

On Independent Spanning Trees in Random and Pseudorandom Graphs

Nemanja Draganić* Keith Frankston† Michael Krivelevich‡ Alexey Pokrovskiy§
Liana Yepremyan¶

October 1, 2025

Abstract

In 1989, Zehavi and Itai conjectured that every k -connected graph contains k independent spanning trees rooted at any prescribed vertex r . That is, for each vertex v , the unique r - v paths within these k spanning trees are internally disjoint. This fundamental problem has received much attention, in part motivated by its applications to network reliability, but despite that has only been resolved for $k \leq 4$ and certain restricted graph families.

We establish the conjecture for almost all graphs of essentially any relevant density. Specifically, we prove that there exists a constant $C > 1$ such that, with high probability, the random graph $G(n, p)$ contains $\delta(G)$ independent spanning trees rooted at any vertex whenever $C \log n/n \leq p < 0.99$. Since the lower bound on p coincides (up to the constant C) with the connectivity threshold of $G(n, p)$, this result is essentially optimal. In addition, we show that (n, d, λ) -graphs with fairly mild bounds on the spectral ratio d/λ contain $(1 - o(1))d$ independent spanning trees rooted at each vertex, thereby settling the conjecture asymptotically for random d -regular graphs as well.

1 Introduction

Spanning trees T_1, \dots, T_k (with $k \geq 2$) in a graph G , rooted at a vertex r , are said to be *ISTs*, short for *independent spanning trees*, if, for every vertex $v \in V(G)$, the unique r - v paths in the T_i 's are internally vertex-disjoint; that is, they share no vertices other than r and v . These are also commonly known as vertex- or node-independent spanning trees¹. Independent spanning trees surface in an impressively wide range of settings: in computational biology, hypercube ISTs help to detect anomalies in mitochondrial DNA [28]; in high-performance computing, ISTs provide fault-tolerant broadcast and low-diameter routes for multidimensional torus interconnections [29]; carrier-grade IP networks leverage three edge-ISTs for fast rerouting under dual-link failures [14]; and emerging server-centric datacentre fabrics employ so-called completely independent trees to preserve high throughput despite switch outages [21]. We refer the interested reader to an extensive survey [6] by Cheng, Wang and Fan on the broad applications of ISTs in various fields ranging from biology to large-scale computing.

Probably the most famous conjecture in the area was made by Zehavi and Itai in the late '80s [32], essentially at the conception of the notion of ISTs. It is widely referred to as the *Independent Spanning Trees Conjecture*. To state it, we recall that an n -vertex graph G is k -connected if deleting any $k - 1 < n$ vertices does not disconnect the graph; the largest such k is the (vertex-) connectivity of G , denoted by $\kappa(G)$.

*Mathematical Institute, University of Oxford, UK. Research supported by SNSF project 217926. Email: nemanja.draganic@maths.ox.ac.uk

†Center for Communications Research – Princeton, USA. Email: k.frankston@fastmail.com

‡School of Mathematical Sciences, Tel Aviv University, Tel Aviv 6997801, Israel. Research supported in part by NSF-BSF grant 2023688. Email: krivelev@tauex.tau.ac.il

§Department of Mathematics, University College London, UK. Email: dralexeypokrovskiy@gmail.com

¶Department of Mathematics, Emory University, Atlanta, USA. Email: liana.yepremyan@emory.edu. Research is supported by the National Science Foundation grant 2247013: Forbidden and Colored Subgraphs.

¹Edge-independent trees are also widely studied in the literature and are related to edge-connectivity instead (see [6] for more details)

Conjecture 1.1 (Zehavi-Itai, '89). *Every graph G contains $\kappa(G)$ many ISTs rooted at r , for every choice of $r \in V(G)$.*²

This conjecture is also of theoretical importance due to being a qualitative strengthening of Menger's Theorem [23]. Indeed, given any pair of vertices u, v , if we have k independent spanning trees rooted at u , then the paths in these trees from u to v give k internally vertex-disjoint paths as guaranteed by Menger's Theorem.

This problem received a lot of attention, and has been studied in numerous papers and for various classes of graphs. Before the conjecture was formally stated in 1989, Itai and Rodeh [18] showed that the conjecture is true for $k = 2$. Soon after, Cheriyan and Maheshwari [7] showed that the conjecture holds for $k = 3$. Zehavi and Itai [32] independently arrived at the same result, and stated Conjecture 1.1 [32] in the same work. Much later, Curran, Lee, and Yu [8] settled the $k = 4$ case. Further results have been established for graphs with particular underlying topology, such as planar graphs [24] or cube-like graphs [30, 31].

Regarding general bounds, Censor-Hillel, Ghaffari, Giakkoupis, Haeupler and Kuhn [4] showed that k -connected n -vertex graphs contain $\Omega(k/\log^2 n)$ many ISTs via connected dominating sets (CDSs). Observe that one can naturally construct k ISTs from k disjoint connected dominating sets by selecting a spanning tree within each dominating set and attaching the remaining vertices of the graph as leaves. However, the CDS-based approach has inherent limitations: Censor-Hillel, Ghaffari, and Kuhn [5] proved that there exist k -vertex-connected n -vertex graphs containing only $O(k/\log n)$ many disjoint CDSs, so the best bound obtainable by this approach falls short of the conjectured bound by a logarithmic factor in n .

Nevertheless, for d -regular pseudorandom graphs, Draganić and Krivelevich [9] employed CDSs to show that such graphs contain $(1 - o(1))d/\log d$ ISTs, thereby removing the dependence on n for this class. In particular, this implies that random d -regular graphs contain $(1 + o(1))d/\log d$ disjoint CDSs; their method yields the same bound for binomial random graphs $G(n, d/n)$ when $d \gg \log n$. These bounds are asymptotically best possible via the CDS approach, since in $G(n, d/n)$ the smallest dominating set typically has size $(1 + o(1))n \log d/d$, thus there cannot be more than $(1 + o(1))d/\log d$ disjoint CDSs. They further conjectured that random graphs above the connectivity threshold satisfy, at least approximately, the IST Conjecture. Our first main result establishes this conjecture and essentially shows that the Zehavi-Itai Conjecture typically holds for random graphs above the connectivity threshold:

Theorem 1.2. *There exists $C > 1$ such that, for any $C \log n/n \leq p \leq 0.99$, if $G \sim G(n, p)$, then, **whp**, for every vertex r , there are $\delta(G)$ many ISTs rooted at r .*

As a classical result of Bollobás and Thomason [3] shows that typically $\kappa(G) = \delta(G)$ for any $p = p(n)$, this implies that the Zehavi-Itai conjecture holds for binomial random graphs in the relevant range. We note that we prove our result only for $p \leq 0.99$ for simplicity of presentation, and a slightly more careful argument would extend this to $p = 1 - o(1)$. For $p \ll \log n/n$ the random graph $G(n, p)$ is typically not connected, so the Zehavi-Itai conjecture holds trivially.

Very recently, independent from our work, an approximate version of Theorem 1.2 was established by Hollom, Lichev, Mond, Portier, and Wang [16], where, instead of finding $\delta(G)$ many ISTs, they showed that for every $\varepsilon > 0$ there exists $C = C(\varepsilon)$ such that, if $p \geq C \log n/n$ and $G \sim G(n, p)$, then, **whp**, for every vertex $r \in V(G)$, there exist $(1 - \varepsilon)\delta(G)$ such trees rooted at r in G .

We note that attaining the exact bound, as opposed to an asymptotic solution, requires a considerably more delicate argument. To give some intuition, in a collection of ISTs, every vertex not adjacent to the root must have at least $\delta = \delta(G)$ incident edges belonging to distinct trees. Consequently, for vertices whose degree is close to δ in $G(n, d/n)$, one must carefully plan how their edges are distributed among the trees. On the other hand, if one aims only to construct $(1 - \varepsilon)d$ ISTs, these complications are far less severe, and our proof can be considerably shortened in this case.

Our second main result resolves the Zehavi-Itai conjecture asymptotically for pseudorandom graphs, in particular, for (n, d, λ) -graphs.

Definition 1.3. An n -vertex graph G whose adjacency matrix has eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ is called an (n, d, λ) -graph if G is d -regular and $\max\{\lambda_2, |\lambda_n|\} \leq \lambda$.

²For convenience, we will sometimes simply say the Zehavi-Itai conjecture holds when we mean G has this property.

For readers less familiar with this notion, we refer to Section 4, in particular Lemma 4.7 (the Expander Mixing Lemma), which provides a concrete pseudorandomness corollary sufficient for our purposes. These graphs are known for their strong connectivity and expansion properties, making them well-suited for applications such as fault-tolerant network design [17]. Despite being defined deterministically, they closely mimic the behavior of random graphs in many key respects. The reader is invited to consult comprehensive survey [20] for a thorough discussion of properties of (n, d, λ) -graphs.

In Section 4, we establish the following asymptotic result for (n, d, λ) -graphs.

Theorem 1.4. *Fix $\varepsilon > 0$, and let G be an (n, d, λ) -graph with $d/\lambda = \omega(\log d)$ and $d = o(n/\log^2 n)$. Then G contains $(1 - \varepsilon)d$ vertex-independent spanning trees for each choice of root.*

Random d -regular graphs are typically (n, d, λ) -graphs with spectral ratio $d/\lambda = \Theta(\sqrt{d})$ (see [27]), hence our results imply that such random graphs with $d = o(n/\log^2 n)$ satisfy the Zehavi-Itai conjecture asymptotically as well. For $d = \omega(\log n)$, $G_1 \sim G(n, d)$ can be coupled with $G_2 \sim G(n, p)$ with $p = (1 - o(1))d/n$ such that whp $G_2 \subseteq G_1$ (see [13] and its bibliography). Thus Theorem 1.2 implies that the conjecture holds asymptotically for all larger d as well.

Theorem 1.5. *Fix $\varepsilon > 0$. Let G be a random d -regular graph for $d \geq d_0(\varepsilon)$. For every vertex $r \in V(G)$, there exist $(1 - \varepsilon)d$ edge-disjoint independent spanning trees rooted at r .*

Hollom et al. [16] also established the existence of $(1 - \varepsilon)d$ ISTs for random regular graphs $G(n, d)$, but their result requires $d \gg \log n$, whereas we achieve it for $d = \Omega(1)$. In the same work, they further showed that $\lfloor d/4 \rfloor$ ISTs exist for all d , with the caveat that the root r is not arbitrary but can be chosen from most of the vertices of G .

2 Preliminaries

2.1 Notation

Throughout this paper we adopt the usual graph-theoretic conventions. For a graph G , we write $V(G)$ for its vertex set and $E(G)$ for its edge set. For $S, B \subseteq V(G)$, $G[S]$ denotes the subgraph of G induced by S , $N_G(S)$ denotes the *external* neighbourhood of S , while $\Gamma_B(S)$ denotes all vertices in B adjacent to at least one vertex in S . Whenever it is clear from the context, we drop the subscript. Given two subsets $S_1, S_2 \subseteq V(G)$, $e_G(S_1, S_2)$ counts the edges with one endpoint in each. We write $d_G(v, S)$ to denote the number of neighbours of vertex v in the set S . We write $\delta(G)$ and $\Delta(G)$ for the minimum and maximum degree of G , respectively. Let $G(n, p)$ denote the binomial random graph on n vertices in which each of the $\binom{n}{2}$ possible edges appears independently with probability p . We say that an event \mathcal{E} holds *with high probability* (abbreviated **whp**) if $\Pr(\mathcal{E}) \rightarrow 1$ as $n \rightarrow \infty$. All logarithms are natural (base e) unless explicitly indicated otherwise.

2.2 Properties of random graphs

We first state Chernoff's inequality, before we proceed to collect some useful properties that hold for $G(n, p)$ **whp**:

Theorem 2.1 (Chernoff bounds). *Let $X \sim B(n, p)$ be a binomial random variable with n trials and success probability p . Then the following hold.*

- $P[X > (1 + \delta)np] \leq e^{-\delta^2 np / (2 + \delta)}$ for $\delta > 0$.
- $P[X < (1 - \delta)np] \leq e^{-\delta^2 np / 2}$ for $0 < \delta < 1$.

The next lemma states some typical properties of $G(n, p)$ that we will use in our proofs. Various expressions in the lemma are tailored for our needs and are not optimized.

Lemma 2.2. *Given sufficiently large $C > 1$, for any $C \log n/n \leq p \leq 0.99$, let $G \sim G(n, p)$. Denote by S the set of vertices of degree less than $np - 0.9\sqrt{2np \log n}$. Then, with the stated probabilities, G satisfies the following properties:*

- (a) **Min and max degree:** we have **whp** $\delta(G) = np - (1 \pm \varepsilon)\sqrt{2np(1-p) \log n}$, where $\varepsilon \rightarrow 0$ as $C \rightarrow \infty$; furthermore $\Delta(G) \leq 2pn$ with probability at least $1 - o(n^{-10})$.
- (b) **Low-degree set:** For $p = o(1)$, $|S| \leq n^{0.2}$ with probability $1 - o(n^{-3})$.
- (c) **Sparse boundary:** For $p \leq n^{-0.49}$, every vertex outside of S has at most \sqrt{np} neighbours in $S \cup N(S)$ with probability at least $1 - o(n^{-3})$.
- (d) **Common neighbors:** For $p \leq n^{-0.49}$, every pair of vertices u, v has at most \sqrt{np} common neighbours with probability at least $1 - o(n^{-3})$.
- (e) **Independence of low-degree set:** For $p \leq n^{-0.49}$ the set S is independent **whp**.
- (f) **Connectivity:** G is $\delta(G)$ -connected **whp**.

Proof.

- (a) See Chapter 3 in [2].
- (b) Consider a fixed set S of size $|n^{0.2}|$. Apply Chernoff bounds to get that the probability that a vertex has degree at most $np - 0.9\sqrt{2np(1-p) \log n}$ outside of S is at most $e^{-0.81(1-p) \log n} = n^{-0.81(1-p)}$. Note that these events are independent for all vertices in S . Hence the probability that all vertices in S have small degree outside is at most $(n^{-0.81(1-p)})^{|S|}$. Thus, by a union bound, we get that the probability of such a set existing is at most

$$\binom{n}{n^{0.2}} \cdot n^{-0.809n^{0.2}} \leq (en^{0.8})^{n^{0.2}} n^{-0.809n^{0.2}} \leq n^{-n^{0.2}/1000} \leq n^{-4}$$

Thus with probability $1 - o(n^{-3})$ there is no set S which has small neighbourhood outside of S , and in particular does not have a small neighbourhood in general.

- (c) Fix a vertex v . Expose all pairs in $G - v$. By (b), we have that the set S' of vertices of degree at most $np - 0.9\sqrt{2np(1-p) \log n}$ in $G - v$ is at most $n^{0.21}$, and thus $|N(S')| \leq 2n^{1.21}p$ with probability at least $1 - o(n^{-9})$, by the first part. Now expose the neighbours of v . The expected number of neighbours in $S' \cup N(S')$ is at most $pn^{1.21} \cdot p = n^{1.21}p^2 = o(\sqrt{np})$. Thus v has at most $o(\sqrt{np})$ neighbours in $S \cup N(S)$ with probability $1 - o(n^{-3})$ in case $np = \omega(\log^2 n)$. Otherwise, the probability that v has at least 100 neighbours in $S' \cup N(S')$ is at most $|S' \cup N(S')|^{100} p^{100} \leq n^{-5}$. Hence we are done by a union bound over all vertices v in both cases.
- (d) For a given pair of vertices u, v , the size of their common neighbourhood $N(u, v)$ is binomially distributed with parameters $(n - 2, p^2)$. Therefore, the claim follows from Chernoff-type bounds for the binomial distribution and the union bound over all pairs of vertices.
- (e) Fix a pair $u, v \in V(G)$. The probability that u (or v respectively) has at most $np - 0.9\sqrt{2np(1-p) \log n}$ neighbours in $V(G) - \{u, v\}$ is at most $n^{0.21}/n$ by Property (b). Thus the event that uv is an edge in S has probability at most $(n^{-0.79})^2 p = o(n^{-2})$. By a union bound over all pairs of edges, there is no edge in S **whp**.

□

A balanced random bipartite graph typically contains a perfect matching when above the connectivity threshold. Below is a proof with a bound on the success probability.

Lemma 2.3. *Let $G \sim G(n, n, p)$ be a binomial random bipartite graph with parts of size n , and edge probability $p \geq 40 \log n/n$. Then G has a perfect matching with probability at least $1 - n^{-\frac{pn}{2 \log n} + 4} \geq 1 - o(n^{-10})$.*

Proof. Denote the parts by A, B . If $G[A, B]$ has no perfect matching, then by Hall's theorem there is a set $A_0 \subset A$ of cardinality $|A_0| \leq n/2$ with $|N(A_0, B)| < |A_0|$, or a set $B_0 \subset B$ of cardinality $|B_0| \leq n/2$ with $|N(B_0, A)| < |B_0|$. The probability of this can be bounded from above by

$$\sum_{k=1}^{n/2} \binom{n}{k}^2 (1-p)^{k(n-k)} \leq \sum_{k=1}^{n/2} \left(\frac{en}{k}\right)^{2k} e^{-pkn/2} \leq \sum_{k=1}^{n/2} (3n)^{2k} n^{-pkn/(2 \log n)} \leq n (3n)^2 n^{-pn/(2 \log n)},$$

which implies the required bound as $n \rightarrow \infty$. \square

Lemma 2.4. *Let $G \sim G(n, p)$ where $p \geq 10^5 \log n/n$, let $k \leq n^{0.51}/2$, and let Q be a subset of $V(G)$ of cardinality k . Then with probability at least $1 - o(n^{-1})$ there are k many vertex disjoint paths each of length $n/50k$ such that every vertex of Q is an endpoint of one of these paths.*

Proof. We expose the random edges of $G(n, p)$ in bunches. Start with $Q_0 = Q$ and expose all the edges between Q_0 and $V(G) \setminus Q_0$, call the graph induced by these edges G_1 . Notice that $|V(G) \setminus Q_0| \geq n - n^{0.51}/2$, thus by Lemma 2.3 (by adding dummy vertices to Q_0 to be of the same size as $V(G) \setminus Q_0$) with probability $1 - o(n^{-2})$ we can find a matching M_1 saturating Q_0 and going into $V(G) \setminus Q_0$ in $G_1 \subset G$. We now let Q_1 be the endpoints of M_1 disjoint from Q_0 . Expose the edges between Q_1 and $V(G) \setminus (Q_0 \cup Q_1)$, and call the graph induced by these edges $G_2 \subset G$. Just like in the previous step, with probability $1 - o(n^{-2})$ there exists a matching M_2 between Q_1 and $V(G) \setminus (Q_0 \cup Q_1)$ in G_2 . Repeat this process $\ell = n/50k$ times. At the j -th step, we will still have $|V(G) \setminus \cup_{i=1}^{j-1} Q_i| \geq 49n/50$, thus again with probability $1 - o(n^{-2})$ there is a matching M_j saturating Q_{j-1} and going into $V(G) \setminus \cup_{i=1}^{j-1} Q_i$ in the subgraph $G_j \subset G$. Taking a union bound over all ℓ rounds, we get that with probability $1 - o(n^{-1})$ the union of these matchings $\{M_i\}_{i=1}^{\ell}$ gives us the desired collection of paths. \square

We will use the following randomized version of the previous lemma:

Lemma 2.5. *Let $G \sim G(n, p)$ where $p \geq 10^6 \log n/n$, let $k \leq n^{0.51}$, $\ell \leq n/100k$. For any $Q \subseteq W \subseteq V(G)$ with $|Q| = k$, $|W| \geq n/2$, we can define a randomized set $P \subseteq V(G)$ with the following properties:*

- *With probability at least $1 - o(n^{-1})$ there are k many vertex disjoint paths each of length ℓ such that every vertex of Q is an endpoint of one of these paths.*
- *The distribution of $P \setminus Q$ is that of a uniformly chosen random subset of W of order $k\ell$.*

Proof. Let σ be a random permutation of W chosen uniformly at random from all such permutations that fix Q . Note that σ^{-1} also has the same distribution. Let $G' = \sigma(G[W])$, noting that for distinct $x, y \in W$, xy is an edge of G' independently with probability p i.e. G' has the distribution of another Erdős-Renyi random graph on vertex set W . By Lemma 2.4, with probability $1 - o(n^{-1})$, G' contains a set k many vertex disjoint paths each of length $n/100k$ such that every vertex of Q is an endpoint of one of these paths. When this occurs, shorten these to have length ℓ , and let P' be the vertex set of these paths noting $|P'| = k\ell$. In outcomes when this doesn't occur, let P' be an arbitrary subset of W of order $k\ell$. Let $P = \sigma^{-1}(P')$. Since σ^{-1} is a uniformly chosen random permutation fixing Q , every subset of W of order $k\ell$ is equally likely to end up as P . \square

3 Proof of Theorem 1.2

We start by defining the notion of a *nice* collection of trees, which will be useful for building ISTs:

Definition 3.1. Given a graph G , let $\mathcal{S} = \{S_i\}_{i \in [t]}$ be a collection of trees rooted at the same vertex r , but pairwise disjoint otherwise. For every $v \in V(G)$, let $I(v) = \{i : v \notin S_i \cup N(S_i)\}$ be the set of indices of

trees which do not contain v or any of its neighbours. The collection \mathcal{S} is *nice* if, for all $v \in V(G)$ we can, for each $i \in I(v)$, find $u_i \notin \bigcup_{j \in [t]} S_j$ and $w_i \in S_i$ such that the u_i are all distinct and $v-u_i-w_i$ is a two-edge path in G .

The next result shows how to construct ISTs from a nice collection of trees.

Lemma 3.2. *If a graph G contains a nice collection of trees $\mathcal{T} = \{S_i\}_{i \in [t]}$ then it contains a collection of ISTs, $\{T_i\}_{i \in [t]}$, such that $S_i \subseteq T_i$ for each i .*

Proof. For each tree S_i , add all vertices from $N(S_i)$ to S_i by attaching each of them as a leaf to an arbitrary neighbour in S_i , to obtain tree S'_i .

Next, for every vertex v outside S'_i , add the edge vu_i guaranteed by the definition of a nice collection of trees where $u_i \notin \bigcup_{i \in [t]} S_i$. Note that that by doing this, for each i we produced a spanning tree T_i , as we were only adding leaves. To avoid any confusion, note that the particular edge $u_i w_i$ from Definition 3.1 need not appear in T_i ; in the first step of the proof, u_i may instead connect to some other vertex in S_i ; the definition only guarantees that such an edge exists. Let us prove now that T_1, \dots, T_t is a vertex-independent collection.

For this, fix two trees T_i, T_j , and let us prove that for every v the paths between r and v in T_i and T_j are internally vertex disjoint. We have the following cases:

- $v \in S_i \cup N(S_i)$, and $v \in S_j \cup N(S_j)$: In this case all internal vertices in the $r-v$ path in T_i are contained in S_i , while the internal vertices in the $r-v$ path in T_j are contained in S_j , completing this case.
- $v \in S_i \cup N(S_i)$, and $v \notin S_j \cup N(S_j)$: All internal vertices of the first path are in S_i , while the second path has all internal vertices in $S_j \cup \{u_j\}$ where $u_j \notin \bigcup_{i \in [t]} S_i \supset S_i$, proving this case.
- $v \notin S_i \cup N(S_i) \cup S_j \cup N(S_j)$: Apart from vertices u_i, u_j , the internal vertices of the two paths lie in S_i and S_j respectively. Since u_i and u_j are distinct by construction, this completes the proof.

□

Proof of Theorem 1.2. We consider two cases depending on p . **The dense regime:** $\log^2 n / \sqrt{n} \leq p \leq 0.99$.

This is the simpler case, so it serves as a good warm-up. We prove that the following claim holds, which will imply the statement in this case.

Claim 3.1. *The following holds with probability $1 - o(1/n)$ for every integer $k \in [np/2, np]$. For every ordered pair of non-adjacent vertices (u, v) in G , let K be the set of u 's first $\min\{k, d(u)\}$ neighbours (according to the natural ordering on $[n]$). There exists a matching between $K \setminus N(v)$ and $N(v) \setminus K$ which covers the smaller side.*

Before we show the proof of the claim, let us see how it implies the statement of the theorem in this case. Suppose G is a graph for which the conclusion of the claim holds. Denote $\delta = \delta(G)$, and note that by Lemma 2.2 we have that **whp** $\delta \in [np/2, np]$ with room to spare.

Let u be an arbitrary vertex in G , fixed to be the root. Apply the claim with $k = \delta$, and let $K = \{v_1, \dots, v_\delta\}$ be the first δ neighbours of u . Define $S_i = \{u, v_i\}$. Then the claim exactly shows that the S_i 's are a nice collection of trees, which completes the proof by Lemma 3.2. Indeed, to verify that Definition 3.1 applies to our tree collection, note that if $v \in N(u)$ then $I(v) = \emptyset$, as v is in the neighbourhood of the root u . Otherwise, let $I(v) = \{i : v_i \notin N(v)\}$. Note that since $|N(v)| \geq \delta = |K|$, we have $|K \setminus N(v)| \leq |N(v) \setminus K|$, so there is a matching in G which covers $K \setminus N(v)$ while the other endpoints are in $N(v) \setminus K \subseteq V(G) \setminus \bigcup_{i \in [\delta]} S_i$. For each $v_i \in K \setminus N(v)$, let $u_i \in N(v) \setminus K$ be the vertex it is matched to. Noting that $K \setminus N(v) = \{v_i : i \in I(v)\}$, we have now defined distinct vertices u_i for all $i \in I(v)$ such that $v-u_i-v_i$ is a two-edge path from v to $v_i \in S_i$.

Proof of Claim 3.1. Fix u and v , and expose the neighbourhood of u . If v is adjacent to u then there is nothing to prove. By assumption on p and Chernoff bounds, we have that with probability at least $1 - o(n^{-8})$ we have $d(u) \geq np/2$. Let K be the set of u 's first $\min\{k, d(u)\} \geq \sqrt{n} \log n$ neighbours. We

now expose the neighbours of v ; by the upper bound on p , we know that $|N(v) \setminus K| \geq d(v)/10^4 \geq \sqrt{n} \log n$ with probability at least $1 - o(n^{-8})$. Similarly $|K \setminus N(v)| \geq |K|/10^4 \geq \sqrt{n} \log n$. Now, expose all edges between these two disjoint sets $K \setminus N(v)$ and $N(v) \setminus K$. We claim that there is a matching saturating the smaller side with probability at least $1 - o(n^{-4})$, which completes the proof by a union bound over all pairs u, v and all of the at most n choices for k . To see this, choose A and B to be subsets of $N(v) \setminus K$ and $K \setminus N(v)$ respectively of size $n_0 = |A| = |B| = \min\{|N(v) \setminus K|, |K \setminus N(v)|\}$. Since $n_0 \geq \sqrt{n} \log n$ and $p \geq \log^2 n / \sqrt{n} \geq 40 \log n_0 / n_0$ we have by Lemma 2.3 that there is a required matching between A and B with probability at least $1 - o(n_0^{-10}) > 1 - o(n^{-5})$. \square

The sparse regime: $C \log n / n \leq p \leq \log^2 n / \sqrt{n}$.

Let $p_2 = \frac{10^6 \log n}{n}$, $p_3 = p/1000$, and p_1 be defined by $1 - p = (1 - p_1)(1 - p_2)(1 - p_3)$, notice that $p_1 \sim 0.999p$. Let $G_i \sim G(n, p_i)$, and observe that if $G \sim G(n, p)$ then $G = G_1 \cup G_2 \cup G_3$.

Proof outline. We first outline the proof for a fixed root $r \notin S \cup N_G(S)$; the full argument requires slightly more careful probability estimates, since in the end we perform a union bound over all vertices. We expose the edges of G in stages, one graph G_i at a time for a chosen subset of pairs. First we expose the whole graph G_1 , which typically contains 99.9% of all edges. We then identify the set S of vertices of small degrees in G_1 . Since G_1 contains a big majority of the edges, we can say that after we eventually reveal all of $G_2 \cup G_3$, the vertices with smallest degrees in G are likely to be in S . Next we reveal the edges of $G_2 \cup G_3$ touching S . Notice that all edges of G touching S have now been revealed. Denote $\delta = \min\{d_G(v) \mid v \in S\}$. As indicated before, we expect δ to be the minimum degree of G .

Next we pick an arbitrary vertex $r \notin S \cup N_G(S)$, and expose all edges containing it in $G_2 \cup G_3$. We will argue that **whp** the G -degree of r outside $S \cup N_G(S)$ is at least δ , and we pick δ neighbours to be the initial edges of the required δ trees. Then we expose the edges of G_2 outside of $V(G) - (S \cup N_G(S))$. The density of G_2 is sufficient to find **whp** a family of δ many paths in G_2 , denoted by \mathcal{P} , all rooted at r and pairwise disjoint otherwise, and of length $\Omega(\log n / np^2)$. We hope to show that these paths (viewed as trees) are nice in G **whp**. Since the set P of all vertices in these paths can essentially be viewed as a random set of a given size, by standard concentration bounds, typically no vertex has a lot of its G_1 -degree into P . After that we expose the neighbours in G_3 of every vertex $v \in V - S$ towards $R = V(G) - S - P$. As $S \cup P$ is small, we can show that for all such v **whp** $d_{G_1 \cup G_3}(v, R) \geq \delta$. For every $v \in V(G) - S$, denote its neighbourhood in R by B_v . Also, for every $v \in S$, denote $B_v := N_G(v)$. Thus, **whp** for every $v \in V(G)$, B_v is a set of cardinality of size at least δ outside of P . Now, for each $v \in V(G)$ reveal the edges of G_3 between P and B_v . Let \mathcal{P} be the set of all paths, and H_v be the auxiliary bipartite graph between \mathcal{P} and B_v , with an edge connecting P_i to $u \in B_v$ if there is an edge between P_i and u in G_3 . As the number of paths is $|\mathcal{P}| = \delta \leq |B_v|$, and the expected degree of each vertex in this auxiliary graph is at least $50 \log n$, by a standard Hall-type argument, there is a matching covering \mathcal{P} with probability $1 - o(1/n)$. By a union bound, this holds for every vertex. This is exactly the setup for Lemma 3.2, which will complete the proof.

Now we fill in the details of the proof stage by stage.

Exposing G_1 and identifying small degree vertices and their neighbours. We first expose all edges in G_1 . By Lemma 2.2(b) we have that the set of vertices S of degree at most $np_1 - 0.9\sqrt{2np_1 \log n}$ is of size at most $|S| \leq n^{0.2}$.

Note that, by Chernoff bounds and a union bound, with probability $1 - ne^{-4.1 \log n / 2} = 1 - o(n^{-1})$ all vertices in $V(G)$ have degree at least

$$n(p_2 + p_3) - \sqrt{4.1n(p_2 + p_3) \log n}$$

in $G_2 \cup G_3$. Hence with the same probability vertices outside S will have degree in G at least

$$np_1 - 0.9\sqrt{2np_1 \log n} + n(p_2 + p_3) - \sqrt{4.1n(p_2 + p_3) \log n} \geq np - 0.95\sqrt{2np \log n}, \quad (1)$$

where we used $0.9 \cdot \sqrt{0.999} + \sqrt{2.05 \cdot 0.001} < 0.95$. Note that here we just calculated the probability of the mentioned event, we did not yet expose the randomness of G_2 and G_3 .

Denote $\delta = \min\{d_G(v) \mid v \in S\}$, anticipating that the lowest degree vertex will **whp** be in S , as by Lemma 2.2 the minimum degree in G will be $np - (1 \pm o_C(1))\sqrt{2np \log n}$.

Now expose all edges in $G_2 \cup G_3$ touching S . By Lemma 2.2(a) we can assume that the set $N_G(S) = N_{G_1 \cup G_2 \cup G_3}(S)$ has size at most $2np|S|$, and by (e) that S is independent.

Fixing a root and finding a potentially nice collection of trees. So far we only exposed G_1 , and the edges touching S in G . Fix an arbitrary root $r \in V(G)$. Our aim is to show that with probability at least $1 - o(n^{-1})$, the vertex r is a root of a collection of δ many ISTs. Then we would be done by a union bound over all n vertices.

If r is not in S , also expose all edges containing r in $G_2 \cup G_3$. In that case, since by Lemma 2.2(c) with probability $1 - o(n^{-1})$ the root r has at most \sqrt{np} neighbours in $S \cup N(S)$, by Equation (1) with the same probability r has least δ neighbours outside $S \cup N(S)$. Fix a subset Q of δ such neighbours if $r \notin S$, and otherwise if $r \in S$, then let Q be a set of δ arbitrary vertices in $N(r)$; note that in the latter case $N(r) \subseteq N(S)$ as S is independent by Lemma 2.2(e). The edges from r to Q will be edges which belong to distinct trees in the nice collection we are about to find.

Next, we expose all edges of G_2 not touching $S \cup N(S)$ or r . We define the set W of vertices which we initially use to find our nice collection of trees: $W := V(G) - (S \cup N(S) \cup \{r\}) + Q$.

Apply Lemma 2.5 to $G = G(n, p_2)$, $W, Q, k = \delta, \ell = \lceil \frac{10^5 \log n}{np^2} \rceil$, to get a set P such that with probability $1 - o(n^{-1})$, there is collection \mathcal{P} of vertex-disjoint paths P_1, \dots, P_δ , each having exactly one endpoint in Q (denote by r_i the endpoint of P_i), of length $\lceil \frac{10^5 \log n}{np^2} \rceil$, and having $V(\mathcal{P}) = P \cup Q$. Note that

$$|P| = \delta \left\lceil \frac{10^5 \log n}{np^2} \right\rceil \leq \max \left\{ \frac{2 \cdot 10^5 \log n}{p}, np \right\} \leq 2 \cdot 10^5 n / C,$$

since $\delta \leq np$, where the second inequality considers the two cases when the expression under the ceiling is either less or more than 1.

From Lemma 2.5, the set P is a uniformly at random chosen set of size $\delta \lceil \frac{10^5 \log n}{np^2} \rceil$ in $W - Q$. More formally, if E is the event that in the random graph $G_2[W]$ the collection \mathcal{P} exists, then we have that $Pr[P = S_1 \mid E] = Pr[P = S_2 \mid E]$ for every two sets S_1, S_2 of size $\delta \lceil \frac{10^5 \log n}{np^2} \rceil$, as every vertex is equally likely to be included in P .

Thus in the previously exposed graph G_1 , with probability $1 - o(n^{-1})$ every vertex $v \in V(G)$ has at most $\max\{10^6 \log n, 2np^2\}$ neighbours in $P - Q$. This is because the number of neighbours of v in P is distributed hypergeometrically $\text{Hyp}(|W| - |Q|, K, m)$ where $K < 1.1np$, and $m \leq \max\{\frac{2 \cdot 10^5 \log n}{p}, np\}$; note that $|W| = (1 - o(1))n$. Therefore, the expected number of such neighbours is at most $1.2np \cdot \max\{\frac{2 \cdot 10^5 \log n/p}{n}, p\} \leq \max\{3 \cdot 10^5 \log n, 1.2np^2\}$, so the claim follows by Chernoff-type bounds for the hypergeometric distribution (see Theorem 2.10 in [19]) and a union bound over all vertices. Also by Lemma 2.2(d), for every vertex $v \in V(G) - \{r\}$, the degree of v into Q is at most \sqrt{np} .

Exposing G_3 to prove niceness In this part of the proof, we will use the remaining randomness to prove that the collection of paths $\{P_i + \{r_i, r\}\}_{i \in [\delta]}$ is nice, which by Lemma 3.2 is enough to complete the proof. Consider a vertex $v \in V(G) - \{r\}$. If it is not in S already, expose its G_3 -neighbours in $R_v := V(G) - P - Q - S - N_{G_1}(v)$. Since $|R_v| \geq n - 2 \cdot 10^5 \log n / p - n^{0.2} - 4np$, we get by Chernoff bounds

that with probability $1 - o(n^{-1})$ for every such v its G_3 -neighbourhood in R_v is of size at least

$$np_3 - \sqrt{4.1np_3 \log n} - \max \left\{ 10 \left(\frac{2 \cdot 10^5 \log n}{p} + n^{0.2} + 4np \right) p_3, 100 \log n \right\} \\ \geq np_3 - \sqrt{4.2np_3 \log n},$$

where we used $\frac{C \log n}{n} < p < \frac{\log^2 n}{\sqrt{n}}$, and the $100 \log n$ term ensures that the Chernoff bound gives the required error probability even in cases where the other expression inside the maximum is small. Thus the combined degree of v in $G_1 \cup G_3$ into R_v is at least

$$np_1 - 0.9\sqrt{2np_1 \log n} - (10^6 \log n + 2np^2 + 2\sqrt{np}) \\ + np_3 - \sqrt{4.2np_3 \log n} \geq np - 0.99\sqrt{2np \log n} > \delta(G),$$

where we used that the degree of v into $P \cup Q$ is at most $10^6 \log n + 2np^2 + \sqrt{np}$ and the degree of v into S is at most \sqrt{np} .

For each $v \notin S \cup \{r\}$, denote $B_v = R_v \cap N_{G_1 \cup G_3}(v)$; hence for such v we have $|B_v| \geq \delta$; secondly, for each $v \in S - \{r\}$, denote by $B_v := N_G(v) \setminus N_G(r)$. Thus, in the first case $|B_v| \geq \delta \geq |I(v)|$ (recall the definition of $I(v)$ from Definition 3.1), and in the second one clearly $|B_v| \geq \delta - |N_G(r) \cap N_G(v)| \geq |I(v)|$, as every vertex in $N_G(r)$ already belongs to a distinct tree in our collection.

Fix any $v \in V(G) - \{r\}$. If $v \notin S$, denote $I = [\delta]$; if $v \in S$, let $I \subseteq [\delta]$ be the set of indices i for which $N_G(v)$ does not contain r_i . Now consider the auxiliary bipartite graph H_v with one part $A = \{P_i\}_{i \in I}$ and the other part B_v ; recall again that $|B_v| \geq |A|$. Note also that by Lemma 2.2(d) we have $|A| \geq \delta - \sqrt{np} \sim \delta$. There is an edge between $P_i \in A$ and $u \in B_v$ in H_v if there is any edge between P_i and u in the graph G_3 . For each such pair, we expose the edges between P and B_v in G_3 .

We want to find a matching in H_v which covers A , and then we would be done. Indeed, the collection $\{P_i \cup \{r_i, r\}\}_{i \in [\delta]}$ would satisfy Definition 3.1, as we would have the appropriate paths of length two for each $i \in I(v)$. Consider a set $B \subseteq B_v$ of size $n_0 = |B| = |A| \sim \delta \geq (1 - o_C(1))np$, and let us argue that with probability $1 - o(n^{-2})$ there is a perfect matching in $H_v[A, B]$; then we would be done by a union bound over all n choices of v .

There is *no edge* in our auxiliary graph between P_i and w with probability

$$1 - p_0 := (1 - p_3)^{|P_i|} \leq 1 - p_3|P_i|/2 \leq 1 - 50 \log n / np.$$

Thus each edge in H_v is there with probability at least $p_0 \geq 50 \log n / np \geq 40 \log n / n_0$ independently. Hence we can invoke Lemma 2.3 to get a perfect matching in $H_v[A, B]$ with probability at least

$$1 - n_0^{n_0 p_0 / (2 \log n_0) - 4} \geq 1 - e^{40 \log n / 2 - 4 \log n_0} \geq 1 - o(n^{-2})$$

with room to spare. This completes the proof, as each one of the steps we performed holds with probability at least $1 - o(n^{-1})$, as required for a union bound over all choices of roots r . \square

4 Pseudorandom graphs: Proof of Theorem 1.4

In this section we prove our asymptotically optimal result for pseudorandom graphs, Theorem 1.4. Before we start with the main part of the proof, we will need to introduce some machinery related to embedding trees into expander graphs.

4.1 Trees in expanding graphs

Here we describe the tree-embedding method based on work of Friedman and Pippenger, often called the extendability approach. Drawing on the foundational embedding theorems of Friedman–Pippenger [12] and

Haxell [15], it allows us to build trees of size linear in the number of vertices inside an expander graph. Our main lever is Theorem 4.3, which ensures the existence of such embeddings, and we augment it with the straightforward Lemma 4.5, which shows how to prune leaves from the tree without breaking its extendability. Maintaining extendability in this way is what allows us to apply Theorem 4.3 repeatedly. This extendability framework has been the key ingredient in settling many classic open problems in graph theory (see, e.g., [10, 11, 25]). Here we follow the presentation of Montgomery [25]; further details are given below.

We use the following connectivity condition.

Definition 4.1. A graph G is m -joined for $m > 0$, if for every two disjoint vertex subsets $A, B \subseteq V(G)$ of size at least m we have that $e(A, B) > 0$.

We also need the notion of an (m, D) -extendable embedding from [25].

Definition 4.2. Let $m, D \in \mathbb{N}$ satisfy $D \geq 3$ and $m \geq 1$, let G be a graph, and let $S \subset G$ be a subgraph of G . We say that S is (m, D) -extendable if S has maximum degree at most D and

$$|\Gamma_G(U) \setminus V(S)| \geq (D-1)|U| - \sum_{x \in U \cap V(S)} (d_S(x) - 1) \quad (2)$$

for all sets $U \subset V(G)$ with $|U| \leq 2m$.

The next theorem is the key technical tool in the section. It allows us to extend an (m, D) -extendable embedding of a graph by attaching to one of its vertices a tree of certain size and maximum degree at most $D/2$.

Theorem 4.3 (Corollary 3.7 in [25]). *Let $m, D \in \mathbb{N}$ satisfy $D \geq 3$ and $m \geq 1$, and let T be a tree with maximum degree at most $D/2$, which contains the vertex $t \in V(T)$. Let G be an m -joined graph and suppose R is an (m, D) -extendable subgraph of G with maximum degree $D/2$. Let $v \in V(R)$ and suppose $|R| + |T| \leq |G| - 2Dm - 3m$. Then, there is a copy S of T in $G - (V(R) \setminus \{v\})$, in which t is copied to v , such that $R \cup S$ is (m, D) -extendable in G .*

The following result shows that adding an edge between two vertices of an (m, D) -extendable graph maintains its extendability.

Lemma 4.4 (Lemma 3.9 in [25]). *Let $m, D \in \mathbb{N}$ satisfy $D \geq 3$ and $m \geq 1$, let G be a graph, and let S be an (m, D) -extendable subgraph of G . If $s, t \in V(S)$ with $d_S(s), d_S(t) \leq D-1$ and $st \in E(G)$, then $S + st$ is (m, D) -extendable in G .*

Finally, the last result we need is another simple corollary of the definition of (m, D) -extendability which allows for leaf removal.

Lemma 4.5 (Lemma 3.8 in [25]). *Let $m, D \in \mathbb{N}$ satisfy $D \geq 3$ and $m \geq 1$, let G be a graph, and let S be a subgraph of G . Furthermore, suppose there exist vertices $s \in V(S)$ and $y \in N_G(S)$ so that the graph $S + ys$ is (m, D) -extendable. Then S is (m, D) -extendable. \square*

Remark. Let us remark here that Theorem 4.3, Lemma 4.4 and Lemma 4.5 are not algorithmic in the form they are stated. But, in [10] an algorithmic version is presented for (n, d, λ) -graphs. Hence, every step of the proof which uses these three results can also be made into a polynomial time algorithm. More precisely, removing and deleting edges while preserving the property of (m, D) -extendability can be implemented in polynomial time. We will also use the Lovasz Local Lemma in a few instances to prove the existence of certain sets — these too can be made into a deterministic polynomial time algorithm [26].

4.2 Disjoint paths with prescribed vertex-subsets in (n, d, λ) -graphs

We'll need the Lovasz Local Lemma (see Lemma 5.1.1 in [1])

Lemma 4.6 (Asymmetric Lovasz Local Lemma). *Let $\mathcal{A} = \{A_1, \dots, A_n\}$ be a finite set of events in the probability space. For $A \in \mathcal{A}$ let $\Gamma(A)$ denote the neighbours of A in the dependency graph (In the dependency graph, event A is not adjacent to events which are mutually independent). If there exists an assignment of reals $x : \mathcal{A} \rightarrow [0, 1)$ to the events such that*

$$\forall A \in \mathcal{A} : \Pr(A) \leq x_A \prod_{B \in \Gamma(A)} (1 - x_B)$$

then the probability of avoiding all events in \mathcal{A} is positive.

We will use the following version of the Expander Mixing Lemma for (n, d, λ) -graphs (see for example, Corollary 9.2.5 in [1]) which asserts that the number of edges between any two vertex sets asymptotically matches the expected count in a random graph.

Lemma 4.7 (Expander Mixing Lemma). *Let G be an (n, d, λ) -graph and let $A, B \subseteq V(G)$. Then*

$$\left| e(A, B) - \frac{|A||B|d}{n} \right| \leq \lambda \sqrt{|A||B|}.$$

Here $e(A, B)$ represents the count of ordered pairs (a, b) where $a \in A, b \in B$, and ab forms an edge in graph G . It is worth noting that A and B can overlap; specifically, $e(A, A)$ is twice the number of edges within any subset $A \subseteq V(G)$. The next statement is a corollary of the Expander Mixing Lemma. An elementary proof can be found in [9].

Lemma 4.8. *Let G be an (n, d, λ) -graph. Then the following hold:*

- (a) *For each two disjoint sets of vertices X, Y of size $|X|, |Y| > \frac{\lambda n}{d}$ it holds that $e(X, Y) > 0$.*
- (b) *Fix $0 < \varepsilon \leq 1$. Let $B, Z \subseteq V(G)$ be such that every vertex in B has at least $\varepsilon d/3$ neighbours in Z . Then for every k with $1 < k < (\frac{d}{12\lambda})^2$, and for every subset $X \subseteq B$ of size at most $\frac{n}{12k}$, it holds that $|\Gamma_Z(X)| \geq \varepsilon^2 k |X|$.*

The following result, the goal of this subsection, states that we can connect subsets of vertices in a pseudorandom graph through a subset of vertices which contains a fraction of the neighbourhood of every vertex. The proof closely follows the one in [9].

Lemma 4.9. *Let $0 < \varepsilon < 10^{-3}$, and let $C \geq C_0(\varepsilon)$. Let G be an (n, d, λ) -graph with $d/\lambda \geq C$. Let R, S_1, \dots, S_t be pairwise disjoint vertex-subsets of size at least two, such that each vertex in $S_1 \cup \dots \cup S_t \cup R$ has at least εd neighbours in R . If $\sum_{i=1}^t |S_i| \log(n/|S_i|) \leq \varepsilon n/100$, then there exist t vertex-disjoint paths P_i , where P_i contains S_i , and $V(P_i) \subseteq S_i \cup R$.*

Proof. By Lemma 4.8 applied with $B := S_1 \cup \dots \cup S_t \cup R$, and $Z := R$ and $k = 1/\varepsilon^3$, we get that $|\Gamma_R(X)| \geq |X|/\varepsilon$ for every $X \subseteq B$ of size at most $|X| \leq \varepsilon^3 n/12$.

Consider the graph $G' = G[R \cup S_1 \cup \dots \cup S_t]$, and observe that $|G'| \geq \varepsilon n/2$ by the Expander Mixing Lemma, as

$$|G'| \varepsilon d \leq e(V(G'), V(G')) \leq |G'|^2 \frac{d}{n} + \lambda |G'|,$$

which together with $\lambda < \varepsilon d/2$ implies the claim. Note that $G' \subset G$ is m -joined for

$$m = \varepsilon^5 n \geq \lambda n/d + 1$$

by Lemma 4.8. Furthermore, by the above we have that the independent set $I := I[S_1 \cup \dots \cup S_t]$ is an $(m, 10)$ -extendable subgraph of G' . Indeed, by the first paragraph of the proof

$$|\Gamma_G(U) \setminus V(I)| = |\Gamma_{G'}(U)| \geq |U|/\varepsilon \geq 10|U| - \sum_{x \in U \cap V(I)} (d_I(x) - 1),$$

for all $|U| \leq 2m$, since $d_I(x) \geq 0$ for all $x \in V(I)$ (in fact, $d_I(x) = 0$). Hence, we can use Theorem 4.3, Lemma 4.5, and Lemma 4.4 to attach and remove leaves from $V(I)$ in G' , as long as the total size of the current forest is at most

$$s := \varepsilon n/4 \leq |G'| - 20m - 3m.$$

We connect the sets S_i one by one through R using the given machinery, by always adding a path between two vertices in S_i . We will now show how to construct the required path for S_1 , and later we will argue how to do the same for the remaining S_i 's analogously.

First, we extend the $(m, 10)$ -extendable subgraph I in G' by attaching disjoint binary trees of size $\frac{s}{3|S_1|}$ (and of depth at most $\log_2 \frac{s}{3|S_1|}$) to each vertex in S_1 . Note that this can be done using Theorem 4.3, since the total size of the constructed forest (including all vertices in $S := S_1 \cup \dots \cup S_t$ whose number is at most $\varepsilon n/100$) is at most $s/3 + |S| \leq s/2$. Now, consider the set L_1 of the vertices of an arbitrary choice of $\lfloor |S_1|/2 \rfloor$ of the trees, and L_2 the vertices of the remaining trees. Since both sets contain at least $s/9 > \lambda n/d$ vertices each, there is an edge between the two sets by Lemma 4.8(a). Adding this edge to the current $(m, 10)$ -extendable graph creates a path P of length at most $2 \log_2 \frac{s}{3|S_1|} + 1$ between two vertices in S_1 , while the constructed graph stays $(m, 10)$ -extendable by Lemma 4.4. Now we can use Lemma 4.5 to remove all vertices (leaf-by-leaf) of the attached trees, except those in the path P , so that we are again left with an $(m, 10)$ -extendable subgraph.

We now repeat this process, until we have connected all vertices in S_1 . More precisely, we perform the following inductively described procedure. Suppose that after the i -th step, for $0 < i < |S_1| - 1$, we have constructed internally vertex-disjoint paths P_1, P_2, \dots, P_i with endpoints in S_1 , that satisfy the following invariants:

- the internal vertices are in R ;
- the length of P_j for all $j \leq i$ is at most $2 \log_2 \left(\frac{n}{|S_1| - j + 1} \right) + 1$;
- The graph $I[S] \cup P_1 \cup \dots \cup P_i$ is a linear forest and an $(m, 10)$ -extendable subgraph of G' ;
- the number of components (i.e., paths) in $I[S_1] \cup P_1 \cup \dots \cup P_i$ is $|S_1| - i$.

In the $(i + 1)$ -th step we want to construct a path P_{i+1} that maintains the invariants. In the end, after the $(|S_1| - 1)$ -th step, we will then have the required path which contains all vertices in S_1 .

For each path in the linear forest $I[S_1] \cup P_1 \cup \dots \cup P_i$, choose one endpoint, and let S'_1 denote the set of these endpoints, noting that $|S'_1| = |S_1| - i$. As before, extend the $(m, 10)$ -extendable subgraph $I[S] \cup P_1 \cup \dots \cup P_i$ by attaching a binary tree to each vertex in S'_1 , each one of size $\frac{s}{3|S'_1|}$, and hence of depth at most $\log_2 \left(\frac{s}{|S'_1|} \right) \leq \log_2 \left(\frac{n}{|S_1| - i} \right)$. We can indeed use Theorem 4.3, as the current forest consists of S and the now constructed trees (which together contain at most $s/2$ vertices), and the previously found paths P_1, \dots, P_i , whose total number of vertices is at most:

$$\begin{aligned}
\sum_{1 \leq j \leq i} \left(2 \log_2 \left(\frac{n}{|S_1| - j + 1} \right) + 1 \right) &\leq |S_1| + 2 \sum_{j=1}^{|S_1|-2} \log_2 \left(\frac{n}{|S_1| - j + 1} \right) \\
&= |S_1| + 2 \sum_{j=3}^{|S_1|} \log_2 \left(\frac{n}{j} \right) \\
&\leq |S_1| + 2|S_1| \log_2 n - 2 \sum_{j=3}^{|S_1|} \log_2 j \\
&= |S_1| + 2|S_1| \log_2 n - 2 \log_2 ((|S_1|)!/2) \\
&\leq |S_1| + 2|S_1| \log_2 n - 2|S_1| \log_2 \left(\frac{|S_1|}{e} \right) \\
&\leq |S_1| (1 + 2 \log_2 \left(\frac{ne}{|S_1|} \right)) \leq 3|S_1| \log_2 \left(\frac{n}{|S_1|} \right) \leq 3\epsilon n/100. \quad (3)
\end{aligned}$$

Hence, in total, the forest at any point contains less than s vertices, so we can successfully apply Theorem 4.3.

As $i < |S_1| - 1$, by the induction hypothesis we have at least $|S_1| - i \geq 2$ trees, so we can group them into two collections each containing at least a third of all vertices. Since the total number of vertices in all trees is $s/3$, each of the two collections contains at least $\frac{s/3}{2} > m$ vertices. Thus, there must be an edge between them. This edge closes a path P_{i+1} connecting two endpoints in the linear forest $I[S_1] \cup P_1 \cup \dots \cup P_i$. Adding this edge via Lemma 4.4 and removing the unused edges (leaf-by-leaf) using Lemma 4.5, we are left with an $(m, 10)$ -extendable subgraph $I[S_1] \cup P_1 \cup \dots \cup P_{i+1}$ of G' . Let us verify that all invariants are still satisfied.

By construction, all internal vertices of the new path are in R . The length of the new path is at most twice the depth of the trees plus one, i.e. it is of length at most $2 \log_2 \left(\frac{n}{|S_1| - i} \right) + 1$. By construction, the subgraph $I[S_1] \cup P_1 \cup \dots \cup P_i \cup P_{i+1}$ is still $(m, 10)$ -extendable. Furthermore, as P_{i+1} connects two endpoints of a linear forest through a path which is internally vertex-disjoint from $I[S_1] \cup P_1 \cup \dots \cup P_i$, we obtain a new linear forest with one less component than before, resulting in $|S_1| - i - 1$ components/paths, as required.

Continuing from the final $(m, 10)$ -extendable embedding (which now connects all vertices in S_1), we can now follow the same procedure for the sets S_2, \dots, S_t , one by one, always extending the current $(m, 10)$ -extendable subgraph. Here we stress that when we are about to do the procedure for the set S_i , we continue extending the $(m, 10)$ -extendable subgraph left after having completed treating S_{i-1} , repeating the exact same argument as before, only bearing in mind that we have already used up some vertices in R for the previous steps of the embedding.

The only thing we need to verify is that the total number of vertices in the forest which we construct is always at most s . But this is straightforward to see, as the total number of vertices in the current binary trees which we use in every step is at most $s/2$ by construction. On the other hand, the total number of vertices in paths used for all S_i , by (3), is at most

$$3 \sum_{i=1}^t |S_i| \log_2 \left(\frac{n}{|S_i|} \right) \leq 3\epsilon n/100 < s/2,$$

so we are done. □

4.3 ISTs in pseudorandom graphs

We now prove our second main result, restated here for the reader's convenience.

Theorem 1.4. *Fix $\epsilon > 0$, and let G be an (n, d, λ) -graph with $d/\lambda = \omega(\log d)$ and $d = o(n/\log^2 n)$. Then*

G contains $(1 - \varepsilon)d$ vertex-independent spanning trees for each choice of root.

Proof. We will prove the claim with 10ε instead of ε ; we may also assume that $\varepsilon < 10^{-3}$ since the lemma only gets stronger by decreasing ε . Define $k = (1 - 10\varepsilon)d$. Fix an arbitrary root r , and k of its neighbours $L := \{v_1, \dots, v_k\}$. The proof follows from the next two claims.

Claim 4.1. *There exists a partition of $V(G) \setminus (\{r\} \cup L) := U \cup R \cup S_1 \cup \dots \cup S_k$ such that the following holds for every $v \in V(G)$:*

- Let $I(v)$ be the set of $i \in [k]$ such that v is not a neighbour of v_i . Then for each $i \in I(v)$ there are edges $u_i y_i$, where $u_i \in N_U(v)$ are distinct vertices, and $y_i \in S_i$;
- $|N_R(v)| \geq \varepsilon^3 d$;
- $|S_i| \leq \varepsilon^{10} n / (d \log d)$.

Claim 4.2. *Given a partition as in Claim 4.1, there exist disjoint paths P_1, \dots, P_k such that*

$$S_i \cup \{v_i\} \subseteq V(P_i) \subseteq S_i \cup R \cup \{v_i\}.$$

Before we prove the two claims, we observe that they suffice to prove the theorem. Define $T_i := P_i + (r, v_i)$. Since in Claim 4.1 we have $u_i \in U$, and $U \subset V(G) \setminus \bigcup_{i \in [k]} T_i$, using the first bullet point of Claim 4.1 we have that $\{T_i\}_{i \in [k]}$ is a nice collection of trees, so we are done by Lemma 3.2.

Proof of Claim 4.1. Consider an arbitrary vertex $v \neq r$. Let $t_v := k - |N_L(v)|$, and denote $B_v := N(v) \setminus N_L(v) \cup \{r\}$, noting that $|B_v| = t_v + 10\varepsilon d$. At most εd vertices in B_v have at least εd neighbours in $B_v \cup L$ as by the Expander Mixing Lemma we have

$$e(B_v, B_v \cup L \cup \{r\}) \leq d \cdot 2d \cdot d/n + \lambda \sqrt{d \cdot 2d} \leq \varepsilon^2 d^2$$

where we used $d = o(n)$ and $\lambda = o(d)$. Thus, we can take a subset $B'_v \subseteq B_v$ with $|B'_v| = t_v + \varepsilon d$, of vertices with at least εd neighbours outside $W := B_v \cup L \cup \{r\}$. Thus every subset $X \subseteq B'_v$ has $|N_{G-W}(X)| \geq C \log^2 d |X|$ by Lemma 4.8 by setting $k_{4.7} = C \log^2 d$, $Z_{4.7} = V(G) \setminus W$, for $C > C_0(\varepsilon)$ large enough so that the rest of the argument goes through. Hence, by the generalised Hall's theorem we can find a collection of $t_v + \varepsilon d$ vertex disjoint stars Q_i of size $C \log^2 d$ each, where $i \in [t_v + \varepsilon d]$, with centers in B'_v . The need for size $C \log^2 d$ is the bottleneck that dictates the assumed bound on d/λ .

Now we construct the required partition. For each vertex v in $V(G) - L - \{r\}$, randomly assign it to one of the following sets, with given probability:

- with probability $\varepsilon^{11} / (d \log d)$ put v into S_i , for each $i \in [k]$;
- with probability $\varepsilon/100$ put v into the *reservoir* R ;
- with the remaining probability, put v into U , the set of unassigned vertices.

For each $v \in V(G)$, denote by A_v the intersection of the following events:

- (1) There are at least $t_v + \varepsilon d/100$ vertices in $B'_v \cap U$. By possibly relabelling, let $Q_1, \dots, Q_{t_v + \varepsilon d/100}$ be their previously chosen stars.
- (2) There is an injection $\sigma : I(v) \rightarrow [t_v + \varepsilon d/100]$, such that $Q_{\sigma(i)}$ contains a leaf in S_i , for all $i \in I(v)$.
- (3) $|N_R(v)| \geq \varepsilon^3 d$.

Note that if $\bigcap_v A_v$ holds, then the first two bullets of Claim 4.1 hold (with u_i the center of $Q_{\sigma(i)}$ and y_i the leaf which guaranteed in (2)). Let us estimate the probability $P[A_v]$. First, note that the first property is satisfied with probability at least $1 - e^{-\Theta(d)}$, as $\mathbb{E}[|B'_v \cap U|] \geq |B'_v|(1 - \varepsilon/3) \geq (1 - \varepsilon/3)t_v + \varepsilon d/2 \geq t_v + \varepsilon d/6$, so a Chernoff bound applies.

Conditioning on the first property, for the second one consider the auxiliary bipartite graph H_v with first part $I(v)$ of size t_v , and second part $[t_v + \varepsilon d/100]$, where (i, j) is an edge if there is a vertex $w \in S_i \cap Q_j$. To show that H_v has a matching covering $I(v)$ with sufficiently high probability, set $A := I(v)$ and $B := [t_v + \varepsilon d/100]$. We claim the following.

Proposition 4.10. *With probability at least $1 - d^{-6}$ there is a matching between A and B that saturates A .*

Proof. If $H_v[A, B]$ does not have a required matching, then by Hall's theorem there is a set $A_0 \subset A$ of cardinality $|A_0| \leq t_v$ with $|N(A_0, B)| < |A_0|$. For A_0 to have no neighbours in $B_0 \subseteq B$, it is necessary that, for every $i \in B_0$, all leaves of the stars Q_i lie outside the sets S_j with $j \in A_0$. The probability of this event, for fixed sets A_0 and B_0 , where $|B_0| = |B| - |A_0|$, is

$$\left(1 - \frac{|A_0|\varepsilon^{11}}{d \log d}\right)^{|B_0|C \log^2 d} \leq e^{-\frac{|A_0|\varepsilon^{11}}{d \log d}(t_v + \varepsilon d/100 - |A_0|)C \log^2 d} \leq d^{-\frac{\varepsilon^{12}C|A_0|}{100}}$$

where we used that $|A_0| \leq t_v$. Thus, by a union bound, the probability that Hall's condition for side A is violated in H_v is at most

$$\begin{aligned} \sum_{\ell=1}^{t_v} \binom{t_v}{\ell} \binom{t_v + \varepsilon d}{\ell} d^{-\frac{\varepsilon^{12}C\ell}{100}} &\leq \sum_{\ell=1}^{t_v} \left(\frac{e(t_v + \varepsilon d)}{\ell}\right)^{2\ell} d^{-\frac{\varepsilon^{12}C\ell}{100}} \\ &\leq \sum_{\ell=1}^{t_v} (3d)^{2\ell} d^{-100\ell} \leq d^{-6}, \end{aligned}$$

where we used that $d > t_v$. This completes the proof of the Proposition 4.10. \square

Finally, the third property also holds with probability at least $1 - e^{-\Theta(d)}$, as the expected number of neighbours in R is at least $\varepsilon d \cdot \varepsilon/100$. Hence we conclude that $P[A_v] \geq 1 - 3d^{-6}$. Note that each event A_v depends on at most d^4 other events, since this bounds the number of vertices within distance 4 of v ; for vertices w at distance at least 5, the events A_w are independent of A_v .

Let E_i be the event given by the third bullet of Claim 4.1, i.e. the event that $|S_i| \leq \varepsilon^{10}n/(d \log d)$. By Chernoff bounds, we have $P[E_i] \geq 1 - e^{-\Theta(n/(d \log d))} \geq 1 - n^{-3}$, as by assumption $d = o(n/\log^2 n)$. Note that each E_i might depend on each event A_v (the number of which is n) and each other E_j (of which we have at most d).

Thus, by the Asymmetric Local Lemma (Lemma 4.6), there is a choice of sets R, S_1, \dots, S_k such that each A_v and each E_i hold. Indeed, by setting $x_v = d^{-5}$ for all $v \neq r$, and $x_i = n^{-2}$ for all $i \in [k]$, the condition of the lemma is satisfied as:

$$P[\overline{A_v}] \leq 3d^{-6} \leq d^{-5}(1 - d^{-5})^{d^4} \cdot (1 - n^{-2})^d \sim d^{-5},$$

and:

$$P[\overline{E_i}] = e^{-\Theta(n/(d \log d))} \leq n^{-2}(1 - d^{-5})^n(1 - n^{-2})^d \sim e^{-\Theta(n/d^5)}.$$

This completes the proof. \square

Now we turn to the second claim.

Proof of Claim 4.2. Denote $S'_i := S_i + v_i$. We apply Lemma 4.9, to the sets R, S'_1, \dots, S'_k . Indeed, we can do that since each term in $\sum |S'_i| \log(n/|S'_i|)$ is maximized when $|S'_i|$ is as large as possible; by Claim 4.1 we can assume that $|S_i| \leq \varepsilon^{10}n/(d \log d)$, thus we have

$$\sum_{i=1}^k |S'_i| \log(n/|S'_i|) \leq d \cdot \frac{\varepsilon^{10}n}{d \log d} 2 \log d = 2\varepsilon^{10}n.$$

This connects the set $S'_i = S_i + v_i$ via path P_i through R , such that the P_i 's are pairwise vertex-disjoint. This completes the proof of the second claim. \square

As we have proven the two required claims, this completes the proof of Theorem 1.4 \square

5 Concluding Remarks

We have proven the Zehavi-Itai IST conjecture for random graphs $G(n, p)$ for essentially the whole range of relevant values of the edge probability $p(n)$, specifically for $C \log n/n \leq p(n) \leq 0.99$, for a large enough absolute constant $C > 0$. We have also shown that the conjecture holds asymptotically (namely, one can find $(1 - o(1))\kappa(G)$ ISTs) for a broad class of weakly pseudorandom graphs (specifically, (n, d, λ) -graphs, under a rather mild assumption on the eigenvalue ratio $d/\lambda = \omega(\log d)$). As a consequence, the conjecture also holds asymptotically for random d -regular graphs when d is at least a large constant. In other words, d -regular graphs typically contain a family of $(1 - o(1))d$ ISTs.

Several problems remain open. For general k -connected graphs, the best known lower bound [4] guarantees only $\Omega(k/\log^2 n)$ independent spanning trees, and improving this remains an exciting challenge. Our techniques offer some insight in this direction, and it would be worthwhile to investigate whether one can exploit a dichotomy between good expansion and small cuts to obtain stronger lower bounds.

The conjecture is known to hold for small values of k , specifically for $k \leq 4$. It would be interesting to explore the other end of the spectrum, namely graphs with linear connectivity, and to determine whether one can prove the conjecture in this setting or at least to improve significantly the lower bound from [4].

Another immediate challenge is to strengthen Theorem 1.4 by replacing the hypothesis $d/\lambda = \omega(\log d)$ with $d/\lambda \geq C(\varepsilon)$ for a sufficiently large constant $C(\varepsilon)$ depending only on ε . We can notice that this strengthening holds under the additional assumption that G is $K_{s,t}$ -free for fixed integers $s, t \geq 2$: a constant spectral gap already yields polynomial expansion in this case. In particular, in our proof one can then find stars of size d^c for some $c > 0$, whereas the proof requires only stars of order $C \log^2 d$, which we obtained using the Expander Mixing Lemma. Indeed, combining the Expander Mixing Lemma with a Kővári–Sós–Turán-type argument (see, e.g., [22]) gives the needed polynomial expansion and hence the required stars.

Finally, it would be interesting to address the remaining range of p for binomial random graphs as well, and in particular at the weak threshold for connectivity, when $\log n/n \leq p = O(\log n/n)$.

Although our result is stated existentially, our proof yields a polynomial-time algorithm, as can be verified by examining each step of the argument.

Acknowledgement. Part of this research was conducted at Princeton University during the workshop on Graph expansion and its applications in April-May 2025. We would like to thank Matija Bucić for organizing the workshop, and Princeton University for great hospitality.

References

- [1] N. Alon and J. H. Spencer. *The Probabilistic Method*. John Wiley & Sons, 2008.
- [2] B. Bollobás. *Random graphs*, volume 73 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, second edition, 2001.
- [3] B. Bollobás and A. Thomason. Random graphs of small order. In *North-Holland Mathematics Studies*, volume 118, pages 47–97. Elsevier, 1985.
- [4] K. Censor-Hillel, M. Ghaffari, G. Giakkoupis, B. Haeupler, and F. Kuhn. Tight bounds on vertex connectivity under sampling. *ACM Transactions on Algorithms (TALG)*, 13(2):1–26, 2017.
- [5] K. Censor-Hillel, M. Ghaffari, and F. Kuhn. A new perspective on vertex connectivity. *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms, SIAM*, pages 546–561, 2014.

- [6] B. Cheng, D. Wang, and J. Fan. Independent spanning trees in networks: a survey. *ACM Computing Surveys*, 55(14s):1–29, 2023.
- [7] J. Cheriyan and S. N. Maheshwari. Finding nonseparating induced cycles and independent spanning trees in 3-connected graphs. *Journal of Algorithms*, 9(4):507–537, 1988.
- [8] S. Curran, O. Lee, and X. Yu. Finding four independent trees. *SIAM Journal on Computing*, 35(5):1023–1058, 2006.
- [9] N. Draganić and M. Krivelevich. Disjoint connected dominating sets in pseudorandom graphs. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 329–335, 2025.
- [10] N. Draganić, M. Krivelevich, and R. Nenadov. Rolling backwards can move you forward: on embedding problems in sparse expanders. *Transactions of the American Mathematical Society*, 375(7):5195–5216, 2022.
- [11] N. Draganić, R. Montgomery, D. M. Correia, A. Pokrovskiy, and B. Sudakov. Hamiltonicity of expanders: optimal bounds and applications. *arXiv preprint arXiv:2402.06603*, 2024.
- [12] J. Friedman and N. Pippenger. Expanding graphs contain all small trees. *Combinatorica*, 7:71–76, 1987.
- [13] P. Gao, M. Isaev, and B. D. McKay. Sandwiching random regular graphs between binomial random graphs. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 690–701. SIAM, 2020.
- [14] A. Gopalan and S. Ramasubramanian. IP fast rerouting and disjoint multipath routing with three edge-independent spanning trees. *IEEE/ACM Transactions on Networking*, 24(3):1336–1349, 2016.
- [15] P. E. Haxell. Tree embeddings. *Journal of Graph Theory*, 36(3):121–130, 2001.
- [16] L. Hollom, L. Lichev, A. Mond, J. Portier, and Y. Wang. Approximate Itai-Zehavi conjecture for random graphs. *arXiv preprint arXiv:2506.23970*, 2025.
- [17] S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.
- [18] A. Itai and M. Rodeh. The multi-tree approach to reliability in distributed networks. *Information and Computation*, 79(1):43–59, 1988.
- [19] S. Janson, T. Łuczak, and A. Ruciński. *Random graphs*. John Wiley & Sons, 2000.
- [20] M. Krivelevich and B. Sudakov. Pseudo-random graphs. In *More sets, graphs and numbers: A Salute to Vera Sos and András Hajnal*, pages 199–262. Springer, 2006.
- [21] X.-Y. Li, W. Lin, X. Liu, C.-K. Lin, K.-J. Pai, and J.-M. Chang. Completely independent spanning trees on BCCC data center networks with an application to fault-tolerant routing. *IEEE Transactions on Parallel and Distributed Systems*, 33(8):1939–1952, 2021.
- [22] H. Liu and R. Montgomery. A proof of Mader’s conjecture on large clique subdivisions in C_4 -free graphs. *Journal of the London Mathematical Society*, 95(1):203–222, 2017.
- [23] K. Menger. Zur allgemeinen Kurventheorie. *Fund. Math.*, pages 96–115, 1927.
- [24] K. Miura, D. Takahashi, S.-I. Nakano, and T. Nishizeki. A linear-time algorithm to find four independent spanning trees in four-connected planar graphs. *International Journal of Foundations of Computer Science*, 10(2):195–210, 1999.
- [25] R. Montgomery. Spanning trees in random graphs. *Advances in Mathematics*, 356:106793, 2019.
- [26] R. A. Moser and G. Tardos. A constructive proof of the general Lovász local lemma. *Journal of the ACM (JACM)*, 57(2):1–15, 2010.
- [27] A. Sarid. The spectral gap of random regular graphs. *Random Structures & Algorithms*, 63(2):557–587, 2023.
- [28] E. S. d. Silva and H. Pedrini. Inferring patterns in mitochondrial dna sequences through hypercube independent spanning trees. *Computers in Biology and Medicine*, 70:51–57, 2016.
- [29] S.-M. Tang, J.-S. Yang, Y.-L. Wang, and J.-M. Chang. Independent spanning trees on multidimensional torus networks. *IEEE Transactions on Computers*, 59(1):93–102, 2010.
- [30] J. Werapun, S. Intakosum, and V. Boonjing. An efficient parallel construction of optimal independent spanning trees on hypercubes. *Journal of Parallel and Distributed Computing*, 72(12):1713–1724, 2012.

- [31] J.-S. Yang, S.-M. Tang, J.-M. Chang, and Y.-L. Wang. Parallel construction of optimal independent spanning trees on hypercubes. *Parallel Computing*, 33(2):73–79, 2007.
- [32] A. Zehavi and A. Itai. Three tree-paths. *Journal of Graph Theory*, 13(2):175–188, 1989.