Resilience of Perfect Matchings and Hamiltonicity in Random Graph Processes

Rajko Nenadov[*]   Angelika Steger[†]   Miloš Trujić[†,‡]

Abstract

Let $\{G_i\}$ be the random graph process: starting with an empty graph G_0 with n vertices, in every step $i \geq 1$ the graph G_i is formed by taking an edge chosen uniformly at random among the non-existing ones and adding it to the graph G_{i-1}. The classical 'hitting-time' result of Ajtai, Komlós, and Szemerédi, and independently Bollobás, states that asymptotically almost surely the graph becomes Hamiltonian as soon as the minimum degree reaches 2, that is if $\delta(G_i) \geq 2$ then G_i is Hamiltonian. We establish a resilience version of this result. In particular, we show that the random graph process almost surely creates a sequence of graphs such that for $m \geq (\frac{1}{6} + o(1))n \log n$ edges, the 2-core of the graph G_m remains Hamiltonian even after an adversary removes $(\frac{1}{2} - o(1))$-fraction of the edges incident to every vertex. A similar result is obtained for perfect matchings.

1   Introduction

The theory of random graphs originated in 1959 with the two seminal papers by Erdős and Rényi [13] and Gilbert [16]. It is now a well-established research area with many applications in theoretical computer science, statistical physics, and other branches of mathematics, cf. [6, 14, 19]. Typical questions in random graph theory traditionally concern the *existence* of certain (sub)structures. A more recent trend, introduced by Sudakov and Vu [33], is the study of their *resilience* properties. Formally, resilience of a graph G with respect to some property $\mathcal{P}$ is defined as follows:

Definition 1.1. Let $G = (V, E)$ be a graph. We say that G is α-*resilient* with respect to the property $\mathcal{P}$, for some $\alpha \in [0, 1]$, if for every spanning subgraph $H \subseteq G$ such that $\deg_H(v) \leq \alpha \deg_G(v)$ for every $v \in V$, we have $G - H \in \mathcal{P}$.

In other words, α-resilience implies that the property $\mathcal{P}$ cannot be destroyed even if an adversary is allowed to remove an (arbitrary) α-fraction of all edges incident to each vertex. For instance, Dirac's celebrated theorem [11] states that K_n is $(1/2)$-resilient with respect to containing a Hamilton cycle. The study of resilience of complete graphs is one of the central topics in extremal combinatorics (see, e.g. [9, 17, 20, 26] for some of the cornerstones).

In this paper we show that two seminal results of random graph theory hold in a resilient fashion. More precisely, we consider two famous 'hitting-time' results—one by Erdős and Rényi [12] and Bollobás and Thomason [7] regarding perfect matchings, and one by Ajtai, Komlós, and Szemerédi [1] and, independently, Bollobás [5] regarding Hamiltonicity. Let $n \in \mathbb{N}$ and consider

the following process: set G_0 to be an empty graph with n vertices and for each $i \in \{1, \ldots, \binom{n}{2}\}$ form the graph G_i by choosing an edge $e_i \notin G_{i-1}$ uniformly at random and adding it to G_{i-1}. In this way we obtain a sequence of nested graphs $\{G_i\}_{i=0}^{N}$ for $N = \binom{n}{2}$, where G_0 is an empty graph and G_N is a complete graph. It is an easy exercise to show that for every $m \in \{1, \ldots, N\}$ the graph G_m has the same distribution as the Erdős-Rényi random graph $G_{n,m}$, a graph chosen uniformly at random among all labelled graphs with n vertices and exactly m edges. A necessary condition for a graph to contain a perfect matching is that the minimum degree is at least 1 (assuming that n is even). Similarly, having minimum degree at least 2 is a necessary condition to be Hamiltonian. The aforementioned results show that, perhaps surprisingly, these conditions are also sufficient: as soon as $\delta(G_m) \geq 1$ the graph G_m contains a perfect matching, and as soon as $\delta(G_m) \geq 2$ the graph contains a Hamilton cycle, both asymptotically almost surely[1]. This kind of results are usually referred to as *hitting-time* results.

In this paper we show that, in fact, as soon as the minimum degree reaches 1 the graph does not only contain a perfect matching, but it does so *robustly*—meaning it is $(1/2 - o(1))$-resilient with respect to this property. Similarly, as soon as the minimum degree reaches 2 the graph G_m becomes robustly Hamiltonian, that is $(1/2 - o(1))$-resilient with respect to being Hamiltonian. To see why the parameter $1/2$ is asymptotically best possible, first recall that every graph G contains a partition $V_1 \cup V_2$ such that every vertex in one part has at least half of its neighbours in the other. In the case when such a partition is not balanced, deleting all the edges within V_1 and V_2 clearly prevents the resulting graph from having either a perfect matching or a Hamilton cycle. Otherwise, both parts have size $n/2$ and it is an easy exercise to show that if G is a random graph with $m \gg n$ edges, which will be the case here, at least one vertex in V_1 also has roughly half of its neighbours in V_1. Moving such a vertex to V_2 results in an unbalanced partition, and the rest of the argument is the same as before.

Even though the study of resilience properties of random graphs has attracted considerable attention over the last years, cf. e.g. [2, 8, 18, 23] and the recent survey [32], this is the first result which determines precisely when the random graph process becomes resilient with respect to some well-studied property.

1.1 Results

The classical hitting-time result for perfect matchings by Bollobás and Thomason [7], and a previously obtained result by Erdős and Rényi [12] which establishes the correspondence between minimum degree one and perfect matchings in a slightly weaker form, follow quite easily from well-known sufficient conditions for the existence of perfect matchings. On the other hand, the hitting time result for Hamiltonicity by Ajtai, Komlós, and Szemerédi [1] and Bollobás [5] is significantly more intricate and builds on a result of Pósa [31] and subsequent refinements.

The situation is similar for our results. While Hamiltonicity is more interesting and challenging (and our main result), we use the case of perfect matchings to demonstrate some of our main ideas and the general approach.

Perfect matchings. Our first result shows a bit more than discussed earlier: as long as the number of added edges is not too small, removing isolated vertices, an obvious obstacle in obtaining a perfect matching, results in a graph which is resilient with respect to containing a perfect matching[2].

---

[1]An event is said to hold asymptotically almost surely (a.a.s. for short) if the probability that it holds approaches 1 as $n \to \infty$.

[2]In order to avoid having to distinguish between an odd and an even number of vertices of the given graph, we use the term 'perfect matching' to indicate that the graph contains a matching that covers all but at most one of

Theorem 1.2. *Let $\varepsilon > 0$ be a constant and consider the random graph process $\{G_i\}$. Then a.a.s. for every $m \geq \frac{1+\varepsilon}{4} n \log n$ we have that the graph obtained from G_m by deleting all isolated vertices is $(1/2 - \varepsilon)$-resilient with respect to containing a perfect matching.*

This improves a result of Sudakov and Vu [33] who showed that a random graph $G_{n,m}$ is $(1/2-\varepsilon)$-resilient with respect to containing a perfect matching if $m \geq Cn \log n$, for some large constant $C(\varepsilon) > 100$. Note that the bound on m in Theorem 1.2 is asymptotically optimal. The reason being that for $m = \frac{1-\varepsilon}{4} n \log n$ there exist many *cherries* (see, e.g. [7]), pairs of vertices of degree 1 which have a common neighbour. Clearly, for each cherry only one vertex can be part of a matching and a perfect matching thus cannot exist. The standard non-resilience version of Theorem 1.2 with somewhat more precise lower bound on m was established in [7].

In fact, we obtain a slightly stronger resilience statement. That is, even if we allow the adversary to delete a graph H that contains almost all the edges incident to 'atypical' vertices (vertices whose degree deviates significantly from the average degree), the resulting graph still contains a perfect matching. The precise statement can be found in Section 3.

Theorem 1.2 also immediately implies a resilience version of the hitting-time result for perfect matchings.

Theorem 1.3. *Let $\varepsilon > 0$ be a constant. Consider the random graph process $\{G_i\}$ with an even number of vertices n and let $m_1 = \min\{m \colon \delta(G_m) = 1\}$ denote the step in which the last isolated vertex disappears. Then a.a.s. we have that G_{m_1} is $(1/2 - \varepsilon)$-resilient with respect to containing a perfect matching.*

Hamiltonicity. Prior to proving that the point when $\delta(G_m)$ becomes 2 coincides with the point when G_m is Hamiltonian, Komlós and Szemerédi [21] showed that if $m = \frac{n}{2}(\log n + \log \log n + c_n)$ then $G_{n,m}$ contains a Hamilton cycle if $c_n \to \infty$. On the other hand, if $c_n \to -\infty$ then it is a.a.s. not Hamiltonian (see [21]), precisely because of the existence of a vertex with degree less than 2. However, something can still be said for such a sparse graph, or rather its 2-*core*. Given a graph G we define the 2-*core* of G to be the graph obtained from G by successively removing vertices of degree at most 1. In other words, the 2-core of a graph G is its largest subgraph which satisfies the necessary minimum degree condition for being Hamiltonian. Łuczak [28] showed that, in the case of a random graph $G_{n,m}$, such subgraphs are indeed Hamiltonian as long as $m \geq \frac{1+\varepsilon}{6} n \log n$. Here we extend this by showing that, starting from that point, the 2-core of every graph G_m is resilient with respect to containing a Hamilton cycle. Moreover, the bound on m is asymptotically optimal, as a.a.s. the 2-core of $G_{n,m}$ is not Hamiltonian for $m \leq \frac{1-\varepsilon}{6} n \log n$ (see [28]).

Theorem 1.4. *Let $\varepsilon > 0$ be a constant and consider the random graph process $\{G_i\}$. Then a.a.s. for every $m \geq \frac{1+\varepsilon}{6} n \log n$ we have that the 2-core of G_m is $(1/2 - \varepsilon)$-resilient with respect to being Hamiltonian.*

Similarly to the case of perfect matchings we actually prove a slightly stronger resilience result that allows deletion of almost all edges incident to 'atypical' vertices, cf. Section 4 and Theorem 4.8 for details.

The question of resilience of random graphs with respect to Hamiltonicity has attracted considerable attention in the last years. Sudakov and Vu [33] showed that $G_{n,m}$ is $(1/2-o(1))$-resilient with respect to Hamiltonicity for $m \geq n \log^4 n$, and conjectured that already $m \gg n \log n$ should suffice. The lower bound on m was gradually improved in a series of paper: Frieze and Krivelevich [15] showed that $G_{n,m}$ is α-resilient for $p \geq Kn \log n$, for some small $\alpha > 0$ and sufficiently large K; Ben-Shimon, Krivelevich, and Sudakov [3] showed $(1/6-o(1))$-resilience and soon after

---

its vertices.

improved it to $(1/3 - o(1))$-resilience [4] for $m \geq \frac{1+\varepsilon}{2} n \log n$; finally, Lee and Sudakov [27] showed the optimal $(1/2 - o(1))$-resilience for $m \geq Kn \log n$.

While the last result settles the original conjecture of Sudakov and Vu, the gap between the value of m for which $G_{n,m}$ is typically Hamiltonian and for which we know is $(1/2 - o(1))$-resilient with respect to Hamiltonicity remains. A simple corollary of Theorem 1.4 closes this gap by showing that the two coincide. More precisely, we obtain that as soon as $\delta(G_m) \geq 2$ the graph G_m is resilient with respect to Hamiltonicity. This provides a resilience version of the hitting-time result of Ajtai, Komlós, and Szemerédi [1] and Bollobás [5].

Theorem 1.5. *Let $\varepsilon > 0$ be a constant. Consider the random graph process $\{G_i\}$ and let $m_2 = \min\{m : \delta(G_m) = 2\}$ denote the step in which the last vertex of degree at most one disappears. Then a.a.s. we have that G_{m_2} is $(1/2 - \varepsilon)$-resilient with respect to being Hamiltonian.*

We note that Theorem 1.5 (but not 1.4) was proven independently also by Montgomery [30] using a different approach.

Structure of the paper. In the next section we give some necessary definitions and tools used throughout the paper. In addition to providing standard facts, we introduce the notion of *tiny* and *atypical* vertices and a strengthened notion of α-resilience. In Section 3 we prove (a strengthened version) of Theorem 1.2 that is based on this new version of α-resilience. Even though the proof strategy is simple on a high level, it requires many intricate details related to tiny and atypical vertices. In Section 4 we use some of these ideas and combine them with the Pósa's *rotation-extension* technique to prove a strengthened version of Theorem 1.4.

2 Preliminaries

Our graph theoretic notation is standard and follows the one from [10]. In particular, given a graph G and (not necessarily disjoint) subsets $X, Y \subseteq V(G)$, we denote by $e_G(X, Y)$ the number of edges of G with one endpoint in X and the other in Y. Note that every edge which lies in the intersection of X and Y is counted twice. We denote by $e_G(X)$ the number of edges with both endpoints in X. Furthermore, given a vertex v and an integer $\ell \geq 1$, we denote by $N_G(v)$ its set of neighbours and by $N_G^\ell(v)$ the set of all vertices which are at distance at most ℓ from v, excluding v. Given graphs G and H with the same vertex set, we define $G - H$ as the graph with vertex set $V(G)$ and edge set $E(G) \setminus E(H)$. Moreover, for a graph G and a set of edges E on the same vertex set, we write $G + E$ to denote the graph on the vertex set $V(G)$ and edge set $E(G) \cup E$. For a positive integer n and a function $0 \leq p := p(n) \leq 1$ we let $G_{n,p}$ denote the probability space of graphs with vertex set $[n] = \{1, \ldots, n\}$ where each pair of vertices forms an edge of $G \sim G_{n,p}$ with probability p, independently of all other pairs. We make use of the standard asymptotic notation o, O, ω, and Ω. In addition, we write O_ε or Ω_ε to emphasise that the hidden constant depends on a parameter ε. For two functions a and b we write $a \ll b$ to indicate $a = o(b)$ and $a \gg b$ to indicate $a = \omega(b)$. In all occurrences $\log$ denotes the natural logarithm. We mostly suppress floors and ceilings whenever they are not crucial. Given $\varepsilon, x, y \in \mathbb{R}$, we write $x \in (1 \pm \varepsilon)y$ to denote $(1 - \varepsilon)y \leq x \leq (1 + \varepsilon)y$. Finally, we use subscripts with constants such as $C_{2.5}$ to indicate that $C_{2.5}$ is a constant given by Claim/Lemma/Proposition/Theorem 2.5.

Throughout the paper we make use of the following standard estimate on tail probabilities of a binomial random variable $\text{Bin}(n, p)$ with parameters n and p, see e.g. [14].

Lemma 2.1 (Chernoff bounds). *Let $X \sim Bin(n, p)$ and let $\mu := \mathbb{E}[X]$. Then for all $0 < \delta < 1$:*

- $\Pr[X \geq (1 + \delta)\mu] \leq e^{-\frac{\delta^2 \mu}{3}}$, *and*

- $\Pr[X \leq (1 - \delta)\mu] \leq e^{-\frac{\delta^2 \mu}{2}}$.

Remark 2.1. This result remains true if X has a hypergeometric distribution, cf. [19, Theorem 2.10].

2.1 Tiny and atypical vertices

Even though our main results concern the random graph process rather than the binomial random graph $G_{n,p}$, in the course of the proof we extensively rely on some of its properties. For small values of p, in particular those which give rise to a number of edges as in Theorem 1.2 and Theorem 1.4, a random graph $G_{n,p}$ contains many vertices whose degrees deviate significantly from the average degree. In our proofs those vertices require special attention. The following definition captures them formally.

Definition 2.2. Given $\delta, p \in [0, 1]$ and a graph G with n vertices, we define the following sets of vertices:

$$\text{TINY}_{p,\delta}(G) = \{v \in V(G)\colon \deg_G(v) < \delta np\},$$
$$\text{ATYP}_{p,\delta}(G) = \{v \in V(G)\colon \deg_G(v) \notin (1 \pm \delta)np\}.$$

We refer to the vertices in $\text{TINY}_{p,\delta}(G)$ as *tiny* and to the vertices in $\text{ATYP}_{p,\delta}(G)$ as *atypical*.

Note that in the above definition G is an arbitrary graph, not necessarily a $G_{n,p}$. This is the reason why p appears in the index as well. In our applications p is usually chosen such that np is roughly (but not necessarily exactly) the average degree of G. Thus, as the names indicate, tiny vertices have a degree significantly smaller than the average degree and atypical ones are bounded away from the average by a small constant factor. If $G \sim G_{n,p}$ then both of these sets become empty as soon as p is large enough: for $p \geq (1 + \varepsilon) \log n / n$ there are a.a.s. no tiny vertices for sufficiently small δ (depending on ε), and similarly for $p \geq C \log n / n$ there are a.a.s. no atypical vertices if C is sufficiently large (depending on δ).

With this definition in mind we refine the notion of α-resilience.

Definition 2.3. Given a graph $G = (V, E)$, a graph property $\mathcal{P}$, constants $\alpha, \delta_t, \delta_a \in [0, 1]$, and integers $K_t, K_a \in \mathbb{N}$, we say that G is $(\alpha, \delta_t, K_t, \delta_a, K_a)$-*resilient* with respect to $\mathcal{P}$ if for every spanning subgraph $H \subseteq G$ such that

$$\deg_H(v) \leq \begin{cases} \deg_G(v) - K_t, & \text{if } v \in \text{TINY}_{p,\delta_t}(G), \\ \deg_G(v) - K_a, & \text{if } v \in \text{ATYP}_{p,\delta_a}(G) \setminus \text{TINY}_{p,\delta_t}(G), \\ \alpha \deg_G(v), & \text{otherwise,} \end{cases}$$

for every $v \in V$, where $p = |E|/\binom{|V|}{2}$, we have $G - H \in \mathcal{P}$.

Under some natural conditions, one easily sees that $(\alpha, \delta_t, K_t, \delta_a, K_a)$-resilience implies α-resilience.

Lemma 2.4. *Let $G = (V, E)$ be a graph with n vertices and $\alpha, \delta_t, \delta_a \in [0, 1]$ and $K_t, K_a \in \mathbb{N}_0$ constants such that $(1 - \alpha)\delta_t np \geq K_a \geq K_t$, where $p = |E|/\binom{n}{2}$. If the minimum degree d of G is such that $d - \lfloor \alpha d \rfloor \geq K_t$, then $(\alpha, \delta_t, K_t, \delta_a, K_a)$-resilience of G implies α-resilience of G.* $\quad\square$

In particular, in our applications we have that K_t is either 1 (for perfect matchings) or 2 (for Hamilton cycles) and K_a is a (large) constant. A random graph $G_{n,p}$ with $p \gg 1/n$ then a.a.s. trivially fulfils the first condition in the lemma above. In addition, for $\alpha = 1/2 - \varepsilon$, the second condition is also satisfied if we constrain $G_{n,p}$ to the subgraph induced by vertices with degree 1 (perfect matchings) resp. the 2-core (Hamilton cycles).

Next, we establish a couple of properties of tiny and atypical vertices in random graphs.

Lemma 2.5. *Given $\delta > 0$, if $p \geq \log n/(3n)$ then $G \sim G_{n,p}$ a.a.s. satisfies:*

$$|\mathrm{ATYP}_{p,\delta}(G)| \leq n/(\log n)^3.$$

Proof. The left hand side consists of all vertices $v \in V(G)$ with degree $\deg(v) \notin (1 \pm \delta)np$. By Chernoff bounds, we have that

$$\Pr[\deg(v) \notin (1 \pm \delta)np] \leq e^{-\Omega_\delta(np)} \leq (\log n)^{-4},$$

with room to spare. Thus, Markov's inequality implies that the probability that there are at least $n/(\log n)^3$ vertices with degree not in $(1 \pm \delta)np$ is at most $(\log n)^{-1} = o(1)$, as claimed. $\square$

The crucial observation in our proof strategy is that tiny and atypical vertices cannot be clumped. In particular, no vertex has too many atypical vertices in its proximity.

In the graph process $\{G_i\}$ the property of being tiny or atypical is not monotone. Nevertheless, we expect that these vertices only change slightly over short periods of time. In order to capture this formally, we make use of the fact that one can generate a subsequence of the graph process as follows: choose p_0 such that $G^- \sim G_{n,p_0}$ has a.a.s. just a bit less than m edges, and choose p' such that $G_{n,p'}$ has just a bit more than εm edges. Then the union $G^+ = G^- \cup G_{n,p'}$ contains all graphs from $\{G_i\}$ with $m \leq i \leq (1 + \varepsilon)m$. Moreover, as $p' \ll p_0$ we expect that the tiny and atypical vertices of G^- are still scattered, even if we include all the edges from G^+. The following lemma makes this precise.

Lemma 2.6. *Let $k \geq 2$ be an integer and $\varepsilon > 0$ a constant. There exist positive constants $\delta(\varepsilon)$ and $L(\varepsilon, k)$, such that if $p_0 \geq (1 + \varepsilon) \log n/(kn)$ and $p' \leq \varepsilon p_0$ then a.a.s. the following holds. Let $G^- \sim G_{n,p_0}$ and set $G^+ = G^- \cup G_{n,p'}$ and $p_1 = 1 - (1 - p_0)(1 - p')$. Then:*

(i) for every $v \in V(G^+)$ we have $\left| N^3_{G^+}(v) \cap (\mathrm{TINY}_{p_0,\delta}(G^-) \cup \mathrm{TINY}_{p_1,\delta}(G^+)) \right| \leq k - 1$,

(ii) for every $v \in V(G^+)$ we have $\left| N^3_{G^+}(v) \cap (\mathrm{ATYP}_{p_0,\delta}(G^-) \cup \mathrm{ATYP}_{p_1,\delta}(G^+)) \right| \leq L$,

(iii) for every cycle $C \subseteq G^+$ with $v(C) \leq 2k$ we have

$$\left| V(C) \cap (\mathrm{TINY}_{p_0,\delta}(G^-) \cup \mathrm{TINY}_{p_1,\delta}(G^+)) \right| \leq k - 2.$$

Proof. Throughout the proof we assume that $V(K_n) = V(G^+)$. We make use of the fact that G^+ is distributed as G_{n,p_1}.

(*i*) We prove a more general statement: if $T \subseteq G^+$ is a tree with $k \leq v(T) \leq 4k$ vertices, then it contains at most $k - 1$ vertices from $\mathrm{TINY}_{p_0,\delta(G^-)} \cup \mathrm{TINY}_{p_1,\delta}(G^+)$. Suppose that this is the case. Let $v \in V(G^+)$ be an arbitrary vertex and assume, towards a contradiction, that there are k vertices $u \in \mathrm{TINY}_{p_0,\delta}(G^-) \cup \mathrm{TINY}_{p_1,\delta}(G^+)$ in $N^3_{G^+}(v)$. Choose T to be a tree which contains v and a shortest path between v and every such vertex u. Then T has at least $1 + k$ and at most $1 + 3k$ vertices, and contains k vertices from $\mathrm{TINY}_{p_0,\delta}(G^-) \cup \mathrm{TINY}_{p_1,\delta}(G^+)$, which is a contradiction.

Let $T \subseteq K_n$ be a tree with $k \leq v(T) \leq 4k$ vertices and consider a subset $S \subseteq V(T)$ of size exactly k. Let $\mathcal{E}_T$ denote the event that $T \subseteq G^+$ and $\mathcal{E}_{T,S}$ the event that every vertex in S has at most δnp_0 neighbours in G^- which are outside of $V(T)$, or at most δnp_1 neighbours in G^+ outside of $V(T)$. Note that $\mathcal{E}_T$ and $\mathcal{E}_{T,S}$ are independent events. Moreover, if there exists a tree $T \subseteq G^+$ which contains at least k vertices from $\mathrm{TINY}_{p_0,\delta}(G^-) \cup \mathrm{TINY}_{p_1,\delta}(G^+)$ then clearly both events $\mathcal{E}_T$ and $\mathcal{E}_{T,S}$ happen for S being a subset of size k of such vertices. We show that the probability that $\mathcal{E}_T \wedge \mathcal{E}_{T,S}$ happens for any T, S is $o(1)$, which implies the statement.

First, note that for a fixed tree $T \subseteq K_n$ we have $\Pr[\mathcal{E}_T] = p_1^{v(T)-1}$. Next, in order for $\mathcal{E}_{T,S}$ to happen, for each vertex in S we need it to have at most δnp_0 edges incident to it in G^-, or at

6

most $\delta n p_1$ edges in G^+, with the other endpoint being in $V(G^+) \setminus V(T)$. The probability that this happens is at most

$$\Pr[\text{Bin}(n - v(T), p_0) \leq \delta n p_0] + \Pr[\text{Bin}(n - v(T), p_1) \leq \delta n p_1].$$

Elementary calculations show that this sum can be bounded by $e^{-(1-\varepsilon/4)np_0}$, for sufficiently small δ depending on ε. As the probabilities are independent for different vertices in S, we get

$$\Pr[\mathcal{E}_{T,S}] \leq (e^{-(1-\varepsilon/4)np_0})^k.$$

For each $k \leq t \leq 4k$ there are at most $\binom{n}{t}t^{t-2}$ different trees $T \subseteq K_n$ with exactly t vertices, and for each such tree $\binom{t}{k}$ choices for $S \subseteq V(T)$. Therefore, using union bound over all values of t and all such pairs (T, S) we estimate the probability that some $\mathcal{E}_T \wedge \mathcal{E}_{T,S}$ happens as follows:

$$\Pr[\bigcup_{(T,S)} (\mathcal{E}_T \wedge \mathcal{E}_{T,S})] \leq \sum_{t=k}^{4k} \binom{n}{t}\binom{t}{k}t^{t-2} \cdot \Pr[\mathcal{E}_T \wedge \mathcal{E}_{T,S}] \leq \sum_{t=k}^{4k} n^t t^t \cdot p_1^{t-1} \cdot e^{-k(1-\varepsilon/4)np_0}.$$

As $p_1 \leq (1+\varepsilon)p_0$ and $n^t p_0^{t-1} \cdot e^{-k(1-\varepsilon/4)np_0}$ is decreasing in p_0, this implies

$$\Pr[\bigcup_{(T,S)} (\mathcal{E}_T \wedge \mathcal{E}_{T,S})] = O_{\varepsilon,k}(n \cdot (\log n)^{4k} \cdot n^{-1-\varepsilon/2}) = o(1).$$

(*ii*) Similarly to the previous case, we prove a statement for trees that implies the desired result. Let $T \subseteq K_n$ be a tree with $L \leq v(T) \leq 4L$ and let $S \subseteq V(T)$ be a set of L vertices. Let $\mathcal{E}_T$ denote the event that $T \subseteq G^+$ and $\mathcal{E}_{T,S}$ the event that every vertex $v \in S$ satisfies either $|N_{G^-}(v) \setminus V(T)| \notin (1 \pm \delta/2)np_0$ or $|N_{G^+}(v) \setminus V(T)| \notin (1 \pm \delta/2)np_1$. If there exists $T \subseteq G^+$ with at least L vertices belonging to $\text{ATYP}_{p_0,\delta}(G^-) \cup \text{ATYP}_{p_1,\delta}(G^+)$, both $\mathcal{E}_T$ and $\mathcal{E}_{T,S}$ happen (where once again S is any subset of L such vertices). For a fixed vertex $v \in S$, Chernoff bounds show that the probability that v satisfies $|N_{G^-}(v) \setminus V(T)| \notin (1 \pm \delta/2)np_0$ or $|N_{G^+}(v) \setminus V(T)| \notin (1 \pm \delta/2)np_1$ is at most

$$\Pr[\mathcal{E}_{T,S}] \leq \Pr[\text{Bin}(n - v(T), p_0) \notin (1 \pm \delta/2)np_0] + \Pr[\text{Bin}(n - v(T), p_1) \notin (1 \pm \delta/2)np_1] \leq e^{-\gamma np_0},$$

for some $\gamma = \gamma(\varepsilon) > 0$ (for $\delta = \delta(\varepsilon)$ sufficiently small). Note that these events are independent for different vertices in S and therefore

$$\Pr[\mathcal{E}_{T,S}] \leq (e^{-\gamma np_0})^L.$$

Analogous union bound analysis as in part (*i*) yields

$$\Pr[\bigcup_{(T,S)} (\mathcal{E}_T \wedge \mathcal{E}_{T,S})] = O_{\varepsilon,k}(n \cdot (\log n)^{4L} \cdot e^{-L\gamma np_0}) = o(1),$$

for L large enough depending on ε and k.

(*iii*) Let $C \subseteq K_n$ be a cycle with $3 \leq v(C) \leq 2k$ vertices and consider a subset $S \subseteq V(C)$ of size exactly $k - 1$. Let $\mathcal{E}_C$ and $\mathcal{E}_{C,S}$ denote the events as in (*i*) with T replaced by C. Assuming that there exists a cycle $C \subseteq G^+$ with at least $k - 1$ vertices from $\text{TINY}_{p_0,\delta}(G^-) \cup \text{TINY}_{p_1,\delta}(G^+)$, both events $\mathcal{E}_C$ and $\mathcal{E}_{C,S}$ happen (again, for S being any subset of $k-1$ such vertices). Analogous analysis as before shows

$$\Pr[\mathcal{E}_{C,S}] \leq (e^{-(1-\varepsilon/4)np_0})^{k-1},$$

On the other hand, we have $\Pr[\mathcal{E}_C] = p_1^{v(C)}$. For each $3 \leq \ell \leq 2k$ there are less than n^ℓ cycles $C \subseteq K_n$ with exactly ℓ vertices, and for each such cycle at most $\binom{\ell}{k-1}$ many choices for

7

$S \subseteq V(C)$. A union bound over all values of ℓ and all possible pairs (C, S), shows that the probability that the property *(iii)* fails is at most

$$\Pr[\bigcup_{(C,S)} (\mathcal{E}_C \wedge \mathcal{E}_{C,S})] \leq \sum_{\ell=3}^{2k} n^\ell \binom{\ell}{k-1} \cdot \Pr[\mathcal{E}_C \wedge \mathcal{E}_{C,S}] \leq \sum_{\ell=3}^{2k} n^\ell \ell^k \cdot p_1^\ell \cdot e^{-(k-1)(1-\varepsilon/4)np_0}.$$

We once again obtain

$$\Pr[\bigcup_{(C,S)} (\mathcal{E}_C \wedge \mathcal{E}_{C,S})] = O_{\varepsilon,k}((\log n)^{2k} \cdot n^{-(k-1)(1+\varepsilon/2)/k}) = o(1),$$

and the property follows. $\qquad\square$

We note that the proof actually shows that we could replace $N_{G+}^3(v)$ with $N_{G+}^\ell(v)$ for any constant $\ell \in \mathbb{N}$. For our purposes, the third neighbourhood suffices.

2.2 Properties of random graphs

In this subsection we establish some further, more general properties of random graphs which we rely on in the later sections. The following is a well known bound on the number of edges between sets of vertices in random graphs (see, e.g. [25, Corollary 2.3]).

Lemma 2.7. *Let $p = p(n) \leq 0.99$. Then $G \sim G_{n,p}$ a.a.s. has the property that for every two (not necessarily disjoint) subsets $X, Y \subseteq V(G)$ the number of edges with one endpoint in X and the other in Y satisfies:*
$$|e_G(X, Y) - |X||Y|p| \leq c\sqrt{|X||Y|np},$$
for some absolute constant $c > 0$.

Lastly, we show that if p is not too large then the neighbourhood of every two vertices is almost disjoint.

Lemma 2.8. *If $p = o(n^{-5/6})$ then $G \sim G_{n,p}$ a.a.s. has the property that every two distinct vertices have at most two common neighbours, i.e. $|N_G(u) \cap N_G(v)| \leq 2$ for all distinct $u, v \in V(G)$.*

Proof. It is sufficient to show that there are no two vertices with a common neighbourhood of size three. The probability that there exists a pair of vertices $u, v \in V(G)$ violating this property is at most
$$\binom{n}{2}\binom{n-2}{3}p^6 \leq n^5 p^6 = o(1).$$

This completes the proof. $\qquad\square$

3 Perfect matchings

Let $\{G_i\}$ denote the random graph process. Recall that Theorem 1.2 states that a.a.s. for every $m \geq \frac{1+\varepsilon}{4}n \log n$ we have that the subgraph of G_m obtained by removing isolated vertices is resilient with respect to containing a perfect matching. As remarked earlier, G_m has the same distribution as $G_{n,m}$ thus an obvious way to prove this statement would be to estimate for each m the probability that $G_{n,m}$ has the resilience property and apply a union bound. Unfortunately, this probability is roughly $1 - e^{-\alpha \cdot 2m/n}$, for some small constant $\alpha > 0$, which is too weak to cover all the values of m which are of order $n \log n$. This is not surprising as typically random

graphs fail to satisfy a property involving spanning subgraphs with a probability that is only exponential in the average degree. We go around this issue by using the following strategy: rather than handling each graph individually, we bundle consecutive graphs in groups of size roughly $\varepsilon n \log n$. The following proposition shows that almost surely *all* graphs in such a group have the desired property. A union bound over constantly many groups thus implies that the theorem holds for all $m \in \{\frac{1+\varepsilon}{4}n \log n, \dots, Cn \log n\}$, for some large constant C of our choice. The remaining values of m can then be treated one at a time using a result of Sudakov and Vu [33], which gives a probability of at least $1 - 1/n^3$ for all $m \geq Cn \log n$.

Proposition 3.1. *For every constant $\varepsilon > 0$ and integer $\frac{1+\varepsilon}{4}n \log n \leq m_0 \leq n(\log n)^2$, there exist positive constants $\delta_t(\varepsilon)$, $\delta_a(\varepsilon)$, and $K(\varepsilon)$ such that the random graph process $\{G_i\}$ a.a.s. has the following property: For any integer $m_0 \leq m \leq (1 + \varepsilon/4)m_0$ the graph obtained by removing all isolated vertices from G_m is $(1/2 - \varepsilon, \delta_t, 1, \delta_a, K)$-resilient with respect to having a perfect matching.*

We remark that the fact that we use the same ε in the lower bound for m_0 as we do for the resilience is without loss of generality, as we can always choose the smaller of the two for both.

Proof overview. The proof itself is somewhat technical, so let us first give a brief overview before we dive into the details. To handle all graphs G_m for $m \in [m_0, (1+\varepsilon/4)m_0]$ simultaneously we change the way G_m is generated (akin to what has been indicated prior to Lemma 2.6). Instead of generating G_m starting from an empty graph, we obtain it as follows: sample $G^- \sim G_{n,p_0}$ and $G_{n,p'}$, and choose a random ordering π of the edges in $G_{n,p'}$. In particular, we choose p_0 and p' such that almost surely G^- has less than m_0 edges (but not too much less) and $G^+ = G^- \cup G_{n,p'}$ has more than $(1 + \varepsilon/4)m_0$ edges (again, not too much more). Now we can generate each G_m as a union of G^- and the first $m - e(G^-)$ edges according to π; it is a simple exercise to show that this is the same as generating G_m from 'scratch'.

It is crucial that all the properties of G^- and G^+ that we use (e.g. upper bound on the number of edges between certain sets in G^+, tiny and atypical vertices being far apart in both G^- and G^+, etc.), are such that they are also satisfied by all graphs 'squeezed' in between them. While this is a standard technique in showing hitting-time results, it raises problems in the resilience setting. For example, vertices can be tiny or atypical in G_m without being tiny or atypical in G^- or G^+.

We circumvent this by defining *tiny* vertices to be those vertices that are tiny in at least one of G^- or G^+ with respect to a parameter δ_t' which is somewhat larger than the value of δ_t in the definition of $(1/2 - \varepsilon, \delta_t, 1, \delta_a, K)$-resilience in Proposition 3.1. In this way we can guarantee that every vertex in G that is tiny with respect to δ_t is also tiny in G^- or G^+ with respect to δ_t'. If we thus allow an adversary to remove all but one incident edge for all vertices that are tiny in G^- or G^+ with respect to δ_t', this may result in the removal of more edges than allowed by the definition of $(1/2 - \varepsilon, \delta_t, 1, \delta_a, K)$-resilience. Nevertheless, Lemma 2.6 still guarantees that all those vertices are sufficiently far apart in G^+ and thus also in $G \subseteq G^+$. This, in turn, implies that all these vertices can be covered in the matching even though they all have only one edge left. Atypical vertices are handled similarly by considering atypical vertices in G^- or G^+ with respect to a parameter δ_a' that is somewhat smaller than the value of δ_a in the definition of $(1/2 - \varepsilon, \delta_t, 1, \delta_a, K)$-resilience. The rest of the proof then follows standard arguments for finding a perfect matching in sparse random graphs (see, e.g. [19, 29]): greedily match all tiny and atypical vertices, in that order, and show that the remaining graph can be equipartitioned in a way that the resulting bipartite graph satisfies Hall's matching criteria.

Proof. Let $p_0 = (1 - \varepsilon/16)m_0/\binom{n}{2}$, $p' = (\varepsilon/2)p_0$, and let G^+ be the union of independent random graphs $G^- \sim G_{n,p_0}$ and $G_{n,p'}$. Then G^+ is distributed as G_{n,p_1}, where $p_1 = 1 - (1 - p_0)(1 - p')$.

By Lemma 2.7 we have that the number of edges in G^- is a.a.s. at most m_0 and the number of edges in G^+ is a.a.s. at least $(1 + \varepsilon/4)m_0$.

Let $\delta'_t = \delta_{2.6}(\varepsilon/2)$, $\delta'_a = \min\{\varepsilon/64, \delta_{2.6}(\varepsilon/2)\}$, $c = c_{2.7}$, $L = \max\{L_{2.6}(\varepsilon/2), L_{3.2}(\varepsilon)\}$. With this at hand, we set the constants given in the statement of the proposition as follows: $\delta_t = \delta'_t/40$, $\delta_a = \max\{\varepsilon/4, 16\delta'_a\}$, and $K = 2L$.

For the rest of the proof consider some $m_0 \le m \le (1 + \varepsilon/4)m_0$, and let $G \subseteq G_m$ be the subgraph obtained by removing all isolated vertices from G_m. Let $V = V(G)$ denote its vertex set and let TINY and ATYP be sets of vertices defined as:

$$\text{TINY} := (\text{TINY}_{p_0, \delta'_t}(G^-) \cup \text{TINY}_{p_1, \delta'_t}(G^+)) \cap V,$$
$$\text{ATYP} := (\text{ATYP}_{p_0, \delta'_a}(G^-) \cup \text{ATYP}_{p_1, \delta'_a}(G^+)) \cap V.$$

By Lemma 2.5 there are at most $n/\log^3 n$ isolated vertices in G^-, thus there are at most that many in G_m as well. Our choice of constants then implies that for all $m_0 \le m \le (1 + \varepsilon/4)m_0$ we have

$$\text{TINY}_{p, \delta_t}(G) \subseteq \text{TINY} \qquad \text{and} \qquad \text{ATYP}_{p, \delta_a}(G) \subseteq \text{ATYP},$$

where $p = m/\binom{v(G)}{2}$ denotes the density of G (note that $E(G) = E(G_m)$ as we only remove isolated vertices). As noted in the discussion before the proof, if we show that for every graph H of the form

$$\deg_H(v) \le \begin{cases} \deg_G(v) - 1, & \text{if } v \in \text{TINY}, \\ \deg_G(v) - K, & \text{if } v \in \text{ATYP} \setminus \text{TINY}, \\ (1/2 - \varepsilon)np_1, & \text{otherwise}, \end{cases} \tag{1}$$

the graph $G - H$ contains a perfect matching, then this implies that G is $(1/2 - \varepsilon, \delta_t, 1, \delta_a, K)$-resilient with respect to having a perfect matching. Note that this follows from the fact that typical vertices in G have degree at most $(1 + \delta_a)np$ and thus removing an $(1/2 - \varepsilon)$-fraction of their degree is less than removing $(1/2 - \varepsilon)np_1$ incident edges, since $p_1 > (1 + \varepsilon/4)p$. In the remainder of the proof we show that the graph $G - H$ indeed contains a perfect matching.

First, note that G satisfies a series of properties:

(M1) the maximum degree of G is at most $\Delta(G) \le \log^3 n$,

(M2) for all $X, Y \subseteq V$ we have $e_G(X, Y) \le |X||Y|p_1 + c\sqrt{|X||Y|np_1}$,

(M3) for all $v, u \in V$ we have $|N_G(v) \cap N_G(u)| \le 2$,

(M4) for all $v \in V$ we have $|N_G^2(v) \cap \text{TINY}| \le 1$ and $|N_G^2(v) \cap \text{ATYP}| \le L$.

(M5) $|\text{ATYP}| \le \frac{2n}{\log^3 n}$.

We show that the properties hold in G^+, and hence in any subgraph $G \subseteq G^+$. Indeed, we have $p \le 10\log^2 n/n$ (by our assumptions on m_0) thus (M1) follows from simple bounds on the binomially distributed random variable and a union bound over all vertices. Property (M2) is given by Lemma 2.7 and (M3) by Lemma 2.8. Next, (M4) holds by our choice of L and by Lemma 2.6 applied with $k = 2$ and $\varepsilon/2$ as ε. Note that we can indeed apply Lemma 2.6 with these parameters as $p_0 \ge (1 + \varepsilon/2)\log n/(2n)$. Lastly, (M5) holds by Lemma 2.5.

Consider graph H which satisfies (1), and let $G' = G - H$ and $U := V \setminus \text{ATYP}$. If G has an odd number of vertices then we remove a vertex $v \in \text{TINY}$ which is incident to $u \notin \text{TINY}$, and if such a vertex does not exist, then we remove one from $V \setminus \text{TINY}$. Owing to the property (M4) this does not decrease the degree of any vertex from TINY and decreases the degree of a vertex from $V \setminus \text{TINY}$ by at most one. Note that all vertices $v \in U$ have degree in G^- at least

$(1 - \delta'_a)np_0$, hence by (M4) we have:

$$\begin{aligned}
\deg_{G'}(v, U) &\geq \deg_{G^-}(v) - \deg_H(v) - \deg_G(v, \text{ATYP}) - 1 \\
&\geq (1 - \delta'_a)np_0 - (1/2 - \varepsilon)np_1 - L - 1 \\
&\geq (1/2 - \delta'_a + \varepsilon/2)np_1 - L - 1 \\
&\geq (1/2 + \varepsilon/4)np_1,
\end{aligned}$$

as $\delta'_a \leq \varepsilon/64$ and $p_0 \geq p_1/(1 + \varepsilon/2)$.

Due to property (M4) it is straightforward to greedily find a matching saturating all vertices from TINY. Next, we match all remaining vertices from ATYP. Again we proceed greedily. Consider an arbitrary, still unmatched vertex $w \in \text{ATYP} \setminus \text{TINY}$. Note that $\deg_{G'}(w) \geq K - 1$ (here -1 comes from the fact that we might have removed one vertex from G to achieve an even number of vertices). By (M4) and our choice of K we have $|N^2_{G'}(w) \cap \text{ATYP}| \leq L \leq K/2$. Therefore, at most $K/2$ neighbours of w have been matched so far, thus there is an available one.

Let $V(M)$ be the set of all vertices saturated in this partial matching. In particular $V(M)$ contains all vertices of ATYP. Set $U_1 := U \setminus V(M)$ and let $G'' := G'[U_1]$ be the subgraph induced by the remaining vertices. Observe that by property (M4) the degree of any vertex in U_1 decreases by at most $2L$ with respect to its degree in G', hence for all vertices $v \in U_1$ we have

$$\deg_{G''}(v) \geq \deg_{G'}(v, U) - \deg_{G'}(v, V(M)) \geq (1/2 + \varepsilon/4)np_1 - 2L \geq (1/2 + \varepsilon/5)np_1.$$

Take $U_1 = A \cup B$ to be a uniformly at random chosen balanced bipartition of the set U_1. We now define the set of all vertices which have significantly less than the expected degree in either A or B as

$$D := \{v \in U_1 \colon \deg_{G''}(v, A) < (1/2 + \varepsilon/7)|A|p_1 \text{ or } \deg_{G''}(v, B) < (1/2 + \varepsilon/7)|B|p_1\}$$

and call all such vertices *degenerate*.

Firstly, we show that there are not many degenerate vertices. As $|A|, |B| \geq (1/2 - o(1))n$ by property (M5), we have from Chernoff bounds for hypergeometrically distributed random variables that a fixed vertex v is degenerate with probability

$$\Pr[v \text{ is degenerate}] \leq 2\Pr\left[\deg_{G''}(v, A) < (1/2 + \varepsilon/7)|A|p_1\right] \leq e^{-2\gamma np_1} \leq n^{-2\gamma},$$

for some $\gamma = \gamma(\varepsilon) > 0$. Consequently, by Markov's inequality there are at most $n^{1-\gamma}$ degenerate vertices, i.e. $|D| \leq n^{1-\gamma}$. Next we show that the degenerate vertices cannot be 'too close' in G''.

Claim 3.2. *There exists a positive constant $L(\varepsilon)$ such that a.a.s. for every $v \in U_1$ we have $|N^2_{G''}(v) \cap D| < L$.*

Proof. Consider some vertex $v \in V(G'')$ and a subset $D_v \subseteq N^2_{G''}(v)$ of size L. What is the probability that a random equipartition $A \cup B$ of $V(G'')$ makes all the vertices in D_v degenerate?

First, by property (M3) we know that for every vertex $u \in D_v$ there are at least

$$(1/2 + \varepsilon/5)np_1 - 2L \geq (1/2 + \varepsilon/6)np_1 \tag{2}$$

vertices in its neighbourhood in G'' which do not belong to the neighbourhood of any other vertex in $D_v \setminus \{u\}$. Let us denote such vertices with N^*_u.

If all vertices in D_v are degenerate, then there has to exist a subset $D'_v \subseteq D_v$ of size at least $L/2$ such that all vertices in D'_v either have too few neighbours into A or too few neighbours into B.

By symmetry we may assume that this set is A. That is, we have $|A \cap N_u^*| \leq (1/2 + \varepsilon/7)|A|p_1$ for every $u \in D_v'$. Therefore,

$$\left| A \cap \bigcup_{u \in D_v'} N_u^* \right| \leq (1/2 + \varepsilon/7)|A|p_1|D_v'|.$$

Let $\mathcal{E}_{D_v'}$ denote the event of this happening. From (2) we have

$$\left| \bigcup_{u \in D_v'} N_u^* \right| \geq (1/2 + \varepsilon/6)np_1|D_v'|.$$

Thus, as A is a randomly chosen subset of linear size, Chernoff bounds for hypergeometrically distributed random variables show

$$\Pr[\mathcal{E}_{D_v'}] \leq e^{-\gamma L np_1}, \tag{3}$$

for some $\gamma = \gamma(\varepsilon) > 0$ not depending on L. To summarise, we can bound the probability that all the vertices in D_v are degenerate by applying (3) together with a union bound over all possible choices for D_v'. Finally, we take a union bound over all vertices v and sets D_v. There are n choices for v and at most $\binom{(\log^3 n)^2}{L}$ choices for $D_v \subseteq N_{G''}^2(v)$ of size L (which follows from (M1)). The expected number of vertices for which there exists a set D_v of size L consisting solely of degenerate vertices is then at most

$$n \cdot (\log n)^{6L} \cdot 2^L \cdot e^{-\gamma L np_1} = o(1),$$

for $L > 0$ sufficiently large. In other words, for L large enough no such set exists with probability $1 - o(1)$, as claimed. $\qquad\square$

Similarly as before, we greedily construct a partial matching M_D that saturates all the degenerate vertices. For an arbitrary degenerate vertex $v \in D$ we have $\deg_{G''}(v) \geq (1/2 + \varepsilon/5)np_1$ and as there cannot be more than L degenerate vertices in $N_{G''}^2(v)$ by Claim 3.2, there is a vertex available to match v to.

Let $V(M_D)$ be the set of all vertices saturated in this partial matching and let $A' := A \setminus V(M_D)$ and $B' := B \setminus V(M_D)$. Again by Claim 3.2, for all $v \in (A \cup B) \setminus V(M_D)$ we get

$$\deg_{G''}(v, A') \geq (1/2 + \varepsilon/7)|A|p_1 - 2L \geq (1/4 + \varepsilon/16)np_1,$$

and analogously

$$\deg_{G''}(v, B') \geq (1/2 + \varepsilon/7)|B|p_1 - 2L \geq (1/4 + \varepsilon/16)np_1,$$

as $|A|, |B| > (1/2 - o(1))n$. However, we might not have $|A'| = |B'|$ any more. Assume w.l.o.g. that $|A'| > |B'|$ and note that $|A'| \leq |B'| + |V(M_D)|$. In order to find a balanced bipartition we thus need to redistribute at most $|D| \leq n^{1-\gamma}$ vertices, for some $\gamma > 0$. To achieve this we build a 2-*independent set* in A', i.e. an independent set in which no two vertices have a common neighbour, of size at least $n^{1-\gamma}$. Recall, from property (M1) we have $\Delta(G) \leq \log^3 n$, thus a straightforward greedy construction shows that there exists a 2-independent set of size at least $n/\log^7 n \geq n^{1-\gamma}$ in A', which is more than enough for our purposes.

Let A'' and B'' be the two sets after moving the vertices belonging to the 2-independent set into B'. Then $|A''| = |B''|$ (by construction) and vertices in A'' (respectively, B'') have degree at least $(1/4 + \varepsilon/16)np_1 - 1 \geq (1/4 + \varepsilon/17)np_1$ in B'' (respectively, A''), as we moved a 2-independent set.

It remains to find a perfect matching in the bipartite graph $G''[A'', B'']$. This is easily achieved through Hall's matching criteria (see, e.g. [10]). Recall its statement: if each subset $S \subseteq A''$ of

size at most $|A''|/2$ has at least $|S|$ neighbours in B'' (and similarly for $S \subseteq B''$), then $G''[A'', B'']$ contains a perfect matching. We now verify that this is indeed the case.

Let $S \subseteq A''$ be an arbitrary subset of size s. The number of edges between any two disjoint subsets of G'' of size $s \leq n/4$ is at most $s(np_1/4 + c\sqrt{np_1})$, by property (M2). On the other hand however, the minimum degree condition in G'' yields that the number of edges between S and B'' is at least $s(1/4 + \varepsilon/17)np_1$, which implies $|N(S)| > s$. This completes the proof of the proposition. $\qquad\square$

Having Proposition 3.1 at hand we now complete the proof. Let us first, as promised in the introduction, restate Theorem 1.2 in terms of stronger resilience.

Theorem 3.3. *Let $\varepsilon > 0$ be a constant and consider the random graph process $\{G_i\}$. There exist positive constants $\delta_t(\varepsilon)$, $\delta_a(\varepsilon)$, and $K(\varepsilon)$ such that a.a.s. for every $m \geq \frac{1+\varepsilon}{4}n \log n$ we have that the graph obtained from G_m by deleting all isolated vertices is $(1/2 - \varepsilon, \delta_t, 1, \delta_a, K)$-resilient with respect to containing a perfect matching.*

Proof. Let C be a sufficiently large constant. If $m \geq Cn \log n$ then the random graph $G_{n,m}$ does not contain atypical vertices. Therefore, by [33, Theorem 3.1] the statement holds for every such (fixed) m with probability at least $1 - e^{-\alpha \cdot C \log n}$, for some small constant $\alpha > 0$, thus a union bound implies that it holds for all $m \geq Cn \log n$ simultaneously. As for the rest, consider intervals of the form

$$\left[\frac{1 + i\varepsilon}{4} n \log n, \frac{1 + (i+1)\varepsilon}{4} n \log n \right)$$

for $i \in \{1, \ldots, C_\varepsilon\}$, where C_ε is such that the last interval contains $Cn \log n$. For each interval the conclusion of Proposition 3.1 holds with probability $1 - o(1)$, thus a union bound shows that it a.a.s. holds for all intervals. This concludes the proof. $\qquad\square$

Proof of Theorem 1.2. The assertion follows directly from Lemma 2.4 and Theorem 3.3. $\qquad\square$

4 Hamiltonicity

The so-called Pósa's *rotation-extension* method introduced in [31] is nowadays a standard approach for constructing Hamilton cycles in random graphs, cf. also [4, 14, 15, 22, 24]. The problem with having edge probabilities near (or even below) the connectivity threshold, is that then the random graph does not satisfy the expansion properties needed to apply the method. We go around this by partitioning the vertex set into typical and atypical vertices. The subgraph induced by the vertices of typical degree satisfies the expansion properties required to apply the rotation-extension technique and we thus get a Hamilton cycle in this subgraph by a standard approach. Our new contribution is to show how to extend this cycle to also contain all atypical vertices.

Towards this goal we make use of the facts given by Lemma 2.6, that is that atypical vertices do not clump. This allows us to modify the rotation-extension procedure from [27] such that it only uses the expansion properties of typical vertices. In the next section we review some basic notions and state necessary lemmas which are then used in the subsequent section to derive Theorem 1.4.

4.1 Backbone graphs and boosters

The central notion of the rotation-extension method is that of *boosters*. A booster is a non-edge in a graph G whose existence would increase the length of a longest path in G or close a Hamilton

path to a Hamilton cycle. The idea behind the rotation-expansion technique is that a graph which is not Hamiltonian contains so many boosters that adding a few random edges is highly likely to contain one of them. The name 'rotation-extension' comes from the way how boosters are obtained (*rotation*) and the fact that a longest path of every non-Hamiltonian graph can be increased by adding an edge in place of a booster (*extension*). We now make this precise.

Definition 4.1 (Boosters). Given a graph Γ, we say that a non-edge $\{u, v\} \notin E(\Gamma)$ is a *booster* with respect to Γ, if either $\Gamma + \{u, v\}$ is Hamiltonian or adding $\{u, v\}$ to Γ increases the length of a longest path. For a vertex $v \in V(\Gamma)$, we denote by $B_\Gamma(v)$ the set of boosters associated with v:

$$B_\Gamma(v) = \{u \in V(\Gamma) \setminus (N_\Gamma(v) \cup \{v\}) \colon \{v, u\} \text{ is a booster}\}.$$

A standard strategy for implementing the rotation-extension technique is to split the given graph into two graphs: a 'backbone' graph responsible for obtaining boosters, and the remainder responsible for finding real edges corresponding to boosters. As we are dealing with subgraphs of random graphs, rather than random graphs themselves, it is convenient to capture the main pseudorandom properties which are used.

Definition 4.2 (Backbone graph). Given $\alpha, q \in (0, 1)$ and an integer $K \geq 1$, we say that a graph Γ with n vertices is an (α, K, q)-*backbone graph* if there exists a partition of its vertex set $V(\Gamma) = U \cup W_1 \cup W_2$ such that the following holds:

(P1) $|W_1 \cup W_2| \leq \frac{n}{\log^2 n}$,

(P2) for every $v \in W_1$ we have $\deg_\Gamma(v) = 2$ and for every $u \in W_2$ we have $\deg_\Gamma(u) \geq 2K$,

(P3) for every $v \in W_1$ we have $N_\Gamma^2(v) \cap W_1 = \varnothing$,

(P4) for every $v \in V(\Gamma)$ we have $|N_\Gamma^2(v) \cap W_1| \leq 2$ and $|N_\Gamma^2(v) \cap W_2| \leq K$, and

(P5) for all $S \subseteq U$ we have

$$|N_\Gamma(S)| \geq \begin{cases} |S|\sqrt{nq}, & \text{if } |S| < K/q, \\ (1/2 + \alpha/2)n, & \text{if } |S| \geq K/q. \end{cases}$$

The role of the sets W_1 and W_2 and properties (P1)–(P4) is to capture properties of tiny and atypical vertices in a random graph with density q. Property (P5) states that the subgraph induced by typical vertices has good expansion properties.

The next lemma shows that a backbone graph contains many boosters. It can be proven similarly as [27, Lemma 3.2], with slight modifications which allow us to deal with the vertices in W_1 and W_2. We defer the proof to Section 4.3.

Lemma 4.3. *For every $\alpha, \delta > 0$, there exists a positive constant $K(\alpha)$ such that the following holds for $q \geq \delta \log n / n$. Let G_q be a graph with n vertices satisfying the property of Lemma 2.7 for q (as p) and some constant c' (as c), and $H_q \subseteq G_q$ a graph with $\Delta(H_q[U]) \leq (1/2 - 2\alpha)nq$, for some $U \subseteq V(G_q)$. Then an (α, K, q)-backbone graph Γ with a witness partition $V(\Gamma) = U \cup W_1 \cup W_2$, such that $\Gamma[U] = G_q[U] - H_q[U]$, is either Hamiltonian or there are at least $(1/2 + \alpha)n$ vertices $v \in U$ such that $|B_\Gamma(v) \cap U| \geq (1/2 + \alpha)n$.*

Note that rather than asking for $\Gamma = G_q - H_q$, we only require the subgraph of Γ induced by U to be given by $G_q[U] - H_q[U]$. The way we exploit such a weaker requirement becomes apparent in the next section. As the notation q in the previous lemma already indicates, we obtain an (α, K, q)-backbone graph Γ by considering a random subgraph with density q of some graph G.

The next lemma states that given a sufficiently sparse graph Γ with many boosters, a random graph with appropriate density is likely to contain plenty of edges corresponding to them.

Lemma 4.4. *For every $\alpha > 0$ there exists a positive constant $\mu(\alpha)$ such that for $p \geq \log n/(3n)$, the random graph $G \sim G_{n,p}$ a.a.s. satisfies the following. Let Γ be a graph with $e(\Gamma) \leq \mu n^2 p$ and $U \subseteq V(G)$ a subset of vertices such that $\Gamma[U] \subseteq G$. If there are at least $(1/2 + \alpha)n$ vertices $v \in U$ such that $|B_\Gamma(v) \cap U| \geq (1/2 + \alpha)n$, then there exists a vertex $v \in U$ satisfying $|N_G(v, B_\Gamma(v) \cap U)| > np/2$.*

The proof of the lemma is fairly standard and goes along the lines of the proof of [27, Lemma 3.5]; we include it for completeness in Section 4.3.

4.2 Proof of Theorem 1.4

Similarly as in the case of perfect matchings, instead of proving Theorem 1.4 directly, we first show a proposition which considers only a small range of m, until $m = Cn \log n$.

Proposition 4.5. *For every constant $\varepsilon > 0$ and integer $m_0 \geq (1 + \varepsilon)\frac{n \log n}{6}$ there exist positive constants $\delta_t(\varepsilon)$, $\delta_a(\varepsilon)$, and $K(\varepsilon)$ such that the random graph process $\{G_i\}$ a.a.s. has the following property: For every integer $m_0 \leq m \leq (1 + \varepsilon/6)m_0$, the 2-core of G_m is $(1/2 - \varepsilon, \delta_t, 2, \delta_a, K)$-resilient with respect to having a Hamilton cycle.*

The general strategy of the proof is similar as in the case of perfect matchings. We first sample two random graphs $G^- \sim G_{n,p_0}$ and $G_{n,p'}$ and look at their union $G^+ = G^- \cup G_{n,p'}$. The choice of densities p_0 and p' is such that for all values of $m \in [m_0, (1 + \varepsilon/6)m_0]$ the graph G_m is a.a.s. 'in between', that is $G^- \subseteq G_m \subseteq G^+$.

Let G denote the 2-core of G_m and let H be a graph removed by the adversary, as in Definition 2.3. We split the graph $G - H$ into a sparse graph Γ' which mostly contains edges from $G^- - H$ (edges outside of G^- are borrowed to handle atypical vertices), and the rest $G - H - \Gamma'$. Ideally, we would like that Γ' is a backbone graph. For m close to $n \log n/6$, the graph G contains short paths between two tiny vertices, some of which exist in Γ' as well preventing it from satisfying (P3). We circumvent this by constructing a new graph Γ from Γ' by *contracting* all paths of length at most two between tiny vertices. As tiny vertices do not clump, this has only a mild impact on the structure of the graph and all the other properties remain satisfied. Due to these contractions the graph Γ is not a subgraph of $G - H$ any more, however, the subgraph induced by typical vertices which are not a part of contracted paths, is. This allows us to apply Lemma 4.3 to conclude that such Γ is either Hamiltonian in which case we are done by 'unrolling' the contracted vertices, or contains many boosters between 'real' vertices of G. Using Lemma 4.4, we subsequently show that many such boosters appear in $G - H$ which allows us to complete a Hamilton cycle.

We point out that by requiring m just a bit larger, that is $m \geq (1 + \varepsilon)n \log n/4$, would be enough for (P3) to hold already in Γ' and the contraction process would not be necessary.

Proof. Let $p_0 = (1 - \varepsilon/16)m_0/\binom{n}{2}$, $p' = (\varepsilon/2)p_0$, and let G^+ be the union of independent random graphs $G^- \sim G_{n,p_0}$ and $G_{n,p'}$. Then G^+ is distributed as G_{n,p_1}, where $p_1 = 1 - (1 - p_0)(1 - p')$. By Lemma 2.7 we a.a.s. have both $e(G^-) \leq m_0$ and $e(G^+) \geq (1 + \varepsilon/6)m_0$.

Let $\delta_t' = \delta_{2.6}(\varepsilon/2)$, $\delta_a' = \min\{\varepsilon/64, \delta_{2.6}(\varepsilon/2)\}$, $c = c_{2.7}$, $L = \max\{L_{2.6}(\varepsilon/2), L_{4.6}(\varepsilon)\}$, $\mu = \mu_{4.4}(\varepsilon/4)$, $c' = c'(c, \mu)$ sufficiently large (cf. (4)), and set the constants given in the statement of the proposition as follows:

$$\delta_t = \delta_t'/40, \quad \delta_a = \max\{\varepsilon/4, 16\delta_a'\}, \quad \text{and} \quad K = \max\{40L, (64c'/\varepsilon)^2\}.$$

For the rest of the proof consider some $m_0 \leq m \leq (1 + \varepsilon/6)m_0$. Let G be the 2-core of G_m and set $V = V(G)$. Note that G can be obtained using the following procedure: initially set

15

$G = G_m$, and as long as G contains a vertex of degree at most one, remove it. Let R denote the set of removed vertices. The definition of the procedure then implies that there are at most $|R|$ edges incident to R in G_m. By Lemma 2.6 applied to G^+ (with $k = 3$), each vertex has at most two neighbours in $\text{TINY}_{p_0, \delta'_t}(G^-)$ thus no vertex which is not in $\text{TINY}_{p_0, \delta'_t}(G^-)$ can ever be removed. This also implies that the degree of every vertex decreases by at most 2. As thus $R \subseteq \text{TINY}_{p_0, \delta'_t}(G^-)$, Lemma 2.5 implies $|R| = O(n/\log^3 n)$.

Similarly as in the proof of Proposition 3.1, we define TINY and ATYP as follows:

$$\text{TINY} := (\text{TINY}_{p_0, \delta'_t}(G^-) \cup \text{TINY}_{p_1, \delta'_t}(G^+)) \cap V,$$

$$\text{ATYP} := (\text{ATYP}_{p_0, \delta'_a}(G^-) \cup \text{ATYP}_{p_1, \delta'_a}(G^+)) \cap V.$$

Using the previous observation on the number of removed vertices and edges, it is easy to see that for all $m_0 \le m \le (1 + \varepsilon/6)m_0$ we have

$$\text{TINY}_{p, \delta_t}(G) \subseteq \text{TINY} \qquad \text{and} \qquad \text{ATYP}_{p, \delta_a}(G) \subseteq \text{ATYP},$$

where $p = e(G)/\binom{v(G)}{2}$. Consider a graph $H \subseteq G$ such that

$$\deg_H(v) \le \begin{cases} \deg_G(v) - 2, & \text{if } v \in \text{TINY}, \\ \deg_G(v) - K, & \text{if } v \in \text{ATYP} \setminus \text{TINY}, \\ (1/2 - \varepsilon)np_1, & \text{otherwise.} \end{cases}$$

We show that $G - H$ contains a Hamilton cycle, which in turn implies that G is $(1/2 - \varepsilon, \delta_t, 2, \delta_a, K)$-resilient with respect to Hamiltonicity (as in the case of perfect matchings).

The key to our proof is the fact that tiny and atypical vertices cannot be clumped up in G, captured by the following properties:

(H1) for all $v \in V$ we have $|N_G^3(v) \cap \text{TINY}| \le 2$ and $|N_G^2(v) \cap \text{ATYP}| \le L$,

(H2) every cycle $C \subseteq G$ of size at most 6 contains at most one vertex from TINY.

By Lemma 2.6 applied for $k = 3$ and $\varepsilon/2$ as ε, one easily checks that properties (H1) and (H2) hold in G^+, and hence in any subgraph $G \subseteq G^+$. Note that we can indeed apply Lemma 2.6 with such parameter, as $p_0 \ge (1 + \varepsilon/2)\log n/(3n)$.

In the first step of the proof we carefully construct a backbone graph Γ. Set $q = (\mu/4)p_0$, where $\mu = \mu_{4.4}(\varepsilon/4)$, and consider a graph G_q obtained by keeping every edge of G^- independently with probability $\mu/4$. Moreover, remove from G_q all vertices which are not in V. By the discussion from the beginning of the proof, removing such vertices decreases the degree of every vertex in G_q by at most 2. We now show that G_q a.a.s. has certain properties.

As G^- a.a.s. satisfies the assertion of Lemma 2.7, we also have that G_q a.a.s. satisfies

$$|e_{G_q}(X, Y) - |X||Y|q| \le c'\sqrt{|X||Y|nq}, \tag{4}$$

for every two subsets $X, Y \subseteq V$, for some sufficiently large constant c' depending on c and μ. Next, we claim that most of the vertices in $V \setminus \text{ATYP}$ have degree at least $(1 - 2\delta'_a)nq$ in G_q, as well as degree at most $(1/2 - \varepsilon/2)nq$ in $H_q = G_q \cap H$. Fix $v \in V \setminus \text{ATYP}$. By Chernoff bounds we get

$$\Pr[\deg_{G_q}(v) < (1 - 2\delta'_a)nq] + \Pr[\deg_{H_q}(v) > (1/2 - \varepsilon/2)nq] \le e^{-\Omega_\varepsilon(nq)} \le n^{-2\gamma},$$

for some $\gamma = \gamma(\varepsilon) > 0$. Therefore, by Markov's inequality there are at most $n^{1-\gamma}$ vertices which satisfy at least one of these two conditions. We call all such vertices *degenerate* and denote them by D. Similarly as atypical vertices, such vertices cannot be clumped in G, and thus neither in G_q. This is formalised in the following claim.

16

Claim 4.6. *There exists a positive constant $L(\varepsilon)$ such that a.a.s. for every $v \in V$ we have $|N_G^2(v) \cap D| < L$.*

Proof. The proof goes along the lines of that of Claim 3.2. We omit the details. $\qquad\square$

In conclusion, G_q a.a.s. satisfies (4), every vertex in $V \setminus (\text{ATYP} \cup D)$ has degree at least $(1-2\delta'_a)nq$ in G_q and at most $(1/2 - \varepsilon/2)nq$ in H_q, and no vertex has more than L vertices from D in its second neighbourhood. From now on we choose one such graph G_q. In particular, (4) implies $e(G_q) \leq (\mu/4)n^2 p_0$.

We now construct a graph Γ' which is 'almost' a backbone graph, and then convert it into a true backbone graph, thereby overcoming the issue of two tiny vertices being on a short path, as discussed previously. Let $W_1' := \text{TINY}$, $W_2' := (\text{ATYP} \cup D) \setminus \text{TINY}$, and $U' := V \setminus (W_1' \cup W_2')$. Take Γ' to be a graph on the vertex set V obtained by taking all edges in $G_q[U'] - H_q[U']$ and adding some of the edges of $G - H$ incident to vertices in $W_1' \cup W_2'$ such that for every $v \in W_1'$ we have $\deg_{\Gamma'}(v) = 2$, and for every $u \in W_2$ we have $\deg_{\Gamma'}(u) \geq K - 2$ (by requiring that all $v \in W_1$ have $\deg_{\Gamma'}(v) = 2$ we potentially remove at most two incident edges to a vertex $u \in W_2'$). Note that $e(\Gamma') \leq (\mu/2)n^2 p_0$. Recall that property (P3) does not necessarily hold in Γ', thus we cannot claim it is a backbone graph. We take care of this issue as follows: let Γ be a graph obtained from Γ' by contracting every uv-path of length at most two, where $u, v \in W_1'$, and keeping exactly two edges incident to the newly obtained vertex, namely the ones incident to u resp. v in Γ' that are not a part of the uv-path. We also remove all multi-edges and loops. Observe that Γ is well-defined since there cannot exist a vertex $w \in V$ which belongs to two such paths uv-paths, due to property (H1).

In order to show that Γ is a backbone graph, we first define a witness partition $U \cup W_1 \cup W_2$. Let X be the set of vertices of Γ obtained by contractions, Y be the set of all vertices in $W_2' \cup U'$ that are inner vertices of some contracted path, and Z the set of all vertices of W_1' that are the endpoints of such paths. We now define

$$W_1 := (W_1' \setminus Z) \cup X, \qquad W_2 := W_2' \setminus Y, \qquad \text{and} \qquad U := U' \setminus Y.$$

In other words, W_1 consists of all newly formed vertices (obtained by contractions) as well as all remaining vertices of W_1', and W_2 and U consist of all vertices in W_2' resp. U' that are not a part of any contracted path.

We are now ready to show that Γ is an $(\varepsilon/4, K/10, q)$-backbone graph. Recall that $|V| \geq n - n/\log^2 n$. For every vertex $v \in U \cup W_2$ we have

$$\deg_\Gamma(v) \geq \deg_{\Gamma'}(v) - 1, \tag{5}$$

by using property (H1). We now check all the requirements of Definition 4.2:

(P1) follows by Lemma 2.5 and an upper bound on the size of D:

$$|W_1 \cup W_2| \leq |\text{ATYP}| + |D| \leq \frac{2n}{\log^3 n} + n^{1-\gamma} \leq \frac{n}{2\log^2 n}.$$

(P2) follows by construction. Indeed, if for any $v \in W_1$ we have $\deg_\Gamma(v) < 2$, then one easily checks that we arrive to a contradiction with either (H1) or (H2). The second part of the property follows from (5).

(P3) also holds by construction: if there still exists a path of length at most two between two vertices from W_1, then this would contradict (H1).

(P4) follows from (H1), Claim 4.6, and our choice of K, since all the edges incident to the vertices in W_1 already exist in Γ' and as the number of vertices from W_2 that are in $N_G^2(v)$ of any vertex v can at most double in Γ.

Lastly, we show the expansion properties of vertices in U, i.e. property (P5). By (5) and the fact that $N_{\Gamma'}(S) \subseteq N_{\Gamma'}(S')$ whenever $S \subseteq S'$, it suffices to show

$$|N_{\Gamma'}(S)| \geq \begin{cases} |S|\sqrt{nq} + |S|, & \text{if } |S| < K/(10q), \\ (1/2 + \varepsilon/8)n + |S|, & \text{if } |S| = \lfloor K/(10q) \rfloor. \end{cases}$$

For every $v \in U$ we have

$$\begin{aligned} \deg_{\Gamma'}(v, U) &= \deg_{G_q}(v) - \deg_{H_q}(v) - \deg_{G_q}(v, W_1 \cup W_2) \\ &\geq (1 - 2\delta_a')nq - (1/2 - \varepsilon/2)nq - K/10 - 2 \\ &\geq (1/2 + \varepsilon/4)nq, \end{aligned}$$

by (H1), Claim 4.6, and our choice of K and δ_a'. Take $S \subseteq U$ to be an arbitrary subset of size $|S| \leq K/(10q)$ and let $T := N_{\Gamma'}(S) \cap U$. From the previously obtained bound on $\deg_{\Gamma'}(v, U)$, we have

$$e_{\Gamma'}(S, T) \geq |S|(1/2 + \varepsilon/4)nq.$$

Assume towards a contradiction that $|T| < 2|S|\sqrt{nq}$. Then, from $\Gamma'[U] \subseteq G_q[U]$ and (4) we derive

$$|S|(1/2 + \varepsilon/4)nq \leq e_{\Gamma'}(S, T) \leq |S||T|q + c'\sqrt{|S||T|nq} \leq 2|S|^2(nq)^{1/2}q + 2c'|S|(nq)^{3/4},$$

which is a contradiction. On the other hand, if $|S| = \lfloor K/(10q) \rfloor$ then, assuming $|T| < (1/2 + \varepsilon/8)n + \varepsilon n/16$, again from (4) we have

$$|S|(1/2 + \varepsilon/4)nq \leq e_{\Gamma'}(S, T) \leq |S||T|q + c'\sqrt{|S||T|nq} < |S|(1/2 + 3\varepsilon/16)nq + \varepsilon|S|nq/16,$$

where the last inequality follows from our choice of K. We have a contradiction once again.

To conclude, we obtained an $(\varepsilon/4, K/10, q)$-backbone graph Γ with the witness partition $V(\Gamma) = U \cup W_1 \cup W_2$ and at most $(\mu/2)n^2 p_0$ edges, and graphs G_q and H_q applicable by Lemma 4.3.

Note that, by the construction of Γ, the following is now true: if for any subset of edges $E' \subseteq E(G^-[U] - H)$ the graph $\Gamma + E'$ is Hamiltonian, then $G - H$ is Hamiltonian as well. Indeed, let $x_{uv} \in X$ be some vertex obtained by contracting a uv-path in Γ'. The two edges incident to x_{uv} in $\Gamma + E'$ necessarily lie on a Hamilton cycle. Moreover, as they correspond to two edges in $\Gamma' \subseteq G - H$, one incident to u and the other to v, by splitting every such vertex x_{uv} back into the uv-path we obtain a Hamilton cycle in $G - H$.

The following claim, akin to [4, Lemma 3.4], allows us to complete the proof.

Claim 4.7. [4] *If for every subset $E' \subseteq E(G^-[U] - H)$ of $|E'| \leq n$ edges such that $\Gamma + E'$ is not Hamiltonian there is a vertex $v \in U$ satisfying*

$$|N_G(v) \cap B_{\Gamma+E'}(v)| > \deg_H(v),$$

then $G - H$ is Hamiltonian. □

Let $E' \subseteq E(G^-[U] - H)$ be a set of edges of size at most $|E'| \leq n$. It is not too difficult to see that $\Gamma + E'$ is an $(\varepsilon/4, K/4, q)$-backbone graph. Indeed, none of the properties (P1)–(P3) can be violated by adding edges with both endpoints in U. Similarly, the expansion property (P5) is not affected by addition of edges. Lastly, property (P4) holds as $\Gamma + E' \subseteq G^-[U] - H$ and by referring to (H1). By Lemma 4.3 applied with G_q (for $\varepsilon/4$ as α) we get that the set of vertices $v \in U$ such that $|B_{\Gamma+E'}(v)| \geq (1/2 + \varepsilon/4)n$ is of size at least $(1/2 + \varepsilon/4)n$. As $e(\Gamma + E') \leq \mu n^2 p$, by Lemma 4.4 applied with $\varepsilon/4$ as α and G^- as G, there exists a vertex $v \in U$ with $|N_{G^-}(v) \cap B_{\Gamma+E'}(v)| > np_0/2 > \deg_H(v)$. Finally, Claim 4.7 implies that the graph $G - H$ is Hamiltonian. □

Having Proposition 4.5, the proof of the following theorem and Theorem 1.4 are identical to the proofs of Theorem 3.3 and Theorem 1.2, using [27, Theorem 1.1] instead of [33, Theorem 3.1] to handle $m \geq Cn \log n$.

Theorem 4.8. *Let $\varepsilon > 0$ be a constant and consider the random graph process $\{G_i\}$. There exist positive constants $\delta_t(\varepsilon)$, $\delta_a(\varepsilon)$, and $K(\varepsilon)$ such that a.a.s. for every $m \geq \frac{1+\varepsilon}{6} n \log n$ we have that the 2-core of G_m is $(1/2 - \varepsilon, \delta_t, 2, \delta_a, K)$-resilient with respect to being Hamiltonian.* □

4.3   Proof of Lemma 4.3 and Lemma 4.4

Let us first give a brief outline of how to apply Pósa's rotation-extension technique to backbone graphs with the goal of constructing long paths.

Assume Γ is a backbone graph and let $P = v_0, \ldots, v_\ell$ be a path in Γ. If $\{v_0, v_\ell\}$ is an edge in Γ then such a path can be closed into a cycle. If the obtained cycle does not cover all vertices, then one easily checks that properties (P2)–(P5) imply connectivity of Γ and we can thus *extend* P into a path longer than P.

Suppose that P cannot be extended and let $\{v_0, v_i\}$ be an edge in Γ for some $2 \leq i \leq \ell - 1$. Then the path $P' = v_{i-1}, \ldots, v_0, v_i, \ldots, v_\ell$ is another path in Γ of the same length ℓ. We say that P' is obtained from P by a *rotation* around the *endpoint v_0*, with *pivot point v_i*, and *broken edge* $\{v_{i-1}, v_i\}$ (see, Figure 1). Observe that by performing a rotation we can now possibly obtain a cycle by adding the edge $\{v_{i-1}, v_\ell\}$ as well. Otherwise, we perform more rotations to obtain more boosters. The rotation is repeated until we find a closing edge in Γ.
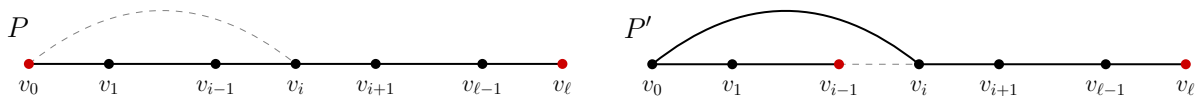

Figure 1: *Rotation of the the path P around the fixed endpoint v_0, with pivot point v_i, and the broken edge $\{v_{i-1}, v_i\}$. Pairs of red vertices denote the endpoints of the paths, i.e. boosters.*

In our setting, we have to be careful with vertices of low degree. This is why we need properties (P2)–(P4). We illustrate in particular why we need property (P3). Suppose $P = v_0, \ldots, v_\ell$ is a longest path in Γ such that $v_0, v_{i-1} \in W_1$ for some v_i which is the only neighbour (other than v_1) of v_0 on the path P. Rotating around v_0 with pivot point v_i yields a path $P' = v_{i-1}, \ldots, v_0, v_i, \ldots, v_\ell$. Now, the only rotation that can be performed around v_{i-1} brings the path P' back to where we started from. In conclusion, in such a situation we cannot prove that Γ contains many boosters. Property (P3) guarantees that this cannot occur.

Proof of Lemma 4.3. Let $P = v_0, \ldots, v_\ell$ be a longest path in Γ. For a subset of vertices $Z \subseteq V(P)$ we write $Z^+ := \{v_{i+1} \colon v_i \in Z\}$ and $Z^- := \{v_{i-1} \colon v_i \in Z\}$. For a vertex z we abbreviate $\{z\}^+$ to z^+ and $\{z\}^-$ to z^-.

We first show that there exists a longest path with both endpoints in U. Suppose $v_0 \in W_2$. Then by (P2) and (P4) we know that there are at least K neighbours $u \in N_\Gamma(v)$ on the path P which have $u^- \in U$. Thus, by performing a single rotation around v_0 with pivot point u we get a path P' with an endpoint in U. Similarly, if $v_0 \in W_1$ its only other neighbour on the path u has u^- belonging to either U or W_2, by (P3). In any case, by performing at most two rotations we reach a path P' with an endpoint in U. Repeating the same argument for v_ℓ yields the claimed path P^* with both endpoints in U.

Phase 1: Initial rotations

The first phase consists of repeatedly rotating the first endpoint as in [27, Lemma 3.2 **Step 1**] in order to obtain a set of at least $n/4$ new endpoints of paths of length ℓ that all contain the

same vertices and all end in v_ℓ. Compared to [27], the only difference in our argument is that we ignore all pivot points u with at least one of u, u^+, and u^- not belonging to U. As every vertex can have at most K such neighbours u, one easily sees that all the calculations of [27] essentially remain the same. Indeed, starting with $X_0 = \{v_0\}$ and denoting by X_i the set of endpoints obtained by exactly i rotations, we get

$$|X_{i+1}| \geq \frac{1}{2}\Big(|N_\Gamma(X_i)| - K|X_i| - 3\sum_{j=0}^{i} |X_i|\Big).$$

Using identical calculations as in [27] yields $|X_{i+1}| \geq (nq/4)^{(i+1)/2}$, for all $i \geq 0$ with $|X_i| \leq K/q$. Once we reach a set $|X_m| = \max\{1, \lfloor K/q \rfloor\}$, which is easily observed to happen after at most $O(\log n / \log \log n)$ many steps, the above bound on the size of X_{i+1}, together with property (P5), immediately implies $|X_{m+1}| \geq n/4$.

Phase 2: Terminal rotation

At the end of the first phase we have a set X of at least $n/4$ possible endpoints. Let Y denote the set of endpoints that can be generated by exactly one more rotation starting from X.

In [27, Lemma 3.2 **Step 2**] it is shown that by partitioning the path P into appropriately many intervals, denoted by P_i, one can define pairs of vertex sets (X_i, Y_i) such that there is no edge between X_i and Y_i in $G_q - H_q$, for every i. From Lemma 2.7 we know lower bounds on the number of edges in $E_{G_q}(X_i, Y_i)$ and thus all these need to belong to H_q. In [27, Lemma 3.2 **Step 2**] it is shown that $|Y| < (1/2 + \alpha)n$ implies a contradiction to the degree assumption of H_q. The only difference in our case is that we again need to ignore all pivot points $v^+ \in P_i$ with $v \notin U$ (which should belong to Y_i now), and similarly v^-. However, as every interval P_i (in line 4 of the proof of **Step 2**) is of size at least $n(\log \log n)^{1/2}/(4 \log n)$ and by (P1) each interval P_i has at most $o(|P_i|)$ such 'bad' vertices, we can simply ignore them since this contributes $o(n^2 q)$ many edges. All remaining calculations from [27, Lemma 3.2] remain the same and we obtain a set of new endpoints of size at least $(1/2 + \alpha)n$.

Phase 3: Rotating the other endpoint

Exactly as in [27, Lemma 3.2 **Step 3**] for every of the newly obtained $(1/2 + \alpha)n$ endpoints, we analogously rotate the other endpoint to obtain the intended result. $\qed$

We wrap-up by giving a proof of Lemma 4.4.

Proof of Lemma 4.4. Consider a graph Γ with at most $\mu n^2 p$ edges and a subset $U \subseteq V(G)$ such that $\Gamma[U] \subseteq G$ and
$$A := \{v \in U \colon |B_\Gamma(v) \cap U| \geq (1/2 + \alpha)n\}$$
is of size $|A| \geq (1/2 + \alpha)n$. Take $A' \subseteq A$ to be a set of size $|A'| = \alpha n/2$, and for every $v \in A'$ let $B'(v) := (B_{\Gamma'}(v) \cap U) \setminus A'$ and note that $|B'(v)| \geq (1/2 + \alpha/2)n$. By Chernoff bounds we have that the probability of a fixed vertex $v \in A'$ having fewer than $np/2$ neighbours inside of the set $B'(v)$ in the graph G is at most $e^{-\Omega_\alpha(np)}$. Since A' is disjoint from all the sets $B'(v)$, these events are independent for different vertices $v, u \in A'$. Therefore, the probability that no vertex $v \in A'$ has more than $np/2$ neighbours inside its respective set $B'(v)$ in G is bounded by $e^{-\Omega_\alpha(n^2 p)}$. Hence, the probability of failure is upper bounded by

$$\Pr \leq e^{-\Omega_\alpha(n^2 p)} \cdot \sum_{e(\Gamma') \leq \mu n^2 p} \Pr[\Gamma' \subseteq G].$$

The probability for a fixed graph with t edges to be a subgraph of G is p^t and thus by a union

bound over all such graphs, we get

$$\Pr \le e^{-\Omega_\alpha(n^2 p)} \cdot \sum_{t=1}^{\mu n^2 p} \binom{\binom{n}{2}}{t} p^t \le e^{-\Omega_\alpha(n^2 p)} \sum_{t=1}^{\mu n^2 p} \left(\frac{en^2 p}{t} \right)^t .$$

One easily sees that for $0 < \mu \le 1$ the right hand side is increasing for $1 \le t \le \mu n^2 p$ and hence we may substitute $t = \mu n^2 p$ to conclude

$$\Pr \le e^{-\Omega_\alpha(n^2 p)} \cdot (\mu n^2 p) \cdot \left(\frac{e}{\mu} \right)^{\mu n^2 p} = e^{-\Omega_\alpha(n^2 p)} \cdot e^{O(\mu \log(1/\mu)n^2 p)} = o(1),$$

for sufficiently small μ depending on α. □

References

[1] M. Ajtai, J. Komlós, and E. Szemerédi. First occurrence of Hamilton cycles in random graphs. *North-Holland Mathematics Studies*, 115:173–178, 1985.

[2] J. Balogh, B. Csaba, and W. Samotij. Local resilience of almost spanning trees in random graphs. *Random Structures & Algorithms*, 38(1-2):121–139, 2011.

[3] S. Ben-Shimon, M. Krivelevich, and B. Sudakov. Local resilience and Hamiltonicity Maker–Breaker games in random regular graphs. *Combinatorics, Probability and Computing*, 20(02):173–211, 2011.

[4] S. Ben-Shimon, M. Krivelevich, and B. Sudakov. On the resilience of Hamiltonicity and optimal packing of Hamilton cycles in random graphs. *SIAM Journal on Discrete Mathematics*, 25(3):1176–1193, 2011.

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

[6] B. Bollobás. Random graphs. In *Modern Graph Theory*, pages 215–252. Springer, 1998.

[7] B. Bollobás and A. Thomason. Random graphs of small order. *North-Holland Mathematics Studies*, 118:47–97, 1985.

[8] J. Böttcher, Y. Kohayakawa, and A. Taraz. Almost spanning subgraphs of random graphs after adversarial edge removal. *Combinatorics, Probability and Computing*, 22(5):639–683, 2013.

[9] J. Böttcher, M. Schacht, and A. Taraz. Proof of the bandwidth conjecture of Bollobás and Komlós. *Mathematische Annalen*, 343(1):175–205, 2009.

[10] R. Diestel. *Graph theory*. Springer-verlag, 2000.

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

[12] P. Erdős and A. Rényi. On the existence of a factor of degree one of a connected random graph. *Acta Math. Acad. Sci. Hungar.*, 17:359–368, 1966.

[13] P. Erdős and A. Rényi. On random graphs I. *Publicationes Mathematicae Debrecen*, 6:290–297, 1959.

[14] A. Frieze and M. Karoński. *Introduction to random graphs.* Cambridge University Press, 2015.

[15] A. Frieze and M. Krivelevich. On two Hamilton cycle problems in random graphs. *Israel Journal of Mathematics*, 166(1):221–234, 2008.

[16] E. N. Gilbert. Random graphs. *The Annals of Mathematical Statistics*, 30(4):1141–1144, 1959.

[17] A. Hajnal and E. Szemerédi. Proof of a conjecture of P. Erdős. *Combinatorial theory and its applications*, 2:601–623, 1970.

[18] H. Huang, C. Lee, and B. Sudakov. Bandwidth theorem for random graphs. *Journal of Combinatorial Theory, Series B*, 102(1):14–37, 2012.

[19] S. Janson, T. Łuczak, and A. Ruciński. *Random graphs*, volume 45. John Wiley & Sons, 2011.

[20] J. Komlós, G. N. Sárközy, and E. Szemerédi. On the Pósa-Seymour conjecture. *Journal of Graph Theory*, 29(3):167–176, 1998.

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

[22] M. Krivelevich. Long paths and Hamiltonicity in random graphs. *arXiv preprint arXiv:1507.00205*, 2015.

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

[24] M. Krivelevich, C. Lee, and B. Sudakov. Robust Hamiltonicity of Dirac graphs. *Transactions of the American Mathematical Society*, 366(6):3095–3130, 2014.

[25] M. Krivelevich and B. Sudakov. Pseudo-random graphs. *More sets, graphs and numbers*, pages 199–262, 2006.

[26] D. Kühn and D. Osthus. The minimum degree threshold for perfect graph packings. *Combinatorica*, 29(1):65–107, 2009.

[27] C. Lee and B. Sudakov. Dirac's theorem for random graphs. *Random Structures & Algorithms*, 41(3):293–305, 2012.

[28] T. Łuczak. On matchings and Hamiltonian cycles in subgraphs of random graphs. *North-Holland Mathematics Studies*, 144:171–185, 1987.

[29] T. Łuczak and A. Ruciński. Tree-matchings in graph processes. *SIAM Journal on Discrete Mathematics*, 4(1):107–120, 1991.

[30] R. Montgomery. Hamiltonicity in random graphs is born resilient. *arXiv preprint arXiv:1710.00505*, 2017.

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

[32] B. Sudakov. Robustness of graph properties. *Surveys in Combinatorics 2017*, 440:372, 2017.

[33] B. Sudakov and V. H. Vu. Local resilience of graphs. *Random Structures & Algorithms*, 33(4):409–433, 2008.