Cycle lengths in the percolated hypercube

Michael Anastos [*] Sahar Diskin [†] Joshua Erde [‡] Mihyun Kang [‡]
Michael Krivelevich [†] Lyuben Lichev [§]

Abstract

Let Q^d_p be the random subgraph of the d-dimensional binary hypercube obtained after edge-percolation with probability p. It was shown recently by the authors that, for every $\varepsilon > 0$, there is some $c = c(\varepsilon) > 0$ such that, if $pd \geq c$, then typically Q^d_p contains a cycle of length at least $(1-\varepsilon)2^d$. We strengthen this result to show that, under the same assumptions, typically Q^d_p contains cycles of all even lengths between 4 and $(1-\varepsilon)2^d$.

1 Introduction

The study of *Hamiltonicity* of graphs is a central topic in probabilistic combinatorics. Whilst in general the problem of determining whether a graph is Hamiltonian is known to be computationally hard [28], classic structural criteria for the existence of Hamilton cycles in dense graphs due to Dirac [16], Ore [35], Chvátal [13], and Bondy and Chvátal [12] have been known for the last half a century. Yet, understanding the property in satisfactory generality is highly non-trivial and has given birth to some fascinating mathematical advances. Pósa [36] pioneered the study of Hamiltonicity in random graphs, developing the *rotation-extension method* and using it to show that $G(n,p)$, the random graph on n vertices where each edge appears independently with probability p, is typically Hamiltonian when $np = C\log n$ for a suitably large constant C. Variations of this method have proved instrumental in a number of works on the subject. Pósa's result was further improved by Korshunov [30], Bollobás [9], and Komlós and Szemerédi [29]. In the sparser regime, Ajtai, Komlós, and Szemerédi [1] showed that, for any constant $c > 1$, $G(n,c/n)$ typically contains a cycle of length $(1-o_c(1))n$ (see also the independently obtained result of Fernandez de la Vega [22]).

A natural research direction generalising the advances on Hamiltonicity is the study of *pancyclicity*. Denoting by $\mathcal{L}(H)$ the set of cycle lengths in a given graph H, also known as the *cycle spectrum* of H, we say that H is pancyclic if $\mathcal{L}(H) = \{3,\ldots,|V(H)|\}$ that is, if H contains cycles of all possible lengths. In the setting of $G(n,p)$, Cooper and Frieze [15] established a hitting time result for pancyclicity in the random graph process, showing that minimum degree two is typically both a necessary and sufficient condition. In the sparser setting, Alon, the fifth author, and Lubetzky [4] showed that, for sufficiently large $c > 0$ and for any $\ell = \ell(n)$ tending arbitrary slowly to infinity, typically $\mathcal{L}(G(n,c/n)) \supseteq \{\ell, \ell+1, \ldots, (1-o(1))L_{max}\}$, where L_{max} is the length of a longest cycle in $G(n,c/n)$. This was further improved by the first author [5] who calculated the probability that, for $c > 0$ sufficiently large, $G(n,c/n)$ is *weakly pancyclic*, that is, contains cycles of all lengths between 3 and the length of a longest cycle.

The binomial random graph $G(n,p)$ is perhaps the simplest example of a model of *p-bond percolation* where every edge of a host graph G is retained independently with probability p, thus producing a random subgraph G_p. Thus, $G(n,p)$ can be seen as p-bond percolation on an K_n, the complete graph on n vertices. Studying foundational topics like Hamiltonicity and pancyclicity in other models involving bond percolation is a natural and inviting question.

[*]Institute of Science and Technology Austria (ISTA), Klosterneurburg 3400, Austria. Email: michael.anastos@ist.ac.at.
[†]School of Mathematical Sciences, Tel Aviv University, Tel Aviv 6997801, Israel. Emails: sahardiskin@mail.tau.ac.il, krivelev@tauex.tau.ac.il.
[‡]Institute of Discrete Mathematics, Graz University of Technology, Steyrergasse 30, 8010 Graz, Austria. Emails: erde@math.tugraz.at, kang@math.tugraz.at.
[§]Institute of Statistics and Mathematical Methods in Economics, Technical University of Vienna, A-1040 Vienna, Austria. Email: lyuben.lichev@tuwien.ac.at.

Perhaps the second most studied host graph after the complete graph is the high-dimensional binary hypercube. For an integer $d \geq 1$, the *d-dimensional hypercube* Q^d is the graph with vertex set $\{0,1\}^d$ where every two vertices differing in a single coordinate form an edge. The non-trivial lattice-like geometry of the hypercube Q^d has proved a significant obstacle for adapting the methods used in the analysis of $G(n,p)$ to the percolated hypercube Q^d_p. For example, a phase transition for the existence of a giant component in percolated hypercubes was established by Ajtai, Komlós and Szemerédi [2], see also Bollobás, Kohayakawa and Łuczak [11], more than two decades after the original work of Erdős and Rényi on $G(n,p)$ [21]. More than that, thresholds for the existence of long cycles and of Hamiltonian cycles in Q^d_p were resolved only recently [7, 14]. In a breakthrough, Condon, Espuny Díaz, Girão, Kühn and Osthus [14] established a sharp threshold for Hamiltonicity in Q^d_p at $p = 1/2$. In fact, they even showed a stronger *hitting time* result, and with minor modifcations of their approach, one can obtain a hitting time result for *even pancyclicity* as well. In the (much) sparser regime, the authors [7] showed that, for large enough $c > 0$, **whp**[1] $Q^d_{c/d}$ contains a cycle of length $(1 - o_c(1))2^d$.

In this paper, we consider the cycle spectrum of Q^d_p. Note that, since the hypercube is bipartite, Q^d (and Q^d_p) does not contain any cycles of odd length. Our main result is an analogue of the result of [4] in the setting of percolated hypercubes.

Theorem 1. *For every $\varepsilon \in (0,1)$, there exists a constant $c(\varepsilon) > 0$ such that, for all $c \geq c(\varepsilon)$, with $p = p(d) = c/d$,* ***whp*** *Q^d_p simultaneously contains all cycles of even length between 4 and $(1-\varepsilon)2^d$.*

We note that the dependency of c in ε which we obtain is inverse polynomial. Let us further note that Bollobás (see [10], §4.1), and separately Karoński and Ruciński [27] proved that for fixed ℓ the number of cycles of length ℓ in $G(n,c/n)$ converges to Poisson$(c^\ell/(2\ell))$, hence $G(n,c/n)$ does not contain a cycle of length ℓ with probability uniformly bounded away from 0. On the other hand, a second moment argument shows that typically $Q^d_{c/d}$ does contain small cycles of a given even length.

Our work leaves some natural open problems. It would be very interesting to determine if, analogous to the case of $G(n,p)$, when pd is sufficiently large, Q^d_p is typically *weakly even-pancyclic*, that is, it contains cycles of all even lengths up to L_{max}. Settling a well-known conjecture of the fifth author and Sudakov [33], a recent work by Draganić, Montgomery, Munhá Correia, Pokrovskiy and Sudakov [20] showed that *C-expanders* are Hamiltonian. Results for cycle lengths have also been shown for *pseudorandom graphs* [23, 26] (which are, in particular, C-expanders). It would be interesting to know if these results are *robust* enough to still hold after percolation, as is known to be the case for the existence of long cycles, see [19]. The fifth author, Lee and Sudakov [32] studied the *local resilience* of pancyclicity in $G(n,p)$. Lee and Samotij [34] considered the *global resilience* of pancyclicity in $G(n,p)$. Similar questions about resilience, both local and global, can now be asked in the setting of the percolated hypercube.

Structure of the paper. In Section 2, we set out notation and present some preliminary results. Section 3 is dedicated to the proof of Theorem 1.

2 Preliminaries

2.1 Notation

For a positive integer n, we write $[n] = \{1,\dots,n\}$. Rounding notation is systematically omitted for better readability whenever it does not affect the validity of our arguments. All logarithms are with respect to the natural basis e. We denote by $\mathbb{N}$ the set of natural numbers, and by $2\mathbb{N}$ the set of *even* natural numbers.

[1] With high probability, that is, with probability tending to one as d tends to infinity.

Given a graph G, we denote its vertex set by $V(G)$ and its edge set by $E(G)$. For a set $A \subseteq V(G)$, we denote by $N_G(A)$ the external neighbourhood of A in G, that is, the neighbours of A in $V(G) \setminus A$. Given $A, B \subseteq V(G)$ with $A \cap B = \varnothing$, we denote by $e_G(A, B)$ the number of edges in G with one endpoint in A and the other endpoint in B. If the graph G is clear from context, we often omit the subscript. Furthermore, we denote by $G[A]$ the subgraph of G induced by A. Given $v \in V(G)$ and $r \in \mathbb{N}$, we denote by $B(v, r)$ the ball of radius r centred at v, that is, the set of vertices at distance at most r from v.

In the hypercube Q^d, the *i-th layer* consists of the vertices of Q^d with exactly i coordinates equal to 1. A *monotone path* in the hypercube is a path containing at most one vertex on each layer of that hypercube. A monotone path is *maximal* when its length (that is, number of edges) coincides with the dimension of the host hypercube.

Given $p \in [0, 1]$, we form Q^d_p by retaining every edge of Q^d independently and with probability p. Given $q, p \in [0, 1]$, we define the *mixed-percolated hypercube* $Q^d_p(q)$ as the graph $Q^d_p[V_q]$ where $V_q \subseteq V(Q^d)$ is a random set obtained by retaining every vertex independently with probability q.

Throughout the paper, when considering subgraphs $H \subseteq Q^d$, we slightly abuse notation and write H_p for the random subgraph $H \cap Q^d_p$. This naturally couples the graphs H_p and H'_p for different $H, H' \subseteq Q^d$.

2.2 Auxiliary results

We will use the following classical Chernoff bound (see, for example, [3, Theorem A.1.12]).

Lemma 2.1. *For any binomial random variable X and $a \in [0, \mathbb{E}X]$,*

$$\mathbb{P}\left(|X - \mathbb{E}X| \geq a\right) \leq 2\exp\left(-\frac{a^2}{3\mathbb{E}X}\right).$$

We will often use Lemma 2.1 when estimating a one-sided tail bound of $X \sim Bin(n, p)$, noting that

$$\mathbb{P}\left[X \leq a\right] \leq \mathbb{P}\left[|X - \mathbb{E}X| \geq a - \mathbb{E}X\right], \quad \text{and} \quad \mathbb{P}[X \geq a] \leq \mathbb{P}[|X - \mathbb{E}X| \geq a - \mathbb{E}X].$$

Next, we state a simplified version of the main theorem in [6], estimating the probability of existence of maximal monotone paths in supercritical percolation on high-dimensional hypercubes.

Theorem 2.2 (see Theorem 1 in [6])**.** *For all sufficiently large integers D, the following holds. Let $\rho = \rho(D) > \mathrm{e}/D$ then, Q^D_ρ contains a maximal monotone path with probability at least $1/2$.*

We will also utilise the following estimate on the probability of the existence of maximal monotone paths in percolated hypercubes.

Lemma 2.3. *Let $D \geq 1$ be an integer and $\rho = \rho(D) \leq 1/D$. Then, Q^D_ρ contains a maximal monotone path with probability at least $(\rho D/(2\mathrm{e}))^D$.*

Proof. We analyse the following random greedy algorithm. Denote by v_0 the all-zero vertex of Q^D, and set $\mathbf{i} = 0$. At each step $0 \leq \mathbf{i} \leq D - 1$, consider the vertex $v_{\mathbf{i}}$ and expose the edges from $v_{\mathbf{i}}$ to the $(\mathbf{i}+1)$-st layer in the hypercube Q^D_ρ. If no such edge exists, the algorithm terminates. Otherwise, fix the first such edge (according to an arbitrary order), call its other endpoint $v_{\mathbf{i}+1}$ and increment $\mathbf{i}$ by 1.

Note that, in this algorithm, no edge is exposed more than once. Moreover, since every vertex on the i-th layer contains $D-i$ edges towards the $(i+1)$-st layer for all $i \in [0, D-1]$, conditionally on the event that the algorithm does not terminate before reaching layer i, the probability that it reaches layer $i+1$ is $\mathbb{P}(\mathrm{Bin}(D-i,\rho)\geq 1) = 1-(1-\rho)^{D-i}$. Hence, the probability that the algorithm reaches layer D, which implies the existence of a maximal monotone path, is

$$\prod_{i=0}^{D-1}(1-(1-\rho)^{D-i}) = \prod_{j=1}^{D}(1-(1-\rho)^{j}) \geq \prod_{j=1}^{D}(1-\mathrm{e}^{-j\rho}) \geq D!\left(\frac{\rho}{2}\right)^{D} \geq \left(\frac{\rho D}{2\mathrm{e}}\right)^{D},$$

where the first inequality above uses that $1-t\leq \mathrm{e}^{-t}$ for every $t\geq 0$, the second inequality holds since $1-\mathrm{e}^{-t}\geq t/2$ for every $t\in[0,1]$, and the third inequality uses that $D!\geq (D/\mathrm{e})^D$ for every integer $D\geq 1$. □

Our proof also critically relies on the aforementioned recent result from [7] on the existence of a nearly spanning cycle in the (mixed) percolated hypercube.

Theorem 2.4 (Theorem 1 and Remark 1 in [7])**.** *For every fixed $q\in(0,1]$ and $\varepsilon\in(0,1)$, there exists a constant $c_0=c_0(q,\varepsilon)>0$ such that, for every $c\geq c_0$, **whp** a longest cycle in $Q^d_{c/d}(q)$ has length at least $(1-\varepsilon)q\cdot 2^d$.*

We require some results on the component structure and typical properties of the giant component in bond percolation on the hypercube. We recall that the *order* of a connected graph is the number of vertices in it.

Theorem 2.5. *Let $c>1$ and $p=c/d$. Then, there exist constants $C=C(c)>0$ and $y=y(c)>0$ such that the following properties hold **whp**.*

(a) *There exists a unique giant component in Q^d_p whose order is at least $y2^d$. All other components have order at most Cd.*

(b) *The diameter of the giant component is at most $Cd(\log d)^2$.*

Theorem 2.5(a) follows from the classical results of Ajtai, Komlós and Szemerédi [2] and Bollobás, Kohayakawa and Łuczak [11] (see [31] for a simple and self-contained proof). Theorem 2.5(b) follows from [17, Theorem 6(a)] (note that, whilst [17, Theorem 6(a)] is stated for c sufficiently close to 1, as discussed in [17, Page 747, first paragraph], the results naturally extend to any constant $c>1$ when allowing the constant C to depend on c).

We also require some variant of the above results (and other typical properties) in the setting of mixed percolation.

Theorem 2.6. *Let $q\in(0,1]$. Then, there exists a constant $c_0=c_0(q)$ such that, for all $c\geq c_0$, there are $y=y(c,q)>0$ and $C=C(c,q)>0$ such that the following properties hold **whp**.*

(a) *There exists a unique giant component in $Q^d_{c/d}(q)$ whose order is at least $y2^d$. All other components have order at most Cd.*

(b) *The diameter of the giant component is at most $2^d/d^8$.*

Let us stress here that Theorem 2.6(b) is far from being tight, however this crude estimate will suffice for our needs. The proof of Theorem 2.6(a) follows the proof of [31, Theorem 2] with very straightforward modifications in Lemmas 4–8 therein due to the vertex percolation.

We turn our attention to the proof of Theorem 2.6(b), which requires some preparation. The following lemma provides an upper bound on the number of k-vertex trees in a graph G of maximum degree d. It is an immediate consequence of [8, Lemma 2].

Lemma 2.7. *Let G be a graph of maximum degree at most d, $v\in V(G)$ and $k\geq 1$ be an integer. Denote by $t_k(G,v)$ the number of trees on k vertices in G rooted at v. Then,*

$$t_k(G,v)\leq (\mathrm{e}d)^{k-1}.$$

We will also use a simplified version from [24] of the well-known Harper's vertex-isoperimetric inequality [25].

Lemma 2.8 (Lemma 6.2 in [24]). *There exists a constant $c_1 > 0$ such that, for every set $S \subseteq V(Q^d)$ of size $|S| \le d^{10}$, one has $|N(S)| \ge c_1 d|S|$.*

With these two lemmas at hand, we can show that sets $S \subseteq V(Q^d)$ of size $k \in [d, d^{10}]$ typically have a constant vertex-expansion in $Q^d_p(q)$, assuming they span a connected subgraph of $Q^d_p(q)$.

Lemma 2.9. *Let $q \in (0,1]$. Then, there exist positive constants $c_2 = c_2(q)$ and $a = a(c_2)$ such that, for all $c \ge c_2$, with $p = c/d$, the following holds* **whp**. *For every $S \subseteq V(Q^d_p(q))$ of size $k \in [d, d^{10}]$ such that $Q^d_p(q)[S]$ is connected, we have that $|N_{Q^d_p(q)}(S)| \ge ak$.*

Proof. Let c_1 be the constant whose existence is guaranteed by Lemma 2.8.

On the event that there exists a set $S \subseteq V(Q^d_p(q))$ of size $k \in [d, d^{10}]$ such that $Q^d_p(q)[S]$ is connected and $|N_{Q^d_p(q)}(S)| < ak$, the following holds. There exist $k \in [d, d^{10}]$, a vertex $v \in V(Q^d)$ and a tree T on k vertices in Q^d such that each of the following holds:

- $v \in V(T)$;
- T is a subgraph of $Q^d_p(q)$;
- out of the $|N_{Q^d}(V(T))|$ neighbours of $V(T)$ in Q^d, at most ak belong to $V(Q^d_p(q))$ and are connected to $V(T)$ in $Q^d_p(q)$.

There are 2^d choices for v and, thereafter, by Lemma 2.7, at most $(ed)^{k-1}$ many choices for a tree $T \subseteq Q^d$ on k vertices that contains v. Each such tree T is present in $Q^d_p(q)$ with probability $q^k p^{k-1}$. Conditionally on $T \subseteq Q^d_p(q)$, each vertex in $N_{Q^d}(V(T))$ belongs to $V(Q^d_p(q))$ and is connected to T in $Q^d_p(q)$ independently with probability at least pq. Moreover, Lemma 2.8 and the fact that $k \le d^{10}$ imply that $|N_{Q^d}(V(T))| \ge c_1 kd$. Thus,

$$\mathbb{P}\left(|N_{Q^d_p(q)}(V(T))| < ak\right) \le \mathbb{P}\left(\mathrm{Bin}(c_1 kd, pq) \le c_1 kdpq/2\right),$$

where we assumed that $a \le c_1 c_2 q/2 \le c_1 cq/2 = c_1 dpq/2$. Thus, by the union bound, the probability that there exists a set $S \subseteq V(Q^d_p(q))$ of size $k \in [d, d^{10}]$ such that $Q^d_p(q)[S]$ is connected and $|N_{Q^d_p(q)}(S)| < ak$ is at most

$$\begin{aligned}
\sum_{k=d}^{d^{10}} 2^d (ed)^{k-1} p^{k-1} q^k \mathbb{P}(\mathrm{Bin}(c_1 kd, qp) \le c_1 kdqp/2) &\le \sum_{k=d}^{d^{10}} 2^d (edpq)^k 2 \exp\left\{-\frac{c_1 kdqp}{12}\right\} \\
= \sum_{k=d}^{d^{10}} 2^{d+1} \left(ecq \cdot \exp\left\{-\frac{c_1 cq}{12}\right\}\right)^k &\le \sum_{k=d}^{d^{10}} 2^{d+1} \left(2^{-2}\right)^k \le d^{10} 2^{-d+1} = o(1),
\end{aligned}$$

where the first inequality follows from Chernoff's bound (Lemma 2.1) and the second inequality uses that $c \ge c_2(q)$ is suitably large with respect to q. This finishes the proof. □

We are now ready to prove Theorem 2.6(b).

Proof of Theorem 2.6(b). Fix $p = c/d$. By Lemma 2.9, **whp** every connected set in $Q^d_p(q)$ of size between d and d^{10} has vertex-expansion by a factor of at least a for some constant $a > 0$. We assume in the sequel that this property holds.

We may assume that the giant component of $Q^d_p(q)$, which we denote by L_1, satisfies that $|V(L_1)| \ge \frac{2^d}{d^{10}} \ge d$. Thus, for every $v \in V(L_1)$, we have that $|B(v,d)| \ge d$. Hence, by our vertex-expansion assumption, for every $v \in V(L_1)$ we have that $|B(v, 2d)| \ge d^{10}$.

Fix two vertices u, v in L_1 and let $P = \{v_0, \dots, v_t\}$ be a shortest path (in L_1) between them, where $u = v_0$ and $v = v_t$. Let $x_1, \dots, x_{t/5d}$ be a set of $\frac{t}{5d}$ vertices along P, such that the distance between x_i and x_j is at least $5d$, for any $i \neq j$. Hence, for every $i \neq j$, we have that $B(x_i, 2d) \cap B(x_j, 2d) = \varnothing$. Thus, $\sum_{i=1}^{t/5d} |B(x_i, 2d)| \leq |L_1| \leq 2^d$. On the other hand, by the above, $\sum_{i=1}^{t/5d} |B(x_i, 2d)| \geq \frac{t}{5d} \cdot d^{10}$. Therefore, $t \leq \frac{5 \cdot 2^d}{d^9} < \frac{2^d}{d^8}$. This completes the proof. □

3 Proof of Theorem 1

The proof proceeds in four steps, which separately guarantee that typically cycles in Q^d_p with any length in the following sets $I_1, \dots, I_4$ exist:

- (Very short cycles) $I_1 = [4, d/5] \cap 2\mathbb{N}$;
- (Short cycles) $I_2 = [d/5, d^{10}] \cap 2\mathbb{N}$;
- (Medium cycles) $I_3 = [d^{10}, 2^{d-4}] \cap 2\mathbb{N}$;
- (Long cycles) $I_4 = [2^{d-4}, (1-\varepsilon)2^d] \cap 2\mathbb{N}$.

Since the union $I_1 \cup I_2 \cup I_3 \cup I_4 = [4, (1-\varepsilon)2^d] \cap 2\mathbb{N}$, this is enough to derive Theorem 1.

Finding very short cycles. First, we show the typical existence of all cycles with lengths in I_1. In fact, we will show a slightly stronger result which will be useful in the subsequent cases.

Lemma 3.1. *Let D be a sufficiently large integer and $\rho = \rho(D) \geq 10/D$. Then, with probability at least $1 - \exp\{-1.1^D\}$, Q^D_ρ contains simultaneously all cycles of even length in the interval $[4, D/8]$.*

Proof. Since the property of containing cycles of every even length in $[4, D/8]$ is monotone with respect to ρ, we may (and do) assume that $\rho = 10/D$. Fix an integer $\ell = \ell(D) \in [1, D/16 - 1]$. Then, by Lemma 2.3, the probability that Q^ℓ_ρ contains a maximal monotone path is bounded from below by $\eta = \eta(\ell) := (\ell\rho/(2\mathrm{e}))^\ell$.

Divide the hypercube Q^D into a family $\mathcal{Q}$ of $2^{D-\ell-1}$ disjoint copies of $Q^{\ell+1}$. For every copy $Q \in \mathcal{Q}$ of $Q^{\ell+1}$, we split Q into disjoint copies Q_0 and Q_1 of Q^ℓ, which are joined by a matching. Note that if $(Q_0)_\rho$ and $(Q_1)_\rho$ simultaneously contain maximal monotone paths P_0, P_1 whose starting and ending points are also connected in Q^D_ρ, then Q^D_ρ contains a cycle of length $2\ell + 2$. Note further that these events are independent for every two $Q, Q' \in \mathcal{Q}$. Thus, the probability that Q^D_ρ does not contain a cycle of length $2\ell + 2$ is at most

$$
\begin{aligned}
(1 - \rho^2\eta^2)^{2^{D-\ell-1}} &\leq \exp\left\{-\rho^2\eta^2 2^{D-\ell-1}\right\} = \exp\left\{-\left(\frac{10}{D}\right)^2 \left(\frac{\ell\rho}{2\mathrm{e}}\right)^{2\ell} 2^{D-\ell-1}\right\} \\
&\leq \exp\left\{-D^{-2}\left(\frac{10\ell}{3\mathrm{e}D}\right)^{2\ell} 2^D\right\} \\
&\leq \exp\left\{-D^{-2}\left(\frac{2}{(3\mathrm{e})^{1/8}}\right)^D\right\} \leq \exp\left\{-1.2^D\right\}.
\end{aligned}
$$

The union bound over the $D/16 - 1$ cycle lengths completes the proof. □

Proof of Theorem 1 for the lengths in I_1. Follows immediately from Lemma 3.1 with $D = d$, and assuming that $c \geq 10$. □

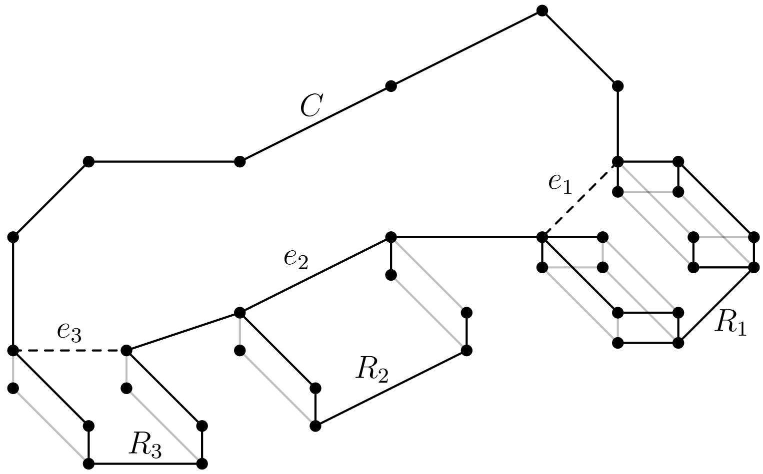

Figure 1: Illustration of the proof of Lemma 3.2. Here, $k_1 = 3$ and $k_2 = 2$. Edges in the cubes R_1, R_2 and R_3 missing from Q^d_p are depicted in light grey. We fail to extend the cycle $\mathcal{C}$ in R_2: indeed, the right part of R_2 does not contain a maximal monotone path after percolation. However, we manage to replace the (dashed) edges e_1 and e_3 by paths of length $2k_1 + 1$ and $2k_2 + 1$ in $(R_1)_p$ and $(R_3)_p$, respectively.

Finding short cycles. The next part of the proof consists of arguing about cycles of all lengths in I_2. To show this, we will utilise Lemma 3.1 in an inductive way. The following lemma will serve as the key inductive step.

Lemma 3.2. *Let D be a sufficiently large integer and $\rho \geq 2^7\mathrm{e}/D$. Let $H \subseteq Q^D$ be the subcube of dimension $D/2$ obtained by fixing the last $D/2$ coordinates to be 0. Condition on the graph H_ρ and the event that it contains a cycle C of length $2L \in [D/100, D^{11}]$. Then, with probability at least $1-\exp\{-\Omega(D)\}$, we have that $(Q^D \setminus H)_\rho \cup C$ contains simultaneously all cycles of even length in the interval $[2L + 2^{-5}D, 2^{-8}LD]$.*

Proof. Fix $2\ell \in [2L + 2^{-5}D, 2^{-8}LD]$ for $\ell \in \mathbb{N}$. Let k be such that $2L + 2k = 2\ell$ and note that $2^{-6}D \leq k \leq 2^{-9}LD$. Note further that k can be written as $k = t \cdot k_1 + k_2$, where $k_1 = 2^{-5}D$, $t \leq 2^{-4}L$, and $k_2 \in [2^{-6}D, 3 \cdot 2^{-6}D]$ (in particular, $2\ell = 2L + 2tk_1 + 2k_2$).

Fix L vertex-disjoint edges in C and denote them by $e_1 = u_1u_1', \ldots, e_L = u_Lu_L'$. For each $i \in [L]$, let R_i be the subcube obtained by fixing the *first* $D/2-1$ coordinates on which $\{u_i, u_i'\}$ agree, and letting the other $D/2+1$ coordinates vary. Note that, since H has the *last* $D/2$ coordinates fixed to be 0 and $C \subseteq H$, for each $i \in [L]$, we have that $H \cap R_i = e_i$. Furthermore, for every $i \neq j \in [L]$, we have that u_j, u_j' disagree on some of the first $D/2$ coordinates with each of u_i, u_i', and thus $R_i \cap R_j = \varnothing$.

Now, for each $i \in [L]$, let $R_{u_i} \subseteq R_i$ be the subcube obtained by fixing the first $D/2$ coordinates to be as in u_i, and let $R_{u_i'}$ be the subcube obtained by fixing the first $D/2$ coordinates to be as in u_i'. Note that $V(R_i) = V(R_{u_i}) \cup V(R_{u_i'})$, $R_{u_i} \cap R_{u_i'} = \varnothing$, and that each of R_{u_i} and $R_{u_i'}$ shares no edges with H. For every $i \in [L]$, we orient the cubes R_{u_i} and $R_{u_i'}$ such that u_i, u_i' are their all-zero coordinate. Thus, the j-th layer in R_{u_i} (resp. in $R_{u_i'}$) is the set of vertices in R_{u_i} (resp. in $R_{u_i'}$) at distance j from u_i (resp. from u_i'). There are $\binom{D/2}{j}$ vertices in these layers.

Then, for every $i \in [L/2]$, we expose the edges of $(Q^D \setminus H)_\rho$ between the k_1-th layer in R_{u_i} and the k_1-th layer in $R_{u_i'}$. Since D is sufficiently large, for each i, the probability that none of these edges are in $(Q^D \setminus H)_\rho$ is $(1-\rho)^{\binom{D/2}{k_1}} \leq \frac{1}{10}$.

For each i such that at least one such edge exist, we choose one of these edges arbitrarily and denote it by $f_i = w_iw_i'$ where $w_i \in V(R_{u_i}), w_i' \in V(R_{u_i'})$. For each such i, let $Q(u_i, w_i) \subseteq R_{u_i}$ be the unique subcube of dimension k_1 which contains u_i and w_i. Note that there is a natural isomorphism from $Q(u_i, w_i)$ to Q^{k_1} which maps u_i to the all-zero vector; thus, when we talk about monotone paths in $Q(u_i, w_i)$, we mean a path whose image under this isomorphism is monotone. Define $Q(u_i', w_i') \subseteq R_{u_i'}$ in a similar manner.

We note that $u_iu_i', w_iw_i' \in E((Q^D \setminus H)_\rho \cup C)$, and that all the edges of $Q(u_i, w_i), Q(u_i', w_i')$ have not been exposed so far. Also, if there is a monotone path from u_i to w_i in $Q(u_i, w_i)_\rho$, and from u_i' to w_i' in $Q(u_i', w_i')_\rho$, then replacing e_i with these two paths and f_i increases the length of C by $2k_1$ (see Figure 1). Since $\rho k_1 = \rho 2^{-5}D > e$, by Theorem 2.2, the probability that $Q(u_i, w_i)_\rho$ contains a (maximal) monotone path from u_i to w_i is at least $1/2$, and the same holds for the probability of the existence of a maximal monotone path in $Q(u_i', w_i')_\rho$. Hence, the number of $i \in [L/2]$ for which both f_i and the pair of maximal monotone paths in $Q(u_i, w_i)$ and $Q(u_i', w_i')$ are contained in Q_ρ^D stochastically dominates $\text{Bin}\left(L/2, \frac{9}{10} \cdot \frac{1}{4}\right)$: indeed, the edge f_i is present in $(R_i)_\rho$ with probability at least $\frac{9}{10}$ and, conditionally on it, the said two monotone paths exist independently with probability at least $1/2$ each. By Lemma 2.1 we have that, with probability $1 - \exp\{-\Omega(D)\}$, there are at least $L/9$ such i. Choosing $t \le 2^{-4}L < L/9$ such i allows us to increase the length of C by $2tk_1$.

We can now repeat the same argument for $i \in [L/2+1, L]$, this time seeking only one suitable i for k_2 (instead of t for k_1), which follows verbatim. This produces, with probability at least $1 - \exp\{-\Omega(D)\}$, a cycle of length $2L + 2tk_1 + 2k_2 = 2\ell$ (see Figure 1). A union bound over the at most $LD \le D^{12}$ possible choices of k completes the proof. □

With Lemma 3.2 at hand, we are ready to show the typical existence of cycles of lengths in I_2.

Proof of Theorem 1 for the lengths in I_2. For every $i \in \{0, \dots, 10\}$, define $Q_{(i)}$ to be the subcube of Q^d obtained by fixing the last $(1 - 2^{i-12})d$ coordinates to be zero, and letting the remaining $2^{i-12}d$ coordinates vary. Let $H_i = (Q_{(i)})_p$.

We claim that there exist constants $b_0, b_1, \dots, b_{10} > 0$ such that **whp** H_i contains simultaneously every cycle of even length in the interval $[4, b_i d^{i+1}]$. We prove the claim inductively. As a base case, since c is sufficiently large, by Lemma 3.1 applied to $H_0 = (Q_{(0)})_p$ (that is, with $D = 2^{-12}d$), **whp** H_0 contains simultaneously every cycle of even length in the interval $[4, 2^{-15}d]$, and so we can take $b_0 = 2^{-15}$. Moreover, conditionally on the latter event, a direct application of Lemma 3.2 shows the claim for $i = 1$.

Now, fix $i \in [9]$ and suppose that there exists a constant $b_i > 0$ such that **whp** H_i contains simultaneously every cycle of even length in the interval $[4, b_i d^{i+1}]$. We expose the edges of H_i and condition on this event.

Fix a cycle C_i of length $b_i d^{i+1} - d/2$ in H_i. Then, applying Lemma 3.2 with $D = 2^{i-11}d$ to H_i, C_i in the ambient cube Q_{i+1}, we see that there exists a constant $b_{i+1} > 0$ such that **whp** $(H_{i+1} \setminus H_i) \cup C_i$ contains simultaneously every cycle of even length in the interval

$$[b_i d^{i+1} - d/2 + 2^{i-16}d, 2^{i-19}b_i d^{i+2}] \supseteq [b_i d^{i+1}, b_{i+1} d^{i+2}],$$

where we take $b_{i+1} = 2^{i-19}b_i$. Hence, since $H_i \subseteq H_{i+1}$, **whp** H_{i+1} contains simultaneously every cycle of even length in the interval $[4, b_i d^{i+1}] \cup [b_i d^{i+1}, b_{i+1} d^{i+2}] = [4, b_{i+1} d^{i+2}]$.

As a result, there exists a constant $b_{10} > 0$ such that **whp** $H_{10} \subseteq Q_p^d$ contains simultaneously every cycle of even length in the interval $[4, b_{10} d^{11}] \supseteq I_2$, as required. □

Finding cycles with length in the middle range. We now turn to finding cycles with lengths in I_3. We will build upon the cycle whose typical existence is guaranteed by Theorem 2.4 with $q = 1$ and c sufficiently large.

Formally, let us partition Q^d into four disjoint subcubes $(Q_{i,j})_{i,j \in \{0,1\}}$ according to the values of the first two coordinates. Note that, by Theorem 2.4, **whp** there exists a cycle of length at least 2^{d-3} in $(Q_{0,0})_p$. In the sequel, we condition on a cycle $C \subseteq Q_{0,0}$ of length at least 2^{d-3}.

For every $\boldsymbol{j} \in \{0,1\}^{d/2-2}$, let $Q_{(0,0)}(\boldsymbol{j})$ be the subcube obtained by fixing the first two coordinates to be zero, the vector consisting of the next $d/2-2$ coordinates to be $\boldsymbol{j}$, and letting the last $d/2$ coordinates to vary. Note that $\mathcal{Q}_0 := \{Q_{(0,0)}(\boldsymbol{j}) \colon \boldsymbol{j} \in \{0,1\}^{d/2-2}\}$ is a partition of $Q_{0,0}$ into $2^{d/2-2}$ disjoint hypercubes of dimension $d/2$. Similarly, for every $\boldsymbol{j} \in \{0,1\}^{d/2-2}$, we define $Q_{1,0}(\boldsymbol{j})$ to be the subcube obtained by fixing the first two coordinates to be $\{1,0\}$, the vector consisting of the next $d/2-2$ coordinates to be $\boldsymbol{j}$, and letting the last $d/2$ coordinates to vary In particular, we have that $\mathcal{Q}_1 := \{Q_{1,0}(\boldsymbol{j}) \colon \boldsymbol{j} \in \{0,1\}^{d/2-2}\}$ is a partition of $Q_{1,0}$ into $2^{d/2-2}$ disjoint hypercubes of dimension $d/2$.

For any length $2\ell \in I_3$, by a straightforward greedy approach, one can find (possibly intersecting) paths $P_1(\ell), \ldots, P_{d^4}(\ell)$ with the following properties:

(P1) each of these paths is contained in C,

(P2) the length of each path is equal to $2\ell - d^2$,

(P3) for every distinct $i, j \in [d^4]$, every endpoint u of $P_i(\ell)$ and every endpoint v of $P_j(\ell)$, u and v are in distinct subcubes in $\mathcal{Q}_0$,

(indeed, note that there are $\Omega(2^d)$ paths of length $2\ell - d^2$ contained in C, and there are $2^{d/2}$ vertices in each subcube in $\mathcal{Q}_0$). Given $2\ell \in I_3$ and $i \in [d^4]$, we denote by $u_{i,\ell}$ and $v_{i,\ell}$ the two endpoints of $P_i(\ell)$, and by $u'_{i,\ell}$ and $v'_{i,\ell}$ the neighbours in $Q_{1,0}$ of $u_{i,\ell}$ and $v_{i,\ell}$, respectively. Let $\boldsymbol{j}(u_{i,\ell}), \boldsymbol{j}(v_{i,\ell})$ be the coordinates such that $u_{i,\ell} \in Q_{0,0}(\boldsymbol{j}(u_{i,\ell})), v_{i,\ell} \in Q_{0,0}(\boldsymbol{j}(v_{i,\ell}))$. Then, by construction of $\mathcal{Q}_0, \mathcal{Q}_1$, we have that $u'_{i,\ell} \in Q_{1,0}(\boldsymbol{j}(u_{i,\ell})) =: Q(u'_{i,\ell})$ and $v'_{i,\ell} \in Q_{1,0}(\boldsymbol{j}(v_{i,\ell})) =: Q(v'_{i,\ell})$. In particular, property (P3) implies that, for distinct $i, j \in [d^4]$, the two sets $\{Q(u'_{i,\ell}), Q(v'_{i,\ell})\}$ and $\{Q(u'_{j,\ell}), Q(v'_{j,\ell})\}$ are disjoint. However, note that either set may nevertheless consist of a single subcube.

We require the following lemma.

Lemma 3.3. *Let $C \subseteq Q_{0,0}$ be a cycle of length at least 2^{d-3}, $c > 0$ be sufficiently large and $p = c/d$. For each $2\ell \in I_3$ let $P_1(\ell), \ldots, P_{d^4}(\ell)$ be a family of paths satisfying properties (P1)–(P3). Then,* ***whp****, for every $2\ell \in I_3$, there is an index $i \in [d^4]$ such that each of the following properties hold simultaneously:*

(i) each of the edges $u_{i,\ell}u'_{i,\ell}$ and $v_{i,\ell}v'_{i,\ell}$ belongs to Q^d_p,

(ii) the vertices $u'_{i,\ell}, v'_{i,\ell}$ belong to components of order at least $2^{d/2}/d$ in $Q(u'_{i,\ell})_p$ and $Q(v'_{i,\ell})_p$, respectively,

(iii) the shortest path $P'_{i,\ell}$ between $v'_{i,\ell}$ and $u'_{i,\ell}$ in $(Q_{1,0})_p$ has length at most $d(\log d)^3$.

Proof. Fix a length $2\ell \in I_3$. Then, by Theorem 2.5, for every $i \in [d^4]$, the probability that each of (i) and (ii) holds for $P_i(\ell)$ is at least $p^2 \cdot (1-o(1))y^2 \geq p^2y^2/2$, where $y = y(c/2) > 0$ (as the dimension of the relevant cubes is $d/2$) is the constant guaranteed by Theorem 2.5(a). Indeed, if $Q(u'_{i,\ell}) \neq Q(v'_{i,\ell})$, then this is clear by independence and, if $Q(u'_{i,\ell}) = Q(v'_{i,\ell})$, then the claim holds by Harris' inequality (see, e.g., [3, Theorem 6.3.3]) and the observation that the property of belonging to a component of order at least $2^{d/2}/d$ in $Q(u'_{i,\ell})_p = Q(v'_{i,\ell})_p$ is increasing. Moreover, by property (P3), these events are jointly independent for different $i \in [d^4]$. Therefore, we have that, with probability at least $1 - \mathbb{P}(\mathrm{Bin}(d^4, p^2y^2/2) = 0) = 1 - o(2^{-d})$, there is some $i \in [d^4]$ such that (i)-(ii) hold. In particular, by a union bound, this is true for all $2\ell \in I_3$.

Finally, we note that, by Theorem 2.5, **whp** $(Q_{1,0})_p$ contains a unique component L of order at least $2^d/d$, and all the other components have order $O(d)$. Hence every vertex that lies in a component of size at least $2^{d/2}/d$ belongs to L. In addition, by Theorem 2.5, L has diameter at most $d(\log d)^3$. In particular, **whp**, for any $i \in [d^4]$ such that (i)-(ii) hold, $v'_{i,\ell}$ and $u'_{i,\ell}$ lie in L and are thus joined by a path $P'_{i,\ell}$ of length at most $d(\log d)^3$ in $(Q_{1,0})_p$. $\square$

Note that, if the conclusion of Lemma 3.3 holds for some $2\ell \in I_3$ and $i \in [d^4]$, then there is a cycle $C' = u_{i,\ell}P_i(\ell)v_{i,\ell}v'_{i,\ell}P'_{i,\ell}u'_{i,\ell}u_{i,\ell}$ in $(Q_{0,0} \cup Q_{1,0})_p$ of length at most $2\ell - d^2 + d(\log d)^3 + 2$. Furthermore, $C' \cap Q_{0,0} = P_i(\ell)$.

With Lemma 3.3 at hand, we are now ready to complete the proof of Theorem 1 for cycles of lengths in I_3.

Proof of Theorem 1 for the lengths in I_3. First, we expose the edges of $(Q_{0,0} \cup Q_{1,0})_p$. By Theorem 2.4, **whp** there exists a cycle C of length at least 2^{d-3} in $(Q_{0,0})_p$. For each $2\ell \in I_3$, let $P_1(\ell), \dots, P_{d^4}(\ell)$ be a family of paths satisfying properties (P1)-(P3). Then, by Lemma 3.3, **whp**, for each $2\ell \in I_3$, we have that $(Q_{0,0} \cup Q_{1,0})_p$ contains a cycle C' of length $|C'(\ell)| \le 2\ell - d^2 + d(\log d)^3 + 2$ such that $C'(\ell) \cap Q_{0,0}$ is a path of length $2\ell - d^2$.

For each $2\ell \in I_3$, we fix vertex-disjoint edges $e_1(\ell), \dots, e_{\ell/3}(\ell)$ in $C'(\ell) \cap Q_{0,0}$. For each such edge $e_i(\ell)$, there is a unique path $T_i(\ell)$ of length 3 with the same endpoints as $e_i(\ell)$ whose middle edge lies in the hypercube $Q_{0,1}$. Note that, by construction, the paths $T_1(\ell), \dots, T_{\ell/3}(\ell)$ are disjoint from each other and from the set $Q_{0,0} \cup Q_{1,0}$ of exposed edges, and the probability that any of them is entirely included in Q_p^d is p^3. Hence, a standard application of Chernoff's bound (Lemma 2.1) together with a union bound over the at most 2^d possible values of $2\ell \in I_3$ implies that, **whp**, for every $2\ell \in I_3$, at least $d^{-4}\ell \ge d^2$ of the paths $T_1(\ell), \dots, T_{\ell/3}(\ell)$ are contained in Q_p^d. By replacing $(2\ell - |C'(\ell)|)/2 \le d^2$ edges $e_i(\ell)$ with $i \in [\ell/3]$ with the corresponding paths $T_i(\ell)$, this produces a cycle of length exactly 2ℓ, as desired. □

Finding long cycles. Finally, it remains to ensure that cycles of lengths in I_4 exist in Q_p^d. Given $\delta \in [0,1]$, partition $V(Q^d)$ into sets $V_1 = V_1(\delta), V_2 = V_2(\delta)$ and $V_3 = V_3(\delta)$ where a vertex is assigned to V_1 with probability $q_1 = 1 - \delta/2$, and to any of V_2, V_3 with probability $q_2 = q_3 = \delta/4$, independently for different vertices. We will make use of the following lemma.

Lemma 3.4. *For every $\varepsilon \in (0,1)$, there exist constants $c(\varepsilon) > 0$ and $\delta \in [0,1]$ such that, with V_1, V_2, V_3 as above, for all $c \ge c(\varepsilon)$ with $p = p(d) = c/d$, each of the following holds* ***whp****.*

(P1') *$Q_p^d[V_1]$ contains a cycle C_1 of length at least $(1-\varepsilon)2^d$.*

(P2') *$Q_p^d[V_2]$ contains a unique component L_2 of size at least $2^d/d$ whose diameter is at most $2^d/d^8$.*

(P3') *All but $o(2^d/d^{10})$ vertices have at least one neighbour (in Q^d) in L_2.*

(P4') *For all but $o(2^d/d^{10})$ edges $uv \in E(Q^d)$, there are vertices $u', v' \in V_3$ such that u, v, v', u' (in this order) form a 4-cycle in Q^d.*

Proof. Let $\delta \in (0,1)$ be such that $(1-\delta)(1-\delta/2) = 1-\varepsilon$. Theorem 2.4, with $q = 1-\delta$, implies that **whp** $Q_p^d[V_1]$ contains a cycle of length $(1-\delta)(1-\delta/2)2^d = (1-\varepsilon)2^d$ and hence (P1') holds.

(P2') follows from Theorem 2.6, with L_2 being the unique giant component of $Q_p^d[V_2]$.

For (P3'), fix a vertex v and some δd of its neighbours in Q^d, denoted $w_1, \dots, w_{\delta d}$. Consider δd vertex-disjoint subcubes $Q(1), \dots, Q(\delta d)$, each of dimension at least $(1-\delta)d$ and each containing exactly one of the w_i (such exist by, for example, Claim 2.2 in [18]). Then, by symmetry of the hypercube and by Theorem 2.6, for every $i \in [\delta d]$, with probability α for some constant $\alpha > 0$, w_i belongs to a giant component in $(Q(i))_p[V_2]$ (which is, in fact, **whp** part of the giant component of $Q_p^d[V_2]$ by the discrepancy in the sizes of the components ensured in Theorem 2.6(a)). Thus, assuming Theorem 2.6(a), the probability that none of the neighbours of v is in the giant component of $Q_p^d[V_2]$ is at most $(1-\alpha)^{\delta d} \le \exp\{-\alpha\delta d\}$. Thus, by Markov's inequality, all but $o(2^d/d^{10})$ vertices in Q^d have at least one neighbour in the giant component of $Q_p^d[V_2]$ (where we note that the exponent 10 in the power of d is rather arbitrary).

For (P4'), fix an edge $uv \in E(Q^d)$. There are $d-1$ distinct pairs of vertices u', v' such that $uvv'u'$ form a 4-cycle in Q^d. The probability that none of these pairs belongs to V_3 is $(1-q_3^2)^{d-1} \le \exp\{-\delta^2 d/20\}$. Thus, the expected number of edges $uv \in E(Q^d)$ such that there are *no* vertices $u', v' \in V_3$ with $uvv'u'$ forming a 4-cycle is at most $\exp\{-\delta^2 d/20\} d 2^d$. Thus, by Markov's inequality, **whp** the latter event holds for $o(2^d/d^{10})$ pairs only, as desired. $\square$

We are ready to complete the proof of Theorem 1.

Proof of Theorem 1 for the lengths in I_4. Let V_1, V_2, V_3 be as in Lemma 3.4 and $c > 0, \delta \in [0,1]$ be such that (P1')-(P4') hold **whp**. We begin by exposing V_1, V_2, V_3 and $Q_p^d[V_1] \cup Q_p^d[V_2]$. We condition on (P1')-(P4'). We denote by C_1 a cycle of length at least $(1-\varepsilon)2^d$ in $Q_p^d[V_1]$, given by (P1'), and by L_2 the giant component of $Q_p^d[V_2]$, given by (P2').

Fix $2\ell \in [2^{d-3}, |C|]$ and $k = 2^d/d^8$. We enumerate $C = \{v_1, \dots, v_{|C|}\}$ and let $S_1 = \{v_1, \dots, v_k\}$ and $S_2 = \{v_{2\ell-2k+1}, \dots, v_{2\ell-k}\}$. Note that these sets are well-defined and disjoint.

By (P3'), the number of vertices $u \in S_1$ with a neighbour v in the giant component of $Q_p^d[V_2]$ is at least $k/2$. Next, expose the edges of Q_p^d between V_1 and V_2. Then, the probability that there are no edges in Q_p^d between S_1 and L_2 is at most $(1-p)^{k/2} \le \exp\left\{-2^{d-1}/d^9\right\}$. A similar argument shows that the probability that there are no edges in Q_p^d between S_2 and L_2 is also at most $\exp\left\{-2^{d-1}/d^9\right\}$.

Hence, with probability at least $1 - 2\exp\left\{-2^{d-1}/d^9\right\}$, there exist $u_1 \in S_1, u_2 \in S_2$, and $w_1, w_2 \in L_2$, such that $u_1w_1, u_2w_2 \in E(Q_p^d)$. By (P2'), there is a path $P \subseteq Q_p^d[V_2]$ of length at most k between w_1 and w_2. Denote by P' the path between u_1 and u_2 along the cycle C containing v_{k+1}. Then, with probability at least $1 - 2\exp\{-2^d/d^8\}$, there is a cycle $C' \coloneqq P' \cup P \cup \{u_1w_1, u_2w_2\}$ whose length is between $2\ell - 3k + 2$ and $(2\ell - k - 2) + k + 2 = 2\ell$ (see Figure 2).

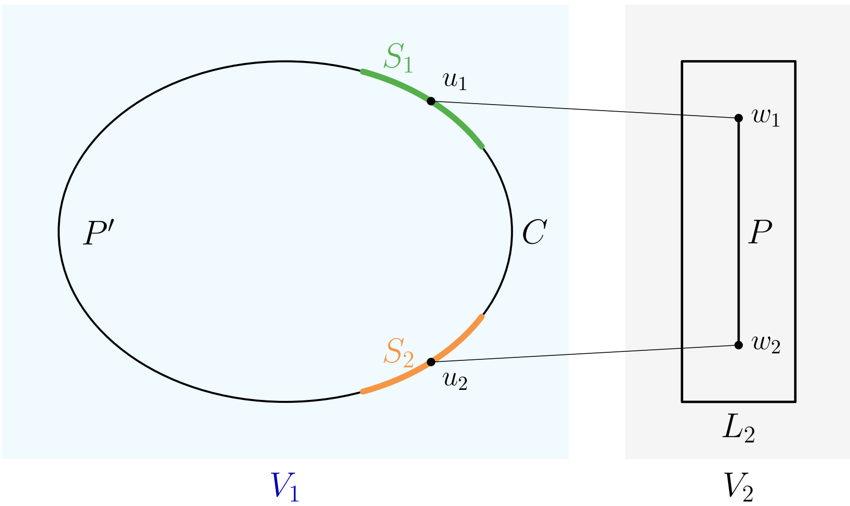

Figure 2: The construction of the cycle C' using the paths $P' \subseteq C$, $P \subseteq L_2$ and the edges u_1w_1, u_2w_2 connecting the endpoints of P' and P.

By (P4'), there are at least $\frac{2\ell-3k+2}{2} - o(2^d/d^{10}) \ge \ell/2$ vertex-disjoint edges uv in C', such that there are vertices $u', v' \in V_3$ such that $uvv'u'$ is a 4-cycle in Q^d. Since every vertex in V_3 can participate in at most $\binom{d}{2} \le d^2/2$ such cycles, there are at least $\ell/(2d^2)$ vertex-disjoint 4-cycles $uvv'u'$ of this form in Q^d.

Finally, we expose the edges in Q_p^d between V_1 and V_3 and inside V_3. Then, each of the said 4-cycles belongs to Q_p^d with probability p^3. Hence, by Chernoff's bound (Lemma 2.1), the probability that there are less than $2^d/d^6$ vertex-disjoint edges uv in P' such that there are $u', v' \in V_3$ with $uvv'u'$ forming a 4-cycle in Q_p^d is at most $1 - \exp\{-2^d/d^6\}$. Replacing

$2\ell - |\hat{C}| \leq 3k = o(2^d/d^6)$ of these edges with the vertex-disjoint 3-paths going through V_3, we obtain that, with probability at least

$$1 - 2\exp\{-2^{d-1}/d^9\} - 2\exp\{-2^d/d^6\} = 1 - o(2^{-d}),$$

a cycle of length exactly 2ℓ exists. A union bound over the at most 2^d choices of ℓ completes the proof. □

Acknowledgement

This research was funded in part by the Austrian Science Fund (FWF) [10.55776/P36131 (J. Erde), 10.55776/F1002 (M. Kang), 10.55776/ESP624 (L. Lichev), 10.55776/ESP3863424 (M. Anastos)], and by NSF-BSF grant 2023688 (M. Krivelevich). For open access purposes, the authors have applied a CC BY public copyright license to any author accepted manuscript version arising from this submission.

References

[1] M. Ajtai, J. Komlós, and E. Szemerédi. The longest path in a random graph. *Combinatorica*, 1:1–12, 1981.

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

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

[4] Y. Alon, M. Krivelevich, and E. Lubetzky. Cycle lengths in sparse random graphs. *Random Structures Algorithms*, 61(3):444–461, 2022.

[5] M. Anastos. A note on long cycles in sparse random graphs. *Electron. J. Combin.*, 30(2):Paper No. 2.21, 24, 2023.

[6] M. Anastos, S. Diskin, D. Elboim, and M. Krivelevich. Climbing up a random subgraph of the hypercube. *Electron. Commun. Probab.*, 29:13, 2024. Id/No 70.

[7] M. Anastos, S. Diskin, J. Erde, M. Kang, M. Krivelevich, and L. Lichev. Nearly spanning cycle in the percolated hypercube. *arXiv preprint arXiv:2505.04436*, 2025.

[8] A. Beveridge, A. Frieze, and C. McDiarmid. Random minimum length spanning trees in regular graphs. *Combinatorica*, 18(3):311–333, 1998.

[9] B. Bollobás. The evolution of sparse graphs. In *Graph theory and combinatorics (Cambridge, 1983)*, pages 35–57. Academic Press, London, 1984.

[10] B. Bollobás. *Random graphs*, volume 73 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, second edition, 2001.

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

[12] J. A. Bondy and V. Chvátal. A method in graph theory. *Discrete Math.*, 15(2):111–135, 1976.

[13] V. Chvátal. On Hamilton's ideals. *J. Combinatorial Theory Ser. B*, 12:163–168, 1972.

[14] P. Condon, A. Espuny Díaz, A. Girão, D. Kühn, and D. Osthus. Hamiltonicity of random subgraphs of the hypercube. *Mem. Amer. Math. Soc.*, 304(1534):v+132, 2024.

[15] C. Cooper and A. M. Frieze. Pancyclic random graphs. In *Random graphs '87 (Poznań, 1987)*, pages 29–39. Wiley, Chichester, 1990.

[16] G. A. Dirac. Some theorems on abstract graphs. *Proceedings of the London Mathematical Society*, 3(1):69–81, 1952.

[17] S. Diskin, J. Erde, M. Kang, and M. Krivelevich. Isoperimetric inequalities and supercritical percolation on high-dimensional graphs. *Combinatorica*, 44(4):741–784, 2024.

[18] S. Diskin and M. Krivelevich. Supercritical site percolation on the hypercube: small components are small. *Combin. Probab. Comput.*, 32(3):422–427, 2023.

[19] S. Diskin and M. Krivelevich. Expansion in supercritical random subgraphs of expanders and its consequences. *Random Structures Algorithms*, 65(3):576–600, 2024.

[20] N. Draganić, R. H. Montgomery, D. Munhá Correia, A. Pokrovskiy, and B. Sudakov. Hamiltonicity of expanders: optimal bounds and applications. *arXiv preprint arXiv:2402.06603*, 2024.

[21] P. Erdős and A. Rényi. On the evolution of random graphs. *Magyar Tud. Akad. Mat. Kutató Int. Közl.*, 5:17–61, 1960.

[22] W. Fernandez de la Vega. Long paths in random graphs. *Studia Sci. Math. Hungar.*, 14(4):335–340, 1979.

[23] L. Friedman and M. Krivelevich. Cycle lengths in expanding graphs. *Combinatorica*, 41(1):53–74, 2021.

[24] D. Galvin. On homomorphisms from the Hamming cube to **Z**. *Israel J. Math.*, 138:189–213, 2003.

[25] L. H. Harper. Optimal numberings and isoperimetric problems on graphs. *J. Combinatorial Theory*, 1:385–393, 1966.

[26] D. Hefetz, M. Krivelevich, and T. Szabó. Hamilton cycles in highly connected and expanding graphs. *Combinatorica*, 29(5):547–568, 2009.

[27] M. Karoński and A. Ruciński. On the number of strictly balanced subgraphs of a random graph. In *Graph theory (Łagów, 1981)*, volume 1018 of *Lecture Notes in Math.*, pages 79–83. Springer, Berlin, 1983.

[28] R. M. Karp. Reducibility among combinatorial problems. In *Complexity of computer computations. Proceedings of a symposium on the complexity of computer computations.*, pages 85–103. New York-London: Plenum Press, 1972.

[29] J. Komlós and E. Szemerédi. Limit distribution for the existence of Hamiltonian cycles in a random graph. *Discrete Math.*, 43(1):55–63, 1983.

[30] A. D. Koršunov. Solution of a problem of P. Erdős and A. Rényi on Hamiltonian cycles in nonoriented graphs. *Diskret. Analiz*, (31):17–56, 90, 1977.

[31] M. Krivelevich. Component sizes in the supercritical percolation on the binary cube. *arXiv preprint arXiv:2311.07210*, 2023.

[32] M. Krivelevich, C. Lee, and B. Sudakov. Resilient pancyclicity of random and pseudorandom graphs. *SIAM J. Discrete Math.*, 24(1):1–16, 2010.

[33] M. Krivelevich and B. Sudakov. Sparse pseudo-random graphs are Hamiltonian. *J. Graph Theory*, 42(1):17–33, 2003.

[34] C. Lee and W. Samotij. Pancyclic subgraphs of random graphs. *J. Graph Theory*, 71(2):142–158, 2012.

[35] O. Ore. Note on Hamilton circuits. *Amer. Math. Monthly*, 67:55, 1960.

[36] L. Pósa. Hamiltonian circuits in random graphs. *Discrete Math.*, 14(4):359–364, 1976.