Hamiltonicity of random subgraphs of the hypercube*

Padraig Condon† Alberto Espuny Díaz‡ António Girão† Daniela Kühn†
Deryk Osthus†

20th October 2020

Abstract

We study Hamiltonicity in random subgraphs of the hypercube $\mathcal{Q}^n$. Our first main theorem is an optimal hitting time result. Consider the random process which includes the edges of $\mathcal{Q}^n$ according to a uniformly chosen random ordering. Then, with high probability, as soon as the graph produced by this process has minimum degree $2k$, it contains k edge-disjoint Hamilton cycles, for any fixed $k \in \mathbb{N}$. Secondly, we obtain a perturbation result: if $H \subseteq \mathcal{Q}^n$ satisfies $\delta(H) \geq \alpha n$ with $\alpha > 0$ fixed and we consider a random binomial subgraph $\mathcal{Q}^n_p$ of $\mathcal{Q}^n$ with $p \in (0,1]$ fixed, then with high probability $H \cup \mathcal{Q}^n_p$ contains k edge-disjoint Hamilton cycles, for any fixed $k \in \mathbb{N}$. In particular, both results resolve a long standing conjecture, posed e.g. by Bollobás, that the threshold probability for Hamiltonicity in the random binomial subgraph of the hypercube equals $1/2$. Our techniques also show that, with high probability, for all fixed $p \in (0,1]$ the graph $\mathcal{Q}^n_p$ contains an almost spanning cycle. Our methods involve branching processes, the Rödl nibble, and absorption.

1 Introduction and results

The n-dimensional *hypercube* $\mathcal{Q}^n$ is the graph whose vertex set consists of all n-bit 01-strings, where two vertices are joined by an edge whenever their corresponding strings differ by a single bit. The hypercube and its subgraphs have attracted much attention in graph theory and computer science, e.g. as a sparse network model with strong connectivity properties. It is well known that hypercubes contain spanning paths (also called *Gray codes* or *Hamilton paths*) and, for all $n \geq 2$, they contain spanning cycles (also referred to as *cyclic Gray codes* or *Hamilton cycles*). Classical applications of Gray codes in computer science are described in the surveys of Savage [46] and Knuth [34]. Applications of hypercubes to parallel computing are discussed in the monograph of Leighton [43].

1.1 Spanning subgraphs in hypercubes.

The systematic study of spanning paths, trees and cycles in hypercubes was initiated in the 1970s. There is by now an extensive literature about subtrees of the hypercube; see, for instance, results of Bhatt, Chung, Leighton and Rosenberg [5] about embedding subdivided trees (instigated by processor allocation in distributed computing systems).

As a generalization of Hamilton paths, Caha and Koubek [18] considered the problem of finding a collection of spanning vertex-disjoint paths, given a prescribed set of endpoints. After several improvements [20, 29], this problem was recently resolved by Dvořák, Gregor and Koubek [23].

The applications of hypercubes as networks in computer science inspired questions about the reliability of its properties. This led to considering 'faulty' hypercubes in which some edges or vertices are missing. For instance, Chan and Lee [19] showed that, if $\mathcal{Q}^n$ has at most $2n-5$ faulty edges and every vertex has (non-faulty) degree at least 2, then there is a Hamilton cycle in $\mathcal{Q}^n$ which avoids all faulty edges (and this condition is best possible). They also showed that the general problem of determining the Hamiltonicity of $\mathcal{Q}^n$ with a larger number of faulty edges is NP-complete. More generally, Dvořák and Gregor [22] studied the existence of spanning collections of vertex-disjoint paths with prescribed endpoints in faulty hypercubes. These can be seen as extremal results about the *robustness* of the hypercube with respect to containing spanning collections of paths and cycles.

1.2 Hamilton cycles in binomial random graphs.

One of the most studied random graph models is the binomial random graph $G_{n,p}$. Here we have a

*This project has received partial funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 786198, P. Condon, A. Espuny Díaz, D. Kühn and D. Osthus). The research leading to these results was also partially supported by the EPSRC, grant nos. EP/N019504/1 (A. Girão and D. Kühn) and EP/S00100X/1 (D. Osthus), as well as the Royal Society and the Wolfson Foundation (D. Kühn).

†University of Birmingham.
‡Technische Universität Ilmenau.

(labelled) set of n vertices and we include each edge with probability p independently of all other edges.

Given some monotone increasing graph property $\mathcal{P}$, a function $p^* = p^*(n)$ is said to be a (coarse) *threshold* for $\mathcal{P}$ if $\mathbb{P}[G_{n,p} \in \mathcal{P}] \to 1$ whenever $p/p^* \to \infty$ and $\mathbb{P}[G_{n,p} \in \mathcal{P}] \to 0$ whenever $p/p^* \to 0$. One can define the stronger notion of a sharp threshold similarly: $p^* = p^*(n)$ is said to be a *sharp threshold* for $\mathcal{P}$ if, for all $\varepsilon > 0$, we have that $\mathbb{P}[G_{n,p} \in \mathcal{P}] \to 1$ whenever $p \geq (1+\varepsilon)p^*$ and $\mathbb{P}[G_{n,p} \in \mathcal{P}] \to 0$ whenever $p \leq (1-\varepsilon)p^*$. The problem of finding the threshold for the containment of a Hamilton cycle was solved independently by Pósa [45] and Koršunov [37]. Furthermore, Koršunov [37] determined the sharp threshold for Hamiltonicity to be $p^* = \log n/n$. These results were later made even more precise by Komlós and Szemerédi [36]. It is worth noting that $p^* = \log n/n$ is also the sharp threshold for the property of having minimum degree at least 2. In this sense, the results about Hamilton cycles in $G_{n,p}$ can be interpreted as saying that the natural obstruction of having sufficiently high minimum degree is also an 'almost sufficient' condition.

A property that generalises Hamiltonicity is that of containing k edge-disjoint Hamilton cycles, for some $k \in \mathbb{N}$. We will present more results in this direction in Section 1.4; for now, let us simply note that the sharp threshold for the containment of k edge-disjoint Hamilton cycles in $G_{n,p}$, for some $k \in \mathbb{N}$ independent of n, is $p^* = \log n/n$, i.e., the same as the threshold for Hamiltonicity.

The study of *robustness* of graph properties has also attracted much attention recently. For instance, given a graph G which is known to satisfy some property $\mathcal{P}$, consider a random subgraph G_p obtained by deleting each edge of G with probability $1-p$, independently of all other edges. The problem then is to determine the range of p for which G_p satisfies $\mathcal{P}$ with high probability. In this setting, a result of Krivelevich, Lee and Sudakov [39] asserts that, for any n-vertex graph G with minimum degree at least $n/2$, the graph G_p is asymptotically almost surely Hamiltonian whenever $p \gg \log n/n$. This can be viewed as a robust version of Dirac's theorem on Hamilton cycles.

1.3 Hamilton cycles in binomial random subgraphs of the hypercube. Throughout this paper, we will consider random subgraphs of the hypercube and show that the hypercube is robustly Hamiltonian in the above sense. We will denote by $\mathcal{Q}_p^n$ the random subgraph of the hypercube obtained by removing each edge of $\mathcal{Q}^n$ with probability $1-p$ independently of every other edge.

The random graph $\mathcal{Q}_p^n$ was first studied by Burtin [17], who proved that the sharp threshold for connectivity is $1/2$. This result was later made more precise by Erdős and Spencer [25] and Bollobás [7]. As a related problem, Dyer, Frieze and Foulds [24] determined the sharp threshold for connectivity in subgraphs of $\mathcal{Q}^n$ obtained by removing both vertices and edges uniformly at random. Later, Bollobás [9] proved that $1/2$ is also the sharp threshold for the containment of a perfect matching in $\mathcal{Q}_p^n$. As with the $G_{n,p}$ model, this also coincides with the threshold for having minimum degree at least 1.

The main goal of this paper is to study the analogous problem for Hamiltonicity in random subgraphs of the hypercube. There is a folklore conjecture that the sharp threshold for Hamiltonicity in $\mathcal{Q}_p^n$ should be $1/2$, i.e., the same as the threshold for having minimum degree at least 2. This question was explicitly asked by Bollobás [10] at several conferences in the 1980s, in the ICM surveys of Frieze [27] and Kühn and Osthus [41], as well as in the recent survey of Frieze [28]. A special case of our first result resolves this problem.

THEOREM 1.1. *For any $k \in \mathbb{N}$, the sharp threshold for the property of containing k edge-disjoint Hamilton cycles in $\mathcal{Q}_p^n$ is $p^* = 1/2$.*

For $k = 1$, this can be seen as a probabilistic version of the result on faulty hypercubes [19], and also as a statement about the robustness of Hamiltonicity in the hypercube.

While, for $p < 1/2$, with high probability $\mathcal{Q}_p^n$ will not contain a Hamilton cycle, it turns out that the reason for this is mostly due to local obstructions (e.g., vertices with degree zero or one). More precisely, we prove that, for any constant $p \in (0, 1/2)$, a.a.s. the random graph $\mathcal{Q}_p^n$ contains an almost spanning cycle.

THEOREM 1.2. *For any $\delta, p \in (0,1]$, a.a.s. the graph $\mathcal{Q}_p^n$ contains a cycle of length at least $(1-\delta)2^n$.*

We believe that the probability bound is far from optimal, in the sense that random subgraphs of the hypercube where edges are picked with vanishing probability should also satisfy this property.

CONJECTURE 1.1. *Suppose that $p = p(n)$ satisfies that $pn \to \infty$. Then, a.a.s. $\mathcal{Q}_p^n$ contains a cycle of length $(1-o(1))2^n$.*

Similarly, it would be interesting to determine which (long) paths and (almost spanning) trees can be found in $\mathcal{Q}_p^n$. Moreover, our methods might also be useful to embed other large subgraphs, such as F-factors.

CONJECTURE 1.2. *Suppose $\varepsilon > 0$ and an integer $\ell \geq 2$ are fixed and $p \geq 1/2 + \varepsilon$. Then, a.a.s. $\mathcal{Q}^n_p$ contains a C_{2^ℓ}-factor, that is, a set of vertex-disjoint cycles of length 2^ℓ which together contain all vertices of $\mathcal{Q}^n$.*

1.4 Hitting time results. Remarkably, the above intuition that having the necessary minimum degree is an 'almost sufficient' condition for the containment of edge-disjoint perfect matchings and Hamilton cycles can be strengthened greatly via so-called hitting time results. These are expressed in terms of random graph processes. The general setting is as follows. Let G be an n-vertex graph with $m = m(n)$ edges, and consider an arbitrary labelling $E(G) = \{e_1, \dots, e_m\}$. The *$G$-process* is defined as a random sequence of nested graphs $\tilde{G}(\sigma) = (G_t(\sigma))_{t=0}^m$, where σ is a permutation of $\{1, \dots, m\}$ chosen uniformly at random and, for each $i \in \{0, \dots, m\}$, we set $G_i(\sigma) = (V(G), E_i)$, where $E_i \coloneqq \{e_{\sigma(j)} : j \in \{1, \dots, i\}\}$. Given any monotone increasing graph property $\mathcal{P}$ such that $G \in \mathcal{P}$, the *hitting time* for $\mathcal{P}$ in the above G-process is the random variable $\tau_{\mathcal{P}}(\tilde{G}(\sigma)) \coloneqq \min\{t \in \{0, \dots, m\} : G_t(\sigma) \in \mathcal{P}\}$.

Let us denote the properties of containing a perfect matching by $\mathcal{PM}$, Hamiltonicity by $\mathcal{HAM}$, and connectivity by $\mathcal{CON}$, respectively. For any $k \in \mathbb{N}$, let $\boldsymbol{\delta} k$ denote the property of having minimum degree at least k, and let $\mathcal{HM}k$ denote the property of containing $\lfloor k/2 \rfloor$ edge-disjoint Hamilton cycles and, if k is odd, one matching of size $\lfloor n/2 \rfloor$ which is edge-disjoint from these Hamilton cycles. With this notion of hitting times, many of the results about thresholds presented in Sections 1.2 and 1.3 can be strengthened significantly. For instance, Bollobás and Thomason [13] showed that a.a.s. $\tau_{\mathcal{CON}}(\tilde{K}_n(\sigma)) = \tau_{\boldsymbol{\delta}1}(\tilde{K}_n(\sigma))$ and, if n is even, then a.a.s. $\tau_{\mathcal{PM}}(\tilde{K}_n(\sigma)) = \tau_{\boldsymbol{\delta}1}(\tilde{K}_n(\sigma))$. Ajtai, Komlós and Szemerédi [1] and Bollobás [8] independently proved that a.a.s. $\tau_{\mathcal{HAM}}(\tilde{K}_n(\sigma)) = \tau_{\boldsymbol{\delta}2}(\tilde{K}_n(\sigma))$. This was later generalised by Bollobás and Frieze [11], who proved that, given any $k \in \mathbb{N}$, for n even a.a.s. $\tau_{\mathcal{HM}k}(\tilde{K}_n(\sigma)) = \tau_{\boldsymbol{\delta}k}(\tilde{K}_n(\sigma))$.

A hitting time result for the property of having k edge-disjoint Hamilton cycles when k is allowed to grow with n is still not known, even in K_n-processes. As a slightly weaker notion, consider property $\mathcal{H}$, where we say that a graph G satisfies property $\mathcal{H}$ if it contains $\lfloor \delta(G)/2 \rfloor$ edge-disjoint Hamilton cycles, together with an additional edge-disjoint matching of size $\lfloor n/2 \rfloor$ if $\delta(G)$ is odd. Knox, Kühn and Osthus [33], Krivelevich and Samotij [40] as well as Kühn and Osthus [42] proved results for different ranges of p which, together, show that $G_{n,p}$ a.a.s. satisfies property $\mathcal{H}$.

For graphs other than the complete graph, Johansson [32] recently obtained a robustness version of the hitting time results for Hamiltonicity. In particular, for any n-vertex graph G with $\delta(G) \geq (1/2+\varepsilon)n$, he proved that a.a.s. $\tau_{\mathcal{HAM}}(\tilde{G}(\sigma)) = \tau_{\boldsymbol{\delta}2}(\tilde{G}(\sigma))$. This was later extended to a larger class of graphs G and to hitting times for $\mathcal{HM}2k$, for all $k \in \mathbb{N}$ independent of n, by Alon and Krivelevich [2].

In the setting of random subgraphs of the hypercube, Bollobás [9] determined the hitting time for perfect matchings by showing that a.a.s. $\tau_{\mathcal{PM}}(\tilde{\mathcal{Q}}^n(\sigma)) = \tau_{\mathcal{CON}}(\tilde{\mathcal{Q}}^n(\sigma)) = \tau_{\boldsymbol{\delta}1}(\tilde{\mathcal{Q}}^n(\sigma))$. One of our main results (which implies Theorem 1.1) is a hitting time result for Hamiltonicity (and, more generally, property $\mathcal{HM}k$) in $\mathcal{Q}^n$-processes. Again, this question was raised by Bollobás [10] at several conferences.

THEOREM 1.3. *For all $k \in \mathbb{N}$, a.a.s. $\tau_{\mathcal{HM}k}(\tilde{\mathcal{Q}}^n(\sigma)) = \tau_{\boldsymbol{\delta}k}(\tilde{\mathcal{Q}}^n(\sigma))$, that is, the hitting time for the containment of a collection of $\lfloor k/2 \rfloor$ Hamilton cycles and $k - 2\lfloor k/2 \rfloor$ perfect matchings, all pairwise edge-disjoint, in $\mathcal{Q}^n$-processes is a.a.s. equal to the hitting time for the property of having minimum degree at least k.*

We also wonder whether this is true if k is allowed to grow with n, and propose the following conjecture which, if true, would be an approximate version of the results of [33, 40, 42] in the hypercube.

CONJECTURE 1.3. *For all $p \in (1/2, 1]$ and $\eta > 0$, a.a.s. $\mathcal{Q}^n_p$ contains $(1/2-\eta)\delta(\mathcal{Q}^n_p)$ edge-disjoint Hamilton cycles.*

1.5 Randomly perturbed graphs. A relatively recent area at the interface of extremal combinatorics and random graph theory is the study of *randomly perturbed graphs*. Generally speaking, the idea is to consider a deterministic dense n-vertex graph H (usually satisfying some minimum degree condition) and a random graph $G_{n,p}$ on the same vertex set as H. The question is whether H is close to satisfying some given property $\mathcal{P}$ in the sense that a.a.s. $H \cup G_{n,p} \in \mathcal{P}$ for some small p. This line of research was sparked off by Bohman, Frieze and Martin [6], who showed that, if H is an n-vertex graph with $\delta(H) \geq \alpha n$, for any constant $\alpha > 0$, then a.a.s. $H \cup G_{n,p}$ is Hamiltonian for all $p \geq C(\alpha)/n$. Other properties that have been studied in this context are e.g. the existence of powers of Hamilton cycles and general bounded degree spanning graphs [16], F-factors [3] or spanning bounded degree trees [15, 38]. One common phenomenon in this model is that, by considering the union with a dense graph H (i.e., a graph H with linear degrees), the probability threshold of different properties is significantly lower than that in the classical $G_{n,p}$ model. The results for Hamiltonicity [6] were very recently generalised by Hahn-Klimroth, Maesaka, Mogge, Mohr and Parczyk [30] to allow α to tend to 0 with n (that is, to allow graphs H which are not dense).

We consider randomly perturbed graphs in the setting of subgraphs of the hypercube. To be precise, we take an arbitrary spanning subgraph H of the hypercube, with linear minimum degree, and a random subgraph $\mathcal{Q}^n_\varepsilon$, and consider $H \cup \mathcal{Q}^n_\varepsilon$. Note that the minimum degree is required to be linear with respect to the dimension n (and not with respect to the number of vertices, as in the aforementioned results). Note also that $\mathcal{Q}^n_\varepsilon$ is a 'dense' subgraph of $\mathcal{Q}^n$, but for $\varepsilon < 1/2$ it will contain both isolated vertices and vertices of very low degrees. In this setting, we show the following result.

THEOREM 1.4. *For all $\varepsilon, \alpha \in (0,1]$ and $k \in \mathbb{N}$, the following holds. Let H be a spanning subgraph of $\mathcal{Q}^n$ such that $\delta(H) \geq \alpha n$. Then, a.a.s. $H \cup \mathcal{Q}^n_\varepsilon$ contains k edge-disjoint Hamilton cycles.*

We can also allow H to have much smaller degrees, at the cost of requiring a larger probability to find the Hamilton cycles.

THEOREM 1.5. *For every integer $k \geq 2$, there exists $\varepsilon > 0$ such that a.a.s., for every spanning subgraph H of $\mathcal{Q}^n$ with $\delta(H) \geq k$, the graph $H \cup \mathcal{Q}^n_{1/2-\varepsilon}$ contains a collection of $\lfloor k/2 \rfloor$ Hamilton cycles and $k - 2\lfloor k/2 \rfloor$ perfect matchings, all pairwise edge-disjoint.*

Note that Theorem 1.5 can be viewed as a 'universality' result for H, meaning that it holds for all choices of H simultaneously. It would be interesting to know whether such a result can also be obtained for the lower edge probability assumed in Theorem 1.4, i.e., is it the case that, for all $\varepsilon, \alpha \in (0,1]$, a.a.s. $G \sim \mathcal{Q}^n_\varepsilon$ has the property that, for every spanning $H \subseteq \mathcal{Q}^n$ with $\delta(H) \geq \alpha n$, $G \cup H$ is Hamiltonian?

Theorem 1.1 follows straightforwardly from Theorem 1.4, and it follows trivially from Theorem 1.3. In turn, Theorem 1.3 follows from Theorem 1.5. On the other hand, Theorems 1.2, 1.4 and 1.5, while being proved with similar ideas, are incomparable.

1.6 Percolation on the hypercube. To build Hamilton cycles in random subgraphs of the hypercube, we will consider a random process which can be viewed as a branching process or percolation process on the hypercube. With high probability, for constant $p > 0$, this process results in a bounded degree tree in $\mathcal{Q}^n_p$ which covers most of the neighbourhood of every vertex in $\mathcal{Q}^n$, and thus spans almost all vertices of $\mathcal{Q}^n$. The version stated below is a special case of our main result in this setting.

THEOREM 1.6. *For any fixed $\varepsilon, p \in (0,1]$, there exists $D = D(\varepsilon)$ such that a.a.s. $\mathcal{Q}^n_p$ contains a tree T with $\Delta(T) \leq D$ and such that $|V(T) \cap N_{\mathcal{Q}^n}(x)| \geq (1-\varepsilon)n$ for every $x \in V(\mathcal{Q}^n)$.*

Further results concerning the local geometry of the giant component in $\mathcal{Q}^n_p$ for constant $p \in (0, 1/2)$ were proved recently by McDiarmid, Scott and Withers [44].

The random process we consider in the proof of Theorem 1.6 can be viewed as a branching random walk (with a bounded number of branchings at each step). Simpler versions of such processes (with infinite branchings allowed) have been studied by Fill and Pemantle [26] and Kohayakawa, Kreuter and Osthus [35], and we will base our analysis on these. Motivated by our approach, we raise the following question, which seems interesting in its own right.

QUESTION 1.4. *Does a non-returning random walk on $\mathcal{Q}^n$ a.a.s. visit almost all vertices of $\mathcal{Q}^n$?*

More generally, there are many results and applications concerning random walks on the hypercube (but allowing for returns). For example, motivated by a processor allocation problem, Bhatt and Cai [4] studied a walk algorithm to embed large (subdivided) trees into the hypercube. Moreover, the analysis of (branching) random walks is a critical ingredient in the study of percolation thresholds for the existence of a giant component in $\mathcal{Q}^n_p$. These have been investigated e.g. by Bollobás, Kohayakawa and Łuczak [12], Borgs, Chayes, van der Hofstad, Slade and Spencer [14] and van der Hofstad and Nachmias [31].

2 Outline of the main proofs

2.1 Overall outline. The proofs of our main results are quite involved, and thus we cannot present them here. Instead, we now sketch the key ideas for the proofs. Full details can be found in [21].

We will first focus on the proof of Theorem 1.4. We will first prove the case $k = 1$, and later use this to deduce the case when $k > 1$. Recall we are given $H \subseteq \mathcal{Q}^n$ with $\delta(H) \geq \alpha n$, and $G \sim \mathcal{Q}^n_\varepsilon$, with $\alpha, \varepsilon \in (0,1]$. Our aim is to show that a.a.s. $H \cup G$ is Hamiltonian.

Our approach for finding a Hamilton cycle is to first obtain a spanning tree. By passing along all the edges of a spanning tree T (with a vertex ordering prescribed by a depth first search), one can create a closed spanning walk W which visits every edge of T twice. The idea is then to modify such a walk into a Hamilton cycle. (This approach is inspired by the approximation algorithm for the Travelling Salesman Problem which returns a tour of at most twice the optimal length.) More precisely, our approach will be to obtain a near-spanning tree of $\mathcal{Q}^{n-s}$, for some suitable constant s, and to blow up vertices of this tree into s-dimensional cubes (see Figure 1). These

cubes can then be used to move along the tree without revisiting vertices, which will result in a near-Hamilton cycle $\mathfrak{H}$. All remaining vertices which are not included in $\mathfrak{H}$ will be absorbed into $\mathfrak{H}$ via absorbing structures that we carefully put in place beforehand.

In Sections 2.2 to 2.4 we outline in more detail how we find a long cycle in G (Theorem 1.2). Note that in Theorem 1.2 we have $G \sim \mathcal{Q}^n_\varepsilon$, so a.a.s. G will have isolated vertices which prevent any Hamilton cycle occurring as a subgraph. In Section 2.5 we outline how we build on this approach to obtain the case $k = 1$ of Theorem 1.4. In Section 2.6 we sketch how we obtain Theorem 1.5 and, thus, Theorem 1.3.

2.2 Building block I: trees via branching processes. We view each vertex in $\mathcal{Q}^n$ as an n-dimensional 01-coordinate vector. By fixing the first s coordinates, we fix one of 2^s *layers* $L_1, \dots, L_{2^s}$ of the hypercube, where $s \in \mathbb{N}$ will be constant. Thus, $L \cong \mathcal{Q}^{n-s}$ for each layer L. By considering a Hamilton cycle in $\mathcal{Q}^s$, we may assume that consecutive layers differ only by a single coordinate on the unique elements of $\mathcal{Q}^s$ which define them. Let $G \sim \mathcal{Q}^n_\varepsilon$. For each layer L, we let $L(G) \coloneqq G[V(L)]$ and, by momentarily viewing these layers as different subgraphs on the vertex set of $\mathcal{Q}^{n-s}$, we define the *intersection graph* $I(G) \coloneqq \bigcap_{i=1}^{2^s} L_i(G)$. Hence, $I(G) \sim \mathcal{Q}^{n-s}_{\varepsilon^{2^s}}$. We view $I(G)$ as a subgraph of $\mathcal{Q}^{n-s}$. We first show that $I(G)$ contains a near-spanning tree T (see Theorem 2.1 below). Thus, a copy of T is present in each of $L_1(G), \dots, L_{2^s}(G)$ simultaneously.

Since the walk W mentioned in Section 2.1 passes through each vertex x of T a total of $d_T(x)$ times, it will be important later for T to have bounded degree. In order to guarantee this, we run bounded degree branching processes from several far apart 'corners' of the hypercube. Roughly speaking, T will be formed by taking a union of these processes and removing cycles. Crucially, the model we introduce for these processes has a joint distribution with $\mathcal{Q}^{n-s}_{\varepsilon^{2^s}}$, so that T will in fact appear as a subgraph of $I(G)$. In applying our main result about trees in $\mathcal{Q}^n_p$, we obtain a bounded degree tree $T \subseteq I(G)$ which contains almost all of the neighbours of every vertex of $I(G)$. We also obtain a 'small' *reservoir* set $R \subseteq V(I(G))$, which T avoids and which will play a key role later in the absorption of vertices which do not belong to our initial long cycle. At this point, both T and R are now present in every layer of the hypercube simultaneously.

The precise statement of our result is given below. Here, given two vertices $x, y \in V(\mathcal{Q}^n)$ and some $k \in \mathbb{N}$, $\mathrm{dist}(x, y)$ represents the distance between x and y, and $B^k_{\mathcal{Q}^n}(x)$ is the set of vertices $v \in V(\mathcal{Q}^n)$ whose distance to x is at most k. The same notation is used for sets of vertices. Furthermore, given a graph G and $\delta \in [0, 1]$, we let $Res(G, \delta)$ be a probability distribution on subsets of $V(G)$, where $R \sim Res(G, \delta)$ is obtained by adding each vertex $v \in V(G)$ to R with probability δ, independently of every other vertex.

THEOREM 2.1. *Let $0 < 1/D, \delta \ll \varepsilon' \leq 1/2$, and let $\varepsilon, \gamma \in (0, 1]$ and $k \in \mathbb{N}$. Then, the following holds a.a.s. Let $\mathcal{A} \subseteq V(\mathcal{Q}^n)$ with the following two properties:*

(P1) *for any distinct $x, y \in \mathcal{A}$ we have $\mathrm{dist}(x, y) \geq \gamma n$, and*

(P2) *if we let $x_1 \coloneqq \{0\}^n$, $x_2 \coloneqq \{1\}^n$, $x_3 \coloneqq \{1\}^{\lceil n/2 \rceil}\{0\}^{n-\lceil n/2 \rceil}$ and $x_4 \coloneqq \{0\}^{\lceil n/2 \rceil}\{1\}^{n-\lceil n/2 \rceil}$, then $B^{k+2}_{\mathcal{Q}^n}(\mathcal{A}) \cap \{x_1, x_2, x_3, x_4\} = \varnothing$.*

Let $R \sim Res(\mathcal{Q}^n, \delta)$. Then, there exists a tree $T \subseteq \mathcal{Q}^n_\varepsilon - (R \cup B^k_{\mathcal{Q}^n}(\mathcal{A}))$ such that

(T1) *$\Delta(T) < D$,*

(T2) *for all $x \in V(\mathcal{Q}^n) \setminus B^k_{\mathcal{Q}^n}(\mathcal{A})$, we have that*

$$|N_{\mathcal{Q}^n}(x) \cap V(T)| \geq (1 - \varepsilon')n.$$

2.3 Building block II: cube tilings via the nibble. Let $\ell < s$ be fixed (in the proof, these values depend on several other parameters, but are independent of n; for simplicity, here we think of them as large constants, both much larger than the maximum degree of the tree T). In order to gain more local flexibility when traversing the near-spanning tree T, we augment T by locally adding a near-spanning ℓ-cube factor of $I(G)$. One can use classical results on matchings in almost regular uniform hypergraphs of small codegree to show that $I(G)$ contains such a collection of $\mathcal{Q}^\ell$ spanning almost all vertices of $I(G)$. However, we require the following stronger properties, namely that there exists a collection $\mathcal{C}$ of vertex disjoint copies of $\mathcal{Q}^\ell$ in $I(G)$ so that, for each $x \in V(I(G))$,

(i) $\mathcal{C}$ covers almost all vertices in $N_{I(\mathcal{Q}^n)}(x)$;

(ii) the directions spanned by the cubes intersecting $N_{I(\mathcal{Q}^n)}(x)$ do not correlate too strongly with any given set of directions.

The precise statement is given below. Neither (i) nor (ii) follow from existing results on hypergraph matchings and the proofs strongly rely on geometric properties intrinsic to the hypercube.

In the following, $\mathcal{D}(\mathcal{Q}^n)$ is the set of 01-strings of length n which contain exactly one 1 (these are the *directions* of the hypercube). Given any collection $\mathcal{C}$ of vertex disjoint copies of $\mathcal{Q}^\ell$ in $I(G)$, a vertex $x \in V(\mathcal{Q}^n)$

and any $Y \subseteq N_{\mathcal{Q}^n}(x)$, we denote $\mathcal{C}_x(Y) := \{C \in \mathcal{C} : \mathrm{dist}(x, C) = 1, V(C) \cap Y \neq \varnothing\}$. In particular, we write $\mathcal{C}_x := \mathcal{C}_x(N_{\mathcal{Q}^n}(x))$. Finally, given any $S \subseteq \mathcal{D}(\mathcal{Q}^n)$ and any positive $t \in \mathbb{R}$, we denote $\Sigma(\mathcal{C}, S, t) := \{C \in \mathcal{C} : |\mathcal{D}(C) \cap S| \geq t\}$.

THEOREM 2.2. *Let $\varepsilon, \delta, \alpha, \beta \in (0,1)$ and $K, \ell \in \mathbb{N}$ be such that $1/\ell \ll \alpha \ll \beta$. For each $x \in \{0,1\}^n$, let $A_0(x) := N_{\mathcal{Q}^n}(x)$ and, for each $i \in \{1, \dots, K\}$, let $A_i(x) \subseteq A_0(x)$ be a set of size $|A_i(x)| \geq \beta n$. Then, the graph $\mathcal{Q}^n_\varepsilon$ a.a.s. contains a collection $\mathcal{C}$ of vertex-disjoint copies of $\mathcal{Q}^\ell$ such that the following properties are satisfied for every $x \in \{0,1\}^n$:*

(M1) $|A_0(x) \cap V(\mathcal{C})| \geq (1-\delta)n$;

(M2) *for every $\hat{e} \in \mathcal{D}(\mathcal{Q}^n)$ we have*

$$|\Sigma(\mathcal{C}_x, \{\hat{e}\}, 1)| = o(n^{1/2});$$

(M3) *for every $i \in \{0, \dots, K\}$ and every $S \subseteq \mathcal{D}(\mathcal{Q}^n)$ with $\alpha n/2 \leq |S| \leq \alpha n$ we have*

$$|\Sigma(\mathcal{C}_x(A_i(x)), S, \ell^{1/2})| \geq |A_i(x)|/3000.$$

To prove that such a near-spanning ℓ-cube factor of $I(G)$ exists, we build on the so-called Rödl nibble. More precisely, we consider the hypergraph $\mathcal{H}$, with $V(\mathcal{H}) = V(\mathcal{Q}^{n-s})$, where the edge set is given by the copies of $\mathcal{Q}^\ell$ in $I(G)$. We run a random iterative process where at each stage we add a 'small' number of edges from $\mathcal{H}$ to $\mathcal{C}$, before removing all those remaining edges of $\mathcal{H}$ which 'clash' with our selection. A careful analysis and an application of the Lovász local lemma yield the existence of an instance of this process which terminates in the near-spanning ℓ-cube factor with the required properties.

2.4 Constructing a long cycle. Roughly speaking, we will use T as a backbone to provide 'global' connectivity, and will use the near-spanning ℓ-cube factor $\mathcal{C}$ and the layer structure to gain high 'local' connectivity and flexibility. We show a representation of this structure in Figure 1. Let $T \cup \bigcup_{C \in \mathcal{C}} C =: \Gamma' \subseteq I(G)$ and let $\Gamma \subseteq \Gamma'$ be formed by removing all leaves and isolated cubes in Γ'. It follows by our tree and nibble results that almost all vertices of $I(G)$ are contained in Γ. Note that, for each $v \in V(\mathcal{Q}^{n-s}) = V(I(G))$, there is a unique vertex in each of the 2^s layers which corresponds to v. We refer to these 2^s vertices as *clones* of v and to the collection of these 2^s clones as a *vertex molecule*. Similarly, each ℓ-cube $C \in \mathcal{C}$ contained in Γ gives rise to a *cube molecule* (that is, a collection of 2^s ℓ-cubes, one in each layer, all corresponding to the same ℓ-cube in $I(G)$). We construct

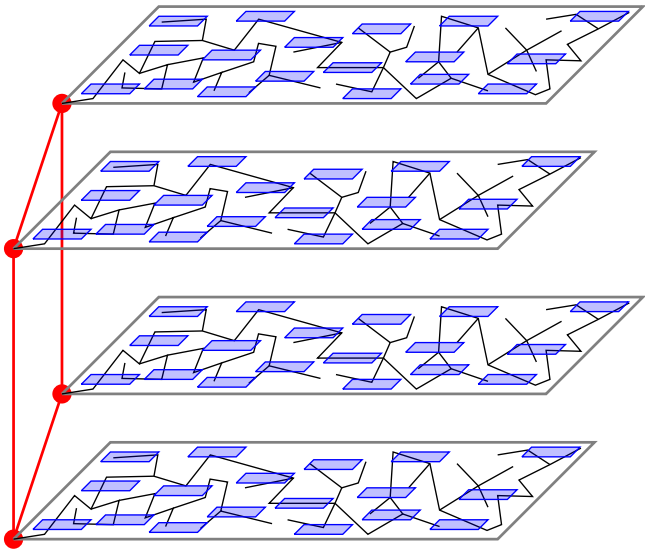

Figure 1: A representation of the main structure used for the proofs. We think of $\mathcal{Q}^n$ as a 'product' of two smaller cubes. Each 'horizontal' cube represents a copy of $\mathcal{Q}^{n-s}$ (a layer), and the red 'vertical' cube represents $\mathcal{Q}^s$. All 'horizontal' cubes contain a copy of the same tree T and the same cube tiling $\mathcal{C}$ (which are consistently distributed with respect to the 'vertical' cube; this gives rise to 'cube molecules'). When finding a long cycle, cube molecules are highly connected and can be covered by few paths, and the tree is used to join cube molecules to one another.

a cycle in G which covers all of the cube molecules (and, therefore, almost all vertices in $\mathcal{Q}^n$).

Let Γ^* be the graph obtained from Γ by contracting each ℓ-cube $C \subseteq \Gamma$ into a single vertex. We refer to such vertices in Γ^* as *atomic vertices*, and to all other vertices as *inner tree vertices*. We run a depth-first search on Γ^* to give an order to the vertices. Next, we construct a *skeleton* which will be the backbone for our long cycle. The skeleton is an ordered sequence of vertices in $\mathcal{Q}^n$ which contains the vertices via which our cycle will enter and exit each molecule. That is, given an *exit vertex* v for some molecule in the skeleton, the vertex u which succeeds v in the skeleton will be an *entry vertex* for another molecule, and such that $uv \in E(G)$. Here, a vertex in the skeleton belonging to an inner tree vertex molecule is referred to as both an entry and exit vertex. (Actually, we will first construct an 'external skeleton', which encodes this information. The skeleton then also prescribes some edges within molecules which go between different layers.) We use the ordering of the vertices of Γ^* to construct the skeleton in a recursive way starting from the lowest ordered vertex. It is crucial that our tree T has bounded degree (much smaller than 2^s), so that no molecule is overused in the skeleton.

Once the skeleton is constructed, we apply our 'connecting lemmas' [21, Lemmas 8.8 and 8.9], for whose

proofs we rely on results of Dvořák and Gregor [22]. These connecting lemmas, applied to a cube molecule with a bounded number of pairs of entry and exit vertices as input (given by the skeleton), provide us with a sequence of vertex-disjoint paths which cover this molecule, where each path has start and end vertices consisting of an input pair. The union of all of these paths combined with all edges in G between the successive exit and entry vertices of the skeleton will then form a cycle $\mathfrak{H} \subseteq G$ which covers all vertices lying in the cube molecules (thus proving Theorem 1.2).

2.5 Constructing a Hamilton cycle. In order to construct a Hamilton cycle in $H \cup G$, we will absorb the vertices of $V(\mathcal{Q}^n) \setminus V(\mathfrak{H})$ into $\mathfrak{H}$. We achieve this via absorbing structures that we identify for each vertex. To construct these absorbing structures, we will need to use some edges of H. Roughly speaking, to each vertex v we associate a left ℓ-cube $C_v^l \subseteq \mathcal{Q}^n$ and a right ℓ-cube $C_v^r \subseteq \mathcal{Q}^n$, where C_v^l, C_v^r are both clones of some ℓ-cubes $C^l, C^r \in \mathcal{C}$ contained in Γ. We choose these cubes so that v will have a neighbour $u \in V(C_v^l)$ and a neighbour $u' \in V(C_v^r)$, to which we refer as *tips* of the absorbing structure. Furthermore, u will have a neighbour $w \in V(C_v^r)$, which is also a neighbour of u'. Our near-Hamilton cycle $\mathfrak{H}$ will satisfy the following properties:

(a) $\mathfrak{H}$ covers all vertices in $C_v^l \cup C_v^r$ except for u, and

(b) $wu' \in E(\mathfrak{H})$.

These additional properties will be guaranteed by our connecting lemmas discussed in Section 2.4. We can then alter $\mathfrak{H}$ to include the segment $wuvu'$ instead of the edge wu', thus absorbing the vertices u and v into $\mathfrak{H}$. The following types of vertices will require absorption.

(i) Every vertex that is not covered by a clone of either some inner tree vertex or of some cube $C \in \mathcal{C}$ which is contained in Γ.

(ii) The cycle $\mathfrak{H}$ does not cover all the clones of inner tree vertices and, thus, the uncovered vertices of this type will also have to be absorbed.

However, we will not know precisely which of the vertices described in (ii) will be covered by $\mathfrak{H}$ and which of these vertices will need to be absorbed until after we have constructed the (external) skeleton. Moreover, many potential absorbing structures are later ruled out as candidates (for example, if they themselves contain vertices that will need to be absorbed). Therefore, it is important that we identify a 'robust' collection of many potential absorbing structures for every vertex in $\mathcal{Q}^n$ at a preliminary stage of the proof. The precise absorbing structure eventually assigned to each vertex will be chosen via an application of our rainbow matching lemma [21, Lemma 5.5] at a late stage in the proof.

We will now highlight the purpose of the reservoir R. Suppose $v \in V(\mathcal{Q}^n)$ is a vertex which needs to be absorbed via an absorbing structure with left ℓ-cube C_v^l and left tip $u \in V(C_v^l)$. Recall that both u and C_v^l are clones of some $u^* \in V(\Gamma)$ and $C^l \in \mathcal{C}$, where $u^* \in V(C^l)$. If u^* has a neighbour w^* in $T - V(C^l)$, then it is possible that the skeleton will assign an edge from u to w for the cycle $\mathfrak{H}$ (where w is the clone of w^* in the same layer as u). Given that u is now incident to a vertex outside of C_v^l, we can no longer use the absorbing structure with u as a (left) tip (otherwise, we might disconnect T). To avoid this problem, we show that most vertices have many potential absorbing structures whose tips lie in the reservoir R (which T avoids). Here we make use of vertex degrees of H. A small number of *scant vertices* will not have high enough degree into R. For these vertices we fix an absorbing structure whose tips do not lie in R, and then alter T slightly so that these tips are deleted from T and reassigned to R. The fact that scant vertices are few and well spread out from each other will be crucial in being able to achieve this (see [21, Lemma 7.20]).

Let us now discuss two problems arising in the construction of the skeleton. Firstly, let $\mathcal{M}_C \subseteq \mathcal{Q}^n$ with $C \in \mathcal{C}$ be a cube molecule which is to be covered by $\mathfrak{H}$. Furthermore, suppose one of the clones C_v^l of C belongs to an absorbing structure for some vertex v. Let u be the tip of C_v^l and suppose that u has even parity. We would like to apply the connecting lemmas to cover $\mathcal{M}_C - \{u\}$ by paths which avoid u. But this would now involve covering one fewer vertex of even parity than of odd parity. This, in turn, has the effect of making the construction of the skeleton considerably more complicated (this construction is simplest when successive entry and exit vertices have opposite parities). To avoid this, we assign absorbing structures in pairs, so that, for each $C \in \mathcal{C}$, either two or no clones of C will be used in absorbing structures. In the case where two clones are used, we enforce that the tips of these clones will have opposite parities, and therefore each molecule $\mathcal{M}_C$ will have the same number of even and odd parity vertices to be covered by $\mathfrak{H}$. We use our robust matching lemma (see [21, Lemma 5.2]) to pair up the clones of absorbing structures in this way. To connect up different layers of a cube molecule, we will of course need to have suitable edges between these. Molecules which do not satisfy this requirement are called 'bondless' and are removed from Γ before the absorption process (so that their vertices are absorbed).

Secondly, another issue related to vertex parities arises from inner tree vertex molecules. Depending

on the degree of an inner tree vertex $v \in V(T)$, the skeleton could contain an odd number of vertices from the molecule $\mathcal{M}_v$ consisting of all clones of v. All vertices in $\mathcal{M}_v$ outside the skeleton will need to be absorbed. But since the number of these vertices is odd, it would be impossible to pair up (in the way described above) the absorbing structures assigned to these vertices. To fix this issue, we effectively impose that $\mathfrak{H}$ will 'go around T twice'. That is, the skeleton will trace through every molecule beginning and finishing at the lowest ordered vertex in Γ^*. It will then retrace its steps through these molecules in an almost identical way, effectively doubling the size of the skeleton. This ensures that the skeleton contains an even number of vertices from each molecule, half of them of each parity.

Finally, once we have obtained an appropriate skeleton, we can construct a long cycle $\mathfrak{H}$ as described in Section 2.4. For every vertex in $\mathcal{Q}^n$ which is not covered by $\mathfrak{H}$ we have put in place an absorbing structure, which is covered by $\mathfrak{H}$ as described in (a) and (b). Thus, as discussed before, we can now use these structures to absorb all remaining vertices into $\mathfrak{H}$ to obtain a Hamilton cycle $\mathfrak{H}' \subseteq H \cup G$, thus proving the case $k = 1$ of Theorem 1.4.

2.6 Hitting time for the appearance of a Hamilton cycle. In order to prove Theorem 1.3, we consider $G \sim \mathcal{Q}^n_{1/2-\varepsilon}$. We show that a.a.s., for any graph H with $\delta(H) \geq 2$, the graph $G \cup H$ is Hamiltonian (Theorem 1.5). The main additional difficulty faced here is that $G \cup H$ may contain vertices having degree as low as 2. For the set $\mathcal{U}$ of these vertices we cannot hope to use the previous absorption strategy: the neighbours of $v \in \mathcal{U}$ may not lie in cubes from $\mathcal{C}$. (In fact, v may not even have a neighbour within its own layer in $G \cup H$.) To handle such small degree vertices, we first prove that they will be few and well spread out [21, Lemma 9.4]. We then define three types of new 'special absorbing structures' (see [21, Section 9.1]). The type of the special absorbing structure $SA(v)$ for v will depend on whether the neighbours a, b of v in H lie in the same layer as v. In each case, $SA(v)$ will consist of a short path P_1 containing the edges av and bv, and several other short paths designed to 'balance out' P_1 in a suitable way. These paths will be incorporated into the long cycle $\mathfrak{H}$ described in Section 2.4. In particular, this allows us to 'absorb' the vertices of $\mathcal{U}$ into $\mathfrak{H}$. To incorporate the paths P_i forming $SA(v)$, we will proceed as follows.

Firstly, we make use of the fact that Theorem 2.1 allows us to choose T in such a way that it avoids a small ball around each $v \in \mathcal{U}$. Thus, (all clones of) T will avoid $SA(v)$, which has the advantage that there will be no interference between T and the special absorbing structures. To link up each $SA(v)$ with the long cycle $\mathfrak{H}$, for each endpoint w of a path in $SA(v)$, we will choose an ℓ-cube in $I(G)$ which suitably intersects T and which contains w (or, more precisely, the vertex in $I(G)$ corresponding to w). Altogether, these ℓ-cubes allow us to find paths between $SA(v)$ and vertices of $\mathfrak{H}$ which are clones of vertices in T. The remaining vertices in molecules consisting of clones of these ℓ-cubes will be covered in a similar way as in Section 2.4. All vertices in these balls around $\mathcal{U}$ which are not part of the special absorbing structures will be absorbed into $\mathfrak{H}$ via the same absorbing structures used in the proof of Theorem 1.4 to once again obtain a Hamilton cycle $\mathfrak{H}'$.

2.7 Edge-disjoint Hamilton cycles. The results on k edge-disjoint Hamilton cycles can be deduced from suitable versions of the case $k = 1$. Those versions are carefully formulated to allow us to repeatedly remove a Hamilton cycle from the original graph. The following is the result which allows to prove Theorem 1.4. Theorem 1.5 follows from a similar but more technical statement [21, Theorem 9.6].

THEOREM 2.3. *For every $\varepsilon, \alpha \in (0, 1]$ and $c > 0$, there exists $\Phi \in \mathbb{N}$ such that the following holds. Let $H \subseteq \mathcal{Q}^n$ be a spanning subgraph with $\delta(H) \geq \alpha n$ and let $G \sim \mathcal{Q}^n_\varepsilon$. Then, a.a.s. there is a subgraph $G' \subseteq G$ with $\Delta(G') \leq \Phi$ such that, for every $F \subseteq \mathcal{Q}^n$ with $\Delta(F) \leq c\Phi$, the graph $((H \cup G) \setminus F) \cup G'$ is Hamiltonian.*

References

[1] M. Ajtai, J. Komlós and E. Szemerédi, First occurrence of Hamilton cycles in random graphs, *Cycles in graphs (Burnaby, B.C., 1982)*, vol. 115 of *North-Holland Math. Stud.*, 173–178, North-Holland, Amsterdam (1985).

[2] Y. Alon and M. Krivelevich, Hitting time of edge disjoint Hamilton cycles in random subgraph processes on dense base graphs, *arXiv e-prints* (2019), `1912.01251`.

[3] J. Balogh, A. Treglown and A. Z. Wagner, Tilings in randomly perturbed dense graphs, *Combin. Probab. Comput.* **28** (2019), 159–176.

[4] S. Bhatt and J.-Y. Cai, Take a walk, grow a tree, *29th Annual Symposium on Foundations of Computer Science (STOC)*, 469–478 (1988).

[5] S. N. Bhatt, F. R. K. Chung, F. T. Leighton and A. L. Rosenberg, Efficient embeddings of trees in hypercubes, *SIAM J. Comput.* **21** (1992), 151–162.

[6] T. Bohman, A. Frieze and R. Martin, How many random edges make a dense graph Hamiltonian?, *Random Structures Algorithms* **22** (2003), 33–42.

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

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

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

[10] B. Bollobás, personal communication (2020).

[11] B. Bollobás and A. M. Frieze, On matchings and Hamiltonian cycles in random graphs, *Random graphs '83 (Poznań, 1983)*, vol. 118 of *North-Holland Math. Stud.*, 23–46, North-Holland, Amsterdam (1985).

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

[13] B. Bollobás and A. Thomason, Random graphs of small order, *Random graphs '83 (Poznań, 1983)*, vol. 118 of *North-Holland Math. Stud.*, 47–97, North-Holland, Amsterdam (1985).

[14] C. Borgs, J. T. Chayes, R. van der Hofstad, G. Slade and J. Spencer, Random subgraphs of finite graphs. III. The phase transition for the n-cube, *Combinatorica* **26** (2006), 395–410.

[15] J. Böttcher, J. Han, Y. Kohayakawa, R. Montgomery, O. Parczyk and Y. Person, Universality for bounded degree spanning trees in randomly perturbed graphs, *Random Structures Algorithms* **55** (2019), 854–864.

[16] J. Böttcher, R. Montgomery, O. Parczyk and Y. Person, Embedding spanning bounded degree graphs in randomly perturbed graphs, *Mathematika* **66** (2020), 422–447.

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

[18] R. Caha and V. Koubek, Spanning multi-paths in hypercubes, *Discrete Math.* **307** (2007), 2053–2066.

[19] M. Y. Chan and S.-J. Lee, On the existence of Hamiltonian circuits in faulty hypercubes, *SIAM J. Discrete Math.* **4** (1991), 511–527.

[20] X.-B. Chen, Paired many-to-many disjoint path covers of the hypercubes, *Inform. Sci.* **236** (2013), 218–223.

[21] P. Condon, A. Espuny Díaz, A. Girão, D. Kühn and D. Osthus, Hamiltonicity of random subgraphs of the hypercube, *arXiv e-prints* (2020), 2007.02891.

[22] T. Dvořák and P. Gregor, Partitions of faulty hypercubes into paths with prescribed endvertices, *SIAM J. Discrete Math.* **22** (2008), 1448–1461.

[23] T. Dvořák, P. Gregor and V. Koubek, Generalized Gray codes with prescribed ends, *Theoret. Comput. Sci.* **668** (2017), 70–94.

[24] M. E. Dyer, A. M. Frieze and L. R. Foulds, On the strength of connectivity of random subgraphs of the n-cube, *Random graphs '85 (Poznań, 1985)*, vol. 144 of *North-Holland Math. Stud.*, 17–40, North-Holland, Amsterdam (1987).

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

[26] J. A. Fill and R. Pemantle, Percolation, first-passage percolation and covering times for Richardson's model on the n-cube, *Ann. Appl. Probab.* **3** (1993), 593–629.

[27] A. Frieze, Random structures and algorithms, *Proceedings of the International Congress of Mathematicians—Seoul 2014. Vol. 1*, 311–340, Kyung Moon Sa, Seoul (2014).

[28] A. Frieze, Hamilton cycles in random graphs: a bibliography, *arXiv e-prints* (2019), 1901.07139.

[29] P. Gregor and T. Dvořák, Path partitions of hypercubes, *Inform. Process. Lett.* **108** (2008), 402–406.

[30] M. Hahn-Klimroth, G. S. Maesaka, Y. Mogge, S. Mohr and O. Parczyk, Random perturbation of sparse graphs, *arXiv e-prints* (2020), 2004.04672.

[31] R. van der Hofstad and A. Nachmias, Hypercube percolation, *J. Eur. Math. Soc. (JEMS)* **19** (2017), 725–814.

[32] T. Johansson, On Hamilton cycles in Erdős-Rényi subgraphs of large graphs, *arXiv e-prints* (2018), 1811.03501.

[33] F. Knox, D. Kühn and D. Osthus, Edge-disjoint Hamilton cycles in random graphs, *Random Structures Algorithms* **46** (2015), 397–445.

[34] D. E. Knuth, *The Art of Computer Programming, Volume 4, Fascicle 2: Generating All Tuples and Permutations (Art of Computer Programming)*, Addison-Wesley Professional (2005).

[35] Y. Kohayakawa, B. Kreuter and D. Osthus, The length of random subsets of Boolean lattices, *Random Structures Algorithms* **16** (2000), 177–194.

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

[37] A. D. Koršunov, Solution of a problem of P. Erdős and A. Rényi on Hamiltonian cycles in undirected graphs, *Metody Diskretn. Anal.* **31** (1977), 17–56.

[38] M. Krivelevich, M. Kwan and B. Sudakov, Bounded-degree spanning trees in randomly perturbed graphs, *SIAM J. Discrete Math.* **31** (2017), 155–171.

[39] M. Krivelevich, C. Lee and B. Sudakov, Robust Hamiltonicity of Dirac graphs, *Trans. Amer. Math. Soc.* **366** (2014), 3095–3130.

[40] M. Krivelevich and W. Samotij, Optimal packings of Hamilton cycles in sparse random graphs, *SIAM J. Discrete Math.* **26** (2012), 964–982.

[41] D. Kühn and D. Osthus, Hamilton cycles in graphs and hypergraphs: an extremal perspective, *Proceedings of the International Congress of Mathematicians—Seoul 2014. Vol. IV*, 381–406, Kyung Moon Sa, Seoul (2014).

[42] D. Kühn and D. Osthus, Hamilton decompositions of regular expanders: applications, *J. Combin. Theory Ser. B* **104** (2014), 1–27.

[43] F. T. Leighton, *Introduction to parallel algorithms and architectures*, Morgan Kaufmann, San Mateo, CA (1992), arrays, trees, hypercubes.

[44] C. McDiarmid, A. Scott and P. Withers, The component structure of dense random subgraphs of the hypercube, *arXiv e-prints* (2018), `1806.06433`.

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

[46] C. Savage, A survey of combinatorial Gray codes, *SIAM Rev.* **39** (1997), 605–629.