On Independent Spanning Trees in Random Graphs*

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

Abstract. A central challenge in network design is ensuring resilience: how can we guarantee multiple, independent, communication pathways between nodes, even when some connections fail in a network? In 1989, Zehavi and Itai formulated a graph-theoretic conjecture that captures the essence of this problem. They proposed that any k-vertex-connected graph contains k *independent spanning trees* rooted at any given root r, which means that for every vertex v in the graph, the unique r–v paths within these k spanning trees are entirely disjoint, apart from their endpoints r and v. Despite decades of effort, this conjecture has only been proven for $k \leq 4$ and for specific graph families using their underlying topological structure, leaving the general case as an open problem in graph theory with substantial consequences in the field of distributed algorithms.

We make significant progress on the Zehavi-Itai conjecture by proving it holds for almost all graphs of relevant densities. More precisely, we show that there exists some constant $C > 1$ such that for all $C \log n/n \leq p < 0.99$, the binomial random graph $G(n,p)$ contains a family of $\delta(G)$ independent spanning trees rooted at any given vertex r with high probability. Note that the lower bound on p up to the constant C matches the standard threshold for connectivity for $G(n,p)$, thus we establish an essentially best possible result for random graphs.

1 Introduction From cloud data centres that must stay online during switch failures, to wireless sensor networks where links flicker with interference, modern communication substrates are expected to *maintain connectivity under failure.* Practical fault models range from benign (single-link outages) to adversarial (targeted node removals or eavesdropping), yet the design principle is uniform: embed *redundant and efficiently exploitable structure* so that traffic can be rerouted quickly, the network is resilient to the failure of a fixed number of paths, and critical services remain uninterrupted. Among the most extensively studied combinatorial tools in this context are *independent spanning trees* (ISTs) — collections of spanning trees that intersect as little as possible along critical paths. Their ability to simultaneously offer path diversity, fast recovery guarantees, and certificates of connectivity has made them a natural object of study in algorithmic network design.

Formally, a collection $T_1, \ldots, T_k$ of $k \geq 2$ spanning trees in a graph G, rooted at a vertex r, are said to be *ISTs*, short for *independent spanning trees*[1], 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. Independent spanning trees surface in an impressively wide range of settings: in *computational biology*, hypercube ISTs help to detect anomalies in mitochondrial DNA [14]; in *high-performance computing*, ISTs provide fault-tolerant broadcast and low-diameter routes for multidimensional torus interconnections [15]; carrier-grade IP networks leverage three edge-ISTs for fast rerouting under dual-link failures [8]; and emerging server-centric datacentre fabrics employ so-called completely independent trees to preserve high throughput despite switch outages [12]. Such diversity, ranging from biology to

*The full version of the paper can be accessed at https://arxiv.org/abs/2509.26401

†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.

[1]These are also commonly known as vertex- or node- independent spanning trees; edge-independent trees are also widely studied in the literature and are related to *edge-connectivity* instead. The interested reader is referred to an excellent survey [4] for further discussion on the edge version.

large-scale computing, underscores the importance of ISTs as a unifying tool for resilient multipath connectivity. We refer the interested reader to an extensive survey [4] by Cheng, Wang and Fan on the subject.

Probably the most famous conjecture in the area was made by Zehavi and Itai in the late '80s [18], 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 $\kappa(G)$.

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. 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 [10] showed that one can construct a collection of ISTs of size two in a 2-connected graph. Further early evidence towards Conjecture 1.1 arrived when Cheriyan and Maheshwari [5] showed that every 3-connected graph contains a collection of three ISTs, providing an explicit $O(n^2)$ algorithm to find such a collection. Zehavi and Itai [18] independently also arrived at the same result, and stated Conjecture 1.1 in that same work. Much later, Curran, Lee, and Yu [6] succeeded in building four ISTs in any 4-connected graph, settling the $k=4$ case. No other general results are known, except for graphs with certain underlying topology such as planar graphs or cube-like graphs. Miura, Takahashi, Nakano, and Nishizeki [13] obtained a linear-time algorithm to find a collection of four ISTs in 4-connected planar graphs. By a collection of results [17, 16] it is known that the n-dimensional hypercube Q_n contains n ISTs, optimally matching the cube's connectivity. Many more results can be found in [4].

Regarding general bounds, Censor-Hillel, Ghaffari, Giakkoupis, Haeupler and Kuhn [3] showed that k-connected n-vertex graphs contain $\Omega(k/\log^2 n)$ many ISTs, via connected dominating sets (CDSs). A *connected dominating set* S in a graph G is a subset of vertices inducing a connected subgraph such that every vertex outside of S has a neighbour in S. Draganić and Krivelevich [7] also employed CDSs to show that d-regular pseudorandom graphs contain about $d/\log d$ many ISTs, removing the dependence on n from the general case. They conjectured that random graphs above the connectivity threshold satisfy the IST conjecture, or at least an approximate version of it.

Our central result is the following theorem essentially showing that the Zehavi-Itai conjecture typically holds for random graphs above the connection threshold:

THEOREM 1.2. *There exists $C>1$ such that, for any $C\log n/n \le p \le 0.99$, 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 Thomasson [2] shows that typically $\kappa(G)=\delta(G)$ for any $p=p(n)$, this implies that the Zehavi-Itai conjecture holds for almost all binomial random graphs in the relevant range.

Remark. Very recently, Hollom, Lichev, Mond, Portier and Wang [9] posted an asymptotic solution of the Zehavi-Itai conjecture for random graphs $G(n,p)$ with $p \gg \log n/n$, and for random regular graphs $G(n,d)$ with $d \gg \log n$. Note that our result provides an *exact* solution of the Zehavi-Itai conjecture for random graphs $G(n,p)$ with $p=\Omega(\log n/n)$, i.e., for an essentially optimal range of $p(n)$. They also showed the existence of $d/4$ ISTs in random d-regular graphs for large enough d, which we improve to $(1-o(1))d$ in the full version of our paper.

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 $S \subseteq V(G)$, $G[S]$ denotes the subgraph of G induced by S and $N_G(S)$ denotes the *external* neighbourhood of S. Whenever G 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*

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

probability (abbreviated **whp**) if $\Pr(\mathcal{E}) \to 1$ as $n \to \infty$. 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.*

- $P[X > (1+\delta)np] \le e^{-\delta^2 np/(2+\delta)}$ *for* $\delta > 0$.
- $P[X < (1-\delta)np] \le e^{-\delta^2 np/2}$ *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 $C \log n/n \le p \le 0.99$, 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 $\varepsilon \to 0$ as $C \to \infty$; furthermore $\Delta(G) \le 2pn$ with probability at least $1 - o(n^{-10})$.*

(b) ***Low-degree set:*** *For $p = o(1)$, $|S| \le n^{0.2}$ with probability $1 - o(n^{-3})$.*

(c) ***Sparse boundary:*** *For $p \le n^{-0.49}$, 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 $p \le n^{-0.49}$, 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 $p \le n^{-0.49}$ the set S is independent* **whp**.

(f) ***Connectivity:*** *G is $\delta(G)$-connected* **whp**.

Proof.

(a) See Chapter 3 in [1].

(b) Consider a fixed set S of size $\lfloor n^{0.2} \rfloor$. 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

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

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 $|N(S')| \le 2n^{1.21}p$ 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 $|S' \cup N(S')|^{100}p^{100} \le n^{-5}$. 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 Item (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 $n \to \infty$. ∎

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. ∎

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. ∎

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 $\mathscr{S} = \{S_i\}_{i\in[t]}$ 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 $\mathscr{S}$ 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 $\mathscr{T} = \{S_i\}_{i\in[t]}$ 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_iw_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: $\log^2 n/\sqrt{n} \le p \le 0.99$. 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 3.3. *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 $|N(v)| \ge \delta = |K|$, we have $|K \setminus N(v)| \le |N(v) \setminus K|$, 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 3.3. 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 $d(u) \ge np/2$. Let K be the set of u's first $\min\{k, d(u)\} \ge \sqrt{n}\log n$ neighbours. We now expose the neighbours of v; by the upper bound on p, we know that $|N(v) \setminus K| \ge d(v)/10^4 \ge \sqrt{n}\log n$ with probability at least $1 - o(n^{-8})$. Similarly $|K \setminus N(v)| \ge |K|/10^4 \ge \sqrt{n}\log n$. 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 $1-o(n_0^{-10}) > 1-o(n^{-5})$. ∎

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 $\mathscr{P}$ be the set of all paths, and H_v be the auxiliary bipartite graph between $\mathscr{P}$ 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 $|\mathscr{P}| = \delta \leq |B_v|$, 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 $\mathscr{P}$ 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

$$np_1 - 0.9\sqrt{2np_1\log n} + n(p_2+p_3) - \sqrt{4.1n(p_2+p_3)\log n} \geq np - 0.95\sqrt{2np\log n}, \tag{3.1}$$

where we used $0.9\cdot\sqrt{0.999} + \sqrt{2.05\cdot 0.001} < 0.95$.

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 (3.1) 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 $\mathscr{P}$ of vertex-disjoint paths $P_1, \dots, P_\delta$, 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 $V(\mathscr{P}) = P \cup Q$. Note that

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

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 $\mathscr{P}$ 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 $\mathrm{Hyp}(|W| - |Q|, K, m)$ 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 [11]) 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

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

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

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

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

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

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.

4 Concluding Remarks We proved the Zehavi-Itai IST conjecture for random graphs $G(n, p)$ with $p \geq C \log n/n$. In the full version of our paper, we also show that the conjecture holds asymptotically for a broad class of weakly pseudorandom graphs (specifically, (n, d, λ)-graphs, which we do not define here, under a rather mild assumption on the eigenvalue ratio d/λ). As a consequence, the conjecture also holds asymptotically for random d-regular graphs when d is a large constant. In other words, d-regular graphs typically contain $(1 - o(1))d$ ISTs.

Several problems remain open. For general k-connected graphs, the best known lower bound [3] 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, in particular 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 to improve significantly the lower bound from [3].

Although our result is stated existentially, our proof yields a polynomial-time algorithm 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. We would also like to thank the referees for their helpful remarks.

References

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

[2] B. Bollobás and A. Thomason. Random graphs of small order. In *North-Holland Mathematics Studies*, volume 118, pages 47–97. Elsevier, 1985.

[3] 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.

[4] B. Cheng, D. Wang, and J. Fan. Independent spanning trees in networks: a survey. *ACM Computing Surveys*, 55(14s):1–29, 2023.

[5] 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.

[6] S. Curran, O. Lee, and X. Yu. Finding four independent trees. *SIAM Journal on Computing*, 35(5):1023–1058, 2006.

[7] 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.

[8] 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.

[9] L. Hollom, L. Lichev, A. Mond, J. Portier, and Y. Wang. Approximate Itai-Zehavi conjecture for random graphs. *arXiv preprint arXiv:2506.23970*, 2025.

[10] A. Itai and M. Rodeh. The multi-tree approach to reliability in distributed networks. *Information and Computation*, 79(1):43–59, 1988.

[11] S. Janson, T. Łuczak, and A. Ruciński. *Random graphs*. John Wiley & Sons, 2000.

[12] 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.

[13] 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.

[14] 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.

[15] 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.

[16] 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.

[17] 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.

[18] A. Zehavi and A. Itai. Three tree-paths. *Journal of Graph Theory*, 13(2):175–188, 1989.