Climbing up a random subgraph of the hypercube

Michael Anastos [*] Sahar Diskin [†] Dor Elboim [‡]
Michael Krivelevich [†]

Abstract

Let Q^d be the d-dimensional binary hypercube. We say that $P = \{v_1, \dots, v_k\}$ is an increasing path of length $k-1$ in Q^d, if for every $i \in [k-1]$ the edge v_iv_{i+1} is obtained by switching some zero coordinate in v_i to a one coordinate in v_{i+1}.
Form a random subgraph Q^d_p by retaining each edge in $E(Q^d)$ independently with probability p. We show that there is a phase transition with respect to the length of a longest increasing path around $p = \frac{e}{d}$. Let α be a constant and let $p = \frac{\alpha}{d}$. When $0 < \alpha < e$, then there exists a $\delta \in (0,1)$ such that **whp** a longest increasing path in Q^d_p is of length at most $(1-\delta)d$. On the other hand, when $\alpha > e$, **whp** there is a path of length $d-2$ in Q^d_p, and in fact, whether it has length $d-2, d-1$, or d depends on whether the vertices $(0,\dots,0)$ and $(1,\dots,1)$ are in the giant connected component.

Keywords: Bond percolation; Hypercube; Phase transition.
MSC2020 subject classifications: NA.

1 Introduction

1.1 Background and main result

The *d-dimensional hypercube*, Q^d, is the graph whose vertex set is $V(Q^d) := \{0,1\}^d$, and where two vertices are adjacent if they differ in exactly one coordinate. The hypercube and its subgraphs arise naturally in many contexts and have received much attention in combinatorics, probability, and computer science.

The *random subgraph* Q^d_p is obtained by retaining each edge in $E(Q^d)$ independently with probability p. Several phase transitions have been observed in Q^d_p. Indeed, the study of Q^d_p was initiated by Sapoženko [13] and Burtin [7], who showed that the threshold for connectivity is $\frac{1}{2}$: when $p < \frac{1}{2}$, **whp**[1] Q^d_p is disconnected, whereas for $p > \frac{1}{2}$, **whp** Q^d_p is connected. This result was subsequently strengthened by Erdős and Spencer [10] and by Bollobás [4]. Bollobás further showed [5] that $p = \frac{1}{2}$ is the threshold for the existence of a perfect matching. Recently, resolving a long-standing open problem, Condon, Espuny

[*]Institute of Science and Technology Austria (ISTA), Klosterneurburg 3400, Austria. Email: michael.anastos@ist.ac.at.
Research supported by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 101034413
[†]School of Mathematical Sciences, Tel Aviv University, Tel Aviv 6997801, Israel. Emails: sahardiskin@mail.tau.ac.il, krivelev@tauex.tau.ac.il.
[‡]School of Mathematics, Institute for Advanced Study, Princeton, N.J. 08540, U.S.A.. Email: dorelboim@gmail.com.
[1]With high probability, that is, with probability tending to one as d tends to infinity.

Díaz, Girão, Kühn, and Osthus [8] showed that $p = \frac{1}{2}$ is also the threshold for the existence of a Hamilton cycle in Q^d_p.

In the sparser regime, Erdős and Spencer asked [10] whether Q^d_p undergoes a phase transition with respect to its component structure around $p = \frac{1}{d}$, similar to that of $G(n,p)$ around $p = \frac{1}{n}$. This was confirmed by Ajtai, Komlós, and Szemerédi [1], with subsequent work by Bollobás, Kohayakawa, and Łuczak [6]. See also the recent work [3] and the references therein for the behaviour of the critical regime. Given $\alpha > 1$, let ζ_α be the unique solution in $(0,1)$ of the equation

$$\zeta_\alpha = 1 - \exp(-\alpha\zeta_\alpha). \tag{1.1}$$

Then, Ajtai, Komlós, and Szemerédi [1], and Bollobás, Kohayakawa, and Łuczak [6] showed that when $p = \frac{1-\epsilon}{d}$, for $\epsilon > 0$, **whp** all components of Q^d_p have order $O_\epsilon(d)$, and when $p = \frac{1+\epsilon}{d}$ **whp** Q^d_p contains a unique giant component of asymptotic order $\zeta_{1+\epsilon}2^d$, and all other components have order $O_\epsilon(d)$. Here, we use the convention $f(x) = O_r(g(x))$ to say that there exist a constant C, which may depend on r, such that for all sufficiently large x we have $f(x) \le C|g(x)|$. We note that ζ_α is equal to the survival probability of a Galton-Watson tree with offspring distribution $\mathrm{Po}(\alpha)$.

In this paper, we show a phase transition in the hypercube, which occurs when $p = \frac{e}{d}$. Before stating our result, let us introduce some notation. For a vertex $v \in Q^d$, we denote by $I(v) \subseteq [d]$ the set of coordinates of v which are 1. For $v_1, v_2 \in Q^d$, we say that $v_1 < v_2$ if $I(v_1) \subsetneq I(v_2)$. Given a path $P = \{v_1, \ldots, v_k\}$ in Q^d (or Q^d_p), we say that P is an *increasing path* of length $k-1$ in the hypercube, if for every $i \in [k-1]$, $v_i < v_{i+1}$ and $|I(v_{i+1})| - |I(v_i)| = 1$. Note that the longest increasing path in Q^d is of length d. Given a subgraph $H \subseteq Q^d$, let $\ell(H)$ be the length of a longest increasing path in H. Our result is as follows.

Theorem 1.1. *Let α be a constant, and let $p = \frac{\alpha}{d}$. Then, the following holds.*

*(a) For every $0 < \alpha < e$, there exists $\delta := \delta(\alpha)$, $0 < \delta < 1$, such that **whp** $\ell(Q^d_p) \le (1-\delta)d$.*

*(b) For every $\alpha > e$, **whp** $\ell(Q^d_p) \ge d-2$. Furthermore,*

$$\mathbb{P}\big(\ell(Q^d_p) = d\big) = (1+o(1))\zeta_\alpha^2,$$
$$\mathbb{P}\big(\ell(Q^d_p) = d-1\big) = (1+o(1))2\zeta_\alpha(1-\zeta_\alpha),$$

where ζ_α is defined according to (1.1).

In a sense, the above shows that for $\alpha > e$ **whp** a longest increasing path is of length at least $d-2$, and whether it is of length $d-2, d-1$, or d depends on whether the all-0-vertex and the all-1-vertex 'percolate'; the probability that a vertex percolates is approximated by the probability that an appropriate branching process survives, and is asymptotically equal to the probability the vertex belongs to the largest component (see Proposition 2.2 for a more precise statement).

Let us mention a related result of Pinsky [12], who considered (among other things) the *number* of increasing paths of length d in Q^d_p, where $p = \frac{\alpha}{d}$. In [12], he showed that if $\alpha < e$, then **whp** the number of such paths is 0. Theorem 1.1(a) shows that, in fact, when $\alpha < e$, typically a longest path is smaller by a multiplicative constant. Pinsky further showed that if $\alpha > e$, then the probability there are any such paths is bounded away from 0 and 1. Theorem 1.1(b) gives a detailed description of the typical length of a longest path when $\alpha > e$.

We finish this section by noting that while some of our lemmas extend to the case where $\alpha = e$, our overall proof does not. It would be interesting to see what the behaviour is at this critical point.

Question 1.2. Let $p = \frac{e}{d}$. Form Q^d_p by retaining each edge of Q^d independently with probability p. What can be said about a longest increasing path in Q^d_p?

1.2 Proof outline

The proof of Theorem 1.1(a) follows from a first moment argument. The proof of Theorem 1.1(b), on the other hand, is far more delicate.

For proving Theorem 1.1(b), we begin by showing, through a careful second-moment argument, that it is not *very* unlikely to have a path between the all-0-vertex and the all-1-vertex in Q^d_p — in fact, we give a lower bound for the probability of that event which is inverse-polynomial in d. Then, our goal is to show that if the all-0-vertex and the all-1-vertex 'percolate', then **whp** we can find polynomially many vertex disjoint subcubes, so that the 0 antipodal point of each of the subcubes is connected by a decreasing path in Q^d_p to the all-0-vertex in Q^d, and the 1 antipodal point of each of the subcubes is connected by an increasing path in Q^d_p to the all-1-vertex in Q^d. Since these subcubes are vertex disjoint, the events that the 0 antipodal point is connected by an increasing path to the 1 antipodal point in each of the subcubes are independent for different subcubes. Thus, since for each subcube the probability that there is an increasing path between the 0 antipodal point and the 1 antipodal point in Q^d_p is at least inverse polynomial in d, having polynomially many such subcubes we will be able to conclude that **whp** in at least one of these subcubes there is an increasing path from the 0 antipodal point to the 1 antipodal point, which then extends to an increasing path between the all-0-vertex and the all-1-vertex in Q^d_p.

Finding these subcubes is the most involved part of the paper, and therein lie several novel ideas. We note that key parts of this lie in the Tree Construction algorithm and its properties, given in Section 2.1.

Roughly, we show that if the all-0-vertex 'percolates', then **whp** we can construct a 'good' tree in Q^d_p, denote it by T_0, which is rooted at the all-0-vertex. We note that T_0 is monotone increasing, that is, the layers of the tree correspond to the layers of the hypercube. We construct this tree so that it has polynomially in d many leaves, all of which reside in a relatively low layer (that is, $O(\log d)$). Furthermore, we assign a set of coordinates C_v to every leaf $v \in V(T_0)$, where $C_v \cap I(v) = \varnothing$. While constructing this tree, we ensure that this list is of size at most polylogarithmic in d, and that for every other leaf $v' \in V(T_0)$, there is some $i \in C_v$ such that $i \in I(v')$ (we informally say that the leaves of the tree are 'easily distinguishable'). Let us denote by V_0 the set of leaves of T_0.

Then, we show that if the all-1-vertex 'percolates', then **whp** we can construct a 'good' tree in Q^d_p, denote it by T_1, which is rooted at the all 1-vertex. Here T_1 is monotone decreasing, that is, the layers of the tree correspond (in decreasing order) to the layers of the hypercube. The set of leaves of this tree, V_1, resides in a relatively high layer (that is, $d - O(\log d)$), and we have that $|V_1| = d|V_0|$. Moreover, we will construct the tree so that all of its vertices are above the vertices of T_0, that is, for every $v_0 \in V_0$ and $v_1 \in V_1$, we have that $I(v_0) \subset I(v_1)$. While this might seem a stringent requirement, note that given $p = \frac{\alpha}{d}$ with $\alpha > e$, we have that $p \cdot \frac{d}{2} > 1$, and thus we can restrict the growth of the trees to half of the coordinates while remaining supercritical. Indeed, at a certain point through the construction of the trees, we will grow T_0 only on the first $\frac{d}{2}$ coordinates, and T_1 only on the last $\frac{d}{2}$ coordinates. After constructing T_1, we will arbitrarily associate disjoint sets of d vertices in V_1 to every vertex in V_0.

Since the leaves of T_0 are easily distinguishable, we utilise a projection-type argument to argue that given a set of, say, d leaves in T_1 and a vertex $v_0 \in V(T_0)$, **whp** we can grow in Q^d_p a tree of small height, rooted at one of these leaves, such that it has a vertex v_1 which is above v_0, but is not above any other leaf of T_0. We do so by growing a tree, rooted at one of the d leaves which were assigned to v_0 (which is fixed), so that all its leaves are above v_0, and to a sufficient height such that for every $i \in C_{v_0}$, we have that $i \in I(v_1)$. This, in turn, will allow us to find pairs of vertices, which will form the antipodal points of the vertex disjoint subcubes we seek to construct.

2 Preliminaries

We begin with some notation and terminology which will be of use for us throughout the paper. Given a graph G, and a vertex $v \in V(G)$, we let $N(v)$ be the set of neighbours of v in G, that is, $N(v) := \{u \in V(G) : uv \in E(G)\}$. Furthermore, given a subset $S \subseteq V(G)$, we let $N_S(v) := N(v) \cap S$, that is, the neighbourhood of v in S. With a slight abuse of notation, given a subgraph $H \subseteq G$ and $v \in V(G)$, we define $N_H(v) := N_{V(H)}(v)$. Finally, given a probability $p \in [0,1]$, we denote by G_p the random subgraph of G obtained by retaining each edge $e \in E(G)$ independently with probability p.

Recall that for every $v \in Q^d$, we denote by $I(v) \subseteq [d]$ the set of coordinates which are 1. Furthermore, for $v_1, v_2 \in Q^d$, we say that $v_1 \le v_2$ if $I(v_1) \subseteq I(v_2)$, and $v_1 < v_2$ if $I(v_1) \subsetneq I(v_2)$. Given a path $P = \{v_1, \ldots, v_k\}$ in Q^d, we say that P is an *increasing path* in the hypercube, if for every $i \in [k-1]$, $v_i < v_{i+1}$ and $|I(v_{i+1})| - |I(v_i)| = 1$. We say that a path $P' = \{v_k, \ldots, v_1\}$ is *decreasing* if $P = \{v_1, \ldots, v_k\}$ is increasing. Also, given a set of vertices $S \subseteq V(Q^d)$, let $I(S) := \cup_{v \in S} I(v)$.

Given two vertices $u, v \in Q^d$ such that $uv \in E(Q^d)$, we denote by $c(u,v)$ the coordinate in $[d]$ in which these two vertices differ. Given a tree T rooted in $r \in V(Q^d)$ and a vertex $v \in V(T)$, let

$$C_T(v) := \big\{c(x,w) : x \neq v, xw \in E(T), x \text{ is on the path } P \text{ from } r \text{ to } v \text{ in } T, w \notin P\big\}.$$

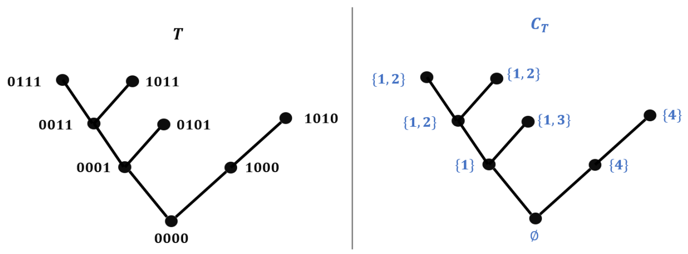

Figure 1: Illustration of $C_T(v)$ for a fixed tree T in Q^4, rooted at the all-0-vertex. On the left side, the tree T and its vertices are presented. On the right side, the values of $C_T(v)$ for every vertex of T appear in blue.

We say that a tree T is increasing if every path from the root to a leaf is increasing. We similarly define a decreasing tree. Throughout the paper we will only consider trees that are monotone (increasing or decreasing). When the tree is rooted at 0 it will be increasing and when it is rooted at 1 it will be decreasing.

Define the sets $C_T(v,0) := C_T(v) \setminus I(v)$ and $C_T(v,1) := C_T(v) \cap I(v)$.

Given $v_1, v_2 \in V(Q^d)$ such that $v_1 \le v_2$, let $Q(v_1, v_2)$ be the induced subcube of Q^d whose vertex set is given by

$$V(Q(v_1,v_2)) := \{u : v_1 \le u \le v_2\}.$$

Note that $Q(v_1, v_2)$ is isomorphic to a hypercube of dimension $|I(v_2)| - |I(v_1)|$. Furthermore, note that given u_1, u_2 and v_1, v_2, we have that $Q(u_1, u_2)$ and $Q(v_1, v_2)$ are guaranteed to be vertex disjoint if $I(u_1) \not\subseteq I(v_2)$, equivalently if $I(u_1) \setminus I(v_2) \neq \varnothing$. We denote by $\mathbf{0}$ the all-0-vertex in Q^d, and by $\mathbf{1}$ the all-1-vertex in Q^d. For $i \in [0, d]$, we denote by L_i the i-th layer in the hypercube, that is, the set of vertices with exactly i ones. Throughout the paper, we omit rounding signs for the sake of clarity of presentation. All the logarithms are assumed to be the natural logarithm.

We will use the following Chernoff-type bounds on the tail probabilities of the binomial distribution (see, for example, [2, Appendix A]).

Lemma 2.1. *Let $N \in \mathbb{N}$, let $p \in [0,1]$ and let $X \sim Bin(N,p)$.*

- *For every positive t, $\mathbb{P}(X > tNp) \leq \left(\frac{e}{t}\right)^{tNp}$.*
- *For every $0 < b \leq \frac{Np}{2}$, $\mathbb{P}(X < Np - b) \leq \exp\left(-\frac{b^2}{4Np}\right)$.*

2.1 Exploring a tree with 'easily distinguishable' leaves

We will utilise a heavily modified variant of the Breadth First Search (BFS) algorithm, which we will refer to as the *Tree Construction algorithm*. The Tree Construction algorithm is fed with the following as input.

(In1) A subcube $H := Q(u_0, u_1) \subseteq Q^d$, with $u_0 \leq u_1$, together with an order σ on $V(H)$;
(In2) a vertex $r \in \{u_0, u_1\}$, which will serve as the root;
(In3) a set of coordinates $C \subseteq [d]$, which are to be avoided;
(In4) a layer $L_i \subseteq V(Q^d)$, with $|I(u_0)| \leq i \leq |I(u_1)|$, at which this algorithm is truncated;
(In5) a sequence of independent Bernoulli(p) random variables, $\{X_e\}_{e \in E(H)}$.

The algorithm outputs a BFS-type tree T, rooted at r in H. To that end, the algorithm maintains three sets of vertices: B, the set of vertices whose exploration is complete (and are part of the BFS tree); A, the active vertices currently being explored, kept in a *queue*; and Y, the vertices that have not been explored yet. The algorithm starts with T being the root r, $B = \varnothing$, $A = \{r\}$ and $Y = V(H) \setminus \{r\}$. As long as A is not empty, the algorithm proceeds as follows.

Let x be the first vertex in A and let $j(x) \in [0,d]$ be such that x belongs to the layer $L_{j(x)}$. Let

$$C_x := \begin{cases} C_T(x,0), & \text{if } r = u_0 \\ C_T(x,1), & \text{if } r = u_1 \end{cases}, \quad \text{and} \quad j_x := \begin{cases} j(x)+1, & \text{if } r = u_0 \\ j(x)-1, & \text{if } r = u_1 \end{cases},$$

where we stress here that in $C_T(x,0)$ and $C_T(x,1)$, we have that T is the tree-in-progress at the point in time that x becomes the first vertex in A.

If $L_{j_x} = L_i$ then let $Y_x = \emptyset$. Otherwise, let

$$Y_x := \{y \in Y \cap L_{j_x} : xy \in E(H) \quad \& \quad c(x,y) \in [d] \setminus (C \cup C_x)\},$$

that is, the set of neighbours of x in $Y \cap L_{j_x}$ that differ from x in a coordinate in $[d] \setminus (C \cup C_x)$. Then, for $y \in Y_x$, we query the edge yx. For this, we reveal the random variable X_{xy}. If $X_{xy} = 1$ then the edge xy belongs to H_p, otherwise it does not. In the case $X_{xy} = 1$, if so far we have identified fewer than $\log d$ edges xy' with $y' \in Y_x$, that is, we have had fewer than $\log d$ random variables $X_{xy'} = 1$, then we move y from Y to the end of A. If we have identified $\log d$ such edges (or if $X_{xy} = 0$), we move to the next edge. Once all the edges xy for every $y \in Y_x$ have been queried, we move x from A to B. The algorithm terminates once A is empty. It then outputs the tree T, rooted at r, and spanned by the edges the algorithm has detected.

We will utilise the following properties of the Tree Construction algorithm.

Proposition 2.2. *The following properties of the Tree Construction algorithm hold.*

(a) If we start with u_0, that is $r = u_0$, then for every $u \in V(T)$ we have that $|C_T(u,0)| \leq i \cdot \log d$. Similarly, if we start with u_1, that is $r = u_1$, then for every $u \in V(T)$ we have that $|C_T(u,1)| \leq (d-i)\log d$.

(b) *If we start with u_0, then for every two leaves $w_1, w_2 \in V(T)$ we have that $I(w_2) \cap C_T(w_1, 0) \neq \varnothing$. Similarly, if we start with u_1, then for every two leaves $w_1, w_2 \in V(T)$ we have that $([d] \setminus I(w_2)) \cap C_T(w_1, 1) \neq \varnothing$.*

(c) *Suppose that we are at the first moment where x is the first vertex in A and let A_x, B_x be the sets A, B at that moment. Let $y \in N_H(x) \cap L_{j_x}$. If $y \in A_x \cup B_x$ then $c(x, y) \in C_x$.*

(d) *Suppose that H has dimension $(1 - o(1))d$, $|C| \leq \frac{d}{2} + \sqrt{d}$, the layer L_i is at distance ℓ from the root, where $\ell = o(\log^5 d), \ell = \omega(1)$, and that $p = \frac{\alpha}{d}$ for some constant $\alpha > e$. Then, with probability at least $(1 + o(1))\zeta_{p(d-|C|)}$, where ζ_c is defined as in (1.1), the following holds. The Tree Construction algorithm outputs a tree T, such that the number of leaves in the layer $L_{\ell-1}$ if we start from u_0, and in the layer $L_{\ell+1}$ if we start from u_1, is between $(p(d - |C|))^{0.9\ell}$ and $(pd)^{1.1\ell}$.*

Proof. For the first item, the claim follows as we truncate at layer L_i, and the number of direct descendants of every $u \in V(T)$ in T, allowed by the Tree Construction algorithm, is at most $\log d$.

For the second item, if we start with u_0, then there exists some common ancestor $w \in V(T)$ of w_1 and w_2, and vertices w_1^-, w_2^-, such that $w < w_1^- \leq w_1$, $w < w_2^- \leq w_2$ and $ww_1^-, ww_2^- \in E(T)$. Let $\ell = c(w, w_2^-)$. Then $\ell \in I(w_2^-) \subseteq I(w_2)$. On the other hand, by definition $\ell \in C_{w_1^-}$. Since for any $s \in A$ we do not allow traversing from s in T along the coordinates of C_s, we have that $\ell \in C_s$ for every descendant s of w_1^- in T, and hence ℓ does not belong to $I(s)$ for any such s. In particular, $\ell \notin I(w_1)$ and $\ell \in C_{w_1}$, and hence $\ell \in C_T(w_1, 0)$. Thus $I(w_2) \cap C_T(w_1, 0) \neq \varnothing$. The statement for when starting with u_1 follows by symmetry.

For the third item, let us consider the case that $r = u_0$, noting that the case $r = u_1$ follows from symmetric arguments. Let $y \in N_H(x) \cap L_{j_x}$ and suppose that $y \in A_x \cup B_x$ at that moment. Let w be the unique vertex on the $x - y$ path spanned by T which satisfies $w < x$ and $w < y$. Furthermore, let x^-, y^- be such that $w < x^- \leq x$, $w < y^- \leq y$, and $wx^-, wy^- \in E(T)$ (see Figure 2). Then $c(w, y^-) \in I(y^-) \subseteq I(y)$. In addition, by definition $c(w, y^-) \in C_{x^-}$. Since for any $s \in A_x$ we do not allow traversing from s in T along the coordinates of C_s, we have that $c(w, y^-) \in C_s$ for every descendant s of x^- in T, and hence $c(w, y^-)$ does not belong to $I(s)$ for any such s. In particular $c(w, y^-) \notin I(x)$. The above implies that $c(x, y) = c(w, y^-) \in C_x$ which completes the proof.

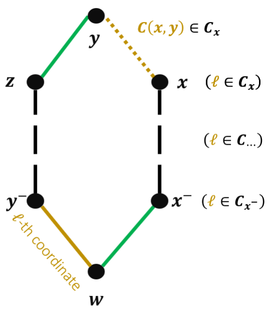

Figure 2: Illustration of the proof of Proposition 2.2(c). Here $c(w, y^-) = \ell$. As for every $s \in A$, we do not traverse from s on the coordinates of C_s in T, we have that ℓ is in C_s for every descendant s of x^-. As both zy and xy are in $E(H)$, along any path from x to y the ℓ-th coordinate must be traversed, and therefore $c(x, y) = \ell \in C_x$.

For the fourth item, for every integer $k \in [0, d]$ and every $x \in L_k$ one has that $|N_{L_{k+1}}(x)| = d - k$ and $|N_{L_{k-1}}(x)| = k$. Thus, suppose this is the first moment where x is the first vertex in A, let d_H be the dimension of H, and let

$$d' = \begin{cases} d_H - i - |C|, & \text{if } v = u_0 \text{ and } x \in \bigcup_{j=0}^{i-2} L_j, \\ d_H - (d_H - i) - |C|, & \text{if } v = u_1 \text{ and } x \in \bigcup_{j=i+2}^{d} L_j, \\ 0 & \text{otherwise.} \end{cases}$$

Then, by Proposition 2.2(c), we have that $|Y_x| \geq d' - |C_x|$. Let Z_x be the number of direct descendants of x in the above process. Then, Z_x stochastically dominates $\min\{\log d, Bin(d' - |C_x|, p)\}$. By Proposition 2.2(a), since L_i is at distance at most $\ell = o(\log^5 d)$ from the root, we have that $|C_x| = o(\log^6 d)$. Since $d_H = (1 + o(1))d$ and $\ell = o(\log^5 d)$ we have that Z_x stochastically dominates

$$Z \sim \min\{\log d, Bin((1 - o(1))(d - |C|), p)\}.$$

By Lemma 2.1, $\mathbb{P}(Bin(d,p) \geq \log d) \leq \left(\frac{\alpha e}{\log d}\right)^{\log d} < \frac{1}{d^5}$. Thus,

$$\mathbb{E}[Z] \geq (1 - o(1))(d - |C|)p - d \cdot \frac{1}{d^5} = (1 - o(1))(d - |C|)p.$$

Since $|C| \leq \frac{d}{2} + \sqrt{d}$ and $\alpha > e$, we have that $\mathbb{E}[Z] > 1$. Furthermore, recall that $\ell = \omega(1)$. Therefore, standard results (see, for example, [11, Page 1] and [9, Chapter 4.3.4]) imply that with probability at least $(1 - o(1))\zeta_{(d-|C|)p}$, the number of leaves after exposing $\ell - 1$ layers (that is, in the layer $L_{\ell-1}$ if we start from u_0, and in the layer $L_{\ell+1}$ if we start from u_1) is at least $(p(d - |C|))^{0.9\ell}$. Similarly, Z_x is stochastically dominated by $Bin(d,p)$, and thus the number of leaves after exposing $\ell - 1$ layers is at most $(pd)^{1.1\ell}$. □

3 Proof of Theorem 1.1(a)

The proof follows from a first-moment argument. Let $X_{(1-\delta)d}$ be the number of increasing paths on $(1-\delta)d$ vertices in Q_p^d. We have at most $\binom{d}{(1-\delta)d}$ ways to choose the coordinates which change along the increasing path and $((1-\delta)d)!$ ways to order them. We then have $2^{\delta d}$ ways to choose the values of the coordinates which are fixed along the path. This determines a path on $(1-\delta)d$ vertices, whose edges appear in Q_p^d with probability $p^{(1-\delta)d-1}$. Therefore,

$$\mathbb{E}\left[X_{(1-\delta)d}\right] = 2^{\delta d}\binom{d}{(1-\delta)d}((1-\delta)d)!p^{(1-\delta)d-1} = 2^{\delta d}\frac{d!}{(\delta d)!}\left(\frac{\alpha}{d}\right)^{(1-\delta)d-1}.$$

By Stirling's approximation and since $\alpha < e$, we have that

$$\mathbb{E}\left[X_{(1-\delta)d}\right] \leq \frac{1}{\alpha\sqrt{\delta}} \cdot d \cdot 2^{\delta d} \cdot \frac{(d/e)^d}{(\delta d/e)^{\delta d}} \cdot \left(\frac{\alpha}{d}\right)^{(1-\delta)d} \leq \frac{1}{\alpha\sqrt{\delta}} \cdot d \cdot \left[\frac{\alpha}{e} \cdot \left(\frac{2}{\delta}\right)^{\delta/(1-\delta)}\right]^{(1-\delta)d}.$$

Since $\lim_{\delta\to 0^+}\left(\frac{2}{\delta}\right)^{\frac{\delta}{(1-\delta)}} = 1$, for any constant $\alpha < e$, we can choose a constant $\delta \in (0,1)$, sufficiently close to 0, such that $\mathbb{E}\left[X_{(1-\delta)d}\right] = o(1)$. Thus, **whp**, there is no increasing path on at least $(1-\delta)d$ vertices.

4 Proof of Theorem 1.1(b)

We focus mainly on showing that $\mathbb{P}\left(\ell(Q_p^d) = d\right) = (1 + o(1))\zeta_\alpha^2$, as the other parts follow with a slight modification of the arguments (which we will argue for at the end of the section). We begin by showing, through a careful second-moment argument, that the existence of a path of length d is not very unlikely (at least inverse-polynomial in d in probability).

To the task at hand, we start with the second-moment argument.

Lemma 4.1. *Let $p = \frac{\alpha}{d}$ with $\alpha \geq e$. Then $\mathbb{P}\left(\ell(Q_p^d) = d\right) \geq \frac{1}{d^5}$.*

Proof. Let X be the random variable counting the number of increasing paths of length d in Q_p^d. By Paley-Zygmund, $\mathbb{P}(X > 0) \geq \frac{\mathbb{E}[X]^2}{\mathbb{E}[X^2]}$. We have that $\mathbb{E}[X] = d!p^d$. Let us turn our attention to $\mathbb{E}[X^2]$. Let Π be the set of all increasing paths of length d in Q^d. Define $id \in \Pi$ by $id = \{v_0v_1, \dots, v_{d-1}v_d\}$, where v_i is the edge vertex whose first i coordinates are one, and whose others are zero. Furthermore, given $\pi_1, \pi_2 \in \Pi$, we stress that $\pi_1 \cap \pi_2, \pi_1 \cup \pi_2$, and $\pi_1 \setminus \pi_2$ are all with respect to the edges of π_1 and π_2. We then have:

$$\mathbb{E}[X^2] = \sum_{\pi_1,\pi_2\in\Pi} \mathbb{P}\left(\pi_1 \in Q_p^d \wedge \pi_2 \in Q_p^d\right) = \sum_{\pi_1,\pi_2\in\Pi} p^{|\pi_1\cup\pi_2|} = \sum_{\pi_1\in\Pi} p^d \cdot \sum_{\pi_2\in\Pi} p^{|\pi_2\setminus\pi_1|}$$
$$= d!p^d \cdot \sum_{\pi\in\Pi} p^{|\pi\setminus id|} = d!p^d \sum_{\pi\in\Pi} p^{d-|\pi\cap id|} = d!p^{2d} \sum_{\pi\in\Pi} p^{-|\pi\cap id|}. \tag{4.1}$$

Now,

$$\sum_{\pi\in\Pi} p^{-|\pi\cap id|} = \sum_{k=0}^{d} \left(\frac{d}{\alpha}\right)^k \cdot Y_k, \tag{4.2}$$

where Y_k is the number of increasing paths which intersect with id on exactly k edges.

Let us now estimate Y_k. When $k = d$, we have that $Y_k = 1$. Suppose that $k < d$. Assume $\pi \in \Pi$ intersects with id on exactly k edges. Let $\ell \geq 1$ be the number of maximal segments of the path id that are edge-disjoint from π. Observe that there are at most $\binom{d}{2\ell}$ ways to choose where these segments lie. Indeed, in order to determine these segments it is enough to choose the first and last edges of each of the segments out of the d edges in the path id. Here, we used that each of these segments contains at least 2 edges. Let us denote the number of edges in each of these segments by $x_1, \dots, x_\ell$, where $\sum_{i=1}^{\ell} x_i = d - k$ (note that when choosing where the segments lie, we determined $x_1, \dots, x_\ell$). As noted above, we have that $x_i \geq 2$ for $1 \leq i \leq \ell$ and therefore $1 \leq \ell \leq \frac{d-k}{2}$. There are at most $x_i!$ ways to form each segment. Therefore, for $k < d$,

$$Y_k \leq \sum_{\ell=1}^{\frac{d-k}{2}} \binom{d}{2\ell} \max_{\substack{x_1,\dots,x_\ell\geq 2\\ \sum_{i=1}^{\ell} x_i = d-k}} \left\{\prod_{i=1}^{\ell} x_i!\right\}.$$

We claim that in this range of variables $\prod_{i=1}^{\ell} x_i!$ is maximised when $x_i = 2$ for all $i \in [\ell]$ except for one of them which is equal to $d - k - 2(\ell - 1)$. Indeed, for any sequence $x_1, \dots, x_\ell$, if there are $i \neq j$ for which $x_i \geq x_j > 2$ then one can replace x_i, x_j with $x_i + 1$ and $x_j - 1$ and increase the product $\prod_{i=1}^{\ell} x_i!$ (since the ratio between these products is $(x_i + 1)!(x_j - 1)!/(x_i!x_j!) = (x_i + 1)/x_j > 1$). Hence, for $k < d$, $Y_k \leq \sum_{\ell=1}^{\frac{d-k}{2}} \binom{d}{2\ell} 2^{\ell-1}(d - k - 2\ell + 2)!$.

Returning to (4.2), we have that

$$\sum_{k=0}^{d} \left(\frac{d}{\alpha}\right)^k \cdot Y_k \leq d! + \sum_{\ell=1}^{\frac{d}{2}} 2^\ell \binom{d}{2\ell} \sum_{k=0}^{d-2\ell} \left(\frac{d}{\alpha}\right)^k (d - k - 2\ell + 2)!\,, \tag{4.3}$$

where the first summand, $d!$, corresponds to the case where $k = d$ — indeed, we use the fact that $\left(\frac{d}{\alpha}\right)^d \leq \left(\frac{d}{e}\right)^d \leq d!$ and that $Y_k = 1$. Now,

$$\sum_{k=0}^{d-2\ell} \left(\frac{d}{\alpha}\right)^k (d-k-2\ell+2)! = \left(\frac{d}{\alpha}\right)^{d-2\ell} \sum_{m=0}^{d-2\ell} \left(\frac{\alpha}{d}\right)^m (m+2)! \leq \left(\frac{d}{\alpha}\right)^{d-2\ell} d^2 \sum_{m=0}^{d-2\ell} \left(\frac{\alpha}{d}\right)^m m!$$
$$\leq \left(\frac{d}{\alpha}\right)^{d-2\ell} d^3 \sum_{m=0}^{d-2\ell} \left(\frac{\alpha}{d}\right)^m \left(\frac{m}{e}\right)^m \leq \left(\frac{d}{\alpha}\right)^{d-2\ell} d^4 \cdot \left(\frac{\alpha}{e}\right)^{d-2\ell} \leq \left(\frac{d}{e}\right)^{d-2\ell} d^4,$$

where in the first inequality we used the fact that $\ell \geq 1$ and in the penultimate inequality we used our assumption $\alpha \geq e$ and the fact that $m \leq d$. Returning to (4.3), we now have

$$\sum_{k=0}^{d}\left(\frac{d}{\alpha}\right)^k \cdot Y_k \leq d! + \sum_{\ell=1}^{\frac{d}{2}} 2^\ell \binom{d}{2\ell}\left(\frac{d}{e}\right)^{d-2\ell} d^4 \leq d! + \left(\frac{d}{e}\right)^d d^4 \sum_{\ell=1}^{\frac{d}{2}}\left(\frac{ed}{\ell}\cdot\frac{e}{d}\right)^{2\ell} \leq d!\cdot d^5.$$

Substituting the above in (4.1), we obtain that $\mathbb{E}[X^2] \leq \left(p^d d!\right)^2 d^5$. By Paley-Zygmund, and the above inequalities, $\mathbb{P}\left(X > 0\right) \geq \frac{1}{d^5}$. □

We note that in the proof above, we did not attempt to optimise the exponent in d^{-5}, and instead aimed for simplicity and clarity — indeed, with more careful calculations one could obtain a better exponent, however, that does not affect the rest of the proof.

We now turn our attention to showing that if **0** and **1** both 'percolate', then **whp** we can find d^{10} pairwise disjoint subcubes, $Q(v_{1,0}, v_{1,1}), \ldots, Q(v_{d^{10},0}, v_{d^{10},1})$, of dimension $(1-o(1))d$, such that there is an increasing path in Q_p^d between **0** and every $v_{i,0}$, and a decreasing path between **1** and every $v_{i,1}$. Showing the existence of such subcubes requires a delicate construction. Throughout the rest of this section, we assume that $p = \frac{\alpha}{d}$ for some $\alpha > e$, and recall that ζ_α is defined according to (1.1).

We begin by showing that with probability at least $(1-o(1))\zeta_\alpha^2$, both **0** and **1** 'percolate' in Q_p^d and we can find a tree rooted at **1** which is 'above' the tree rooted at **0**. Formally,

Lemma 4.2. *With probability at least $(1-o(1))\zeta_\alpha^2$, the following holds. There exist two trees $T_0', T_1' \in Q_p^d$ such that T_0' and T_1' have height $2\log\log d$, and:*

1. ***0** $\in V(T_0')$, **1** $\in V(T_1')$; and,*
2. *$|L_{2\log\log d} \cap V(T_0')| \geq \log d$, $|L_{d-2\log\log d} \cap V(T_1')| \geq \log d$; and,*
3. *$|I(L_{2\log\log d} \cap V(T_0'))| \leq \log^{2\alpha} d$, $|[d] \setminus I(L_{d-2\log\log d} \cap V(T_1'))| \leq \log^{2\alpha} d$; and,*
4. *for every two leaves of T_0', $w_{1,0}, w_{2,0} \in V(T_0')$ with $w_{1,0} \neq w_{2,0}$, we have that $I(w_2) \cap C_{T_0'}(w_1, 0) \neq \varnothing$; and,*
5. *for every two leaves of T_1', $w_{1,1}, w_{2,1}$ with $w_{1,1} \neq w_{2,1}$, we have that $([d] \setminus I(w_2)) \cap C_{T_1'}(w_1, 1) \neq \varnothing$; and,*
6. *for every $v_0 \in V(T_0')$, we have that $|C_{T_0'}(v_0, 0)| \leq \log^2 d$, and for every $v_1 \in V(T_1')$ we have that $|C_{T_1'}(v_1, 1)| \leq \log^2 d$; and,*
7. *for every $v_0 \in V(T_0')$ and $v_1 \in V(T_1')$, $v_0 < v_1$.*

Proof. Run the Tree Construction algorithm described in Section 2.1 with the following inputs. Let (In1), the subcube, be Q^d, let (In2), the root, be **0**, let (In3), the set of coordinates which we avoid, be $\varnothing$, let (In4), the layer at which we truncate, be $L_{2\log\log d+1}$, and let (In5) be a sequence of independent Bernoulli(p) random variables. Let T_0' be the tree this Tree Construction algorithm outputs.

Item 4 follows deterministically from Proposition 2.2(b), and item 6, with respect to T_0', follows deterministically from Proposition 2.2(a). Note the dimension of the subcube in the algorithm's input is d, that we do not avoid any coordinates, and that the distance of the layer we truncate at from the root is $2\log\log d + 1$. Hence, we may apply Proposition 2.2(d). As $pd = \alpha$ and $\alpha^{0.9\cdot 2\log\log d} \geq \log d$ and $\alpha^{1.1\cdot 2\log\log d} \leq \log^{2\alpha} d$, the part of items 1, 2, and 3 which concerns T_0', follows from Proposition 2.2(d). Let us denote the event that such T_0' exists by $\mathcal{A}_0$, where we have that $\mathbb{P}(\mathcal{A}_0) \geq (1-o(1))\zeta_\alpha$ by 2.2(d).

Let $I_0 := I\left(V(T_0) \cap L_{2\log\log d}\right)$. Conditioned on $\mathcal{A}_0$, we have that $|I_0| \leq \log^{2\alpha} d$. We now run the Tree Construction algorithm described in Section 2.1 with the following inputs. Let (In1), the subcube, be Q^d, let (In2), the root, be **1**, let (In3), the set of coordinates which we avoid, be I_0, let (In4), the layer at which we truncate, be $L_{d-2\log\log d-1}$, and let (In5) be a sequence of independent Bernoulli(p) random variables. Let T_1' be the tree this Tree Construction algorithm outputs.

Note that item 7 holds by construction. Item 5 follows deterministically from Proposition 2.2(b), and item 6, with respect to T_1', follows deterministically from Proposition 2.2(a). Note that the dimension of the subcube in the algorithm's input is d, that $|I_0| \leq \frac{d}{2} + \sqrt{d}$, and that the distance of the layer we truncate at from the root is $2\log\log d + 1$. Thus, we may apply Proposition 2.2(d). As $(d - |I_0|)p = (1 + o(1))\alpha$, the part of items 1, 2, and 3 which concerns T_1', follow from Proposition 2.2(d). Let us denote the event that such T_1' exists by $\mathcal{A}_1$, and we have that, conditional on $\mathcal{A}_0$, the probability that $\mathcal{A}_1$ holds is at least $(1 - o(1))\zeta_{(1-o(1))\alpha} = (1 - o(1))\zeta_\alpha$ by 2.2(d).

Furthermore, the sets of edges explored during the construction of T_0' and that during the construction of T_1' are disjoint, and therefore the probability of both events $\mathcal{A}_0$ and $\mathcal{A}_1$ holding is at least $(1 - o(1))\zeta_\alpha^2$. □

We now turn to show that **whp** we can extend T_0' to a tree of height $O(\log d)$, with, say, d^{11} leaves, such that the leaves are easily distinguishable. That is, our goal is that every leaf of T_0' will be uniquely identified with a set of coordinates, whose order is, say, $150\log^2 d$. We also aim to extend T_1' similarly, maintaining all of its vertices above T_0'. We do so, roughly, by growing T_0' on the first $d/2$ coordinates, and T_1' on the last $d/2$ coordinates, utilising Proposition 2.2(d). More precisely,

Lemma 4.3. *Suppose T_0' and T_1' satisfy the properties in the statement of Lemma 4.2. Then, **whp**, there exist trees $T_0, T_1 \in Q_p^d$ such that T_0 and T_1 have height $150\log d$, and:*

1. *$\mathbf{0} \in V(T_0)$ and $\mathbf{1} \in V(T_1)$; and,*
2. *$|L_{150\log d} \cap V(T_0)| \geq d^{11}$ and $|L_{d-150\log d} \cap V(T_1)| \geq d^{11}$; and,*
3. *for every $v \in L_{150\log d} \cap V(T_0)$, we have that $|C_{T_0}(v, 0)| \leq 150\log^2 d$; and,*
4. *for every $v \in L_{d-150\log d} \cap V(T_1)$, we have that $|C_{T_1}(v, 1)| \leq 150\log^2 d$; and,*
5. *for every $u, v \in L_{150\log d} \cap V(T_0)$ with $u \neq v$, we have that $I(u) \cap C_{T_0}(v, 0) \neq \varnothing$; and,*
6. *for every $u, v \in L_{d-150\log d} \cap V(T_1)$ with $u \neq v$ we have that $C_{T_1}(v, 1) \setminus I(u) \neq \varnothing$; and,*
7. *for every $v_0 \in V(T_0)$ and $v_1 \in V(T_1)$, $v_0 < v_1$.*

Proof. Note that when constructing T_0' and T_1', we only considered edges up to the $2\log\log d$-th, and $(d - 2\log\log d)$-th layer, respectively. We may thus assume that edges crossing the other layers have not been exposed yet.

We begin by showing that **whp** T_0 exists. Let $I_1 := I\left(V(T_1') \cap L_{d-2\log\log d}\right)$, where we note that by assumption $|I_1| \geq d - \log^{2\alpha} d$. Furthermore, since T_0' satisfies the properties of Lemma 4.2, we have that $\mathbf{0} \in V(T_0')$ and $|L_{2\log\log d} \cap V(T_0')| \geq \log d$. Let $U_0 = \{u_1, \ldots, u_{\log d}\}$ be a set of arbitrary $\log d$ vertices from $L_{2\log\log d} \cap V(T_0')$. For every $i \in [\log d]$, let w_i be defined by $I(w_i) = [d] \setminus C_{T_0'}(u_i, 0)$. Since $I(u_i) \cap C_{T_0'}(u_i, 0) = \varnothing$, we have that $w_i > u_i$. By our assumption, for every two leaves $u_i \neq u_j \in V(T_0') \cap L_{2\log\log d}$, we have that $I(u_i) \cap C_{T_0'}(u_j, 0) \neq \varnothing$. Furthermore, by our assumption, for every $i \in [\log d]$, we have that $|C_{T_0'}(u_i, 0)| \leq \log^2 d$. Thus, $Q(1) := Q(u_1, w_1), \ldots, Q(\log d) := Q(u_{\log d}, w_{\log d})$ form a set of $\log d$ pairwise disjoint subcubes of dimension at least $d - \log d - \log^2 d \geq d - 2\log^2 d$. We claim that with probability bounded away from zero, there exists an $i \in [\log d]$ such that in $Q(i)_p$ there exists a tree B_i, such that B_i together with the path from $\mathbf{0}$ to u_i in Q_p^d form a suitable choice for T_0.

To that end, for every $i \in [\log d]$, we run the Tree Construction algorithm given in Section 2.1 with the following inputs. Let (In1), the subcube, be $Q(i)$, let (In2), the root, be u_i, let (In3), the set of coordinates which we avoid, be $([d] \setminus I_1) \cup ([d] \setminus [d/2])$, let (In4), the layer at which we truncate, be $L_{150\log d}$, and let (In5), be a sequence of independent Bernoulli(p) random variables. Let B_i be the tree this algorithm outputs.

Note that the subcube in the algorithm's input is of dimension $(1 - o(1))d$, that the set of coordinates we avoid, $([d] \setminus I_1) \cup ([d] \setminus [d/2])$, is of size at most $\log^{2\alpha} d + \frac{d}{2} \leq \frac{d}{2} + \sqrt{d}$, and that the distance of the layer we truncate at from the root is $150\log d - 2\log\log d$. Thus, we may apply Proposition 2.2(d), and obtain that with probability at least c, for some constant $c > 0$, $|L_{150\log d} \cap V(B_i)| \geq \left(\frac{2e}{5}\right)^{135\log d} \geq d^{11}$.

In the event $|L_{150\log d} \cap V(B_i)| \geq d^{11}$, note that letting T_0 be B_i with the path from $\mathbf{0}$ to u_i in Q_p^d, for every $v \in V(B_i)$, we have that $C_{T_0}(v, 0) = C_{B_i}(v, 0)$, and thus by Proposition 2.2(a), $|C_{T_0}(v, 0)| \leq 150\log^2 d$. Since we restrict the queries to $[d/2] \cap I_1$, for every $v_0 \in V(T_0)$, we have that $I(v_0) \subseteq I_1$ and $I(v_0) \subseteq [d/2] \cup I(V(T_0'))$, where the first further implies that for every $v_1 \in V(T_1')$, we have $v_0 < v_1$. Finally, we need to verify that for every $u, v \in L_{150\log d} \cap V(T_0)$ with $u \neq v$, we have that $I(u) \cap C_{T_0}(v, 0) \neq \varnothing$ — indeed, this follows from Proposition 2.2(b). Therefore, with probability at least $c > 0$, the path from $\mathbf{0}$ to u_i in Q_p^d appended with B_i forms a suitable choice for T_0.

Since these events are independent for every $i \in [\log d]$, the probability there is no such T_0 is at most $(1-c)^{\log d} = o(1)$.

The existence of T_1 follows similarly, where we note that $I(V(T_0)) \subseteq I(V(T_0')) \cup [d/2]$ and $|I(V(T_0'))| \leq \log^{2\alpha} d = o(d)$. □

We are now ready to show that **whp** we can extend T_1, so that we can find d^{10} pairwise disjoint subcubes, of dimension $(1-o(1))d$, with their antipodal points connected by increasing paths to $\mathbf{0}$ and $\mathbf{1}$, respectively.

Lemma 4.4. *Suppose T_0 and T_1 satisfy the properties of Lemma 4.3. Then, **whp**, we can find subsets $V_0 = \{v_{1,0}, \ldots, v_{d^{10},0}\} \subseteq V(T_0) \cap L_{150\log d}$ and $V_1 = \{v_{1,1}, \ldots, v_{d^{10},1}\}$, with V_1 being at layer at least $d - 15^3\log^6 d$, such that the following holds.*

1. *For every $v_0 \in V_0$ and $v_1 \in V_1$, there is a path in Q_p^d between $\mathbf{0}$ and v_0 and between $\mathbf{1}$ and v_1.*
2. *for every $i \in [d^{10}]$, we have that $v_{i,0} < v_{i,1}$; and,*
3. *for every $i \in [d^{10}]$, we have that $C_{T_0}(v_{i,0}, 0) \cap I(v_{i,1}) = \varnothing$.*

Proof. Once again, we stress that we have not queried any of the edges between the $150\log d$-th and the $d - 150\log d$-th layers thus far.

Let $V_0 = \{v_{1,0}, \ldots, v_{d^{10},0}\}$ be an arbitrary set of d^{10} vertices in $V(T_0) \cap L_{150\log d}$. Note that $|I(v_{i,0})| = 150\log d$ for every $i \in [d^{10}]$.

Let $M = \{m_1, \ldots, m_{d^{11}}\}$ be an arbitrary set of d^{11} vertices in $L_{d-150\log d} \cap V(T_1)$. We arbitrarily split them to sets $S_1, \ldots, S_{d^{10}}$, each of order d. For every $i \in [d^{10}]$, we associate S_i with $v_{i,0}$. We now turn to show that **whp**, for every $i \in [d^{10}]$ there exists a tree $T(S_i)$ in Q_p^d, rooted at some vertex in S_i, such that at least one of its vertices $y \in V(T(S_i)) \cap \bigcup_{s=d-15^3\log^6 d}^{d} L_s$, satisfies that $v_{i,0} < y$ and $I(y) \cap C_{T_0}(v_{i,0}) = \varnothing$.

Fix $i \in [d^{10}]$. Denote the vertices of S_i by $\{u_1, \ldots, u_d\}$. For every $k \in [d]$, let w_k be the vertex in Q^d defined by $I(w_k) = C_{T_1}(u_k, 1)$. Similar to the proof of Lemma 4.3, we have that $Q(1) \coloneqq Q(w_1, u_1), \ldots, Q(d) \coloneqq Q(w_d, u_d)$ form d pairwise vertex disjoint subcubes. By our assumption, $|C_{T_1'}(u_k, 1)| \leq 150\log^2 d$, and thus each of these subcubes is of dimension at least $d - 200\log^2 d$. Fix $k \in [d]$ and consider $Q(k)$. Run the Tree Construction algorithm given in Section 2.1 with the following inputs. Let (In1), the subcube, be $Q(k)$, let (In2), the root, be u_k, let (In3), the set of coordinates which we avoid, be $I(v_{i,0})$, let (In4), the layer at which we truncate, be $L_{d-10\log^4 d-1}$, and let (In5) be a sequence of independent Bernoulli(p) random variables. Let B_k be the tree this algorithm outputs. Similarly to previous arguments, by Proposition 2.2(d), since $(d - 20\log^4 d)p \geq \frac{\alpha}{2} > \frac{13}{10}$, with probability at least c, for some positive constant c, we have $|V_{d-\log^4 d} \cap V(B_k)| \geq \left(\frac{13}{10}\right)^{9\log^4 d}$. Note that every vertex in B_k is above $v_{i,0}$. Thus, with probability at least $1 - (1-c)^d = 1 - o\left(d^{-11}\right)$, there exists a set of $\left(\frac{13}{10}\right)^{9\log^4 d}$ vertices in layer $L_{d-10\log^4 d}$, all of which are connected to $\mathbf{1}$ via an increasing path in Q_p^d and all of which are above $v_{i,0}$. Denote this set of vertices by X.

Initialise $X(0) \coloneqq X$ and $C(0) \coloneqq C_{T_0}(v_{i,0})$. Now, at each iteration $j \in [|C(0)|]$, we proceed as follows. Let ℓ be the first (smallest) coordinate in $C(j-1)$. If at least $\frac{|X(j-1)|}{d}$ vertices $x \in X(j-1)$ have that $\ell \in [d] \setminus I(x)$, we set $X(j)$ to be these set of vertices and update $C(j) = C(j-1) \setminus \{\ell\}$. Otherwise, there are at least $\left(1-\frac{1}{d}\right)|X(j-1)|$ vertices $x \in X(j-1)$, such that $\ell \in I(x)$. Each such vertex x has exactly one neighbour x' in the layer below, such that $\ell \in [d] \setminus I(x')$. The probability that xx' is in Q_p^d is p, and these are independent trials for every vertex x. Thus, the number of x' such that $\ell \in [d] \setminus I(x')$ and that there is $x \in X(j-1)$ such that $xx' \in E(Q_p^d)$ stochastically dominates $Bin\left(\left(1-\frac{1}{d}\right)|X(j-1)|, p\right)$. By Lemma 2.1, with probability at least $1-\exp\left(-\frac{|X(j-1)|}{10d}\right)$, we have that the number of such x' is at least $\frac{|X(j-1)|}{2d}$, where we stress that if $|X(j-1)| \geq d^2$, this holds with probability at least $1-o(d^{-11})$. We then let $X(j)$ be the set of such x', and update $C(j) = C(j-1) \setminus \{\ell\}$. Repeating the above process for $|C(0)| \leq 150\log^2 d$ iterations, by the union bound we have that with probability at least $1-o(\log^2 d \cdot d^{-11})$ at each iteration, $|X(j)| \geq \frac{|X(j-1)|}{2d}$, and in particular, $|X(|C(0)|)| \geq \frac{\left(\frac{13}{10}\right)^{9\log^4 d}}{(2d)^{150\log^2 d}} \geq d^2$. Thus, with probability at least $1-o(d^{-10})$ there exists a proper choice of $v_{i,1}$ at layer at least $d - 10\log^4 d \cdot 150\log^2 d \geq d - 15^3\log^6 d$. Union bound over the d^{10} choices of i completes the proof. □

We are now ready to prove Theorem 1.1(b).

Proof of Theorem 1.1(b). By Lemmas 4.2 through 4.4, with probability at least $(1-o(1))\zeta_\alpha^2 - o(1) = (1-o(1))\zeta_\alpha^2$, we can find subsets $V_0 = \{v_{1,0}, \ldots, v_{d^{10},0}\} \subseteq V(T_0') \cap L_{150\log d}$ and $V_1 = \{v_{1,1}, \ldots, v_{d^{10},1}\}$, with V_1 being at layer at least $d - 15^3\log^6 d$, such that the following holds.

1. For every $v_0 \in V_0$ and $v_1 \in V_1$, we have that there is a path in Q_p^d between $\mathbf{0}$ and v_0 and between $\mathbf{1}$ and v_1.
2. for every $i \in [d^{10}]$, we have that $v_{i,0} < v_{i,1}$; and,
3. for every $i, j \in [d^{10}]$ with $i \neq j$, we have that $C_{T_0}(v_{i,0}, 0) \cap I(v_{i,1}) = \varnothing$.

By our construction of T_0, for every $u, v \in V(T_0)$ we have that $I(u) \cap C_{T_0}(v, 0) \neq \varnothing$. This implies that $I(u) \setminus I(v) \neq \varnothing$. Therefore, for every $i, j \in [d^{10}]$ with $i \neq j$, we have that the subcubes $Q(v_{i,0}, v_{i,1})$ and $Q(v_{j,0}, v_{j,1})$ are disjoint.

For every $i \in [d^{10}]$, let $H(i) = Q(v_{i,0}, v_{i,1})$. By Lemma 4.1, with probability at least $\frac{1}{d^5}$, there is a path in $H(i)_p$ between $v_{i,0}$ and $v_{i,1}$. Thus, the probability that there is no path between $\mathbf{0}$ and $\mathbf{1}$ in Q_p^d is at most $1-(1-o(1))\zeta_\alpha^2 + \left(1-\frac{1}{d^5}\right)^{d^{10}} = 1-(1-o(1))\zeta_\alpha^2$. Standard results (see, for example, [9]) implies that there does not exist a path in Q_p^d from $\mathbf{0}$ to any vertex in $L_{\log d}$ with probability at least $(1+o(1))(1-\zeta_\alpha)$. Similarly, the probability there is no path from $\mathbf{1}$ to any vertex in $L_{d-\log d}$ is at least $(1+o(1))(1-\zeta_\alpha)$. Noting that these two events are independent, and since the function $f(x) = x^2 + 2x(1-x)$ is increasing in the interval $(0,1)$, the probability that at least one such path does not exist is at least $(1-\zeta_\alpha)^2 + 2\zeta_\alpha(1-\zeta_\alpha) - o(1) = (1+o(1))\left(1-\zeta_\alpha^2\right)$. Therefore, $\mathbb{P}\left(\ell(Q_p^d) = d\right) = (1+o(1))\zeta_\alpha^2$.

The statements for $d-1, d-2$ follow similarly. For $d-1$, in Lemma 4.2 we can first grow T_0'. Conditioned on the event that the vertex $\mathbf{0}$ 'percolates' we have that **whp** there exists a tree T_0' that satisfies the parts of items 1, 2, and 3 that concern T_0' in the statement in Lemma 4.2, call this event $\mathcal{A}_0$. Then we continue and grow T_1' using the Tree Construction algorithm with the following adjustment. To the first level of T_1', in addition to the neighbours of $\mathbf{1}$ in Q_p^d, we add $\log d$ random vertices that lie above T_0'. It follows that with this small adjustment, conditioned on $\mathcal{A}_0$, the statement of Lemma 4.2 is satisfied **whp**. The rest of the arguments follow in an identical manner. We then obtain **whp** an increasing path of length d that may end at one of the $\log d$ added edges. This corresponds to the existence, **whp**, of an increasing path of length $d-1$ in Q_p^d. By symmetry, we get that conditioned on at least one of the vertices **0** or **1** percolating, $\ell(Q_p^d) \geq d-1$ **whp**. Thus $\ell(Q_p^d) \geq d-1$ with probability $1-(1+o(1))(1-\zeta_a)^2$. As $\Pr(\ell(Q_p^d) = d) = (1+o(1))\zeta_a^2$ we have that $\Pr(\ell(Q_p^d) = d-1) = (1+o(1))2\zeta_a(1-\zeta_a)$. As for $d-2$, we can grow both T_0' and T_1' by adding to each of them, at their first level, $\log d$ extra random vertices. This gives that **whp** Q_p^d contains an increasing path of length at least $d-2$. In extension, the probability that a longest increasing path has length at least $d-2$ is $(1+o(1))(1-\zeta_a)^2$. □

References

[1] M. Ajtai, J. Komlós, and E. Szemerédi. Largest random component of a k-cube. *Combinatorica*, 2(1):1–7, 1982.

[2] N. Alon and J. H. Spencer. *The probabilistic method*. Hoboken, NJ: John Wiley & Sons, fourth edition, 2016.

[3] A. Blanc-Renaudie, N. Broutin, and A. Nachmias. The scaling limit of critical hypercube percolation. *arXiv preprint arXiv:2401.16365*, 2024.

[4] B. Bollobás. The evolution of the cube. In *Combinatorial mathematics (Marseille-Luminy, 1981)*, volume 75 of *North-Holland Math. Stud.*, pages 91–97. North-Holland, Amsterdam, 1983.

[5] B. Bollobás. Complete matchings in random subgraphs of the cube. *Random Structures Algorithms*, 1(1):95–104, 1990.

[6] B. Bollobás, Y. Kohayakawa, and T. Łuczak. The evolution of random subgraphs of the cube. *Random Structures Algorithms*, 3(1):55–90, 1992.

[7] Ju. D. Burtin. The probability of connectedness of a random subgraph of an n-dimensional cube. *Problemy Peredači Informacii*, 13(2):90–95, 1977.

[8] P. Condon, A. Espuny Díaz, A. Girão, D. Kühn, and D. Osthus. Hamiltonicity of random subgraphs of the hypercube. *Mem. Amer. Math. Soc.*, to appear.

[9] R. Durrett. *Probability: theory and examples*. Cambridge University Press, Cambridge, 2019.

[10] P. Erdős and J. Spencer. Evolution of the n-cube. *Comput. Math. Appl.*, 5(1):33–39, 1979.

[11] R. Lyons, R. Pemantle, and Y. Peres. Conceptual proofs of $L \log L$ criteria for mean behavior of branching processes. *Ann. Probab.*, 23(3):1125–1138, 1995.

[12] R. Pinsky. Detecting tampering in a random hypercube. *Electron. J. Probab.*, 18:no. 28, 12, 2013.

[13] A. A. Sapoženko. Metric properties of almost all functions of the algebra of logic. *Diskret. Analiz*, 10:91–119, 1967.

Acknowledgments. The authors wish to thank Ross Pinsky for his comments on an earlier version of the paper, and for bringing reference [12] to our attention. The authors are grateful to the anonymous referees for their helpful comments and suggestions.