

Finding Planted Cycles in a Random Graph

Julia Gaudio* Colin Sandon[†] Jiaming Xu[‡] Dana Yang[§]

November 7, 2025

Abstract

In this paper, we study the problem of finding a collection of planted cycles in an Erdős–Rényi random graph $G \sim \mathcal{G}(n, \lambda/n)$, in analogy to the famous Planted Clique Problem. When the cycles are planted on a uniformly random subset of δn vertices, we show that almost-exact recovery (that is, recovering all but a vanishing fraction of planted-cycle edges as $n \rightarrow \infty$) is information-theoretically possible if $\lambda < \frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$ and impossible if $\lambda > \frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$. Moreover, despite the worst-case computational hardness of finding long cycles, we design a polynomial-time algorithm that attains almost exact recovery when $\lambda < \frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$. This stands in stark contrast to the Planted Clique Problem, where a significant computational-statistical gap is widely conjectured.

Keywords— Random graphs, Planted cycles, Phase transitions, Generating functions, Branching processes

*Northwestern University, Department of Industrial Engineering and Management Sciences; julia.gaudio@northwestern.edu

[†]École Polytechnique Fédérale de Lausanne; colin.sandon@epfl.ch

[‡]Duke University, The Fuqua School of Business; jiaming.xu868@duke.edu

[§]Cornell University, Department of Statistics and Data Science; xy374@cornell.edu

1 Introduction

The problem of finding the longest cycle in a graph $G = ([n], E)$ is a classical and fundamental challenge in theoretical computer science. It is well-known to be NP-hard in the worst case. A celebrated color-coding technique can find a cycle of length k , if it exists, in $e^{O(k)} n |E|$ expected time [2]. However, this method only gives a polynomial-time algorithm to find a cycle of length $O(\log n)$.

Given the problem's computational intractability in worst-case instances, a natural direction is to study it in random graph models, where one can hope for more tractable behavior. This line of inquiry dates back to the foundational work of Erdős and Rényi [11], and has since grown into a rich area of research. A celebrated result of [7] shows that in the Erdős–Rényi random graph $\mathcal{G}(n, \lambda/n)$, the length of the longest cycle satisfies $L_{\lambda, n} \geq (1 - \lambda^6 e^{-\lambda})n$ for $\lambda > 0$ sufficiently large. Subsequent works [12] have further refined these bounds, culminating in the result of [3, 4] proving that $L_{\lambda, n}/n \rightarrow f(\lambda)$ almost surely for $\lambda \geq 20$, where $f(\lambda) = 1 - \sum_{k=1}^{\infty} p_k(\lambda) e^{-k\lambda}$ and p_k is a polynomial in λ . From a computational standpoint, the simple depth-first-search algorithm can find a cycle of length $\Theta(n)$ efficiently [17]. However, whether the longest cycle can be found in polynomial time remains an open question.

A similar phenomenon is observed for the problem of finding the largest clique in a graph. It is known that the largest clique in $\mathcal{G}(n, p)$ for constant p is of size $(2 \pm o(1)) \log_{1/p}(n)$ with high probability [6], yet the best known polynomial-time algorithms can only achieve half of this size [15]. In response to this gap, a prominent line of research initiated by [14, 18] considers the problem of finding a clique which is *planted* inside an Erdős–Rényi random graph $G \sim \mathcal{G}(n, p)$. Letting k be the size of the planted clique, it is known that one can recover a planted clique of size $k = \Theta(\log n)$ information-theoretically; however, the best known polynomial-time algorithms can only recover a planted clique of size $k = \Omega(\sqrt{n})$ (see e.g. [1]). We study the analogous question for cycles:

Under what conditions can we (efficiently) find cycles that are planted in a random graph?

Concretely, we generate a graph according to the following planted cycles model.¹

Definition 1.1 (Planted cycles model). To generate a graph G from the planted cycles model, denoted by $\mathcal{G}(n, \lambda, \delta)$, first generate a background Erdős–Rényi graph $G_0 = ([n], E) \sim \mathcal{G}(n, \lambda/n)$ by connecting each pair of two vertices independently at random with probability λ/n . Then choose a set of δn^2 vertices $V \subseteq [n]$, uniformly at random. Finally, select a labeled 2-factor H^* on V , a vertex-disjoint union of cycles spanning all vertices in V , uniformly at random. Let $G = G_0 \cup H^*$ be the union of the background graph and the planted cycles.

The goal is to recover the edges in the planted 2-factor H^* based on the observation G . Note that crucially, the vertex set V of the planted 2-factor H^* is unknown. We focus on the regime where $\lambda > 0$ is a constant. It turns out that exactly recovering H^* with high probability is information-theoretically impossible for constant λ (Theorem B.1). We therefore study the problem of almost exact recovery; that is, recovering all but an $o(1)$ fraction of edges in the planted cycles in expectation. More formally, let $\hat{H}$ be an estimator of H^* ; that is, a set of edges on the complete graph K_n . The reconstruction error, namely the fraction of misclassified edges, is

$$\ell(H^*, \hat{H}) = \frac{|H^* \Delta \hat{H}|}{|H^*|}, \quad (1.1)$$

where Δ denotes the symmetric set difference. We say that $\hat{H}$ achieves almost exact recovery if $\mathbb{E}[\ell(H^*, \hat{H})] = o(1)$, or equivalently, $\mathbb{P}\left\{\ell(H^*, \hat{H}) = o(1)\right\} = 1 - o(1)$, as $n \rightarrow \infty$.

1.1 Main results

Our main results precisely characterize the threshold for almost exact recovery of the planted cycles.

Theorem 1.1. *Consider the planted cycles model $\mathcal{G}(n, \lambda, \delta)$, where $\delta \in (0, 1]$. If $\lambda < \frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$, then almost exact recovery is possible. Moreover, there exists a polynomial-time algorithm that achieves almost exact recovery. Conversely, if $\lambda > \frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$, then almost exact recovery is information-theoretically impossible for any algorithm, regardless of computation time.*

¹We note that our planted cycles problem is distinct from the similarly-named *planted dense cycle problem* [19, 20]. In the planted dense cycle problem, the vertices of the background graph are spatially embedded on a circle. Additional edges are then added, connecting pairs of vertices randomly with a probability that decays with the circular distance.

²Strictly speaking we choose $\lceil \delta n \rceil$ vertices, but for ease of notation we generally just say δn .

The case $\delta = 1$ corresponds to a planted 2-factor on the entire vertex set $[n]$, and was previously studied in [13]. In that setting, almost exact recovery is possible if $\lambda \leq \frac{1}{2}$, while it is impossible for $\lambda > \frac{1}{2}$. Interestingly, the threshold $\frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$ is non-monotonic in δ , attaining its maximum value 1 at $\delta = 0$ and minimum value $1/3$ at $\delta = 2/3$. Thus, when $\lambda < \frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$, we are always in the subcritical regime of the Erdős–Rényi random graph, which contains only short cycles of bounded length.

A salient challenge in our problem is that the vertex set V of the planted cycles H^* is unknown. If V were known, the problem would reduce to recovering a planted 2-factor within the induced subgraph $G[V]$, which is distributed as $\mathcal{G}(n', \lambda'/n')$ with $n' = \delta n$ and $\lambda' = \delta \lambda$. Thus, in this known- V setting, the sharp threshold for almost exact recovery is $\lambda' = 1/2$, or equivalently $\lambda = 1/(2\delta)$. Since this threshold is always higher than $\frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$ for $\delta < 1$, this comparison highlights that the unknown support V renders the recovery task statistically strictly harder.

From a computational perspective, the unknown vertex set V also introduces substantial challenges, as enumerating over all $\binom{n}{\delta n}$ possible subsets requires exponential time. Nevertheless, and perhaps surprisingly, we establish that the sharp recovery threshold can still be attained by a polynomial-time algorithm; hence, there is no statistical-computational gap. This stands in stark contrast to the planted clique problem. Information-theoretically, it is possible to recover a planted clique of size $\Theta(\log n)$. However, it is conjectured that no polynomial-time algorithm can find a planted clique of size $o(\sqrt{n})$, and numerous computational hardness results, such as sum-of-squares lower bounds [5], have been derived.

Remark 1.1. Note that under the planted cycles model, H^* is a random 2-factor and thus likely consists of multiple cycles. We can consider a variation of the model where H^* is conditioned to be a single Hamiltonian cycle on δn randomly chosen vertices, and our main results continue to hold. In particular, our achievability result in Theorem 1.1 holds conditionally on any H^* , so in particular it holds when H^* is a single cycle (see Theorem 2.4). Our algorithmic result similarly holds conditionally on H^* (see Theorem 4.1). As for the impossibility result, we leverage a reduction argument along with the observation that a random 2-factor H^* is a cycle with probability at least $1/(\delta n)$. The detailed arguments are included in Section C, Lemma C.3.

1.2 Analytical and algorithmic innovations

To delineate the achievability and impossibility results, we characterize the structure of the difference graph $H^* \Delta H$ between the planted cycles H^* and any other 2-factor H supported on δn vertices. Since $|H^*| = |H|$, exactly half of the edges in $H^* \Delta H$ are planted. Moreover, this difference graph can be decomposed into a vertex-disjoint union of circuits—closed trails³—where edges must alternate between planted and unplanted edges at shared vertices in $V(H^*) \cap V(H)$.

The core challenge arises in analyzing how the presence of these circuits affects the recoverability of H^* . When $\delta = 1$, the possibility of recovering H^* is driven by the presence of balanced circuits. Intuitively, if the number of balanced circuits—those with equal numbers of planted and unplanted edges—is small, then any competing 2-factor H must largely overlap with H^* , making almost-exact recovery feasible. Conversely, if there exists a diverging number of balanced circuits, then there are many plausible 2-factors that differ significantly from H^* , rendering almost-exact recovery impossible. However, beyond the special case $\delta = 1$, the situation becomes significantly more complex: individual circuits in the decomposition may be imbalanced, as long as the collection as a whole remains balanced (i.e., has half of its edges planted). This added layer of structural flexibility necessitates more refined tools for counting and characterizing such circuits.

Generating function approach To address this complexity, we develop a novel generating function approach that systematically counts balanced and imbalanced circuits. It turns out that the relevant object to count is the set of (a, b) -trails—trails consisting of a planted and b unplanted edges, alternating at shared vertices in $V(H^*) \cap V(H)$ (See the complete definition in Definition 2.1). We show that when the number of (a, b) -trails with $b \geq a$ is small, then almost-exact recovery is information-theoretically possible. Conversely, the presence of many perfectly balanced trails (with $a = b$) implies impossibility.

Specifically, let $c_{a,b}$ be a constant such that the expected number of (a, b) -trails from a given vertex to vertices in $V(H^*)$ is approximately equal to $c_{a,b}$. We show that the generating function of the sequence $(c_{a,b})_{a,b=1}^\infty$ is

$$g(x, y) \triangleq \sum_{a,b=1}^{\infty} c_{a,b} x^a y^b = \sum_{k=1}^{\infty} \left(\frac{2x}{1-x} \cdot \frac{\delta \lambda y}{1 - (1-\delta)\lambda y} \right)^k.$$

³A trail is a sequence of distinct edges $(e_1, e_2, \dots, e_{n-1})$ where $e_i = (v_i, v_{i+1})$ for $1 \leq i \leq n-1$ and $n \geq 1$. A circuit is a closed trail with $v_n = v_1$.

This generating function admits a natural interpretation. We assign weight x to each planted edge and weight y to each unplanted edge, and then enumerate the weighted (a, b) -trails from a given vertex to vertices in $V(H^*)$ in expectation as follows:

- The factor $\frac{\delta \lambda y}{(1 - (1 - \delta) \lambda y)} = \delta \lambda y (1 + (1 - \delta) \lambda y + ((1 - \delta) \lambda y)^2 + \dots)$ enumerates an unplanted segment, where the exponent of y records the number of unplanted edges. In such a segment, all but the last edge terminate at unplanted vertices. Hence, the factor $(1 - \delta) \lambda$ captures the expected number of unplanted edges incident to an unplanted vertex, while $\delta \lambda$ is the expected number of unplanted edges incident to a planted vertex.
- The factor $\frac{2x}{1 - x} = 2x (1 + x + x^2 + \dots)$ enumerates a planted segment, where the exponent of x records the number of planted edges. The coefficient 2 arises from the choice of two directions when traversing a planted cycle.
- Finally, the summation over $k \geq 1$ accounts for the possibility of switching between unplanted and planted segments arbitrarily many times.

To bound the number of circuits with more unplanted than planted edges, we choose $0 < x < 1 < y$ such that $xy = 1$. In this case, $g(x, y)$ upper bounds $\sum_{a \leq b} c_{a, b}$, and more generally, we have $\sum_{b \geq a + \ell} c_{a, b} \leq g(x, y)(x/y)^{\ell/2}$ for any $\ell \geq 0$,⁴ which decays exponentially in ℓ . Hence, if there exists such a pair (x, y) for which $g(x, y)$ is finite, we can conclude that the total number of (a, b) -trails with $b \geq a$ is bounded—implying almost-exact recovery is possible. We show that this occurs precisely when $\lambda < \frac{1}{(\sqrt{2\delta} + \sqrt{1 - \delta})^2}$.

Conversely, when $\lambda > \frac{1}{(\sqrt{2\delta} + \sqrt{1 - \delta})^2}$, the generating function $g(x, y)$ diverges for all choices of x, y satisfying $0 < x < 1 < y$ and $xy = 1$. While this divergence does not imply that $c_{a, a}$ is large for every a , we prove that there exists an m^* such that $c_{m^*, m^*} > 1$ (cf. Lemma 3.4). This enables us to construct a supercritical branching process that recursively grows many balanced (a, b) -paths with $a = b$, rooted at any given vertex. These paths, in turn, can be stitched together into long cycles, ultimately forming exponentially many 2-factors that significantly deviate from the planted H^* . This establishes the impossibility of almost-exact recovery in this regime.

Greedy search algorithm When $\lambda < \frac{1}{(\sqrt{2\delta} + \sqrt{1 - \delta})^2}$, we further design a novel polynomial-time greedy algorithm to achieve almost-exact recovery of the planted cycles. While enumerating all cycles in a graph is computationally intractable, our algorithm circumvents this bottleneck by leveraging structural insights from the generating function analysis and the decomposition of the symmetric difference $H^* \Delta H$.

A central observation is that when the generating function is bounded, any subgraph consisting of a disjoint union of cycles and a small number (i.e., $o(n)$) of paths, with a total of $\delta n - o(n)$ edges, differs from the planted cycles H^* by at most $o(n)$ edges. Guided by this observation, our algorithm iteratively grows a subgraph that consists of a disjoint union of cycles and a small number of paths. It begins with an empty graph H and incrementally grows it by searching through the collection of trails of length less than $\log n$ to identify “good” candidates P to update H to $H \Delta P$ (i.e., the XOR operation). These “good” trails, when XOR’ed onto H , either strictly increase $|H|$ without introducing new degree-1 vertices or increase $|H|$ by at least $\sqrt{\log n}$ while introducing at most 2 new degree-1 vertices. Such “good” trails are guaranteed to exist thanks to the trail decomposition of $H^* \Delta H$: One can always take P as either a full trail with more planted edges than unplanted ones or a segment of a trail with $\sqrt{\log n}$ more planted edges than unplanted ones. Importantly, these “good” trails can be efficiently identified through exhaustive search, as the total number of trails of length less than $\log n$ is $n^{O(1)}$ with high probability.

It is instructive to contrast this setting with the planted clique problem, where similar greedy algorithms are expected to fail. In that context, attempting to grow a clique C from an empty graph by iteratively adding vertices connected to all current members typically stalls at size $\Theta(\log n)$: at that point, many vertices in C are likely not part of the planted clique, and the chance of finding a vertex connected to all of C becomes negligible, due to the independence of edge connections outside the planted structure. The algorithm thus becomes trapped in a local optimum. By contrast, the planted cycles problem offers a more benign combinatorial landscape. Even if the current subgraph H contains many unplanted edges, the trail decomposition of $H \Delta H^*$ ensures that “good” trails always exist to continue growing H , until getting close to H^* .

1.3 Organization

The rest of this paper is structured as follows. Section 2 includes the achievability proofs, with Theorem 2.4 corresponding to the achievability statement of Theorem 1.1. Section 3 treats the impossibility results, with Theorem 3.1 corresponding to the impossibility statement of Theorem 1.1. Finally, Section 4 proposes a computationally efficient estimator that achieves the almost exact recovery threshold, as stated in Theorem 4.1. Section 5 includes several directions for future work.

⁴Note that $\sum_{b \geq a + \ell} c_{a, b} = \sum_{b \geq a + \ell} c_{a, b}(xy)^{(a+b)/2} = \sum_{b \geq a + \ell} c_{a, b} x^a y^b (x/y)^{(b-a)/2} \leq (x/y)^{\ell/2} \sum_{b \geq a + \ell} c_{a, b} x^a y^b$.

2 Generating function bounds and proof of possibility of recovery

The following lemma characterizes the difference graph of two 2-factors, both on (possibly different) subsets of m vertices. The proof follows from a general result in [16, Theorem 1], which characterizes the existence of a Eulerian circuit that always traverses from one edge class to another edge class at every vertex.

Lemma 2.1. *Let H^* be a planted 2-factor on m vertices, and let H be another 2-factor on a possibly different set of m vertices. Then*

- *The difference graph $H^* \Delta H$ is a graph where every vertex has degree 2 or 4. Furthermore, every degree-4 vertex has two incident red (planted) edges and two incident blue (unplanted) edges.*
- *Furthermore, every connected component of $H^* \Delta H$ has a Eulerian circuit whose edge color must alternate between red (planted) and blue (unplanted) at vertices in $V(H^*) \cap V(H)$.*

The second part ensures that for any vertex in a circuit that is in the planted 2-factor, at least one of the adjacent edges in the circuit is a planted edge.

Proof of Lemma 2.1. To prove the first part, observe that for any $v \in [n]$, the sum of the degrees of v in H^* and H is either 0, 2, or 4. The degree of v in $H^* \Delta H$ is equal to the sum of degrees of v in H^* and H , minus twice the number of edges incident to v in $H^* \cap H$. It follows that v has a degree of either 0, 2, or 4 in $H^* \Delta H$. Observe that any degree-4 vertex in $H^* \Delta H$ must have no edges incident to v in $H^* \cap H$, which implies that it has two red (planted) incident edges with the remaining two edges being blue (unplanted).

To prove the second part, we appeal to a general result of Kotzig [16]. Theorem 1 therein considers the setting of a connected graph $G = (V, E)$ where every vertex has an even degree and is associated with a labeled set of its incident edges. That is, a vertex v is associated with a labeled partition of its incident edges $Q_v = (Q_{v,1}, Q_{v,2}, \dots)$, where $(v, w) \in Q_{v,i}$ means that the edge (v, w) is given label i relative to v . The result states that if $|Q_{v,i}| \leq \frac{1}{2} \deg(v)$ for all $v \in V$, then there exists a Eulerian circuit of G such that for every pair of adjacent edges of the form $(u, v), (v, w)$, we have that (u, v) and (v, w) have different labels relative to v .

In our setting, we apply [16, Theorem 1] to each connected component of $H^* \Delta H$. There are four different types of vertices, depending on the colors of the incident edges: 2 red, 2 blue, 1 red and 1 blue, or 2 red and 2 blue. The vertices belong to $V(H^*) \cap V(H)$ in the latter two cases but not in the first two. In the first three cases, we simply classify the two incident edges into two label classes. For the third case, we put the red edges in one class and the blue edges in another class. The result then follows immediately from [16, Theorem 1]. $\square$




Figure 1: An example where H^* is a 4-cycle with red solid edges, H is a 4-cycle with blue dashed edges, and $H^* \Delta H$ has a $(3, 3)$ -Eulerian circuit.

In light of Lemma 2.1, $H^* \Delta H$ can be represented as a vertex-disjoint union of a certain class of circuits. To specify that class of circuits, we introduce the definition of (a, b) -circuit, or more generally, (a, b) -trail. See Figure 1 for a graphical illustration.

Definition 2.1. Consider a graph G with a planted 2-factor on a subset of its vertices. For $a \geq 0, b \geq 1$, an (a, b) -trail in G is a trail in G containing exactly a planted edges and b unplanted edges, such that the first edge is unplanted, the last edge is planted (if $a \geq 1$), no edge is repeated, and no two consecutive unplanted edges meet at a vertex that is contained in the planted 2-factor. An (a, b) -path is an (a, b) -trail that does not repeat vertices. An (a, b) -circuit is a closed (a, b) -trail.

The following lemma gives an upper bound on the expected number of (a, b) -trails, which will be used later to bound $\mathbb{E}[|H^* \Delta H|]$.

Lemma 2.2. *Let $G \sim \mathcal{G}(n, \lambda, \delta)$ be the graph generated from the planted cycles model and let H^* be the planted 2-factor on the set of δn vertices. Let v and v' be (not necessarily distinct) vertices. Fix $a \geq 1, b \geq 1$. Then, conditioned on H^* , the expected number of (a, b) -trails from v to v' in G is at most $c_{a,b}/(\delta n)$, where*

$$c_{a,b} \triangleq (\lambda(1-\delta))^b \sum_{k=1}^{\infty} (2\delta/(1-\delta))^k \binom{a-1}{k-1} \binom{b-1}{k-1}. \quad (2.1)$$

The expected number of $(0, b)$ trails from v to v' in G is at most $(1-\delta)^{b-1} \lambda^b/n$.

Proof. We first handle the case where $a = 0$. We need to bound the expected number of trails containing exactly b unplanted edges that start at v and end at v' . Furthermore, the $b-1$ intermediate vertices must be unplanted. Since there are at most $((1-\delta)n)^{b-1}$ such potential trails, and each one is present with probability $(\frac{\lambda}{n})^b$, the expected number of $(0, b)$ trails is at most $((1-\delta)n)^{b-1} \cdot (\frac{\lambda}{n})^b = (1-\delta)^{b-1} \lambda^b/n$.

Next, we consider $a \geq 1$. For any $k > 0$ and $a_1, \dots, a_k, b_1, \dots, b_k > 0$, let an $((a_1, \dots, a_k), (b_1, \dots, b_k))$ -trail be an $(\sum_i a_i, \sum_i b_i)$ -trail that starts with b_1 unplanted edges followed by a_1 planted edges, then b_2 unplanted edges, a_2 planted edges and so on. Let $\mathcal{W}_{\vec{a}, \vec{b}}(v, v')$ denote the set of $(\vec{a}, \vec{b})$ -trails from v to v' in the complete graph K_n (where edges in the complete graph are red if they are contained in H^* , and they are blue otherwise). Then we let

$$\mathcal{W}_{a,b}(v, v') = \bigcup_{k \in \mathbb{N}_+} \bigcup_{\vec{a}, \vec{b} \in \mathbb{N}_+^k : \|\vec{a}\|_1 = a, \|\vec{b}\|_1 = b} \mathcal{W}_{\vec{a}, \vec{b}}(v, v')$$

denote the set of (a, b) -trails from v to v' in the complete graph K_n . For each (a, b) -trail T in $\mathcal{W}_{a,b}(v, v')$, $\mathbb{P}\{T \subset G\} = (\lambda/n)^b$. Thus,

$$\mathbb{E} \left[\sum_{T \in \mathcal{W}_{a,b}(v, v')} \mathbf{1}_{\{T \subset G\}} \right] = |\mathcal{W}_{a,b}(v, v')| \left(\frac{\lambda}{n} \right)^b.$$

We claim that for each $\vec{a}, \vec{b} \in \mathbb{N}_+^k$,

$$|\mathcal{W}_{\vec{a}, \vec{b}}(v, v')| \leq 2(2\delta n)^{k-1} ((1-\delta)n)^{\|\vec{b}\|_1 - k}. \quad (2.2)$$

To see this, we first choose a set of directed paths of length a_i in H^* for each $1 \leq i \leq k$. For each of the first $k-1$ directed paths, there are at most δn choices for the starting vertex, and two possible directions to traverse the planted cycles it is in. For the last directed path of length a_k in H^* , since it must end at a given vertex v' in H^* , the only choice is which direction to traverse the cycle that v' is in. Thus there are at most $2^k (\delta n)^{k-1}$ choices in total for the set of planted directed paths. Now, given such a set, we are left to determine the set of directed paths of length b_i consisting of unplanted edges for each $1 \leq i \leq k$. For each of these directed paths, since the starting and ending vertices have already been fixed, we only need to choose the unplanted vertices in the middle. There are in total $b-k$ unplanted vertices in the middles of paths, and thus there are at most $[(1-\delta)n]^{b-k}$ different choices. In total, this results in at most $2^k (\delta n)^{k-1} [(1-\delta)n]^{b-k}$ distinct $(\vec{a}, \vec{b})$ -trails from v to v' in the complete graph K_n . Therefore,

$$\begin{aligned} \mathbb{E} \left[\sum_{T \in \mathcal{W}_{a,b}(v, v')} \mathbf{1}_{\{T \subset G\}} \right] &\leq \sum_{k=1}^{\infty} \sum_{\vec{a}, \vec{b} \in \mathbb{N}_+^k : \|\vec{a}\|_1 = a, \|\vec{b}\|_1 = b} 2^k (\delta n)^{k-1} ((1-\delta)n)^{b-k} \cdot (\lambda/n)^b \\ &= \sum_{k=1}^{\infty} \sum_{\vec{a}, \vec{b} \in \mathbb{N}_+^k : \|\vec{a}\|_1 = a, \|\vec{b}\|_1 = b} \frac{1}{\delta n} (2\delta/(1-\delta))^k (\lambda(1-\delta))^b \\ &= \frac{1}{\delta n} (\lambda(1-\delta))^b \sum_{k=1}^{\infty} (2\delta/(1-\delta))^k \binom{a-1}{k-1} \binom{b-1}{k-1} = \frac{c_{a,b}}{\delta n}, \end{aligned} \quad (2.3)$$

where the second equality holds because the number of distinct $\vec{a} \in \mathbb{N}_+^k$ such that $\|\vec{a}\|_1 = a$ is $\binom{a-1}{k-1}$. Therefore, the expected number of (a, b) -trails from v to v' in G is at most $c_{a,b}/\delta n$ for any choice of H^* . $\square$

To study $c_{a,b}$, we consider its generating function; that is,

$$\sum_{a,b=1}^{\infty} c_{a,b} x^a y^b = \sum_{a,b=1}^{\infty} \sum_{k=1}^{\infty} \sum_{\vec{a}, \vec{b} \in \mathbb{N}_+^k : \|\vec{a}\|_1 = a, \|\vec{b}\|_1 = b} (2\delta/(1-\delta))^k (\lambda(1-\delta))^b x^a y^b$$

$$\begin{aligned}
&= \sum_{k=1}^{\infty} \sum_{\vec{a}, \vec{b} \in \mathbb{N}_+^k} (2\delta/(1-\delta))^k (\lambda(1-\delta))^{\|\vec{b}\|_1} x^{\|\vec{a}\|_1} y^{\|\vec{b}\|_1} \\
&= \sum_{k=1}^{\infty} \prod_{i=1}^k \left(\sum_{a_i, b_i=1}^{\infty} \frac{2\delta}{1-\delta} x^{a_i} ((1-\delta)\lambda y)^{b_i} \right) \\
&= \sum_{k=1}^{\infty} \left(\frac{2x}{1-x} \cdot \frac{\delta\lambda y}{1-(1-\delta)\lambda y} \right)^k \triangleq g(x, y).
\end{aligned} \tag{2.4}$$

If there exist x and y with $0 < x < 1$, $0 < y < 1/[\lambda(1-\delta)]$, and $xy = 1$ for which $\frac{2x}{1-x} \cdot \frac{\delta\lambda y}{1-(1-\delta)\lambda y} < 1$ then $\sum_{k=1}^{\infty} \left(\frac{2x}{1-x} \cdot \frac{\delta\lambda y}{1-(1-\delta)\lambda y} \right)^k$ converges for these values of x and y . That means that the expected number of (a, b) -circuits with $a = b$ is upper bounded by $g(x, y)$, and the expected number of (a, b) -circuits with $|a - b| = O(1)$ is at most $O(g(x, y))$. With some more work, we can use that to prove that with probability $1 - o(1)$, every 2-factor on δn vertices is nearly equal to the planted one.

Conversely, if $\frac{2x}{1-x} \cdot \frac{\delta\lambda y}{1-(1-\delta)\lambda y} > 1$ for all $x < 1$, $y < 1/[\lambda(1-\delta)]$ with $xy = 1$ then $\sum_{k=1}^{\infty} \left(\frac{2x}{1-x} \cdot \frac{\delta\lambda y}{1-(1-\delta)\lambda y} \right)^k$ diverges for all such x and y . That can be used to show that the expected number of (a, b) -circuits with $a = b$ diverges as n increases. This brings us to the question of when these cases hold, which has the following answer.

Lemma 2.3. *If $\lambda < \frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$, then there exist x and y with $0 < x < 1$, $1 < y < 1/[\lambda(1-\delta)]$, and $xy = 1$ such that $\frac{2x}{1-x} \cdot \frac{\delta\lambda y}{1-(1-\delta)\lambda y} < 1$. On the other hand, if $\lambda > \frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$, then $\frac{2x}{1-x} \cdot \frac{\delta\lambda y}{1-(1-\delta)\lambda y} > 1$ for all x and y with $0 < x < 1$, $1 < y < 1/[\lambda(1-\delta)]$, and $xy = 1$.*

Proof. If $xy = 1$ then

$$\frac{2x}{1-x} \cdot \frac{\delta\lambda y}{1-(1-\delta)\lambda y} = \frac{2}{1-x} \cdot \frac{\delta\lambda}{1-(1-\delta)\lambda/x} = \frac{2\delta\lambda x}{(1-x)(x-(1-\delta)\lambda)}.$$

At this point, the question of whether there is $(1-\delta)\lambda < x < 1$ for which this is less than 1 is equivalent to the question of whether there is $(1-\delta)\lambda < x < 1$ for which $(1-x)(x-(1-\delta)\lambda) > 2\delta\lambda x$. Let

$$f(x) = 2\delta\lambda x - (1-x)(x-(1-\delta)\lambda) = x^2 + ((3\delta-1)\lambda-1)x + (1-\delta)\lambda.$$

Note that $f(1) = 2\delta\lambda \geq 0$, $f((1-\delta)\lambda) = 2\delta\lambda^2(1-\delta) \geq 0$, and $f(x)$ achieves its minimum value of $f^* \triangleq (1-\delta)\lambda - ((3\delta-1)\lambda-1)^2/4$ at $x = x^* \triangleq (1-(3\delta-1)\lambda)/2$. Thus there exists $(1-\delta)\lambda < x < 1$ for which $f(x) < 0$, provided that

$$f^* < 0 \text{ and } (1-\delta)\lambda \leq x^* \leq 1. \tag{2.5}$$

Note that $f^* < 0$ is equivalent to $(3\delta-1)^2\lambda^2 - (2\delta+2)\lambda + 1 > 0$, which happens when $\lambda < \lambda_-$ or $\lambda > \lambda_+$, where

$$\begin{aligned}
\lambda_{\pm} &= \frac{(2\delta+2) \pm \sqrt{(2\delta+2)^2 - 4(3\delta-1)^2}}{2(3\delta-1)^2} = \frac{\delta+1 \pm \sqrt{(\delta+1)^2 - (3\delta-1)^2}}{(3\delta-1)^2} = \frac{\delta+1 \pm \sqrt{8\delta-8\delta^2}}{(3\delta-1)^2} \\
&= \frac{(\sqrt{2\delta} \pm \sqrt{1-\delta})^2}{(3\delta-1)^2} = \frac{1}{(\sqrt{2\delta} \mp \sqrt{1-\delta})^2}.
\end{aligned}$$

Also, $(1-\delta)\lambda \leq x^* \leq 1$ is equivalent to $(1+\delta)\lambda \leq 1$. Note that $\lambda_- \leq \frac{1}{1+\delta}$ and $\lambda_+ \geq \frac{1}{1+\delta}$. Thus, (2.5) holds provided that $\lambda < \lambda_- = \frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$.

Conversely, if $\lambda > \lambda_- = \frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$, then either $f^* > 0$ or $(1-\delta)\lambda \geq x^*$. Thus $f(x) > 0$ for all $(1-\delta)\lambda < x < 1$. $\square$

The previous lemma finally allows us to prove that we can achieve almost exact recovery whenever $\lambda < \frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$.

Theorem 2.4. *Let $\delta \in (0, 1]$ and $0 < \lambda < \frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$. Suppose $G \sim \mathcal{G}(n, \lambda, \delta)$ is the graph generated from the planted cycles model conditional $H^* = h$. There exists C independent of n such that $\mathbb{E}[|H^* \Delta H|] \leq C$ for any 2-factor H on δn vertices in G .*

Proof. Recall that $H^* \Delta H$ can be represented as a vertex-disjoint union of (a, b) -circuits such that exactly half of the total edges are planted. However, each individual (a, b) -circuit may be unbalanced, i.e., having an unequal number of planted and unplanted edges. Nevertheless, we will argue that the expected number of (a, b) -circuits with more than $(1/2 - \epsilon)$ fraction of their edges unplanted is bounded.

Towards this end, for each (a, b) -circuit C in graph G , define its excess $\text{ex}(C) = b - (1/2 - \epsilon)(a + b)$, where $\epsilon > 0$ will be specified later. Then define the total excess over all possible valid circuits with positive excess as

$$\Gamma = \sum_{C \subset G} \text{ex}(C) \mathbf{1}_{\{\text{ex}(C) > 0\}},$$

Given any 2-factor on δn vertices in G , suppose that $H^* \Delta H$ is the vertex-disjoint union of (a_i, b_i) -circuits C_i for $1 \leq i \leq m$. Then by definition,

$$\begin{aligned} \frac{1}{2} |H^* \Delta H| &= \sum_{i=1}^m b_i = \sum_{i=1}^m [b_i - (1/2 - \epsilon)(a_i + b_i)] + (1/2 - \epsilon) |H^* \Delta H| \\ &= \sum_{i=1}^m \text{ex}(C_i) + (1/2 - \epsilon) |H^* \Delta H| \\ &\leq \Gamma + (1/2 - \epsilon) |H^* \Delta H|. \end{aligned}$$

It follows that $|H^* \Delta H| \leq \Gamma / \epsilon$.

It remains to bound $\mathbb{E}[\Gamma]$. In order to do that, we decompose $\Gamma = \Gamma_0 + \Gamma_1$, where Γ_0 sums over the $(0, b)$ -circuits C and Γ_1 sums over the (a, b) -circuits C for $a \geq 1$.

By Lemma 2.2, the expected number of $(0, b)$ -trails from v to v' in G is at most $\lambda^b (1 - \delta)^{b-1} / n$. Since a $(0, b)$ -circuit is a special case of a $(0, b)$ -trail that starts and ends at a vertex v not in H^* , it follows that the expected number of $(0, b)$ -circuits is at most $\lambda^b (1 - \delta)^{b-1} / n \times (1 - \delta)n = \lambda^b (1 - \delta)^b$. Therefore,

$$\mathbb{E}[\Gamma_0] \leq (1/2 + \epsilon) \sum_{b \geq 1} b (\lambda(1 - \delta))^b = (1/2 + \epsilon) \frac{\lambda(1 - \delta)}{(1 - \lambda(1 - \delta))^2} \quad (2.6)$$

where the last equality holds because $\sum_{b \geq 1} b x^b = x / (1 - x)^2$ for $|x| < 1$ and $\lambda(1 - \delta) < 1$.

Next, we bound $\mathbb{E}[\Gamma_1]$. By Lemma 2.2, for $a, b \geq 1$, the expected number of (a, b) -trails from v to v' in G is at most $c_{a,b} / (\delta n)$. Since an (a, b) -circuit is a special case of an (a, b) -trail that starts and ends at a vertex v in H^* , it follows that the expected number of (a, b) -circuits is at most $c_{a,b} / (\delta n) \times \delta n = c_{a,b}$. It follows that

$$\mathbb{E}[\Gamma_1] \leq \sum_{a, b \geq 1} [b - (1/2 - \epsilon)(a + b)] c_{a,b} \cdot \mathbf{1}_{\{b > (1/2 - \epsilon)(a + b)\}}.$$

Now, we bound the RHS of the last display equation using the generating function of $c_{a,b}$. Recall that by Lemma 2.3, there exist x_0, y_0 such that $0 < x_0 < 1$, $0 < y_0 < \frac{1}{\lambda(1 - \delta)}$, $x_0 y_0 = 1$, and $\frac{2x_0}{1 - x_0} \cdot \frac{\delta \lambda y_0}{1 - (1 - \delta) \lambda y_0} < 1$. Next, observe that there must exist $y > y_0$ such that $\frac{2x_0}{1 - x_0} \cdot \frac{\delta \lambda y}{1 - (1 - \delta) \lambda y} < 1$. Also, let $x = x_0$. At this point, $x^2 < 1 < xy$ so there must exist some $0 < \epsilon < 1/2$ such that $x^{1+2\epsilon} y^{1-2\epsilon} = 1$. Therefore, for all $b \geq (1/2 - \epsilon)(a + b)$,

$$b - (1/2 - \epsilon)(a + b) \leq \frac{y/x}{y/x - 1} (y/x)^{b - (1/2 - \epsilon)(a + b)} = \frac{y/x}{y/x - 1} x^a y^b,$$

where the inequality holds because $t^{\alpha+1} \geq 1 + (\alpha + 1)(t - 1) \geq \alpha(t - 1)$ for $t \geq 1$ and $\alpha \geq 0$; the equality holds due to $x^{1+2\epsilon} y^{1-2\epsilon} = 1$. Combining the last displayed equations gives that

$$\begin{aligned} \mathbb{E}[\Gamma_1] &\leq \frac{y/x}{y/x - 1} \sum_{a, b \geq 1} c_{a,b} x^a y^b \\ &= \frac{y/x}{y/x - 1} \sum_{k=1}^{\infty} \left(\frac{2x}{1 - x} \cdot \frac{\delta \lambda y}{1 - (1 - \delta) \lambda y} \right)^k \\ &\leq \frac{y/x}{y/x - 1} \frac{1}{1 - \frac{2x}{1 - x} \cdot \frac{\delta \lambda y}{1 - (1 - \delta) \lambda y}}, \end{aligned}$$

In conclusion, $\mathbb{E}[|H^* \Delta H|] \leq \mathbb{E}[\Gamma] / \epsilon \leq C$ for some constant C only depending on $\epsilon, \lambda, \delta$.

□

3 Generating functions to prove impossibility of recovery

In this section, we will prove the following theorem, establishing that when $\lambda > \frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$, any estimator must suffer at least constant error.

Theorem 3.1. *If*

$$\lambda > \frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}, \quad (3.1)$$

then there exists a constant $\epsilon > 0$ depending only on λ and δ such that for any estimator $\hat{H}$,

$$\mathbb{P}\left(\ell(H^*, \hat{H}) \geq \epsilon/2\right) \geq 1 - n^{-\Omega(\log n)}. \quad (3.2)$$

It follows that for all sufficiently large n ,

$$\mathbb{E}\left[\ell(\hat{H}, H^*)\right] \geq \epsilon/4.$$

The proof of Theorem 3.1 is similar in spirit to the proof of [13, Theorem 2.3]. Let μ_G denote the posterior distribution:

$$\mu_G(H) \triangleq \mathbb{P}\{H^* = H \mid G\} = \frac{1}{|\mathcal{H}(G)|} \mathbf{1}_{\{H \in \mathcal{H}(G)\}}, \quad \forall H \in \mathcal{H}(G), \quad (3.3)$$

where $\mathcal{H}(G)$ denotes the set of degree-2 subgraphs on δn vertices. That is, $\mu_G(\cdot)$ is the uniform distribution over degree-2 subgraphs on δn vertices. A simple yet crucial observation is that while a random draw $\tilde{H}$ from the posterior distribution (3.3) may not minimize the reconstruction error, its reconstruction error can be related to that of any estimator $\hat{H}$ as follows: for any $D > 0$,

$$\mathbb{P}\left(\ell(H^*, \hat{H}) < D\right) \leq \sqrt{\mathbb{P}\left(\ell(H^*, \tilde{H}) < 2D\right)}. \quad (3.4)$$

(see [13, Eq. 3.2] for a short proof). Therefore, it suffices to prove (3.2) holds for the posterior sample $\tilde{H}$, which further reduces to demonstrating that the observed graph G contains many more $H \in \mathcal{H}$ that are far from H^* than those close to H^* .

We begin by defining two sets, one containing subgraphs H that are similar to H^* , and the other containing subgraphs H that are markedly different from H^* . Concretely, given $\epsilon > 0$, define

$$\begin{aligned} \mathcal{H}_{\text{good}} &= \{H \in \mathcal{H} : \ell(H, H^*) < \epsilon, H \subset G\} \\ \mathcal{H}_{\text{bad}} &= \{H \in \mathcal{H} : \ell(H, H^*) \geq \epsilon, H \subset G\}. \end{aligned}$$

Using the fact that the posterior distribution μ_G is the uniform distribution over all possible k -factors contained in the observed graph G , we have

$$\mathbb{P}\left\{\ell(\tilde{H}, H^*) < \epsilon \mid G, H^*\right\} = \frac{|\mathcal{H}_{\text{good}}|}{|\mathcal{H}_{\text{good}}| + |\mathcal{H}_{\text{bad}}|}.$$

The following lemmas give high probability bounds on $|\mathcal{H}_{\text{good}}|$ and $|\mathcal{H}_{\text{bad}}|$.

Lemma 3.2. *For $\epsilon < \min\{1, 2\lambda/\delta\}$, conditioned on H^* , with probability at least $1 - e^{-\Omega(n)}$,*

$$|\mathcal{H}_{\text{good}}| \leq \frac{\epsilon \delta n}{2} \left(\frac{2e}{\epsilon} \cdot \frac{4e\lambda}{\epsilon\delta} \right)^{\epsilon \delta n/2}.$$

Proof. With probability $1 - e^{-\Omega(n)}$, there are at most $2\lambda n$ unplanted edges. Conditioned on this event, we bound the cardinality of $\mathcal{H}_{\text{good}}$. Observe that for $H \in \mathcal{H}_{\text{good}}$, $|H| = \delta n$ and $|H \Delta H^*| \leq \epsilon |H^*| = \epsilon \delta n$. To determine H , it suffices to choose the planted edges for $H^* \setminus H$ and unplanted edges for $H \setminus H^*$. Therefore, there are at most $\binom{\delta n}{i} \binom{2\lambda n}{i}$ such graphs H with set difference exactly $2i$ relative to H^* . Hence, we can bound $|\mathcal{H}_{\text{good}}|$ by

$$\begin{aligned} \sum_{i=1}^{\epsilon \delta n/2} \binom{\delta n}{i} \binom{2\lambda n}{i} &\leq \frac{\epsilon \delta n}{2} \binom{\delta n}{\epsilon \delta n/2} \binom{2\lambda n}{\epsilon \delta n/2} \\ &\leq \frac{\epsilon \delta n}{2} \left(\frac{2e}{\epsilon} \cdot \frac{4e\lambda}{\epsilon\delta} \right)^{\epsilon \delta n/2}. \end{aligned}$$

Here the first inequality holds since by assumption $\epsilon < \min\{1, 2\lambda/\delta\}$, which implies that the maximal summand is $\binom{\delta n}{\epsilon \delta n/2} \binom{2\lambda n}{\epsilon \delta n/2}$. The second step uses $\binom{n}{k} \leq (en/k)^k$. $\square$

Lemma 3.3. *Assume that (3.1) holds. Then there exist constants c_1 and c_2 that only depend on λ, δ , such that for all $\epsilon \leq c_2/\delta$, conditioned on H^* containing at most $(\log_2 n)^2$ cycles, with probability at least $1 - e^{-\Omega(n)}$,*

$$|\mathcal{H}_{\text{bad}}| \geq e^{c_1 n}. \quad (3.5)$$

Proof of Theorem 3.1. Combining Lemmas 3.2 and 3.3, we obtain that conditioned on H^* containing at most $(\log_2 n)^2$ cycles, with probability $1 - e^{-\Omega(n)}$,

$$\frac{|\mathcal{H}_{\text{good}}|}{|\mathcal{H}_{\text{good}}| + |\mathcal{H}_{\text{bad}}|} \leq \frac{|\mathcal{H}_{\text{good}}|}{|\mathcal{H}_{\text{bad}}|} \leq \frac{\epsilon \delta n}{2} \left(\frac{2e}{\epsilon} \cdot \frac{4e\lambda}{\epsilon \delta} \right)^{\epsilon \delta n/2} \cdot e^{-c_1 n}.$$

Setting $\epsilon > 0$ to be sufficiently small, we can bound the right-hand side by $n \cdot e^{-c_1 n/2}$. Therefore,

$$\begin{aligned} \left[\mathbb{P} \left\{ \ell(H^*, \hat{H}) < \epsilon/2 \right\} \right]^2 &\leq \mathbb{P} \left\{ \ell(H^*, \tilde{H}) < \epsilon \right\} \\ &= \mathbb{E}_{G, H^*} \left[\mathbb{P} \left\{ \ell(H^*, \tilde{H}) < \epsilon \mid G, H^* \right\} \right] \\ &= \mathbb{E}_{G, H^*} \left[\frac{|\mathcal{H}_{\text{good}}|}{|\mathcal{H}_{\text{good}}| + |\mathcal{H}_{\text{bad}}|} \right] \\ &\leq \left(1 - e^{-\Omega(n)} \right) n e^{-c_1 n/2} + e^{-\Omega(n)} + \mathbb{P} \{ \mathcal{E}^c \} = n^{-\Omega(\log n)}, \end{aligned}$$

where the first inequality comes from (3.4); $\mathcal{E}$ denotes the event that H^* contains at most $(\log_2 n)^2$ cycles; and the last inequality holds because $\mathbb{P} \{ \mathcal{E}^c \} \leq n^{-\Omega(\log n)}$ in view of Lemma C.1. $\square$

The remainder of this section is devoted to proving Lemma 3.3. Building on the circuit decomposition in Lemma 2.1, our goal is to show that with high probability, there exist exponentially many $\Theta(n)$ -length cycles containing an equal number of planted and unplanted edges—referred to as balanced cycles. Perhaps unexpectedly, these long balanced cycles can be assembled from constant-length paths that are themselves balanced, containing the same number of planted and unplanted edges.

We will search for these constant-length, balanced paths by constructing a linear number of trees, each containing many constant-length paths. Specifically, in Lemma 3.4, we first show that there exists a constant m^* for which $c_{m^*, m^*} > 1$. Moreover, in Lemma 3.6, we show that for any $a, b \geq 1$, the expected number of (a, b) -paths starting at a given vertex converges to $c_{a, b}$. Therefore, we can define a branching process, where each step goes from a planted vertex to all other planted vertices reachable from it by an (m^*, m^*) -path. The expected number of children of each node in this process will be greater than 1, enabling us to show that its probability of dying out is bounded away from 0. That allows us to construct many trees where each path to a leaf has an equal number of planted and unplanted edges.

Then, we apply a sprinkling argument, connecting the trees to form an exponentially large number of cycles that have equal numbers of planted and unplanted edges. Similar sprinkling constructions appeared in [9, 13], with [13] being the most similar to the present argument.

In [13], which studies the case of $\delta = 1$, the impossibility of almost exact recovery coincides with the presence of many long *alternating* cycles. Accordingly, their tree construction procedure searches for many trees that alternate between planted and unplanted edges across layers. In contrast, when $\delta < 1$, impossibility is driven instead by the presence of many long *balanced* cycles, which need not be alternating. Fortunately, we find that there still exist sufficiently many balanced cycles containing segments of length $2m^*$ comprising m^* planted edges and m^* unplanted edges, where m^* is some constant depending on λ and δ . If $\delta = 1$, one can simply take $m^* = 1$, recovering the case of alternating cycles. When $\delta < 1$, however, we need to take $m^* > 1$. Our tree construction procedure therefore attempts to find many trees built from such balanced segments (referred to as *path layers*) of length $2m^*$. The challenge is that adding a path layer at each step requires exploring neighborhoods of radius $2m^*$, which introduces dependencies and significantly complicates the analysis, as we will see.

In what follows, we describe our tree construction procedure in Section 3.1 and analyze it in Section 3.2. The sprinkling argument is presented in Section 3.3.

3.1 Tree construction

Recall that when $\lambda > \frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$, we have that $\frac{2x}{1-x} \cdot \frac{\delta \lambda y}{1-(1-\delta)\lambda y} > 1$ for all $x \in (0, 1)$ and $y \in (0, 1/(1-\delta)\lambda)$ such that $xy = 1$ due to Lemma 2.3. The following lemma further establishes that the power series $\sum_{k=1}^{\infty} \left(\frac{2x}{1-x} \cdot \frac{\delta \lambda y}{1-(1-\delta)\lambda y} \right)^k$ has a coefficient c_{m^*, m^*} satisfying $c_{m^*, m^*} > 1$.

Lemma 3.4. *Let $0 < \delta \leq 1$ and $\lambda > \frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$. Then there exists m^* such that $c_{m^*, m^*} > 1$. In particular, when $\delta = 1$, and $\lambda > 1/2$, we have $c_{a,a} = (2\lambda)^a$ and m^* can be chosen as 1.*

Proof. For the special case of $\delta = 1$, by the definition in (2.1), $c_{a,a} = (2\lambda)^a$. Thus when $\lambda > 1/2$, m^* can be chosen as 1. In general, note that $c_{a,b}$ is increasing in λ , so without loss of generality, we assume $\frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2} < \lambda < \frac{1}{1-\delta}$. Let

$$f(x, y) = \frac{2x}{1-x} \cdot \frac{\delta\lambda y}{1 - (1-\delta)\lambda y}$$

and $h(x) = f(x, 1/x)$. Let x^* be a value of x in $((1-\delta)\lambda, 1)$ that minimizes $h(x)$. By Lemma 2.3, $h(x^*) > 1$. Next, define the finite series expansion of $f(x, y)$ as

$$f_m(x, y) \triangleq 2 \sum_{i=1}^m x^i \cdot \sum_{j=1}^m ((1-\delta)\lambda y)^{j-1} \delta\lambda y = \frac{2\delta}{1-\delta} \sum_{i=1}^m x^i \cdot \sum_{j=1}^m ((1-\delta)\lambda y)^j$$

and $h_m(x) = f_m(x, 1/x)$. Note that as $m \rightarrow \infty$, $h_m(x)$ converges uniformly to $h(x)$ on any closed interval contained in $((1-\delta)\lambda, 1)$. So, there exists m_0 such that $h_{m_0}(x)$ has a local minimum strictly greater than 1 at some $x_0 \in ((1-\delta)\lambda, 1)$; that is, $h'_{m_0}(x_0) = 0$ and $h_{m_0}(x_0) > 1$. Note that

$$\begin{aligned} h'_m(x) &= \frac{\partial}{\partial x} f_m(x, y)|_{y=1/x} + \frac{\partial}{\partial y} f_m(x, y)|_{y=1/x} \cdot \frac{d(1/x)}{dx} \\ &= \frac{1}{x} \left(x \frac{\partial}{\partial x} - y \frac{\partial}{\partial y} \right) f_m(x, y)|_{y=1/x} \\ &= \frac{2\delta}{x(1-\delta)} \sum_{i,j=1}^m (i-j)x^i ((1-\delta)\lambda y)^j |_{y=1/x}. \end{aligned}$$

In particular, $h'_{m_0}(x_0) = 0$ implies that

$$\sum_{i,j=1}^{m_0} (i-j)x^i ((1-\delta)\lambda y)^j |_{x=x_0, y=1/x_0} = 0. \quad (3.6)$$

Now, consider picking a random $(\mathbf{I}, \mathbf{J})$ with probability given by

$$\mathbb{P}[(\mathbf{I}, \mathbf{J}) = (i, j)] = \frac{1}{h_{m_0}(x_0)} \frac{2\delta}{1-\delta} x^i ((1-\delta)\lambda y)^j |_{x=x_0, y=1/x_0}$$

for all $1 \leq i, j \leq m_0$. Then it follows from (3.6) that $\mathbb{E}[\mathbf{I}] = \mathbb{E}[\mathbf{J}]$. That means that the exponents of x and y in $f_{m_0}(x, y)$ are in some sense balanced, but ultimately we want to prove the existence of a large coefficient on a single term with equal powers of x and y . To this end, we consider a sequence of i.i.d. copies of $(\mathbf{I}, \mathbf{J})$, that is, $(\mathbf{a}_1, \mathbf{b}_1), \dots, (\mathbf{a}_\ell, \mathbf{b}_\ell) \stackrel{\text{i.i.d.}}{\sim} (\mathbf{I}, \mathbf{J})$. Let $\mathbf{a} = \sum_{i=1}^\ell \mathbf{a}_i$ and $\mathbf{b} = \sum_{i=1}^\ell \mathbf{b}_i$. By Chebyshev's inequality, with probability at least $1/2$,

$$|\mathbf{a} - \ell \mathbb{E}[\mathbf{I}]| < c\sqrt{\ell} \quad \text{and} \quad |\mathbf{b} - \ell \mathbb{E}[\mathbf{J}]| < c\sqrt{\ell}, \quad \text{where } c = 2\sqrt{\max(\text{Var}[\mathbf{a}_1], \text{Var}[\mathbf{b}_1])}.$$

Let $m^* = \lceil \ell \mathbb{E}[\mathbf{I}] + c\sqrt{\ell} \rceil$. Recall from (2.3) that

$$\begin{aligned} c_{m^*, m^*} &= (\lambda(1-\delta))^{m^*} \sum_{k=1}^{\infty} (2\delta/(1-\delta))^k \sum_{\vec{a}, \vec{b} \in \mathbb{N}_+^k : \|\vec{a}\|_1 = m^*, \|\vec{b}\|_1 = m^*} 1 \\ &\geq (\lambda(1-\delta))^{m^*} (2\delta/(1-\delta))^{\ell+1} \sum_{a,b=1}^{m^*-1} \sum_{\vec{a}, \vec{b} \in \mathbb{N}_+^\ell : \|\vec{a}\|_1 = a, \|\vec{b}\|_1 = b} 1 \\ &= (2\delta/(1-\delta))^{\ell+1} (\lambda(1-\delta))^{m^*} \sum_{a,b=1}^{m^*-1} \sum_{\vec{a}, \vec{b} \in \mathbb{N}_+^\ell : \|\vec{a}\|_1 = a, \|\vec{b}\|_1 = b} x^{m^*} y^{m^*} |_{x=x_0, y=1/x_0} \\ &\geq (2\delta/(1-\delta))^{\ell+1} (x\lambda(1-\delta)y)^{\lceil 2c\sqrt{\ell} \rceil} \sum_{a,b=m^* - \lceil 2c\sqrt{\ell} \rceil}^{m^*-1} \sum_{\vec{a}, \vec{b} \in \mathbb{N}_+^\ell : \|\vec{a}\|_1 = a, \|\vec{b}\|_1 = b} x^a (\lambda(1-\delta)y)^b |_{x=x_0, y=1/x_0} \\ &\geq \frac{2\delta}{1-\delta} (h_{m_0}(x_0))^\ell (\lambda(1-\delta))^{\lceil 2c\sqrt{\ell} \rceil} \cdot \mathbb{P}\left\{|\mathbf{a} - \ell \mathbb{E}[\mathbf{I}]| < c\sqrt{\ell}, |\mathbf{b} - \ell \mathbb{E}[\mathbf{J}]| \leq c\sqrt{\ell}\right\} \end{aligned}$$

$$\geq \frac{2\delta}{1-\delta} (h_{m_0}(x_0))^\ell (\lambda(1-\delta))^{\lceil 2c\sqrt{\ell} \rceil} \cdot \frac{1}{2},$$

where the first inequality holds by fixing $k = \ell + 1$ and choosing $a_{\ell+1} = m^* - a$ and $b_{\ell+1} = m^* - b$; the second inequality holds by restricting the sum over a, b to the range $a, b \geq m^* - 2c\sqrt{\ell}$, and using the facts that $x_0 < 1, \lambda(1-\delta)/x_0 < 1$. Finally, since $h_{m_0}(x_0) > 1$, by choosing ℓ sufficiently large, we have $c_{m^*, m^*} > 1$. $\square$

Later, in Lemma 3.6, we show that for any $a, b \geq 1$, the expected number of (a, b) -paths starting at a given planted vertex is asymptotically $c_{a,b}$. Combining this with Lemma 3.4, we conclude that when $\lambda > \frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$, the expected number of (m^*, m^*) -paths from any given planted vertex exceeds 1. This observation allows us to construct a suitable collection of trees, referred to as *path trees*, as illustrated in Figure 2. In particular, when $\delta = 1$, m^* can be chosen as 1, and each purple path in the path tree consists of one unplanted edge followed by one planted edge, thereby recovering the tree construction in [13].




Figure 2: A schematic representation of a path tree with a depth of 2. The red nodes are called “hub” nodes, and the dashed purple lines represent (m^*, m^*) -paths, where m^* is given in Lemma 3.4. A purple path along with the hub node below it is referred to as a “path layer” of the path tree.

The construction of path trees is given in Algorithm 1. The algorithm operates on a pre-selected set of *available* vertices, owing to a pre-processing step that reserves edges for later use in linking trees together (Algorithm 2). The available vertices are precisely those that are not endpoints of reserved edges.

Algorithm 1 Tree Construction

Input: Graph $G_0 \cup H^*$ on n vertices, available vertices $\mathcal{A}$ with $|\mathcal{A}| = (1 - 2\gamma)n$ for some $\gamma > 0$, length m^* , size parameter $\ell \in \mathbb{N}$

Output: A set of two-sided trees where each side has at least 2ℓ hub nodes.

```
1: Set  $\mathcal{T} = \emptyset$ .
2: for  $t \in \{1, 2, \dots, K := \frac{\gamma n}{\ell}\}$  do
3:   Select a random planted edge  $(u_0, u'_0)$  where  $u_0, u'_0 \in \mathcal{A}$ . If no such edge exists, return FAIL.
4:   Initialize  $T$  to be a two-sided tree containing only the center edge  $(u_0, u'_0)$ .
5:   Remove  $u_0$  and  $u'_0$  from  $\mathcal{A}$ .
6:   (Grow the left tree rooted at  $u_0$ .) Initialize the leaf queue to be  $\mathcal{L} \leftarrow \{u_0\}$ , and the cumulative size
   to be  $s \leftarrow 1$ .
7:   while  $\mathcal{L} \neq \emptyset$  and  $s < 2\ell$  do
8:     Let  $u \leftarrow \mathcal{L}.\text{pop}$ .
9:     (Find the descendants of  $u$ .) Let  $\mathcal{C}_u$  be the set of all vertices  $v$  such that there is a non-shortcutted
    $(m^*, m^*)$ -path from  $u$  to  $v$  with all of its vertices except  $u$  in  $\mathcal{A}$ .
10:    Add the path layer (the unique  $(m^*, m^*)$ -path) from  $u$  to every vertex in  $\mathcal{C}_u$  into the current tree
    $T$ .
11:    (Prune available vertices) Let  $\mathcal{A}_0$  be a copy of  $\mathcal{A}$ . Prune the set  $\mathcal{A}$  by removing every vertex that
   is reachable from  $u$  by a path of length at most  $2m^*$  contained in  $\mathcal{A}_0$ .
12:    Set  $s \leftarrow s + |\mathcal{C}_u|$ , and update  $\mathcal{L}$  as  $\mathcal{L}.\text{push}(\mathcal{C}_u)$ 
13:  end while
14:  if  $s \geq 2\ell$  then
15:    Grow the right tree rooted at  $u'_0$  by following Lines 6 to 13 with  $u_0$  replaced by  $u'_0$ .
16:    If the right tree also reaches a size of at least  $2\ell$ , then let  $T$  be the resulting two-sided tree, and
   set  $\mathcal{T} \leftarrow \mathcal{T} \cup \{T\}$ .
17:  end if
18: end for
19: Return  $\mathcal{T}$ .
```

During the process of constructing trees, one must take care to ensure that any blue edges remain independent Bern(λ/n) random variables. Line 11 ensures that whenever the neighborhood of a hub u is explored, that neighborhood will not be accessible to exploration from future hubs, and thus ensures the required independence of blue edges in future explorations. A difficulty arises when u, v are hub nodes as in Figure 3. In the situation described therein, the blue neighbors of v would be removed by Line 11, and thus the required independence for exploration from v would be lost. A similar issue arises when there are at least two paths of length $2m^*$ from u to v (even if only one is an (m^*, m^*) -path). Therefore, to avoid this issue, in Line 9, we only consider “non-shortcutted” paths, meaning that any other path connecting the same endpoints must be strictly longer. In other words, a path is “shortcutted” if there exists another path between its two endpoints that is of the same or shorter length.

Some restraint is required when building the trees; it turns out that we cannot simply take the number of hub nodes ℓ to be linear, in the hopes of building cycles from a few trees. The value of ℓ is limited due to the pruning step in Line 11. In particular, taking ℓ to be linear (or even super-constant) would remove too many vertices for every tree, leading to few trees being constructed. In other words, it is better to create many small trees rather than a few large trees.

3.2 Analysis of tree construction

In this subsection, we prove that Algorithm 1 successfully builds $\Theta(n)$ many two-sided trees, conditioned on H^* not containing too many cycles.

We first record some bounds on the moments of the number of (a, b) -paths, beginning with the expected number of shortcutted paths. These moment bounds will be used to show that the two-sided tree construction can be coupled to a supercritical branching process.

Lemma 3.5. *Conditional on $H^* = h$ where h has at most $(\log_2 n)^2$ cycles, given a fixed vertex v_0 and constants $a, b \geq 1$ the expected number of shortcutted (a, b) -paths starting from v_0 is $O(\log^2(n)/n)$.*

Proof. Given an (a, b) -path $v_0, \dots, v_{a+b}$ and another path $v'_0 = v_0, \dots, v'_\ell = v_{a+b}$ with $\ell \leq a + b$, choose the smallest i




Figure 3: The solid purple line represents a path of length strictly less than $2m^*$.

such that $v'_i \neq v_i$ and the smallest $j \geq i$ such that $v'_j \in \{v_0, \dots, v_{a+b}\}$. Then $v'_{i-1}, \dots, v'_j$ is a path of length at most $a+b$ such that $v'_{i-1}, v'_j \in \{v_0, \dots, v_{a+b}\}$ and this path does not share any other vertices or edges with $v_0, \dots, v_{a+b}$. So, it suffices to bound the expected number of (a, b) -paths that have such a side path.

Let G' be the graph consisting of all the vertices and edges in the paths $v_0, \dots, v_{a+b}$ and $v'_{i-1}, \dots, v'_j$, where the vertex labels are to be specified later. Note that there are at most $(a+b+1)^3$ possible choices of i, j , and the value of j' for which $v'_j = v_{j'}$, and at most 2^{2a+2b} options for which edges in G' are planted. Given i, j, j' , and the choice for the set of planted edges in G' , let G'_r denote the subgraph of G' formed by deleting all the unplanted edges. We call a connected component of G'_r (which could be a single vertex) a *planted component* of G' ; let c be the number of planted components of G' . To determine G' , it remains to label the vertices in all its planted components. Because G' consists of a path with a side path, there is at most one cyclic planted component, and all the other planted components must be paths. We consider the following two cases separately, depending on whether there exists a cyclic planted component.

Case 1: There is a cyclic planted component $\mathcal{C}$. In this case, $\mathcal{C}$ must be a cycle in H^* . Since H^* contains at most $(\log_2 n)^2$ cycles, there are at most $(\log_2 n)^2$ different choices for $\mathcal{C}$ and at most $2(2a+2b)$ options for which of $\mathcal{C}$'s vertices are v_{i-1} and $v_{j'}$. Each other planted component must be a path contained in a cycle of H^* . The length of the path has been pre-determined when we specify which edges in G' are planted. Thus, to determine the planted path component, it remains to specify the starting vertex of the path and which direction to traverse in a cycle of H^* . Therefore, each planted component other than $\mathcal{C}$ has at most $2n$ choices. Moreover, note that v_0 is not contained in $\mathcal{C}$, because (v_0, v_1) is required to be unplanted by the definition of an (a, b) -path. Therefore, the planted component containing v_0 is a singleton. Since v_0 has been pre-fixed, there are at most 2 different choices for the planted component containing v_0 . Let m_r and m_b denote the number of planted and unplanted edges in G' , respectively. Recalling that c denotes the number of planted components in G' , we have

$$m_r + m_b = |V(G')| = |V(G'_r)| = m_r + c - 1,$$

where the first equality holds because G' contains a single cycle, and the last equality holds because all the connected components of G'_r are paths except for one being a cycle. It follows that $m_b = c - 1$.

Case 2: All planted components are paths. Since v_0 has been pre-fixed, there are at most 2 different choices for the planted component containing v_0 . All the other planted components have at most $2n$ choices. Moreover, we have

$$m_r + m_b = |V(G')| = |V(G'_r)| = m_r + c.$$

It follows that $m_b = c$.

Since every unplanted edge exists in the observed graph with probability λ/n , combining the above two cases, we deduce that the expected number of (a, b) -paths that have such a side path is at most

$$(a+b+1)^3 \times 2^{2a+2b} \times \sum_{c=1}^{2a+2b} \left(2(2a+2b)(\log_2 n)^2 \times (2n)^{c-2} \times \left(\frac{\lambda}{n}\right)^{c-1} + 2 \times (2n)^{c-1} \left(\frac{\lambda}{n}\right)^c \right) = O(\log^2(n)/n).$$

□

The following lemma gives an almost matching lower bound to Lemma 2.2, even when a sufficiently small fraction of vertices is excluded and we are restricting to non-shortcutted paths.

Lemma 3.6. *Let $G \sim \mathcal{G}(n, \lambda, \delta)$ and H^* be the hidden planted 2-factor on the set of δn vertices. Let B be a subset of the vertices and $v \in B^c$ that are chosen in such a way that $|B| \leq \zeta n$ for some ζ , and conditional on H^*, B, v , the unplanted edges between vertices not in B are still i.i.d. $\text{Bern}(\lambda/n)$. Consider the subgraph of G induced by B^c , denoted by $G[B^c]$. Then conditional on B, v , and $H^* = h$ where h has at most $(\log_2 n)^2$ cycles, the expected number of non-shortcutted (a, b) -paths in $G[B^c]$ starting from v is at least $\left(1 - \frac{a(a+1)\zeta}{\delta} - \frac{(b-1)\zeta}{1-\delta} - o(1)\right) c_{a,b}$ for any fixed constants $a, b \geq 1$, where $\frac{b-1}{1-\delta} = 0$ when $b = 1$ and $\delta = 1$.*

Proof. In the following proof, we condition on $H^* = h$ where h has at most $(\log_2 n)^2$ cycles. Given $\vec{a}, \vec{b} \in \mathbb{N}_+^k$, let $\mathcal{W}_{\vec{a}, \vec{b}}^*(v)$ denote the set of $(\vec{a}, \vec{b})$ -paths starting from $v \notin B$ in the complete graph K_n that do not go through any vertex in B (where edges in the complete graph are red if they are contained in H^* , and they are blue otherwise). We will establish that

$$|\mathcal{W}_{\vec{a}, \vec{b}}^*(v)| \geq \left(1 - \frac{(a+1)\zeta}{\delta} - o(1)\right)^k \left(1 - \frac{\zeta}{1-\delta} - o(1)\right)^{b-k} \cdot (2\delta n)^k ((1-\delta)n)^{b-k}. \quad (3.7)$$

First, we enumerate sets of directed paths in H^* which do not intersect with each other or B . For any vertex $v'' \in H^*$ and any ℓ , there exist two distinct paths starting at v'' and consisting of ℓ planted edges, unless v'' is contained in a cycle of length at most ℓ . Since there are at most $(\log_2 n)^2$ cycles in H^* , the total number of length- ℓ directed planted paths in H^* is at least $2(\delta n - \ell(\log_2 n)^2)$. Also, given a path of length ℓ' in H^* , there are at most $2(\ell' + \ell + 1)$ directed paths of length ℓ in H^* that intersect it, and there are at most $2(\ell + 1)$ directed paths of length ℓ in H^* that go through any given point. So, for any positive $(a_1, \dots, a_k)$ with sum a , any $k' \leq k$ and any choice of paths of lengths $a_1, a_2, \dots, a_{k'-1}$ in H^* , the number of directed paths of length $a_{k'}$ in H^* that do not intersect v , any of the first $k' - 1$ paths, or any vertex in B is at least

$$\begin{aligned} & 2\delta n - 2a_{k'}(\log_2 n)^2 - (|B| + 1) \cdot 2(a_{k'} + 1) - \sum_{i=1}^{k'-1} 2(a_{k'} + a_i + 1) \\ & \geq 2\delta n - 2a_{k'}(\log_2 n)^2 - 2(k' + \zeta n)(a_{k'} + 1) - 2 \sum_{i=1}^{k'-1} a_i \\ & \geq 2\delta n - 2a(\log_2 n)^2 - 2(k + \zeta n)(a + 1) - 2a. \end{aligned}$$

Therefore, the number of ways to select directed paths of lengths $a_1, \dots, a_k$ in H^* , none of which intersect or contain v , is at least

$$(2\delta n - 2a(\log_2 n)^2 - 2(k + \zeta n)(a + 1) - 2a)^k = (1 - (a + 1)\zeta/\delta - o(1))^k \cdot (2\delta n)^k.$$

Given such a set of paths and any $(b_1, \dots, b_k)$ with sum b , there are at least

$$\prod_{i=0}^{b-k-1} (n - \delta n - \zeta n - i) = \left(1 - \frac{\zeta}{1-\delta} - o(1)\right)^{b-k} ((1-\delta)n)^{b-k}$$

ways to choose an ordered set of $b - k$ vertices not in $H^* \cup B$. Putting $b_i - 1$ of these vertices between the i th and $(i + 1)$ th planted paths for each i always yields an $((a_1, \dots, a_k), (b_1, \dots, b_k))$ -path starting from v in the complete graph, thus establishing (3.7).

Each of these paths appears in G with probability $(\lambda/n)^b$. For any given $a \geq 1$, there are $\binom{a-1}{k-1}$ distinct sequences of positive integers $(a_1, \dots, a_k)$ which sum to a , and similarly for b . Thus, the expected number of (a, b) -paths from v to v' in G avoiding B is at least

$$\begin{aligned} & \sum_{k=1}^{\infty} \sum_{\vec{a}, \vec{b} \in \mathbb{N}_+^k : \|\vec{a}\|_1 = a, \|\vec{b}\|_1 = b} |\mathcal{W}_{\vec{a}, \vec{b}}^*(v)| \cdot (\lambda/n)^b \\ & \geq \left(1 - \frac{(a+1)\zeta}{\delta} - o(1)\right)^a \left(1 - \frac{\zeta}{1-\delta} - o(1)\right)^{b-1} \sum_{k=1}^{\infty} \binom{a-1}{k-1} \binom{b-1}{k-1} \cdot (2\delta/(1-\delta))^k (\lambda(1-\delta))^b \\ & \geq \left(1 - \frac{a(a+1)\zeta}{\delta} - \frac{(b-1)\zeta}{1-\delta} - o(1)\right) c_{a,b}. \end{aligned}$$

Finally, we apply Lemma 3.5 to upper-bound the expected number of shortcutted (a, b) -paths in $G[B^c]$. Let G' be G with all vertex pairs $\{(u, u') \notin H^* : u \in B \text{ or } u' \in B\}$ resampled according to i.i.d. $\text{Bern}(\lambda/n)$. By Lemma 3.5, the expected number of shortcutted (a, b) -paths starting from v in G' is $O(\log^2(n)/n)$. Since $G[B^c] = G'[B^c] \subset G'$, the

expected number of shortcutted (a, b) -paths starting from v in $G[B^c]$ is $O(\log^2(n)/n)$. As a consequence, the expected number of non-shortcutted (a, b) -paths in $G[B^c]$ is at least

$$\left(1 - \frac{a(a+1)\zeta}{\delta} - \frac{(b-1)\zeta}{1-\delta} - o(1)\right) c_{a,b} - O(\log^2(n)/n) = \left(1 - \frac{a(a+1)\zeta}{\delta} - \frac{(b-1)\zeta}{1-\delta} - o(1)\right) c_{a,b}.$$

□

The following lemma gives an upper bound on the variance of the number of (a, b) -paths, which will be used in the tree construction to prove the impossibility of almost exact recovery.

Lemma 3.7. *Fix $a, b \in \mathbb{Z}_+$ and a vertex $v \in V(H^*)$. Let $X_{a,b}(v)$ denote the number of (a, b) -paths starting from v in G . Then $\mathbb{E}[X_{a,b}^2 | H^*] \leq C$ for some constant C that depends only on λ, a , and b , and is independent of n .*

Proof. Let $\mathcal{W}_{a,b}(v)$ denote the set of (a, b) -paths that start from $v \in V(H^*)$ in the complete graph K_n (where edges in the complete graph are red if they are contained in H^* , and they are blue otherwise). Then

$$X_{a,b} = \sum_{P \in \mathcal{W}_{a,b}(v)} \mathbf{1}_{\{P \subset G\}}.$$

By definition,

$$\mathbb{E}[X_{a,b}^2 | H^*] = \sum_{P, P' \in \mathcal{W}_{a,b}(v)} \mathbb{P}\{P \subset G, P' \subset G | H^*\} = \sum_{P, P' \in \mathcal{W}_{a,b}(v)} \left(\frac{\lambda}{n}\right)^{2b - |P_b \cap P'_b|},$$

where P_b denotes the blue (unplanted) segments of path P . Fixing a path $P \in \mathcal{W}_{a,b}(v)$, the number of paths $P' \in \mathcal{W}_{a,b}$ for which $|P_b \cap P'_b| = e$ is at most

$$\sum_{k > 0, (a_1, \dots, a_k), (b_1, \dots, b_k) : \sum a_i = a, \sum b_i = b} \binom{b}{e} \binom{b}{e} 2^e n^{b-e} 2^k,$$

where the first $\binom{b}{e}$ counts the different choices of e overlapping blue edges from P ; the second $\binom{b}{e}$ counts the different locations to put those e overlapping blue edges when we list the blue edges of P' from left to right starting from vertex v ; 2^e accounts for the different directions to put those e overlapping blue edges; n^{b-e} counts the vertex labels of the right endpoints of the additional blue edges of P' after fixing the e overlapping edges; finally, to further determine the red edges and the remaining vertex labels, we just need to follow the sequences $(a_1, \dots, a_k)$ and $(b_1, \dots, b_k)$, traverse a cycle of the 2-factor H^* for a direction, insert the red edges accordingly, and there are at most 2^k different directions to choose from.

Moreover, recall from (2.2) that $|\mathcal{W}_{a,b}(v)| \leq C'_{a,b} n^b$ for some constant C' that only depends on a, b . It follows that

$$\mathbb{E}[X_{a,b}^2 | H^*] \leq C''_{a,b} \sum_{e=0}^b n^b n^{b-e} \left(\frac{\lambda}{n}\right)^{2b-e} = C_{a,b,\lambda},$$

where $C_{a,b,\lambda}$ is some constant that only depends on a, b, λ . □

We now demonstrate that sufficiently many trees are constructed. As a first step, we establish a lower bound on the number of available vertices throughout the execution of Algorithm 1.

Lemma 3.8. *Let $t > 0$, $\gamma \in (0, 1)$, and $\ell \in \mathbb{N}$ be constants. Conditioned on H^* , throughout the first t iterations of tree construction in Algorithm 1, the number of available vertices satisfies $|\mathcal{A}| \geq n - 2\gamma n - 6\ell(2\lambda + 4)^{2m^*} t$ with probability $1 - e^{-\Theta(t)}$.*

Proof. Recall that initially $|\mathcal{A}| = (1 - 2\gamma)n$. Let $\mathcal{A}_i$ be the set of available vertices after the i^{th} pruning stage when Line 11 executes, which happens whenever we find the (m^*, m^*) -paths from a given hub vertex. Here $i \in \{0, \dots, 4\ell t\}$, since each tree has two sides and each side executes at most 2ℓ pruning stages. We let $\mathcal{A}_0$ contain all vertices apart from the reserved vertices, and let $\mathcal{A}_i = \mathcal{A}_T$ for all $i > T$ where T is the (random) number of pruning stages. Let X_i be the number of vertices removed during the i^{th} pruning stage, so that $X_i = |\mathcal{A}_{i-1}| - |\mathcal{A}_i|$. In order to bound the X_i 's, we introduce a sequence of independent branching processes B_i with offspring distribution $2 + \text{Binom}(n, p \triangleq \lambda/n)$. Let Y_i be the number of nodes among the first $2m^* + 1$ levels of the i^{th} branching process B_i . We claim that X_i is stochastically dominated by Y_i .

To prove the claim, we condition on the realization of H^* . If (u, v) is not an edge in H^* , then we call (u, v) a potential blue edge. We call a potential blue edge (u, v) “unseen” after the i^{th} pruning stage if at least one of its

endpoints is available; that is, $|\{u, v\} \cap \mathcal{A}_i| \geq 1$. Let U_i be the set of unseen potential blue edges after the i^{th} pruning stage. Equivalently, U_i contains all potential blue edges outside of the graph induced by $[n] \setminus \mathcal{A}_i$. Then conditioned on the graph induced by $[n] \setminus \mathcal{A}_i$, the edges within U_i exist independently with probability λ/n . For each $i \geq 1$, X_i is equal to the total number of nodes in the $2m^*$ -neighborhood of a hub vertex u at stage i within $\mathcal{A}_{i-1}$, which is completely determined by U_{i-1} . Thus conditional on the graph induced by $[n] \setminus \mathcal{A}_{i-1}$, the $2m^*$ -neighborhood of u within $\mathcal{A}_{i-1}$ is stochastically dominated by the i^{th} branching process B_i .

Having established the branching process domination, it suffices to bound $Y \triangleq \sum_{i=1}^{4\ell t} Y_i$, which we accomplish with a Chernoff bound strategy. For any $r, s > 0$, the Markov inequality yields

$$\mathbb{P}(Y \geq s) = \mathbb{P}\left(e^{rY} \geq e^{rs}\right) \leq e^{-rs} \mathbb{E}\left[e^{rY}\right] = e^{-rs} \left(\mathbb{E}\left[e^{rY_1}\right]\right)^{4\ell t}. \quad (3.8)$$

To compute the MGF of Y_1 , let Z_j be the total progeny of a depth- j branching process with offspring distribution $2 + \text{Binom}(n, p)$. We claim that Z_{j+1} is equal in distribution to

$$1 + \sum_{i=1}^W V_{j,i}, \quad V_{j,i} \stackrel{\text{i.i.d.}}{\sim} Z_j, \quad W \sim 2 + \text{Binom}(n, p),$$

where W is independent of the $V_{j,i}$'s. To see this, observe that we can find the number of offspring through level $j+1$ by adding up the number of descendants through level j relative to each of the children of the root, and then adding 1 for the root itself. Each offspring of the root has a number of descendants through level j (including itself) with distribution Z_j . Thus, the MGF of Z_{j+1} satisfies the following recurrence relation:

$$\begin{aligned} \mathbb{E}\left[e^{rZ_{j+1}}\right] &= \mathbb{E}\left[e^{r(1+\sum_{i=1}^W V_{j,i})}\right] \\ &= e^r \mathbb{E}\left[\left(\mathbb{E}\left[e^{rZ_j}\right]\right)^W\right] \\ &= e^r \left(\mathbb{E}\left[e^{rZ_j}\right]\right)^2 \left(1 - p + p \cdot \mathbb{E}\left[e^{rZ_j}\right]\right)^n \\ &\leq e^r \left(\mathbb{E}\left[e^{rZ_j}\right]\right)^2 \left(\exp\left(-p + p\mathbb{E}\left[e^{rZ_j}\right]\right)\right)^n \\ &= \left(\mathbb{E}\left[e^{rZ_j}\right]\right)^2 \exp\left(r - \lambda + \lambda\mathbb{E}\left[e^{rZ_j}\right]\right). \end{aligned} \quad (3.9)$$

Fix any $\beta > 0$ such that $\alpha \triangleq \frac{e^\beta - 1}{\beta} \geq \frac{3}{2}$. We prove via induction that

$$\mathbb{E}\left[e^{rZ_j}\right] \leq \exp\left([\alpha(\lambda + 2)]^j r\right), \quad \forall j \in [2m^*], \forall 0 \leq r \leq \beta/[\alpha(\lambda + 2)]^{2m^*}. \quad (3.11)$$

Since $Z_0 = 1$, the base case holds trivially. We next assume the inductive hypothesis (3.11) for some $j \in \{1, 2, \dots, 2m^* - 1\}$. Using the recurrence relation (3.10), we have for all $0 \leq r \leq \beta/[\alpha(\lambda + 2)]^{2m^*}$,

$$\begin{aligned} \mathbb{E}\left[e^{rZ_{j+1}}\right] &\leq \exp\left\{2[\alpha(\lambda + 2)]^j r + r - \lambda + \lambda e^{[\alpha(\lambda + 2)]^j r}\right\} \\ &\leq \exp\left\{2[\alpha(\lambda + 2)]^j r + r - \lambda + \lambda \left(1 + \alpha[\alpha(\lambda + 2)]^j r\right)\right\} \\ &= \exp\left\{[\alpha(\lambda + 2)]^j r (2 + \lambda\alpha) + r\right\}, \end{aligned} \quad (3.12)$$

where the first inequality follows by plugging the inductive hypothesis for j , and the second inequality holds due to $e^r \leq 1 + \alpha r$ for all $0 \leq r \leq \beta$ and the fact that $0 \leq [\alpha(\lambda + 2)]^j r \leq \beta$.

Recalling that $\alpha \geq \frac{3}{2}$, we have

$$2r(1 - \alpha)[\alpha(\lambda + 2)]^j + r \leq 2r(1 - \alpha) + r = r(3 - 2\alpha) \leq 0,$$

so that

$$2r(1 - \alpha)[\alpha(\lambda + 2)]^j + r \leq 0.$$

Rearranging, the above inequality is equivalent to

$$[\alpha(\lambda + 2)]^j r (2 + \lambda\alpha) + r \leq [\alpha(\lambda + 2)]^{j+1} r. \quad (3.13)$$

Substituting (3.13) into (3.12), we obtain

$$\mathbb{E}\left[e^{rZ_{j+1}}\right] \leq \exp\left\{[\alpha(\lambda + 2)]^{j+1} r\right\},$$

thus establishing (3.11).

Having established (3.11) for all $j \in \{1, 2, \dots, 2m^*\}$, we get that

$$\mathbb{E} \left[e^{rY_1} \right] = \mathbb{E} \left[e^{rZ_{2m^*}} \right] \leq \exp \left([\alpha(\lambda + 2)]^{2m^*} r \right), \quad \forall 0 \leq r \leq \beta/[\alpha(\lambda + 2)]^{2m^*}.$$

Substituting into (3.8), we get that

$$\begin{aligned} \mathbb{P} \left(Y \geq 6\ell t [\alpha(\lambda + 2)]^{2m^*} \right) &\leq \exp \left(-6r\ell t [\alpha(\lambda + 2)]^{2m^*} + 4\ell t [\alpha(\lambda + 2)]^{2m^*} r \right) \\ &= \exp \left(-2r\ell t [\alpha(\lambda + 2)]^{2m^*} \right) = \exp(-2\ell t \beta), \end{aligned}$$

where the last equality holds by choosing $r = \beta/[\alpha(\lambda + 2)]^{2m^*}$. Finally, we arrive at the desired conclusion by picking $\beta = 1$ and noting that $\alpha \leq 2$. $\square$

Next, we will show that Algorithm 1 results in enough trees for the purposes of creating balanced cycles.

Lemma 3.9. *Let $\delta, \lambda > 0$ satisfy $\lambda > \frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$ and condition on H^* containing at most $(\log_2 n)^2$ cycles. Let m^* be the value from Lemma 3.4 with $c_{m^*, m^*} > 1$, where $m^* = 1$ when $\delta = 1$. Let $\gamma > 0$ be a small constant such that*

$$\zeta \triangleq 2\gamma + 6\gamma(2\lambda + 4)^{2m^*} < \frac{c_{m^*, m^*} - 1}{2c_{m^*, m^*}} \cdot \frac{1}{m^*(m^* + 1)/\delta + (m^* - 1)/(1 - \delta)}. \quad (3.14)$$

Then for all $\ell \in \mathbb{N}$, with probability $1 - e^{-\Omega(n)}$, Algorithm 1 results in at least $K_1 = C'K$ two-sided trees, where $K = \frac{\gamma n}{\ell}$ as in Line 2, and $C' > 0$ is a function of δ and λ . Specifically, $C' = p^2/2$ with $p = \frac{\mu^2 - \mu}{\mu^2 - \mu + C + 1/4}$, where $\mu = (c_{m^, m^*} + 1)/2$ and C is the constant from Lemma 3.7.*

Proof. Algorithm 1 iterates through $t = 1, \dots, K$; at iteration t , it successfully outputs a two-sided tree if both sides of the tree grow to contain at least 2ℓ hub nodes. Ideally, we would like to argue that at iteration t , conditional on the history of the algorithm, the probability of success is lower-bounded by some constant, hence a constant proportion of the K iterations succeed with high probability. However, this argument does not go through given any history of the algorithm. In particular, if the algorithm explores too many vertices before the t^{th} iteration, there may not be enough remaining vertices to grow large enough trees. To resolve this, we use Proposition 3.8 to argue that with high probability, there will be enough vertices available for the algorithm to explore throughout the entire construction. Below is the precise argument.

Let E be the event that throughout all $K = \gamma n/\ell$ iterations of Algorithm 1, the number of available vertices is at least $(1 - \zeta)n$. Since by assumption $1 - \zeta > 1 - \frac{\delta}{2}$, the event E implies that there exists at least one planted edge among the available vertices and hence the FAIL condition in Line 3 is not triggered. By Proposition 3.8, $\mathbb{P}(E) = 1 - e^{-\Theta(n)}$.

Let S_{2t-1} (resp. S_{2t}) denote the events that the left (resp. right) side of the t^{th} two-sided tree grows to contain at least 2ℓ hub nodes. We couple the growth on each side of this two-sided tree with an auxiliary tree construction process which mimics the true one until the event E fails. More precisely, whenever we are about to add the offspring to a vertex in the tree, we check whether the number of available vertices is at least $(1 - \zeta)n$. If so, we directly copy the number of offspring to the auxiliary process. If not, the number of offspring added to the auxiliary process is instead sampled from any distribution supported on $\mathbb{N}$ with mean μ and variance at most C , where $\mu = (c_{m^*, m^*} + 1)/2$ and C is the constant from Lemma 3.7. Note that μ and C only depend on δ and λ .

Let $\tilde{S}_{2t-1}$ and $\tilde{S}_{2t}$ denote the corresponding events for the coupling process. We have

$$\begin{aligned} \mathbb{P} \left\{ \sum_{t \leq K} \mathbb{1}\{S_{2t-1} \cap S_{2t}\} < K_1 \right\} &\leq \mathbb{P} \left(\left\{ \sum_{t \leq K} \mathbb{1}\{\tilde{S}_{2t-1} \cap \tilde{S}_{2t}\} < K_1 \right\} \cap E \right) + \mathbb{P}(E^c) \\ &\leq \mathbb{P} \left(\sum_{t \leq K} \mathbb{1}\{\tilde{S}_{2t-1} \cap \tilde{S}_{2t}\} < K_1 \right) + \mathbb{P}(E^c). \end{aligned} \quad (3.15)$$

The first inequality holds because on event E , $S = \tilde{S}$. Next, we lower-bound the survival probability of the auxiliary trees given the history of the algorithm. For simplicity, let us first focus on the growth of the left tree. Let $\mathcal{F}_{2t-2}$ denote the history of the auxiliary trees up to iteration t . We will show that $\mathbb{P}(\tilde{S}_{2t-1} | \mathcal{F}_{2t-2}) \geq p$, where $p = \frac{\mu^2 - \mu}{\mu^2 - \mu + C + 1/4}$.

To see this, denote Z_i as the number of offspring of the i^{th} node in breadth-first order in the left auxiliary tree. Let $\mathcal{F}_{2t-2, i-1}$ be the union of $\mathcal{F}_{2t-2}$ and the history of the auxiliary process up to before Z_i is generated. Let E_i be the event that there are at least cn available vertices right before Z_i is generated. Then

$$\mathbb{E}(Z_i | \mathcal{F}_{2t-2, i-1}) \geq \min \{ \mathbb{E}(Z_i | \mathcal{F}_{2t-2, i-1}, E_i), \mathbb{E}(Z_i | \mathcal{F}_{2t-2, i-1}, E_i^c) \}.$$

On E_i , the growth of the auxiliary tree and the original tree coincide exactly up to (and including) Z_i . By Lemma 3.6, we have that conditional on $\mathcal{F}_{2t-2, i-1}$ and E_i , the expected number of non-shortcutted (m^*, m^*) -paths that originate from the i^{th} hub node and only contain available vertices is at least

$$\left(1 - \frac{m^*(m^* + 1)\zeta}{\delta} - \frac{(m^* - 1)\zeta}{1 - \delta} - o(1)\right) c_{m^*, m^*}.$$

By the choice of ζ as per (3.14), we have

$$\left(1 - \frac{m^*(m^* + 1)\zeta}{\delta} - \frac{(m^* - 1)\zeta}{1 - \delta}\right) c_{m^*, m^*} > \left[1 - \frac{1}{2} \cdot \frac{c_{m^*, m^*} - 1}{c_{m^*, m^*}}\right] c_{m^*, m^*} = (c_{m^*, m^*} + 1)/2 = \mu,$$

which shows that $\mathbb{E}(Z_i | \mathcal{F}_{2t-2, i-1}, E_i) \geq \mu$ for n sufficiently large. On the other hand, $\mathbb{E}(Z_i | \mathcal{F}_{2t-2, i-1}, E_i^c) \geq \mu$ from the definition of the auxiliary process. Thus $\mathbb{E}(Z_i | \mathcal{F}_{2t-2, i-1}) \geq \mu$ for all i .

Similarly, we can apply Lemma 3.7 to bound the second moment as follows.

$$\mathbb{E}((Z_i)^2 | \mathcal{F}_{2t-2, i-1}) \leq \max\{\mathbb{E}((Z_i)^2 | \mathcal{F}_{2t-2, i-1}, E_i), \mathbb{E}((Z_i)^2 | \mathcal{F}_{2t-2, i-1}, E_i^c)\} \leq C.$$

The first and second moment control allows us to apply Lemma A.1, which lower-bounds the survival probability of such a history-dependent branching process. The lemma ensures that

$$\mathbb{P}\{\tilde{S}_{2t-1} | \mathcal{F}_{2t-2}\} \geq p = \frac{\mu^2 - \mu}{\mu^2 - \mu + C + 1/4}.$$

The same argument applies to the growth of the right auxiliary tree, so that $\mathbb{P}\{\tilde{S}_{2t} | \mathcal{F}_{2t-1}\} \geq p$, where $\mathcal{F}_{2t-1}$ is the union of $\mathcal{F}_{2t-2}$ and growth of the left auxiliary tree (at iteration t). Thus

$$\mathbb{P}\{\tilde{S}_{2t-1} \cap \tilde{S}_{2t} | \mathcal{F}_{2t-2}\} = \mathbb{P}\{\tilde{S}_{2t-1} | \mathcal{F}_{2t-2}\} \mathbb{P}\{\tilde{S}_{2t} | \mathcal{F}_{2t-1}\} \geq p^2. \quad (3.16)$$

From here, we can bound $\mathbb{P}(\sum_{i \leq K} \mathbb{1}\{\tilde{S}_{2t-1}\} \mathbb{1}\{\tilde{S}_{2t}\} < K_1)$ by coupling $\{\mathbb{1}\{\tilde{S}_{2t-1} \cap \tilde{S}_{2t}\}\}_{t \leq K}$ with a sequence $\{X_t\}_{t \leq K}$ of i.i.d. $\text{Bern}(p^2)$ random variables. In more detail, when $\mathbb{1}\{\tilde{S}_{2t-1} \cap \tilde{S}_{2t}\} = 0$, let $X_t = 0$; when $\mathbb{1}\{\tilde{S}_{2t-1} \cap \tilde{S}_{2t}\} = 1$, draw X_t from the Bernoulli distribution with success probability $p^2 / \mathbb{P}(\tilde{S}_{2t-1} \cap \tilde{S}_{2t} | \mathcal{F}_{2t})$ (Note that by (3.16) this is a valid probability). By construction, we have $X_t \leq \mathbb{1}\{\tilde{S}_{2t-1} \cap \tilde{S}_{2t}\}$ for all t , and

$$\mathbb{P}(X_t = 1 | X_1, \dots, X_{t-1}) = \mathbb{P}(\tilde{S}_{2t-1} \cap \tilde{S}_{2t} | \mathcal{F}_{2t-2}) \cdot \frac{p^2}{\mathbb{P}(\tilde{S}_{2t-1} \cap \tilde{S}_{2t} | \mathcal{F}_{2t-2})} = p^2.$$

Thus the X_t 's are distributed i.i.d. $\text{Bern}(p^2)$, and

$$\mathbb{P}\left(\sum_{t \leq K} \mathbb{1}\{\tilde{S}_{2t-1} \cap \tilde{S}_{2t}\} < K_1\right) \leq \mathbb{P}\left(\sum_{t \leq K} X_t < K_1\right) = e^{-\Omega(n)}$$

by choosing $K_1 = C'K$ with constant $C' = p^2/2$ only depending on δ and λ . Combined with (3.15), and the fact that $\mathbb{P}(E^c) = e^{-\Theta(n)}$, we have shown that

$$\mathbb{P}\left\{\sum_{t \leq K} \mathbb{1}\{S_{2t-1} \cap S_{2t}\} < K_1\right\} = e^{-\Theta(n)}.$$

In other words, with probability $1 - e^{-\Omega(n)}$, Algorithm 1 yields at least K_1 two-sided trees. □

3.3 Cycle construction via sprinkling

We now give the details of the preprocessing procedure, which reserves edges and creates the set of available vertices. The reserved edges will be used for the purposes of “sprinkling” to connect trees into cycles.

Algorithm 2 Reserve Edges

Input: Graph H^* on n vertices, $\gamma \in (0, \delta/5]$

Output: A set E^* of γn vertex-disjoint red edges of G , available vertices $\mathcal{A}$

- 1: Let $E^* = \emptyset$, $S = E(H^*)$.
 - 2: **for** $i \in \{1, 2, \dots, \gamma n\}$ **do**
 - 3: Choose an arbitrary edge e from S and add it to E^* .
 - 4: Remove e and all edges at distance at most 2 from S (that is, edges which have an endpoint that is connected by an edge to an endpoint of e).
 - 5: **end for**
 - 6: Let V_1 be the set of endpoints of edges in E^* . Let $\mathcal{A} = [n] \setminus V_1$ be the set of *available* vertices.
-

Since Step 4 of Algorithm 2 removes at most five edges, the final set E^* contains γn edges in total, provided that $\gamma \leq \delta/5$. Note that any set of parameters satisfying the conditions of Lemma 3.9 will always have $\gamma \leq \delta/5$, and will thus be suitable for this algorithm.

Algorithm 3 Cycle Construction

Input: Graph G on n vertices with red subgraph H^* on δn vertices, reserved edge parameter γ , path length parameter $m \in \mathbb{N}$, tree size parameter $\ell \in \mathbb{N}$, and degree parameter d

Output: A set of alternating cycles $\mathcal{C}$ on G

- 1: Apply Algorithm 2 to input (H^*, γ) , obtaining the set of γn reserved edges E^* and the set of $(1 - 2\gamma)n$ available vertices $\mathcal{A}$. We assume $|E^*|$ is even.
 - 2: Let $\mathcal{T} = (L_i, R_i)_{i=1}^{K_1}$ be the output of Algorithm 1 on input $(G, \mathcal{A}, m, \ell)$.
 - 3: Randomly partition E^* into two equally-sized sets (E_L^*, E_R^*) . For each $(u, v) \in E^*$ with $u < v$, designate u as the “tree-facing” vertex and designate v as the “linking” vertex.
 - 4: Initialize an empty (bipartite) graph $\overline{G}$.
 - 5: **for** $i \in [K_1]$ **do**
 - 6: **if** L_i is blue-connected to at least d unmarked tree-facing endpoints among E_L^* and the same is true for R_i with respect to E_R^* **then**
 - 7: Let the first d of these edges be denoted $\mathcal{E}(L_i) \subset E_L^*$ and $\mathcal{E}(R_i) \subset E_R^*$.
 - 8: Mark all edges among $\mathcal{E}(L_i) \cup \mathcal{E}(R_i)$.
 - 9: Include i as a vertex on both sides of $\overline{G}$, and connect them by a red edge.
 - 10: **end if**
 - 11: **end for**
 - 12: **for** $i \in [K_1]$ **do**
 - 13: **for** $j \in [K_1]$ **do**
 - 14: **if** both i and j are vertices in $\overline{G}$, and some linking endpoint in $\mathcal{E}(L_i)$ is connected to some linking endpoint in $\mathcal{E}(R_j)$ by a blue edge **then**
 - 15: Connect i and j by a blue edge in $\overline{G}$.
 - 16: **end if**
 - 17: **end for**
 - 18: **end for**
 - 19: Find the set of alternating cycles in $\overline{G}$, and return their preimage in G .
-

Algorithm 3 connects the trees into cycles via reserved edges, similarly to Algorithm 3 in [13]. At a high level, a cycle is formed by linking some right tree R_{i_1} to a left tree L_{i_2} , then linking R_{i_2} to L_{i_3} , and so on, back to L_{i_1} . In order to retain some independence for the purpose of linking, trees are connected to each other via a “five edge construction:” a hub node in R_i connects to a reserved red edge in E_R^* while a hub node in L_j connects to a red edge in E_L^* ; in turn, the red edges are connected by a blue edge (see Figure 4). By design, the endpoints of the reserved edges can only be connected by blue edges, due to Algorithm 2, Line 4. Since the five edges comprise three blue edges and two red edges, and each two-sided tree has a central red edge, the overall construction preserves the balancedness of blue and red edges.

We finally show that Algorithm 3 produces many long, balanced cycles; Lemma 3.3 then follows directly by observing that every balanced cycle of length t corresponds to a distinct 2-factor H on δn edges with $|H \Delta H^*| = t$.




Figure 4: Five edge construction (reproduced from [13]).

Lemma 3.10. *Let $0 < \delta \leq 1$, $\lambda > \frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$, and condition on H^* containing at most $(\log_2 n)^2$ cycles. Let $\mathcal{C}$ be the output of Algorithm 3 on input (G, γ, m, ℓ, d) , where $\gamma > 0$ satisfies the requirements in Lemma 3.9, the value $m = m^*$ is as defined in Lemma 3.4, $\ell = \frac{2^{14} \log(32e)}{\lambda^2 \gamma^2} \alpha$, and $d = \frac{2^{11} \log(32e)}{\lambda \gamma} \alpha$, where $\alpha \geq 1$ is a sufficiently large constant. Then there exist constants $c_1, c_2 > 0$ such that $\mathcal{C}$ contains at least $e^{c_1 n}$ balanced cycles of length at least $c_2 n$, with probability $1 - e^{-\Omega(n)}$.*

Lemma 3.10 requires the following result of [9].

Lemma 3.11. *[9, Lemma 7] Let G be a bi-colored bipartite graph on $[k] \times [k]'$ whose k red edges are defined by a perfect matching, and blue edges are generated from a bipartite Erdős-Rényi graph with edge probability $\frac{D}{k}$. If $k \geq 525$ and $D \geq 256 \log(32e)$, then with probability at least $1 - \exp(-\frac{Dk}{2^{14}})$, G contains $\exp(\frac{k}{20})$ distinct alternating cycles of length at least $\frac{3k}{4}$.*

Proof of Lemma 3.10. The proof is analogous to the proof of [13, Lemma C.6]. By Lemma 3.9, at least $K_1 = C'K = C' \frac{\gamma n}{\ell}$ two-sided trees $\{(L_i, R_i)\}_{i=1}^{K_1}$ are produced, with probability $1 - e^{-\Omega(n)}$. If L_i, R_j each connect to at least d edges among E_L^*, E_R^* respectively, then L_i and R_j connect in the five-edge construction with probability at least

$$1 - \left(1 - \frac{\lambda}{n}\right)^{d^2} \geq \frac{\lambda d^2}{2n}.$$

Let $c = \frac{512 \log(32e)}{\lambda d^2}$. We want to show that at least cn trees are such that both sides are connected to at least d tree-facing endpoints (and we call such trees “well-connected”). We use the fact that when identifying the tree-facing neighbors of a tree (L_i, R_i) with $i \in [K_1]$, there are at least $\frac{\gamma n}{2} - cnd$ tree-facing vertices that are not yet connected to a tree, as long as we have not yet found cn well-connected trees. By our choice of d and ℓ , we have

$$d = \frac{2^{11} \log(32e)}{\lambda \gamma} \alpha = \frac{\ell \lambda}{2} \cdot \frac{\gamma}{4} \leq \frac{\ell \lambda}{2} \left(\frac{\gamma}{2} - \frac{512 \log(32e)}{\lambda d} \right) = \frac{1}{2} \cdot \frac{\lambda}{n} \cdot \ell \cdot \left(\frac{\gamma n}{2} - cnd \right)$$

Thus, d is upper-bounded by half of the *expected* number of tree-facing neighbors of L_i or R_i , and the number of tree-facing neighbors of a given tree stochastically dominates the $\text{Bin}(\ell(\frac{\gamma n}{2} - cnd), \lambda/n)$ distribution, which in turn dominates the $\text{Bin}(2dn/\lambda, \lambda/n)$ distribution. Since the mean and median of any binomial random variable differ by at most 1, we have that the median of the $\text{Bin}(2dn/\lambda, \lambda/n)$ distribution is at least $2d - 1 \geq d$. Thus, the probability that L_i is connected to at least d tree-facing endpoints is at least $1/2$, and the same is true for R_i . It follows that a particular tree is connected on both sides to d tree-facing vertices with probability at least $(\frac{1}{2})^2$.

Due to independence across trees and the Chernoff bound, we see that for any constant $\eta > 0$, there are at least $(1 - \eta) \frac{K_1}{4}$ trees which are connected to at least d tree-facing endpoints on both sides, with probability $1 - e^{-\Omega(n)}$. Since both d and ℓ are linear in α , for sufficiently large constant α ,

$$cn = \frac{512 \log(32e)}{\lambda d^2} n \leq \frac{1 - \eta}{8} \left(\frac{\mu^2 - \mu}{\mu^2 - \mu + C + 1/4} \right)^2 \cdot \frac{\gamma n}{\ell} = (1 - \eta) \frac{K_1}{4},$$

where $\mu = (c_{m^*, m^*} + 1)/2$ and C is the constant from Lemma 3.7. It follows that with probability $1 - e^{-\Omega(n)}$, there are at least cn trees which are connected to at least d tree-facing endpoints on both sides.

Now we can couple to a bi-colored bipartite graph with cn vertices on each side, representing cn two-sided trees. The bipartite graph contains a perfect red matching, as well as independent blue edges which exist with probability $\frac{\lambda d^2}{2n}$. To apply Lemma 3.11, observe that the expected number of blue edges which are adjacent to a given vertex in the bipartite graph is $\frac{\lambda d^2}{2n} \cdot cn = 256 \log(32e)$. The conclusion follows from Lemma 3.11. $\square$

4 Computationally efficient estimators

In this section, we design a computationally efficient estimator that attains the information-theoretic threshold for almost exact recovery. The estimator is given in Algorithm 4. Instead of searching for all cycles (which is computationally intractable), the algorithm builds a subgraph from short trails, combining or trimming them in a greedy manner to make the output subgraph as large as possible.

Theorem 4.1. *Consider the planted cycles model $\mathcal{G}(n, \lambda, \delta)$ conditional on $H^* = h$, where $\delta \in (0, 1]$. If $\lambda < \frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$, then Algorithm 4 outputs a set H of edges such that $\mathbb{E}[|H \Delta H^*|] = o(n)$, achieving almost exact recovery. Moreover, Algorithm 4 runs in time $O(n^{3+\log(2+\lambda)})$ with high probability.*

We prove the computational and statistical guarantees of Algorithm 4 in the following Lemmas 4.2–4.6. In particular, Theorem 4.1 follows by combining Lemma 4.3, Lemma 4.4, and Lemma 4.6. By design, the algorithm returns a subgraph H which consists of cycles and possibly a small number of paths. In Lemma 4.2, we first generalize Lemma 2.1 to show that the difference graph $H \Delta H^*$ can be decomposed into a collection of edge-disjoint trails with the desired alternating properties at vertices in $V(H) \cap V(H^*)$. Building upon such a decomposition, in Lemma 4.3, we show that any such subgraph H attains almost exact recovery of the true planted cycles H^* whenever information-theoretically possible, as long as H is nearly as large as H^* and does not contain too many paths. We then prove in Lemma 4.4 that Algorithm 4 will indeed not output a large number of paths, and that through its greedy updates, the size of the output subgraph must be at least almost as large as H^* . Lemmas 4.3, 4.4 combined imply that Algorithm 4 attains the information-theoretic threshold for almost exact recovery. Finally, in Lemma 4.6, we prove that with high probability, Algorithm 4 runs in polynomial time.

Before stating and proving the lemmas, below is a brief description of Algorithm 4 and the intuition behind its design. From an empty subgraph H , the algorithm searches through the collection of short trails to find “good” candidates P to update H to $H \Delta P$ (i.e. the XOR operation), while making sure that the updated H remains a collection of vertex-disjoint cycles and paths. That is, no vertex has degree larger than 2 after the update. To clarify what qualifies a trail P to be “good,” recall that the algorithm aims to construct an H that 1) contains many edges, and 2) contains few paths (i.e. few vertices of degree 1). Accordingly, two types of trails are considered “good”:

1. “Cost-free” trails, which when XOR’ed onto H , strictly increases $|H|$ (number of edges) without increasing the number of degree-1 vertices.
2. “Cost-effective” trails, which when XOR’ed onto H , increase the number of degree-1 vertices by at most 2, but increase $|H|$ by at least $\sqrt{\log n}$.

In each iteration, the greedy algorithm invokes subroutines A and B to search for “good” trails of these two types. The algorithm terminates only when neither subroutine is able to find a suitable trail to update H .

First, we generalize Lemma 2.1 and prove that the difference graph $H \Delta H^*$ admits a trail decomposition such that each trail alternates between red (planted) and blue (unplanted) edges when it goes in and out of a vertex in $V(H) \cap V(H^*)$.

Lemma 4.2. *Let H^* be a planted 2-factor on m vertices, and let H be a graph where every vertex has degree at most 2 and there are $2m$ vertices of degree 1 in H . Then*

- *The difference graph $H^* \Delta H$ is a graph where every vertex has degree at most 4. Every degree-4 vertex has two incident red (planted) edges and two incident blue (unplanted) edges, and every degree-3 vertex has two incident red (planted) edges and one incident blue (unplanted) edge.*
- *Furthermore, $H^* \Delta H$ can be decomposed into a collection of edge-disjoint trails such that each trail must alternate between red and blue edges whenever it goes in and out of a vertex in $V(H^*) \cap V(H)$, and there are exactly m open trails whose endpoints must have degree 1 in H .*

Proof. The first part of the statement readily follows from the fact that every vertex in H^* has degree 2 and every vertex in H has degree at most 2. Moreover, a vertex has degree 1 or 3 in $H^* \Delta H$ if and only if it has degree 1 in H . Thus, there are $2m$ vertices of degree 1 or 3 in $H^* \Delta H$.

Algorithm 4 Approximate algorithm for finding a disjoint union of cycles

Input: A graph G

Output: A vertex-disjoint union of cycles and paths, H

```

1: Let  $n$  be the number of vertices in  $G$ .
2: Let  $H = \emptyset$ , boolean  $can\_grow = True$ .
3: Let  $S$  be the set of all trails (closed or open) in  $G$  of length less than  $\log n$ .
4: while  $can\_grow$  do
    // Subroutine A: search for "cost-free" candidate trails
5:   for  $P \in S$  do
6:     if  $|H\Delta P| > |H|$ , there is no vertex of degree greater than 2 in  $H\Delta P$ , and there are at least as
       many degree 1 vertices in  $H$  as in  $H\Delta P$ , then
7:       Set  $H = H\Delta P$ 
8:     end if
9:   end for
    // Subroutine B: search for "cost-effective" candidate trails
10:  Find  $P \in S$  such that  $H\Delta P$  has no vertex of degree greater than 2 maximizing the cardinality of
     $H\Delta P$ 
11:  if Step 10 finds a  $P$  and  $|H\Delta P| \geq |H| + \sqrt{\log n}$  then
12:    Set  $H = H\Delta P$ 
13:  end if
14:  if  $H$  was not updated in lines 5-12 then
15:    Set  $can\_grow = False$ .
16:  end if
17: end while
18: return  $H$ 

```

To prove the second part of the statement, we transform the difference graph $H\Delta H^*$ by splitting the vertices of degree at least 3 as follows. If vertex v has degree at least 3, we split it into two copies, and split the incident edges to v such that one red edge and one blue edge are incident to one copy, while the rest are incident to the other copy. After the transformation, every vertex in $H\Delta H^*$ has degree at most 2, and there are $2m$ vertices of degree 1, which exactly correspond to the $2m$ vertices of degree 1 in H . Thus, the difference graph can be decomposed into a collection of edge-disjoint trails, including exactly m open trails whose endpoints must have degree 1 in H . Moreover, whenever a trail goes in and out of a vertex in $V(H^*) \cap V(H)$, it must alternate between red and blue edges. To get the trail decomposition for the original $H\Delta H^*$, we simply merge the copies of vertices. $\square$

Building upon the above trail decomposition, we show that any H consisting of cycles and $o(n)$ paths with a total of $\delta n - o(n)$ edges achieves almost exact recovery of H^* .

Lemma 4.3. *Let $\delta \in (0, 1]$, $0 < \lambda < \frac{1}{(\sqrt{2\delta} + \sqrt{1-\delta})^2}$, and $G \sim \mathcal{G}(n, \lambda, \delta)$ conditional on $H^* = h$. Let H be a (possibly random) subgraph of G such that every vertex in H has degree at most 2, $\mathbb{E}[|H|] = \delta n - o(n)$, and the expected number of degree-1 vertices in H is $o(n)$. Then $\mathbb{E}[|H^* \Delta H|] = o(n)$.*

Proof. The proof is similar to Theorem 2.4, in that it first decomposes $H\Delta H^*$ into trails (closed or open), and then argues that not many long trails can have more than $(1/2 - \epsilon)$ fraction of their edges unplanted. That allows us to bound the number of unplanted edges in $H\Delta H^*$, which in turn yields a bound on the total number of edges in $H\Delta H^*$, since $H\Delta H^*$ is almost balanced under the assumptions on H .

Formally, we start by applying Lemma 4.2 to decompose $H\Delta H^*$ into a collection of trails $\{P_i\}_{1 \leq i \leq m}$ with a_i planted and b_i unplanted edges respectively, such that the number of open trails is exactly equal to half of the number of degree-1 vertices in H , and each trail alternates between red and blue edges whenever it goes in and out of a vertex in $V(H) \cap V(H^*)$.

For a trail P with a planted edges and b unplanted edges, define the excess unplanted edges of P to be $\text{ex}(P) = b - (1/2 - \epsilon)(a + b)$, where ϵ is a constant to be specified later. On one hand, we have

$$|H \setminus H^*| = \sum_{i=1}^m b_i = \sum_{i=1}^m \text{ex}(P_i) + \left(\frac{1}{2} - \epsilon\right) |H\Delta H^*|.$$

On the other hand,

$$|H\Delta H^*| = |H \setminus H^*| + |H^* \setminus H| = 2|H \setminus H^*| + (|H^*| - |H|),$$

where the last equality is because $|H^* \setminus H| - |H \setminus H^*| = |H^*| - |H|$. Combining the two displays above gives

$$\epsilon |H\Delta H^*| = \sum_{i=1}^m \text{ex}(P_i) - |H \setminus H^*| + \frac{1}{2} |H\Delta H^*| = \sum_{i=1}^m \text{ex}(P_i) + \frac{1}{2} (|H^*| - |H|)$$

so that

$$|H\Delta H^*| = \frac{1}{\epsilon} \left(\sum_{i=1}^m \text{ex}(P_i) + \frac{1}{2} (|H^*| - |H|) \right).$$

By assumption, $\mathbb{E}[|H^*| - |H|] = \delta n - \mathbb{E}[|H|] = o(n)$. We only need to bound the expected number of excess unplanted edges; namely, $\mathbb{E}[\sum_{i \leq m} \text{ex}(P_i)]$. Let $I \subset [m]$ index those trails P_i that are closed; i.e., those which are circuits. Then $\mathbb{E}[\sum_{i \in I} \text{ex}(P_i)]$ can be bounded in the same way as in the proof of Theorem 2.4. To recap, we have

$$\sum_{i \in I} \text{ex}(P_i) \leq \sum_{C \subset G} \text{ex}(C) \mathbf{1}_{\{\text{ex}(C) > 0\}} \triangleq \Gamma,$$

where we sum over all (a, b) -circuits C in G with $a \geq 0$ and $b \geq 1$. Moreover, in the proof of Theorem 2.4, we have shown that $\mathbb{E}[\Gamma] \leq C$ for some constant C independent of n . Therefore,

$$\mathbb{E} \left[\sum_{i \in I} \text{ex}(P_i) \right] \leq \mathbb{E}[\Gamma] \leq C.$$

The same argument does not work for those P_i that are open, since open trails start and end at two different vertices, so that G may contain $\Theta(n)$ open trails with large excess. For a tighter analysis, we need to utilize the assumption that there are not many degree-1 vertices in H .

For each $i \in I^c$, if P_i both starts and ends with an all-red segment, then we remove the initial red segment; if P_i both starts and ends with an all-blue segment, then we split P_i into two trails, where the last segment becomes a new trail by itself. Let $P'_1, \dots, P'_M$ denote the resulting collection of trails for some M . Crucially, each P'_i (reversing the traversal direction if necessary) must start with unplanted edges and end with planted edges if there are any; thus, it is a valid (a, b) -trail for some $a \geq 0$ and $b \geq 1$. Moreover, since each P'_i contains at least one distinct degree-1 vertex in H , M is upper bounded by the number of degree-1 vertices in H and hence $\mathbb{E}[M] = o(n)$. Therefore

$$\sum_{i \in I^c} \text{ex}(P_i) \leq \sum_{i=1}^M \text{ex}(P'_i) \leq \sum_{i=1}^M \sum_{\ell=0}^{\infty} \mathbf{1}_{\{\text{ex}(P'_i) > \ell\}} = \sum_{\ell=0}^{\infty} \sum_{i=1}^M \mathbf{1}_{\{\text{ex}(P'_i) > \ell\}} \leq \sum_{\ell=0}^{\infty} \min \{M, N_\ell\},$$

where $N_\ell \triangleq \sum_{P \subset G} \mathbf{1}_{\{\text{ex}(P) > \ell\}}$ and the sum is over all possible (a, b) -trails P in G with $a \geq 0$ and $b \geq 1$. Taking the expected value on both sides yields that

$$\mathbb{E} \left[\sum_{i \in I^c} \text{ex}(P_i) \right] \leq \sum_{\ell=0}^{\infty} \mathbb{E}[\min \{M, N_\ell\}] \leq \sum_{\ell=0}^{\infty} \min \{\mathbb{E}[M], \mathbb{E}[N_\ell]\}, \quad (4.1)$$

where the second inequality follows from Jensen's inequality as $\min\{x, y\}$ is a concave function of (x, y) . By Lemma 2.2, for any two given vertices v, v' , the expected number of (a, b) -trails from v to v' is at most $c_{a,b}/(\delta n)$ if $a \geq 1, b \geq 1$ and $(1 - \delta)^{b-1} \lambda^b / n$ if $a = 0, b \geq 1$. Therefore,

$$\begin{aligned} \mathbb{E}[N_\ell] &\leq n^2 \sum_{a, b \geq 1} \frac{c_{a,b}}{\delta n} \mathbf{1}_{\{b - (1/2 - \epsilon)(a+b) > \ell\}} + n^2 \sum_{b \geq 1} \frac{(1 - \delta)^{b-1} \lambda^b}{n} \mathbf{1}_{\{(1/2 + \epsilon)b > \ell\}} \\ &\leq \frac{n}{\delta} \sum_{a, b \geq 1} c_{a,b} \mathbf{1}_{\{b - (1/2 - \epsilon)(a+b) > \ell\}} + \frac{n}{(1 - \delta)} \frac{[(1 - \delta)\lambda]^\ell}{1 - (1 - \delta)\lambda}, \end{aligned}$$

where the second inequality holds because $(1 - \delta)\lambda < 1$.

As in the proof of Theorem 2.4, we choose $0 < \epsilon < 1/2$ and $x < 1, y < \frac{1}{\lambda(1 - \delta)}$ such that $\frac{2x}{1 - x} \cdot \frac{\delta \lambda y}{1 - (1 - \delta)\lambda y} < 1$ and $x^{1+2\epsilon} y^{1-2\epsilon} = 1$. Such ϵ, x, y always exist when $0 < \lambda < \frac{1}{(\sqrt{2\delta} + \sqrt{1 - \delta})^2}$. Therefore,

$$\sum_{a, b \geq 1} c_{a,b} \mathbf{1}_{\{b - (1/2 - \epsilon)(a+b) > \ell\}} = \sum_{a, b \geq 1} c_{a,b} \mathbf{1}_{\{b - (1/2 - \epsilon)(a+b) > \ell\}} \left(x^{1/2 + \epsilon} y^{1/2 - \epsilon} \right)^{a+b}$$

$$\begin{aligned}
&= \sum_{a,b \geq 1} c_{a,b} \mathbf{1}_{\{b - (1/2 - \epsilon)(a+b) > \ell\}} x^a y^b \left(\frac{x}{y}\right)^{b - (1/2 - \epsilon)(a+b)} \\
&\leq \left(\frac{x}{y}\right)^\ell \sum_{a,b \geq 1} c_{a,b} x^a y^b \\
&\leq \left(\frac{x}{y}\right)^\ell \frac{1}{1 - \frac{2x}{1-x} \cdot \frac{\delta \lambda y}{1 - (1-\delta)\lambda y}},
\end{aligned}$$

where the first inequality holds because $x/y < 1$; and the second inequality follows from the generating function of $c_{a,b}$ given by (2.4).

Combining the last two displayed inequalities yields that $\mathbb{E}[N_\ell] \leq c_1 c_2^\ell n$ for some constants $c_1 > 0$ and $0 < c_2 < 1$ that are independent of n, ℓ . Plugging it back into (4.1) and defining $\ell_0 \triangleq \left\lceil \frac{\log\left(\frac{c_1 n}{\mathbb{E}[M]}\right)}{\log(1/c_2)} \right\rceil$ yields that

$$\begin{aligned}
\mathbb{E}\left[\sum_{i \in I^c} \text{ex}(P_i)\right] &\leq \sum_{\ell=0}^{\infty} \min\left\{\mathbb{E}[M], c_1 c_2^\ell n\right\} \\
&\leq \mathbb{E}[M] \ell_0 + \sum_{\ell \geq \ell_0} c_1 c_2^\ell n \\
&\leq \mathbb{E}[M] \ell_0 + \frac{\mathbb{E}[M]}{1 - c_2} = o(n),
\end{aligned}$$

where the last inequality holds because $c_1 c_2^{\ell_0} n \leq \mathbb{E}[M]$ and the last equality holds because $\mathbb{E}[M] = o(n)$ and $\ell_0 = O(\log(n/\mathbb{E}[M])) = o(n/\mathbb{E}[M])$.

Finally, we conclude that

$$\begin{aligned}
\mathbb{E}[|H \Delta H^*|] &= \frac{1}{\epsilon} \left(\mathbb{E}\left[\sum_{i \in I} \text{ex}(P_i)\right] + \mathbb{E}\left[\sum_{i \notin I} \text{ex}(P_i)\right] + \frac{1}{2} \mathbb{E}[|H^*| - |H|] \right) \\
&= O(1) + o(n) + o(n) = o(n).
\end{aligned}$$

□

Lemma 4.4. *Let G be a graph that contains a 2-factor H^* on δn vertices. The H output by Algorithm 4 on input G will have at least $\delta n - O(n/\sqrt{\log n})$ edges and at most $2n/\sqrt{\log n}$ vertices of degree 1.*

Proof. To prove that Algorithm 4 outputs a large enough subgraph, we argue that until H contains $\delta n - O(n/\sqrt{\log n})$ edges, in each iteration of the main loop of the algorithm, we can find some trail (closed or open) to XOR onto H to make it bigger, through either subroutine A or B. To show that the output does not contain too many degree-1 vertices, we will argue that only subroutine B can increase the number of degree-1 vertices in H . The increase is by at most 2, and must also increase H by at least $\sqrt{\log n}$.

In more detail, assume we are at a certain iteration of the while loop of the algorithm with an iterate of H . By the algorithm's construction, H does not contain any vertices of degree greater than 2. In other words, H is a vertex-disjoint union of cycles and paths. On a high level, we prove the lemma by arguing that there exists a subgraph P of $H \Delta H^*$ that makes a “good” candidate for the XOR operation, updating H to $H \Delta P$. In particular, we will invoke Lemma 4.2 to decompose $H \Delta H^*$ into a collection $\mathcal{C}$ of trails. Next, we show that some $P \in \mathcal{C}$, or slight variation of it, qualifies as a candidate trail to update H .

To begin with, we state some simple facts on $H \Delta P$ that will be handy in our later arguments.

Fact 4.5. (a) *For each $P \in \mathcal{C}$, the XOR operation $H \Delta P$ simply adds the planted edges of P to H , and removes the unplanted edges of P from H .*

(b) *For $P \in \mathcal{C}$, all degree-1 vertices in $H \Delta P$ are also of degree 1 in H , and $H \Delta P$ does not contain any vertices of degree greater than 2.*

Proof of Fact 4.5. Fact (a) holds simply because P is a subgraph of $H \Delta H^*$, thus the planted edges in P are in $H^* \setminus H$ and the unplanted edges are in $H \setminus H^*$.

To prove Fact (b), we first argue that vertices of even degree in H remain even degree in $H \Delta P$. That is because if the degree of a vertex v is even in H and odd in $H \Delta P$, then v must have an odd degree in P , which is only possible if P is an open trail and v is one of its endpoints. However, by Lemma 4.2, all endpoints of open trails must be degree-1 vertices in H , forming a contradiction.

Since H only contains vertices of degree 1 or 2, $H\Delta P$ does not contain any new degree-1 vertices. To prove that it cannot create any vertices of degree greater than 2, we only need to show that degree-1 vertices in H cannot have degree 3 in $H\Delta P$, and degree-2 vertices in H cannot have degree 4 in $H\Delta P$. For either case to occur, the vertex v must be of degree 3 or 4 in $H\Delta H^*$, and the XOR operation must be adding the two planted edges incident to v to H . However, Lemma 4.2 ensures that whenever $P \in \mathcal{C}$ goes in and out of a vertex of degree 3 or 4 in $H\Delta H^*$, it goes through a planted and an unplanted edge. Therefore the XOR operation can never only add 2 planted edges without removing any unplanted edges. Thus no vertex in $H\Delta P$ can be of degree more than 2. $\square$

By Fact 4.5(b), all $P \in \mathcal{C}$ are automatically “cost-free,” in the sense that they can be XOR’ed onto H without increasing the number of degree-1 vertices or creating any vertices of degree greater than 2. However, this does not guarantee that H is updated in subroutine A, since Line 6 is only run if the subroutine finds a “cost-free” trail that is a) of length less than $\log n$, and b) strictly increases $|H|$. We discuss the following two cases. Either there is some member of $\mathcal{C}$ that satisfies both constraints, in which case H is guaranteed to be updated in subroutine A; or no $P \in \mathcal{C}$ satisfies both constraints, in which case we will argue that H must be updated in subroutine B if subroutine A leaves H unchanged.

Case 1: There exists some $P \in \mathcal{C}$ such that $|P| < \log n$ and $|H\Delta P| > |H|$. We claim that in this case, subroutine A successfully finds at least one “cost-free” candidate trail P .

Since H only contains edges in G by construction, and P is from the decomposition of $H\Delta H^*$, we have that P belongs to the set S of all trails in G of length smaller than $\log n$. In subroutine A, the algorithm searches through all members of S in a for-loop. At some point, it will encounter $P \in S$. If H has already been updated in the for-loop before it reaches P , then we know that H is updated at least once in the for-loop; otherwise P is XOR’ed onto H . Either way, we have that at the end of the for loop, subroutine A successfully finds at least one “cost-free” candidate trail, and H gets updated. Therefore $|H|$ strictly increases.

Case 2: There exists no $P \in \mathcal{C}$ such that $|P| < \log n$ and $|H\Delta P| > |H|$. In this case, we claim that if subroutine A leaves H unchanged, subroutine B is guaranteed to find a “cost-effective” trail. The high-level reasoning is as follows: in Case 2, all “short” trails of length smaller than $\log n$ in $\mathcal{C}$ do not increase $|H|$ when XOR’ed onto H , meaning that they all contain at least as many unplanted as planted edges. However, before H grows to contain $\delta n - O(n/\sqrt{\log n})$ edges, there are still a large number of planted edges not picked up by H . Thus, the collection $\mathcal{C}$ (which comes from the decomposition of $H\Delta H^*$) must contain some long trail with a lot more planted than unplanted edges. As a result, as subroutine B searches through trails of length up to $\log n$, there exists some segment of a long trail in $\mathcal{C}$ that qualifies as “cost-effective,” ensuring that H is updated in subroutine B. Now, we present the rigorous proof below.

Since

$$(\# \text{ of planted edges in } H\Delta H^*) - (\# \text{ of unplanted edges in } H\Delta H^*) = |H^* \setminus H| - |H \setminus H^*| = |H^*| - |H|,$$

the overall collection $\mathcal{C}$ contains $|H^*| - |H|$ more planted edges than unplanted edges. Also, recall that the “short” trails contain at least as many unplanted edges as planted ones. Therefore, those $P \in \mathcal{C}$ of length at least $\log n$ must have at least $|H^*| - |H|$ more planted edges than unplanted edges. For each $P \in \mathcal{C}$ of length at least $\log n$, we divide it into segments of length $\log n/2$. If there are edges left over, we combine it with the last segment. This way, all long trails in $\mathcal{C}$ are chopped into trails whose lengths are in $[\lceil \log n/2 \rceil, \log n)$. As long as $|H| < \delta n - cn/\sqrt{\log n}$ for a large enough constant c , we claim that at least one of these chopped trails contains at least $2\sqrt{\log n}$ more planted than unplanted edges. Otherwise, the total number of chopped trails must be at least $(|H^*| - |H|)/(2\sqrt{\log n}) > cn/(2\log n)$, amounting to a total of at least $\frac{cn}{2\log n} \cdot \frac{\log n}{2} = cn/4$ edges. For large enough c ($c = 9$ suffices) this is impossible since the number of edges in $\mathcal{C}$ is bounded by $|H\Delta H^*| \leq |H| + |H^*| \leq 2\delta n$.

Now we have found a trail P' that contains at least $2\sqrt{\log n}$ more planted than unplanted edges, which means that it can be XOR’ed onto H so that $|H\Delta P'| \geq |H| + 2\sqrt{\log n}$. For Line 12 of Algorithm 4 to be run, we also need to check that $H\Delta P'$ does not contain any vertices of degree greater than 2. Note that since P' is only a segment of some $P \in \mathcal{C}$, Fact 4.5(b) does not apply directly to $H\Delta P'$. However, Lemma 4.2 entails that whenever P goes in and out of a vertex $v \in V(H) \cap V(H^*)$, it must go through a planted edge e and an unplanted edge e' . Therefore, the only way for v to have degree 3 or more in $H\Delta P'$ is if P is chopped at v , and P' contains the planted edge e . In other words, P' starts or ends at v through the planted edge e . In this case, we simply remove the planted edge e from P' . Doing the same for both endpoints of P' results in a trail $\tilde{P}$, which is P' with at most two planted edges removed. For each vertex in $\tilde{P}$, its degree in $H\Delta \tilde{P}$ is bounded by its degree in $H\Delta P'$, which is at most 2. Moreover, since P' contains at least $2\sqrt{\log n}$ more planted than unplanted edges, we have that $\tilde{P}$ contains at least $2\sqrt{\log n} - 2 \geq \sqrt{\log n}$ more planted than unplanted edges. Thus $|H\Delta \tilde{P}| \geq |H| + \sqrt{\log n}$.

To summarize, in Case 2, there exists some trail that when XOR’ed onto H , does not create any vertex of degree greater than 2, and increases $|H|$ by at least $\sqrt{\log n}$. Therefore, if H is not updated in subroutine A, it must be updated in subroutine B. That is, Line 10 of Algorithm 4 is run, the update increases $|H|$ by at least $\sqrt{\log n}$, and increases the number of degree-1 vertices in H by at most 2. That is because updating H to be $H\Delta P$ can only create a new degree 1 vertex if that vertex is an endpoint of P .

Combining Cases 1 and 2, we have that until H contains $\delta n - O(n/\sqrt{\log n})$ edges, each while loop iteration updates H to be bigger. Thus the algorithm outputs an H with at least $\delta n - cn/\sqrt{\log n}$ edges for some large enough constant c . Moreover, for the update to increase the number of degree-1 vertices by 1 or 2, the update must be through subroutine B, which simultaneously increases H by at least $\sqrt{\log n}$. Given that the output H has degree at most 2, we have $|H| \leq n$ and thus it contains at most $2n/\sqrt{\log n}$ degree-1 vertices. $\square$

Lemma 4.6. *Let $\delta \in (0, 1]$, $\lambda > 0$ be a constant, and $G \sim \mathcal{G}(n, \lambda, \delta)$ conditional on H^* . Then Algorithm 4 on input G runs in time $O\left(n^{3+\log(2+\lambda)}\right)$ with high probability.*

Proof. The runtime of Algorithm 4 is the sum of the runtime required to find the set S comprising the trails in G of length less than $\log n$, along with the runtime required to execute the while loop.

For convenience, we use the adjacency list representation of G , where we record the list of neighbors of each vertex. In order to find the set of trails in G of length less than $\log n$ which begin at a particular vertex v , we employ a breadth-first traversal. In expectation, at most $\sum_{k=1}^{\log n} (2+\lambda)^k$ trails beginning at v will be discovered, where $(2+\lambda)^k$ is an upper bound on the expected number of trails at depth exactly k beginning at v . Therefore, Markov's inequality implies that with high probability, at most $n^{1+\epsilon} \sum_{k=1}^{\log n} (2+\lambda)^k \leq n^{1+\epsilon} \log n \cdot (2+\lambda)^{\log n}$ trails will be discovered in total for any $\epsilon > 0$. Each trail requires visiting up to $\log n$ edges, and we can keep track of already-visited edges by modifying the adjacency list representation (delete visited edges and then add them back once a trail is complete). Therefore, finding these trails requires $O\left(n^{1+\epsilon} \log^2 n \cdot (2+\lambda)^{\log n}\right)$ work.

Since Algorithm 4 increases the number of edges in H whenever it updates H and no vertex in H ever has degree greater than 2, there are at most n updates. Each update requires searching through the set S and performing $O(\log n)$ work for each $P \in S$. It follows that the total runtime of the while loop is

$$O(n \cdot |S| \cdot \log n) = O\left(n^{2+\epsilon} \log^2 n \cdot (2+\lambda)^{\log n}\right) = O\left(n^{3+\log(2+\lambda)}\right).$$

Since the while loop dominates the runtime, we conclude that the overall runtime of Algorithm 4 is $O\left(n^{3+\log(2+\lambda)}\right)$ with high probability. $\square$

5 Discussion and future work

In this paper, we have characterized a sharp threshold of almost exact recovery for the planted cycles model $\mathcal{G}(n, \lambda, \delta)$. A precursor paper by the authors [13] examined the special case $\delta = 1$ —where the planted subgraph spans the entire vertex set—and considered more generally planted k -factors for $k \in \mathbb{N}$. Parameterizing the edge probability in the background graph as $p_n = \lambda_n/n$, [13] discovers a curious “All-Something-Nothing” transition in recovery as λ_n increases: when $\lambda_n = o(1)$, exact recovery is possible; when $\Omega(1) = \lambda_n \leq 1/k$, almost exact recovery remains feasible; when $1/k < \lambda_n \leq O(1)$, only partial recovery is achievable; and once $\lambda_n = \omega(1)$, even partial recovery becomes impossible. This multi-phase recovery landscape contrasts sharply with the “All-or-Nothing” transition established by an earlier work [22], for the recovery of a planted graph H^* within a background graph $G \sim \mathcal{G}(n, p_n)$, whenever H^* satisfies certain properties related to size, density, and balancedness. In those cases, the recovery transition between almost exact recovery (“all”) and impossibility of partial recovery (“nothing”) occurs at a threshold that depends on the first-moment thresholds of subgraphs $J \subset H^*$. “All-or-Nothing” transitions are also known to occur in a variety of other high-dimensional inference problems, including sparse linear regression [24], sparse tensor PCA [23], group testing [8, 25], and graph matching [26].

Our results taken together identify the “All” regime (Theorems 1.1 and B.1) and partially identify the “Nothing” regime (Theorem B.2). Yet a compelling open question remains: is there a “Something” regime in between? That is, for constant $\lambda > (\sqrt{2\delta} + \sqrt{1-\delta})^{-2}$, can we achieve partial recovery in the sense that $\mathbb{E}[\ell(H^*, \hat{H})] \leq 1 - \Omega(1)$? In the special case of $\delta = 1$, one can partially recover the planted 2-factor with no errors by simply collecting edges incident to degree-2 vertices. However, for $\delta < 1$, this simple structural cue disappears, and recovery becomes subtler. A naïve alternative of randomly sampling δn edges does recover a constant fraction of planted edges, but also mistakenly includes many unplanted ones, resulting in a high overall error and failing to achieve partial recovery.

We list some additional open problems below.

- If the $\mathcal{G}(n, \lambda, \delta)$ model admits a “something” phase, what is the optimal recovery rate in the partial recovery regime, and can it be achieved by an efficient algorithm?
- Consider a general $\mathcal{G}(n, \lambda, \delta, k)$ model in which a partial k -factor is planted on a subset of δn vertices in the background graph, so that the case $k = 2$ reduces to the planted cycles model $\mathcal{G}(n, \lambda, \delta)$. When $k = 1$ and

$\delta \in (0, 1)$, this corresponds to the planted partial matching problem, for which there is no sharp phase transition for almost-exact recovery: if $\lambda = o(1)$, recovery is possible since there are only $o(n)$ unplanted edges, whereas if $\lambda = \Omega(1)$, recovery becomes impossible because a partial matching of size $\Theta(n)$ formed by unplanted edges will exist, and there is no way to tell apart the edges in this partial matching from the edges in the planted one. This raises the question: what are the recovery thresholds in this generalized model for $k \geq 3$? When $k \geq 3$, the transition from almost-exact recovery to impossibility may be driven by the emergence of k -factors which differ significantly from the planted one, rather than k -factors that differ slightly.

- Similarly, what can we say about planted subgraph recovery thresholds for graphs other than k -factors?
- What can be said for weighted graphs, where planted edges have weights drawn from a distribution $\mathcal{P}$ while unplanted edges have weights drawn from a distribution $\mathcal{Q}$? For the special case of a weighted planted matching where $\mathcal{P}$ and $\mathcal{Q}$ follow exponential distributions, the recovery rate has been precisely characterized [21].

A History-dependent branching process

In order to characterize the sizes of the trees constructed by Algorithm 1, we compare the trees to branching processes. At a high level, the probability that a given tree reaches a prescribed depth can be related to the survival probability of the branching process. We need the following auxiliary result about the survival of a supercritical branching process (used in the proof of Lemma 3.9).

Lemma A.1. *Suppose a branching process has offspring distribution with expected value μ and variance σ^2 for some $\mu > 1$, we have*

$$\mathbb{P}\{\text{Branching process survives}\} \geq \frac{\mu^2 - \mu}{\mu^2 - \mu + \sigma^2}. \quad (\text{A.1})$$

A similar bound holds for an inhomogeneous, possibly history-dependent branching process. Specifically, let $X_1, X_2, \dots$ be the number of children born from vertices in breadth-first order, and let $\mathcal{F}_1 = \sigma(X_1), \mathcal{F}_2 = \sigma(X_1, X_2), \dots$ be the natural filtration. Suppose that $\mathbb{E}[X_i | \mathcal{F}_{i-1}] \geq \mu > 1$ and $\text{Var}(X_i | \mathcal{F}_{i-1}) \leq \sigma^2$ uniformly for all $i \in \{1, 2, \dots\}$. Then

$$\mathbb{P}\{\text{Branching process survives}\} \geq \frac{\mu^2 - \mu}{\mu^2 - \mu + \sigma^2 + \frac{1}{4}}.$$

The proof of Lemma A.1 requires the following fact about distributions related in stochastic order.

Proposition A.2. *Let P be a finite discrete distribution with support $S \subseteq \{0, 1, 2, \dots, s\}$. Let μ and σ^2 respectively denote the mean and variance of P . Then for any $\mu' \in [0, \mu]$, there exists a distribution Q with support $S' \subseteq \{0, 1, 2, \dots, s\}$, mean μ' , and variance σ'^2 such that $Q \preceq P$ and $\sigma'^2 \leq \sigma^2 + \frac{1}{4}$.*

Proof. We construct the distribution Q from P by iteratively “shifting” mass from larger elements of S to smaller ones. Observe that any distribution shifted from P in this way is automatically a stochastic lower bound on P . Shifting always reduces the mean, and may increase or decrease the variance.

Without loss of generality, P places nonzero mass on s (otherwise, we could redefine S). Suppose we shift the probability mass x from s to $s-1$, where $x \in [0, P(s)]$, and let Q' denote the resulting probability distribution. Next, we bound the change of the variance. To this end, we define a coupling (X, X') , where $(X, X') = (i, i)$ with probability $P(i)$ for all $i \leq s-1$, $(X, X') = (s, s)$ with probability $P(i) - x$, and $(X, X') = (s, s-1)$ with probability x . It follows that X is distributed as P and X' is distributed as Q' . Let (Y, Y') be an independent copy of (X, X') . It follows that

$$\begin{aligned} \text{Var}(X') - \text{Var}(X) &= \frac{1}{2} \mathbb{E} \left[(X' - Y')^2 - (X - Y)^2 \right] \\ &= \frac{1}{2} \mathbb{E} \left[\left((X' - Y')^2 - (X - Y)^2 \right) \mathbf{1}_{\{X=s, X'=s-1, Y=Y' \text{ or } X=X', Y=s, Y'=s-1\}} \right] \\ &= \mathbb{E} \left[\left((s-1 - X')^2 - (s - X)^2 \right) \mathbf{1}_{\{X=X'\}} \right] \cdot x \\ &\leq \mathbb{P}\{X = X' = s\} \cdot x = (P(i) - x) \cdot x, \end{aligned}$$

where the first equality holds by the definition of variance; the second equality holds because $X \neq X'$ if and only if $X = s$ and $X' = s-1$, and the same holds for Y, Y' ; the third equality holds because (X, X') and (Y, Y') are independent copies, and $Y = s, Y' = s-1$ with probability x ; the last inequality holds because $(s-1 - X')^2 \leq (s - X)^2$ when $X = X' \leq s-1$ and $(s-1 - X')^2 - (s - X)^2 = 1$ when $X = X' = s$. Therefore, if $x = P(i)$, then the variance of Q' is always no larger than that of P ; if $x < P(i)$, then the variance increases at most by $x(1-x) \leq 1/4$.

Therefore, we can iteratively shift all the probability mass $x = P(s)$ from s to $s-1$, where s decreases by 1 in each iteration, until the very last iteration, where we only shift $x \leq P(s)$ so that the mean decreases to exactly μ' . By

the above argument, the variance keeps non-increasing until the very last iteration, in which it may increase by at most $1/4$. Therefore, we arrive at the final distribution $Q \preceq P$ with mean μ' and variance at most $\sigma^2 + 1/4$. $\square$

Proof of Lemma A.1. We first consider the homogeneous case. In that case, the proof mostly follows the derivations in [10, Chapter 2.1]. Let Z_m denote the number of vertices in generation m . Given Z_{m-1} , the conditional first and second moments of Z_m satisfy

$$\begin{aligned}\mathbb{E}[Z_m | Z_{m-1}] &= \mu Z_{m-1}, \\ \mathbb{E}[Z_m^2 | Z_{m-1}] &= \mu^2 Z_{m-1}^2 + Z_{m-1} \sigma^2.\end{aligned}$$

Taking expected values on both sides and iterating and noting $Z_0 = 1$, we have

$$\mathbb{E}[Z_m] = \mu^m,$$

and

$$\begin{aligned}\mathbb{E}[Z_m^2] &= \mu^{2m} + \sigma^2 \sum_{j=m-1}^{2m-2} \mu^j \\ &\leq \mu^{2m} + \sigma^2 \frac{\mu^{2m-2}}{1 - \mu^{-1}}.\end{aligned}\tag{A.2}$$

By the Paley–Zygmund inequality, the probability that the branching process survives to iteration m is

$$\mathbb{P}\{Z_m \geq 1\} \geq \frac{\mathbb{E}[Z_m]^2}{\mathbb{E}[Z_m^2]} \geq \frac{\mu^{2m}}{\mu^{2m} + \sigma^2 \frac{\mu^{2m-2}}{1 - \mu^{-1}}} = \frac{\mu^2 - \mu}{\mu^2 - \mu + \sigma^2}.$$

Take $m \rightarrow \infty$ to finish the proof.

For the inhomogeneous, possibly history-dependent case, we will construct an auxiliary branching process, where $\{\bar{Z}_m\}$ denotes the number of vertices in generation m in the auxiliary branching process. Let $\bar{X}_i$ be the number of children of the i^{th} node in the auxiliary tree. The initial children $\bar{Z}_1$ are sampled from the shifted distribution based on X_1 , guaranteeing $\bar{Z}_1 = \bar{X}_1 \leq X_1$. The children of the two processes are arbitrarily matched, and their offspring are coupled. For $i \in \{1, 2, \dots, \bar{X}_1\}$, conditioned on $\mathcal{F}_i$, the number of offspring X_{i+1} is coupled to $\bar{X}_{i+1}$, such that $\bar{X}_{i+1} \leq X_{i+1}$ and $\bar{X}_{i+1}$ has mean μ and variance at most $\sigma^2 + 1/4$, as in Proposition A.2. We can continue matching pairs across trees and coupling their offspring. It follows that $\bar{Z}_m \leq Z_m$. Moreover,

$$\mathbb{E}[\bar{Z}_m] = \mu^m,$$

analogously to the homogeneous case. To compute the conditional second moment $\mathbb{E}[\bar{Z}_m^2 | \bar{Z}_{m-1}]$, we claim that

$$\text{Var}(\bar{Z}_m | \bar{Z}_{m-1}) \leq \left(\sigma^2 + \frac{1}{4}\right) \bar{Z}_{m-1}.$$

To see this, let $N = \bar{Z}_{m-1}$, and let $Y_1, \dots, Y_N$ denote the numbers of offspring of the $(m-1)^{\text{th}}$ generation, so that $\bar{Z}_m = \sum_{i=1}^N Y_i$. Let $\mathcal{F}$ denote the history of the auxiliary branching process up to the time when the offspring of the $(N-1)$ -th node at the $(m-1)$ -th generation is added. Then

$$\begin{aligned}\text{Var}(\bar{Z}_m | \bar{Z}_{m-1}) &= \text{Var}\left(\sum_{i=1}^N Y_i\right) \\ &= \mathbb{E}\left[\text{Var}\left(\sum_{i=1}^N Y_i \mid \mathcal{F}\right)\right] + \text{Var}\left(\mathbb{E}\left[\sum_{i=1}^N Y_i \mid \mathcal{F}\right]\right) \\ &= \mathbb{E}[\text{Var}(Y_N | \mathcal{F})] + \text{Var}\left(\sum_{i=1}^{N-1} Y_i + \mu\right) \\ &\leq \left(\sigma^2 + \frac{1}{4}\right) + \text{Var}\left(\sum_{i=1}^{N-1} Y_i\right),\end{aligned}$$

where the second equality holds by the law of total variance; the third equality holds because $\mathbb{E}[Y_N | \mathcal{F}] = \mu$; the last inequality holds due to $\text{Var}(Y_N | \mathcal{F}) \leq \sigma^2 + 1/4$. Thus, we have shown

$$\text{Var}\left(\sum_{i=1}^N Y_i\right) \leq \left(\sigma^2 + \frac{1}{4}\right) + \text{Var}\left(\sum_{i=1}^{N-1} Y_i\right).$$

Recursively applying the above inequality yields that

$$\text{Var}(\bar{Z}_m | \bar{Z}_{m-1}) = \text{Var}\left(\sum_{i=1}^N Y_i\right) \leq \left(\sigma^2 + \frac{1}{4}\right) N = \left(\sigma^2 + \frac{1}{4}\right) \bar{Z}_{m-1}.$$

It follows that

$$\mathbb{E}[\bar{Z}_m^2 | \bar{Z}_{m-1}] \leq \mu^2 \bar{Z}_{m-1}^2 + \left(\sigma^2 + \frac{1}{4}\right) \bar{Z}_{m-1}.$$

Taking expected values and iterating as in the homogeneous case, we obtain

$$\mathbb{E}[\bar{Z}_m^2] \leq \mu^{2m} + \left(\sigma^2 + \frac{1}{4}\right) \frac{\mu^{2m-2}}{1 - \mu^{-1}}.$$

Finally, we conclude that

$$\begin{aligned} \mathbb{P}(Z_m \geq 1) &\geq \mathbb{P}(\bar{Z}_m \geq 1) \geq \frac{\mathbb{E}[\bar{Z}_m]^2}{\mathbb{E}[\bar{Z}_m^2]} \geq \frac{\mu^{2m}}{\mu^{2m} + \left(\sigma^2 + \frac{1}{4}\right) \frac{\mu^{2m-2}}{1 - \mu^{-1}}} \\ &= \frac{\mu^2 - \mu}{\mu^2 - \mu + \sigma^2 + \frac{1}{4}}. \end{aligned}$$

□

B Exact and partial recovery

We say that $\hat{H}$ achieves

- *exact recovery*, if $\ell(H^*, \hat{H}) = 0$ with high probability;
- *partial recovery*, if $\mathbb{E}[\ell(H^*, \hat{H})] \leq 1 - \Omega(1)$.

In this section, we show that exact recovery is possible if and only if $\lambda = o(1)$ (Theorem B.1) and partial recovery is impossible when $\lambda = \omega(1)$ (Theorem B.2).

Theorem B.1. *Let $\delta \in (0, 1]$. Suppose $G \sim \mathcal{G}(n, \delta, \lambda)$ is the graph generated from the planted cycles model, and H^* is the hidden 2-factor on the set of δn vertices. If $\lambda = o(1)$, then with probability $1 - o(1)$, H^* is the only 2-factor on δn vertices contained in G . Conversely, if $\lambda = \Omega(1)$, then with probability $\Omega(1)$, G contains another 2-factor on δn vertices which is different from H^* .*

Proof. For any 2-factor H on δn vertices, $H^* \Delta H$ can be decomposed into a union of (a, b) -circuits such that exactly half of the edges are planted. So, if there exists a 2-factor $H \neq H^*$ on δn vertices then there must exist at least one (a, b) -circuit with $a \leq b$. The probability that such a circuit exists is at most equal to the expected number of such circuits, and by Lemma 2.2, that is at most

$$\begin{aligned} &\sum_{b=1}^{\infty} (1 - \delta)^{b-1} \lambda^b + \sum_{a=1}^{\infty} \sum_{b=a}^{\infty} c_{a,b} \\ &= \sum_{b=1}^{\infty} (1 - \delta)^{b-1} \lambda^b + \sum_{a=1}^{\infty} \sum_{b=a}^{\infty} (\lambda(1 - \delta))^b \sum_{k=1}^{\infty} (2\delta/(1 - \delta))^k \binom{a-1}{k-1} \binom{b-1}{k-1} \\ &= o(1) + \sum_{a=1}^{\infty} \sum_{b=a}^{\infty} (\lambda(1 - \delta))^b \sum_{k=1}^a (2\delta/(1 - \delta))^k \binom{a-1}{k-1} \binom{b-1}{k-1} \\ &\leq o(1) + \sum_{a=1}^{\infty} \sum_{b=a}^{\infty} (\lambda(1 - \delta))^b \sum_{k=1}^a (2\delta/(1 - \delta))^k 2^{a+b-2} \\ &\leq o(1) + \sum_{a=1}^{\infty} a \max(1, (2\delta/(1 - \delta))^a) 2^{a-2} \sum_{b=a}^{\infty} (\lambda(1 - \delta))^b 2^b \end{aligned}$$

$$\begin{aligned}
&= o(1) + \sum_{a=1}^{\infty} a \max(1, (2\delta/(1-\delta))^a) 2^{a-2} \frac{(\lambda(1-\delta))^a 2^a}{1-2\lambda(1-\delta)} \\
&= o(1) + \frac{1}{4-8\lambda(1-\delta)} \sum_{a=1}^{\infty} a \max(4\lambda(1-\delta), 8\lambda\delta)^a \\
&= o(1) + \frac{1}{4-8\lambda(1-\delta)} \cdot \frac{\max(4\lambda(1-\delta), 8\lambda\delta)}{(1-\max(4\lambda(1-\delta), 8\lambda\delta))^2} \\
&= o(1).
\end{aligned}$$

Next, suppose $\lambda = \Omega(1)$. We assume that the set of δn planted vertices is known. Then we can reduce to the problem of recovering a planted 2-factor on $m_n = \delta n$ vertices, where the edge probability in the background graph is $\frac{\delta\lambda}{\delta n} = \frac{\delta\lambda}{m_n}$. Since $\lambda = \Omega(1)$, it follows that $\delta\lambda = \Omega(1)$. Thus, the result follows directly from [13, Theorem 2.1]. $\square$

Theorem B.2. *Consider the planted 2-factor model on δn vertices, where $\delta \in (0, 1)$. When $\lambda = \omega(1)$, partial recovery is impossible.*

Proof. As in the proof of Theorem B.1, we assume that the estimator has access to the vertex set $V(H^*)$. An analogous parameter scaling argument along with [13, Theorem 2.5] completes the proof. $\square$

C Recovery of a single planted cycle

We now turn to the problem of recovering a single planted cycle on δn vertices, chosen uniformly at random, and planted within $G \sim \mathcal{G}(n, p_n)$. Recall Remark 1.1, which notes that the achievability results (Theorem 2.4 and Theorem 4.1) apply to the recovery of a single planted cycle. As for the impossibility direction, we observe that a random 2-factor H^* has a significant probability of being a cycle, so anything that holds with sufficiently high probability for a random planted 2-factor carries over to a planted cycle. In particular, we have the following result, which generalizes [13, Lemma 1.1].

Lemma C.1. *Let H^* be a random 2-factor on m vertices with $m \geq 3$. For any $1 \leq \ell \leq m$, with probability at least $1/(1+m2^{-\ell})$, H^* has at most ℓ cycles. In particular, with probability at least $1/m$, H^* is a cycle.*

Proof. Every 2-factor can be converted to the cycle decomposition of a permutation of the vertices. Specifically, we assign each of its cycles a direction and then we have that the permutation maps each vertex to the next vertex in its cycle. Therefore, every 2-factor with more than ℓ cycles can be converted to a permutation of the vertices in at least $2^{\ell+1}$ different ways by assigning each of its cycles a direction, and no two different 2-factors can yield the same permutation. Thus, there are at most $m!/2^{\ell+1}$ possible 2-factors with more than ℓ cycles on m vertices. On the other hand, there are at least $(m-1)!/2$ possible 2-factors on m vertices that consists of a single cycle. Therefore, the probability that a random 2-factor on m vertices has at most ℓ cycles is at least

$$\frac{(m-1)!/2}{m!/2^{\ell+1} + (m-1)!/2} = \frac{1}{1+m2^{-\ell}}.$$

When $\ell = 1$, this probability is $1/(1+m/2) \geq 1/m$ as $m \geq 2$. $\square$

Corollary C.2. *Let P_1 denote the joint distribution of (H^*, G) generated from the planted $\mathcal{G}(n, \lambda, \delta)$ model. Let P_2 denote the distribution of (H^*, G) conditional on H^* being a cycle. Then for any event E defined in terms of (H^*, G) , it holds that $P_2[E] \leq \delta n \cdot P_1[E]$*

Proof. Let C be the event that H^* is a cycle, and note that P_1 conditioned on C is P_2 . Therefore,

$$P_1[E] \geq P_1[E, C] = P_1[E|C] \cdot P_1[C] = P_2[E] \cdot P_1[C] \geq P_2[E]/(\delta n),$$

where the last inequality holds by Lemma C.1. $\square$

The impossibility result follows immediately.

Lemma C.3. *Theorem 3.1 continues to hold when H^* is a planted cycle on δn vertices.*

Proof. Let $E = \{\ell(H^*, \hat{H}) < \epsilon'\}$. Then Theorem 3.1 and Corollary C.2 imply that $\mathbb{P}(E) \leq \delta n \cdot n^{-\Omega(\log n)} = n^{-\Omega(\log n)}$. $\square$

Acknowledgement

J. Gaudio is supported in part by an NSF CAREER award CCF-2440539. J. Xu is supported in part by an NSF CAREER award CCF-2144593.

References

- [1] Noga Alon, Michael Krivelevich, and Benny Sudakov. Finding a large hidden clique in a random graph. *Random Structures & Algorithms*, 13(3-4):457–466, 1998.
- [2] Noga Alon, Rraphy Yuster, and Uri Zwick. Color-coding: a new method for finding simple paths, cycles and other small subgraphs within large graphs. In *Proceedings of the Twenty-sixth Annual ACM Symposium on Theory of Computing*, pages 326–335, 1994.
- [3] Michael Anastos. A note on long cycles in sparse random graphs. *Electronic Journal of Combinatorics*, 30(2), 2023.
- [4] Michael Anastos and Alan Frieze. A scaling limit for the length of the longest cycle in a sparse random graph. *Journal of Combinatorial Theory, Series B*, 148:184–208, 2021.
- [5] Boaz Barak, Samuel Hopkins, Jonathan Kelner, Pravesh K Kothari, Ankur Moitra, and Aaron Potechin. A nearly tight sum-of-squares lower bound for the planted clique problem. *SIAM Journal on Computing*, 48(2):687–735, 2019.
- [6] Béla Bollobás. *Random Graphs*. Springer, 1998.
- [7] Béla Bollobás, TI Fenner, and AM Frieze. Long cycles in sparse random graphs. *Graph Theory and Combinatorics (Cambridge, 1983)*, Academic Press, London, pages 59–64, 1984.
- [8] Amin Coja-Oghlan, Oliver Gebhard, Max Hahn-Klimroth, Alexander S Wein, and Ilias Zadik. Statistical and computational phase transitions in group testing. In *Conference on Learning Theory*, pages 4764–4781. PMLR, 2022.
- [9] Jian Ding, Yihong Wu, Jiaming Xu, and Dana Yang. The planted matching problem: Sharp threshold and infinite-order phase transition. *Probability Theory and Related Fields*, pages 1–71, 2023.
- [10] Richard Durrett. *Random Graph Dynamics*, volume 200. 2007.
- [11] Paul Erdős and Alfréd Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hungar. Acad. Sci.*, 5:17–61, 1960.
- [12] Alan M Frieze. On large matchings and cycles in sparse random graphs. *Discrete Mathematics*, 59(3):243–256, 1986.
- [13] Julia Gaudio, Colin Sandon, Jiaming Xu, and Dana Yang. “All-Something-Nothing” phase transitions in planted k-factor recovery. *arXiv preprint arXiv:2503.08984*, 2025.
- [14] Mark Jerrum. Large cliques elude the metropolis process. *Random Structures & Algorithms*, 3(4):347–359, 1992.
- [15] Richard M Karp. The probabilistic analysis of some combinatorial search algorithms. In *Algorithms and complexity: New Directions and Recent Results*, 1976.
- [16] Anton Kotzig. Moves without forbidden transitions in a graph. *Matematický Časopis*, 18(1):76–80, 1968.
- [17] Michael Krivelevich, Choongbum Lee, and Benny Sudakov. Long paths and cycles in random subgraphs of graphs with large minimum degree. *Random Structures & Algorithms*, 46(2):320–345, 2015.
- [18] Luděk Kučera. Expected complexity of graph partitioning problems. *Discrete Applied Mathematics*, 57(2-3):193–212, 1995.
- [19] Cheng Mao, Alexander S Wein, and Shenduo Zhang. Detection-recovery gap for planted dense cycles. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 2440–2481. PMLR, 2023.

- [20] Cheng Mao, Alexander S Wein, and Shenduo Zhang. Information-theoretic thresholds for planted dense cycles. *IEEE Transactions on Information Theory*, 2024.
- [21] Mehrdad Moharrami, Cristopher Moore, and Jiaming Xu. The planted matching problem: Phase transitions and exact results. *The Annals of Applied Probability*, 31(6):2663–2720, 2021.
- [22] Elchanan Mossel, Jonathan Niles-Weed, Youngtak Sohn, Nike Sun, and Ilias Zadik. Sharp thresholds in inference of planted subgraphs. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 5573–5577. PMLR, 2023.
- [23] Jonathan Niles-Weed and Ilias Zadik. The All-or-Nothing phenomenon in sparse tensor PCA. *Advances in Neural Information Processing Systems*, 33:17674–17684, 2020.
- [24] Galen Reeves, Jiaming Xu, and Ilias Zadik. The all-or-nothing phenomenon in sparse linear regression. *Mathematical Statistics and Learning*, 3(3):259–313, 2021.
- [25] Lan V Truong, Matthew Aldridge, and Jonathan Scarlett. On the All-or-Nothing behavior of Bernoulli group testing. *IEEE Journal on Selected Areas in Information Theory*, 1(3):669–680, 2020.
- [26] Yihong Wu, Jiaming Xu, and Sophie H. Yu. Settling the sharp reconstruction thresholds of random graph matching. *IEEE Transactions on Information Theory*, 68(8):5391–5417, 2022.