

On statistical inference when fixed points of belief propagation are unstable

Siqi Liu*

Sidhanth Mohanty†

Prasad Raghavendra‡

June 3, 2021

Abstract

Many statistical inference problems correspond to recovering the values of a set of hidden variables from sparse observations on them. For instance, in a planted constraint satisfaction problem such as planted 3-SAT, the clauses are *sparse observations* from which the hidden assignment is to be recovered. In the problem of community detection in a stochastic block model, the community labels are hidden variables that are to be recovered from the edges of the graph.

Inspired by ideas from statistical physics, the presence of a stable fixed point for belief propagation has been widely conjectured to characterize the computational tractability of these problems. For community detection in stochastic block models, many of these predictions have been rigorously confirmed.

In this work, we consider a general model of statistical inference problems that includes both community detection in stochastic block models, and all planted constraint satisfaction problems as special cases. We carry out the cavity method calculations from statistical physics to compute the regime of parameters where detection and recovery should be algorithmically tractable. At precisely the predicted tractable regime, we give:

- (i) a general polynomial-time algorithm for the problem of *detection*: distinguishing an input with a planted signal from one without;
- (ii) a general polynomial-time algorithm for the problem of *recovery*: outputting a vector that correlates with the hidden assignment significantly better than a random guess would.

Analogous to the spectral algorithm for community detection [KMM⁺13, BLM15], the detection and recovery algorithms are based on the spectra of a matrix that arises as the derivatives of the belief propagation update rule. To devise a spectral algorithm in our general model, we obtain bounds on the spectral norms of certain families of random matrices with correlated and matrix valued entries. We then demonstrate how eigenvectors of various powers of the matrix can be used to partially recover the hidden variables.

*EECS Department, University of California Berkeley. sliu18@berkeley.edu

†EECS Department, University of California at Berkeley. sidhanthm@berkeley.edu. Supported by Google PhD Fellowship.

‡EECS Department, University of California Berkeley. prasad@cs.berkeley.edu

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 1 |
| 1.1 | Belief Propagation and Cavity Method | 2 |
| 1.2 | Stable Fixed Point Barrier | 4 |
| 1.3 | Related Work | 6 |
| 1.4 | Technical Overview | 7 |
| 1.5 | Discussion and Future Work | 10 |
| 2 | Preliminaries | 12 |
| 2.1 | Observation Model | 12 |
| 2.1.1 | Miscellaneous simplifying notation | 13 |
| 2.2 | Bayesian Inference | 14 |
| 2.3 | Belief Propagation | 15 |
| 2.4 | Stable Fixed Point Barrier | 16 |
| 2.5 | Analyzing Stability | 17 |
| 2.6 | The local distributions of M | 17 |
| 2.7 | The stability condition | 19 |
| 2.8 | Efficient recovery and detection when the fixed point is unstable | 20 |
| 2.9 | The stability condition via a finite linear operator L | 21 |
| 3 | Technical Overview | 22 |
| 3.1 | Algorithm for distinguishing | 22 |
| 3.2 | Matrix concentration vignette | 23 |
| 3.3 | Matrix concentration statement | 24 |
| 4 | A conjectured detection/recovery threshold | 25 |
| 5 | A spectral distinguishing algorithm | 27 |
| 5.1 | Implementation details | 29 |
| 6 | Statistics for the planted model | 30 |
| 6.1 | Statistics for the centered self-avoiding-walk matrix | 31 |
| 6.2 | Bounding contribution of non-self-avoiding walks | 33 |
| 6.3 | Wrapping up estimates of statistics | 34 |
| 7 | Eigenvalue bounds | 35 |
| 7.1 | Proof of Theorem 7.4 | 37 |
| 8 | Weak Recovery | 46 |
| A | Belief propagation for M | 57 |
| B | Proof of Lemma 2.5 | 58 |
| C | The partial derivative matrix | 60 |

| | | |
|----------|-------------------------------|-----------|
| D | Random graph lemmas | 61 |
| D.1 | Proof of Lemma 7.8 | 61 |
| D.2 | Proof of Lemma D.11 | 64 |
| D.3 | Proof of Claim 7.22 | 65 |

1 Introduction

In the PLANTED- q -COLORING problem, a hidden coloring $\mathbf{c} : [n] \rightarrow \{1, \dots, q\}$ is sampled from the uniform distribution over $[q]^n$. A random graph $G = ([n], E)$ is drawn from the Erdős-Renyi distribution conditioned on \mathbf{c} being a legitimate coloring. So every edge (i, j) is included in the graph with probability $\frac{d}{n} \cdot 1[\mathbf{c}(i) \neq \mathbf{c}(j)]$ independently at random. Given the edges E as input, the goal of an inference algorithm is to recover (even partially) the hidden coloring \mathbf{c} .

PLANTED- q -COLORING is the archetypal example of a broad class of statistical inference problems where the goal is to recover a set of hidden variables from sparse observations on it (see [Mon08]). A large number of inference problems ranging from decoding LDPC codes to community detection in random graphs fall into this broad framework. Broadly speaking, the setup in these inference problems is as follows. A set of *hidden variables* $\{c(1), \dots, c(n)\}$ are drawn from a known prior product distribution $\mathbb{P}_{\mathbf{c}}$. A sequence of *observations* (a.k.a. hyperedges) E on these hidden variables are revealed to the algorithm. Each hyperedge (i_1, \dots, i_k) is included with probability $\frac{1}{n^{k-1}} \cdot \Phi(\mathbf{c}(i_1), \dots, \mathbf{c}(i_k))$ for some constant $\Phi(\mathbf{c}(i_1), \dots, \mathbf{c}(i_k))$ that depends on the values of hidden variables $\mathbf{c}(i_1), \dots, \mathbf{c}(i_k)$. Thus the inference algorithm receives $\Theta(n)$ observations with high probability and its goal is to partially recover the values of the hidden coloring.

The key computational task is to recover the values of the hidden variables. In a sparse setup where the number of observations is linear, it is typically impossible to recover the hidden variables exactly. Therefore, one settles for the relaxed goal of *weak* recovery where the algorithm is required to produce an assignment which correlates better than random with hidden variables.

It is often useful to also define a related decision problem of "detection". Here, the algorithm is required to distinguish between a set of observations consistent with a single fixed assignment to hidden variables (planted distribution) or a set of observations each sampled independently by drawing a new assignment to the hidden variables (null distribution).

In this work, we will be considering a more general model that will permit constantly many *types* of variables and observations. The prior distribution of each variable depends on its type, and the probability of sampling an observation depends on the types and values of variables involved. We defer the formal description of our general model to [Section 2.1](#), but instead present a few examples of these problems.

Example 1.1. (Stochastic Block Models) A natural generalization of the PLANTED- q -COLORING problem is the stochastic block model (SBM). The stochastic block model is defined by a parameter q (the number of labels), a distribution $\mathbb{P}_{\mathbf{c}}$ over $[q]$ (the expected fraction of vertices with a specific label), and a matrix $P \in \mathbb{R}^{[q] \times [q]}$ such that $P[c, d]$ gives the probability of an edge between two vertices with labels c and d . In the community detection problem, a hidden labelling $\mathbf{c} : [n] \rightarrow \{1, \dots, q\}$ is sampled from the product distribution $\mathbb{P}_{\mathbf{c}}^n$. Given \mathbf{c} , a random graph $G = ([n], E)$ is drawn by including each edge (u, v) independently with probability $P[\mathbf{c}(u), \mathbf{c}(v)]$ depending on the labels of the endpoints. The goal of the problem is to recover the labelling \mathbf{c} from the graph G .

Example 1.2. (Planted CSPs) In a planted CSP over a domain $[q]$, an assignment $x \in [q]^n$ is chosen at random and clauses are sampled conditioned on being satisfied by the planted assignment x . Depending on the predicate used, one obtains different planted CSPs such as Planted NAE- k -SAT and Planted k -SAT.

Many more examples of problems that fit our framework will be presented in the rest of the paper. Alternatively, this class of problems can be viewed as "*Bayesian CSPs*". Traditionally, a constraint satisfaction problem involves variables taking values over finite domain and a set of local constraints on them.

The goal is to find an assignment that satisfies either all the constraints (exact CSPs) or the largest fraction of constraints (approximate CSP). The key difference in this setup is that there is a prior distribution associated with assignment on the variables and the constraints.

Constraint satisfaction problems (CSP) lie at the bedrock of worst-case complexity theory tracing back all the way to SAT and NP-completeness and by now there is a rich and comprehensive theory that correctly predicts the computational complexity of the traditional CSPs, with (i) the CSP dichotomy conjecture [Sch78, Zhu20] for exact CSPs, which cleanly classifies a constraint satisfaction problem as polynomial-time solvable or NP-hard depending on whether a pair of solutions could be combined to form a third solution via a function called a polymorphism, and (ii) the Unique Games Conjecture for approximate CSPs, which characterizes the best approximation ratio possible in polynomial time with an integrality gap of a semidefinite program [Kho02, KKMO07, Rag08]. There is also a well understood picture of the complexity of refutation of random CSPs from the lens of the Sum-of-Squares semidefinite programming hierarchy [AOW15, RRS17, KMOW17]. On the other hand, our understanding of the complexity of Bayesian CSPs is still in its nascent stages. Bayesian CSPs are a rich and natural class of average case problems, and understanding their complexity would be a good test-bed for average case complexity theory. Indeed, Goldreich’s pseudorandom generator [Gol11] is precisely based on harnessing the computational intractability of certain Bayesian CSPs.

A naive exponential-time algorithm for the problem would be to use the Bayes rule to compute/sample from the conditional distribution $c|E$. The fundamental question here is to understand the limits of efficient algorithms for this class of statistical inference problems. Furthermore, both exact and approximate versions of traditional CSPs exhibit abrupt transitions wherein the computational complexity of the problem changes from polynomial to exponential. It is a compelling question whether Bayesian CSPs also exhibit similar abrupt transitions in computational complexity, and whether there exist broadly applicable optimal algorithms for them.

1.1 Belief Propagation and Cavity Method

A natural candidate for an optimal algorithm for Bayesian CSPs (especially in the sparse case) is belief propagation (BP). BP is often hypothesized to be theoretically optimal, and is also very efficient in practice. There is a vast body of literature on belief propagation (BP) drawing ideas from statistical physics (see [MM09a, Chapter 14] and [ZK16a] for a comprehensive treatment). It is often very difficult to analyze BP as a standalone algorithm and we are quite far from demonstrating its optimality among polynomial-time algorithms. However, there has been a growing body of work in the past decade which suggest a very general and precise theory to predict the computational complexity of Bayesian CSPs.

To the best of our knowledge, it appears that the work of Krzakala and Zdeborova [KZ09] is the first to hypothesize a precise computational phase transition for planted problems based on ideas from statistical physics. Specifically, Krzakala and Zdeborova [KZ09] hypothesized that for a broad class of planted distributions, the problem of distinguishing the planted vs null distributions becomes computationally intractable at a well-defined threshold. In the case of community detection, this threshold coincides with the so-called Kesten-Stigum threshold. More broadly, in this work, we will often refer to this threshold of intractability for Bayesian CSPs as *the stable fixed point barrier* for reasons that will be soon clear.

Building on the ideas from [KZ09], [DKMZ11b, DKMZ11a] made a fascinating set of conjectures on community detection. For example, they conjectured that the k -coloring problem is easy exactly when the average degree of a vertex in the model satisfies $d > k^2$. Their conjectures fuelled a flurry of work, leading to algorithms that match the conjectured computational thresholds [MNS18, Mas14, BLM15, AS15].

The *stable fixed point barrier* suggested by [KZ09, DKMZ11a] is applicable beyond the setting of community detection. For instance, Krzakala and Zdeborova point out that this stable fixed point barrier is shared by problems such as hyper-graph bicoloring and locked CSPs. Here locked CSPs are those wherein every pair of assignments to a predicate have Hamming distance at least 2 (analogous to pairwise-independence leading to approximation resistance [AM09]). More broadly, there is a heuristic cavity method calculation to pinpoint the location of the stable fixed point barrier in general (see Section 2.4 to Section 2.7).

To illustrate the rich and precise predictions of this heuristic calculation yields, we will show three examples here.

Example 1.3. First, consider the problem of planted NAE3SAT wherein there is a uniformly random assignment in $\{0, 1\}^n$ and Not-All-Equal clauses on 3 variables are sampled so that a ρ -fraction of them are satisfied. As one varies the average constraint-degree of a variable d and the approximation ρ , there is an explicit prediction of the region of parameters where the distinguishing/recovery problem is computationally tractable (blue region in Figure 1). Interestingly, the spectral (and basic SDP) refutation threshold for regular NAE3SAT was determined to be 13.5 in [DMO⁺19] and similar techniques point to the threshold being 12.5 for an “Erdős-Rényi” version of the model. However, our results imply a distinguishing algorithm between random and planted NAE3SAT at a much smaller degree of 4.5, which suggests this planting is not “quiet” and raises the question of what a quiet planted distribution is.




Figure 1: Easy region for planted NAE-3SAT shaded in blue. Average degree on x -axis, fraction of clauses satisfied on y -axis.

Example 1.4. Next, we turn our attention to mixed planted CSPs. For concreteness, we consider one particular example: planted NAE-(3, 5)-SAT. In this example, the variables are given a uniformly random assignment in $\{0, 1\}^n$ and Not-All-Equals clauses are sampled to be on 3 variables with probability p and on 5 variables with probability $1 - p$. As one varies the constraint-degree of a variable d and the proportion of NAE3SAT clauses p , we can plot a precise region of parameters where the distinguishing/recovery problem is computationally tractable (blue region in Figure 2).

Example 1.5. Consider the following version of 4-community stochastic block model with communities labeled $(0, 0)$, $(0, 1)$, $(1, 0)$ and $(1, 1)$ and 3 parameters d_0 , d_1 and d_2 . For a pair of vertices u and v from communities x and y we place an edge between u and v with probability $\frac{d_{\text{dist}(x,y)}}{n}$ where $\text{dist}(x, y)$ is the Hamming distance between x and y . For an additional twist, let us suppose that the first coordinate of the




Figure 2: Easy region for planted NAE-(3, 5)-SAT shaded in blue. Average NAE3-degree of vertex on x -axis, average NAE5-degree of vertex on y -axis.

community that every vertex belongs to is also revealed to the algorithm. What is the region of parameters d_0, d_1, d_2 for which an efficient algorithm can partially recover the second coordinate of the community labels? See Figure 3 for the hypothesized transition.




Figure 3: Easy regions for (d_0, d_1) for variety of settings of d_2 .

Unfortunately, we are still far from establishing the veracity of these heuristic predictions. For most of these problems, BP has not been proven to succeed in the blue region of parameters, nor is any other polynomial time algorithm known. There is no roadmap to establishing intractability of these problems when the parameters are chosen in the white region.

Our main result takes a step towards establishing these predictions by giving a spectral algorithm to partially recover the hidden variables whenever the parameters are in the blue region. Specifically, we devise a spectral algorithm that uses a linearization of BP, an approach that has been successfully carried out for the case of community detection in [KMM⁺13, BLM15].

1.2 Stable Fixed Point Barrier

Belief propagation (BP) aims to estimate the marginals of the hidden variables, in our case $c(v)$ for $v \in [n]$. To visualize BP, it will be useful to consider the bipartite graph \mathcal{H} with variables $[n]$ on one side and the factors (a.k.a. observations) E on the other. There is an edge between a variable v and an observation $e \in E$

if $v \in e$. The execution of BP is divided into rounds where in each round, the variable nodes send messages to factor nodes or vice versa.

Let $m^{v \rightarrow e}$ denote the message sent by a variable v to a factor $e \in E$ and let $m^{e \rightarrow v}$ denote the message from a factor $e \in E$ to a variable v . All messages exchanged are distributions over the domain $[q]$, i.e., $m^{v \rightarrow e} = (m_1^{v \rightarrow e}, \dots, m_q^{v \rightarrow e})$ and similarly $m^{e \rightarrow v} = (m_1^{e \rightarrow v}, \dots, m_q^{e \rightarrow v})$. Intuitively speaking, $m_c^{v \rightarrow e}$ is an estimate of the marginal probability that v is assigned the color c when the factor e is absent, and $m_c^{e \rightarrow v}$ is an estimate of the marginal probability that u has color c when all other factors involving u are absent.

The general schema of a BP algorithm is to start BP with some initialization of the messages

$$\{m^{v \rightarrow e}[0], m^{e \rightarrow v}[0]\}_{v \in [n], e \in E}$$

and iteratively update the messages as specified by the functions Y , until the messages stabilize into a fixed point, i.e., a set of messages $\{\hat{m}^{v \rightarrow e}, \hat{m}^{e \rightarrow v}\}$ so that,

$$\begin{aligned}\hat{m}^{v \rightarrow e} &= Y_{v \rightarrow e}(\{\hat{m}^{f \rightarrow v} \mid f \in \partial v \setminus e\}) \\ \hat{m}^{e \rightarrow v} &= Y_{e \rightarrow v}(\{\hat{m}^{u \rightarrow e} \mid u \in \partial e \setminus v\})\end{aligned}$$

There is a canonical starting point \bar{m} for the BP iterations where the messages $m^{e \rightarrow v}$ correspond to uniform distribution over the possible values $[q]$. Conjecturally, this canonical initialization \bar{m} plays a critical role in characterizing the computational complexity of inferring the hidden variables in model M . There appear to be three possible cases with regards to this canonical initialization.

Case 1: \bar{m} is not a fixed point Suppose \bar{m} is not a fixed point for the BP iteration over the model M , then BP iteration can be expected to make progress, thereby yielding a weak recovery of hidden variables.

In fact, we will present a self-contained algorithm that weakly-recovers the hidden coloring in this case. Formally, we will show the following in [Appendix B](#):

Lemma 1.6. *If \bar{m} is not a fixed point for the BP iteration on model M , then there is a polynomial time algorithm \mathcal{A} and an $\epsilon > 0$ such that*

1. *if $(E, \tau) \sim M$: \mathcal{A} outputs a coloring that beats the correlation random guessing achieves with the hidden coloring by ϵ ,*
2. *\mathcal{A} solves the M vs. M^* (the null distribution) distinguishing problem with high probability.*

In light of the above lemma, it is natural to restrict our attention to the case where \bar{m} is a fixed point for the BP iteration.

Case 2: \bar{m} is an unstable fixed point \bar{m} is an *unstable fixed point* if arbitrarily small perturbations of \bar{m} will lead to the BP iteration moving away from the fixed point \bar{m} . This case was marked by the blue region in [Figure 1](#) and [Figure 3](#). In this case, our main algorithmic result is a spectral algorithm to recover a coloring c' that beats the correlation random guessing achieves with the hidden coloring. Alternatively, the spectral algorithm can be used to distinguish between the planted and the null distributions.

Case 3: \bar{m} is a stable fixed point \bar{m} is a *stable fixed point* if there exists a neighborhood U around \bar{m} such that for any initialization $\hat{m} \in U$, BP iteration converges to the canonical fixed point \bar{m} . In this case, the canonical fixed point \bar{m} clearly highlights a potential failure of BP algorithm. The hypothesis of Krzakala and Zdeborova [[KZ09](#)] asserts that existence of this stable fixed point marks the onset of computational intractability in general.

1.3 Related Work

Ideas from statistical physics have long been brought to bear on inference problems. We refer the reader to [Nis01, MM09b, ZK16b, RTSZ19] for an introduction to the phase transitions that mark changes in statistical and computational properties of these problems.

Planted models Special cases of the planted model we consider have appeared extensively in literature. The conditional probability of the hidden vector given the noisy observations takes the form of a graphical model, i.e. factorizes according to an hypergraph whose nodes correspond to variables and hyperedges correspond to noisy observations. Such graphical models have been studied by many authors in machine learning [LMP01] under the name of *conditional random fields*. We highlight a few among the extensive body of literature on information-theoretic and structural properties of these planted models. Montanari [Mon08] characterized the posterior marginals in terms of fixed points of the associated density evolution operator. Subsequently, Abbe and Montanari [AM13] show concentration for the conditional entropy per hidden variable given the observations. More recently, Coja-Oghlan et. al. [COHKL⁺20] study the information theoretic limits to recovery and confirm a conjectured formula for the mutual information between the observations and the planted assignment.

Spectral algorithms via non-backtracking operator The idea of using the spectra of non-backtracking matrix for recovery in planted problems can be traced back to the seminal work of Krzakala et al. [KMM⁺13] in the context of community detection. While this work provided heuristic arguments supporting the correctness of the algorithm, it was rigorously established in the work of Bordenave et. al. [BLM15]. Subsequently, [SLKZ15] devised spectral algorithms for solving the recovery problem in the censored block model, a variant of community detection wherein the edges are weighted and the weights carry the information about the community labels, but the edges don't. Building on the result of [BLM15], this work shows that the eigenvectors of non-backtracking matrix can be used to partially recover the communities, right up to the threshold. Finally, Angelini et al. [ACKZ15] consider a model of sparse hypergraphs that includes planted CSPs as a special case. The paper proposes a spectral algorithm based on a generalization of a non-backtracking matrix to hypergraphs, and gives a heuristic argument that the algorithm solves detection whenever belief propagation succeeds. Unlike our work, the algorithm proposed in [ACKZ15] uses an unweighted non-backtracking matrix that is independent of the prior probabilities. While it is a desirable feature that the algorithm is *non-parametric*, i.e., does not rely on the knowledge of prior distributions generating the instance, it is unclear if such a non-parametric algorithm can achieve detection up to the threshold in general.

Apart from recovery in planted models, the non-backtracking operator and the closely related Bethe-Hessian matrix have also been applied towards computing upper bounds for the log-partition function in ferro-magnetic Ising models on general graphs [SKZ17].

Quiet Planting Planted distributions that are indistinguishable from their random counterparts are often referred to as "quiet planting", though the terminology is not often consistent on whether the distributions are computationally or statistically indistinguishable.

A quiet planting that is statistically indistinguishable from random was used as a technical tool to study random instances in [AC08]. Krzakala and Zdeborova [KZ09] studied the existence of quiet plantings for graph coloring problem and were the first to hypothesize that under the Kesten-Stigum threshold, the planted ensembles are a computationally indistinguishable from random. Subsequently, the authors

[ZK11] considered planted distributions for *locked CSPs*, wherein every pair of assignments to a predicate have Hamming distance at least 2 and showed that problem is easy above a threshold that coincides with the Kesten-Stigum threshold and hypothesize that non-trivial recovery is computationally hard under it. Finally, a statistically quiet planting for the random k -SAT problem has been proposed in [KMZ12].

Community Detection Extensive work on community detection for stochastic block models has led to the confirmation [MNS18, Mas14, BLM15, AS15] of conjectures of Decelle et. al. [DKMZ11b, DKMZ11a]. As predicted, existing algorithms [Mas14, BLM15, AS15] can partially recover community labels up to the Kesten-Stigum threshold, but no lower. For $q = 2$ communities, the Kesten-Stigum threshold also matches the information theoretic threshold beyond which recovery is impossible. However, for $q \geq 3$ communities, the problem is believed to exhibit a statistical-vs-computational gap, in that there is a range of parameters where partial recovery is possible but is computationally intractable. The presence of a gap between the Kesten-Stigum threshold and the information theoretic threshold for all $q > 5$ was established in [Sly09]. More recently, Ricci-Tersenghi et al. [RTSZ19] predicted the existence of such a gap for $q = 4$ communities for some degree distributions, and also identifies a threshold beyond which there is a hard phase in asymmetric SBM. Furthermore, this work predicts the existence of hybrid-hard phases where it is computationally easy to reach a non-trivial inference accuracy, but computationally hard to match the information theoretically optimal one. Specifically, there are stable fixed points for BP that are not the trivial fixed point, but also don't correspond to optimal recovery.

Spectral norm bounds Technically, our work draws on ideas from Bordenave, Lelarge and Massoulié [BLM15] who established spectral norm bounds for non-backtracking matrices associated with Erdős-Renyi random graphs. Closer to our own setup, Stephan et al. [SM20] show eigenvalue bounds for the non-backtracking matrices of random graphs that have independent and bounded edge weights, and bounded model complexity (measured by the rank of the expected adjacency matrix). However, in our model the edges have correlated matrix weights instead of independent scalar weights, so their eigenvalue bounds do not generally apply to our model. Another work we draw several ideas from is that of Bordenave and Collins [BC19], who prove that the spectra of a wide family of random graphs, namely those arising from matrix-weighted noncommutative polynomials of random permutation matrices (see [OW20] for a comprehensive characterization and examples in this family), are roughly contained within the spectrum of an appropriately defined infinite graph. The key techniques useful in our work are the ones they employ to bound the spectral norms of the non-backtracking matrices of random regular graphs whose the edges are endowed with varying matrix weights.

1.4 Technical Overview

We define a general model for sparse observations on a hidden vector, and carry out the cavity method calculations in full generality following [DKMZ11a]. We obtain a criterion for computational tractability of the recovery and detection problems on this model, and provide spectral algorithms for recovery ([Theorem 2.10](#)) and detection ([Theorem 2.11](#)) in the tractable regime. The key technical ingredient in our work is tight eigenvalue bounds for nonbacktracking matrices of sparse random hypergraphs with (possibly varying) matrix-valued edge weights ([Theorem 3.8](#)).

In this section, we will attempt a brief technical outline of our result specialized to the case of distinguishing a random NAE3SAT instance from one with a hidden satisfying assignment. Concretely, consider

the problem distinguishing $\mathcal{D}_{\text{null}}$ from $\mathcal{D}_{\text{planted}}$ where:¹

- An instance $\mathcal{I} \sim \mathcal{D}_{\text{null}}$ is obtained by sampling each triple of distinct vertices (u, v, w) in $[n]^3$ independently with probability $\frac{d}{3!n^2}$ and then placing uniformly random negations $(\sigma_u, \sigma_v, \sigma_w)$ on each variable.
- An instance $\mathcal{I} \sim \mathcal{D}_{\text{planted}}$ is sampled in a two-stage process: (1) sample a hidden assignment $\mathbf{x} \sim \{\pm 1\}^n$, (2) sample each triple of distinct vertices (u, v, w) in $[n]^3$ independently with probability $\frac{d}{3!n^2}$ and place uniformly random negations $(\sigma_u, \sigma_v, \sigma_w)$ *conditioned* on $\text{NotAllEquals}(\sigma_u \mathbf{x}_u, \sigma_v \mathbf{x}_v, \sigma_w \mathbf{x}_w) = 1$.

First, let us map out the statistical physics prediction of the smallest value of d at which the problem becomes computationally tractable. In particular, we need to work out the value of d for which the trivial fixed point for belief propagation is unstable. To this end, one emulates the cavity method heuristic calculations analogous to the one carried out in [DKMZ11a] for stochastic block models. Oversimplifying for the sake of presentation, the cavity method heuristic amounts to carrying out the calculation by treating the neighborhood of each variable to be an infinite tree (see Section 4 for more details).

Concretely, the setup in the cavity method calculation is as follows. The neighborhood of a variable v in the NAE3SAT instance is modelled as an infinite tree with alternating layers of variable and NAE3SAT constraint nodes. The tree is generated by a Galton-Watson process where each variable v picks a degree $d_w \sim \text{Poisson}(d)$ from the Poisson distribution, and has d_w NAE3SAT constraint nodes as children, and each constraint node has exactly 2 children. For each path of length 2, $u \rightarrow C \rightarrow w$ from a variable u to its constraint node C followed by another variable w in the constraint, there is an associated constant sized matrix M_{uCw} depending on the prior distribution. For the case of NAE3SAT, all of the matrices M_{uCw} are given by $M_{uCw} = \sigma_u \sigma_w M$ where

$$M := \begin{bmatrix} -1/6 & 1/6 \\ 1/6 & -1/6 \end{bmatrix}.$$

For any depth t , consider the following quantity ρ_t where the expectation is over the choice of the infinite tree \mathcal{T} ,

$$\rho_t(d) = \mathbf{E}_{\text{tree } \mathcal{T}} \left[\sum_{\text{paths } u_0=v \rightarrow C_0 \rightarrow u_1 \rightarrow C_1 \rightarrow u_2 \rightarrow \dots \rightarrow u_t} \text{Tr} \left(\left(\prod_{i=0}^{t-1} M_{u_i C_i u_{i+1}} \right) \left(\prod_{i=0}^{t-1} M_{u_i C_i u_{i+1}} \right)^* \right) \right]$$

The threshold d^* predicted by the cavity method is precisely the smallest value of d for which $\lim_{t \rightarrow \infty} \rho_t(d) > 1$.

This characterization of d^* is a little unwieldy in that it is not immediate that the value of the threshold d^* is decidable. Fortunately, through ideas from the work of Bordenave and Collins [BC19], the above characterization can be equivalently written in terms of the spectral radius of an associated finite matrix. Specifically, for NAE3SAT, d^* is the smallest d for which the spectral radius of L exceeds 1 where:

$$L = d \cdot \begin{bmatrix} 1/18 & -1/18 & -1/18 & 1/18 \\ -1/18 & 1/18 & 1/18 & -1/18 \\ -1/18 & 1/18 & 1/18 & -1/18 \\ 1/18 & -1/18 & -1/18 & 1/18 \end{bmatrix}.$$

¹Strictly speaking, this model and the distribution over NAE3SAT instances our generic model yields differ slightly. Nevertheless they are contiguous and so the phenomena in one carry to the other.

(see [Section 2.9](#) for an overview of how to construct L in general, and [Section 4](#) for details). Hence, for NAE3SAT, the problem is hypothesized to become algorithmically tractable once $d > 4.5$. Our main results are algorithms for distinguishing the null and planted distributions, and for partially recovering a hidden assignment in the general model we consider whenever the spectral radius $\rho(L)$ of the matrix L corresponding to the model exceeds 1. In the case of NAE3SAT, we prove:

Theorem 1.7. *When $d > 4.5$, given $\mathcal{I} \sim \mathcal{D}_{\text{null}}$ or $\mathcal{D}_{\text{planted}}$:*

1. *There is an efficient algorithm to distinguish $\mathcal{D}_{\text{null}}$ from $\mathcal{D}_{\text{planted}}$ with probability $1 - o(1)$.*
2. *There is an efficient algorithm to produce $\Theta(1)$ unit vectors V where $\langle v, x \rangle \geq \Omega(\sqrt{n})$ for some $v \in V$.*

We now describe the distinguishing algorithm, which is spectral in nature, and briefly survey the techniques to analyze the matrix involved. The matrix we employ is a power of the so-called non-backtracking matrix obtained by linearizing belief propagation. For each clause C and pair of variables u, v in the clause signed by σ_u, σ_v we define matrix $M_{uCv} := \sigma_u \sigma_v M$. The s -th nonbacktracking power matrix is a $n \times n$ block matrix where each block is 2×2 :

$$A^{(s)}[a, b] := \sum_{uC_1 u_1 C_2 u_2 \dots u_{s-1} C_s v} M_{uC_1 u_1} M_{u_1 C_2 u_2} \dots M_{u_{s-1} C_s v}.$$

The algorithm is then fairly simple:

- Let $s = \lfloor \sqrt{\log n} \rfloor$, and let κ be strictly between $\sqrt{\rho(L)}$ and $\rho(L)$.
- If $\|A^{(s)}\| < \kappa^s$ output $\mathcal{D}_{\text{null}}$, otherwise output $\mathcal{D}_{\text{planted}}$.

In order to prove that the algorithm is correct, there are two key technical steps: (1) to prove that in the planted model, the operator norm is large, (2) to prove that in the null model, the operator norm is bounded.

The key insight in proving (1) is that the large operator norm of $A^{(s)}$ arises from the hidden assignment to the planted instance \mathcal{I} itself. In particular, denoting $\mathbf{y} := \mathbf{x} \otimes \begin{bmatrix} 1 \\ -1 \end{bmatrix}$ we prove:

Lemma 1.8. *With probability $1 - o(1)$:*

$$\frac{\langle \mathbf{y}, A^{(s)} \mathbf{y} \rangle}{\|\mathbf{y}\|^2} \geq \Omega(\rho(L)^s).$$

This is proved in full generality in [Section 6](#).

The main technical difficulty is in proving (2) in the general model, as is done in [Section 7](#). We prove:

Lemma 1.9. *With probability $1 - o(1)$:*

$$\|A^{(s)}\| \leq \left((1 + o(1)) \sqrt{\rho(L)} \right)^s.$$

Our proof is largely inspired by the works of [[BLM15](#), [BC19](#), [SM20](#)]. On one hand, [[BLM15](#)] and [[SM20](#)] show tight eigenvalue bounds for the nonbacktracking matrices of sparse (possibly inhomogenous) Erdős-Rényi graphs with scalar edge weights. The proof exploits the commutativity of scalar products, i.e. the product of edge weights along a walk is invariant under reordering. However, the graphs we consider have matrix-valued weights, which in general don't commute under multiplication. Therefore the product of edge weights changes depending on the order of multiplication. On the other hand, random regular

graphs with matrix-valued edge weights is handled in the work of [BC19]. However, the proof in [BC19] heavily exploits the regularity of the model — each vertex has exactly d adjacent edges and these edges have the exact same set of matrix weights. This leads to every vertex having isomorphic neighborhoods, and simplifies the analysis, which does not occur in our setting due to the lack of regularity. Our situation is further complicated by the fact that due to hyperedges of size greater than 2 even the random matrix weights in different blocks are not independent, which introduces mild correlations. The proof follows the general framework of the trace method and gives a more fine-grained analysis for nonbacktracking walks based on their shapes.

While the spectral radius of non-backtracking powers $A^{(s)}$ serve as a distinguisher, recovering the hidden assignment from the eigenvectors is little more subtle. In particular, this requires proving a converse of Lemma 1.8 that the every vector v for which $\langle v, A^{(s)}v \rangle$ is large, is actually correlated with the planted assignment. Instead, we bypass this issue by collating information from eigenvectors of $A^{(s)}$ for a range of values of s (see Section 8 for details).

1.5 Discussion and Future Work

In this work, we have shown that for a very general class of planted problems, the problem is computationally tractable whenever the trivial fixed point is unstable. This establishes the algorithmic side of the predictions of Krzakala and Zdeborova [KZ09] for all these problems. Several compelling open questions remain, we list a few here.

Reductions. From the standpoint of average case complexity, the main open question is to establish or refute the stable fixed point barrier. Given that all Bayesian CSPs have a uniform onset of intractability as specified by the stable fixed point barrier, perhaps these problems are reducible to one another. Traditional CSPs are very amenable to reductions, it is compelling to see if there are reductions between Bayesian CSPs, and stable fixed point barrier can be obtained as a consequence of the intractability of a single Bayesian CSP. The main challenge here is in coming up with reductions between problems that are *distribution-preserving* and we speculate that the ideas in [BBH18, BB20], which are examples of recent successes in reductions between average case problems, might be useful.

Hardness evidence in restricted computational models. Evidence on the stable fixed point barrier would also be very interesting. [HS17] showed that an algorithm based on low-degree polynomials solves the distinguishing problem in community detection up to Kesten-Stigum threshold, and also proves matching hardness in that low-degree polynomials fail to solve the problem under the Kesten-Stigum threshold. Recent work introduced the local statistics SDP hierarchy [BMR19] and showed the same algorithmic result for this class of algorithms and proved a negative result for the degree-2 SOS version of this algorithm. It will be useful to show that low-degree method and local statistics SDP hierarchy fail to solve the detection problem in the general model we consider in the presence of a stable fixed point. It will also be interesting to see if conditional hardness results for the problem can be obtained in other models such as statistical query algorithms [FGR⁺17].

Another direction in the spirit of the recent work of [BBH⁺20] which establishes an equivalence between the predictions of statistical query algorithms and the low-degree polynomials method would be to formally establish the equivalence of the predictions for the stable fixed point based on the cavity method with the other restricted models of computation such as the ones mentioned above.

Goldreich’s PRG for 1-wise independent, balanced, local predicates Goldreich proposed a construction of pseudorandom generators from random CSPs with balanced local predicates [Gol00]. The generator mapping n variables to $\{0, 1\}^m$ is constructed as follows: let E_1, \dots, E_m be a randomly chosen set of constraints on n input variables, then the i -th bit of the output string indicates whether E_i is satisfied by the input or not.

The constraints in Goldreich’s generator can be sampled from the null distribution of the model that we study. Then on any input \mathbf{c} (analogous to the hidden variables in the model), the output of the generator together with the constraints can be viewed as observations from the model’s planted distribution. Roughly speaking, we say that this generator produces pseudorandom strings if and only if the detection problem for this model is intractable.

For any random CSP with 1-wise independent predicates, the cavity method yields a concrete predicate density threshold above which the detection problem should be tractable (indeed the threshold is always of order $O(n)$). Our distinguishing algorithm confirms tractability in this regime, and therefore provides a concrete linear upper bound on the stretch of the Goldreich’s PRG constructed from the random CSP. Indeed, the upper bounds would be tight if the stable fixed point barrier hypothesis holds.

NP problem If the stable fixed point barrier hypothesis holds, then these Bayesian CSPs are excellent examples of average-case hard problems that are easy to sample. Their intractability can be harnessed to build cryptographic and pseudorandom primitives whose security depends on the existence of average-case hard problems. To this end, it is important that the underlying intractable problem is in NP, i.e., given the true hidden assignment, an efficient algorithm must be able to recognize it. Formally, this motivates the following NP-version of the problem:

Problem 1.10. (NP version) Devise an efficient verification algorithm \mathcal{A} that, given observations E from the planted model M and a candidate assignment $\mathbf{c} : [n] \rightarrow [q]$, has the following property:

- If (\mathbf{c}, E) are generated from the model M , the algorithm \mathcal{A} accepts (\mathbf{c}, E) with high probability.
- If the observations E are generated from the null model M , then for every assignment $\mathbf{c} : [n] \rightarrow [q]$, the algorithm rejects (\mathbf{c}, E) with high probability.

Dense models The focus of this paper has been the sparse settings, where in the underlying variable-observation graph is constant degree on average. Stability of trivial fixed point is also hypothesized to indicate computational intractability in dense problems such as spiked Wigner matrix (see [MV17] for some rigorous results). In this setting, it is the stability of fixed points of the *approximate message passing* (AMP) algorithm. A natural open question is whether the spectral algorithm based on linearizing AMP can be shown to generically hold in the dense setting.

Optimal Recovery Finally, in the region where weak recovery is possible, BP is conjectured to achieve the optimal recovery rate, i.e., achieve the maximum possible correlation with the hidden communities. While spectral algorithms provably achieve weak recovery, there has only been partial progress on the problem of achieving the optimal recovery rate [MNS14] — in particular, optimal recovery even in the 2-community block model close to the Kesten-Stigum threshold is open. In analogy with traditional CSPs, stable fixed point barrier marks the onset of “*approximation resistance*” for some problems, while the recovery rate corresponds to the approximation ratio.

2 Preliminaries

2.1 Observation Model

We will now formally define the *observation model* that is used throughout this work. The basic setup consists of a set of hidden variables $c(1), \dots, c(n)$ taking values over a finite domain $[q] = \{1, \dots, q\}$. Borrowing terminology from the PLANTED COLORING problem, we will refer to $[n] = \{1, \dots, n\}$ as the set of variables, $[q]$ as the set of colors and $\mathbf{c} : [n] \rightarrow [q]$ as the *hidden coloring*.

The hidden coloring $\mathbf{c} : [n] \rightarrow [q]$ is drawn from a prior distribution $\mathbb{P}_{\mathbf{c}}$. A sequence of hyperedges E on the vertex set $[n]$ are drawn, and we will refer to these hyperedges as *observations*. More precisely, an observation is effectively a hyperedge $e \in E$ with a type $\tau(e)$.

Definition 2.1. An observation model $M = ([q], T, \mathcal{T}, \{\mathbb{P}_{\tau}\}_{\tau \in T}, \Phi)$ describes a distribution on n -vertex hypergraphs M_n for every $n \geq 1$ and is specified by,

- **(Variable Types T).** A set of types T for the hidden variables and a distribution \mathcal{T} over them. Each variable is assigned a random type sampled from \mathcal{T} and is described by $\tau : [n] \rightarrow T$; in aggregate there are $\approx \mathcal{T}(\tau) \cdot n$ variables of type τ .

- **(Prior Distributions $\{\mathbb{P}_{\tau}\}_{\tau \in T}$).** For each variable of type $\tau \in T$, a prior distribution \mathbb{P}_{τ} .

The prior distribution of hidden coloring $\mathbf{c} : [n] \rightarrow [q]$ is the product distribution,

$$\mathbb{P}_{\mathbf{c}} = \mathbb{P}_{\tau(1)} \times \mathbb{P}_{\tau(2)} \dots \times \mathbb{P}_{\tau(n)}$$

- **(Observation Types Φ).** Set of observation types $\Phi = \{\phi_1, \dots, \phi_F\}$. The arity of a type i observation is denoted by $a(i)$.

Each observation on the variables is a hyperedge with a type from Φ . Specifically, the set of all observations is a set of hyperedges E partitioned as $E = \cup_{i \in [F]} E_i$ where E_i is a set of $a(i)$ -tuples of distinct elements in $[n]$.

- **(Observation Distributions).** For each observation type $\phi_i \in \Phi$, we have a bounded function $\phi_i : T^{a(i)} \times [q]^{a(i)} \rightarrow \mathbb{R}^+$.

For every $a(i)$ -tuple $(v_1, \dots, v_{a(i)})$ of distinct elements in $[n]$, the observation $\phi_i(v_1, \dots, v_{a(i)})$ is included independently with probability

$$\Pr [\phi_i(v_1, \dots, v_{a(i)}) \in E_i] \stackrel{\text{def}}{=} \frac{\phi_i((\tau(v_1), \mathbf{c}(v_1)), \dots, (\tau(v_{a(i)}), \mathbf{c}(v_{a(i)})))}{n^{a(i)-1}}$$

Notice that the probability of drawing an observation $\phi_i(v_1, \dots, v_{a(i)})$ depends both on the types of the variables and their colors.

We refer the reader to the work of Montanari [Mon08], where this model has been previously used for a wealth of concrete examples captured in this framework. Here we will exhibit a few examples.

Example 2.2. (Stochastic Block Model in semi-supervised setting)

In this variant of community detection, a graph $G = (V, E)$ is drawn from a $[q]$ -community SBM and in addition an α -fraction of the vertex labels are revealed. [ZMZ14] study the KS threshold in this model using the cavity method.

To encode this problem into our framework, we will have vertex types $T = [q] \cup \{\perp\}$ wherein the type of a vertex v is $\tau(v) \in [q]$ if the label of v is revealed, and $\tau(v) = \perp$ if it is unrevealed.

We have a single observation type namely the edges of the SBM, and the probability of an edge (u, v) is clearly $\phi((\tau(u), c(u)), (\tau(v), c(v)))/n$ for a function ϕ depending on the types and colors of two vertices.

More generally, the model can encode variants of SBM wherein there is additional attributes revealed about the vertices or edges or both. For example, SBM with labelled edges [HLM12] are subsumed by different types of observations, while SBM with vertex features [DSMM18] are captured by vertex types.

Further, the model can also be used to express geometric SBM [GMPS18] in restricted cases. In a geometric SBM, the vertices are distributed on a compact metric space like the sphere, and the probability of including an edge between vertices u, v is a function of the distance between the two. If the metric space is compact, say a sphere in a constant dimensional space, then one can use an ϵ -net of the compact set as a finite set of vertex types to model the SBM in our framework.

2.1.1 Miscellaneous simplifying notation

Class function Cl: For notational convenience, we will make a modification to our Definition 2.1 that does not affect the generality of our results. We will enforce that each observation type ϕ_i have a fixed tuple of variable types on which it applies. Formally, each observation type ϕ_i has an associated *class* type $\text{Cl}(i) \in T^{a(i)}$ such that all occurrences of the observation ϕ_i have input variable types given by $\text{Cl}(i)$. It is clear that this restriction is a special case of Definition 2.1 with the additional restriction that,

$$\phi_i((\tau_1, c_1), \dots, (\tau_{a(i)}, c_{a(i)})) = 0 \text{ if } (\tau_1, \dots, \tau_{a(i)}) \neq \text{Cl}(i)$$

Conversely, given a general model M as per Definition 2.1, for each observation type ϕ_i and each tuple $\tau = (\tau_1, \dots, \tau_{a(i)})$, introduce an observation type $\phi'_{i,\tau}$ that is identical to ϕ_i , but restricted to variable types τ , i.e., set $\text{Cl}(i) = \tau$. It is easy to see that this transformation creates a model M' that is equivalent to M . Without loss of generality we will henceforth use $\phi_i(c_1, \dots, c_{a(i)})$ to denote $\phi_i((\tau_1, c_1), \dots, (\tau_{a(i)}, c_i))$ where $\tau = \text{Cl}(i)$.

Average factor density: We will use $\overline{\phi}_i$ to denote the average density of a factor:

$$\overline{\phi}_i := \sum_{(c_1, \dots, c_{a(i)}) \in [q]^{a(i)}} \left(\prod_{k=1}^{a(i)} \mathbb{P}_{\text{Cl}(i)_k} \right) \cdot \phi_i(c_1, \dots, c_{a(i)}).$$

Bipartite view: Given a collection of sampled observations $E = \cup_{i=1}^F E_i$, we associate a bipartite graph G where the left vertex set is given by the variables $[n]$ and the right vertex set is given by the collection of all (i, γ) for γ in E_i .

Index function: For $e = (i, (v_1, \dots, v_{a(i)}))$ we define $i_e(v_s)$ as s and $e[s]$ as v_s . When e is clear from context we will drop the e and just use $i(v_s)$.

Definition 2.3. For an observation model $M = ([q], T, \mathcal{T}, \{\mathbb{P}_\tau\}_{\tau \in T}, \Phi)$ the corresponding *null model* M^* is the observation distribution where for every $a(i)$ -tuple $(v_1, \dots, v_{a(i)})$ a hidden coloring \mathbf{c} is sampled independently, and the observation $\phi_i(v_1, \dots, v_{a(i)})$ is included with probability

$$\Pr [\phi_i(v_1, \dots, v_{a(i)}) \in E_i] \stackrel{\text{def}}{=} \frac{\phi_i((\tau(v_1), \mathbf{c}(v_1)), \dots, (\tau(v_{a(i)}), \mathbf{c}(v_{a(i)})))}{n^{a(i)-1}}.$$

Equivalently, in the null model for every $(v_1, \dots, v_{a(i)})$ the observation $\phi_i(v_1, \dots, v_{a(i)})$ is included independently with probability:

$$\Pr [\phi_i(v_1, \dots, v_{a(i)}) \in E_i] \stackrel{\text{def}}{=} \frac{\bar{\phi}_i}{n^{a(i)-1}}.$$

Remark 2.4. For a model M we will refer to it as the *planted model* and we will refer to M^* as the *null model*. Two computational problems we are interested in are distinguishing whether a sample is drawn from M or M^* , and inferring the hidden coloring for a sample drawn from M .

2.2 Bayesian Inference

Given the variable types $\tau : [n] \rightarrow T$ and the observations \mathbf{E} , the canonical algorithm to infer the hidden coloring \mathbf{c} is to use the Bayes rule to compute the conditional distribution $\mathbb{P}_{\mathbf{c}|\mathbf{E}}$. Formally, the probability that a model $M = ([q], T, \mathcal{T}, \{\mathbb{P}_\tau\}_{\tau \in T}, \Phi)$ generates a hidden coloring \mathbf{c} and observations \mathbf{E} is

$$\begin{aligned} \Pr[\mathbf{E}_1, \dots, \mathbf{E}_F, \mathbf{c} \mid \boldsymbol{\tau}] &= \Pr[\mathbf{c} \mid \boldsymbol{\tau}] \cdot \Pr[\mathbf{E}_1, \dots, \mathbf{E}_F \mid \mathbf{c}, \boldsymbol{\tau}] \\ &= \left(\prod_{v \in [n]} \mathbb{P}_{\tau(v)}(\mathbf{c}(v)) \right) \cdot \\ &\quad \prod_{i \in [F]} \left(\prod_{(v_j)_{j \in [n]^{a(i)}}} \left[\frac{\phi_i(\mathbf{c}(v_1), \dots, \mathbf{c}(v_{a(i)}))}{n^{a(i)-1}} \right]^{\mathbf{1}_{(v_j)_{j \in [n]^{a(i)}} \in E_i}} \left[1 - \frac{\phi_i(\mathbf{c}(v_1), \dots, \mathbf{c}(v_{a(i)}))}{n^{a(i)-1}} \right]^{\mathbf{1}_{(v_j)_{j \in [n]^{a(i)}} \notin E_i}} \right). \end{aligned}$$

By applying Bayes rule,

$$\Pr[\mathbf{c} \mid \mathbf{E}_1, \dots, \mathbf{E}_F, \boldsymbol{\tau}] = \frac{\Pr[\mathbf{E}_1, \dots, \mathbf{E}_F, \mathbf{c} \mid \boldsymbol{\tau}]}{\sum_{\mathbf{c}^*} \Pr[(\mathbf{E}_1, \dots, \mathbf{E}_F, \mathbf{c}^* \mid \boldsymbol{\tau})]}.$$

Ignoring the normalizing constant, we can write

$$\Pr[\mathbf{c} \mid \mathbf{E}_1, \dots, \mathbf{E}_F, \boldsymbol{\tau}] \propto e^{-H(\mathbf{c}|\mathbf{E}_1, \dots, \mathbf{E}_F, \boldsymbol{\tau})},$$

where

$$\begin{aligned} \mathbf{H}(\mathbf{c} \mid \mathbf{E}_1, \dots, \mathbf{E}_F, \boldsymbol{\tau}) &= - \sum_{i \in [F]} \sum_{(v_j)_{j \in [n]^{a(i)}}} \left[\mathbf{1}_{(v_j)_{j \in [n]^{a(i)}} \in E_i} \log \left(\frac{\phi_i(\mathbf{c}(v_1), \dots, \mathbf{c}(v_{a(i)}))}{n^{a(i)-1}} \right) + \mathbf{1}_{(v_j)_{j \in [n]^{a(i)}} \notin E_i} \log \left(1 - \frac{\phi_i(\mathbf{c}(v_1), \dots, \mathbf{c}(v_{a(i)}))}{n^{a(i)-1}} \right) \right] \\ &\quad - \sum_{v \in [n]} \log \mathbb{P}_{\tau(v)}(\mathbf{c}(v)). \end{aligned}$$

The function $\mathbf{H}(\mathbf{c}|\mathbf{E}_1, \dots, \mathbf{E}_F, \boldsymbol{\tau})$ is referred to as the Hamiltonian, and the distribution is the Boltzmann distribution with Hamiltonian \mathbf{H} and inverse temperature $\beta = 1$.

Since in our setting, the hypergraph is sparse, i.e., $\bar{\phi}_i = O(1)$, the terms

$$\log \left(1 - \frac{\phi_i(\mathbf{c}(v_1), \dots, \mathbf{c}(v_{a(i)}))}{n^{a(i)-1}} \right) \approx 0 \tag{1}$$

for all $(v_j)_j \notin E_i$. So these terms can be dropped to simplify the Hamiltonian to

$$\mathbf{H}(\mathbf{c} \mid \mathbf{E}_1, \dots, \mathbf{E}_F, \boldsymbol{\tau}) = - \sum_{i \in [F]} \sum_{(v_j)_j \in E_i} \log(\phi_i(\mathbf{c}(v_1), \dots, \mathbf{c}(v_{a(i)}))) - \sum_{v \in [n]} \log \mathbb{P}_{\boldsymbol{\tau}(v)}(\mathbf{c}(v)). \quad (2)$$

The Hamiltonian \mathbf{H} is a sum of *local* terms each depending on a constant number of variables. The observations \mathbf{E} and the variables \mathbf{c} together form what is termed as *factor graphs* (see [MM09a]), where each observation is a *factor* of the Boltzmann distribution. Recall that the Boltzmann distribution is given by

$$\Pr[\mathbf{c} \mid (\mathbf{E}_1, \dots, \mathbf{E}_F, \boldsymbol{\tau})] = \frac{e^{-\mathbf{H}(\mathbf{c} \mid \mathbf{E}_1, \dots, \mathbf{E}_F, \boldsymbol{\tau})}}{\sum_{\mathbf{c}'} e^{-\mathbf{H}(\mathbf{c}' \mid \mathbf{E}_1, \dots, \mathbf{E}_F, \boldsymbol{\tau})}}$$

The normalization term in the denominator is called the partition function of the distribution and is denoted $Z(\mathbf{M})$. Notice that a naive algorithm to infer the hidden coloring via the Bayes rule as described above would take exponential time.

2.3 Belief Propagation

The algorithm of choice to infer the hidden variables in a sparse factor model would be belief propagation. We refer the reader to [MM09a] for a detailed exposition of belief propagation, and restrict ourselves to a broad outline.

Belief propagation (BP) aims to estimate the marginals of the hidden variables, in our case $\mathbf{c}(v)$ for $v \in [n]$. BP draws its inspiration from a dynamic programming algorithm to compute the marginals when the underlying factor graph is a tree, and is broadly applicable to sparse settings where the local neighborhood of a vertex is tree-like. In particular, while BP computes the marginals exactly on a tree, it is very successful in practice over sparse factor models that are locally tree-like.

To visualize BP, it will be useful to consider the bipartite graph \mathcal{H} with variables $[n]$ on one side and the factors (a.k.a. observations) \mathbf{E} on the other. There is an edge between a variable v and an observation $e \in \mathbf{E}$ if $v \in e$. The execution of BP is divided into rounds where in each round, the variable nodes send messages to factor nodes or vice versa.

Let $m^{v \rightarrow e}$ denote the message sent by a variable v to a factor $e \in \mathbf{E}$ and let $m^{e \rightarrow v}$ denote the message from a factor $e \in \mathbf{E}$ to a variable v . All messages exchanged are marginal distributions over the domain $[q]$, i.e., $m^{v \rightarrow e} = (m_1^{v \rightarrow e}, \dots, m_q^{v \rightarrow e})$ and similarly $m^{e \rightarrow v} = (m_1^{e \rightarrow v}, \dots, m_q^{e \rightarrow v})$. Intuitively speaking, $m_c^{v \rightarrow e}$ is an estimate of the marginal probability that v is assigned the color c when the factor e is absent, and $m_c^{e \rightarrow v}$ is an estimate of the marginal probability that u has color c when all other factors involving u are absent.

BP specifies an update rule for every variable/factor node to update its outgoing messages each round, depending on its incoming messages. Let ∂e denote the set of variables incident a factor e and let ∂v denote the set of factors incident on a variable v . BP specifies functions $Y_{v \rightarrow e}, Y_{e \rightarrow v}$ so that if $\{m^{v \rightarrow e}[t], m^{e \rightarrow v}[t]\}$ denote the messages in round t , then the updated messages are given by

$$m^{v \rightarrow e}[t+1] = Y_{v \rightarrow e}(\{m^{f \rightarrow v}[t] \mid f \in \partial v \setminus e\}) \quad (3)$$

$$m^{e \rightarrow v}[t+1] = Y_{e \rightarrow v}(\{m^{u \rightarrow e}[t] \mid u \in \partial e \setminus v\}) \quad (4)$$

We will describe the specific form of the functions Y in [Appendix A](#), but there are two salient details that we would like to highlight at this time. First, the functions Y are smooth rational functions that map marginals over $[q]$ to a marginal distribution over $[q]$. Second, the updated outgoing message $m^{v \rightarrow e}[t+1]$

depends on all messages incoming to variable v **except** the message $m^{e \rightarrow v}[t]$. Similarly, the updated outgoing message $m^{e \rightarrow v}[t+1]$ is independent of the incoming message $m^{v \rightarrow e}[t]$.

The general schema of a BP algorithm is to start BP with some initialization of the messages

$$\{m^{v \rightarrow e}[0], m^{e \rightarrow v}[0]\}_{v \in [n], e \in E}$$

and iteratively update the messages as specified by the functions Y , until the messages stabilize into a fixed point, i.e., a set of messages $\{\hat{m}^{v \rightarrow e}, \hat{m}^{e \rightarrow v}\}$ so that,

$$\begin{aligned}\hat{m}^{v \rightarrow e} &= Y_{v \rightarrow e}(\{\hat{m}^{f \rightarrow v} \mid f \in \partial v \setminus e\}) \\ \hat{m}^{e \rightarrow v} &= Y_{e \rightarrow v}(\{\hat{m}^{u \rightarrow e} \mid u \in \partial e \setminus v\})\end{aligned}$$

While it can often be difficult at times to show convergence to a fixed point, BP is very successful in practice over locally tree-like factor models.

2.4 Stable Fixed Point Barrier

A natural starting point for BP iteration for a model M is given by the following:

$$\bar{m}^{v \rightarrow e} \stackrel{\text{def}}{=} \text{prior distribution } P_{\tau(v)} \tag{5}$$

$$\bar{m}^{e \rightarrow v} \stackrel{\text{def}}{=} \text{uniform distribution over support of } P_{\tau(v)} \tag{6}$$

Conjecturally, this canonical initialization \bar{m} plays a critical role in characterizing the computational complexity of inferring the hidden variables in model M .

There appear to be three possible cases with regards to this canonical initialization.

Case 1: \bar{m} is not a fixed point Suppose \bar{m} is not a fixed point for the BP iteration over the model M , then BP iteration can be expected to make progress, thereby yielding a weak recovery of hidden variables.

In fact, we will present a self-contained algorithm that weakly-recovers the hidden coloring in this case. Formally, we will show the following in [Appendix B](#):

Lemma 2.5. *If \bar{m} is not a fixed point for the BP iteration on model M , then there is a polynomial time algorithm \mathcal{A} and an $\epsilon > 0$ such that*

1. *if $(E, \tau) \sim M$: \mathcal{A} outputs a coloring that beats the correlation random guessing achieves with the hidden coloring by ϵ ,*
2. *\mathcal{A} solves the M vs. M^* distinguishing problem with high probability.*

In light of the above lemma, it is natural to restrict our attention to the case where \bar{m} is a fixed point for the BP iteration. \bar{m} being a fixed point of BP is equivalent to a “detailed balance” condition holding (in the sense of [\(45\)](#)).

Case 2: \bar{m} is an unstable fixed point \bar{m} is an *unstable fixed point* if arbitrary small perturbations of \bar{m} will lead to the BP iteration moving away from the fixed point \bar{m} . BP is conjectured to succeed in weak-recovery of hidden coloring and distinguishing between M vs. M^* in this case, and this has been extensively demonstrated experimentally [[DKMZ11a](#), [ZMZ14](#)].

Case 3: \bar{m} is a stable fixed point \bar{m} is a *stable fixed point* if there exists a neighborhood U around \bar{m} such that for any initialization $\hat{m} \in U$, BP iteration converges to the canonical fixed point \bar{m} . In this case, the canonical fixed point \bar{m} clearly highlights a potential failure of BP algorithm. A priori, it is conceivable that by using BP with an alternative starting point or an entirely different algorithm, one could still efficiently infer the hidden coloring in this case.

Surprisingly, it is conjectured that the existence of this canonical fixed point that is stable marks the onset of computational intractability! Inspired by ideas from statistical physics, Krzakala and Zdeborova [KZ09] were the first to hypothesize that the existence of a trivial fixed point that is stable marks computational intractability. Building on these intuitions, Decelle et. al. [DKMZ11a] outlined a fascinating set of conjectures on community detection problem which fuelled a flurry of activity, resulting in algorithms matching the conjectured computational thresholds [MNS18, Mas14, BLM15, AS15].

2.5 Analyzing Stability

The stability of the canonical fixed point \bar{m} under BP iteration can be analyzed using derivatives of the BP update rule. Suppose Γ denote the map associated with running two rounds of BP iteration to produce the messages, i.e.,

$$\{m^{v \rightarrow e}[t+2]\}_{v \in [n], e \ni v} = \Gamma(\{m^{v \rightarrow e}[t]\}_{v \in [n], e \ni v})$$

In other words, Γ is given by the composition of the functions in (3) and (4). If \bar{m} is a fixed point of BP, then we will have,

$$\Gamma(\{\bar{m}^{v \rightarrow e}\}) = \{\bar{m}^{v \rightarrow e}\}$$

To analyze the stability of the fixed point \bar{m} , one uses the linear approximation of Γ in a neighborhood of \bar{m} , by setting

$$\Gamma(\bar{m} + \epsilon) = \bar{m} + B\epsilon$$

where B is the matrix of partial derivatives, i.e.,

$$B[m^{u \rightarrow e}, m^{u' \rightarrow e'}] = \frac{\partial \Gamma(m)^{u' \rightarrow e'}}{\partial m^{u \rightarrow e}} \Big|_{\bar{m}}$$

With this linear approximation $\Gamma^\ell(\bar{m} + \epsilon) \approx \bar{m} + B^\ell \epsilon$. Therefore, the stability of the fixed point is characterized by the spectral radius of the operator B .

Specifically, \bar{m} is a stable fixed point if and only if $\rho(B) \leq 1$ where $\rho(B) \stackrel{\text{def}}{=} \max_i |\lambda_i(B)|$ is the largest magnitude of an eigenvalue of B .

Notice that B is an asymmetric random matrix depending on the set of observations E . The *cavity method* is a heuristic to guess the spectral radius of a typical derivative matrix B in terms of the spectral radius of some constant sized linear operator L . In the rest of the section we first use the cavity method to obtain a precise condition on M for $\rho(B) \leq 1$, then state our main theorem that the distinguishing problem and the weak recovery problem are efficiently solvable when $\rho(B) > 1$, and finally define the operator L whose spectral bound λ_L satisfies that $\rho(B) = \lambda_L^{1/2}$.

2.6 The local distributions of M

Before diving into the calculation, we define a few local distributions of M that would be used later.

The color assignment distribution μ_i For each factor $\phi_i \in \Phi$, define a local distribution μ_i over $[q]^{a(i)}$ as,

$$\mu_i(c_1, \dots, c_{a(i)}) \propto \left(\prod_{j \in [a(i)]} \mathbb{P}_{\text{Cl}(i)_j}(c_j) \right) \cdot \phi_i(\mathbf{c}) \quad (7)$$

For each $\phi_i \in \Phi$ and $a, b \in [a(i)]$ define a matrix $\Psi_{i,a|b} \in \mathbb{R}^{[q] \times [q]}$ by fixing,

$$\Psi_{i,a|b}(\alpha, \beta) \stackrel{\text{def}}{=} \Pr_{(c_1, \dots, c_{a(i)}) \sim \mu_i} [c_a = \alpha | c_b = \beta] \quad (8)$$

This says that conditioned on that ϕ_i is in the observations \mathbf{E} , the matrix $\Psi_{i,a|b}$ encodes the color distribution of a conditioned on the color of b . Finally for $\phi_i \in \Phi$ and $a, b \in [a(i)]$ we define a matrix that is useful later,

$$\bar{\mathbf{M}}_{i,a|b} = (\mathbf{I} - \mathbb{P}_{\text{Cl}(i)_a} \mathbf{1}^T) \Psi_{i,a|b}. \quad (9)$$

The neighbor factor distribution of a variable We now take a closer look at a type τ variable's neighbor factor distribution. Here a variable's neighbor factors refer to all factors that are connected to the variable in the factor graph.

To study this neighborhood distribution, we first define random variables $\text{deg}_{i,j}(\tau)$ for a type τ variable v .

Definition 2.6. For $\tau \in T$, $\phi_i \in \Phi$, $\text{deg}_{i,j}(\tau)$ is the random variable denoting the number of type ϕ_i factors in the neighborhood of the type τ variable v such that the index of v in all these factor is j .

From the definition, we see that each $\text{deg}_{i,j}(\tau)$ is the sum of many binomial variables each of which indicates whether a specific type ϕ_i factor exists in the factor graph. We formally define these binomial variables.

Definition 2.7. For $\tau \in T$, $(v_1, \dots, v_{a(i)}) \in [n]^{a(i)}$, $b_\tau^{v_1, \dots, v_{a(i)}}$ is the indicator variable of whether the type ϕ_i factor e whose j -th variable is v_j for all $j \in [a(i)]$ is in the observations \mathbf{E} .

We can compute the probability of $b_\tau^{v_1, \dots, v_{a(i)}} = 1$ in M_n .

$$\begin{aligned} \Pr_{M_n} [b_\tau^{v_1, \dots, v_{a(i)}} = 1] &= \frac{1}{\mathcal{T}(\tau)} \cdot \prod_{j=1}^{a(i)} \mathcal{T}(\text{Cl}(i)_j) \cdot \sum_{(c_1, \dots, c_{a(i)}) \in [q]^{a(i)}} \prod_{j=1}^{a(i)} \mathbb{P}_{\text{Cl}(i)_j}(c_j) \cdot \frac{\phi_i(c_1, \dots, c_{a(i)})}{n^{a(i)-1}} \\ &= \frac{1}{\mathcal{T}(\tau)} \cdot \frac{\bar{\phi}_i}{n^{a(i)-1}} \cdot \prod_{j=1}^{a(i)} \mathcal{T}(\text{Cl}(i)_j) \end{aligned}$$

Thus $b_\tau^{v_1, \dots, v_{a(i)}}$ has distribution Binomial $\left(\frac{\bar{\phi}_i}{\mathcal{T}(\tau) n^{a(i)-1}} \cdot \prod_{j=1}^{a(i)} \mathcal{T}(\text{Cl}(i)_j) \right)$.

Now we can express $\text{deg}_{i,j}(\tau)$ as the sum of $n^{a(i)-1}$ binomial random variables.

$$\text{deg}_{i,j}(\tau) = \sum_{(v_j) \in [n]^{a(i)} | v_j = v} b_\tau^{v_1, \dots, v_{a(i)}}.$$

We also note that most of the $b_\tau^{v_1, \dots, v_{a(i)}}$ s are independent. Two random variables $b_\tau^{v_1, \dots, v_{a(i)}}$ and $b_\tau^{v'_1, \dots, v'_{a(i)}}$ are not independent only if there exist $j, j' \in [a(i)]$ such that $v_j = v'_{j'}$ but $\text{Cl}(i)_j \neq \text{Cl}(i)_{j'}$. That is the two factors share some variable but require the variable to have different types. However, only $O(n^{-(a(i)-1)})$ fraction of the pairs are correlated. Thus, when n is large we can treat the $n^{a(i)-1}$ random variables as being independent. Then each $\text{deg}_{i,j}(\tau)$ has a Poisson distribution.

Claim 2.8. $\deg_{i,j}(\tau) \sim \text{Poisson}\left(\frac{\bar{\phi}_i}{\mathcal{T}(\tau)} \cdot \prod_{j=1}^{a(i)} \mathcal{T}(\text{Cl}(i)_j)\right)$.

For similar reason as above, when n is large we can treat the random variables $\{\deg_{i,j}(\tau)\}_{i \in [F], j \in [a(i)]}$ of a variable v as independent. Therefore for large n , the neighbor factor distribution of a type τ variable v is very close to the product distribution of the random variables $\{\deg_{i,j}(\tau)\}_{i \in F, j \in [a(i)]}$.

2.7 The stability condition

Now we continue to explore the condition on M that makes \bar{m} stable. We focus on sparse models whose average factor degrees $\bar{\phi}_i = O(1)$ for all $\phi_i \in \Phi$. In such models, a variable node is contained in constant number of factors, and its $o(\log n)$ -neighborhood is locally tree-like with high probability. Set the tree depth ℓ be a function such that $\ell(n) \in o(\log n)$, and consider the distance $(2\ell + 1)$ neighborhood of a variable v_0 . Assume each level- $(2\ell + 1)$ factor node e_ℓ 's outgoing message to some level- 2ℓ variable v_ℓ is perturbed to $m^{e_\ell \rightarrow v_\ell} = \bar{m}^{e_\ell \rightarrow v_\ell} + \epsilon_{e_\ell}$. Recall that $\bar{m}^{e_\ell \rightarrow v_\ell}$ is the trivial fixed point message, and ϵ_{e_ℓ} is the random perturbation that is independent across different edges $e_\ell \rightarrow v_\ell$. We want to compute the expected influence of the perturbations on the messages to the root v_0 .

We first consider the distance $2\ell + 1$ neighborhood of a variable v_0 . The treelike neighborhood can be constructed by the following process.

1. Sample the type of the root variable $\tau(v_0)$ from \mathcal{T} .
2. Sample the level-1 factors: for each type of factor $\phi_i \in \Phi$ sample the number of type ϕ_i neighbor factors of v_0 by sampling $\{\deg_{i,j}(\tau(v_0))\}_{i \in F, j \in [a(i)]}$ independently. Add v_0 's neighbor factors to level 1. Add the other variables in these factors to the next level of the tree, assuming that there is no shared variables other than v_0 . Note that these variables already have types.
3. Repeat step 2 for the new variables until we get a depth- $(2\ell + 1)$ tree \mathbb{T}_ℓ .

We next use the tree \mathbb{T}_ℓ to give a precise condition on M for the fixed point \bar{m} to be stable.

In a \mathbb{T}_ℓ , a leaf node e_ℓ is connected to the root node v_0 via a path $e_\ell, v_\ell, e_{\ell-1}, \dots, v_1, e_0, v_0$. A perturbation on the leaf message $m^{e_\ell \rightarrow v_\ell}$ influence the next level message $m^{e_{\ell-1} \rightarrow v_{\ell-1}}$ via the partial derivative matrix $\frac{\partial \Gamma(m)^{e_{\ell-1} \rightarrow v_{\ell-1}}}{\partial m^{e_\ell \rightarrow v_\ell}}$. We can express the partial derivative matrix evaluated at the fixed point using the matrix defined in (9).

Claim 2.9 (Claim 4.1).

$$\frac{\partial \Gamma(m)^{e_{\ell-1} \rightarrow v_{\ell-1}}}{\partial m^{e_\ell \rightarrow v_\ell}} \Big|_{\bar{m}} = \bar{M}_{\theta(e_{\ell-1}, i_{e_{\ell-1}}(v_{\ell-1}) | i_{e_{\ell-1}}(v_\ell))}.$$

This claim is proved in [Appendix C](#). When writing the matrix $\bar{M}_{\theta(e), i_e(v) | i_e(v')}$, it's clear that the index function is associated with the factor e , so we drop the subscript e in i_e for simplicity.

Using the chain rule, we can compose the partial derivative matrices along a path, and conclude that each path influences the root message by

$$\left(\prod_{j=1}^{\ell} \bar{M}_{\theta(e_{j-1}), i(v_{j-1}) | i(v_j)} \right) \epsilon_{e_\ell}.$$

Thus the influence of all paths in \mathbb{T}_ℓ is

$$\sum_{(e_\ell, v_\ell, \dots, e_0, v_0) \in \mathbb{T}_\ell} \left(\prod_{j=1}^{\ell} \bar{M}_{\theta(e_{j-1}), i(v_{j-1}) | i(v_j)} \right) \epsilon_{e_\ell}$$

To decide if the fixed point is stable, we compute the variance α_ℓ of this influence.

$$\begin{aligned}
& \mathbf{E}_{\mathbb{T}_\ell, \epsilon} \left[\left\| \sum_{(e_\ell, v_\ell, \dots, e_0, v_0) \in \mathbb{T}_\ell} \left(\prod_{j=1}^{\ell} \overline{\mathbf{M}}_{\theta(e_{j-1}, i(v_{j-1}) | i(v_j))} \right) \epsilon_{e_\ell} \right\|^2 \right] \\
&= \mathbf{E}_{\mathbb{T}_\ell, \epsilon} \left[\sum_{(e_\ell, v_\ell, \dots, e_0, v_0) \in \mathbb{T}_\ell} \epsilon_{e_\ell}^\top \left(\prod_{j=1}^{\ell} \overline{\mathbf{M}}_{\theta(e_{j-1}, i(v_{j-1}) | i(v_j))} \right)^* \left(\prod_{j=1}^{\ell} \overline{\mathbf{M}}_{\theta(e_{j-1}, i(v_{j-1}) | i(v_j))} \right) \epsilon_{e_\ell} \right] \\
&= \mathbf{E}_{\mathbb{T}_\ell} \left[\sum_{(e_\ell, v_\ell, \dots, e_0, v_0) \in \mathbb{T}_\ell} \text{Tr} \left(\left(\prod_{j=1}^{\ell} \overline{\mathbf{M}}_{\theta(e_{j-1}, i(v_{j-1}) | i(v_j))} \right) \left(\prod_{j=1}^{\ell} \overline{\mathbf{M}}_{\theta(e_{j-1}, i(v_{j-1}) | i(v_j))} \right)^* \right) \right] \cdot \frac{\mathbf{E}_{\epsilon_{e_\ell}} [\|\epsilon_{e_\ell}\|^2]}{q}
\end{aligned}$$

So the squared norm of the perturbations ϵ_{e_ℓ} is amplified by

$$\alpha_\ell \stackrel{\text{def}}{=} \mathbf{E}_{\mathbb{T}_\ell} \left[\sum_{(e_\ell, v_\ell, \dots, e_0, v_0) \in \mathbb{T}_\ell} \text{Tr} \left(\left(\prod_{j=1}^{\ell} \overline{\mathbf{M}}_{\theta(e_{j-1}, i(v_{j-1}) | i(v_j))} \right) \left(\prod_{j=1}^{\ell} \overline{\mathbf{M}}_{\theta(e_{j-1}, i(v_{j-1}) | i(v_j))} \right)^* \right) \right].$$

We note that for a model M , by definition of α_ℓ and the operator B in [Section 2.5](#), $\rho(B) = \lim_{n \rightarrow \infty} \alpha_\ell^{1/2\ell}$ (recall that ℓ is a function of n). Thus when $\lim_{n \rightarrow \infty} \alpha_\ell^{1/2\ell} \leq 1$, \overline{m} is a stable fixed point.

2.8 Efficient recovery and detection when the fixed point is unstable

For a model M whose fixed point \overline{m} is unstable, it is conjectured that the BP algorithm can successfully weak-recover the hidden coloring. We provide a BP-inspired spectral algorithm that solves the weak-recovery problem in this regime. We state the result somewhat informally below; the full formal statement can be found in [Theorem 8.1](#). However, before we state the result we go on a small digression on how to set the benchmark for weak recovery. A first attempt might be:

For a fixed type $\tau \in T$ and color $c \in [q]$, produce a vector $\overline{u} \in \mathbb{R}^n$ such that w correlates with the following ‘‘centered indicator vector’’ of (τ, c) : $\underline{\chi}^{\tau, c}$ where $\underline{\chi}^{\tau, \alpha}$ is an n -dimensional vector with i -th coordinate $\mathbf{1}[\tau(i) = \tau] \cdot (\mathbf{1}[c(i) = c] - \mathbb{P}_\tau(c))$.

However, this benchmark is unattainable since for a problem such as PLANTED- q -COLORING there is no way to statistically discriminate between a given coloring and a different coloring obtained by permuting the names of the colors. Thus, to account for this complication we consider the following modification of the above benchmark, which we first state in words.

Produce a vector $\overline{u} \in \mathbb{R}^n$ such that after some permutation is applied to the names of the colors, for some type $\tau \in T$ and color $c \in [q]$, \overline{u} correlates with the centered indicator vector of (τ, c) .

More formally:

Theorem 2.10. *If a model M has $\lim_{\ell \rightarrow \infty} \alpha_\ell^{1/2\ell} > 1$, there is a spectral algorithm A that solves the weak-recovery problem. A bit more concretely, for $G \sim M_n$, the algorithm $A(G)$ produces $O_M(1)$ vectors $\{\overline{u}_1, \dots, \overline{u}_r\}$ such that one of these vectors \overline{u}_j has constant correlation with the planted coloring in the following sense:*

There is a type $\tau \in T$, and a color $\alpha \in [q]$ such that if we construct $\underline{\chi}^{\tau, \alpha} \in \mathbb{R}^n$ as

$$\underline{\chi}^{\tau, \alpha}[i] = \mathbf{1}[\tau(i) = \tau](\mathbf{1}[c(i) = \alpha] - \mathbb{P}_\tau(\alpha))$$

then

$$\langle \overline{u}_j, \underline{\chi}^{\tau, \alpha} \rangle \geq \Omega_M(1) \cdot \sqrt{n}.$$

For a model M with \bar{m} as an unstable fixed point, it is conjectured that the BP algorithm can successfully distinguish it from the null model M^* . We provide a BP-inspired spectral algorithm that solves the detection problem in this regime. We state the result below but leave the proof sketch to the technical overview section.

Theorem 2.11. *If a planted model M has $\lim_{\ell \rightarrow \infty} \alpha_\ell^{1/2\ell} > 1$, there is a spectral algorithm $A : \text{factor graphs} \rightarrow \{P, N\}$ that solves the detection problem in the following sense*

$$\Pr_{M_n}[A(\text{factor graph}) = P] = 1 - o_n(1) \quad \text{and} \quad \Pr_{M_n}[A(\text{factor graph}) = N] = 1 - o_n(1).$$

2.9 The stability condition via a finite linear operator L

In this part we define a finite linear operator L whose spectral radius gives a criterion for when $\lim_{\ell \rightarrow \infty} \alpha_\ell^{1/2\ell}$ is greater than 1 or less than 1.

Naively, computing $\lim_{\ell \rightarrow \infty} \alpha_\ell^{1/2\ell}$ requires us to consider trees whose size grows with ℓ . However we can simplify the expression for $\alpha_\ell^{1/2\ell}$ via an insight of [BC19] by observing that the tree \mathbb{T}_ℓ is constructed recursively. For any even level variable node v_k in the tree, the distribution of its children factor nodes depends only on v_k 's type. Furthermore, the factor node distribution $\{\deg_{i,j}(\tau(v_k))\}_{i \in F, j \in [a(i)]}$ also fully describes the distance 2 neighborhood of v_k . For a type τ variable v , define the random variable $\text{num}_\tau(i, j, j')$ to be the number of variables u that are connected to v via some type ϕ_i factor, and additionally u have index j' in the factor and v has index j . These random variables give a way to concisely express the total influence of the distance 2 type τ' variables to the type τ variable. This influence is

$$\sum_{i, j, j' | \text{Cl}(i)_{j'} = \tau'} \text{num}_\tau(i, j, j') \cdot \bar{\mathbf{M}}_{i, j | j'} \in \mathbb{R}^{[q] \times [q]}.$$

Then we can build a 2 step quadratic influence operator $L : \mathbb{R}^{T \cdot [q] \times T \cdot [q]} \rightarrow \mathbb{R}^{T \cdot [q] \times T \cdot [q]}$ such that the (τ, τ) block of $L(M)$ is

$$L(M)^{\tau, \tau} = \sum_{\tau'} \sum_{i \in F, j, j' \in [a(i)]} \sum_{i, j, j' | \text{Cl}(i)_{j'} = \tau'} \text{num}_\tau(i, j, j') \cdot \bar{\mathbf{M}}_{i, j | j'} M^{\tau', \tau'} \bar{\mathbf{M}}_{i, j | j'}^*,$$

and the off-diagonal blocks are $L(M)^{\tau, \tau'} = \mathbf{0}$.

Suppose the input M is such that each $M^{\tau', \tau'}$ captures the quadratic influence of some path $v_1, e_1, \dots, v_\ell, e_\ell$ such that the endpoint v_1 has type τ' ,

$$\left(\prod_{j=2}^{\ell} \bar{\mathbf{M}}_{\theta(e_{j-1}), i(v_{j-1}) | i(v_j)} \right) \left(\prod_{j=2}^{\ell} \bar{\mathbf{M}}_{\theta(e_{j-1}), i(v_{j-1}) | i(v_j)} \right)^*.$$

And all other blocks in M are 0. Then after applying the operator, every diagonal block $L(M)^{\tau, \tau}$ captures the expected quadratic influence of all paths $v_0, e_0, v_1, e_1, \dots, v_\ell, e_\ell$ which are 2-step extensions of the path $v_1, e_1, \dots, v_\ell, e_\ell$ and whose endpoint v_0 has type τ .

Using this operator L we can rewrite α_ℓ as

$$\alpha_\ell = \mathbf{E}_{\mathbb{T}_\ell} \left[\sum_{(e_\ell, v_\ell, \dots, e_0, v_0) \in \mathbb{T}_\ell} \text{Tr} \left(\left(\prod_{j=1}^{\ell} \bar{\mathbf{M}}_{\theta(e_{j-1}), i(v_{j-1}) | i(v_j)} \right) \left(\prod_{j=1}^{\ell} \bar{\mathbf{M}}_{\theta(e_{j-1}), i(v_{j-1}) | i(v_j)} \right)^* \right) \right] = \text{Tr} (L^\ell(\text{Diag}(\mathbf{P}))),$$

where $\text{Diag}(\mathbf{P}) \in \mathbb{R}^{T \cdot [q] \times T \cdot [q]}$ is a diagonal matrix whose $((\tau, c), (\tau, c))$ entry has value $\mathbb{P}_\tau(c)$.

Use λ_L to denote the maximum eigenvalue of L . Then by standard linear algebra fact,

$$\lim_{n \rightarrow \infty} \text{Tr} \left(L^\ell (\text{Diag}(\mathbb{P})) \right) \rightarrow \lambda_L^\ell.$$

Therefore we obtain the following equivalence relation between α_ℓ and λ_L .

Lemma 2.12. *For a model M , $\lim_{\ell \rightarrow \infty} \alpha_\ell^{1/2\ell} = \lambda_L^{1/2}$. Thus the fixed point \bar{m} of M is stable if and only if $\lambda_L \leq 1$.*

3 Technical Overview

3.1 Algorithm for distinguishing

We now describe our algorithm for distinguishing if an instance G was sampled from the null distribution M^\times from the planted distribution M . Our algorithm constructs a matrix M_G obtained from linearizing ℓ rounds of the belief propagation algorithm at the uninformative fixed point on input G and tests if its largest eigenvalue exceeds a chosen threshold κ . If it does then the algorithm declares that G came from the planted distribution, and otherwise claims G was sampled from the null distribution. A bulk of the technical work is in proving that this particular matrix M_G has all its eigenvalues bounded by the chosen threshold κ when G is sampled from the null model, and in illustrating that M_G has an ‘‘outlier’’ eigenvalue exceeding κ otherwise. In this section, we delve more into the description of M_G and then give a brief description of how we prove the statements about the eigenvalues of M_G in the null and planted models.

More concretely, given a random instance G sampled either from M^\times or M , we set M_G as the following matrix $\underline{A}_G^{(\ell)}$ called the *length- ℓ centered nonbacktracking walk power* of G , for which we provide a slightly informal description below.

Definition 3.1 (Centered nonbacktracking power (slightly informal)). $\underline{A}_G^{(\ell)}$ is a $nq \times nq$ matrix which we treat as a $n \times n$ grid of $q \times q$ blocks. The block rows and columns are indexed by $[n]$. In the (i, j) -th block, we place the following $q \times q$ matrix:

$$\sum_{\substack{i_{e_1} v_1 e_2 v_2 \dots e_\ell j \in \\ \text{all nonbacktracking walks} \\ \text{from } i \text{ to } j \text{ in complete} \\ \text{factor graph}}} \bar{M}_{e_1, i|v_1} \cdot \bar{M}_{e_2, v_1|v_2} \cdots \bar{M}_{e_\ell, v_{\ell-1}|j} \cdot (\mathbf{1}[e_1 \in G] - \text{Pr}_{M^\times}[e_1 \in G]) \cdots (\mathbf{1}[e_\ell \in G] - \text{Pr}_{M^\times}[e_\ell \in G]).$$

Recall the matrix L from [Section 2.8](#) which the stability prediction of belief propagation was based on, and let λ_L denote its largest eigenvalue. The two main technical theorems we prove about $\underline{A}_G^{(\ell)}$ in service of proving that our algorithm is correct with high probability are:

Theorem 3.2 (Local statistics in planted model). *When $G \sim M$, with probability $1 - o_n(1)$: $\lambda_{\max}(\underline{A}_G^{(\ell)}) \geq \frac{\lambda_L^\ell}{q}$.*

The proof of this is carried out in [Section 6](#) and uses two ingredients: the first is recognizing that $\underline{A}_G^{(\ell)}$ is self-adjoint under a certain inner product $\langle \cdot, \cdot \rangle_{\mathbb{H}}$, due to which for any vector x :

$$\langle x, \underline{A}_G^{(\ell)} x \rangle_{\mathbb{H}} \leq \lambda_{\max} \left(\underline{A}_G^{(\ell)} \right) \cdot \langle x, x \rangle_{\mathbb{H}}.$$

The second ingredient is in identifying a vector x depending on the planted solution the instance G was sampled with which makes the above quadratic $\langle x, \underline{A}_G^{(\ell)} x \rangle_{\mathbb{H}}$ larger than the desired lower bound of $\frac{\lambda_L^\ell}{q}$ with high probability.

The second main technical theorem, which is proved in [Theorem 3.3](#) is:

Theorem 3.3 (Eigenvalue bound in null model). *When $G \sim M^\times$ and $(\log \log n)^2 \leq \ell \leq \frac{\log n}{(\log \log n)^2}$, for every constant $\varepsilon > 0$ with probability $1 - o_n(1)$, all eigenvalues of $\underline{A}_G^{(\ell)}$ are bounded in magnitude by $((1 + \varepsilon)\sqrt{\lambda_L})^\ell$.*

When $\lambda_L > 1$, we choose $\delta > 0$ so that $1 + \delta < \sqrt{\lambda_L}$, ℓ as $(\log \log n)^2$, and κ as $((1 + \delta)\sqrt{\lambda_L})^\ell$. Then as an immediate consequence of [Theorem 3.2](#) and [Theorem 3.3](#) we know that the algorithm correctly distinguishes between M^\times and M with high probability.

We now elaborate on [Theorem 3.3](#) and elucidate the exact random matrix concentration statement.

3.2 Matrix concentration vignette

Consider an Erdős-Rényi graph H sampled from $G(n, \frac{d}{n})$. Pick a random vertex v in H and observe a “large” radius neighborhood around v . Typically, this neighborhood around v will be a tree, and additionally, and in the large- n limit the distribution of this tree is a *Galton-Watson process* – a random (possibly infinite) tree \mathbb{T} generated by starting at a root vertex r , attaching $\text{Poisson}(d)$ children to r , and then attaching $\text{Poisson}(d)$ children to each child of r and so on.² So this tells us that there is some sense in which \mathbb{T} “approximates” the finite random graph. This intuition is spectrally articulated by a theorem which is (implicitly) due to [\[BLM15\]](#) (see also [\[FM17\]](#) and [\[BMR19\]](#)). Before we state the theorem, we bring up a natural quantity to associate to the random tree: the *growth rate* which is defined as

$$\lim_{\ell \rightarrow \infty} \mathbf{E}[\# \text{ of vertices at depth-}\ell]^{1/\ell},$$

which is equal to d for the aforementioned Galton-Watson process.

Theorem 3.4 ([\[BLM15, FM17, BMR19\]](#)). *Let $\underline{A}_H := A_H - \mathbf{E}A_H$ be the centered adjacency matrix of H . Suppose $d > 1$, then:*

$$|\lambda|_{\max}(\underline{A}_H^{(\ell)}) \leq ((1 + o(1))\sqrt{d})^\ell = ((1 + o(1))\sqrt{\text{growth rate of } \mathbb{T}})^\ell$$

for $\ell \in \left[(\log \log n)^2, \frac{\log n}{\log \log n} \right]$.

Now, let’s add a small twist: sample two Erdős-Rényi graphs $H_1 \sim G(n, \frac{d_1}{n})$ and $H_2 \sim G(n, \frac{d_2}{n})$ and consider the weighted graph $H = 0.9H_1 - H_2$. The random tree that H locally resembles is the following *different* Galton-Watson process \mathbb{T}' : start at a root vertex r , connect $\text{Poisson}(d_1)$ children with edges of weight 0.9 and $\text{Poisson}(d_2)$ children with edges of weight -1 to r , then repeat the same for each child vertex, and keep going. The following quantity is the correct generalization of growth rate to weighted graphs, which we call the *weighted growth rate* of the tree:

$$\text{wgr}(\mathbb{T}') := \lim_{\ell \rightarrow \infty} \mathbf{E} \left[\sum_{\substack{P \in \text{length-}\ell \text{ paths} \\ \text{starting at root}}} \prod_{e \in P} w_e^2 \right].$$

For illustrative purposes, one subcase of our matrix concentration result is:

Theorem 3.5. *Let $\underline{A}_H := A_H - \mathbf{E}A_H$ be the centered adjacency matrix of H . Suppose $\text{wgr}(\mathbb{T}') > 1$, then:*

$$|\lambda|_{\max}(\underline{A}_H^{(\ell)}) \leq ((1 + o(1))\sqrt{\text{wgr}(\mathbb{T}')})^\ell$$

for $\ell \in \left[(\log \log n)^2, \frac{\log n}{(\log \log n)^2} \right]$

²The reader is advised to not pay too much attention to the fact that the number of children are distributed according to a *Poisson* random variable. The important property is that a vertex has d children on average.

We now discuss our full matrix concentration theorem which captures both of the above mentioned theorems. Before doing so, it is worth noting that the picture for random graphs being spectrally approximated by infinite graphs is far more well understood in the setting of models of random *regular* graphs through works of [Fri03, Bor19, BC19, MOP20, OW20] but we defer the readers to [OW20] for an extensive discussion of what is known in that setting.

3.3 Matrix concentration statement

Let $\langle \cdot, \cdot \rangle_v$ be an inner product on \mathbb{R}^{nq} and let M^* denote the adjoint of a matrix M under this inner product. We consider $nq \times nq$ random matrices sampled according to the following model (whose notation the reader should treat independently from the preceding notation related to distinguishing instances from the null and planted distributions).

Definition 3.6 (Random matrix model (slightly informal)). The model has an underlying *left vertex set* which is equal to $[n]$. First, every vertex v is assigned a *type* $\tau(v)$ in $[T]$ sampled according to a distribution π . There are F types of right vertices, given by set $[F]$. Each right vertex type i comes with an *arity* k_i , which is a positive integer, a *profile* χ_i which is a tuple in $[T]^{k_i}$, a collection of $k_i(k_i - 1)$ matrices $\{M_{i,(a,b)}\}_{(a,b) \in [k_i]^2: a \neq b}$, and a *density* ϕ_i . A random instance H is sampled in the following way: for every $(i, (v_1, \dots, v_{k_i}))$ for $i \in [F]$ and tuple (v_1, \dots, v_{k_i}) in $[n]^{k_i}$ of distinct elements such that $(\tau(v_1), \dots, \tau(v_{k_i})) = \chi_i$ we add $(i, (v_1, \dots, v_{k_i}))$ as a right vertex with probability $\frac{\phi_i}{n^{k_i-1}}$, connect edges to v_1, \dots, v_{k_i} and mark the edge to v_t with number t . We use \mathcal{K}_n to refer to the bipartite graph with left vertex set $[n]$, and the right vertex set containing every potential right vertex. Now let $\gamma = (i, (v_1, \dots, v_{k_i}))$; for a two-step $v_a \rightarrow \gamma \rightarrow v_b$ in the complete graph for $a \neq b$ we use $M_{v_a \gamma v_b}$ to denote the matrix $M_{i,(a,b)}$. The random matrix we are interested in, which we denote $\underline{A}_H^{(\ell)}$, is the matrix where the uv entry contains:

$$\underline{A}_H^{(\ell)}[i, j] := \sum_{\substack{i\gamma_1 v_1 \dots \gamma_\ell j \\ \text{nonbacktracking walks in } \mathcal{K}}} M_{i\gamma_1 v_1} \dots M_{v_{\ell-1} \gamma_\ell j} \cdot (1[\gamma_1 \in H] - \Pr[\gamma_1 \in H]) \dots (1[\gamma_\ell \in H] - \Pr[\gamma_\ell \in H]).$$

Definition 3.7 (Galton-Watson tree approximating random matrix (informal)). For a given setting of parameters for the random model from Definition 3.6, the bipartite Galton-Watson tree \mathbb{T} which “locally resembles” an instance H sampled from the model is as follows:

1. Start with a left root vertex r and assign it type $\tau(r) \sim \pi$.
2. For each $i \in [F]$ and each $j \in [k_i]$ such that $(\chi_i)_j = \tau(r)$ sample

$$n_{i,j} \sim \text{Poisson} \left(\frac{\phi_i}{\pi_{\tau(v)}} \prod_{t=1}^{k_i} \pi_{(\chi_i)_t} \right),$$

and attach $n_{i,j}$ right vertices of type i to r and mark the corresponding edge with j . Then to each such right vertex, attach $k_i - 1$ (left vertex) children and mark the edges with numbers from $[k_i] \setminus \{j\}$. To each added child vertex v with edge marked with t , assign it type $(\chi_i)_t$.

3. Repeat [step 2](#) for each added left vertex child.

We define the *matrix weighted growth rate* of \mathbb{T} to be the following:

$$\text{mwgr}(\mathbb{T}) := \lim_{\ell \rightarrow \infty} \mathbf{E} \left[\text{Tr} \left(\sum_{\substack{a\gamma_1 v_1 \dots \gamma_\ell b \\ \text{nonbacktracking walks in } \mathcal{K}}} M_{a\gamma_1 v_1} \dots M_{v_{\ell-1} \gamma_\ell b} M_{b\gamma_\ell v_{\ell-1}} \dots M_{v_1 \gamma_1 a} \right) \right]^{1/\ell}.$$

We prove:

Theorem 3.8 (Main matrix concentration theorem (slightly informal)). *Let \mathbf{H} be a random instance of a setting of parameters for the model from Definition 3.6 and let \mathbb{T} be the tree which locally approximates \mathbf{H} in the sense of Definition 3.7. Suppose:*

1. *For every right vertex $\gamma = (i, (v_1, \dots, v_{k_i}))$ in \mathcal{K}_n and every $1 \leq a, b \leq k_i$ for distinct a, b the $nq \times nq$ matrix obtained by placing $M_{v_a \gamma v_b}$ in the (v_a, v_b) block and zeros everywhere else is adjoint to the $nq \times nq$ matrix obtained by placing $M_{v_b \gamma v_a}$ in the (v_b, v_a) block and zeros everywhere else under $\langle \cdot, \cdot \rangle_v$.*
2. *There is a constant C such that for any nonbacktracking walk $v_0 \gamma_1 v_1 \dots \gamma_s v_s$ in \mathcal{K}_n :*

$$\|M_{v_0 \gamma_1 v_1} \dots M_{v_{s-1} \gamma_s v_s}\| \leq C.$$

3. $\text{mwgr}(\mathbb{T}) \geq 1$.

Then if $(\log \log n)^2 \leq \ell \leq \frac{\log n}{(\log \log n)^2}$, with probability $1 - o_n(1)$:

$$|\lambda|_{\max}(\underline{A}_{\mathbf{H}})^{(\ell)} \leq ((1 + o(1))\sqrt{\text{mwgr}(\mathbb{T})})^\ell.$$

The full formal set-up for Theorem 3.8 along with its proof is in Section 7.

4 A conjectured detection/recovery threshold

In Section 2.9 we see the connection between λ_L and α_ℓ of a model \mathbf{M} . In this section we prove this connection more rigorously, and conclude with a conjectured weak-recovery threshold in terms of λ_L . We start by quickly going through the definitions of the partial derivative matrices $\bar{\mathbf{M}}_{\theta(e_j), i_{e_j}(v_j) | i_{e_j}(v_{j+1})}$, the color distribution matrix \mathbf{D}_τ , the influence variance α_ℓ , and their connections.

Claim 4.1. The following matrices satisfy:

1. $\bar{\mathbf{M}}_{\theta(e_j), i_{e_j}(v_j) | i_{e_j}(v_{j+1})} = (\mathbf{I} - \mathbf{P}_{\tau(v_j)} \mathbf{1}^\top) \Psi_{\theta(e_j), i_{e_j}(v_j) | i_{e_j}(v_{j+1})}$.
2. Let $\mathbf{D}_\tau := \text{Diag}(\mathbf{P}_\tau)$. Then $\mathbf{D}_{\tau(v_j)}^\dagger \bar{\mathbf{M}}_{\theta(e_j), i(v_j) | i(v_{j+1})} = \bar{\mathbf{M}}_{\theta(e_j), i(v_{j+1}) | i(v_j)}^\top \mathbf{D}_{\tau(v_{j+1})}^\dagger$ where the \dagger in the superscript denotes the pseudoinverse of the matrix.

This claim is proved in Appendix C. From here on we drop the subscript e from the index function i_e whenever it is clear from the context which factor is being considered.

Using these notations we can express the influence variance, or the *amplification factor*, α_ℓ as follows:

$$\alpha_\ell := \mathbf{E}_{\mathbb{T}} \left[\text{Tr} \left(\sum_{(e_\ell, v_\ell, \dots, e_0, v_0) \in \mathbb{T}} \left(\prod_{i=1}^{\ell} \bar{\mathbf{M}}_{\theta(e_{i-1}), i(v_{i-1}) | i(v_i)} \right) \left(\prod_{i=1}^{\ell} \bar{\mathbf{M}}_{\theta(e_{i-1}), i(v_{i-1}) | i(v_i)} \right)^* \right) \right].$$

The key property of the above amplification factor we are interested in is its limiting behavior as k goes to infinity. We say that the uninformative fixed point of the belief propagation update rule is *stable* if $\lim_{n \rightarrow \infty} \alpha_\ell^{1/2\ell} \leq 1$, and *unstable* otherwise. Furthermore, if the fixed point is stable, the problem of weak-recovering the hidden coloring of \mathbf{M} is conjectured to be hard, and if the fixed point is unstable, this problem is conjectured to be easy.

In the remainder of this section, we focus on obtaining a simpler criterion for stability by simplifying the expression for the amplification factor. In particular, we give a constant dimensional linear transformation whose top eigenvalue is greater than 1 if the fixed point is stable and is less than 1 if the fixed point is unstable. For a leaf-to-root path $e_\ell v_\ell e_{\ell-1} v_{\ell-1} \dots e_0 v_0$ we say its type is

$$\theta_\ell \rightarrow \text{out}_\ell \rightarrow \tau_\ell \rightarrow \text{in}_{\ell-1} \rightarrow \theta_{\ell-1} \rightarrow \text{out}_{\ell-1} \rightarrow \dots \rightarrow \text{out}_1 \rightarrow \tau_1 \rightarrow \text{in}_0 \rightarrow \theta_0 \rightarrow \text{out}_0 \rightarrow \tau_0$$

where θ_t is the factor type of e_t , τ_t is the variable type of v_t , out_t is the index of v_t in e_t , and in_t is the index of v_{t+1} in e_t . The amplification factor can then be written as:

$$\alpha_\ell = \text{Tr} \left(\sum_{\theta_\ell \rightarrow \dots \rightarrow \tau_0} M(\theta_\ell \rightarrow \dots \rightarrow \tau_0) \right)$$

where the sum is enumerated over all leaf-to-root path types and $M(\theta_\ell \rightarrow \dots \rightarrow \tau_0)$ is defined as follows:

$$\begin{aligned} M(\theta_\ell \rightarrow \dots \rightarrow \tau_0) &= \mathcal{T}(\tau_0) \cdot \left(\prod_{t=0}^{\ell-1} \frac{\phi_{\theta_t}}{\mathcal{T}(\tau_t)} \cdot \prod_{s=1}^{a(\theta_t)} \mathcal{T}(\text{Cl}(\theta_t)_s) \cdot \mathbf{1}[\text{Cl}(\theta_t)_{\text{out}_t} = \tau_t] \cdot \mathbf{1}[\text{Cl}(\theta_t)_{\text{in}_t} = \tau_{t+1}] \right) \cdot \\ &\quad \left(\prod_{t=0}^{\ell-1} \overline{\mathbf{M}}_{\theta_t, \text{out}_t | \text{in}_t} \right) \left(\prod_{t=0}^{\ell-1} \overline{\mathbf{M}}_{\theta_t, \text{out}_t | \text{in}_t} \right)^* \cdot \frac{\phi_{\theta_\ell}}{\mathcal{T}(\tau_\ell)} \prod_{s=1}^{a(\theta_\ell)} \mathcal{T}(\text{Cl}(\theta_\ell)_s) \cdot \mathbf{1}[\text{Cl}(\theta_\ell)_{\text{out}_\ell} = \tau_\ell]. \end{aligned}$$

Now, let's define $V_{\ell, \tau}$ as:

$$V_{\ell, \tau} := \sum_{\substack{\theta_\ell \rightarrow \dots \rightarrow \tau_0 \\ \tau_0 = \tau}} M(\theta_\ell \rightarrow \dots \rightarrow \tau_0)$$

and V_ℓ as the following $|T| \cdot q \times |T| \cdot q$ block diagonal matrix comprised of $q \times q$ -dimensional blocks with block rows and columns indexed by T :

$$V_\ell[\tau, \tau] = V_{\ell, \tau}.$$

Finally, we define a linear transformation L on the space of $|T| \cdot q \times |T| \cdot q$ matrices. To define $L(M)$ we treat M as a block matrix comprised of $q \times q$ -dimensional blocks with blocks rows and columns indexed by T .

$$\begin{aligned} L(M)[\tau, \tau] &:= \mathcal{T}(\tau) \sum_{\tau' \rightarrow \text{in} \rightarrow \theta \rightarrow \text{out} \rightarrow \tau} \frac{\phi_\theta}{\mathcal{T}(\tau)} \prod_{s=1}^{a(\theta)} \mathcal{T}(\text{Cl}(\theta)_s) \cdot \mathbf{1}[\text{Cl}(\theta)_{\text{out}} = \tau] \cdot \mathbf{1}[\text{Cl}(\theta)_{\text{in}} = \tau'] \cdot \\ &\quad \overline{\mathbf{M}}_{\theta, \text{out} | \text{in}} \cdot M[\tau', \tau'] \cdot \overline{\mathbf{M}}_{\theta, \text{out} | \text{in}}^*. \end{aligned}$$

$$L(M)[\tau_1, \tau_2] := 0 \quad \text{for } \tau_1 \neq \tau_2.$$

Now observe that $V_{\ell+1} = L(V_\ell)$; consequently $V_\ell = L^\ell(V_0)$ and $\alpha_\ell = \text{Tr}(L^\ell(V_0))$.

We now connect the limiting behavior of the amplification factor α_ℓ to the eigenvalues of L . We start by making a few observations:

Observation 4.2. If M is a positive semidefinite matrix, then $L(M)$ is also a positive semidefinite matrix.

Observation 4.3. V_0 is a diagonal matrix with strictly positive entries on its diagonal and hence is positive definite.

Observation 4.4. Since $L^\ell(V_0)$ is positive semidefinite:

$$\|L^\ell(V_0)\|_F \leq \text{Tr}(L^\ell(V_0)) \leq \sqrt{q \cdot |T|} \|L^\ell(V_0)\|_F.$$

Our final ingredient is a lemma that appears in [BC19, Theorem 16, part (ii)].

Lemma 4.5. *If L is a linear operator on $r \times r$ matrices with spectral radius λ_L such that for any positive semidefinite matrix M , $L(M)$ is also positive semidefinite, then for any positive definite M' ,*

$$\lambda_L = \lim_{\ell \rightarrow \infty} \|L^\ell(M')\|_F^{1/\ell}.$$

By [Observation 4.2](#), [Observation 4.3](#) and [Lemma 4.5](#): $\lambda_L = \lim_{\ell \rightarrow \infty} \|L^\ell(V_0)\|_F^{1/\ell}$. [Observation 4.4](#) lets us conclude:

$$\lambda_L = \lim_{\ell \rightarrow \infty} \text{Tr}(L^\ell(V_0))^{1/\ell} = \lim_{\ell \rightarrow \infty} \alpha_\ell^{1/\ell}.$$

Thus we prove [Lemma 2.12](#), and make the following conjecture:

Conjecture 4.6. *If $\lambda_L > 1$, then $\lim_{\ell \rightarrow \infty} \alpha_\ell$ goes to ∞ we conjecture that recovery is easy and if $\lambda_L \leq 1$, then $\lim_{\ell \rightarrow \infty} \alpha_\ell = 0$ and we conjecture that it is hard.*

5 A spectral distinguishing algorithm

We now describe the spectral distinguisher we use, which is based on linearizing the belief propagation algorithm outlined in [Section 4](#). Recall that given G sampled from either M^\times or M our goal is to output “null” if $G \sim M^\times$ and “planted” if $G \sim M$ with probability $1 - o(1)$. Further, the messages given by [\(5\)](#) and [\(6\)](#) are a fixed point for the BP update rule for M (which is equivalent to the detailed balanced condition [\(45\)](#) holding). The sample G is given by the tuple $([n], E_1, \dots, E_F, \tau)$. Our algorithm constructs a matrix called the *null-centered nonbacktracking power matrix* and thresholds on its largest eigenvalue against a particular value t which is a function of the null and planted models and outputs “planted” if the largest eigenvalue exceeds t and “null” otherwise.

Recall the definition of the matrix $\underline{A}_G^{(\ell)}$.

$$\sum_{\substack{i e_1 v_1 e_2 v_2 \dots e_\ell j \in \\ \text{all nonbacktracking walks} \\ \text{from } i \text{ to } j \text{ in complete} \\ \text{factor graph}}} \bar{M}_{e_1, i|v_1} \cdot \bar{M}_{e_2, v_1|v_2} \cdots \bar{M}_{e_\ell, v_{\ell-1}|j} \cdot (\mathbf{1}[e_1 \in G] - \Pr_{M^\times|\tau}[e_1 \in G]) \cdots (\mathbf{1}[e_\ell \in G] - \Pr_{M^\times|\tau}[e_\ell \in G])$$

where the complete factor graph is defined as follows.

Definition 5.1. We define the *complete factor graph* $\mathcal{K}_n = ([n], E_1, \dots, E_F)$ where E_i denotes the collection of all potential type- i factors that could appear in G .

We now describe our algorithm.³

- Compute a matrix representation of the linear operator L from the statement of [Conjecture 4.6](#).
- Let λ_L be the spectral radius of L .
- Choose κ strictly in between $\sqrt{\lambda_L}$ and λ_L .
- Let $s = \lceil \sqrt{\log n} \rceil$. Compute $\underline{A}_G^{(s)}$ and compute its largest eigenvalue ρ .
- If $\rho > \kappa^s$, output “planted”, otherwise output “null”.

³The details for why each step can be carried out efficiently are briefly discussed at the end of this section.

To prove that the above algorithm works it suffices to prove that when G is sampled from the null distribution, all its eigenvalues are all less κ^s and when G is sampled from the planted distribution there is an eigenvalue greater than κ^s . Henceforth we assume $\lambda_L > 1$. To prove both of these facts under the hypothesis that $\lambda_L > 1$, one ingredient we need is that the matrix $\underline{A}_G^{(s)}$ is self-adjoint under an appropriate inner product.

Given a vector in \mathbb{R}^{nq} we treat it as a block vector comprising of n blocks of dimension q each where each block corresponds to a vertex in $[n]$. Now, we define a $nq \times nq$ -dimensional positive diagonal matrix \mathbf{H}_τ where the (v, v) block is equal to:

$$\mathbf{H}_{\tau, (v, v)}[c, c] := \begin{cases} \mathbb{P}_{\tau(v)}(c) & \text{if } \mathbb{P}_{\tau(v)}(c) > 0 \\ 1 & \text{otherwise.} \end{cases}$$

We will use the following inner product on \mathbb{R}^{nq} :

$$\langle x, y \rangle_{\mathbf{H}} := x^\top \mathbf{H}_\tau^{-1} y.$$

Remark 5.2. Ideally, we would like to simply place $\mathbb{P}_{\tau(v)}(c)$ in every diagonal entry $[(v, c), (v, c)]$ and use the pseudoinverse of \mathbf{H} instead. But doing so leads to some complications of the defined bilinear form not necessarily satisfying strict positive definiteness required of an inner product. The choice of 1 is arbitrary and does not influence any of the statements since all the vectors we work with have zeros in the (v, c) coordinates where $\mathbb{P}_{\tau(v)}(c)$ is 0, and also that coordinate subspace is in the kernel of every matrix we work with.

Claim 5.3. The matrix $\underline{A}_G^{(s)}$ is self-adjoint under $\langle \cdot, \cdot \rangle_{\mathbf{H}}$.

Proof. For any vectors x, y :

$$\begin{aligned} \langle x, \underline{A}_G^{(s)} y \rangle_{\mathbf{H}} &= \sum_{u, v \in [n]} \sum_{\substack{u_{e_1} v_1 e_2 v_2 \\ \dots e_\ell v \in \mathcal{K}_n}} x[u]^\top \mathbf{H}^{-1} \overline{\mathbf{M}}_{e_1, u|v_1} \cdot \overline{\mathbf{M}}_{e_2, v_1|v_2} \cdots \overline{\mathbf{M}}_{e_\ell, v_{\ell-1}|v} y[v]. \\ &= (\mathbf{1}[e_1 \in G] - \Pr_{M^*|\tau}[e_1 \in G]) \cdots (\mathbf{1}[e_\ell \in G] - \Pr_{M^*|\tau}[e_\ell \in G]) \end{aligned}$$

From [Part 2 of Claim 4.1](#),

$$x[u]^\top \mathbf{H}^{-1} [u, u] \overline{\mathbf{M}}_{e_1, u|v_1} \cdots \overline{\mathbf{M}}_{e_\ell, v_{\ell-1}|v} y[v] = x[u]^\top (\overline{\mathbf{M}}_{e_\ell, v|v_{\ell-1}} \cdots \overline{\mathbf{M}}_{e_1, v_1|u})^\top \mathbf{H}^{-1} [v, v] y[v].$$

Plugging this back into the above gives:

$$\begin{aligned} &= \sum_{u, v \in [n]} \sum_{\substack{u_{e_1} v_1 e_2 v_2 \\ \dots e_\ell v \in \mathcal{K}_n}} x[u]^\top (\overline{\mathbf{M}}_{e_\ell, v|v_{\ell-1}} \cdots \overline{\mathbf{M}}_{e_1, v_1|u})^\top \mathbf{H}^{-1} [v, v] y[v]. \\ &= (\mathbf{1}[e_1 \in G] - \Pr_{M^*|\tau}[e_1 \in G]) \cdots (\mathbf{1}[e_\ell \in G] - \Pr_{M^*|\tau}[e_\ell \in G]) \\ &= \sum_{u, v \in [n]} \sum_{\substack{u_{e_1} v_1 e_2 v_2 \\ \dots e_\ell v \in \mathcal{K}_n}} (\overline{\mathbf{M}}_{e_\ell, v|v_{\ell-1}} \cdots \overline{\mathbf{M}}_{e_1, v_1|u} x[u])^\top \mathbf{H}^{-1} [v, v] y[v]. \\ &= (\mathbf{1}[e_1 \in G] - \Pr_{M^*|\tau}[e_1 \in G]) \cdots (\mathbf{1}[e_\ell \in G] - \Pr_{M^*|\tau}[e_\ell \in G]) \\ &= \langle \underline{A}_G^{(s)} x, y \rangle_{\mathbf{H}} \end{aligned}$$

which proves the claim. □

We first focus on obtaining spectral norm bounds on $\underline{A}_G^{(s)}$ in the null model. We obtain these bounds from [Theorem 7.4](#) so we verify that the matrix $\underline{A}_G^{(s)}$ indeed meets the hypothesis of the theorem statement. [Condition 1](#) is satisfied due to [Claim 5.3](#). As a consequence of the first part of [Claim 4.1](#), all the matrices $M(p)$ are Markov transition matrices with an eigenspace projected away, and hence have all their entries bounded by 1. Since these matrices have dimension $q \times q$, their operator norm is bounded by some constant C depending only on q , and hence [Condition 2](#) is also satisfied. Next, $\rho(\text{Bl}, M^*)$ is exactly equal to $\sqrt{\lambda_L}$, which by our assumption is greater than 1. Finally, we chose s in the range handled by the theorem statement and thus [Theorem 7.4](#) implies:

Theorem 5.4. *Suppose $G \sim M^*$. For every constant $\varepsilon > 0$, with probability $1 - o(1)$:*

$$|\lambda|_{\max} \left(\underline{A}_G^{(s)} \right) \leq ((1 + \varepsilon)\sqrt{\lambda_L})^s.$$

We can choose ε small enough so that $|\lambda|_{\max} \left(\underline{A}_G^{(s)} \right) \leq \kappa^s$ for $G \sim M^*$ whp.

Finally, to prove that there is an eigenvalue greater than κ^s when $G \sim M$, by [Claim 5.3](#) it suffices to illustrate a vector $x \in \mathbb{R}^{nq}$ such that $\frac{\langle x, \underline{A}_G x \rangle_{\mathbb{H}}}{\langle x, x \rangle_{\mathbb{H}}} \geq \kappa^s$. Then as a direct consequence of [Theorem 6.1](#):

Theorem 5.5. *Suppose $G \sim M$. There is an absolute constant C such that with probability $1 - o(1)$:*

$$|\lambda|_{\max} \left(\underline{A}_G^{(s)} \right) \geq C\lambda_L^s.$$

Since s is super-constant and κ is strictly less than λ_L , it is indeed true that $|\lambda|_{\max} \left(\underline{A}_G^{(s)} \right) \geq \kappa^s$ for $G \sim M$ whp. Consequently, we can summarize our main theorem on distinguishing the null distribution from the planted distribution:

Theorem 5.6. *When $\lambda_L > 1$, the task of distinguishing M^* from M with high probability can be done in polynomial time.*

5.1 Implementation details

Our first goal is to explain how to efficiently choose κ which strictly between $\sqrt{\lambda_L}$ and λ_L when $\lambda_L > 1$. First note that there is a small enough ε such that if $\tilde{\lambda}_L$ is an additive ε -approximation of λ_L , then $\frac{\tilde{\lambda}_L + \sqrt{\lambda_L}}{2}$ lies strictly in between $\sqrt{\lambda_L}$ and λ_L . By [Lemma 4.5](#) there is large enough constant C such that $\|L^C(\mathbf{I})\|_F^{1/C}$ is ε -close to λ_L . Thus, if our estimator $\tilde{\lambda}_L$ is $\|L^{\log n}(\mathbf{I})\|_F^{1/\log n}$, then for large enough n , choosing κ as $\frac{\tilde{\lambda}_L + \sqrt{\lambda_L}}{2}$ would give us a number strictly between $\sqrt{\lambda_L}$ and λ_L .

Our second goal is to explain how to efficiently compute the matrix $\underline{A}_G^{(s)}$. Towards doing so, define the *nonbacktracking walk generator matrix* as the matrix with rows and columns indexed by vev' for variable vertices v and v' and constraint vertex e' :

$$B_G[(v_1 e_1 v_2), (v_3 e_2 v_4)] = \begin{cases} \overline{\mathbf{M}}_{e_2, v_3 | v_4} \cdot (\mathbf{1}[e_2 \in G] - \Pr_{G \sim M | \tau}[e_2 \in G]) & \text{if } e_1 \neq e_2, v_2 = v_3 \\ 0 & \text{otherwise.} \end{cases}$$

Let S_G be the matrix with rows indexed by variables in $[n]$ and columns indexed by all vev' where the (v, vev') entry contains $\overline{\mathbf{M}}_{e, v | v'} \cdot (\mathbf{1}[e \in G] - \Pr_{G \sim M | \tau}[e \in G])$ and the remaining entries contain 0, and let T_G be the matrix with rows indexed by all vev' and columns indexed by variables in $[n]$ where the (vev', v')

entry is \mathbf{I} and the remaining entries contain 0. Then $S_G B_G^{s-1} T_G = \underline{A}_G^{(s)}$, and it is apparent that the LHS can be computed efficiently.

Finally, since $\underline{A}_G^{(s)}$ is self-adjoint under $\langle \cdot, \cdot \rangle_{\mathbf{H}}$ by [Claim 5.3](#) its largest eigenvalue can be efficiently computed via standard methods such as the power iteration method to a precision necessary for the distinguishing algorithm.

6 Statistics for the planted model

Consider the planted model $M_n = (n, T, \mathcal{T}, C, \mathbb{P}, \phi)$. Use $G = ([n], E_1, \dots, E_F, \tau, \mathbf{c})$ to denote a sample from $M_p^{[n]}$. Recall that in this section our goal is to prove a lower bound on the spectral radius of $\underline{A}_G^{(s)}$ for $s \leq \frac{\log n}{(\log \log n)^2}$ by illustrating a “witness” vector with large quadratic form.

We define the *local statistics vector* associated to G denoted $g \in \mathbb{R}^{q^n}$ to be the concatenation of vectors $g^v = u_{c_v}$ for all $v \in [n]$, where $u_c \in \mathbb{R}^q$ is the indicator of vector of color c . We will shorten $\langle g, g \rangle_{\mathbf{H}}$ to $\|g\|^2$ in this section.

In this section we prove:

Theorem 6.1. *Let $s \leq \frac{\log n}{(\log \log n)^2}$. There exists a constant γ such that with probability $1 - o_n(1)$:*

$$\frac{\langle g, \underline{A}_G^{(s)} g \rangle_{\mathbf{H}}}{\|g\|^2} = \gamma \lambda_L^s.$$

To prove [Theorem 6.1](#) we will introduce and recall some notation to streamline the proofs. First recall that we defined \mathcal{K}_n in [Definition 5.1](#) as the instance on variable set $[n]$ with all potential factors that could appear in a graph sampled from M_n . Let p be some length- $2s$ walk ($v_0 \rightarrow e_0 \rightarrow v_1 \dots v_{s-1} \rightarrow e_{s-1} \rightarrow v_s$) in \mathcal{K}_n starting and ending at a variable node. Use p_e to denote the factor nodes in p and p_v to denote the variable nodes in ∂p_e .⁴ We use $p_v[i]$ to denote v_i and $p_e[i]$ to denote e_i . Recall that we use $\theta(e)$ to denote the type of a factor node e . We use $\chi(p)$ to denote the number of excess edges in p . Concretely:

Definition 6.2. $\chi(p) = |p_v| - \sum_{j \in [s]} (a(\theta(p_e[j])))$.

Definition 6.3. Let $\text{NB}(s)$ denote the set of all length- s nonbacktracking walks in \mathcal{K}_n .

Definition 6.4. $\bar{\mathbf{M}}_p := \prod_{j=0}^{s-1} \bar{\mathbf{M}}_{\theta(p_e[j]), i(p_v[j]) | i(p_v[j+1])}$.

Definition 6.5. $\Psi_p := \prod_{j=0}^{s-1} \Psi_{\theta(p_e[j]), i(p_v[j]) | i(p_v[j+1])}$.

For $G \sim M_n$, we define the following notation:

Definition 6.6. $\mathbf{w}(G, p) := \bar{\mathbf{M}}_p \prod_{j=0}^{s-1} (\mathbf{1}[p_e[j] \in G] - \Pr_{G \sim M | \tau}[p_e[j] \in G])$.

Definition 6.7. We will use $\text{wt}(p)$ to denote $\prod_{j=0}^{s-1} (\mathbf{1}[p_e[j] \in G] - \Pr_{G \sim M | \tau}[p_e[j] \in G])$.

Remark 6.8. In the above language, the centered nonbacktracking power matrix of G is:

$$\underline{A}_G^{(s)}[u, v] = \sum_{p \in \text{NB}(s) | p_v[0]=u, p_v[s]=v} \mathbf{w}(G, p).$$

⁴We would like to stress that p_v also includes vertices that are *not* walked on, but are incident to factor nodes which are walked on.

Recall that a block (u, v) in the centered nonbacktracking power matrix $\underline{A}_G^{(s)}$ captures the matrix weight of all length- $2s$ nonbacktracking walks from u to v . Although the paths are nonbacktracking, there could be variables that are visited multiple times and there could also be off-path variables in p_v connected to multiple factors. It is hard to pinpoint the statistics precisely for this kind of walks, but luckily their contribution is negligible. Thus, we first remove these “bad” walks from $\underline{A}_G^{(s)}$ and analyze the resulting matrix $\overline{A}_G^{(s)}$, and then bound the contribution of these bad walks.

We give a formal definition of the “nice” walks that are kept in $\overline{A}_G^{(s)}$.

Definition 6.9. A path p is *self-avoiding* if $\chi(p) = |p_v| - \sum_{j \in [s]} (a(\text{Cl}_p(j)) - 1) = 1$.

In any self-avoiding path, every variable node in the interior of the path has degree 2 (i.e. is contained in 2 factors in p) and each of the other variable nodes has degree 1 (i.e. is contained in 1 factor in p).

Definition 6.10. Let $\text{SA}(s)$ denote the set of all length- $2s$ self-avoiding walks in \mathcal{K}_n .

Definition 6.11. The centered self-avoiding-walk matrix of G is $\overline{A}_G^{(s)} \in \mathbb{R}^{q^n \times q^n}$ where the (u, v) -th block of the matrix is

$$\overline{A}_G^{(s)}[u, v] = \sum_{p \in \text{SA}(s) | p_v[0]=u, p_v[s]=v} \mathbf{w}(G, p).$$

6.1 Statistics for the centered self-avoiding-walk matrix

Claim 6.12. For any self-avoiding path p in \mathcal{K}_n ,

$$\mathbf{E}_G \left[\langle g^{p_v[0]}, \mathbf{w}(G, p) g^{p_v[s]} \rangle_{\mathbb{H}} \right] = \Pr_{G \sim \mathcal{M} | \tau} [p \in G] \cdot \text{Tr} \left(\overline{\mathbf{M}}_p \overline{\mathbf{M}}_p^* \right).$$

Proof. The proof is via a chain of equalities. To lighten notation we use e_j to denote $p_e[j]$ and v_j to denote $p_v[j]$. We use $\text{Int}(p_v)$ to denote the interior vertices of p .

$$\begin{aligned} \mathbf{E}_{G | \tau} \left[\langle g^{v_0}, \mathbf{w}(G, p) g^{v_s} \rangle_{\mathbb{H}} \right] &= \sum_{c: p_v \rightarrow [q]} \prod_{w \in p_v} \mathbb{P}_{\tau(w)}(c(w)) \cdot \mathbf{E}_{G | \tau, c} \left[\mathbf{H}^{-1}[(v_0, c(v_0)), (v_0, c(v_0))] \cdot \right. \\ &\quad \left. \prod_{j=0}^{s-1} (\mathbf{1}[e_j \in G] - \Pr_{G \sim \mathcal{M} | \tau} [e_j \in G]) \overline{\mathbf{M}}_p[c(v_0), c(v_s)] \right] \\ &= \sum_{c: p_v \rightarrow [q]} \prod_{w \in p_v} \mathbb{P}_{\tau(w)}(c(w)) \cdot \prod_{j=0}^{s-1} (\Pr_{G \sim \mathcal{M} | \tau, c} [e_j \in G] - \Pr_{G \sim \mathcal{M} | \tau} [e_j \in G]) \cdot \\ &\quad \frac{1}{\mathbb{P}_{\tau(v_s)}(c(v_s))} \cdot \overline{\mathbf{M}}_p^*[c(v_s), c(v_0)] \\ &= \sum_{c: p_v \rightarrow [q]} \prod_{w \in \text{Int}(p_v)} \frac{1}{\mathbb{P}_{\tau(w)}(c(w))} \cdot \prod_{j=0}^{s-1} \Pr_{G \sim \mathcal{M} | \tau} [e_j \in G] \cdot \left(\mu_{e_j}(c(\partial e_j)) - \prod_{w \in \partial e_j} \mathbb{P}_{\tau(w)}(c(w)) \right) \cdot \\ &\quad \frac{1}{\mathbb{P}_{\tau(v_s)}(c(v_s))} \cdot \overline{\mathbf{M}}_p^*[c(v_s), c(v_0)] \\ &= \Pr_{G \sim \mathcal{M} | \tau} [p \in G] \sum_{c: \{v_0, \dots, v_s\} \rightarrow [q]} \prod_{j=0}^{s-1} (\Psi_{e_j, v_j | v_{j+1}}[c(v_j), c(v_{j+1})] - \mathbb{P}_{\tau(v_j)}) \cdot \overline{\mathbf{M}}_p^*[c(v_s), c(v_0)] \\ &= \Pr_{G \sim \mathcal{M} | \tau} [p \in G] \sum_{c(v_0), c(v_s)} \overline{\mathbf{M}}_p[c(v_0), c(v_s)] \cdot \overline{\mathbf{M}}_p^*[c(v_s), c(v_0)] \\ &= \Pr_{G \sim \mathcal{M} | \tau} [p \in G] \cdot \text{Tr} \left(\overline{\mathbf{M}}_p \overline{\mathbf{M}}_p^* \right) \end{aligned}$$

□

Now we give precise estimates for the statistics of the centered self-avoiding walk matrix.

Lemma 6.13.

$$\mathbf{E}_G \left[\left\langle g, \overline{A}_G^{(s)} g \right\rangle_{\mathbf{H}} \right] = (1 - o_n(1)) \lambda_L^s n.$$

Proof. Expanding out $\overline{A}_G^{(s)}$ as a sum by its definition in [Definition 6.11](#) gives:

$$\begin{aligned} \mathbf{E}_G \left[\left\langle g, \overline{A}_G^{(s)} g \right\rangle_{\mathbf{H}} \right] &= \sum_{p \in \text{SA}(s)} \mathbf{E}_G \left[\left\langle g^{p_v[0]}, \mathbf{w}(G, p) g^{p_v[s]} \right\rangle_{\mathbf{H}} \right] \\ &= \sum_{p \in \text{SA}(s)} \Pr_{G \sim \mathcal{M} | \tau} [p \in G] \cdot \text{Tr} \left(\overline{\mathbf{M}}_p \overline{\mathbf{M}}_p^* \right) \quad (\text{via Claim 6.12}) \\ &= n \cdot (1 - o_n(1)) \lambda_L^s \quad (\text{by the definition of } \lambda_L). \end{aligned}$$

□

We next bound the variance of the local statistics.

Lemma 6.14. *There is an absolute constant C such that:*

$$\mathbf{E}_G \left[\left\langle g, \overline{A}_G^{(s)} g \right\rangle_{\mathbf{H}}^2 \right] - \mathbf{E}_G \left[\left\langle g, \overline{A}_G^{(s)} g \right\rangle_{\mathbf{H}} \right]^2 \leq n(Cs)^{Cs}.$$

Proof. If two walks p and \tilde{p} do not share any vertices, then their contribution to $\mathbf{E}_G \left[\left\langle g, \overline{A}_G^{(s)} g \right\rangle_{\mathbf{H}}^2 \right]$ satisfies

$$\begin{aligned} \text{contribution of } p, \tilde{p} &= \mathbf{E}_G \left[\left\langle g^{p_v[0]}, \mathbf{w}(G, p) g^{p_v[s]} \right\rangle_{\mathbf{H}} \left\langle g^{\tilde{p}_v[0]}, \mathbf{w}(G, \tilde{p}) g^{\tilde{p}_v[s]} \right\rangle_{\mathbf{H}} \right] \\ &= \mathbf{E}_G \left[\left\langle g^{p_v[0]}, \mathbf{w}(G, p) g^{p_v[s]} \right\rangle_{\mathbf{H}} \right] \mathbf{E}_G \left[\left\langle g^{\tilde{p}_v[0]}, \mathbf{w}(G, \tilde{p}) g^{\tilde{p}_v[s]} \right\rangle_{\mathbf{H}} \right]. \end{aligned}$$

So this contribution cancels out with the identical term in $\mathbf{E}_G \left[\left\langle g, \overline{A}_G^{(s)} g \right\rangle_{\mathbf{H}} \right]^2$. Thus it suffices for us to consider self-avoiding walks p and \tilde{p} that share some vertices.

We write $p \parallel \tilde{p}$ if they share some variable or factor nodes and the shared nodes have consistent types and use p^\cup to denote the union of the two walks. Now, note that:

1. Conditioned on τ and c every factor node with arity k is chosen independently with probability at most $\frac{\alpha}{n^{k-1}}$ for some constant α . Thus, for any subgraph of \mathcal{K}_n on e edges and r factor nodes, the probability of it occurring in G is at most $\alpha^r n^{e-r}$.
2. The matrix weight $\overline{\mathbf{M}}_p$ of any self-avoiding path has entries bounded in magnitude by 1 since it is a product of projected stochastic matrices.

Using the above facts, a straightforward calculation tells us:

$$\left| \mathbf{E}_G \left[\left\langle g^{p_v[0]}, \mathbf{w}(G, p) g^{p_v[s]} \right\rangle_{\mathbf{H}} \left\langle g^{\tilde{p}_v[0]}, \mathbf{w}(G, \tilde{p}) g^{\tilde{p}_v[s]} \right\rangle_{\mathbf{H}} \right] \right| \leq (\alpha')^s n^{-(\sum_{e \in p_e^\cup} a(e)-1)}. \quad (10)$$

We say $p_1^\cup \sim p_2^\cup$ if the subgraphs induced by them are isomorphic. \sim partitions the space of all p^\cup into equivalence classes. We use $[p^\cup]$ to denote the equivalence class of p^\cup . The number of equivalence classes can be bounded by $(C's)^{C's}$ for some constant $C' > 1$ (since the graph representing the equivalence class of

$p \cup \tilde{p}$ is on $O(s)$ vertices can be specified by a list of $O(s)$ edges). Due to the shared vertices, p^\cup is connected and $\chi(p^\cup) \leq 1$. Thus, we now bound the variance as follows:

$$\begin{aligned}
& \mathbf{E}_G \left[\left\langle g, \overline{A}_G^{(s)} g \right\rangle_{\mathbf{H}}^2 \right] - \mathbf{E}_G \left[\left\langle g, \overline{A}_G^{(s)} g \right\rangle_{\mathbf{H}} \right]^2 \\
&= \sum_{p \parallel \tilde{p}} \mathbf{E}_G \left[\left\langle g^{p_v[0]}, \mathbf{w}(G, p) g^{p_v[s]} \right\rangle_{\mathbf{H}} \left\langle g^{\tilde{p}_v[0]}, \mathbf{w}(G, \tilde{p}) g^{\tilde{p}_v[s]} \right\rangle_{\mathbf{H}} \right] \\
&\leq \sum_{p^\cup} \sum_{p \parallel \tilde{p} : p \cup \tilde{p} = p^\cup} (\alpha')^s n^{-\sum_{e \in p_e^\cup} a(e)-1} \\
&\leq \sum_{[p^\cup]} \sum_{p_i^\cup \in [p^\cup]} (9\alpha' s)^s n^{-\sum_{e \in p_e^\cup} a(e)-1} \\
&= \sum_{[p^\cup]} (9\alpha' s)^s n^{|p^\cup| - \sum_{e \in p_e^\cup} a(e)-1} \\
&= \sum_{[p^\cup]} (9\alpha' s)^s n^{\chi(p^\cup)} \\
&\leq (9C' \alpha' s^2)^{C'} s n.
\end{aligned}$$

Thus, the claim follows. \square

6.2 Bounding contribution of non-self-avoiding walks

In comparison to the self-avoiding walks, the non-self-avoiding walks have negligible contributions to the expectation and the variance of the statistics. We prove this statement using the following claim.

Claim 6.15. We can bound the statistics of non-self-avoiding walks as follows:

$$\begin{aligned}
& \left| \mathbf{E}_G \left[\sum_{p \text{ non-self-avoiding}} \left\langle g^{p_v[0]}, \mathbf{w}(G, p) g^{p_v[s]} \right\rangle_{\mathbf{H}} \right] \right| \leq (Cs)^{Cs}. \\
& \mathbf{E}_G \left[\sum_{\substack{p \parallel \tilde{p} \\ p \text{ or } \tilde{p} \text{ non-self-avoiding}}} \left\langle g^{p_v[0]}, \mathbf{w}(G, p) g^{p_v[s]} \right\rangle_{\mathbf{H}} \cdot \left\langle g^{\tilde{p}_v[0]}, \mathbf{w}(G, \tilde{p}) g^{\tilde{p}_v[s]} \right\rangle_{\mathbf{H}} \right] \leq (Cs)^{Cs}
\end{aligned}$$

for some absolute constant C .

Proof. The first part is derived by applying [Item 1](#) in the proof of [Lemma 6.14](#).

$$\begin{aligned}
& \left| \mathbf{E}_G \left[\sum_{p \text{ non-self-avoiding}} \left\langle g^{p_v[0]}, \mathbf{w}(G, p) g^{p_v[s]} \right\rangle_{\mathbf{H}} \right] \right| \leq \sum_{p \text{ non-self-avoiding}} \left| \mathbf{E}_G \left[\left\langle g^{p_v[0]}, \mathbf{w}(G, p) g^{p_v[s]} \right\rangle_{\mathbf{H}} \right] \right| \\
&\leq \sum_{p \text{ non-self-avoiding}} \Pr_G[p \in G] \cdot O(1) \\
&\leq \sum_{p \text{ non-self-avoiding}} (\alpha')^s n^{\chi(p) - |p_v|} \\
&\leq \sum_{p \text{ non-self-avoiding}} (\alpha')^s n^{-|p_v|} \\
&\leq (Cs)^{Cs}.
\end{aligned}$$

where the equality from the second to third line is a consequence of [Item 1](#) in the proof of [Lemma 6.17](#), and the last inequality is due to $\chi(p) \leq 0$ for a non-self-avoiding walk.

The second expression is equal to:

$$\sum_{\substack{p \parallel \tilde{p} \\ p \text{ or } \tilde{p} \text{ non-self-avoiding}}} \mathbf{E}_G \left[\left\langle g^{p_v[0]}, \mathbf{w}(G, p) g^{p_v[s]} \right\rangle_{\mathbf{H}} \cdot \left\langle g^{\tilde{p}_v[0]}, \mathbf{w}(G, \tilde{p}) g^{\tilde{p}_v[s]} \right\rangle_{\mathbf{H}} \right].$$

By (10) the above can be bounded by:

$$\sum_{\substack{p \parallel \tilde{p} \\ p \text{ or } \tilde{p} \text{ non-self-avoiding}} (\alpha')^s n^{\chi(p^\cup) - |p_v^\cup|}$$

where p^\cup , recall, is the union of p and \tilde{p} . Since p and \tilde{p} share vertices, $\chi(p^\cup) \leq \min\{\chi(p), \chi(\tilde{p})\}$, and since at least one of the two walks is non-self-avoiding, $\chi(p^\cup) \leq 0$. This lets us bound the above by:

$$\begin{aligned} \sum_{\substack{p \parallel \tilde{p} \\ p \text{ or } \tilde{p} \text{ non-self-avoiding}} (\alpha')^s n^{-|p_v^\cup|} &\leq \sum_{p^\cup} \sum_{p \parallel \tilde{p} : p \cup \tilde{p} = p^\cup} (\alpha')^s n^{-|p_v^\cup|} \\ &= \sum_{[p^\cup]} \sum_{p_i^\cup \in [p^\cup]} (9\alpha' s)^s n^{-|p_v^\cup|} \\ &= \sum_{[p^\cup]} (9\alpha' s)^s \\ &\leq (9C' \alpha' s^2)^{C' s} \end{aligned}$$

which gives us the desired bound in the second part of the statement. \square

6.3 Wrapping up estimates of statistics

The following claim about the statistics of the centered non-backtracking-walk matrix $\underline{A}_G^{(s)}$ is an immediate consequence of combining Lemma 6.13 and the first part of Claim 6.15.

Lemma 6.16. For $s \leq \frac{\log n}{(\log \log n)^2}$:

$$\mathbf{E}_G \left[\left\langle g, \underline{A}_G^{(s)} g \right\rangle_{\mathbf{H}} \right] = (1 \pm o_n(1)) \lambda_L^s n.$$

Lemma 6.17. For some absolute constant C :

$$\mathbf{E}_G \left[\left\langle g, \underline{A}_G^{(s)} g \right\rangle_{\mathbf{H}}^2 \right] - \mathbf{E}_G \left[\left\langle g, \underline{A}_G^{(s)} g \right\rangle_{\mathbf{H}} \right]^2 \leq n(Cs)^{Cs}.$$

Proof. Using the observation that only pairs of walks that share some vertices contribute to the variance, we have:

$$\begin{aligned} &\mathbf{E}_G \left[\left\langle g, \underline{A}_G^{(s)} g \right\rangle_{\mathbf{H}}^2 \right] - \mathbf{E}_G \left[\left\langle g, \underline{A}_G^{(s)} g \right\rangle_{\mathbf{H}} \right]^2 \\ &= \mathbf{E}_G \left[\left\langle g, \overline{A}_G^{(s)} g \right\rangle_{\mathbf{H}}^2 \right] - \mathbf{E}_G \left[\left\langle g, \overline{A}_G^{(s)} g \right\rangle_{\mathbf{H}} \right]^2 \\ &\quad + \sum_{\substack{p \parallel \tilde{p} \\ p \text{ or } \tilde{p} \text{ non-self-avoiding}}} \mathbf{E}_G \left[\left\langle g^{\eta(p_v^0)}, \mathbf{w}(G, p, \eta) g^{\eta(p_v^s)} \right\rangle_{\mathbf{H}} \cdot \left\langle g^{\eta(\tilde{p}_v^0)}, \mathbf{w}(G, \tilde{p}, \eta) g^{\eta(\tilde{p}_v^s)} \right\rangle_{\mathbf{H}} \right] \end{aligned}$$

We can conclude the desired bound immediately from Lemma 6.14 and the second part of Claim 6.15. \square

Finally, to establish [Theorem 6.1](#) we first use Chebyshev's inequality to conclude that when $s \leq \frac{\log n}{(\log \log n)^2}$,

$$\left\langle g, \underline{A}_G^{(s)} g \right\rangle_{\mathbf{H}} = (1 \pm o_n(1)) \lambda_L^s n$$

with probability $1 - o_n(1)$. Now, since the $\|g\|^2$ is $(1 \pm o_n(1)) \frac{n}{\gamma}$ with probability $1 - o_n(1)$ for some constant γ , we can conclude that with probability $1 - o_n(1)$,

$$\frac{\left\langle g, \underline{A}_G^{(s)} g \right\rangle_{\mathbf{H}}}{\|g\|^2} = (1 \pm o_n(1)) \gamma \lambda_L^s$$

thereby finishing the proof.

7 Eigenvalue bounds

In this section we show an eigenvalue upper bound for the centered nonbacktracking-walk matrix in the null model. We first describe the matrix distribution in detail.

Recall the definition of a null model M^* ([Definition 2.3](#)). Let $\mathbf{H} := \bigcup_{i=1}^F E_i$ be an observation sampled from M^* . As discussed in [Section 2.1.1](#), the observation \mathbf{H} has an associated bipartite graph which we denote $\text{Bip}(\mathbf{H})$. The left vertex set is given by the variables $[n]$ and the right vertex set is given by the factors $\gamma \in \mathbf{H}$. We will use $L(\mathbf{H})$ and $R(\mathbf{H})$ to denote the left and right vertex sets of $\text{Bip}(\mathbf{H})$.

Next associate with each triple $v\gamma u$, where $v, u \in \gamma$ and $v \neq u$, a $q \times q$ matrix $M_{v\gamma u}$. Like before, the value of the matrix only depends on the factor type $\theta(\gamma)$ and the two variables' indices $i(v), i(u)$ in γ . We use Bl to denote the collection of these $q \times q$ matrices $\{M_{a\phi_i b}\}_{i \in [F], a \neq b \in [a(i)]}$. Now we are ready to define the matrix distribution.

Definition 7.1. The matrix distribution is defined as follows. First sample an observation H from some null model M^* . We define the *length- ℓ M^* -centered nonbacktracking power* $\underline{A}_H^{(\ell)}$ is the $n \times n$ block matrix where the (i, j) -block as the following $q \times q$ matrix:

$$\underline{A}_H^{(\ell)}[i, j] := \sum_{\substack{(v_0 \gamma_1 v_1 \dots v_{\ell-1} \gamma_{\ell} v_{\ell}) \\ \in \text{NB}(\mathcal{K}_n, \ell, i, j)}} \prod_{t=1}^{\ell} M_{v_{2t-2} \gamma_t v_{2t-1}} (\mathbf{1}[\gamma_t \in H] - \Pr_{M^*}[\gamma_t \in H]),$$

where $\text{NB}(\mathcal{K}_n, \ell, i, j)$ denote the set of all length- 2ℓ nonbacktracking walks in the complete bipartite factor graph $\text{Bip}(\mathcal{K}_n)$ starting at variable i and ending at variable j .

We are interested in obtaining a high probability upper bound on $\|\underline{A}_H^{(\ell)}\|$ in terms of a particular quantity depending on Bl and M^* , which we denote by $\rho(\text{Bl}, M^*)$. Before giving its definition, we simplify the notation a bit:

Definition 7.2. Given a length- 2ℓ nonbacktracking walk $W = v_0 \gamma_1 v_1 \gamma_2 v_2 \dots v_{\ell-1} \gamma_{\ell} v_{\ell}$ in $\text{Bip}(\mathbf{H})$, we define the weight of the walk with M_W to be

$$M_W := M_{v_0 \gamma_1 v_1} M_{v_1 \gamma_2 v_2} \dots M_{v_{\ell-1} \gamma_{\ell} v_{\ell}}.$$

We use $\gamma_i(W)$ to denote the i -th factor visited by W .

We now define $\rho(\text{Bl}, M^*)$.

Definition 7.3. For a positive integer m , construct a length- m path P with random matrix weights in the following way. Start with a vertex v_0 and assign to it a random label $\tau(v_0)$ sampled from \mathcal{T} . We iteratively construct a path with edges weighed by matrices until its length is equal to m . Suppose we already have a path $v_0 v_1 \dots v_t$ where each v_i has a label $\tau(v_i)$, and for an edge $\{v_i, v_{i+1}\}$ its two directed edges (v_i, v_{i+1}) and (v_{i+1}, v_i) have matrix weights $\mathbf{W}_{i,i+1}$ and $\mathbf{W}_{i+1,i}$ respectively. To grow this path, we sample (s, a) where $s \in [F]$ and $a \in [a(s)]$ with probability proportional to $\overline{\phi}_s \cdot \mathbf{1}[\text{Cl}(s)_a = \tau(v_t)] \cdot \prod_{j=1, j \neq a}^t \mathcal{T}(\text{Cl}(s)_j)$, followed by a uniformly random b in $[a(s)] \setminus \{a\}$. Add vertex v_{t+1} and set $\tau(v_{t+1}) = \text{Cl}(s)_b$. Then add edge $\{v_t, v_{t+1}\}$ and let the matrix weight of the directed edge (v_t, v_{t+1}) be $\mathbf{W}_{t,t+1} = M_{a\phi_s b} \in \text{Bl}$, and matrix weight of directed edge (v_{t+1}, v_t) be $\mathbf{W}_{t+1,t} = M_{b\phi_s a} \in \text{Bl}$.⁵ We then define w_m as:

$$w_m := \text{ETr}(\mathbf{W}_{0,1} \mathbf{W}_{1,2} \dots \mathbf{W}_{m-1,m} \mathbf{W}_{m,m-1} \dots \mathbf{W}_{2,1} \mathbf{W}_{1,0}).$$

Now define $r(\text{Bl}, \mathbf{M}^*)$ as

$$r(\text{Bl}, \mathbf{M}^*) := \limsup_{m \rightarrow \infty} w_m^{\frac{1}{2m}}.$$

Next, define $d(\mathbf{M}^*)$ (which, intuitively, is the average degree of a vertex in a sample from \mathbf{M}^*) as

$$d(\mathbf{M}^*) := \sum_{t=1}^T \mathcal{T}(t) \sum_{i=1}^F \sum_{j=1}^{k_i} \overline{\phi}_i \mathbf{1}[\text{Cl}(i)_j = t].$$

Finally, we define $\rho(\text{Bl}, \mathbf{M}^*)$ as

$$\rho(\text{Bl}, \mathbf{M}^*) := r(\text{Bl}, \mathbf{M}^*) \sqrt{d(\mathbf{M}^*)}.$$

We remark that if the weight collection Bl is defined such that $M_{b\phi_i a} = \overline{M}_{i,a|b}$ (defined in (9)) for all $i \in [F]$, $a \neq b \in [a(i)]$, then $\rho(\text{Bl}, \mathbf{M}^*) = \sqrt{\lambda_L}$.

The main result of this section is that $\|A_H^{(\ell)}\| \leq ((1 + o_n(1))\rho(\text{Bl}, \mathbf{M}^*))^\ell$ for a wide range of ℓ when $\mathbf{H} \sim \mathbf{M}^*$.

Theorem 7.4. *Suppose:*

1. *There is an inner product $\langle \cdot, \cdot \rangle_\nu$ on \mathbb{R}^{nq} such that for every right vertex $\gamma = (v_1, \dots, v_{a(i)})$ in $\mathcal{K}_{\text{Bl},n}$, for all $1 \leq a, b \leq a(i)$, the $nq \times nq$ matrix obtained by placing $M_{v_a \gamma v_b}$ in the (v_a, v_b) block and zeros everywhere else is the adjoint of the $nq \times nq$ matrix obtained by placing $M_{v_b \gamma v_a}$ in the (v_b, v_a) block and zeroes everywhere else under $\langle \cdot, \cdot \rangle_\nu$.*
2. *There is a constant $C \geq 1$ such that the weight M_W every nonbacktracking walk W in \mathcal{K}_n satisfies:*

$$\|M_W\| \leq C$$

where $\|\cdot\|$ is the operator norm induced by $\langle \cdot, \cdot \rangle_\nu$.

3. $\rho(\text{Bl}, \mathbf{M}^*) \geq 1$.

Then for every $\varepsilon > 0$ and $(\log \log n)^2 \leq \ell \leq \frac{\log n}{(\log \log n)^2}$, with probability $1 - o_n(1)$:

$$|\lambda|_{\max} \left(A_H^{(\ell)} \right) \leq ((1 + \varepsilon)\rho(\text{Bl}, \mathbf{M}^*))^\ell.$$

⁵For the sake of intuition the reader should think of the distribution of v_{t+1} as first sampling $\mathbf{H} \sim \mathbf{M}^*$, then choosing a random vertex v with label $\tau(v_t)$ and finally choosing as a random neighbor w of v within \mathbf{H} . $\tau(v_{t+1})$ is then set to the label of w and the matrix weight on edge (v_t, v_{t+1}) is chosen as the matrix weight on (v, w) .

7.1 Proof of Theorem 7.4

The proof of Theorem 7.4 is via the *trace method*. One preliminary observation is that Condition 1 implies that $\underline{A}_H^{(\ell)}$ is self-adjoint and hence all its eigenvalues are real. Consequently, for any positive even integer k :

$$\|\underline{A}_H^{(\ell)}\|^k \leq \text{Tr} \left(\left(\underline{A}_H^{(\ell)} \right)^k \right). \quad (11)$$

Our goal is now to obtain a handle on $S := \text{Tr} \left(\left(\underline{A}_H^{(\ell)} \right)^k \right)$ and obtain a high probability bound on it. We borrow some terminology from [MOP20]:

Definition 7.5 (Linkages). A $(k \times 2\ell)$ -nonbacktracking BI-linkage is a length- $2k\ell$ closed walk in $\text{Bip}(\mathcal{K}_n)$ that starts and ends in the left vertex set and can be expressed as a concatenation of k nonbacktracking walks of length- 2ℓ each. Each length- 2ℓ nonbacktracking segment is called a *link*. We use $\text{Lkgs}(\text{Bl}, n, k, \ell)$ to denote the collection of all $(k \times 2\ell)$ -nonbacktracking BI-linkages.

Definition 7.6. Given a $(k \times 2\ell)$ -nonbacktracking BI-linkage W , we use $L(W)$ to denote the set of left vertices visited by W , $R(W)$ to denote the set of right vertices visited by W , $V(W)$ to denote $L(W) \cup R(W)$, $E(W)$ to denote the set of edges visited by W , and $G(W)$ to denote the graph $(V(W), E(W))$ induced by W .

With the above terminology and notation in hand, we can write S as:

$$S = \sum_{W \in \text{Lkgs}(\text{Bl}, n, k, \ell)} \text{Tr}(M_W) \prod_{t=1}^{k\ell} \left(\mathbf{1}[\gamma_t(W) \in H] - \mathbf{E}_{H|\tau} \mathbf{1}[\gamma_t(W) \in H] \right).$$

A natural strategy to obtaining a high probability bound on S is to bound $\mathbf{E}[S]$ by some Z and use Markov's inequality to conclude that S is bounded by, say, nZ with high probability. However, $\mathbf{E}[S]$ is not as small as we would hope due to blowing up in magnitude owing to the occurrences of certain rare and problematic subgraphs in $\text{Bip}(H)$. So this suggests a natural tweak of conditioning away these rare subgraphs and trying to carry out the same strategy. This tweak is an idea that occurs in many previous papers in the line of work on getting eigenvalue bounds on sparse random matrices [Fri03, BLM15, Bor19, BC19, MOP20]. These problematic subgraphs all share one common trait – having multiple cycles in a small neighborhood.

Definition 7.7. We say a graph Γ is *r-bicycle free* if for every vertex v , the radius- r ball around v contains at most one cycle. We say Γ is an *r-bicycle* if it has at most r edges and has at least two cycles, and we say Γ is an *r-bicycle frame* if it is an *r-bicycle* such that no subgraph of it is an *r-bicycle*.

Lemma 7.8. With probability $1 - o_n(1)$, $\text{Bip}(H)$ is *r-bicycle free* for $r = \frac{\log n}{\log \log n}$.

We refer the reader to Corollary D.9 for a proof of this fact.

Henceforth, we use \mathcal{E} to denote the event that H is *r-bicycle free* for $r = \frac{\log n}{\log \log n}$. Now define $U := S \cdot \mathbf{1}[\mathcal{E}]$. By Lemma 7.8 with probability $1 - o_n(1)$, $S = U$ so if we can prove that $U \leq Z$ with probability $1 - o_n(1)$, we can also show that $S \leq Z$ with probability $1 - o_n(1)$. Thus, we turn our attention to bounding $\mathbf{E}[U]$.

$$U = \sum_{W \in \text{Lkgs}(\text{Bl}, n, k, \ell)} \text{Tr}(M_W) \prod_{t=1}^{k\ell} \left(\mathbf{1}[\gamma_t(W) \in H] - \mathbf{E}_{H|\tau} \mathbf{1}[\gamma_t(W) \in H] \right) \mathbf{1}[\mathcal{E}]. \quad (12)$$

We now study the quantity $\prod_{t=1}^{k\ell} \left(\mathbf{1}[\gamma_t(W) \in H] - \mathbf{E}_{H|\tau} \mathbf{1}[\gamma_t(W) \in H] \right)$ more closely.

Definition 7.9. For a given right vertex γ of \mathcal{K}_n the *multiplicity* $m_W(\gamma)$ of γ in W is the number of times γ is visited by W . $S(W)$ denotes the set of all right vertices that are visited exactly once and are called *singleton right vertices*. $D(W)$ denotes the set of all right vertices that are visited more than once, and are called *duplicative right vertices*.

Henceforth, we shorten $\mathbf{1}[\gamma \in H]$ to $\mathbf{1}_\gamma$ and $\mathbf{E}_{H|\tau}\mathbf{1}[\gamma \in H]$ to μ_γ . Thus:

$$\begin{aligned} \prod_{t=1}^{k\ell} (\mathbf{1}_{\gamma_t(W)} - \mu_{\gamma_t(W)}) &= \prod_{\gamma \in S(W)} (\mathbf{1}_\gamma - \mu_\gamma) \prod_{\gamma \in D(W)} (\mathbf{1}_\gamma - \mu_\gamma)^{m_W(\gamma)} \\ &= \prod_{\gamma \in S(W)} (\mathbf{1}_\gamma - \mu_\gamma) \prod_{\gamma \in D(W)} \left(\mathbf{1}_\gamma \cdot \sum_{i=1}^{m_W(\gamma)} (-\mu_\gamma)^{m_W(\gamma)-i} \binom{m_W(\gamma)}{i} + (-\mu_\gamma)^{m_W(\gamma)} \right) \end{aligned}$$

To lighten notation, we use α_γ to denote $\sum_{i=1}^{m_W(\gamma)} (-\mu_\gamma)^{m_W(\gamma)-i} \binom{m_W(\gamma)}{i}$.

$$\begin{aligned} &= \prod_{\gamma \in S(W)} (\mathbf{1}_\gamma - \mu_\gamma) \prod_{\gamma \in D(W)} (\mathbf{1}_\gamma \alpha_\gamma + (-\mu_\gamma)^{m_W(\gamma)}) \\ &= \prod_{\gamma \in S(W)} (\mathbf{1}_\gamma - \mu_\gamma) \sum_{L \subseteq D(W)} \prod_{\gamma \in L} \mathbf{1}_\gamma \alpha_\gamma \prod_{\gamma \in D(W) \setminus L} (-\mu_\gamma)^{m_W(\gamma)} \\ &= \sum_{L \subseteq D(W)} \prod_{\gamma \in L} \alpha_\gamma \prod_{\gamma \in D(W) \setminus L} (-\mu_\gamma)^{m_W(\gamma)} \prod_{\gamma \in S(W)} (\mathbf{1}_\gamma - \mu_\gamma) \prod_{\gamma \in L} \mathbf{1}_\gamma. \end{aligned} \quad (13)$$

Plugging in (13) into (12) gives:

$$U = \sum_{W \in \text{Lkgs}(\text{Bl}, n, k, \ell)} \text{Tr}(M_W) \sum_{L \subseteq D(W)} \prod_{\gamma \in L} \alpha_\gamma \prod_{\gamma \in D(W) \setminus L} (-\mu_\gamma)^{m_W(\gamma)} \prod_{\gamma \in S(W)} (\mathbf{1}_\gamma - \mu_\gamma) \prod_{\gamma \in L} \mathbf{1}_\gamma \mathbf{1}[\mathcal{E}]$$

We are interested in understanding $\mathbf{E}[U]$. Note that this is equal to $\mathbf{E}_\tau \mathbf{E}_{H|\tau}[U]$. We will first focus our attention on understanding $\mathbf{E}_{H|\tau}[U]$. We have:

$$\begin{aligned} \mathbf{E}_{H|\tau}[U] &= \sum_{W \in \text{Lkgs}(\text{Bl}, n, k, \ell)} \text{Tr}(M_W) \sum_{L \subseteq D(W)} \prod_{\gamma \in L} \alpha_\gamma \prod_{\gamma \in D(W) \setminus L} (-\mu_\gamma)^{m_W(\gamma)} \mathbf{E}_{H|\tau} \left[\prod_{\gamma \in S(W)} (\mathbf{1}_\gamma - \mu_\gamma) \prod_{\gamma \in L} \mathbf{1}_\gamma \mathbf{1}[\mathcal{E}] \right] \\ &\leq \sum_{W \in \text{Lkgs}(\text{Bl}, n, k, \ell)} \text{Tr}(M_W) \sum_{L \subseteq D(W)} \prod_{\gamma \in L} |\alpha_\gamma| \prod_{\gamma \in D(W) \setminus L} \mu_\gamma^{m_W(\gamma)} \left| \mathbf{E}_{H|\tau} \left[\prod_{\gamma \in S(W)} (\mathbf{1}_\gamma - \mu_\gamma) \prod_{\gamma \in L} \mathbf{1}_\gamma \mathbf{1}[\mathcal{E}] \right] \right|. \end{aligned}$$

Notice that $\alpha_\gamma = (1 - \mu_\gamma)^{m_W(\gamma)} - (-\mu_\gamma)^{m_W(\gamma)}$ and hence $|\alpha_\gamma| \leq (1 - \mu_\gamma)^{m_W(\gamma)} + \mu_\gamma^{m_W(\gamma)} \leq (1 - \mu_\gamma) + \mu_\gamma = 1$ where the second inequality is a consequence of $\mu_\gamma \in [0, 1]$. The result is:

$$\mathbf{E}_{H|\tau}[U] \leq \sum_{W \in \text{Lkgs}(\text{Bl}, n, k, \ell)} \text{Tr}(M_W) \sum_{L \subseteq D(W)} \prod_{\gamma \in D(W) \setminus L} \mu_\gamma^{m_W(\gamma)} \left| \mathbf{E}_{H|\tau} \left[\prod_{\gamma \in S(W)} (\mathbf{1}_\gamma - \mu_\gamma) \prod_{\gamma \in L} \mathbf{1}_\gamma \mathbf{1}[\mathcal{E}] \right] \right|. \quad (14)$$

Next, we would like to obtain a bound on $\left| \mathbf{E}_{H|\tau} \left[\prod_{\gamma \in S(W)} (\mathbf{1}_\gamma - \mu_\gamma) \prod_{\gamma \in L} \mathbf{1}_\gamma \mathbf{1}[\mathcal{E}] \right] \right|$.

Towards doing so, we first set up a couple of definitions and an observation.

Definition 7.10 (Closure of subgraph). Given a subgraph Γ of $\text{Bip}(\mathcal{K}_n)$ with right vertex set $R(\Gamma)$, we define its *closure* $\text{Clos}(\Gamma)$ as the induced subgraph on the vertex set $(\bigcup_{\gamma \in R(\Gamma)} N(\gamma)) \cup R(\Gamma)$. We say Γ is *closed* if $\text{Clos}(\Gamma) = \Gamma$.

Definition 7.11 (Excess). Given a graph Γ on e edges, v vertices and c connected components, we define the *excess* of Γ , denoted $\text{Exc}(\Gamma)$, to be $e - v + c$.

The following is immediate from the observation that the excess of a graph cannot decrease on adding a new vertex or a new edge.

Lemma 7.12. *If $\Gamma = (V, E)$ and $\Gamma' = (V', E')$ are two graphs such that Γ is a subgraph of Γ' , i.e. $V \subseteq V'$ and $E \subseteq E'$, then $\text{Exc}(\Gamma') \geq \text{Exc}(\Gamma)$.*

If $\text{Clos}(L)$ is not r -bicycle free, then $\prod_{\gamma \in L} \mathbf{1}_\gamma \mathbf{1}[\mathcal{E}]$ is equal to 0. Otherwise [Lemma D.11](#) then shows that:

$$\left| \mathbf{E}_{H|\tau} \left[\prod_{\gamma \in S(W)} (1_\gamma - \mu_\gamma) \prod_{\gamma \in L} \mathbf{1}_\gamma \mathbf{1}[\mathcal{E}] \right] \right| \leq \prod_{\gamma \in S \cup L} \mu_\gamma \cdot 2^{|S(W)|} \left(\frac{1}{n^5} \right)^{\frac{|S(W)|}{r} - \text{Exc}(\text{Clos}(S(W) \cup L))}.$$

Plugging the above into [\(14\)](#) tells us:

$$\mathbf{E}_{H|\tau}[U] \leq \sum_{W \in \text{Lkgs}(\text{Bl}, n, k, \ell)} \text{Tr}(M_W) \sum_{\substack{L \subseteq D(W) \\ \text{Clos}(L) \text{ } r\text{-bicycle free}}} \prod_{\gamma \in R(W)} \mu_\gamma \prod_{\gamma \in D(W) \setminus L} \mu_\gamma^{m_W(\gamma) - 1} 2^{|S(W)|} \left(\frac{1}{n^5} \right)^{\frac{|S(W)|}{r} - \text{Exc}(\text{Clos}(S(W) \cup L))}.$$

Henceforth, we will shorten $\text{Exc}(\text{Clos}(S(W) \cup D(W)))$ to Exc_W for simplicity of notation. Since $\text{Clos}(S(W) \cup L)$ is a subgraph of $\text{Clos}(S(W) \cup D(W))$, by [Lemma 7.12](#) we have $\text{Exc}_W \geq \text{Exc}(\text{Clos}(S(W) \cup L))$, which means:

$$\begin{aligned} \mathbf{E}_{H|\tau}[U] &\leq \sum_{W \in \text{Lkgs}(\text{Bl}, n, k, \ell)} \text{Tr}(M_W) \sum_{\substack{L \subseteq D(W) \\ \text{Clos}(L) \text{ } r\text{-bicycle free}}} \prod_{\gamma \in R(W)} \mu_\gamma \prod_{\gamma \in D(W) \setminus L} \mu_\gamma^{m_W(\gamma) - 1} 2^{|S(W)|} \left(\frac{1}{n^5} \right)^{\frac{|S(W)|}{r} - \text{Exc}_W} \\ &= \sum_{W \in \text{Lkgs}(\text{Bl}, n, k, \ell)} \text{Tr}(M_W) \cdot 2^{|S(W)|} \left(\frac{1}{n^5} \right)^{\frac{|S(W)|}{r} - \text{Exc}_W} \prod_{\gamma \in R(W)} \mu_\gamma \sum_{\substack{L \subseteq D(W) \\ \text{Clos}(L) \text{ } r\text{-bicycle free}}} \prod_{\gamma \in D(W) \setminus L} \mu_\gamma^{m_W(\gamma) - 1} \end{aligned} \quad (15)$$

Next, we focus on bounding $\sum_{\substack{L \subseteq D(W) \\ \text{Clos}(L) \text{ } r\text{-bicycle free}}} \prod_{\gamma \in D(W) \setminus L} \mu_\gamma^{m_W(\gamma) - 1}$. For starters, observe that:

$$\sum_{\substack{L \subseteq D(W) \\ \text{Clos}(L) \text{ } r\text{-bicycle free}}} \prod_{\gamma \in D(W) \setminus L} \mu_\gamma^{m_W(\gamma) - 1} \leq \sum_{\substack{L \subseteq D(W) \\ \text{Clos}(L) \text{ } r\text{-bicycle free}}} \prod_{\gamma \in D(W) \setminus L} \left(\frac{\bar{\phi}_{\max}}{n} \right)^{m_W(\gamma) - 1}$$

We proceed to bound this in a manner identical to [\[BMR19\]](#). We define a weight function w on subsets of $D(W)$ as follows: $w(K) = \sum_{\gamma \in K} m_W(\gamma) - 1$. Choose $D^*(W)$ as a maximum weight subset (according to w) of W such that $\text{Clos}(D^*(W))$ is r -bicycle free, and let $\Delta(W) := w(D(W)) - w(D^*(W))$. Note that for any $L \subseteq D(W)$ such that $\text{Clos}(L)$ is r -bicycle free, $\Delta(W) \leq w(D(W) \setminus L)$.

$$\sum_{\substack{L \subseteq D(W) \\ \text{Clos}(L) \text{ } r\text{-bicycle free}}} \prod_{\gamma \in D(W) \setminus L} \left(\frac{\bar{\phi}_{\max}}{n} \right)^{m_W(\gamma) - 1} \leq \sum_{\substack{L \subseteq D(W) \\ \text{Clos}(L) \text{ } r\text{-bicycle free}}} \left(\frac{\bar{\phi}_{\max}}{n} \right)^{w(D(W) \setminus L)}$$

Since $L \subseteq D(W)$ every $m_W(\gamma) - 1 \geq 1$. Using this along with $\Delta(W) \leq w(D(W) \setminus L)$ we can bound the above by:

$$\begin{aligned} &\leq \sum_{L \subseteq D(W)} \left(\frac{\bar{\phi}_{\max}}{n} \right)^{\max\{|D(W) \setminus L|, \Delta(W)\}} \\ &= \sum_{i \leq \Delta(W)} \left(\frac{\bar{\phi}_{\max}}{n} \right)^{\Delta(W)} \cdot \binom{|D(W)|}{i} + \sum_{i > \Delta(W)} \left(\frac{\bar{\phi}_{\max}}{n} \right)^i \binom{|D(W)|}{i} \end{aligned}$$

$$\begin{aligned}
&\leq (\Delta(W) + 1) \left(\frac{\bar{\phi}_{\max}}{n} \right)^{\Delta(W)} + \sum_{i > \Delta(W)} \left(\frac{\bar{\phi}_{\max} |D(W)|}{n} \right)^i \\
&\leq (\Delta(W) + 2) \left(\frac{\bar{\phi}_{\max}}{n} \right)^{\Delta(W)} \\
&\leq 2 \left(\frac{2\bar{\phi}_{\max}}{n} \right)^{\Delta(W)}
\end{aligned}$$

Plugging this back into (15) gives us:

$$(15) \leq 2 \sum_{W \in \text{Lkgs}(\text{Bl}, n, k, \ell)} \text{Tr}(M_W) \cdot 2^{|\text{S}(W)|} \cdot \left(\frac{1}{n^5} \right)^{\frac{|\text{S}(W)|}{r} - \text{Exc}_W} \cdot \left(\frac{2\bar{\phi}_{\max}}{n} \right)^{\Delta(W)} \cdot \prod_{\gamma \in R(W)} \mu_\gamma$$

As a first step towards simplifying the above quantity we make the following definition:

Definition 7.13 (Shape of a linkage). Given a $(k \times 2\ell)$ -nonbacktracking linkage $W = v_0 v_1 \dots v_{2k\ell}$ that visits v distinct vertices, we say the *shape* of W denoted $\text{Sh}(W)$ is the $(k \times 2\ell)$ -nonbacktracking linkage on graph on vertex set $[v]$ obtained by first constructing map $\xi : V(W) \rightarrow [2k\ell]$ where $\xi(v) = i$ where v is the i -th distinct vertex visited by W and defining the t -th step of the walk $\text{Sh}(W)$ to be $\xi(v_{t-1})\xi(v_t)$. We say the left vertex set of $\text{Sh}(W)$ is $\xi(L(W))$ and the right vertex set is $\xi(R(W))$. For $a \in V(\text{Sh}(W))$ we will use the notation $W[a]$ to denote $\xi^{-1}(a)$.

We use $\text{Shps}(k, \ell)$ to denote the set of all distinct shapes of linkages in $\text{Lkgs}(\text{Bl}, n, k, \ell)$, and $\text{Shps}(k, \ell, v, e)$ to denote the set of all shapes in $\text{Shps}(k, \ell)$ on v vertices and e edges.

We can rewrite the bound on (15) as:

$$(15) \leq 2 \sum_{\text{Sh} \in \text{Shps}(k, \ell)} 2^{|\text{S}(\text{Sh})|} \cdot \left(\frac{1}{n^5} \right)^{\frac{|\text{S}(\text{Sh})|}{r} - \text{Exc}_{\text{Sh}}} \cdot \left(\frac{2\bar{\phi}_{\max}}{n} \right)^{\Delta(\text{Sh})} \sum_{\substack{W \in \text{Lkgs}(\text{Bl}, n, k, \ell) \\ \text{Sh}(W) = \text{Sh}}} \text{Tr}(M_W) \cdot \prod_{\gamma \in R(W)} \mu_\gamma. \quad (16)$$

For a given linkage W , we begin by deriving an upper bound on $\text{Tr}(M_W)$. A preliminary observation is:

Observation 7.14. $\text{Tr}(M_W) \leq q \|M_W\|$.

Our next step is to decompose W into simpler “subwalks”. This segment of the argument follows [BC19, OW20]

Definition 7.15. We call a vertex v in $L(W)$ a *landmark* of W if it satisfies at least one of the following conditions: (i) is an endpoint of a link, (ii) $\deg_{G(W)}(v) \geq 3$, (iii) $\deg_{G(W)}(w) \geq 3$ for some $w \in R(W)$ which is incident to v within $G(W)$. We refer to the set of all landmark vertices in W as $\text{Lm}(W)$. We call any path between two landmark vertices v_1 and v_2 with no intermediate landmark vertices a *trail*. We call a trail a *forked trail* if it has an intermediate vertex w in $R(W)$ such that $\deg_{G(W)}(w) \geq 3$, and an *unforked trail* otherwise. We use $\text{Trs}(W)$ to denote the collection of all trails in W , $\text{UTrs}(W)$ to denote the collection of all unforked trails in W and $\text{FTrs}(W)$ to denote the collection of all forked trails in W .

Observation 7.16. Any forked trail must be a single two-step of the form $u\gamma v$ where u and v are left vertices and γ is a right vertex.

Any $W \in \text{Lkgs}(\text{Bl}, n, k, \ell)$ is a sequence of nonbacktracking walks on trails. W can be written as the sequence of vertices visited $v_0 \gamma_1 v_1 \dots \gamma_{k\ell} v_{k\ell}$. Let T be the set of all times t such that v_t is a landmark. Using T , we construct a set of *pause times* P in the following way:

For each $t \in T$, if the trail starting or ending at time t is visited for the first or second time, we add t to P .

Recall from [Definition 7.2](#) that M_W is the product of $k\ell$ matrices $M_{v_0 y_1 v_1} \dots M_{v_{k\ell-1} y_{k\ell} v_{k\ell}}$. Let $p_1 < \dots < p_s$ be the sequence of all pause times. By submultiplicativity of the operator norm,

$$\|M_W\| \leq \|M_{v_0 y_1 v_1} \dots M_{v_{p_1-1} y_{p_1} v_{p_1}}\| \cdot \|M_{v_{p_1} y_{p_1+1} v_{p_1+1}} \dots M_{v_{p_2-1} y_{p_2} v_{p_2}}\| \dots \|M_{v_{p_s} y_{p_s+1} v_{p_s+1}} \dots M_{v_{k\ell-1} y_{k\ell} v_{k\ell}}\|.$$

Each segment between consecutive pauses p_i and p_{i+1} falls into one of the following categories:

- $\text{Seg}_{\leq 2}(W)$: the segment is composed of exactly one trail and is the first or second visit to the trail.
- $\text{Seg}_{> 2}(W)$ the segment is a union of trails and each of these trails has already been visited at least twice before.

Rewriting the above upper bound, we now have:

$$\|M_W\| \leq \prod_{\omega \in \text{Seg}_{\leq 2}(W)} \|M_\omega\| \cdot \prod_{\omega \in \text{Seg}_{> 2}(W)} \|M_\omega\|.$$

Given a trail T with endpoints u and v , there are two nonbacktracking walks ω_1 and ω_2 that cover T , one from u to v and another from v to u . By [Condition 1](#), $M_{\omega_1}^* = M_{\omega_2}$ where the $*$ in the superscript refers to the adjoint induced by the inner product $\langle \cdot, \cdot \rangle_v$ and so $\|M_{\omega_1}\| = \|M_{\omega_2}\|$. Henceforth we use $\|M_T\|$ to denote $\|M_{\omega_1}\| = \|M_{\omega_2}\|$. Using the notation $\text{UTrs}_{\geq 2}(W)$ for unforked trails that are visited more than once, we can write the above as:

$$\|M_W\| \leq \prod_{T \in \text{UTrs}_{\geq 2}(W)} \|M_T\|^2 \cdot \prod_{\omega \in \text{Seg}_{> 2}(W) \cup \text{Seg}_1(W) \cup \text{UTrs}_{\geq 2}(W)} \|M_\omega\|.$$

[Condition 2](#) further lets us get the following bound:

$$\|M_W\| \leq C^{|\text{Seg}_1(W)| + |\text{Seg}_{\geq 2}(W)|} \cdot \prod_{T \in \text{UTrs}_{\geq 2}(W)} \|M_T\|^2 \leq C^{|P|+1} \cdot \prod_{T \in \text{UTrs}_{\geq 2}(W)} \|M_T\|^2 \quad (17)$$

We now turn our attention to bounding $|P|$. Given a landmark vertex v and an edge e incident to it, there are at most a_{\max} trails that start at v and tread on e on their first step, and hence the number of distinct trails starting at v is at most $a_{\max} \cdot \deg_{G(W)}(v)$. By the construction of P , the number of pauses at vertex v is at most twice the number of distinct trails starting at v , and hence is at most $2a_{\max} \cdot \deg_{G(W)}(v)$. Thus:

$$|P| \leq 2a_{\max} \sum_{v \in \text{Lm}(W)} \deg_{G(W)}(v) \leq 2a_{\max} \left(2|\text{Lm}_{\leq 2}(W)| + \sum_{v \in \text{Lm}_{\geq 3}(W)} \deg_{G(W)}(v) \right)$$

where $\text{Lm}_{\leq 2}(W)$ and $\text{Lm}_{\geq 3}(W)$ denote the sets of landmark vertices of degree ≤ 2 and degree ≥ 3 respectively. Each vertex in $\text{Lm}_{\leq 2}(v)$ is either an endpoint of a link or a neighbor of a degree ≥ 3 right vertex. There are exactly $k+1$ endpoints of links, and at most $\sum_{v \in G(W): \deg_{G(W)}(v) \geq 3} \deg_{G(W)}(v)$ landmark vertices induced as neighbors of degree ≥ 3 right vertices, and hence:

$$|P| \leq 2a_{\max} \left(2(k+1) + 2 \sum_{v \in G(W): \deg_{G(W)}(v) \geq 3} \deg_{G(W)}(v) + \sum_{v \in G(W): \deg_{G(W)}(v) \geq 3} \deg_{G(W)}(v) \right) \quad (18)$$

It remains to bound $\sum_{v \in G(W): \deg_{G(W)}(v) \geq 3} \deg_{G(W)}(v)$. Let X be a set of edges of size Exc_W such that $T(W, X) := (V(W), E(W) \setminus X)$ is a tree. Since $G(W)$ has at most k leaves, $T(W, X)$ has at most $k + 2\text{Exc}_W$ leaves. We now state the following well known fact about trees and refer the reader to [\[BMR19, Fact 6.35\]](#) for a proof.

Fact 7.17. Let T be a tree with l leaves. Then $3l \geq \sum_{v \in V(T) : \deg_T(v) \geq 3} \deg_T(v)$.

As a consequence of [Fact 7.17](#):

$$\sum_{v \in V(T(W,X)) : \deg_{T(W,X)}(v) \geq 3} \deg_{T(W,X)}(v) \leq 3(k + 2\text{Exc}_W).$$

Now observe that for any graph Γ and graph Γ' obtained by adding a single edge to Γ :

$$\sum_{v \in V(\Gamma') : \deg_{\Gamma'}(v) \geq 3} \deg_{\Gamma'}(v) \leq \left(\sum_{v \in V(\Gamma) : \deg_{\Gamma}(v) \geq 3} \deg_{\Gamma}(v) \right) + 6,$$

and thus

$$\sum_{v \in G(W) : \deg_{G(W)}(v) \geq 3} \deg_{G(W)}(v) \leq 3(k + 2\text{Exc}_W) + 6\text{Exc}_W = 3k + 12\text{Exc}_W. \quad (19)$$

Plugging this back into [\(18\)](#) and using $k \geq 2$ gives:

$$|P| \leq 2a_{\max}(12k + 36\text{Exc}_W) = 24a_{\max}(k + 3\text{Exc}_W).$$

Remark 7.18. Observe that the above also proved the following, which will be of utility later in the proof:

$$\sum_{v \in \text{Lm}(W)} \deg_{G(W)}(v) \leq 12k + 36\text{Exc}_W.$$

Next, plugging the bound on $|P|$ back into [\(17\)](#) along with [Observation 7.14](#) gives:

$$\text{Tr}(M_W) \leq qC^{24a_{\max}(k+3\text{Exc}_W)+1} \cdot \prod_{T \in \text{UTrs}_{\geq 2}(W)} \|M_T\|^2 \leq qC^{25a_{\max}(k+3\text{Exc}_W)} \cdot \prod_{T \in \text{UTrs}_{\geq 2}(W)} \|M_T\|^2.$$

And via [\(16\)](#) and introducing the expected value over the randomness of τ :

$$\begin{aligned} \mathbf{E}_{\tau} \mathbf{E}_{H|\tau}[U] &\leq 2q \sum_{\text{Sh} \in \text{Shps}(k,\ell)} \left(\frac{2^r}{n^5} \right)^{\frac{|\text{Sh}|}{r}} \cdot (n^5 C^{75a_{\max}})^{\text{Exc}_{\text{Sh}}} C^{25a_{\max}k} \cdot \left(\frac{2\bar{\phi}_{\max}}{n} \right)^{\Delta(\text{Sh})} \\ &\quad \sum_{\substack{W : W \in \text{Lkgs}(\text{Bl}, n, k, \ell) \\ \text{Sh}(W) = \text{Sh}}} \prod_{T \in \text{UTrs}_{\geq 2}(W)} \|M_T\|^2 \cdot \mathbf{E}_{\tau} \prod_{\gamma \in R(W)} \mu_{\gamma} \end{aligned} \quad (20)$$

Thus, for a fixed shape Sh we now restrict our attention to bounding:

$$\sum_{\substack{W : W \in \text{Lkgs}(\text{Bl}, n, k, \ell) \\ \text{Sh}(W) = \text{Sh}}} \prod_{T \in \text{UTrs}_{\geq 2}(W)} \|M_T\|^2 \cdot \mathbf{E}_{\tau} \prod_{\gamma \in R(W)} \mu_{\gamma}. \quad (21)$$

We now define the notion of Cl -consistent.

Definition 7.19. Let Γ be a subgraph of the complete bipartite factor graph $\text{Bip}(\mathcal{K}_n)$. We say Γ is Cl -consistent if there exists a τ such that every $\gamma = (v_1, \dots, v_{a(i)}) \in R(\Gamma)$ satisfies $(\tau(v_1), \dots, \tau(v_{a(i)})) = \text{Cl}(\theta(\gamma))$. If Γ is Cl -consistent we use τ_{Γ} to refer to the unique τ such that Γ is (τ, Cl) -consistent.

Observation 7.20. For $W \in \text{Lkgs}(\text{Bl}, n, k, \ell)$, $\prod_{\gamma \in R(W)} \mu_{\gamma}$ is equal to 0 if W is not Cl -consistent.

Thus:

$$(21) = \sum_{\substack{W : W \in \text{Lkgs}(\text{Bl}, n, k, \ell) \\ \text{Sh}(W) = \text{Sh} \\ W \text{ Cl-consistent}}} \prod_{T \in \text{UTrs}_{\geq 2}(W)} \|M_T\|^2 \cdot \prod_{\gamma \in R(W)} \frac{\overline{\phi_{\theta(\gamma)}}}{n^{a(\theta(\gamma))-1}} \cdot \prod_{v \in L(\text{Clos}(W))} \mathcal{T}(\tau_W(v)).$$

Definition 7.21. We call two walks W_1 and W_2 *equivalent* denoted $W_1 \sim W_2$ if

- $\text{Sh}(W_1) = \text{Sh}(W_2) =: \text{Sh}$,
- for any $a \in R(\text{Sh})$, $\theta(W_1[a]) = \theta(W_2[a])$,
- for any edge $\{v, \gamma\}$ in Sh for $v \in L(\text{Sh})$ and $\gamma \in R(\text{Sh})$, $i(W_1[\gamma], W_1[v]) = i(W_2[\gamma], W_2[v])$.

We say W_1 and W_2 are *closure equivalent* denoted $W_1 \sim_{\text{Clos}} W_2$ if $W_1 \sim W_2$ and the graphs induced by $\text{Clos}(W_1)$ and $\text{Clos}(W_2)$ are isomorphic.

The relationship \sim partitions the space of all Cl-consistent W in $\text{Lkgs}(\text{Bl}, n, k, \ell)$ with shape Sh into a collection of equivalence classes \mathcal{C} . We use $[W]$ to denote the equivalence class it is contained in. \sim_{Clos} further partitions each equivalence class $[W] \in \mathcal{C}$ into a collection of sub-equivalence classes $\mathcal{C}_{[W]}$, and we denote the sub-equivalence class of W with $[[W]]$. We use $J(W)$ to denote $|E(\text{Clos}(W))| - |E(W)|$. With these definitions and notation in hand, we can write:

$$\begin{aligned}
(21) &= \sum_{[W] \in \mathcal{C}} \sum_{[[W]] \in \mathcal{C}[W]} \sum_{W' \in [[W]]} \prod_{T \in \text{UTrs}_{\geq 2}(W)} \|M_T\|^2 \cdot \prod_{\gamma \in R(W)} \frac{\overline{\phi_{\theta(\gamma)}}}{n^{a(\theta(\gamma))-1}} \cdot \prod_{v \in L(\text{Clos}(W))} \mathcal{T}(\tau_W(v)) \\
&\leq \sum_{[W] \in \mathcal{C}} \sum_{[[W]] \in \mathcal{C}[W]} \prod_{T \in \text{UTrs}_{\geq 2}(W)} \|M_T\|^2 \cdot \prod_{\gamma \in R(W)} \frac{\overline{\phi_{\theta(\gamma)}}}{n^{a(\theta(\gamma))-1}} \cdot \prod_{v \in L(\text{Clos}(W))} \mathcal{T}(\tau_W(v)) \cdot n^{|\text{Clos}(W)|} \\
&= \sum_{[W] \in \mathcal{C}} \sum_{t=0}^{J(W)} \sum_{\substack{[[W]] \in \mathcal{C}[W] \\ |L(\text{Clos}(W))| = |L(W)| + t}} \prod_{T \in \text{UTrs}_{\geq 2}(W)} \|M_T\|^2 \cdot \prod_{\gamma \in R(W)} \frac{\overline{\phi_{\theta(\gamma)}}}{n^{a(\theta(\gamma))-1}} \cdot \prod_{v \in L(\text{Clos}(W))} \mathcal{T}(\tau_W(v)) \cdot \frac{n^{|\text{Clos}(W)| + |J(W)|}}{n^{|J(W)| - t}}
\end{aligned}$$

To enumerate the innermost sum, first observe that any W can be changed to $\text{Clos}(W)$ where $|L(\text{Clos}(W))| - |L(W)| = t$ via the following procedure. First add $J(W)$ new vertices and to each $\gamma \in R(W)$ attach an edge from γ to $a(\theta(\gamma)) - \deg_W(\gamma)$ of the new vertices. There exists a sequence of $J(W) - t$ “merge” operations on the left vertices, and a labeling of the newly added left vertices in $[n]$ that would result in $\text{Clos}(W)$. The number of possible sequences of merge operations is at most $(a_{\max} k \ell)^{2(|J(W)| - t)}$. Thus, the above sum can be bounded by:

$$\begin{aligned}
&\leq \sum_{[W] \in \mathcal{C}} \sum_{t=0}^{J(W)} \prod_{T \in \text{UTrs}_{\geq 2}(W)} \|M_T\|^2 \cdot \prod_{\gamma \in R(W)} \frac{\overline{\phi_{\theta(\gamma)}}}{n^{a(\theta(\gamma))-1}} \cdot \prod_{v \in L(\text{Clos}(W))} \mathcal{T}(\tau_W(v)) \cdot n^{|\text{Clos}(W)| + |J(W)|} \cdot \left(\frac{(a_{\max} k \ell)^2}{n} \right)^{|J(W)| - t} \\
&= \sum_{[W] \in \mathcal{C}} \sum_{t=0}^{J(W)} \prod_{T \in \text{UTrs}_{\geq 2}(W)} \|M_T\|^2 \cdot \prod_{\gamma \in R(W)} \left(\frac{\overline{\phi_{\theta(\gamma)}}}{n^{a(\theta(\gamma))-1}} \cdot \prod_{v \in L(\text{Clos}(W)) \setminus L(W)} \mathcal{T}(\tau_W(v)) \right) \\
&\quad \prod_{v \in L(W)} \mathcal{T}(\tau_W(v)) \cdot \prod_{v \in L(\text{Clos}(W)) \setminus L(W)} \frac{1}{\mathcal{T}(\tau_W(v))^{\deg_W(v)-1}} \cdot n^{|\text{Clos}(W)| + |J(W)|} \cdot \left(\frac{(a_{\max} k \ell)^2}{n} \right)^{|J(W)| - t} \\
&\leq \sum_{[W] \in \mathcal{C}} \sum_{t=0}^{J(W)} \prod_{T \in \text{UTrs}_{\geq 2}(W)} \|M_T\|^2 \cdot \prod_{\gamma \in R(W)} \left(\frac{\overline{\phi_{\theta(\gamma)}}}{n^{a(\theta(\gamma))-1}} \cdot \prod_{v \in L(\text{Clos}(W)) \setminus L(W)} \mathcal{T}(\tau_W(v)) \right) \\
&\quad \prod_{v \in L(W)} \mathcal{T}(\tau_W(v)) \cdot n^{|\text{Clos}(W)| + |J(W)|} \cdot \left(\frac{(a_{\max} k \ell)^2}{\mathcal{T}_{\min} n} \right)^{|J(W)| - t}
\end{aligned}$$

We can use the bound $\sum_{t=0}^{J(W)} \left(\frac{(a_{\max} k \ell)^2}{\mathcal{T}_{\min} n} \right)^{|J(W)| - t} \leq 2$ to deduce:

$$\leq 2 \sum_{[W] \in \mathcal{C}} \prod_{T \in \text{UTrs}_{\geq 2}(W)} \|M_T\|^2 \cdot \prod_{\gamma \in R(W)} \left(\frac{\overline{\phi_{\theta(\gamma)}}}{n^{a(\theta(\gamma))-1}} \cdot \prod_{v \in L(\text{Clos}(W)) \setminus L(W)} \mathcal{T}(\tau_W(v)) \right) \cdot \prod_{v \in L(W)} \mathcal{T}(\tau_W(v)) \cdot n^{|\text{Clos}(W)| + |J(W)|}$$

$$= 2 \sum_{[W] \in \mathcal{C}} \prod_{T \in \text{UTrs}_{\geq 2}(W)} \|M_T\|^2 \cdot \prod_{\gamma \in R(W)} \left(\overline{\phi_{\theta(\gamma)}} \cdot \prod_{v \in L(\text{Clos}(W)) \setminus L(W)} \mathcal{T}(\tau_W(v)) \right) \cdot \prod_{v \in L(W)} \mathcal{T}(\tau_W(v)) \cdot n^{1 - \text{Exc}_W}. \quad (22)$$

We decompose $\prod_{\gamma \in R(W)} (\overline{\phi_{\theta(\gamma)}} \cdot \prod_{v \in L(\text{Clos}(W)) \setminus L(W)} \mathcal{T}(\tau_W(v)))$ into three parts: the contribution of singleton right vertices $\prod_{\gamma \in S(W)} (\overline{\phi_{\theta(\gamma)}} \cdot \prod_{v \in L(\text{Clos}(W)) \setminus L(W)} \mathcal{T}(\tau_W(v)))$ which can be bounded by $\overline{\phi_{\max}}^{|S(\text{Sh})|}$, the contribution of duplicative degree- ≥ 3 right vertices $\prod_{\substack{\gamma \in D(W) \\ \deg_W(\gamma) \geq 3}} (\overline{\phi_{\theta(\gamma)}} \cdot \prod_{v \in L(\text{Clos}(W)) \setminus L(W)} \mathcal{T}(\tau_W(v)))$ which via (19) can be bounded by $\overline{\phi_{\max}}^{-3k+12\text{Exc}_{\text{Sh}}}$, and finally the contribution of duplicative degree-2 right vertices $\prod_{\substack{\gamma \in D(W) \\ \deg_W(\gamma)=2}} (\overline{\phi_{\theta(\gamma)}} \cdot \prod_{v \in L(\text{Clos}(W)) \setminus L(W)} \mathcal{T}(\tau_W(v)))$. For each γ considered in the final case we can identify a unique $T \in \text{UTrs}_{\geq 2}(W)$ such that γ is in T , and likewise for every $T \in \text{UTrs}_{\geq 2}(W)$, every right vertex γ in T is duplicative and has degree exactly 2 and hence appears in the product. Thus, the third product can be written as

$$\prod_{T \in \text{UTrs}_{\geq 2}(W)} \prod_{\gamma \in R(T)} \left(\overline{\phi_{\theta(\gamma)}} \cdot \prod_{v \in L(\text{Clos}(T)) \setminus L(T)} \mathcal{T}(\tau_W(v)) \right).$$

Next, observe that using the facts that each $\pi_i \in [0, 1]$ and every interior left vertex of a trail occurs in no other trail, $\prod_{v \in L(W)} \mathcal{T}(\tau_W(v))$ can be upper bounded by $\prod_{T \in \text{UTrs}_{\geq 2}(W)} \prod_{v \in L(T)} \mathcal{T}(\tau_W(v)) \cdot \prod_{v \in Lm(W)} \frac{1}{\mathcal{T}(\tau_W(v))^{\deg_{\mathcal{C}(W)}(v)}}$, which by Remark 7.18 is at most $\prod_{T \in \text{UTrs}_{\geq 2}(W)} \prod_{v \in L(T)} \mathcal{T}(\tau_W(v)) \cdot \left(\frac{1}{\mathcal{T}_{\min}}\right)^{12k+36\text{Exc}_{\text{Sh}}}$.

$$\begin{aligned} (22) &\leq 2 \overline{\phi_{\max}}^{|S(\text{Sh})|+3k+12\text{Exc}_{\text{Sh}}} \left(\frac{1}{\mathcal{T}_{\min}}\right)^{12k+36\text{Exc}_{\text{Sh}}} n^{-\text{Exc}_{\text{Sh}}+1}. \\ &\sum_{[W]} \prod_{T \in \text{UTrs}_{\geq 2}(W)} \left(\|M_T\|^2 \cdot \prod_{\gamma \in R(T)} \overline{\phi_{\theta(\gamma)}} \cdot \prod_{v \in L(\text{Clos}(T)) \setminus L(T)} \mathcal{T}(\tau_W(v)) \cdot \prod_{v \in L(T)} \mathcal{T}(\tau_W(v)) \right) \\ &\leq 2 \overline{\phi_{\max}}^{|S(\text{Sh})|+3k+12\text{Exc}_{\text{Sh}}} \left(\frac{1}{\mathcal{T}_{\min}}\right)^{12k+36\text{Exc}_{\text{Sh}}} n^{-\text{Exc}_{\text{Sh}}+1}. \\ &\prod_{T \in \text{UTrs}_{\geq 2}(\text{Sh})} \sum_{[U]: \text{Sh}(U)=T} \left(\|M_U\|^2 \cdot \prod_{\gamma \in R(U)} \overline{\phi_{\theta(\gamma)}} \cdot \prod_{v \in L(\text{Clos}(U)) \setminus L(U)} \mathcal{T}(\tau_W(v)) \cdot \prod_{v \in L(U)} \mathcal{T}(\tau_W(v)) \right) \end{aligned}$$

For any constant ε , there exists a constant C_ε such that the above is at most:

$$\leq 2 \overline{\phi_{\max}}^{|S(\text{Sh})|+3k+12\text{Exc}_{\text{Sh}}} \left(\frac{1}{\mathcal{T}_{\min}}\right)^{12k+36\text{Exc}_{\text{Sh}}} n^{-\text{Exc}_{\text{Sh}}+1} \prod_{T \in \text{UTrs}_{\geq 2}(\text{Sh})} C_\varepsilon \cdot ((1 + \varepsilon)\rho(\text{Bl}, M^\times))^{2|T|}$$

Since $|\text{UTrs}_{\geq 2}(\text{Sh})|$ is at most the sum of degrees of landmark vertices on which we have an upper bound by Remark 7.18, and since $\sum_{T \in \text{UTrs}_{\geq 2}(\text{Sh})} |T| \leq \frac{k\ell}{2}$:

$$\leq 2 \overline{\phi_{\max}}^{|S(\text{Sh})|+3k+12\text{Exc}_{\text{Sh}}} \left(\frac{1}{\mathcal{T}_{\min}}\right)^{12k+36\text{Exc}_{\text{Sh}}} n^{-\text{Exc}_{\text{Sh}}+1} C_\varepsilon^{12k+36\text{Exc}_{\text{Sh}}} \cdot ((1 + \varepsilon)\rho(\text{Bl}, M^\times))^{k\ell}.$$

Since (22) is an upper bound on (21), rearranging the terms in the above gives:

$$(21) \leq 2n \cdot \overline{\phi_{\max}}^{|S(\text{Sh})|} \cdot \left(\frac{\overline{\phi_{\max}}^3 C_\varepsilon^{12}}{\mathcal{T}_{\min}^{12}}\right)^k \cdot \left(\frac{\overline{\phi_{\max}}^{12} C_\varepsilon^{36}}{\mathcal{T}_{\min}^{36} n}\right)^{\text{Exc}_{\text{Sh}}} \cdot ((1 + \varepsilon)\rho(\text{Bl}, M^\times))^{k\ell}.$$

Plugging this upper bound on (21) into (20) gives us:

$$\mathbf{E}_\tau \mathbf{E}_{H|\tau}[U] \leq 4nq((1+\varepsilon)\rho(\text{Bl}, M^*))^{k\ell} \cdot \sum_{\text{Sh} \in \text{Shps}(k, \ell)} \left(\frac{2^r \bar{\phi}_{\max}^r}{n^5} \right)^{\frac{|S(\text{Sh})|}{r}} \left(\frac{C^{75a_{\max}} \bar{\phi}_{\max}^{12} C_\varepsilon^{36}}{\mathcal{T}_{\min}^{36} n^5} \right)^{\text{Exc}_{\text{Sh}}} \left(\frac{C^{25a_{\max}} \bar{\phi}_{\max}^3 C_\varepsilon^{12}}{\mathcal{T}_{\min}^{12}} \right)^k \left(\frac{2\bar{\phi}_{\max}}{n} \right)^{\Delta(\text{Sh})}.$$

To notationally lighten the above, we choose β as a constant larger than $2\bar{\phi}_{\max}$, $\frac{C^{75a_{\max}} \bar{\phi}_{\max}^{12} C_\varepsilon^{36}}{\mathcal{T}_{\min}^{36}}$ and $\frac{C^{25a_{\max}} \bar{\phi}_{\max}^3 C_\varepsilon^{12}}{\mathcal{T}_{\min}^{12}}$. Then:

$$\mathbf{E}_\tau \mathbf{E}_{H|\tau}[U] \leq 4nq((1+\varepsilon)\rho(\text{Bl}, M^*))^{k\ell} \beta^k \cdot \sum_{\text{Sh} \in \text{Shps}(k, \ell)} \left(\frac{\beta^r}{n^5} \right)^{\frac{|S(\text{Sh})|}{r}} \left(\frac{\beta}{n^5} \right)^{\text{Exc}_{\text{Sh}}} \left(\frac{\beta}{n} \right)^{\Delta(\text{Sh})} \quad (23)$$

To obtain a bound on (23) we first bound:

$$\sum_{\text{Sh} \in \text{Shps}(k, \ell)} \left(\frac{\beta^r}{n^5} \right)^{\frac{|S(\text{Sh})|}{r}} \left(\frac{\beta}{n^5} \right)^{\text{Exc}_{\text{Sh}}} \left(\frac{\beta}{n} \right)^{\Delta(\text{Sh})} \quad (24)$$

We proceed by partition all $\text{Sh} \in \text{Shps}(k, \ell)$ into sets where each set of Sh share the same $|S(\text{Sh})|$, Exc_{Sh} , and $\Delta(\text{Sh})$. We bound the sum for each set with the following claim proved in [Appendix D](#).

Claim 7.22. Let $\mathcal{U}_{s,x,\Delta}$ denote the set of all $\text{Sh} \in \text{Shps}(k, \ell)$ with $|S(\text{Sh})| = s$, $\text{Exc}_{\text{Sh}} = x$, and $\Delta(\text{Sh}) = \Delta$. Then

$$\sum_{\text{Sh} \in \mathcal{U}_{s,x,\Delta}} \left(\frac{\beta^r}{n^5} \right)^{\frac{|S(\text{Sh})|}{r}} \left(\frac{\beta}{n^5} \right)^{\text{Exc}_{\text{Sh}}} \left(\frac{\beta}{n} \right)^{\Delta(\text{Sh})} \leq \left(\frac{\beta^r \cdot (4k\ell)^r}{n^5} \right)^{\frac{s}{r}} \left(\frac{\beta \cdot 2(k\ell)^3}{n^5} \right)^x \left(\frac{4\beta \cdot (2k\ell)^2}{n} \right)^\Delta (2k\ell)^{O(k \log k\ell)}.$$

Using this claim we can derive

$$\begin{aligned} (24) &= \sum_{s \in [k\ell], x \in [2k\ell], \Delta \in [k\ell]} \sum_{\text{Sh} \in \mathcal{U}_{s,x,\Delta}} \left(\frac{\beta^r}{n^5} \right)^{\frac{s}{r}} \left(\frac{\beta}{n^5} \right)^x \left(\frac{\beta}{n} \right)^\Delta \\ &\leq \sum_{s \in [k\ell], x \in [2k\ell], \Delta \in [k\ell]} \left(\frac{\beta^r \cdot (4k\ell)^r}{n^5} \right)^{\frac{s}{r}} \left(\frac{\beta \cdot 2(k\ell)^3}{n^5} \right)^x \left(\frac{4\beta \cdot (2k\ell)^2}{n} \right)^\Delta (2k\ell)^{O(k \log k\ell)} \\ &\leq (2k\ell)^{O(k \log k\ell)} k\ell \cdot 2k\ell \cdot k\ell \max_{s \in [k\ell], x \in [2k\ell], \Delta \in [k\ell]} \left(\frac{\beta^r \cdot (4k\ell)^r}{n^5} \right)^{\frac{s}{r}} \left(\frac{\beta \cdot 2(k\ell)^3}{n^5} \right)^x \left(\frac{4\beta \cdot (2k\ell)^2}{n} \right)^\Delta \end{aligned}$$

We set the bicycle-free radius to $r := o(\log n^5 / \log 4k\ell\beta)$ so that all three terms $\frac{\beta^r \cdot (4k\ell)^r}{n^5}$, $\frac{\beta \cdot 2(k\ell)^3}{n^5}$, $\frac{4\beta \cdot (2k\ell)^2}{n}$ are less than 1. Then we observe that

$$(24) \leq 2k^3 \ell^3 (2k\ell)^{O(k \log k\ell)}.$$

Plugging this bound into (23) gives us:

$$\mathbf{E}_\tau \mathbf{E}_{H|\tau}[U] \leq ((1+\varepsilon)\rho(\text{Bl}, M^*))^{k\ell} \cdot 2nq\beta^k 2k^3 \ell^3 (2k\ell)^{O(k \log k\ell)}$$

Set $k\ell = O(\log n \cdot \log \log n)$, and $\ell = \omega(\log k\ell)$. Then we have the bound

$$(n\mathbf{E}_\tau \mathbf{E}_{H|\tau}[U])^{1/k} \leq \left((1+\varepsilon)\rho(\text{Bl}, M^*) \cdot (4q\beta^k \cdot n^2 \cdot (k\ell) \cdot (2k\ell)^{O(\log k\ell)/\ell})^{1/k\ell} \right)^\ell$$

$$\begin{aligned}
&\leq \left((1 + \varepsilon)\rho(\text{Bl}, M^\times) \cdot (8q\beta^k \cdot n^3)^{1/k\ell} \right)^\ell \\
&\leq \left((1 + \varepsilon)\rho(\text{Bl}, M^\times) \cdot \left(1 + \frac{2 \log(8q\beta^k)}{k\ell} + \frac{2 \log(n^3)}{k\ell} \right) \right)^\ell \\
&\leq \left(\left(1 + \varepsilon + O\left(\frac{1}{\log \log n} \right) \right) \rho(\text{Bl}, M^\times) \right)^\ell.
\end{aligned}$$

We now complete the proof of [Theorem 7.4](#).

8 Weak Recovery

We begin the section by briefly describing an algorithm for weak recovery.

Weak Recovery Algorithm

1. Fix $\delta > 0$ such that $\lambda_L \geq (1 + \delta)^4$.
2. $C = O_{M,\delta}(1)$ is a sufficiently large constant depending on model M and δ .
3. For $(\log \log n)^3 \leq t \leq (\log \log n)^5$, $v_t \in \mathbb{R}^{nq}$ be the eigenvector with largest eigenvalue of $\underline{A}_G^{(t)}$ and let Λ_t^t denote the largest eigenvalue. Compute v_t , Λ_t and $\underline{A}_G^{(t)}$ for all t in this range.
4. Find $m \in [(\log \log n)^3, (\log \log n)^5]$ such that for all $s \in [(\log \log n)^3, m]$, we have

$$\|\underline{A}_G^{(s)} v_m\| \leq \Lambda_m^s (1 + \delta)^{m-s}$$

(see [Claim 8.6](#) for proof of existence of m)

5. For each $0 \leq \ell \leq C$, set

$$w_\ell \stackrel{\text{def}}{=} \underline{A}_G^{(m-\ell)} v_m$$

and let $\bar{w}_\ell \stackrel{\text{def}}{=} \frac{1}{\|w_\ell\|} \cdot w_\ell$.

6. Output the set of vectors $\{\bar{u}_{\ell,\beta}\}$ in \mathbb{R}^n for $0 \leq \ell \leq C$, $\beta \in [q]$ defined as,

$$\bar{u}_{\ell,\beta}[i] = \bar{w}_\ell[i, \beta]$$

Theorem 8.1. *There exists a constant $C = O_{M,\delta}(1)$ depending on model M and δ such that the following holds with probability $1 - o_n(1)$: For some $\ell \in \{1, 2, \dots, C\}$ and $\beta \in [q]$, the unit vector $\bar{u}_{\ell,\beta}$ is correlated with the coloring in the following sense: $\exists \tau \in T, \alpha \in [q]$ such that if we construct $\underline{\chi}^{\tau,\alpha} \in \mathbb{R}^n$ as*

$$\underline{\chi}^{\tau,\alpha}[i] = \mathbf{1}[\tau(i) = \tau](\mathbf{1}[\mathbf{c}(i) = \alpha] - \mathbb{P}_\tau(\alpha))$$

then

$$|\langle \bar{u}_{\ell,\beta}, \underline{\chi}^{\tau,\alpha} \rangle| \geq \Omega_M(1) \cdot \sqrt{n}$$

In the rest of the section, we will outline the proof of correctness of the above described weak-recovery algorithm. To this end, we begin by recalling the matrix $\underline{A}_G \in \mathbb{R}^{nq \times nq}$. For all $i \neq j \in [n]$,

$$\underline{A}_G[i, j] \stackrel{\text{def}}{=} \sum_{e \in \mathcal{K}_n, e \ni ij} \overline{\mathbf{M}}_{e, ij} \cdot (\mathbf{1}[e \in G] - \Pr_{\mathbf{M}}[e \in G | \boldsymbol{\tau}])$$

and $\underline{A}_G[i, i] = 0$. Similarly, for all $i \neq j \in [n]$ we set

$$\underline{B}_G[i, j] \stackrel{\text{def}}{=} \sum_{e \in \mathcal{K}_n, e \ni ij} \overline{\mathbf{M}}_{e, ij} \cdot (\mathbf{1}[e \in G] - \Pr_{\mathbf{M}}[e \in G | \boldsymbol{\tau}, \mathbf{c}])$$

and $\underline{B}_G[i, i] = 0$. Finally, let

$$\begin{aligned} \underline{R}_G[i, j] &\stackrel{\text{def}}{=} \underline{A}_G[i, j] - \underline{B}_G[i, j] \\ &= \sum_{e \in \mathcal{K}_n, e \ni ij} \overline{\mathbf{M}}_{e, ij} (\Pr_{\mathbf{M}}[e \in G | \boldsymbol{\tau}, \mathbf{c}] - \Pr_{\mathbf{M}}[e \in G | \boldsymbol{\tau}]) \end{aligned}$$

Let $\vec{x}_G \in \mathbb{R}^{[T]}$ encode the number of variables of each type in G and $\vec{y}_G \in \mathbb{R}^{q[T]}$ encode the number of variables of each type and color in G . Then block $\underline{R}_G[i, j]$ only depends on $\boldsymbol{\tau}(i)$, $\boldsymbol{\tau}(j)$, $\mathbf{c}(i)$, $\mathbf{c}(j)$ and \vec{y}_G . More specifically

$$\underline{R}_G[i, j] = \sum_{e \in \mathcal{K}_n, e \ni ij} \overline{\mathbf{M}}_{e, ij} (\Pr_{\mathbf{M}}[e \in G | \boldsymbol{\tau}(i), \boldsymbol{\tau}(j), \mathbf{c}(i), \mathbf{c}(j), \vec{y}_G] - \Pr_{\mathbf{M}}[e \in G | \boldsymbol{\tau}(i), \boldsymbol{\tau}(j), \vec{x}_G])$$

Remark 8.2. We remark that with probability $1 - o_n(1)$ each entry of \vec{y}_G satisfies $\vec{y}_G[\tau, \alpha] \in (1 \pm \epsilon) \cdot \mathbb{E}_{\mathbf{M}}[\vec{y}_G[\tau, \alpha]]$ for some small constant ϵ . From now on we only consider G that satisfies this condition.

Next we introduce notation for the non-backtracking product of two matrices.

Definition 8.3. For two matrices $A, B \in \mathbb{R}^{nq \times nq}$, define $A \circ B$ to be the non-backtracking product of A and B by setting for all $i, j \in [n]$

$$(A \circ B)[i, j] = \begin{cases} \sum_k A[i, k] B[k, j] & \text{if } i \neq j \\ 0 & \text{otherwise} \end{cases}$$

Inductively define $A^{(s)} \stackrel{\text{def}}{=} A^{(s-1)} \circ A$

Suppose $v_t \in \mathbb{R}^{nq}$ be the eigenvector with largest eigenvalue of $\underline{A}_G^{(t)}$, and let its eigenvalue be Λ_t^t . By [Theorem 6.1](#), we know that for each s , with probability $1 - o_n(1)$,

$$\lambda_{\max}(\underline{A}_G^{(s)}) \geq \Omega_{\mathbf{M}}(1) \cdot \lambda_L^s$$

Therefore,

$$v_s \underline{A}_G^{(s)} v_s \geq \Omega_{\mathbf{M}}(1) \cdot \lambda_L^s$$

On the other hand, by the spectral norm bound for all $s \geq (\log \log n)^3$ in the null model,

$$v_s \underline{B}_G^{(s)} v_s \leq (1 + o(1))^s \left(\sqrt{\lambda_L} \right)^s$$

For simplifying notation, we will ignore the $(1 + o(1))^s$ term in the above bound, here in the rest of the section. Rewriting the difference we get,

$$v_i^T \left(\underline{A}_G^{(t)} - \underline{B}_G^{(t)} \right) v_i = \sum_{s=0}^{t-1} v_i^T \underline{B}_G^{(s)} \circ \underline{R}_G \circ \underline{A}_G^{(t-s-1)} v_i$$

Now we can replace the non-backtracking product in the above expression with the usual matrix product using [Lemma 8.4](#).

Lemma 8.4. For all $A, B, R \in \mathbb{R}^{nq \times nq}$,

$$\|A^{(s)} \circ R \circ B^{(t)} - A^{(s)} R B^{(t)}\| \leq q \|R\|_\infty \cdot \left(\|A^{(s)}\| \|B^{(t-s)}\| + \|A\|_{1 \rightarrow 1} \|A^{(s-1)}\| \|B^{(t)}\| + \|B^T\|_{1 \rightarrow 1} \|A^{(s)}\| \|B^{(t-1)}\| \right)$$

where $\|R\|_\infty = \max_{\ell, \ell' \in [nq]} |R_{\ell, \ell'}|$

We will postpone the proof of this Lemma to later in the section, and proceed with the argument.

Notice that under the condition in [Remark 8.2](#),

$$\|\underline{R}_G\|_\infty \leq O_M(1) \cdot \frac{1}{n} \quad (25)$$

The maximum degree of a variable in the factor graph is $O(\log n)$ with probability $1 - o_n(1)$. Therefore a naive bound on the spectral norm of $\underline{A}_G^{(s)}$ would be

$$\|\underline{A}_G^{(s)}\| \leq O(\log n)^s \leq o(n^{1/4}) \quad (26)$$

for $s \leq o(\log n / \log \log n)$. Similarly, we can bound $\|\underline{B}_G^{(t)}\| \leq o(n^{1/4})$. Using these bounds in [Lemma 8.4](#), we can replace non-backtracking product by the usual product to conclude,

$$v_i^T \left(\underline{A}_G^{(t)} - \underline{B}_G^{(t)} \right) v_i = \sum_{s=0}^{t-1} v_i^T \underline{B}_G^{(s)} \underline{R}_G \underline{A}_G^{(t-s-1)} v_i + o_n(1)$$

We will now rewrite the matrix \underline{R}_G explicitly in terms of the coloring \mathbf{c} . To this end, we make a few definitions. For types $\tau, \tau' \in T$ and colors $\alpha, \alpha' \in [q]$ define $\Gamma_{\alpha, \alpha'}^{\tau, \tau'} \in \mathbb{R}^{[q] \times [q]}$ as,

$$\Gamma_{\alpha, \alpha'}^{\tau, \tau'} \stackrel{\text{def}}{=} \sum_{e \in \mathcal{K}_n, e \ni ij} \bar{\mathbf{M}}_{e, ij} \cdot \Pr_M[e \in G | \tau(i) = \tau, \tau(j) = \tau', \mathbf{c}(i) = \alpha, \mathbf{c}(j) = \alpha', \vec{y}_G]$$

In terms of the matrices $\{\Gamma_{\alpha, \alpha'}^{\tau, \tau'}\}$ we can write for $i \neq j$,

$$\underline{R}_G[i, j] = \sum_{\tau, \tau' \in T} \sum_{\alpha, \alpha' \in [q]} \Gamma_{\alpha, \alpha'}^{\tau, \tau'} \cdot \mathbf{1}[\tau(i) = \tau] \mathbf{1}[\tau(j) = \tau'] \left(\mathbf{1}[\mathbf{c}(i) = \alpha] \mathbf{1}[\mathbf{c}(j) = \alpha'] - \Pr[\mathbf{c}(i) = \alpha | \tau(i), \vec{y}_G] \Pr[\mathbf{c}(j) = \alpha' | \tau(j), \vec{y}_G] \right)$$

For every type τ and a color α , let $\underline{\chi}^{\tau, \alpha}, \mathbf{c}^{\tau, \alpha}, \mu^{\tau, \alpha} \in \mathbb{R}^n$ be defined as follows:

$$\begin{aligned} \underline{\chi}^{\tau, \alpha}[i] &\stackrel{\text{def}}{=} \mathbf{1}[\tau(i) = \tau] \cdot \left(\mathbf{1}[\mathbf{c}(i) = \alpha] - \Pr[\mathbf{c}(i) = \alpha | \tau(i), \vec{y}_G] \right) \\ \chi^{\tau, \alpha}[i] &\stackrel{\text{def}}{=} \mathbf{1}[\tau(i) = \tau] \cdot \left(\mathbf{1}[\mathbf{c}(i) = \alpha] \right) \\ \mu^{\tau, \alpha}[i] &\stackrel{\text{def}}{=} \mathbf{1}[\tau(i) = \tau] \cdot \Pr[\mathbf{c}(i) = \alpha | \tau(i), \vec{y}_G] \end{aligned}$$

Hence for $i \neq j$ we have,

$$\underline{R}_G[i, j] = \sum_{\tau, \tau' \in T} \sum_{\alpha, \alpha' \in [q]} \Gamma_{\alpha, \alpha'}^{\tau, \tau'} \cdot \left(\underline{\chi}^{\tau, \alpha}[i] \underline{\chi}^{\tau', \alpha'}[j] + \underline{\chi}^{\tau, \alpha}[i] \mu^{\tau', \alpha'}[j] \right)$$

Define $\underline{R}[i, i]$ so that we have the equality,

$$\underline{R}_G = \sum_{\tau, \tau', \alpha, \alpha'} \Gamma_{\alpha, \alpha'}^{\tau, \tau'} \otimes \left(\underline{\chi}^{\tau, \alpha} (\underline{\chi}^{\tau', \alpha'})^T + \underline{\chi}^{\tau, \alpha} (\underline{\mu}^{\tau', \alpha'})^T \right) \quad (27)$$

Using (27) for the matrix \underline{R} , we can write,

$$v_t^T \left(\underline{A}_G^{(t)} - \underline{B}_G^{(t)} \right) v_t = \sum_{\tau, \tau', \alpha, \alpha'} v_t^T \sum_{s=1}^t \underline{B}_G^{(s)} \left(\Gamma_{\alpha, \alpha'}^{\tau, \tau'} \otimes \left(\underline{\chi}^{\tau, \alpha} (\underline{\chi}^{\tau', \alpha'})^T + \underline{\chi}^{\tau, \alpha} (\underline{\mu}^{\tau', \alpha'})^T \right) \right) \underline{A}_G^{(t-s-1)} v_t + o_n(1)$$

The second term corresponding to $\underline{\chi}^{\tau, \alpha} (\underline{\mu}^{\tau', \alpha'})^T$ is negligible. Specifically, we will prove the following Lemma.

Lemma 8.5. *With probability $1 - o_n(1)$ the following holds, for all $\tau, \alpha, \tau', \alpha'$ For all $1 \leq s, t \leq \sqrt{\log n}$,*

$$\| \underline{B}_G^{(s)} \left(\Gamma_{\alpha, \alpha'}^{\tau, \tau'} \otimes \underline{\chi}^{\tau, \alpha} (\underline{\mu}^{\tau', \alpha'})^T \right) \underline{A}_G^{(t)} \| \leq \left(\sqrt{\lambda_L} \right)^{s+t} \cdot O_M(1)$$

We postpone the proof to later in the section and proceed with the main argument.

Using Lemma 8.5, we can drop all terms arising from $\underline{\chi}^{\tau, \alpha} (\underline{\mu}^{\tau', \alpha'})^T$ by losing less than $O_M(1) \sum_{s=0}^{t-1} \left(\sqrt{\lambda_L} \right)^s$. Since $\lambda_L > (1 + \delta)^4$, for all t larger than a fixed constant $\Theta_M(1)$, this sum $O_M(1) \sum_{s=0}^{t-1} \left(\sqrt{\lambda_L} \right)^s \leq \frac{1}{10} \lambda_L^t$. Therefore, we arrive at our inequality,

$$\sum_{\tau, \tau', \alpha, \alpha'} v_t^T \sum_{s=0}^{t-1} \underline{B}_G^{(s)} \left(\Gamma_{\alpha, \alpha'}^{\tau, \tau'} \otimes \underline{\chi}^{\tau, \alpha} (\underline{\chi}^{\tau', \alpha'})^T \right) \underline{A}_G^{(t-s-1)} v_t \geq 0.9 \cdot \Lambda_t^t$$

Let us write the matrix $\Gamma_{\alpha, \alpha'}^{\tau, \tau'} = \sum_{\beta, \beta' \in [q]} \Gamma_{\alpha, \alpha'}^{\tau, \tau'} [\beta, \beta'] \cdot \mathbf{e}_\beta (\mathbf{e}_{\beta'})^T$ where $\mathbf{e}_\beta, \mathbf{e}_{\beta'}$ are standard basis vectors in \mathbb{R}^q . Note that $\Gamma_{\alpha, \alpha'}^{\tau, \tau'}$ has entries that are $O_M(1)/n$. There must exist some choice of $\tau, \tau', \alpha, \alpha', \beta, \beta'$ such that,

$$v_t^T \sum_{s=0}^{t-1} \underline{B}_G^{(s)} \left(\mathbf{e}_\beta \otimes \underline{\chi}^{\tau, \alpha} (\mathbf{e}_{\beta'} \otimes \underline{\chi}^{\tau', \alpha'})^T \right) \underline{A}_G^{(t-s-1)} v_t \geq \Omega_M(1) \cdot \Lambda_t^t \cdot n$$

where $\Omega_M(1)$ hides a constant depending on the model M . Rewriting the above inequality,

$$\begin{aligned} \Omega_M(1) \cdot \Lambda_t^t \cdot n &\leq \sum_{s=0}^{t-1} \left\langle v_t^T, \underline{B}_G^{(s)} \mathbf{e}_\beta \otimes \underline{\chi}^{\tau, \alpha} \right\rangle \left\langle \mathbf{e}_{\beta'} \otimes \underline{\chi}^{\tau', \alpha'}, \underline{A}_G^{(t-s-1)} v_t \right\rangle \\ &\leq \sum_{s=0}^{t-1} \| \underline{B}_G^{(s)} \mathbf{e}_\beta \otimes \underline{\chi}^{\tau, \alpha} \| \left| \left\langle \mathbf{e}_{\beta'} \otimes \underline{\chi}^{\tau', \alpha'}, \underline{A}_G^{(t-s-1)} v_t \right\rangle \right| \end{aligned}$$

Using Lemma 8.8 on the fixed vector $\mathbf{e}_\beta \otimes \underline{\chi}^{\tau, \alpha}$ and the planted distribution we get that with probability $1 - o_n(1)$,

$$\Omega_M(1) \cdot \Lambda_t^t \cdot n \leq \sum_{s=0}^{t-1} \left(\sqrt{\lambda_L} \right)^s \cdot \sqrt{n} \left| \left\langle \mathbf{e}_{\beta'} \otimes \underline{\chi}^{\tau', \alpha'}, \underline{A}_G^{(t-s-1)} v_t \right\rangle \right|$$

For notational convenience, let us reparametrize $s \rightarrow t - s$ and conclude,

$$\Omega_M(1) \cdot \Lambda_t^t \cdot \sqrt{n} \leq \sum_{s=1}^t \left(\sqrt{\lambda_L} \right)^{t-s} \left| \left\langle \mathbf{e}_{\beta'} \otimes \underline{\chi}^{\tau', \alpha'}, \underline{A}_G^{(s-1)} v_t \right\rangle \right|$$

With high probability, the maximum degree of a variable is $O(\log n) \ll o(\log^2 n)$, and therefore $\|\underline{A}^{(s)}\| \leq o(\log n)^{2s}$.

Since $\Lambda_t \geq \lambda_L \geq (1 + \delta) \left(\sqrt{\lambda_L} \right)$ we can bound the terms for $s = 1, \dots, (\log \log n)^3$ as follows,

$$\sum_{s=1}^{(\log \log n)^3} \left(\sqrt{\lambda_L} \right)^{t-s} \left| \left\langle \mathbf{e}_{\beta'} \otimes \underline{\chi}^{\tau', \alpha'}, \underline{A}_G^{(s-1)} v_t \right\rangle \right| \leq \sum_{s=1}^{(\log \log n)^3} \frac{\Lambda_t^{t-s}}{(1 + \delta)^{t-s}} \cdot o((\log n)^{2s}) \cdot \sqrt{n} \leq o(1) \cdot \Lambda_t^t \cdot \sqrt{n}$$

where the last inequality holds for $t > (\log \log n)^4$. Deleting terms for small s , we have the correlation inequality,

$$\Omega_M(1) \cdot \Lambda_t^t \cdot \sqrt{n} \leq \sum_{s=(\log \log n)^3}^t \left(\sqrt{\lambda_L} \right)^{t-s} \left| \left\langle \mathbf{e}_{\beta'} \otimes \underline{\chi}^{\tau', \alpha'}, \underline{A}_G^{(s-1)} v_t \right\rangle \right| \quad (28)$$

for $t > (\log \log n)^4$.

For each s , recall that v_s is the top eigenvector of $\underline{A}_G^{(s)}$, and Λ_s^s denotes the largest eigenvalue.

Claim 8.6. There exists a $m \in [(\log \log n)^{4.5}, (\log \log n)^5]$ such that, for all $s \in [(\log \log n)^3, m]$,

$$\|\underline{A}^{(s)} v_m\| \leq \Lambda_m^s (1 + \delta)^{t-s}$$

Before we see the proof of above claim, let us see how it leads to an algorithm. Applying (28) for this choice of m , we conclude that

$$\sqrt{n} \cdot \Omega_M(1) \cdot \Lambda_m^m \leq \sum_{s=(\log \log n)^3}^m \left(\sqrt{\lambda_L} \right)^{m-s} \left| \left\langle \mathbf{e}_{\beta'} \otimes \underline{\chi}^{\tau', \alpha'}, \underline{A}_G^{(s-1)} v_m \right\rangle \right| \quad (29)$$

In (29), we can bound the sum of all terms with $s < t^* - C$ as follows.

$$\sum_{s=(\log \log n)^3}^{m-C} \left(\sqrt{\lambda_L} \right)^{m-s} \left| \left\langle \mathbf{e}_{\beta'} \otimes \underline{\chi}^{\tau', \alpha'}, \underline{A}_G^{(s-1)} v_t \right\rangle \right| \leq \sum_{s=(\log \log n)^3}^{m-C} \left(\sqrt{\lambda_L} \right)^{m-s} \|\mathbf{e}_{\beta'} \otimes \underline{\chi}^{\tau', \alpha'}\| \|\underline{A}_G^{(s-1)} v_t\| \quad (30)$$

$$\leq \sum_{s=(\log \log n)^3}^{m-C} \left(\sqrt{\lambda_L} \right)^{m-s} \Lambda_m^{s-1} \cdot \sqrt{n} \cdot (1 + \delta)^{m-s+1} \quad (31)$$

$$\leq (1 + \delta) \Lambda_m^m \sqrt{n} \cdot \left(\sum_{s=(\log \log n)^3}^{m-C} \left(\frac{\left(\sqrt{\lambda_L} \right) (1 + \delta)}{\Lambda_m} \right)^{m-s} \right) \quad (32)$$

$$= \Lambda_m^m \cdot \sqrt{n} \cdot \left(\frac{1}{\delta (1 + \delta)^{C-2}} \right) \quad (33)$$

Using (29) and (33), for sufficiently large $C = O_{M, \delta}(1)$, we conclude that

$$\Omega_M(1) \cdot \Lambda_m^m \cdot \sqrt{n} \leq \sum_{s=m-C+1}^{m-1} \left(\sqrt{\lambda_L} \right)^{m-s} \left| \left\langle \mathbf{e}_{\beta'} \otimes \underline{\chi}^{\tau', \alpha'}, \underline{A}_G^{(s-1)} v_m \right\rangle \right| \quad (34)$$

Note that there are only $C = O_{M, \delta}(1)$ terms in the sum, so one of them is large. In particular, there exists some $\ell \in [m - C, m]$ such that if we set

$$w_\ell \stackrel{\text{def}}{=} \underline{A}_G^{(\ell)} v_m$$

then,

$$\left| \left\langle \mathbf{e}_{\beta'} \otimes \underline{\chi}^{\tau', \alpha'}, w_\ell \right\rangle \right| \geq \frac{1}{C \left(\sqrt{\lambda_L} \right)^C} \cdot \Lambda_m^m \cdot \Omega_{M, \delta}(1) \cdot \sqrt{n},$$

But note that by the choice of m ,

$$\|w_\ell\| \leq \Lambda_m^\ell \cdot (1 + \delta)^{m-\ell} \leq \Lambda_m^m \cdot (1 + \delta)^C.$$

So there exists some $\ell \in [m - C, m]$, such that the unit vector $\bar{w}_\ell = \frac{w_\ell}{\|w_\ell\|}$ satisfies,

$$\left| \left\langle \mathbf{e}_{\beta'} \otimes \underline{\chi}^{\tau', \alpha'}, \bar{w}_\ell \right\rangle \right| \geq \Omega_{M, \delta}(1) \cdot \sqrt{n}.$$

In other words, for some choice of $\ell \in [m - C, m]$ and $\beta' \in [q]$, if we construct $\bar{u} \in \mathbb{R}^n$ as,

$$\bar{u}[i] = \bar{w}_\ell[i, \beta']$$

then $|\langle \bar{u}, \underline{\chi}^{\tau', \alpha'} \rangle| \geq \Omega_{M, \delta}(1) \cdot \sqrt{n}$. This finishes the proof of [Theorem 8.1](#).⁶

Proof. (Proof of [Claim 8.6](#)) The idea behind the proof is a descent/bootstrap argument to get a contradiction. Let us start with $t = (\log \log n)^5$ as the guess for m . If current value of t satisfies the condition of the claim, we are done. Otherwise, there exists $s < t$ such that,

$$\|A_G^{(s)} v_t\| \geq \Lambda_t^s (1 + \delta)^{t-s}$$

This implies that,

$$\Lambda_s^s \geq \Lambda_t^s (1 + \delta)^{t-s}$$

or equivalently,

$$\log \Lambda_s \geq \log \Lambda_t + \log(1 + \delta) \cdot (t - s)/s$$

Suppose we use s as the new candidate for m and recurse. Let us suppose we iteratively construct a sequence of $t_0 = (\log \log n)^5 > \dots > t_r$ in this manner. The value of $\log \Lambda_{t_i}$ increases along the sequence. By [Fact 8.7](#), if we obtain a sequence of $t_0 = (\log \log n)^5 > \dots > t_r = (\log \log n)^{4.5}$ then we will have,

$$\log \Lambda_r \geq (\log t_0 - \log t_r) \log(1 + \delta) \geq \log(1 + \delta) \cdot \Omega(\log \log \log n)$$

This suggests that $\|A_G^{(t_r)}\|^{1/t_r} \geq \omega(1)$ for some $t_r = \Omega(\log \log n)^{4.5}$. A contradiction, since with probability $1 - o_n(1)$, we will have $\|A_G^{(t_r)}\|^{1/t_r} = O(1)$. This follows from the fact that with probability $1 - o_n(1)$, degree of every vertex in $A_G^{(s)}$ is at most $O(D^s)$ for some constant D for all $s > (\log \log n)^2$. Therefore, the sequence terminates and we find a $t_r \in [(\log \log n)^{4.5}, (\log \log n)^5]$, implying the claim.

Fact 8.7. Given a sequence of positive integers, $a_1 \geq a_2 \geq \dots a_r$,

$$\sum_{i=1}^r \frac{a_i - a_{i-1}}{a_{i-1}} \geq \sum_{i=1}^r \sum_{x=a_{i-1}}^{a_i-1} \frac{1}{x} = \sum_{x=a_r}^{a_1} \frac{1}{x} \approx \ln(a_1) - \ln(a_r) \quad (35)$$

□

⁶Note that the $\underline{\chi}^{\tau', \alpha'}$ here is a bit different from the $\underline{\chi}^{\tau', \alpha'}$ defined in the theorem, but by [Remark 8.2](#) they are within a multiplicative factor of $(1 \pm \epsilon)$ from each other. Thus the inequality still hold for the $\underline{\chi}^{\tau', \alpha'}$ defined in the theorem statement.

Proof. (Proof of [Lemma 8.4](#)) For $s \in \mathbb{N}$, let \mathcal{P}_s be the set of length s non-backtracking walks in complete graph \mathcal{K}_n . So $\underline{\alpha} = (\alpha_0, \alpha_1, \dots, \alpha_s) \in \mathcal{P}_s$ will be a non-backtracking path in \mathcal{K}_n with vertices $\alpha_0, \dots, \alpha_s$.

For $\ell, \ell' \in [n]$, we can write

$$A^{(s)}RB^{(t)}[\ell, \ell'] = \sum_{\substack{\underline{\alpha} \in \mathcal{P}_s, \underline{\beta} \in \mathcal{P}_t \\ \alpha_0 = \ell, \beta_t = \ell'}} \cdot \prod_{i=1}^s A_{\alpha_{i-1}\alpha_i} \cdot R_{\alpha_s, \beta_0} \cdot \prod_{j=1}^t B_{\beta_{j-1}\beta_j} \quad (36)$$

$$A^{(s) \circ R \circ B^{(t)}}[\ell, \ell'] = \sum_{\substack{\underline{\alpha} \in \mathcal{P}_s, \underline{\beta} \in \mathcal{P}_t \\ \alpha_0 = \ell, \beta_t = \ell' \\ \alpha_s \neq \beta_0, \alpha_{s-1} \neq \beta_0, \alpha_s \neq \beta_1}} \cdot \prod_{i=1}^s A_{\alpha_{i-1}\alpha_i} \cdot R_{\alpha_s, \beta_0} \prod_{j=1}^t B_{\beta_{j-1}\beta_j} \quad (37)$$

It is clear that the difference $A^{(s) \circ R \circ B^{(t)}} - A^{(s)}RB^{(t)}$ consists of three different terms.

Term 1: $\alpha_s = \beta_0$ Consider the block-diagonal matrix $\mathcal{D}_1 \in \mathbb{R}^{nq \times nq}$ given by,

$$\mathcal{D}_1[i, j] = \mathbf{1}[i = j] \cdot R_{ij}.$$

then we can write this term as $A^{(s)}\mathcal{D}_1B^{(t)}$. Hence we get the following bound,

$$\|A^{(s)}\mathcal{D}_1B^{(t)}\| \leq q\|R\|_\infty \cdot \|A^{(s)}\| \cdot \|B^{(t-s)}\|$$

Term 2: $\alpha_{s-1} = \beta_0$ Consider the block-diagonal matrix $\mathcal{D}_2 \in \mathbb{R}^{nq \times nq}$ given by,

$$\mathcal{D}_2[i, i] = \sum_j A_{ij}R_{ji}.$$

then we can write this term as $A^{(s-1)}\mathcal{D}_2B^{(t)}$. We have the upper bound,

$$\|\mathcal{D}_2\| \leq q\|R\|_\infty \|A\|_{1 \rightarrow 1}$$

which implies that

$$\|A^{(s-1)}\mathcal{D}_2B^{(t)}\| \leq q\|x\|_\infty \|y\|_\infty \|A\|_{1 \rightarrow 1} \|A^{(s-1)}\| \|B^{(t)}\|$$

Term 3: $\alpha_s = \beta_1$ Consider the block-diagonal matrix $\mathcal{D}_3 \in \mathbb{R}^{nq \times nq}$ given by,

$$\mathcal{D}_3[i, i] = \left(\sum_j R_{ij}B_{ji} \right).$$

then we can write this term as $A^{(s)}\mathcal{D}_3B^{(t-1)}$. Analogous to the previous case, we get an upper bound of,

$$\|A^{(s)}\mathcal{D}_3B^{(t-1)}\| \leq q\|R\|_\infty \|B^T\|_{1 \rightarrow 1} \|A^{(s)}\| \|B^{(t-1)}\|$$

Adding the three terms, we have the claim of the lemma. □

We will need the following theorem about local statistics on the expectation and concentration of local statistics in order to complete the proof of [Lemma 8.5](#).

Lemma 8.8. Fix a vector $x \in \mathbb{R}^{nq}$, such that $\|x\|_\infty = O(1)$. With probability $1 - o_n(1)$, for all $1 \leq s \leq \sqrt{\log n}$ we have,

$$\|\underline{A}_G^{(s)} x\| \leq C \left(\sqrt{\lambda_L} \right)^s \cdot \sqrt{n} \quad (38)$$

for an absolute constant $C \geq 1$.

Proof sketch. This is equivalent to proving $\langle \underline{A}_G^{(s)} x, \underline{A}_G^{(s)} x \rangle \leq C \left(\sqrt{\lambda_L} \right)^{2s} n$. Indeed, this quantity can be rewritten as:

$$\left\langle \left(\underline{A}_G^{(s)} \right)^2, x x^\top \right\rangle.$$

Via similar calculations to the ones done in [Section 6](#), we can show that this quantity is dominated by the contribution of walks that are self-avoiding for the first s steps, and retrace the same steps taken in the next s steps, which in turn can be used to show that this quantity concentrates around the expected total weight of walks in the associated random tree that walk out s steps and walk back s steps, and hence for large enough s is at most $C \left(\sqrt{\lambda_L} \right)^{2s} n$ for an absolute constant C . \square

Proof of [Lemma 8.5](#). Let $\Gamma_{\alpha, \alpha'}^{\tau, \tau'} = \sum_{j \in [q]} u_j v_j^\top$ be the singular decomposition of $\Gamma_{\alpha, \alpha'}^{\tau, \tau'}$.

$$\| \underline{B}_G^{(s)} \Gamma_{\alpha, \alpha'}^{\tau, \tau'} \otimes \underline{\chi}^{\tau, \alpha} (\mu^{\tau', \alpha'})^\top \underline{A}_G^{(t)} \| \leq \sum_j \| \underline{B}_G^{(s)} u_j \otimes \underline{\chi}^{\tau, \alpha} \| \| v_j \otimes (\mu^{\tau', \alpha'})^\top \underline{A}_G^{(t)} \| \quad (39)$$

Applying [Lemma 8.8](#) to the planted model M with the fixed vector $u_j \otimes \underline{\chi}^{\tau, \alpha}$, we conclude that with probability $1 - o_n(1)$,

$$\| \underline{B}_G^{(s)} u_j \otimes \underline{\chi}^{\tau, \alpha} \| \leq C \left(\sqrt{\lambda_L} \right)^s \cdot \| u_j \otimes \underline{\chi}^{\tau, \alpha} \| \leq \left(\sqrt{\lambda_L} \right)^s \cdot \| u_j \| \cdot n^{1/2} \quad (40)$$

Similarly, applying [Lemma 8.8](#) to the null model M^\times with fixed vector $v_j \otimes \mu^{\tau', \alpha'}$, we conclude that with probability $1 - o_n(1)$, for all $1 \leq t \leq \sqrt{\log n}$,

$$\| v_j \otimes (\mu^{\tau', \alpha'})^\top \underline{A}_G^{(t)} \| \leq C \left(\sqrt{\lambda_L} \right)^s \cdot \| v_j \otimes (\mu^{\tau', \alpha'})^\top \| \leq \left(\sqrt{\lambda_L} \right)^t \cdot \| v_j \| \cdot n^{1/2} \quad (41)$$

Finally, note that $\sum_j \| u_j \| \| v_j \| = \| \Gamma_{\alpha, \alpha'}^{\tau, \tau'} \|_{Fr} = O_M(\frac{1}{n})$ where O_M hides a fixed constant depending on the model. Using (40) and (41) in (39), we conclude the proof. \square

Acknowledgments

We would like to thank Tselil Schramm for insightful discussions around belief propagation, and we would like to thank Ryan O'Donnell and Xinyu Wu for conversations that brought clarity to various parts of [\[BC19\]](#). We would also like to thank Lenka Zdeborová and anonymous reviewers for bringing several relevant references to our attention.

References

- [AC08] D. Achlioptas and A. Coja-Oghlan. Algorithmic barriers from phase transitions. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 793–802, 2008. [6](#)
- [ACKZ15] Maria Chiara Angelini, Francesco Caltagirone, Florent Krzakala, and Lenka Zdeborová. Spectral detection on sparse hypergraphs. In *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 66–73. IEEE, 2015. [6](#)
- [AM09] Per Austrin and Elchanan Mossel. Approximation resistant predicates from pairwise independence. *Computational Complexity*, 18(2):249–271, 2009. [3](#)

- [AM13] Emmanuel Abbe and Andrea Montanari. Conditional random fields, planted constraint satisfaction and entropy concentration. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 332–346. Springer, 2013. [6](#)
- [AOW15] Sarah R Allen, Ryan O’Donnell, and David Witmer. How to refute a random CSP. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 689–708. IEEE, 2015. [2](#)
- [AS15] Emmanuel Abbe and Colin Sandon. Community detection in general stochastic block models: Fundamental limits and efficient algorithms for recovery. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 670–688. IEEE, 2015. [2](#), [7](#), [17](#)
- [BB20] Matthew Brennan and Guy Bresler. Reducibility and Statistical-Computational Gaps from Secret Leakage. *arXiv preprint arXiv:2005.08099*, 2020. [10](#)
- [BBH18] Matthew Brennan, Guy Bresler, and Wasim Huleihel. Reducibility and computational lower bounds for problems with planted sparse structure. *arXiv preprint arXiv:1806.07508*, 2018. [10](#)
- [BBH⁺20] Matthew Brennan, Guy Bresler, Samuel B Hopkins, Jerry Li, and Tselil Schramm. Statistical query algorithms and low-degree tests are almost equivalent. *arXiv preprint arXiv:2009.06107*, 2020. [10](#)
- [BC19] Charles Bordenave and Benoît Collins. Eigenvalues of random lifts and polynomials of random permutation matrices. *Annals of Mathematics*, 190(3):811–875, 2019. [7](#), [8](#), [9](#), [10](#), [21](#), [24](#), [27](#), [37](#), [40](#), [53](#)
- [BLM15] Charles Bordenave, Marc Lelarge, and Laurent Massoulié. Non-backtracking spectrum of random graphs: community detection and non-regular Ramanujan graphs. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 1347–1357. IEEE, 2015. [1](#), [2](#), [4](#), [6](#), [7](#), [9](#), [17](#), [23](#), [37](#)
- [BMR19] Jess Banks, Sidhanth Mohanty, and Prasad Raghavendra. Local Statistics, Semidefinite Programming, and Community Detection. *arXiv preprint arXiv:1911.01960*, 2019. [10](#), [23](#), [39](#), [41](#), [63](#)
- [Bor19] Charles Bordenave. A new proof of Friedman’s second eigenvalue Theorem and its extension to random lifts. In *Annales scientifiques de l’Ecole normale supérieure*, 2019. [24](#), [37](#)
- [COHKL⁺20] Amin Coja-Oghlan, Max Hahn-Klimroth, Philipp Loick, Noela Müller, Konstantinos Panagiotou, and Matija Pasch. Inference and mutual information on random factor graphs. *arXiv preprint arXiv:2007.07494*, 2020. [6](#)
- [DKMZ11a] Aurelien Decelle, Florent Krzakala, Cristopher Moore, and Lenka Zdeborová. Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Physical Review E*, 84(6):066106, 2011. [2](#), [3](#), [7](#), [8](#), [16](#), [17](#)
- [DKMZ11b] Aurelien Decelle, Florent Krzakala, Cristopher Moore, and Lenka Zdeborová. Inference and phase transitions in the detection of modules in sparse networks. *Physical Review Letters*, 107(6):065701, 2011. [2](#), [7](#)

- [DMO⁺19] Yash Deshpande, Andrea Montanari, Ryan O’Donnell, Tselil Schramm, and Subhabrata Sen. The threshold for SDP-refutation of random regular NAE-3SAT. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2305–2321. SIAM, 2019. [3](#)
- [DSMM18] Yash Deshpande, Subhabrata Sen, Andrea Montanari, and Elchanan Mossel. Contextual stochastic block models. In *Advances in Neural Information Processing Systems*, pages 8581–8593, 2018. [13](#)
- [FGR⁺17] Vitaly Feldman, Elena Grigorescu, Lev Reyzin, Santosh S Vempala, and Ying Xiao. Statistical algorithms and a lower bound for detecting planted cliques. *Journal of the ACM (JACM)*, 64(2):1–37, 2017. [10](#)
- [FM17] Zhou Fan and Andrea Montanari. How well do local algorithms solve semidefinite programs? In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 604–614, 2017. [23](#), [64](#)
- [Fri03] Joel Friedman. A proof of Alon’s second eigenvalue conjecture. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 720–724, 2003. [24](#), [37](#)
- [GMPS18] Sainyam Galhotra, Arya Mazumdar, Soumyabrata Pal, and Barna Saha. The geometric block model, 2018. [13](#)
- [Gol00] Oded Goldreich. Candidate One-Way Functions Based on Expander Graphs. *IACR Cryptol. ePrint Arch.*, 2000:63, 2000. [11](#)
- [Gol11] Oded Goldreich. Candidate one-way functions based on expander graphs. In *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation*, pages 76–87. Springer, 2011. [2](#)
- [HLM12] Simon Heimlicher, Marc Lelarge, and Laurent Massoulié. Community detection in the labelled stochastic block model. *arXiv preprint arXiv:1209.2910*, 2012. [13](#)
- [HS17] Samuel B Hopkins and David Steurer. Efficient bayesian estimation from few samples: community detection and related problems. In *Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on*, pages 379–390. IEEE, 2017. [10](#)
- [Kho02] Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, pages 767–775, 2002. [2](#)
- [KKMO07] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O’Donnell. Optimal inapproximability results for max-cut and other 2-variable csps? *SIAM Journal on Computing*, 37(1):319–357, 2007. [2](#)
- [KMM⁺13] Florent Krzakala, Cristopher Moore, Elchanan Mossel, Joe Neeman, Allan Sly, Lenka Zdeborová, and Pan Zhang. Spectral redemption in clustering sparse networks. *Proceedings of the National Academy of Sciences*, 110(52):20935–20940, 2013. [1](#), [4](#), [6](#)

- [KMOW17] Pravesh K Kothari, Ryuhei Mori, Ryan O’Donnell, and David Witmer. Sum of squares lower bounds for refuting any CSP. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 132–145, 2017. [2](#)
- [KMZ12] Florent Krzakala, Marc Mézard, and Lenka Zdeborová. Reweighted belief propagation and quiet planting for random k-sat. *Journal on Satisfiability, Boolean Modeling and Computation*, 8(3-4):149–171, 2012. [7](#)
- [KZ09] Florent Krzakala and Lenka Zdeborova. Hiding quiet solutions in random constraint satisfaction problems. *Physical review letters*, 102:238701, 07 2009. [2](#), [3](#), [5](#), [6](#), [10](#), [17](#)
- [LMP01] John Lafferty, Andrew McCallum, and Fernando CN Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. 2001. [6](#)
- [Mas14] Laurent Massoulié. Community detection thresholds and the weak ramanujan property. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 694–703, 2014. [2](#), [7](#), [17](#)
- [MM09a] Marc Mezard and Andrea Montanari. Information, physics, and computation. 2009. [2](#), [15](#)
- [MM09b] Marc Mezard and Andrea Montanari. *Information, physics, and computation*. Oxford University Press, 2009. [6](#)
- [MNS14] Elchanan Mossel, Joe Neeman, and Allan Sly. Belief propagation, robust reconstruction and optimal recovery of block models. In *Conference on Learning Theory*, pages 356–370, 2014. [11](#)
- [MNS18] Elchanan Mossel, Joe Neeman, and Allan Sly. A proof of the block model threshold conjecture. *Combinatorica*, 38(3):665–708, 2018. [2](#), [7](#), [17](#)
- [Mon08] Andrea Montanari. Estimating random variables from random sparse observations. *European Transactions on Telecommunications*, 19(4):385–403, 2008. [1](#), [6](#), [12](#)
- [MOP20] Sidhanth Mohanty, Ryan O’Donnell, and Pedro Paredes. The SDP Value for Random Two-Eigenvalue CSPs. In *37th International Symposium on Theoretical Aspects of Computer Science (STACS 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020. [24](#), [37](#), [65](#), [66](#)
- [MV17] Andrea Montanari and Ramji Venkataramanan. Estimation of low-rank matrices via approximate message passing. *arXiv preprint arXiv:1711.01682*, 2017. [11](#)
- [Nis01] Hidetoshi Nishimori. *Statistical physics of spin glasses and information processing: an introduction*. Number 111. Clarendon Press, 2001. [6](#)
- [OW20] Ryan O’Donnell and Xinyu Wu. Explicit near-fully X-Ramanujan graphs. *arXiv preprint arXiv:2009.02595*, 2020. [7](#), [24](#), [40](#)
- [Rag08] Prasad Raghavendra. Optimal algorithms and inapproximability results for every csp? In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 245–254, 2008. [2](#)

- [RRS17] Prasad Raghavendra, Satish Rao, and Tselil Schramm. Strongly refuting random CSPs below the spectral threshold. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 121–131, 2017. [2](#)
- [RTSZ19] Federico Ricci-Tersenghi, Guilhem Semerjian, and Lenka Zdeborová. Typology of phase transitions in bayesian inference problems. *Physical Review E*, 99(4):042109, 2019. [6](#), [7](#)
- [Sch78] Thomas J Schaefer. The complexity of satisfiability problems. In *Proceedings of the tenth annual ACM symposium on Theory of computing*, pages 216–226, 1978. [2](#)
- [SKZ17] Alaa Saade, Florent Krzakala, and Lenka Zdeborová. Spectral bounds for the ising ferromagnet on an arbitrary given graph. *Journal of Statistical Mechanics: Theory and Experiment*, 2017(5):053403, 2017. [6](#)
- [SLKZ15] Alaa Saade, Marc Lelarge, Florent Krzakala, and Lenka Zdeborová. Spectral detection in the censored block model. In *2015 IEEE International Symposium on Information Theory (ISIT)*, pages 1184–1188. IEEE, 2015. [6](#)
- [Sly09] Allan Sly. Reconstruction for the potts model. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 581–590, 2009. [7](#)
- [SM20] Ludovic Stephan and Laurent Massoulié. Non-backtracking spectra of weighted inhomogeneous random graphs. *arXiv preprint arXiv:2004.07408*, 2020. [7](#), [9](#)
- [Zhu20] Dmitriy Zhuk. A proof of the csp dichotomy conjecture. *Journal of the ACM (JACM)*, 67(5):1–78, 2020. [2](#)
- [ZK11] Lenka Zdeborová and Florent Krzakala. Quiet planting in the locked constraint satisfaction problems. *SIAM Journal on Discrete Mathematics*, 25(2):750–770, 2011. [7](#)
- [ZK16a] Lenka Zdeborová and Florent Krzakala. Statistical physics of inference: Thresholds and algorithms. *Advances in Physics*, 65(5):453–552, 2016. [2](#)
- [ZK16b] Lenka Zdeborová and Florent Krzakala. Statistical physics of inference: Thresholds and algorithms. *Advances in Physics*, 65(5):453–552, 2016. [6](#)
- [ZMZ14] Pan Zhang, Cristopher Moore, and Lenka Zdeborová. Phase transitions in semisupervised clustering of sparse networks. *Physical Review E*, 90(5), Nov 2014. [13](#), [16](#)

A Belief propagation for M

We briefly describe the belief propagation (BP) algorithm that aims to estimate the marginal distribution of $c(v)$, $v \in [n]$ under the Boltzmann distribution μ with Hamiltonian H . Define the messages $\{m_c^{v \rightarrow e}\}_{c \in [q]}$ that a variable v passes to some constraint $e \in E_i$, and the messages $\{m_c^{e \rightarrow u}\}_{c \in [q]}$ that a constraint $e \in E_i$ passes to a variable u . Intuitively speaking, $m_c^{v \rightarrow e}$ is an estimate of the marginal probability that v is assigned the color c when the constraint e is absent, and $m_c^{e \rightarrow u}$ is an estimate of the marginal probability that u has color c when all other constraints involving u are absent. Since the distribution of $c(v)$ under μ depends on the constraints that contain v , we only focus on the messages $m^{v \rightarrow e}$, $m^{e \rightarrow u}$ such that $v \in \partial e$ (i.e. e contains v) and $e \in \partial u$ (i.e. e contains u).

$$m_c^{v \rightarrow e}[t+1] = \frac{1}{Z^{v \rightarrow e}} \mathbb{P}_{\tau(v)}(c) \prod_{f \in \partial v \setminus e} m_c^{f \rightarrow v}[t], \quad (42)$$

where $Z^{v \rightarrow e} = \sum_{c \in [q]} \mathbb{P}_{\tau(v)}(c) \prod_{f \in \partial v \setminus e} m_c^{f \rightarrow v}[t]$.

For factors $e \in \mathbf{E}_i$, the messages are defined as

$$m_c^{e \rightarrow u}[t+1] = \frac{1}{Z^{e \rightarrow u}} \sum_{c_e | c_e(u)=c} \phi_i(\mathbf{c}_e) \prod_{v \in e \setminus u} m_{c_e(v)}^{v \rightarrow e}[t], \quad (43)$$

where $Z^{e \rightarrow u} = \sum_{c \in [q]} \sum_{c_e | c_e(u)=c} \phi_i(\mathbf{c}_e) \prod_{v \in e \setminus u} m_{c_e(v)}^{v \rightarrow e}[t]$.

To obtain an estimate of the marginal probability of the assignment to a variable v , apply the message update rules until reaching some fixed point $\{\hat{m}_c^{v \rightarrow e}, \hat{m}_c^{e \rightarrow u}\}_{c \in [q]}$. The estimate is called the belief and is given by

$$b_c^v = \frac{1}{Z^v} \mathbb{P}_{\tau(v)}(c) \prod_{f \in \partial v} \hat{m}_c^{f \rightarrow v},$$

where $Z^v = \sum_{c \in [q]} \mathbb{P}_{\tau(v)}(c) \prod_{f \in \partial v} \hat{m}_c^{f \rightarrow v}$.

B Proof of Lemma 2.5

Recall that:

$$\bar{\phi}_i = \sum_{(c_1, \dots, c_{a(i)}) \in [q]^{a(i)}} \left(\prod_{k=1}^{a(i)} \mathbb{P}_{\text{Cl}(i)_k}(c_k) \right) \cdot \phi_i(c_1, \dots, c_{a(i)}) \quad (44)$$

We first explain how to solve the distinguishing problem and then explain the recovery algorithm. By definition of the BP update functions $Y_{v \rightarrow e}$ (equation (42)) and $Y_{e \rightarrow v}$ (equation (43)), the set of trivial messages \bar{m} being a BP fixed point is equivalent to:

$$\bar{\phi}_i = \sum_{\substack{(c_1, \dots, c_{a(i)}) \in [q]^{a(i)} \\ c_j = c}} \left(\prod_{k \neq j} \mathbb{P}_{\text{Cl}(i)_k}(c_k) \right) \cdot \phi_i(c_1, \dots, c_{a(i)}) \quad \forall c \in [q] : \mathbb{P}_{\text{Cl}(i)_j}(c) > 0 \quad (45)$$

If the set of trivial messages is not a fixed point of the belief propagation update rule: then there exist $c \in [q]$, $i \in [F]$, and $j \in [a(i)]$ with $\mathbb{P}_{\text{Cl}(i)_j}(c) > 0$ such that

$$\bar{\phi}_i \neq \sum_{\substack{(c_1, \dots, c_{a(i)}) \in [q]^{a(i)} \\ c_j = c}} \left(\prod_{k \neq j} \mathbb{P}_{\text{Cl}(i)_k}(c_k) \right) \cdot \phi_i(c_1, \dots, c_{a(i)}).$$

Let $\text{deg}_{i,j}(v)$ be the number of type- i factors with variable v in the j -th position. Via standard results for Poisson(d) approximating Binom($n, d/n$) we have the following:

- In \mathbf{M}^\times , for any variable v of type $\text{Cl}(i)_j$ and any constant T ,

$$\text{Edeg}_{i,j}(v)^T = \mathbf{E}X^T \pm o_n(1)$$

where $X \sim \text{Poisson}(\lambda)$ and $\lambda = \bar{\phi}_i$.

- On the other hand, in the planted model M:

$$\mathbf{E} \deg_{i,j}(v)^T = \mathbf{E} Y^T \pm o_n(1)$$

where Y is distributed as the mixture of Poisson distributions $p_1 \text{Poisson}(\lambda_1) + \dots + p_s \text{Poisson}(\lambda_s)$ where s is the number of colors which vertex v has nonzero probability of attaining, not all λ_i are equal, and all $p_i > 0$.

By (44) $p_1 \lambda_1 + \dots + p_s \lambda_s = \lambda$. We first recall the following well known fact about Poisson random variables.

Fact B.1. *If $A \sim \text{Poisson}(\mu)$, then $\mathbf{E} A^2 = \mu^2 + \mu$.*

As a consequence of **Fact B.1**: $\mathbf{E} X^2 = \lambda^2 + \lambda$, and $\mathbf{E} Y^2 = p_1(\lambda_1^2 + \lambda_1) + \dots + p_s(\lambda_s^2 + \lambda_s)$.

$$\mathbf{E} Y^2 - \mathbf{E} X^2 = p_1 f(\lambda_1) + \dots + p_s f(\lambda_s) - f(\lambda) = p_1 f(\lambda_1) + \dots + p_s f(\lambda_s) - f(p_1 \lambda_1 + \dots + p_s \lambda_s)$$

Since not all λ_i are equal, all $p_i > 0$ and f is strictly convex, $\mathbf{E} Y^2 - \mathbf{E} X^2$ is equal to a constant δ strictly greater than 0. Suppose $n_{i,j,2}(\mathbf{G}) := \mathbf{E} \sum_{v \in [n]} \deg_{i,j}(v)^2$, then $|\mathbf{E}_{G \sim M^*} n_{i,j,2}(\mathbf{G}) - \mathbf{E}_{G \sim M} n_{i,j,2}(\mathbf{G})| \geq \Omega(n)$. Since $\mathbf{E} Y^4$ and $\mathbf{E} X^4$ are constants, the variance of $n_{i,j,2}(\mathbf{G})$ is $O(n)$ for both $G \sim M^*$ and $G \sim M$. This informs using the following polynomial time distinguisher:

Compute $n_{i,j,2}(\mathbf{G})$ and if $|n_{i,j,2}(\mathbf{G}) - \mathbf{E}_{G \sim M^*} [n_{i,j,2}(\mathbf{G})]| < |n_{i,j,2}(\mathbf{G}) - \mathbf{E}_{G \sim M} [n_{i,j,2}(\mathbf{G})]|$ output “null”; otherwise output “planted”.

We now discuss performing recovery. Recall the inner product $\langle \cdot, \cdot \rangle_{\mathbf{H}}$ from [Section 5](#) which is defined as follows: First, we define a $nq \times nq$ -dimensional positive diagonal matrix \mathbf{H}_{τ} where the (v, v) block is equal to:

$$\mathbf{H}_{\tau, (v, v)}[c, c] := \begin{cases} \mathbb{P}_{\tau(v)}(c) & \text{if } \mathbb{P}_{\tau(v)}(c) > 0 \\ 1 & \text{otherwise.} \end{cases}$$

The inner product on \mathbb{R}^{nq} is then:

$$\langle x, y \rangle_{\mathbf{H}} := x^T \mathbf{H}_{\tau}^{-1} y.$$

And let $\|\cdot\|$ denote the norm induced by the above inner product. Let \mathbf{c} be the hidden coloring. Our goal in recovery is to output a vector v such that $\langle v, \mathbf{c} - \mathbf{E} \mathbf{c} | \tau \rangle_{\mathbf{H}} \geq \epsilon \cdot \|v\| \cdot \|\mathbf{c} - \mathbf{E} \mathbf{c} | \tau\|$. Let c and c' be two colors such that:

$$d_c = \sum_{\substack{(c_1, \dots, c_{a(i)}) \in [q]^{a(i)} \\ c_j = c}} \left(\prod_{k \neq j} \mathbb{P}_{\text{Cl}(i)_k}(c_k) \right) \cdot \phi_i(c_1, \dots, c_{a(i)}) > \sum_{\substack{(c_1, \dots, c_{a(i)}) \in [q]^{a(i)} \\ c_j = c'}} \left(\prod_{k \neq j} \mathbb{P}_{\text{Cl}(i)_k}(c_k) \right) \cdot \phi_i(c_1, \dots, c_{a(i)}) = d_{c'}$$

The distribution of the number of type- i factors that a color c vertex is part of is $\text{Poisson}(d_c)$ and similarly is $\text{Poisson}(d_{c'})$ for a color c' vertex. The following algorithm can then be shown to produce a vector v meeting the aforementioned goal.

For each vertex u of type $\mathcal{T}(\text{Cl}(i)_j)$, let m_u be the number of type- i factors it is part of in the j -th position. If m_u has a higher probability of being sampled from $\text{Poisson}(d_c)$ than $\text{Poisson}(d_{c'})$ then assign the u -th block of vector v to be the indicator of color c . Otherwise assign the u -th block of vector v to be the indicator of color c' .

Since $d_c \neq d_{c'}$ there is a constant $\varepsilon > 0$ such that with high probability $(\frac{1}{2} + \varepsilon) \mathbb{P}_{\text{Cl}(i_j)}(c) \mathcal{T}(\text{Cl}(i_j)) n$ variables of color c and type $\text{Cl}(i_j)$ are assigned the correct color and also $(\frac{1}{2} + \varepsilon) \mathbb{P}_{\text{Cl}(i_j)}(c') \mathcal{T}(\text{Cl}(i_j)) n$ variables of color c' and type $\text{Cl}(i_j)$ are assigned the correct color. Consequently:

$$\langle v, c - \mathbf{E}c | \tau \rangle_{\mathbb{H}} \geq \varepsilon' \cdot \|v\| \cdot \|c - \mathbf{E}c | \tau\|$$

for some $\varepsilon' > 0$.

C The partial derivative matrix

Recall the BP update function Γ defined by equations (42) and (43).

We observe that by definition

$$\bar{\mathbf{M}}_{\theta(e_j), i(v_j) | i(v_{j+1})} = \frac{\partial \Gamma(m)^{v_j \rightarrow e_{j-1}}}{\partial m^{e_j \rightarrow v_j}} \Big|_{\bar{m}} \cdot \frac{\partial \Gamma(m)^{e_j \rightarrow v_j}}{\partial m^{v_{j+1} \rightarrow e_j}} \Big|_{\bar{m}}.$$

Thus we first compute the two derivative matrices. For any pairs of colors $c, d \in [q]$,

$$\begin{aligned} \frac{\partial \Gamma(m)_c^{v_j \rightarrow e_{j-1}}}{\partial m_d^{e_j \rightarrow v_j}} &= \frac{1}{Z^{v_j \rightarrow e_{j-1}}} \mathbb{P}_{\tau(v_j)}(c) \prod_{a \in \partial v_j \setminus \{e_{j-1}, e_j\}} m_c^{a \rightarrow v_j} \cdot \mathbf{1}_{d=c} - \frac{m_c^{v_j \rightarrow e_{j-1}}}{Z^{v_j \rightarrow e_{j-1}}} \cdot \mathbb{P}_{\tau(v_j)}(d) \prod_{a \in \partial v_j \setminus \{e_{j-1}, e_j\}} m_d^{a \rightarrow v_j} \\ &= \frac{1}{m_d^{e_j \rightarrow v_j}} \cdot \frac{1}{Z^{v_j \rightarrow e_{j-1}}} \mathbb{P}_{\tau(v_j)}(c) \prod_{a \in \partial v_j \setminus \{e_{j-1}\}} m_c^{a \rightarrow v_j} \cdot \mathbf{1}_{d=c} - \frac{m_c^{v_j \rightarrow e_{j-1}}}{m_d^{e_j \rightarrow v_j}} \cdot \frac{\mathbb{P}_{\tau(v_j)}(d)}{Z^{v_j \rightarrow e_{j-1}}} \prod_{a \in \partial v_j \setminus \{e_{j-1}\}} m_d^{a \rightarrow v_j} \\ &= \frac{m_c^{v_j \rightarrow e_{j-1}}}{m_d^{e_j \rightarrow v_j}} \cdot \mathbf{1}_{d=c} - \frac{m_c^{v_j \rightarrow e_{j-1}}}{m_d^{e_j \rightarrow v_j}} \cdot m_d^{v_j \rightarrow e_{j-1}} \end{aligned}$$

The last equality is derived from the fixed point identity $m_c^{v \rightarrow e} = \frac{1}{Z^{v \rightarrow e}} \mathbb{P}_{\tau(v)}(c) \prod_{a \in \partial v \setminus e} m_c^{a \rightarrow v}$.

Evaluating the derivative at the factorized fixed point gives the transformation matrix

$$\frac{\partial \Gamma(m)_c^{v_j \rightarrow e_{j-1}}}{\partial m_d^{e_j \rightarrow v_j}} \Big|_{\bar{m}} = \text{support}(\mathbb{P}_{\tau(v_j)}) \cdot (\mathbb{P}_{\tau(v_j)}(c) \cdot \mathbf{1}_{d=c} - \mathbb{P}_{\tau(v_j)}(c) \cdot \mathbb{P}_{\tau(v_j)}(d)) ,$$

where $\text{support}(\mathbb{P}_{\tau(v_j)})$ denote the size of $\mathbb{P}_{\tau(v_j)}$'s support. To write the matrix compactly we define $\mathbf{D}_{\tau} := \text{Diag}(\mathbb{P}_{\tau})$, and derive from the above computation that $\frac{\partial \Gamma(m)^{v_j \rightarrow e_{j-1}}}{\partial m^{e_j \rightarrow v_j}} \Big|_{\bar{m}} = \text{support}(\mathbb{P}_{\tau(v_j)}) \cdot (\mathbf{D}_{\tau(v_j)} - \mathbb{P}_{\tau(v_j)} \mathbb{P}_{\tau(v_j)}^T)$.

For any edge of the form $v_{j+1} \xrightarrow{e_j} v_j$ on this path where $\theta(e_j) = \phi_i$ we have,

$$\begin{aligned} \frac{\partial \Gamma(m)_c^{e_j \rightarrow v_j}}{\partial m_d^{v_{j+1} \rightarrow e_j}} &= \frac{1}{Z^{e_j \rightarrow v_j}} \sum_{\mathbf{c}_{e_j} | \mathbf{c}_{e_j}(v_j, v_{j+1}) = (c, d)} \phi_i(\mathbf{c}_{e_j}) \prod_{w \in e_j \setminus \{v_j, v_{j+1}\}} m_{\mathbf{c}_{e_j}(w)}^{w \rightarrow e_j} \\ &\quad - \frac{m_c^{e_j \rightarrow v_j}}{Z^{e_j \rightarrow v_j}} \sum_{c' \in C} \sum_{\mathbf{c}_{e_j} | \mathbf{c}_{e_j}(v_j, v_{j+1}) = (c', d)} \phi_i(\mathbf{c}_{e_j}) \prod_{w \in e_j \setminus \{v_j, v_{j+1}\}} m_{\mathbf{c}_{e_j}(w)}^{w \rightarrow e_j} \\ &= \frac{1}{Z^{e_j \rightarrow v_j}} \cdot \frac{1}{m_d^{v_{j+1} \rightarrow e_j}} \sum_{\mathbf{c}_{e_j} | \mathbf{c}_{e_j}(v_j, v_{j+1}) = (c, d)} \phi_i(\mathbf{c}_{e_j}) \prod_{w \in e_j \setminus v_j} m_{\mathbf{c}_{e_j}(w)}^{w \rightarrow e_j} \\ &\quad - \frac{m_c^{e_j \rightarrow v_j}}{Z^{e_j \rightarrow v_j}} \cdot \frac{1}{m_d^{v_{j+1} \rightarrow e_j}} \sum_{c' \in C} \sum_{\mathbf{c}_{e_j} | \mathbf{c}_{e_j}(v_j, v_{j+1}) = (c', d)} \phi_i(\mathbf{c}_{e_j}) \prod_{w \in e_j \setminus v_j} m_{\mathbf{c}_{e_j}(w)}^{w \rightarrow e_j} \end{aligned}$$

Recall we defined a distribution μ_i over \mathbf{c}_{e_j} and stochastic matrices $\Psi_{\theta(e_j), i(v_{j+1}) | i(v_j)}$ before. Evaluating the derivative at the factorized fixed point gives the transformation matrix

$$\begin{aligned} \frac{\partial \Gamma(m)^{e_j \rightarrow v_j}}{\partial m^{v_{j+1} \rightarrow e_j}} \Big|_m &= \frac{1}{\mathbb{P}_{\tau(v_{j+1})}(d) \cdot \text{support}(\mathbb{P}_{\tau(v_j)})} \left(\Pr_{\mathbf{c}_{e_j} \sim D_\theta}[\mathbf{c}_{e_j}(v_{j+1}) = d \mid \mathbf{c}_{e_j}(v_j) = c] \right. \\ &\quad \left. - \frac{1}{\text{support}(\mathbb{P}_{\tau(v_j)})} \sum_{c' \in C} \Pr_{\mathbf{c}_{e_j} \sim D_\theta}[\mathbf{c}_{e_j}(v_{j+1}) = d \mid \mathbf{c}_{e_j}(v_j) = c'] \right), \end{aligned}$$

Then $\frac{\partial m^{e_j \rightarrow v_j}}{\partial m^{v_{j+1} \rightarrow e_j}} \Big|_{(m)^{\text{fp}}} = \frac{1}{\text{support}(\mathbb{P}_{\tau(v_j)})} \left(\mathbf{I} - \frac{1}{\text{support}(\mathbb{P}_{\tau(v_j)})} \mathbf{1}\mathbf{1}^\top \right) \Psi_{\theta(e_j), i(v_{j+1}) | i(v_j)}^\top \mathbf{D}_{\tau(v_{j+1})}^\dagger$ where the \dagger in the superscript denotes the pseudoinverse of the matrix.

Then we can write the transformation matrix for the step $v_{j+1} \xrightarrow{e_j} e_{j-1}$ as

$$\begin{aligned} \bar{\mathbf{M}}_{\theta(e_j), i(v_j) | i(v_{j+1})} &= \left(\mathbf{D}_{\tau(v_j)} - \mathbb{P}_{\tau(v_j)} \mathbb{P}_{\tau(v_j)}^\top \right) \left(\mathbf{I} - \frac{1}{\text{support}(\mathbb{P}_{\tau(v_j)})} \mathbf{1}\mathbf{1}^\top \right) \Psi_{\theta(e_j), i(v_{j+1}) | i(v_j)}^\top \mathbf{D}_{\tau(v_{j+1})}^\dagger \\ &= \mathbf{D}_{\tau(v_j)} \left(\mathbf{I} - \mathbf{1} \mathbb{P}_{\tau(v_j)}^\top \right) \left(\mathbf{I} - \frac{1}{\text{support}(\mathbb{P}_{\tau(v_j)})} \mathbf{1}\mathbf{1}^\top \right) \Psi_{\theta(e_j), i(v_{j+1}) | i(v_j)}^\top \mathbf{D}_{\tau(v_{j+1})}^\dagger \\ &= \mathbf{D}_{\tau(v_j)} \left(\mathbf{I} - \mathbf{1} \mathbb{P}_{\tau(v_j)}^\top \right) \Psi_{\theta(e_j), i(v_{j+1}) | i(v_j)}^\top \mathbf{D}_{\tau(v_{j+1})}^\dagger \\ &= \left(\mathbf{I} - \mathbb{P}_{\tau(v_j)} \mathbf{1}^\top \right) \mathbf{D}_{\tau(v_j)} \Psi_{\theta(e_j), i(v_{j+1}) | i(v_j)}^\top \mathbf{D}_{\tau(v_{j+1})}^\dagger \\ &= \left(\mathbf{I} - \mathbb{P}_{\tau(v_j)} \mathbf{1}^\top \right) \Psi_{\theta(e_j), i(v_j) | i(v_{j+1})}. \end{aligned}$$

This establishes the first part of [Claim 4.1](#). To establish the second part, consider the following chain of equalities where the first equality is one we know from the above chain.

$$\begin{aligned} \bar{\mathbf{M}}_{\theta(e_j), i(v_j) | i(v_{j+1})} &= \left(\mathbf{I} - \mathbb{P}_{\tau(v_j)} \mathbf{1}^\top \right) \mathbf{D}_{\tau(v_j)} \Psi_{\theta(e_j), i(v_{j+1}) | i(v_j)}^\top \mathbf{D}_{\tau(v_{j+1})}^\dagger \\ &= \mathbf{D}_{\tau(v_j)} \left(\mathbf{I} - \mathbf{1} \mathbb{P}_{\tau(v_j)}^\top \right) \Psi_{\theta(e_j), i(v_{j+1}) | i(v_j)}^\top \mathbf{D}_{\tau(v_{j+1})}^\dagger \\ &= \mathbf{D}_{\tau(v_j)} \left(\Psi_{\theta(e_j), i(v_{j+1}) | i(v_j)} \left(\mathbf{I} - \mathbb{P}_{\tau(v_j)} \mathbf{1}^\top \right) \right)^\top \mathbf{D}_{\tau(v_{j+1})}^\dagger \\ &= \mathbf{D}_{\tau(v_j)} \left(\Psi_{\theta(e_j), i(v_{j+1}) | i(v_j)} - \mathbb{P}_{\tau(v_{j+1})} \mathbf{1}^\top \right)^\top \mathbf{D}_{\tau(v_{j+1})}^\dagger \\ &= \mathbf{D}_{\tau(v_j)} \left(\left(\mathbf{I} - C_{\tau(v_{j+1})} \mathbf{1}^\top \right) \Psi_{\theta(e_j), i(v_{j+1}) | i(v_j)} \right)^\top \mathbf{D}_{\tau(v_{j+1})}^\dagger \\ &= \mathbf{D}_{\tau(v_j)} \bar{\mathbf{M}}_{\theta(e_j), i(v_{j+1}) | i(v_j)}^\top \mathbf{D}_{\tau(v_{j+1})}^\dagger \end{aligned}$$

If $\mathbb{P}_{\tau(v_j)}(c) = 0$, then the c -th column of $\bar{\mathbf{M}}_{\theta(e_j), i(v_{j+1}) | i(v_j)}$ and the c -th row of $\bar{\mathbf{M}}_{\theta(e_j), i(v_j) | i(v_{j+1})}$ are 0, and hence the second part of [Claim 4.1](#) follows as well.

D Random graph lemmas

D.1 Proof of [Lemma 7.8](#)

Let M^\times be a null model, and let $H \sim M^\times$. We use a_{\max} to denote $\max_{i \in [F]} a(i)$ and $\bar{\phi}_{\max}$ to denote $\max_{i \in [F]} \bar{\phi}_i$.

Define the notion of (τ, Cl) -consistent:

Definition D.1. Given τ and Cl , we say a subgraph Γ of a bipartite factor graph with right vertex set $R(\Gamma)$ is (τ, Cl) -consistent if every $\gamma = (v_1, \dots, v_{a(i)}) \in R(\Gamma)$ satisfies $(\tau(v_1), \dots, \tau(v_{a(i)})) = \text{Cl}(\theta(\gamma))$.

It is easy to see the following.

Observation D.2. Suppose Γ is a subgraph of $\text{Bip}(\mathcal{K}_n)$. Then the probability that Γ is a subgraph of \mathbf{H} is equal to $\prod_{\gamma \in R(\Gamma)} \frac{\overline{\phi}_{\theta(\gamma)}}{n^{a(\theta(\gamma))-1}} \mathbf{1}[\Gamma \text{ is } (\tau, \text{Cl})\text{-consistent}]$.

Definition D.3 (Partially labeled graph). A *partially labeled graph* $\Gamma = (L, R, P, p, E)$ is given by a left vertex set L , a right vertex set R , a distinguished set of left vertices P along with an injective labeling of the distinguished vertices $p : P \rightarrow [n]$, and edge set E .

Definition D.4 (Occurrence of partially labeled graph). An *occurrence* of a partially labeled graph $\Gamma = (L, R, P, p, E)$ in $\text{Bip}(\mathbf{H})$ is a pair of injective functions $f_L : L \rightarrow L(\mathbf{H})$ and $f_R : R \rightarrow R(\mathbf{H})$ such that for all $v \in P$ satisfies $f_L(v) = p(v)$, and if $\{u, v\} \in E$, then $\{f_L(u), f_R(v)\} \in E(\text{Bip}(\mathbf{H}))$.

Given a partially labeled subgraph Γ we are interested in bounding the expected number of occurrences of Γ in $\text{Bip}(\mathbf{H})$.

Lemma D.5. *Given partially labeled graph $\Gamma = (L, R, P, p)$ with no isolated right vertices, the expected number of occurrences of Γ in $\text{Bip}(\mathbf{H})$ is at most*

$$n^{|L|+|R|-|P|-|E|} (F a_{\max} \overline{\phi}_{\max})^{|E|}.$$

Proof. There are at most $n^{|L|-|P|}$ choices for f_L . For each potential choice of f_R , we can associate $t_{f_R} : R \rightarrow [F]$ such that $t_{f_R}(r)$ is the type of $f_R(r)$. There are at most $F^{|R|}$ possible values for t_{f_R} . For each fixed choice of f_L and t , we wish to bound the expected number of f_R such that (f_L, f_R) is an occurrence and $t_{f_R} = t$. The number of such potential f_R is bounded by

$$\prod_{i \in [R]} a(t(i))^{\deg_{\Gamma}(i)} n^{a(t(i)) - \deg_{\Gamma}(i)}$$

and the probability that (f_L, f_R) is a valid occurrence for a given such f_R is at most

$$\prod_{i \in [R]} \frac{\overline{\phi}_i}{n^{a(t(i))-1}},$$

which gives us a bound of

$$\prod_{i \in [R]} \left(\frac{a(t(i))}{n} \right)^{\deg_{\Gamma}(i)} \cdot n \cdot \overline{\phi}_i \leq \left(\frac{a_{\max}}{n} \right)^{|E|} n^{|R|} \overline{\phi}_{\max}^{|R|}.$$

Combining this with the bound on total number of f_L and t_{f_R} gives us a bound of:

$$n^{|L|-|P|} F^{|R|} \left(\frac{a_{\max}}{n} \right)^{|E|} n^{|R|} \overline{\phi}_{\max}^{|R|} = n^{|L|+|R|-|P|-|E|} F^{|R|} a_{\max}^{|E|} \overline{\phi}_{\max}^{|R|}.$$

Since there are no isolated vertices, $|E| \geq |R|$ and hence the above is at most $n^{|L|+|R|-|P|-|E|} (F a_{\max} \overline{\phi}_{\max})^{|E|}$. \square

Definition D.6. For a graph Γ and a subset of its vertices S we use $B_{\Gamma}(S, r)$ to denote the radius- r ball around set S within Γ . We also abuse notation and use $B_{\mathbf{H}}(S, r)$ to mean $B_{\text{Bip}(\mathbf{H})}(S, r)$.

Lemma D.7. *Given a set of vertices S in $\text{Bip}(\mathbf{H})$, the probability that $|E(B_{\mathbf{H}}(S, r))| - |V(B_{\mathbf{H}}(S, r))| + |S| \geq t$ is at most $\left(\frac{(F a_{\max} \overline{\phi}_{\max})^{2(r+1)} (36t^3 r^2)^{|S|}}{n} \right)^t$.*

In preparation to prove [Lemma D.7](#) we will need the following statement about counts of trees with a bounded number of leaves. The statement along with a proof can be found in [[BMR19](#), Lemma 6.33].

Lemma D.8. *The number of nonisomorphic trees on v vertices and L leaves is bounded by $(4Lv)^{2L+1}$.*

Proof of Lemma D.7. Let us call a partially labeled subgraph $\Gamma = (L, R, P, p, E)$ a *candidate witness* if

- $p(P) = S$,
- Γ can be expressed as $F \cup B$ where F is a forest and $B = \{\{u_1, v_1\}, \dots, \{u_t, v_t\}\}$ is a set of t additional edges,
- F has $|P|$ connected components where each connected component contains exactly one $v \in P$ and has depth r when rooted at v .

If $|E(B_H(S, r))| - |V(B_H(S, r))| + |S| \geq t$, then there must be an occurrence of some candidate witness $\Gamma = (L, R, P, p, E)$ within H . We will first find a “simple” subgraph of $\Gamma = F \cup B$, which we call the *trim* of Γ . First let us augment F to \tilde{F} by adding a single vertex w and connecting it to all vertices in P – note that \tilde{F} is a tree. Now let $\text{Trim}(\tilde{F})$ be the tree obtained by only choosing vertices that lie on paths from vertices in $L := \{u_1, \dots, u_t, v_1, \dots, v_t\}$ to w . Since the depth of $\text{Trim}(\tilde{F})$ is $r + 1$ and has at most $2t$ leaves when rooted at w , the number of vertices in $\text{Trim}(\tilde{F})$ is at most $2tr + 1$. Let $\text{Trim}(\Gamma) = (L', R', P, p, E')$ be the graph obtained by deleting w from $\text{Trim}(\tilde{F})$, adding edges $\{u_1, v_1\}, \dots, \{u_t, v_t\}$, and adding vertices in $P \setminus V(\text{Trim}(\tilde{F}))$. Since $\text{Trim}(\Gamma)$ is a subgraph of Γ there must be an occurrence of $\text{Trim}(\Gamma)$ in H .

$\text{Trim}(\Gamma)$ has at most $2tr$ vertices and $|E(\text{Trim}(\Gamma))| - |L(\text{Trim}(\Gamma))| - |R(\text{Trim}(\Gamma))| + |P| \geq t$. Thus, from [Lemma D.5](#) the probability that $\text{Trim}(\Gamma)$ occurs in H is bounded by $\left(\frac{(Fa_{\max}\bar{\phi}_{\max})^{2(r+1)}}{n}\right)^t$. Thus:

$$\begin{aligned} \Pr[|E(B_H(S, r))| - |V(B_H(S, r))| + |S| \geq t] &\leq \Pr[\text{there is a candidate witness } \Gamma \text{ in } H] \\ &\leq \Pr[\text{there is a trim of a candidate witness } \text{Trim}(\Gamma) \text{ in } H] \\ &\leq \sum_{\Gamma' \text{ trim of a candidate witness}} \Pr[\Gamma' \text{ in } H] \\ &\leq \sum_{\Gamma' \text{ trim of a candidate witness}} \left(\frac{(Fa_{\max}\bar{\phi}_{\max})^{2(r+1)}}{n}\right)^t. \end{aligned} \quad (46)$$

Next we bound the number of terms in the above summation. Since each Γ' in the above sum can be specified by taking a tree on at most $2tr + 1$ vertices and at most $2t + 1$ leaves, deleting one vertex, and labeling each neighbor of this deleted vertex with an element of P , from [Lemma D.8](#) and the fact the maximum degree in a tree is bounded by the number of leaves the number of terms is at most:

$$(4(2t + 1)(2tr + 1))^{4t+3} \cdot (2tr + 1) \cdot |P|^{2t+1} \leq ((36t^3 r^2)^5 |P|)^t.$$

Plugging this into (46) and using $|S| = |P|$ gives:

$$\Pr[|E(B_H(S, r))| - |V(B_H(S, r))| + |S| \geq t] \leq \left(\frac{(Fa_{\max}\bar{\phi}_{\max})^{2(r+1)}(36t^3 r^2)^5 |S|}{n}\right)^t$$

□

Corollary D.9. *With probability $1 - o_n(1)$, $\text{Bip}(H)$ is r -bicycle free for $r = \frac{\log n}{\log \log n}$.*

Proof. This is a simple consequence of [Lemma D.7](#). Indeed, by [Lemma D.7](#) the probability that the radius- $(r + 1)$ neighborhood of a single vertex $v \in [n]$ contains more than one cycle is at most $\frac{1}{n^{2-o_n(1)}}$, and hence by a union bound over all vertices the probability of any left vertex containing more than one cycle in its radius- $r + 1$ neighborhood is bounded by $\frac{1}{n^{1-o_n(1)}}$. Since every right vertex is incident to a left vertex, the statement we wish to prove follows. \square

D.2 Proof of [Lemma D.11](#)

We will need the following combinatorial lemma that appears in [[FM17](#), Lemma A.2].

Lemma D.10. *If e distinct edges of a graph Γ belong to a r -bicycle frame, then $\text{Exc}(\Gamma) \geq \frac{e}{r}$.*

Our proof of the statement below follows the same strategy as the proof of a similar statement appearing in [[FM17](#)].

Lemma D.11. *Suppose S and L are disjoint sets of right vertices of \mathcal{K}_n of size at most $\log^2 n$, $\mathbf{1}_\gamma$ is the indicator random variable for whether γ is in \mathbf{H} , μ_γ is the probability that γ is in \mathbf{H} , and \mathcal{E} denotes the event that \mathbf{H} is r -bicycle free for $r = \frac{\log n}{\log \log n}$. Then:*

$$\left| \mathbf{E} \left[\prod_{\gamma \in S} (\mathbf{1}_\gamma - \mu_\gamma) \prod_{\gamma \in L} \mathbf{1}_\gamma \mathbf{1}[\mathcal{E}] \right] \right| \leq \prod_{\gamma \in S \cup L} \mu_\gamma \cdot 2^{|S|} \left(\frac{1}{n^5} \right)^{\frac{|S|}{r} - \text{Exc}(\text{Clos}(S \cup L))}.$$

Proof.

$$\begin{aligned} \left| \mathbf{E} \left[\prod_{\gamma \in S} (\mathbf{1}_\gamma - \mu_\gamma) \prod_{\gamma \in L} \mathbf{1}_\gamma \mathbf{1}[\mathcal{E}] \right] \right| &= \left| \sum_{J \subseteq S} \mathbf{E} \left[\prod_{\gamma \in J} \mathbf{1}_\gamma \cdot \prod_{\gamma \in S \setminus J} (-\mu_\gamma) \cdot \prod_{\gamma \in L} \mathbf{1}_\gamma \mathbf{1}[\mathcal{E}] \right] \right| \\ &= \left| \sum_{J \subseteq S} (-1)^{|S| - |J|} \prod_{\gamma \in S \setminus J} \mu_\gamma \mathbf{E} \left[\prod_{\gamma \in J \cup L} \mathbf{1}_\gamma \mathbf{1}[\mathcal{E}] \right] \right| \\ &= \left| \sum_{J \subseteq S} (-1)^{|J|} \prod_{\gamma \in S \cup L} \mu_\gamma \Pr[\mathcal{E} | \gamma \in \mathbf{H} \forall \gamma \in J \cup L] \right| \\ &= \prod_{\gamma \in S \cup L} \mu_\gamma \left| \sum_{J \subseteq S} (-1)^{|J|} \Pr[\mathcal{E} | \gamma \in \mathbf{H} \forall \gamma \in J \cup L] \right| \end{aligned} \quad (47)$$

Now we focus our attention on understanding the quantity $\left| \sum_{J \subseteq S} (-1)^{|J|} \Pr[\mathcal{E} | \gamma \in \mathbf{H} \forall \gamma \in J \cup L] \right|$. Let \mathbf{g}_0 be $\mathbf{H} \setminus (S \cup L)$.

$$\left| \sum_{J \subseteq S} (-1)^{|J|} \Pr[\mathcal{E} | \gamma \in \mathbf{H} \forall \gamma \in J \cup L] \right| = \left| \mathbf{E}_{\mathbf{g}_0} \sum_{J \subseteq S} (-1)^{|J|} \Pr[\mathcal{E} | \gamma \in \mathbf{H} \forall \gamma \in J \cup L, \mathbf{g}_0] \right| \quad (48)$$

For $K \subseteq S$, define $f_{\mathbf{g}_0}(K)$ as 1 if $\text{Clos}(\mathbf{g}_0 \cup K \cup L)$ has no r -bicycles and 0 otherwise. Suppose there is $s \in S$ that $f_{\mathbf{g}_0}$ does not depend on s – that is, for any $K \subseteq S$, $f_{\mathbf{g}_0}(K) = f_{\mathbf{g}_0}(K \Delta \{s\})$, then for every J which contains s :

$$\Pr[\mathcal{E} | \gamma \in \mathbf{H} \forall \gamma \in J \cup L, \mathbf{g}_0] = \Pr[\mathcal{E} | \gamma \in \mathbf{H} \forall \gamma \in J \cup L \setminus \{s\}, \mathbf{g}_0].$$

This means (48) is equal to:

$$(48) = \left| \mathbf{E}_{\mathbf{g}_0} \mathbf{1}[f_{\mathbf{g}_0} \text{ depends on every } s \in S] \sum_{J \subseteq S} (-1)^{|J|} \Pr[\mathcal{E} | \gamma \in \mathbf{H} \forall \gamma \in J \cup L, \mathbf{g}_0] \right|$$

$$\leq 2^{|S|} \cdot \Pr_{g_0}[f_{g_0} \text{ depends on every } s \in S]. \quad (49)$$

Let E_S be the set of all edges incident to S . If f_{g_0} depends on every $s \in S$, then the function h_{g_0} defined on subsets of E_S which is 1 on input $K \subseteq E_S$ if $\text{Clos}(g_0 \cup L) \cup K$ has no r -bicycles depends on at least $|S|$ edges in E_S . That means:

$$\begin{aligned} (49) &\leq 2^{|S|} \cdot \Pr_{g_0}[h_{g_0} \text{ depends on at least } |S| \text{ edges in } E_S] \\ &\leq 2^{|S|} \cdot \Pr_{g_0}[\text{At least } |S| \text{ edges in } E_S \text{ part of } r\text{-bicycle frame in } \text{Clos}(g_0 \cup S \cup L)] \end{aligned}$$

which, via [Lemma D.10](#), can be bounded by

$$\begin{aligned} &\leq 2^{|S|} \cdot \Pr_{g_0} \left[\text{Exc}(B_{\text{Clos}(g_0 \cup S \cup L)}(\text{Clos}(S), r)) \geq \frac{|S|}{r} \right] \\ &\leq 2^{|S|} \cdot \Pr_{g_0} \left[\text{Exc}(B_{\text{Clos}(g_0)}(\text{Clos}(S \cup L), r)) \geq \frac{|S|}{r} - \text{Exc}(\text{Clos}(S \cup L)) \right] \\ &\leq 2^{|S|} \cdot \Pr_H \left[\text{Exc}(B_H(\text{Clos}(S \cup L), r)) \geq \frac{|S|}{r} - \text{Exc}(\text{Clos}(S \cup L)) \right] \end{aligned}$$

By [Lemma D.7](#), the bounds on size of $|S|$ and $|L|$, and the value of r , we can conclude that the above is at most:

$$\begin{aligned} &\leq 2^{|S|} \min \left\{ \left(\frac{1}{n^5} \right)^{\frac{|S|}{r} - \text{Exc}(\text{Clos}(S \cup L))}, 1 \right\} \\ &\leq 2^{|S|} \left(\frac{1}{n^5} \right)^{\frac{|S|}{r} - \text{Exc}(\text{Clos}(S \cup L))}. \end{aligned}$$

Plugging this back into [\(47\)](#) gives us the desired statement. \square

D.3 Proof of [Claim 7.22](#)

Proof. Consider any $\text{Sh} \in \mathcal{U}_{s,x,\Delta}$. Recall that $S(\text{Sh})$ is the set of singleton vertices in $R(\text{Sh})$ and $D(\text{Sh})$ the set of duplicative vertices in $R(\text{Sh})$. $D^*(\text{Sh})$ is the maximum weight subset of $D(W)$ that makes Sh r -bicycle free, and $\Delta(\text{Sh}) = w(D(\text{Sh})) - w(D^*(\text{Sh}))$. Thus we deduce that $\Delta(\text{Sh}) \geq |D(\text{Sh}) \setminus D^*(\text{Sh})|$. We apply the following procedures to the walk Sh .

1. Break the walk Sh into $\leq s + \Delta + |D(\text{Sh}) \setminus D^*(\text{Sh})| + k$ segments by first removing the vertices $S(\text{Sh})$ and $D(\text{Sh}) \setminus D^*(\text{Sh})$ from Sh and second breaking the remaining segments at endpoints of the k links in Sh . Denote the new union of walks Sh_1 .
2. Since $S(\text{Sh})$ and $D(\text{Sh}) \setminus D^*(\text{Sh})$ are removed from Sh , Sh_1 is singleton free, and the graph on Sh_1 , denoted by $G(\text{Sh}_1)$, is r -bicycle free.
3. We contract the graph $G(\text{Sh}_1)$ by merging all adjacent edges that share a degree-2 vertex. We denote the resulting graph $G(\text{Sh}_1)_c$, and we note that the vertices left in $G(\text{Sh}_1)_c$ are those with degree ≥ 3 in $G(\text{Sh}_1)$.

We make the following observations on the size of $G(\text{Sh}_1)$ and $G(\text{Sh}_1)_c$.

The number of vertices in $G(\text{Sh}_1)$ is $|V(\text{Sh})| - |S(\text{Sh})| - |D(\text{Sh}) \setminus D^*(\text{Sh})|$. The number of edges in $G(\text{Sh}_1)$ is $\leq |E(\text{Sh})| - 2|S(\text{Sh})| - 2|D(\text{Sh}) \setminus D^*(\text{Sh})|$.

To bound the number of vertices in $G(\text{Sh}_1)_c$ is we apply the following lemma from [\[MOP20\]](#).

Lemma D.12 (Lemma 6.18 in [MOP20]). *Let C be a $(k, 2\ell)$ -nonbacktracking, internally 2ℓ -bicycle-free linkage. Assume $\log k\ell = o(\ell)$. Then $G(C)$ has at most $O(k \log k\ell)$ vertices of degree exceeding 2.*

Applying the lemma to the walk Sh_1 , we obtain that the number of degree ≥ 3 vertices in Sh_1 is $O(k \log k\ell)$. Thus the number of vertices in $G(\text{Sh}_1)_c$ is $O(k \log k\ell)$. The number of edges in $G(\text{Sh}_1)_c$ is

$$\begin{aligned} & |E(G(\text{Sh}_1))| - (|V(G(\text{Sh}_1))| - |V(G(\text{Sh}_1)_c)|) \\ & \leq |E(\text{Sh})| - 2|S(\text{Sh})| - 2|D(\text{Sh}) \setminus D^*(\text{Sh})| - |V(\text{Sh})| + |S(\text{Sh})| + |D(\text{Sh}) \setminus D^*(\text{Sh})| + O(k \log k\ell) \\ & = (|E(\text{Sh})| - |V(\text{Sh})|) - |S(\text{Sh})| - |D(\text{Sh}) \setminus D^*(\text{Sh})| + O(k \log k\ell) \\ & \leq x + O(k \log k\ell). \end{aligned}$$

Now to count the number of distinct $\text{Sh} \in \mathcal{U}_{s,x,\Delta}$, it suffices to count 1. the number of distinct sets of breaking points ($S(\text{Sh}), D(\text{Sh}) \setminus D^*(\text{Sh})$, and k link endpoints), 2. the number of distinct graphs $G(\text{Sh}_1)$, 3. given the breaking points and $G(\text{Sh}_1)$, the number of distinct walk segments in $G(\text{Sh}_1)$ with those breaking points. We count each of the three quantities separately and multiply them together to obtain an upper bound on $|\mathcal{U}_{s,x,\Delta}|$.

The number of distinct sets of breaking points: these breaking points breaks Sh into at most $s + 2\Delta + k$ segments. So there are $(2k\ell)^{s+2\Delta+k}$ ways to choose these breaking points.

The number of distinct graphs $G(\text{Sh}_1)$: $G(\text{Sh}_1)$ can be contracted to a graph $G(\text{Sh}_1)_c$ on $O(k \log k\ell)$ vertices and $x + O(k \log k\ell)$ edges. Each edge in $G(\text{Sh}_1)_c$ represents a length- $\leq 2k\ell$ simple path in $G(\text{Sh}_1)$. Thus there are $O(k \log k\ell)^{2(x+O(k \log k\ell))} \cdot (2k\ell)^{x+O(k \log k\ell)}$ distinct graphs $G(\text{Sh}_1)$.

Given the breaking points and $G(\text{Sh}_1)$, the number of distinct walk segments in $G(\text{Sh}_1)$ with those breaking points: since $G(\text{Sh}_1)$ is r -bicycle free with $r \geq 2\ell$ and each segment is of length $\leq 2\ell$, there are only 2 distinct length $\leq 2\ell$ walk between any two vertices in $G(\text{Sh}_1)$. Thus the number of distinct walk segments are $2^{s+2\Delta+k}$.

Combine the three bound together we obtain that

$$|\mathcal{U}_{s,x,\Delta}| \leq (2k\ell)^{s+2\Delta+x+O(k \log k\ell)} O(k \log k\ell)^{2(x+O(k \log k\ell))} 2^{s+2\Delta+k}.$$

From this bound we quickly derive that

$$\begin{aligned} & \sum_{\text{Sh} \in \mathcal{U}_{s,x,\Delta}} \left(\frac{\beta^r}{n^5}\right)^{\frac{|S(\text{Sh})|}{r}} \left(\frac{\beta}{n^5}\right)^{\text{Exc}_{\text{Sh}}} \left(\frac{\beta}{n}\right)^{\Delta(\text{Sh})} \\ & = |\mathcal{U}_{s,x,\Delta}| \cdot \left(\frac{\beta^r}{n^5}\right)^{\frac{s}{r}} \left(\frac{\beta}{n^5}\right)^x \left(\frac{\beta}{n}\right)^\Delta \\ & = \left(\frac{\beta^r \cdot (4k\ell)^r}{n^5}\right)^{\frac{s}{r}} \left(\frac{\beta \cdot 2k\ell \cdot O(k \log k\ell)^2}{n^5}\right)^x \left(\frac{4\beta \cdot (2k\ell)^2}{n}\right)^\Delta 2^k (2k\ell \cdot O(k \log k\ell)^2)^{O(k \log k\ell)} \\ & \leq \left(\frac{\beta^r \cdot (4k\ell)^r}{n^5}\right)^{\frac{s}{r}} \left(\frac{\beta \cdot 2(k\ell)^3}{n^5}\right)^x \left(\frac{4\beta \cdot (2k\ell)^2}{n}\right)^\Delta (2k\ell)^{O(k \log k\ell)} \end{aligned}$$

The last inequality follows since we pick ℓ such that $\log k\ell = o(\ell)$. □