is some $z \in \mathbb{Z}^2$ such that $V_{t-x} \cap S_z \neq \emptyset$ and there exists $X, Y \in \Phi \cap S_z$ such that $|X - Y| \in [t - x, t)$. Hence we have

$$\begin{aligned} \mathbb{P}^{O, \underline{t}}(\text{perc}(O, \underline{t}) \in [t - x, t)) &\leq \left(5 + \sum_{z, \|z - \underline{t}/2\| \geq 2} \mathbb{P}_t^{O, \underline{t}}(V_{t-x} \cap S_z \neq \emptyset)\right) C(1 + t^{2d-1})x \\ &\leq C_1 e^{-C_2 t} x, \end{aligned} \quad (15)$$

where (15) follows from Lemma 4.1 of [10] and the fact that

$$\mathbb{P}_t^{O, \underline{t}}(V_{t-x} \cap S_z \neq \emptyset) \leq \mathbb{P}_t^{O, \underline{t}}\left(d(V_{t-x}) > \left\|z - \frac{\underline{t}}{2}\right\|\right).$$

We now concentrate on the case $t \in (r_c - \epsilon, r_c + \epsilon)$. We define the event

$$A = \{\text{the points of } \Phi \text{ on the axis } e_1 \text{ are in } G^O\}, \quad \text{where } e_1 = (1, 0, \dots, 0).$$

Under $\mathbb{P}_r^{O, \underline{t}}$, with probability one, we have $A = \{\underline{t} \in G^O\}$ and

$$\begin{aligned} \mathbb{P}^{O, \underline{t}}(\text{perc}(O, \underline{t}) \in [t - x, t)) &= \mathbb{P}_t^{O, \underline{t}}(A) - \mathbb{P}_{t-x}^{O, \underline{t}}(A) \\ &= \mathbb{P}_t^{O, \underline{t}}(A) - \mathbb{P}_{t-x}^{O, t-x}(A) + \mathbb{P}_{t-x}^{O, t-x}(A) - \mathbb{P}_{t-x}^{O, \underline{t}}(A). \end{aligned} \quad (16)$$

We first prove that

$$\left| \mathbb{P}_{t-x}^{O, t-x}(A) - \mathbb{P}_{t-x}^{O, \underline{t}}(A) \right| \leq C t^d x. \quad (17)$$

Note that

$$\begin{aligned} \mathbb{P}_{t-x}^{O, \underline{t}}(A) &= \mathbb{P}_{t-x}^O(B(\underline{t}, t-x) \cap G^O \neq \emptyset), \\ \mathbb{P}_{t-x}^{O, t-x}(A) &= \mathbb{P}_{t-x}^O(B(\underline{t-x}, t-x) \cap G^O \neq \emptyset), \end{aligned}$$

where $B(X, r)$ denotes the open ball of radius $r > 0$ centered at $X \in \mathbb{R}^d$. Hence we have

$$\left| \mathbb{P}_{t-x}^{O, t-x}(A) - \mathbb{P}_{t-x}^{O, \underline{t}}(A) \right| \leq \mathbb{P}(\Phi(B(\underline{t}, t-x) \Delta B(\underline{t-x}, t-x)) \geq 1) \leq C t^d x,$$

where $B(\underline{t}, t) \Delta B(\underline{t-x}, t-x)$ denotes the symmetric difference. This is exactly (17).

We then write:

$$\begin{aligned} \frac{1}{x} \int_{r_c - \epsilon}^{r_c + \epsilon} \mathbb{P}^{O, \underline{t}}(\text{perc}(O, \underline{t}) \in [t - x, t)) t^{d-1} dt &= \frac{1}{x} \int_{r_c - \epsilon}^{r_c + \epsilon} \mathbb{P}_t^{O, \underline{t}}(A) t^{d-1} dt - \frac{1}{x} \int_{r_c - \epsilon}^{r_c + \epsilon} \mathbb{P}_{t-x}^{O, t-x}(A) t^{d-1} dt \\ &\quad + \frac{1}{x} \int_{r_c - \epsilon}^{r_c + \epsilon} (\mathbb{P}_{t-x}^{O, t-x}(A) - \mathbb{P}_{t-x}^{O, \underline{t}}(A)) t^{d-1} dt. \end{aligned} \quad (18)$$

With the change of variable $t \mapsto t - x$, the second term on the right-hand side of (18) is decomposed as follows:

$$\frac{1}{x} \int_{r_c - \epsilon}^{r_c + \epsilon} \mathbb{P}_{t-x}^{O, t-x}(A) t^{d-1} dt \leq \frac{1}{x} \int_{r_c - \epsilon - x}^{r_c + \epsilon - x} \mathbb{P}_t^{O, \underline{t}}(A) t^{d-1} dt + K \int_{r_c - \epsilon - x}^{r_c + \epsilon - x} \mathbb{P}_t^{O, \underline{t}}(A) t^{d-2} dt,$$

where K is a constant depending on d .

Hence, the decomposition (18) is further decomposed as

$$\begin{aligned} \left| \frac{1}{x} \int_{r_c-\epsilon}^{r_c+\epsilon} \mathbb{P}^{O,t}(\text{perc}(O, \underline{t}) \in [t-x, t]) t^{d-1} dt \right| &\leq \frac{1}{x} \int_{r_c+\epsilon-x}^{r_c+\epsilon} \mathbb{P}_t^{O,t}(A) t^{d-1} dt + \frac{1}{x} \int_{r_c-\epsilon-x}^{r_c-\epsilon} \mathbb{P}_t^{O,t}(A) t^{d-1} dt \\ &\quad + K \int_{r_c-\epsilon-x}^{r_c+\epsilon-x} \mathbb{P}_t^{O,t}(A) t^{d-2} dt \\ &\quad + \frac{1}{x} \int_{r_c-\epsilon}^{r_c+\epsilon} (\mathbb{P}_{t-x}^{O,t-x}(A) - \mathbb{P}_{t-x}^{O,t}(A)) t^{d-1} dt. \end{aligned}$$

By (17), the last term is bounded by $\int_{r_c-\epsilon}^{r_c+\epsilon} C t^{2d-1} dt = C_1$. It implies that

$$\left| \frac{1}{x} \int_{r_c-\epsilon}^{r_c+\epsilon} \mathbb{P}^{O,t}(\text{perc}(O, \underline{t}) \in [t-x, t]) t^{d-1} dt \right| \leq C_2. \tag{19}$$

Proposition 8 now follows from Lemma 9 and Eq. (19).

3.3. The lower bound in Model 2

We start with a slight extension of Proposition 9 of [3] (see also Theorem 7 in [4]). In what follows, a set of points is identified with its associated geometric graph which is the complete graph over these points with Euclidean distance as edge-lengths.

Lemma 10. *Let Φ_n denote the point process consisting of n points $\{\xi_i, 1 \leq i \leq n\}$ which are independent and have the uniform distribution on the square $[0, n^{1/d}]^d$. For each n , let U_n be chosen independently and uniformly from the set $\{1, \dots, n\}$, and let*

$$\Phi_n^O = \{\xi_i^{(n)} := \xi_i - \xi_{U_n}, 1 \leq i \leq n\}.$$

To each vertex $\xi_i^{(n)}$ of the rooted (at the origin) geometric graph Φ_n^O , we associate the mark $\text{perc}_i^n = \text{perc}(O, \xi_i^{(n)})$ as defined in (1). We denote by $(\Phi_n^O, \text{perc}^n) = \{\xi_i^{(n)}, \text{perc}_i^n\}$ the corresponding marked geometric graph. Then one has joint weak convergence

$$((\Phi_n^O, \text{perc}^n), \text{MST}(\Phi_n^O)) \xrightarrow{d} ((\Phi^O, \text{perc}), \mathcal{F}_\infty), \tag{20}$$

where (Φ^O, perc) is the Palm version of the Poisson process of intensity 1 with the mark $\text{perc}(O, \eta_i)$ associated to point η_i .

Here convergence $\text{MST}(\Phi_n^O) \xrightarrow{d} \mathcal{F}_\infty$ is local weak convergence in the sense of [4].

Proof of Lemma 10. The analog of (20) without marks is Proposition 9 of [3]. By the Skorokhod representation theorem, we can assume that with probability one, we have

$$(\Phi_n^O = \{\xi_i^{(n)}\}, \text{MST}(\Phi_n^O)) \rightarrow (\Phi^O = \{\eta_i\}, \mathcal{F}_\infty). \tag{21}$$

We have to prove that for any $i \geq 1$,

$$\lim_{n \rightarrow \infty} \text{perc}(O, \xi_i^{(n)}) = \text{perc}(O, \eta_i) \quad \text{a.s.}$$

By Lemma 7, we know that $\text{perc}(O, \eta_i) = \max\{\text{len}(e), e \in \pi^*\}$ where π^* is the minimax path from O to η_i . By definition of the metric of local weak convergence, (21) implies that for arbitrary fixed L , we have with $S(L) = [-L, L]^d$,

$$\forall \eta_i \in S(L), \quad \xi_i^{(n)} \rightarrow \eta_i.$$

For L sufficiently large, the path π^* is included in $S(L)$ and let π_n^* be the associated path in Φ_n^O . Since

$$\text{perc}(O, \xi_i^{(n)}) = \min_{\pi: O \rightarrow \xi_i^{(n)}} \max_{e \in \pi} \text{len}(e) \leq \max_{e \in \pi_n^*} \text{len}(e),$$

by the convergence of π_n^* to π^* , we have

$$\limsup_{n \rightarrow \infty} \text{perc}(O, \xi_i^{(n)}) \leq \max\{\text{len}(e), e \in \pi^*\} = \text{perc}(O, \eta_i).$$

Now we need to prove that

$$\liminf_{n \rightarrow \infty} \text{perc}(O, \xi_i^{(n)}) \geq \text{perc}(O, \eta_i). \tag{22}$$

Take $\pi^{(n)}: O \rightarrow \xi_i^{(n)}$ such that

$$\max_{e \in \pi^{(n)}} \text{len}(e) = \text{perc}(O, \xi_i^{(n)}).$$

For $r > 0$, we denote by $G_r(\eta_i)$ (resp. $G_r^n(\xi_i^{(n)})$) the connected component of Φ^O (resp. Φ_n^O) with edge length less than r containing η_i (resp. $\xi_i^{(n)}$). Let $\text{perc}(O, \eta_i) = t$, so that we have $G_t(O) \cap G_t(\eta_i) = \emptyset$ and say $G_t(O)$ is finite (see the uniqueness property in the proof of Lemma 7). We define \tilde{G}_t (resp. \tilde{G}_t^n) to be the subgraph of Φ^O (resp. Φ_n^O) consisting of those edges with length less than $t + 1$ with exactly one of its end vertices in $G_t(O)$ (resp. $G_t^n(O)$). Let $e^* = \arg \max\{\text{len}(e), e \in \pi^*\}$. By Lemma 7, we know that $\text{perc}(O, \eta_i) = \text{len}(e^*)$ and $e^* \in \tilde{G}_t$ is such that $e^* = \arg \min\{\text{len}(e), e \in \tilde{G}_t\}$. Since $G_t(O)$ is finite, we have clearly that \tilde{G}_t is included in $S(L)$ for sufficiently large L . Then we have

$$\max_{e \in \pi^{(n)}} \text{len}(e) \geq \min\{\text{len}(e), e \in \tilde{G}_t^n\} \rightarrow t = \text{perc}(O, \eta_i) \quad \text{as } n \rightarrow \infty,$$

where the last limit follows from the convergence of \tilde{G}_t^n to \tilde{G}_t .

We now return to the proof of the lower bound in Model 2. We start by copying and modifying the argument from Section 2.4. Fix $\delta > 0$. Let ξ_{U_n} be a uniform random vertex from $(\xi_i, 1 \leq i \leq n)$. Consider a pair (T'_n, T_n) attaining the minimum in the definition (3) of $\varepsilon_n(\delta)$. Then

$$\mathbb{E} \sum_i \mathbb{1}\{(\xi_{U_n}, \xi_i) \in T'_n \setminus T_n\} = \frac{2\mathbb{E}|T'_n \setminus T_n|}{n} \geq 2\delta \tag{23}$$

and

$$\begin{aligned} \mathbb{E}\varepsilon_n(\delta) &= \frac{1}{n} \mathbb{E}(\text{len}(T'_n) - \text{len}(T_n)) \\ &\geq \frac{1}{n} \mathbb{E} \sum_{e \in T'_n \setminus T_n} \text{exc}(e) \quad \text{by Lemma 5} \\ &= \frac{1}{2} \mathbb{E} \sum_i \text{exc}(\xi_{U_n}, \xi_i) \mathbb{1}\{(\xi_{U_n}, \xi_i) \in T'_n \setminus T_n\}. \end{aligned} \tag{24}$$

Note that for $0 < L < \infty$

$$\begin{aligned} \mathbb{E} \sum_i \mathbb{1}\{\text{len}(\xi_{U_n}, \xi_i) \geq L, (\xi_{U_n}, \xi_i) \in T'_n\} &= \frac{2}{n} \mathbb{E}|\{e \in T'_n: \text{len}(e) \geq L\}| \\ &\leq \frac{2}{n} \frac{\mathbb{E} \text{len}(T'_n)}{L} \\ &\leq \delta \quad \text{for } L = L(\delta) \text{ sufficiently large} \end{aligned}$$

the last inequality because $\mathbb{E} \text{len}(T'_n) = O(n)$. So fixing such an L , (23) implies

$$\mathbb{E} \sum_i \mathbb{1}\{(\xi_{U_n}, \xi_i) \in T'_n \setminus T_n, \text{len}(\xi_{U_n}, \xi_i) \leq L\} \geq \delta \quad (25)$$

while (24) trivially implies

$$2\mathbb{E}\varepsilon_n(\delta) \geq \mathbb{E} \sum_i \text{exc}(\xi_{U_n}, \xi_i) \mathbb{1}\{(\xi_{U_n}, \xi_i) \in T'_n \setminus T_n, \text{len}(\xi_{U_n}, \xi_i) \leq L\}. \quad (26)$$

The purpose of these representations is to exploit local weak convergence. Consider the near-minimal STs T'_n appearing in (25), (26). By a compactness argument and by passing to a subsequence of n we may assume that they converge to some forest \mathcal{F}'_∞ on $(\tilde{\eta}_i)$; that is, we may assume that (20) remains true when we append T'_n to the left side and \mathcal{F}'_∞ to the right side. We can now take limits in (25) to deduce

$$\sum_i \mathbb{P}\{(O, \eta_i) \in \mathcal{F}'_\infty \setminus \mathcal{F}_\infty, \text{len}(O, \eta_i) \leq L\} \geq \delta.$$

And taking limits in (26) gives

$$2 \liminf_n \mathbb{E}\varepsilon_n(\delta) \geq \mathbb{E} \sum_i (\text{len}(O, \eta_i) - \text{perc}(O, \eta_i)) \mathbb{1}\{(O, \eta_i) \in \mathcal{F}'_\infty \setminus \mathcal{F}_\infty, \text{len}(O, \eta_i) \leq L\}. \quad (27)$$

Writing

$$\begin{aligned} V_i &= \text{len}(O, \eta_i) - \text{perc}(O, \eta_i), \\ A_i &= \{(O, \eta_i) \in \mathcal{F}'_\infty \setminus \mathcal{F}_\infty, \text{len}(O, \eta_i) \leq L\}, \end{aligned}$$

we are precisely in the setting in which Proposition 8 and Lemma 6(b) apply, and the conclusion is that the right side of (27) is $\geq (\beta - o(1))\delta^2$ for small δ , implying the lower bound in Theorem 2. \square

References

- [1] D. J. Aldous and A. G. Percus. Scaling and universality in continuous length combinatorial optimization. *Proc. Natl. Acad. Sci. USA* **100** (2003) 11211–11215. [MR2007286](#)
- [2] D. J. Aldous. The $\zeta(2)$ limit in the random assignment problem. *Random Structures Algorithms* **18** (2001) 381–418. [MR1839499](#)
- [3] D. J. Aldous and J. M. Steele. Asymptotics for Euclidean minimal spanning trees on random points. *Probab. Theory Related Fields* **92** (1992) 247–258. [MR1161188](#)
- [4] D. J. Aldous and J. M. Steele. The objective method: Probabilistic combinatorial optimization and local weak convergence. In *Probability on Discrete Structures* 1–72. H. Kesten (Ed.), Springer, Berlin, 2003. [MR2023650](#)
- [5] K. S. Alexander. Percolation and minimal spanning forests in infinite graphs. *Ann. Probab.* **23** (1995) 87–104. [MR1330762](#)
- [6] K. S. Alexander. Simultaneous uniqueness of infinite clusters in stationary random labeled graphs. *Comm. Math. Phys.* **168** (1995) 39–55. [MR1324390](#)
- [7] G. Chartrand and L. Lesniak. *Graphs and Digraphs*, 2nd edition. Wadsworth, Monterey, CA, 1986. [MR0834583](#)
- [8] L. P. Kadanoff. *Statistical Physics*. World Scientific, River Edge, NJ, 2000. [MR1772071](#)
- [9] W. Krauth and M. Mézard. The cavity method and the travelling-salesman problem. *Europhys. Lett.* **8** (1989) 213–218.
- [10] R. Meester and R. Roy. *Continuum Percolation*. Cambridge Univ. Press, 1996. [MR1409145](#)
- [11] J. M. Steele. *Probability Theory and Combinatorial Optimization*. SIAM, Philadelphia, PA, 1997. [MR1422018](#)
- [12] M. Penrose and J. E. Yukich. Weak laws of large numbers in geometric probability. *Ann. Appl. Probab.* **13** (2003) 277–303. [MR1952000](#)
- [13] J. E. Yukich. *Probability Theory of Classical Euclidean Optimization Problems*. Springer, Berlin, 1998. [MR1632875](#)