

Local Statistics, Semidefinite Programming, and Community Detection

Jess Banks^{*}

Sidhanth Mohanty[†]

Prasad Raghavendra[‡]

November 4, 2019

Abstract

We propose a new hierarchy of semidefinite programming relaxations for inference problems, inspired by recent ideas of ‘pseudocalibration’ in the Sum-of-Squares literature. As a test case, we consider the problem of community detection in a distribution of random regular graphs we’ll call the Degree Regular Block Model, wherein the vertices are partitioned into k communities, and a graph is sampled conditional on a prescribed number of inter- and intra-community edges. The problem of *detection*, where we are to decide with high probability whether a graph was drawn from this model or the uniform distribution on regular graphs, is conjectured to undergo a computational phase transition at a point called the Kesten-Stigum (KS) threshold, and we show (i) that sufficiently high constant levels of our hierarchy can perform detection arbitrarily close to this point, (ii) that our algorithm is robust to $o(n)$ adversarial edge perturbations, and (iii) that below Kesten-Stigum no level constant level can do so.

In the more-studied case of the (irregular) Stochastic Block Model, it is known that efficient algorithms exist all the way down to this threshold, although none are robust to adversarial perturbations of the graph when the average degree is small. More importantly, there is little complexity-theoretic evidence that detection is hard below Kesten-Stigum. In the DRBM with more than two groups, it has not to our knowledge been proven that any algorithm succeeds down to the KS threshold, let alone that one can do so robustly, and there is a similar dearth of evidence for hardness below this point.

Our SDP hierarchy is highly general and applicable to a wide range of hypothesis testing problems.

^{*}University of California, Berkeley

[†]University of California, Berkeley

[‡]University of California, Berkeley

1 Introduction

Community detection in graphs is a canonical and widely applicable problem in computer science and machine learning. The setup is both simple and flexible: we are shown a graph and asked for a coarse-grained description in the form of a partition of the vertices into ‘communities’ with atypically many internal edges. The literature contains innumerable algorithms and approaches for this task, but perhaps the most fruitful has been a Bayesian perspective wherein we treat the graph as the output of some generative model, whose unknown parameters we attempt to estimate. In other words, we assume that there are some true and hidden community labels, and that the graph has been drawn probabilistically in a way that respects this ‘planted’ structure.

Much of the existing literature on community detection concerns the *stochastic block model (SBM)*. For now let us discuss the symmetric setting where we first partition n vertices in to k groups, and include each edge independently and with probability p_{in} or p_{out} depending on whether or not the labels of its endpoints coincide. Research in this area spans several decades, and it will not be fruitful to attempt a thorough review of the literature here; we refer the reader to [Abb17] for a survey. Most salient to us, however, is a rich theory of computational threshold phenomena which has emerged out of the past several years of collaboration between computer scientists, statisticians, and statistical physicists.

The key computational tasks associated with the SBM are *recovery* and *detection*: we attempt either to reconstruct the planted communities from the graph, or to decide whether a graph was drawn from the planted model or the Erdős-Rényi model with the same average degree. A set of fascinating conjectures were posed in Decelle et al. [DKMZ11], regarding these tasks in the case of ‘sparse’ models where $p_{\text{in}}, p_{\text{out}} = O(1/n)$ and the average degree is $O(1)$ as the number of vertices diverges.

It is typical to parametrize the symmetric SBM in terms of k , the average degree

$$d = \frac{np_{\text{in}} + (k-1)np_{\text{out}}}{k},$$

and a ‘signal-to-noise ratio’

$$\lambda \triangleq \frac{np_{\text{in}} - np_{\text{out}}}{kd}.$$

In this setup, it is believed that as we hold k and λ constant, then there is an *information-theoretic threshold* $d_{\text{IT}} \approx \frac{\log k}{k\lambda^2}$, in the sense that when $d < d_{\text{IT}}$ both detection and recovery are impossible for any algorithm. Moreover, Decelle et al. conjecture that efficient algorithms for both tasks exist only when the degree is larger than a point known as the *Kesten-Stigum threshold* $d_{\text{KS}} = \lambda^{-2}$. Much of this picture is now rigorous [MNS18, Mas14, BLM15, ABH16]. Still, fundamental questions remain unanswered. What evidence can we furnish that detection and recovery are indeed intractable in the so-called ‘hard regime’ $d_{\text{IT}} < d < d_{\text{KS}}$? How robust are these thresholds to adversarial noise or small deviations from the model?

Zooming out, this discrepancy between information-theoretic and computational thresholds is conjectured to be quite universal among planted problems, where we are to reconstruct or detect a structured, high-dimensional signal observed through a noisy channel [citations]. The purpose behind our work is to begin developing a framework capable of providing evidence for average case computational intractability in such settings. To illustrate this broader motivation, consider a different average-case problem also conjectured to be computationally intractable: refutation of random 3-SAT. A random instance of 3-SAT with n literals and, say $m = 1000n$ clauses is unsatisfiable with high probability. However, it is widely conjectured that the problem of *certifying* that a given random 3-SAT instance is unsatisfiable is computationally intractable (all the way up to $n^{3/2}$ clauses) [Fei02]. While proving intractability remains out of

reach, the complexity theoretic literature now contains ample evidence in support of this conjecture. Most prominently, exponential lower bounds are known for the problem in restricted computational models such as linear and semidefinite programs [Gri01] and resolution based proofs [BSW01]. Within the context of combinatorial optimization, the Sum-of-Squares (SoS) SDPs yield a hierarchy of successively more powerful and complex algorithms which capture and unify many other known approaches. A lower bound against the SoS SDP hierarchy such as [Gri01] provides strong evidence that this refutation problem is computationally intractable. This paper is a step towards developing a similar framework to reason about the computational complexity of detection and recovery in stochastic block models specifically, and planted problems generally.

A second motivation is the issue of robustness of computational thresholds under adversarial perturbations of the graph. Spectral algorithms based on non-backtracking walk matrix [BLM15] achieve weak-detection as soon as $d > d_{KS}$, but are not robust in this sense. Conversely, robust algorithms for recovery are known, but only when the edge-densities are significantly higher than Kesten-Stigum [GV16, MMV16, CSV17, SVC16]. The positive result that gets closest to robustly achieving the conjectured computational phase transition at d_{KS} is the work of Montanari and Sen [MS15] who observe that their SDP-based algorithm for testing whether the input graph comes from the Erdős-Rényi distribution or a Stochastic Block Model with $k = 2$ communities also works in presence of $o(|E|)$ edge outlier errors. On the negative side, Moitra et al. [Moi12] consider the problem of weak recovery in a SBM with two communities and $p_{in} > p_{out}$ in the presence of *monotone errors* that add edges within communities and delete edges between them. Their main result is a statistical lower bound indicating the phase transition for weak recovery changes in the presence of monotone errors. This still leaves open the question of whether there exist algorithms that weakly recover right at the threshold and are robust to $o(|E|)$ perturbations in the graph.

2 Main Results

We define a new hierarchy of semidefinite programming relaxations for inference problems that we refer to as the *Local Statistics* hierarchy, denoted $\text{LoSt}(D_1, D_2)$ and indexed by parameters $D_1, D_2 \in \mathbb{N}$. This family of SDPs is inspired by the technique of pseudocalibration in proving lower bounds for sum-of-squares (SoS) relaxations, as well as subsequent work of Hopkins and Steurer [HS17] extending it to an SoS SDP based approach to inference problems. The LoSt hierarchy can be defined for a broad range of inference problems involving a joint distribution μ on an observation and hidden parameter.

As a test case, we apply our SDP relaxations to community detection in the *Degree Regular Block Model (DRBM)*, a family of distributions over degree regular graphs with planted community structure. The degree-regularity will simplify some aspects of our analysis, allowing us to illustrate key features of the LoSt hierarchy without a proliferation of technicalities. We will comment later on about the possibilities for extension to the irregular case. As an aside, we cannot help but editorialize briefly that, although the DRBM is less useful in practice than the standard block model discussed above, its combinatorics are intricate and beautiful in their own right, and the related case of d -regular graphs with planted colorings have been quite well-studied.

We will specify the DRBM on n vertices in full generality by several parameters: the number of communities k , degree d , and a $k \times k$ transition matrix M for a reversible Markov chain, with stationary distribution π . In other words, M has row sums equal to one, and $\text{Diag}(\pi)M$ is a symmetric matrix. To sample a graph $\mathbf{G} = (V(\mathbf{G}), E(\mathbf{G}))$ —we will use bold-face type for random objects throughout the

paper—first partition the n vertices randomly into k groups $V_1(\mathbf{G}), \dots, V_n(\mathbf{G})$ with $|V_i(\mathbf{G})| = \pi(i)n$, and then choose a d -regular random graph conditional on there being $\pi(i)M_{i,j}dn$ edges between groups $i \neq j$ and $\pi(i)M_{i,i}dn/2$ internal to each group i . As $\text{Diag}(\pi)M$ is symmetric, this process is well-defined. We will assume always that the parameters are set to make these quantities integer-valued; settings for which this holds infinitely often as $n \rightarrow \infty$ are dense in the parameter space.

Remark 2.1. The DRBM as we have defined it differs from the Regular Stochastic Block Model of [BDG⁺16], in which each vertex has a prescribed number of neighbors in every community. Although superficially similar, the behavior of this ‘equitable’ model (as it is known in the physics literature [NM14]) is quite different from ours. For instance, [BDG⁺16] show that whenever detection is possible, one can recover the community labels *exactly*. This is not true in our case.

The DRBM contains several more familiar distributions as special cases, and the reader is welcome to focus on her favorite for concreteness. When $\pi(i) = 1/k$ for every i , we have the DRBM with equal groups. Setting $M_{i,i} = 0$ and $M_{i,j} = \frac{1}{k-1}$, we are in a somewhat restrictive case of the planted k -coloring model, where each pair of color classes has the same number of edges between them. We will refer to the case when $M_{i,i} = m_{\text{in}}$ and m_{out} otherwise as the symmetric DRBM. As M describes a reversible Markov chain, its spectrum is real, and we will write its eigenvalues as $1 = \lambda_1 \geq |\lambda_2| \geq \dots \geq |\lambda_k|$. The second eigenvalue λ_2 can be thought of as a kind of signal-to-noise ratio, and will be repeatedly important to our analysis. One can verify, for instance, that in the case of the symmetric DRBM, $\lambda_2 = \dots = \lambda_n = m_{\text{in}} - m_{\text{out}}$.

It is widely believed that the threshold behavior of the DRBM is similar to that of the SBM, though the inhomogeneities in group size and edge density we allow for make the situation somewhat more complicated than in the symmetric case discussed earlier. This phenomenology includes an information-theoretic threshold $d_{\text{IT}} \approx \frac{\log k}{k\lambda_2^2}$ for the symmetric DRBM (and a more complicated characterization in general that will not be relevant to us here). In the general model, the Kesten-Stigum threshold for *detection* is $d_{\text{KS}} \triangleq \lambda_2^{-2} + 1$, and we expect recovery of *all* communities once $d > 1/\lambda_k^2 + 1$. However, most formal treatment in the literature has been limited to the distribution of d -regular graphs conditional on having a planted k -coloring, a case not fully captured by our model. Characterization of the information-theoretic threshold, even for the symmetric DRBM remains largely folklore, and in Appendix [ref] we will for good measure provide a few rigorous pieces of the picture.

Our main theorem is that the the Local Statistics hierarchy can robustly solve the detection problem on the DRBM whenever $d > d_{\text{KS}}$, but that otherwise any constant level fails to do so.

Theorem 2.2. *For every $\epsilon > 0$, and set of parameters (d, k, M, π) satisfying $d > d_{\text{KS}} + \epsilon$, there exists $m \in \mathbb{N}$ sufficiently large so that with probability $1 - o(1)$ the $\text{LoSt}(2, m)$ SDP, given an input graph \mathbf{G} , can distinguish in time [need] whether*

- \mathbf{G} is a uniformly random d -regular graph
- \mathbf{G} is sampled from the DRBM with parameters (d, k, M, π)

and is robust to adversarial addition or deletion of $o(n)$ edges. On the other hand, for any constant m and $d < d_{\text{KS}}$, the $\text{LoSt}(2, m)$ SDP fails with probability $1 - o(1)$ to distinguish.

We also prove a stronger robustness guarantee, in particular that that $\text{LoSt}(2, m)$ can tolerate $\rho_m n$ adversarial edge perturbations, although $\rho_m \rightarrow 0$ as we move up the hierarchy. This creates a trade-off

between robustness, which we lose as added information is incorporated to the SDP at each successive level, and fidelity to the threshold, which we approach as $m \rightarrow \infty$.

Theorem 2.3. *For every $\rho > 0$, there exists $\delta > 0$ and m sufficiently large, so that even given a graph $\tilde{\mathbf{G}}$ which is a $\rho|E|$ -perturbation of the edges of some \mathbf{G} , $\text{LoSt}(2, m)$ can be used to distinguish whether \mathbf{G} is a uniformly random d -regular graph or was drawn from a DRBM ϵ -away from the threshold.*

Along the way we will inadvertently prove that standard spectral detection using the adjacency matrix succeeds above d_{KS} , but cannot have the same robustness guarantee. It is a now-classic result of Friedman that, with probability $1 - o_n(1)$, the spectrum of a uniformly random d -regular graph is within $o_n(1)$ of $(-2\sqrt{d-1}, 2\sqrt{d-1}) \cup \{d\}$. Conversely, we show:

Corollary 2.4. *Let \mathbf{G} be drawn from the DRBM with parameters (d, k, M, π) satisfying $d > d_{\text{KS}} + \epsilon$. There exists some $\eta = \eta(\epsilon)$ such that, for each eigenvalue λ of M satisfying $|\lambda| > 1/\sqrt{d-1} + \epsilon$, the adjacency matrix $A_{\mathbf{G}}$ is guaranteed one eigenvalue μ satisfying $|\mu| > 2\sqrt{d-1} + \eta$.*

Regrettably, we do not resolve to similar satisfaction the issue of efficient or robust recovery above Kesten-Stigum. However, in Appendix A we will reduce some central aspects of this issue to the following conjecture regarding the spectrum of $A_{\mathbf{G}}$ for \mathbf{G} drawn from the planted model.

Conjecture 2.5. *Let $\mathcal{P}_{(d,k,M,\pi)}$ be any DRBM with $|\lambda_1|, \dots, |\lambda_\ell| > (d-1)^{-1/2}$. Then, for any η , with high probability, $A_{\mathbf{G}}$ has only ℓ eigenvalues with modulus larger than $2\sqrt{d-1} + \eta$.*

We will discuss in Appendix A that, conditional on this conjecture (or even a weaker version in which we are guaranteed only constantly many eigenvalues outside the bulk), (i) the span of the corresponding eigenvectors is correlated to the community structure, and (ii) the Local Statistics hierarchy can robustly produce vectors with macroscopic correlation to this span. From weak convergence of the empirical spectral distribution of $A_{\mathbf{G}}$ to the Kesten-McKay law, we know that there must be $o(n)$ eigenvalues with modulus larger than $2\sqrt{d-1}$, it will take substantial technical work to push this down to $O(1)$. We believe the most feasible approach is a careful mirror of the techniques in [BLM15], but the execution of this is beyond the scope of this paper. These issues and a related conjecture are discussed in ?? in the context of the DRBM with two groups.

Related Work. Semidefinite programming approaches have been most studied in the dense, irregular case, where exact recovery is possible (for instance [ABH16, AS15]), and it has been shown that an SDP relaxation can achieve the information-theoretically optimal threshold [HWX16]. However, in the sparse regime we consider, the power of SDP relaxations for weak recovery remains unclear. Guedon and Vershynin [GV16] show upper bounds on the estimation error of a standard SDP relaxation in the sparse, two-community case of the SBM, but only when the degree is roughly 10^4 times the information theoretic threshold. More recently, in a tour-de-force, Montanari and Sen [MS15] showed that for two communities, the SDP of Guedon and Vershynin achieves the information theoretically optimal threshold for large but constant degree, in the sense that the performance approaches the threshold if we send the number of vertices, and then the degree, to infinity. Semi-random graph models have been intensively studied in [BS95, FK00, FK01, CO04, KV06, CO07, MMV12, CJSX14, GV16] and we refer the reader to [MMV16] for a more detailed survey. In the logarithmic-degree regime, robust algorithms for community detection are developed in [CL⁺15, KK10, AS12]. Far less is known in the case of regular graphs.

3 Technical Overview

Denote by \mathcal{N} the uniform distribution on n -vertex d -regular graphs, and write $\mathcal{P} = \mathcal{P}_{d,k,M,\pi}$ the DRBM. We will use bold face font for random objects sampled from these distributions. Because we care only about the case when the number of vertices is very large, we will use *with high probability (w.h.p)* to describe any sequence of events with probability $1 - o_n(1)$ in \mathcal{N} or \mathcal{P} as $n \rightarrow \infty$. We will write $[n] = \{1, \dots, n\}$, and in general use the letters u, v, w to refer to elements of $[n]$ and i, j for elements of $[k]$. The identity matrix will be denoted by $\mathbb{1}$, and we will write X^T for the transpose of a matrix X , $\langle X, Y \rangle = \text{tr } X^T Y$ for the standard matrix inner product, and $\|X\|_F$ for the associated Frobenius norm. Positive semidefiniteness will be indicated with the symbol \succeq . The standard basis vectors will be denoted e_1, e_2, \dots , the all-ones vector written as e , and the all-ones matrix as $\mathbb{J} = ee^T$. Finally, let $\text{diag} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}$ be the function extracting the diagonal of a matrix, and $\text{Diag} : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times n}$ be the one which populates the nonzero elements of a diagonal matrix with the vector it is given as input.

3.1 Detection, Refutation, and Sum-of-Squares

We will begin the discussion of the Local Statistics algorithm by briefly recalling Sum-of-Squares programming. Say we have a constraint satisfaction problem presented as a system of polynomial equations in variables $x = (x_1, \dots, x_n)$ that we are to simultaneously satisfy. In other words, we are given a set

$$\mathcal{S} = \{x \in \mathbb{R}^n : f_1(x), \dots, f_m(x) = 0\}$$

and we need to decide if it is non-empty. Whenever the problem is satisfiable, any probability distribution supported on \mathcal{S} gives rise to an operator $\mathbb{E} : \mathbb{R}[x] \rightarrow \mathbb{R}$ mapping a polynomial x to its expectation. Trivially, \mathbb{E} obeys

$$\text{Normalization} \quad \mathbb{E} \mathbb{1} = 1 \quad (1)$$

$$\text{Satisfaction of } \mathcal{S} \quad \mathbb{E} f_i(x) \cdot p(x) = 0 \quad \forall i \in [m], \forall p \in \mathbb{R}[x] \quad (2)$$

$$\text{Positivity} \quad \mathbb{E} p(x)^2 \geq 0 \quad \forall p \in \mathbb{R}[x] \quad (3)$$

In general, we will say that an operator mapping some subset of $\mathbb{R}[x]$ to the reals is *normalized*, *satisfies* \mathcal{S} , or is *positive* if it obeys (1), (2), or (3), respectively, on all polynomials in its domain.

Proving that $\mathcal{S} = \emptyset$, and thus that our problem is unsatisfiable, is equivalent to showing that no operator obeying (1)-(3) can exist. The key insight of SoS is that, at least sometimes, one can do this by focusing only on polynomials of some bounded degree. Writing $\mathbb{R}[x]_{\leq D}$ for the polynomials of degree at most D , we call an operator $\tilde{\mathbb{E}} : \mathbb{R}[x]_{\leq D} \rightarrow \mathbb{R}$ a *degree-D pseudoexpectation* if it is normalized, and for every polynomial in its domain satisfies \mathcal{S} and is positive. It is well-known that one can search for a degree D pseudoexpectation with a semidefinite program of size $O(n^D)$, and if this smaller, relaxed problem is infeasible, we've shown that \mathcal{S} is empty. This is the *degree-D Sum-of-Squares relaxation* of our CSP.

A naive way to employ SoS for hypothesis testing or reconstruction problems such as community detection is to choose some statistic known to distinguish the planted and null distributions, and write down a relaxed sum-of-squares search algorithm for this statistic. In the case of the DRBM, a graph drawn from the planted model is guaranteed a partition of the vertices into groups of sizes $\pi(i)n$, with $\pi(i)M_{i,j}dn$ edges between groups i and j . Let us refer to such a partition $\sigma : [n] \rightarrow [k]$ as *M-good*. A routine first moment calculation shows that when d is sufficiently large, uniformly random d -regular graphs from the null distribution, \mathcal{N} , are exponentially unlikely to have an *M-good* partition.

Proposition 3.1. *With probability $1 - o_n(1)$ (in fact, exponentially close to one) a graph \mathbf{G} from the null model has no M -good partitions whenever*

$$d - 1 > \frac{H(\pi) + H(\pi, M)}{H(\pi) - H(\pi, M)}, \quad (4)$$

where $H(\pi) = -\sum_i \pi(i) \log \pi(i)$ is the standard Shannon entropy, and $H(\pi, M)$ is the average with respect to π of the entropy of the rows of M .

Thus we can solve detection in exponential time above this first moment threshold by exhaustively searching for even one M -good division of the vertices. In other words, detection in this regime is no harder than *refutation* of an M -good partition. This refutation problem can be encoded with kn variables $x_{u,i}$, describing whether each vertex $u \in [n]$ is in group $i \in [k]$, subject to the polynomial constraints

$$\begin{array}{lll} \text{Boolean} & x_{u,i}^2 = x_{u,i} & \forall u \in [n] \text{ and } i \in [k] \\ \text{Single Color} & \sum_i x_{u,i} = 1 & \forall u \in [n] \\ \text{Group size} & \sum_u x_{u,i} = \pi(i)n & \forall i \in [k] \\ \text{M-good} & \sum_{(u,v) \in E} x_{u,i} x_{v,j} = \pi(i)M_{i,j}dn & \forall i, j \in [k] \end{array}$$

It will be useful later to denote by $\mathcal{B}_k \subset \mathbb{R}^{nk}$ the set described by the Boolean and Single Color equations above. Each level of the SoS Hierarchy, applied to the polynomial system described above, immediately gives us a one-sided detection algorithm: if given a graph \mathbf{G} the degree- D SoS relaxation is infeasible, we can be sure that there are no M -good partitions, and thus that graph came from the null model and not the planted one. However, as it is a relaxation, if this SDP is feasible we have not a priori learned anything at all. For a two-sided test we need to prove that with high probability there is no feasible solution for graphs drawn from the null model.

There are two fundamental limitations to this approach. First, statistics like existence of an M -good partition are in some cases not optimal for differentiating the null and planted distributions. Consider for simplicity a less constrained version of the symmetric DRBM, where for a parameter $\lambda < 0$ we partition the vertices into 2 equal sized groups, and sample a d -regular graph conditional on there being $(1 - \lambda)dn/4$ edges among vertices in the same community, with the remaining $(1 + \lambda)dn/4$ connecting vertices in different groups. Both the information theoretic and Kesten-Stigum thresholds in this case occur when $\lambda > 1/\sqrt{d} - 1$. Such graphs are guaranteed to have a maximum cut of at least $(1 + \lambda)dn/4$, so we can distinguish the null and planted models for any λ making this larger than the maximum cut in a d -regular random graph. However, we know from work of Dembo et al. [DMS⁺17] that the maximum cut in d -regular random graphs is, with high probability,

$$\left(1 + \frac{2P_*}{\sqrt{d}} + o_d(\sqrt{d})\right) \frac{dn}{4} + o_n(n),$$

where $2P_* \approx 1.5264$ is twice the vaunted Parisi constant from statistical physics. Thus, when d is large, the maximum cut cannot distinguish the null and planted distributions until roughly $\lambda > 2P_*/\sqrt{d} - 1$, i.e $d > 4P_*^2 d_{KS}$. This same phenomenon holds in the irregular SBM with two groups.

The second issue is that even in regimes where we know detection can be performed by exhaustive search for an M -good partition, low-degree SoS relaxations of this search problem are known to fail. In the

case of the symmetric DRBM, with $m_{\text{in}} < m_{\text{out}}$, a similar first moment bound to the one above shows that at roughly the same threshold, random d -regular graphs are exponentially unlikely to have any k -way cut with the same total number of between-group edges as the hidden partition in the planted model. Banks et al. [BKM17] show that, for the degree-two SoS relaxation of k -way cut, detection is only possible once $d > 4d_{\text{KS}}$: for smaller degree, when \mathbf{G} is sampled from the null model, there exists a feasible degree-two pseudoexpectation. A similar result for a slightly weaker SDP holds in the case of Erdős-Rényi graphs with planted k -colorings [BT19].

This is not the only case where degree-two SoS for refutation does not succeed all the way down to the conjectured computational threshold for detection. Consider for instance the *Rademacher-spiked Wigner model*, where our goal is to distinguish whether an observed matrix X is either (Null) an $n \times n$ Wigner matrix W , with $W_{i,j} \sim \mathcal{N}(0, 1/n)$ and $W_{i,i} \sim \mathcal{N}(0, 2/n)$, or (Planted) of the form $X = \lambda n^{-1} \sigma \sigma^T + W$ for some uniformly random hypercube vector $\sigma \in \{\pm 1\}^n$. Results of Feral and Peche [FP07] tell us that detection is possible simply by examining the spectrum of X , whenever $\lambda > 1$, and Perry et al. [PWBM16] show that this is in fact the information-theoretic threshold. On the other hand, the planted model satisfies $\sigma^T X \sigma \approx \lambda n$, so we could try and solve detection by refuting the existence of a hypercube vector with a large quadratic form. Unfortunately, in the null model $X = W$, degree-two SoS can only refute the existence of some $\tau \in \{\pm 1\}^n$ satisfying $\tau^T X \tau > 2$ [MS15]. Bandeira et al. provide evidence, using ideas of Hopkins and Steurer regarding low-degree test statistics, that there is a fundamental computational barrier to outperforming degree-two SoS at this refutation task [BKW19]; quite recently, [KB19] show that this gap persists for degree-four SoS, and conjecture that refutation of any smaller maximum is impossible for SoS of constant degree.

These results fit into a broader current in the literature probing the nature and origin of computational barriers in random refutation problems. In the preceding discussion, we were attempting to solve detection in the DRBM, for d in the conjectured computationally feasible regime, by refuting the existence of some combinatorial structure in the observed graph. However, refutation is essentially a prior-free task! There are, at least potentially, many planted distributions for producing graphs with M -good partitions—just as there are many ways to produce a Gaussian random matrix whose maximum quadratic form over the hypercube is atypically large—and *they need not all have the same computational phase transition*. The idea is that refutation in the null model is hard exactly when it would allow us to solve *detection* in the computationally hard or information-theoretically impossible regime of some ‘quietly’ planted distribution, whose low degree moments mimic those of the null model (see [BKW19], for example).

All of this is bad news for refutation, but not necessarily for detection. The problem of detection and the related one on reconstruction are in a Bayesian setting, where the prior distribution is completely specified. Yet, the semi-definite programs described above use little information from the prior distribution in their formulation. Why not include information about the prior distribution in our SDP?

3.2 The Local Statistics Hierarchy

Let us regard the planted model as a joint distribution on random variables $\mathbf{x} = \{\mathbf{x}_{u,i}\}$ encoding the group labels, and $\mathbf{G} = \{\mathbf{G}_{u,v}\}$ indexed by $\{u, v\} \subset [n]$ and describing which edges of the graph are present. Instead of our somewhat ad-hoc SDP relaxing the problem of searching for an M -good partition, we will try and find a pseudoexpectation on the variables $\mathbf{x}_{u,i}$ which (i) satisfies \mathcal{B}_k —the Boolean and Single-Color constraints—and (ii) matches certain low-degree moments of the planted distribution. To a first approximation, we will

add constraints of the form

$$\tilde{\mathbb{E}} p(\mathbf{G}, \boldsymbol{\chi}) \simeq \mathbb{E}_{(\mathbf{G}, \boldsymbol{\chi}) \sim \mathcal{P}} p(\mathbf{G}, \boldsymbol{\chi}),$$

for a restricted class of polynomials p in variables $\boldsymbol{\chi} = \{\chi_{u,i}\}_{u \in [n], i \in [k]}$ and $\mathbf{G} = \{\mathbf{G}_{u,v}\}_{u,v \in [n]}$. The exact meaning of \simeq will depend on the concentration of $p(\mathbf{G}, \boldsymbol{\chi})$ with respect to the randomness in $\boldsymbol{\chi}$ and \mathbf{G} ; we will make it precise below.

The DRBM has a natural symmetry: we can freely permute the vertices, and the distribution is unchanged. This gives us an action of \mathfrak{S}_n , the symmetric group on n elements, on the random variables $\boldsymbol{\chi} = \{\chi_{u,i}\}$ and $\mathbf{G} = \{\mathbf{G}_{u,v}\}$ describing our random graphs, and their non-random counterparts $\boldsymbol{\chi} = \{\chi_{u,i}\}$ and $\mathbf{G} = \{\mathbf{G}_{u,v}\}$ appearing in the polynomials in the domain of $\tilde{\mathbb{E}}$. In particular, $\theta \in \mathfrak{S}_n$ acts as $\theta : \chi_{u,i} \mapsto \chi_{\theta(u),i}$ and $\theta : \mathbf{G}_{u,v} \mapsto \mathbf{G}_{\theta(u),\theta(v)}$. It is only meaningful to consider polynomials in $\boldsymbol{\chi}$ and \mathbf{G} that are fixed under this action; these roughly correspond to counting the instances of subgraphs of \mathbf{G} with vertices constrained to have particular labels. Note that unless we are in the case of the symmetric DRBM, the community labels do not have a similar symmetry.

Since the random variables \mathbf{G} are all zero-one indicators, we only need consider polynomials $p(\boldsymbol{\chi}, \mathbf{G})$ that are multilinear in \mathbf{G} . We claim that every such polynomial in $\mathbb{R}[\boldsymbol{\chi}, \mathbf{G}]$ fixed under this action, and with degrees D_1 and D_2 in the $\boldsymbol{\chi}$ and \mathbf{G} variables respectively, is of the following form. Let $H = (V(H), E(H))$ be a graph with at most D_2 edges, $S \subset V(H)$ a designated subset of at most D_1 vertices, and $\tau : S \rightarrow [k]$ a set of labels on these distinguished vertices. Write Φ_H for the set of all injective homomorphisms $\varphi : H \rightarrow \mathbf{G}$, i.e. maps for which (1) $\varphi(a) \neq \varphi(b)$ for every distinct $a, b \in V(H)$ and (2) $(a, b) \in E(H)$ implies $(\varphi(a), \varphi(b)) \in E(\mathbf{G})$. The image of each $\varphi \in \Phi_H$ is a copy of H inside \mathbf{G} . For each, there is a corresponding polynomial

$$p_{H,S,\tau}(\boldsymbol{\chi}, \mathbf{G}) = \sum_{\varphi \in \Phi_H} \prod_{u \in S} \chi_{\varphi(u), \tau(u)}, \quad (5)$$

that counts occurrences H in \mathbf{G} which conform, on the vertices in S , to the labels specified by τ . One can check that these polynomials are a basis for the vector space of polynomials in $\mathbb{R}[\boldsymbol{\chi}, \mathbf{G}]$ fixed under the action above.

Definition 3.2. The degree (D_1, D_2) level of the Local Statistics hierarchy is the following SDP: find a degree- D_1 pseudoexpectation $\tilde{\mathbb{E}}$ satisfying \mathcal{B}_k , such that

$$\tilde{\mathbb{E}} p_{H,S,\tau}(\boldsymbol{\chi}, \mathbf{G}) \approx \mathbb{E}_{(\boldsymbol{\chi}, \mathbf{G}) \sim \mathcal{P}} p_{H,S,\tau}(\boldsymbol{\chi}, \mathbf{G}) \quad (6)$$

for every $|S| \leq D_1$ and $|E(H)| \leq D_2$.

Note that, among many new constraints that this SDP imposes on $\tilde{\mathbb{E}}$, it recovers the conditions on group size and M -good-ness from our earlier SoS relaxation, as

$$\sum_i \chi_{u,i} \quad \text{and} \quad \sum_{(u,v) \in E} \chi_{u,i} \chi_{v,j}$$

are both of the form (5). We obtain the first when H is the graph on one vertex with label i , and the second when H is a single edge, with endpoints labeled i and j .

Remark 3.3. Although we have stated it in the specific context of the DRBM, the local statistics framework extends readily to any planted problem involving a joint distribution μ on pairs $(\boldsymbol{\chi}, \mathbf{G})$ of a hidden structure and observed signal, if we take appropriate account of the natural symmetries in μ . For a broad range of such

problems, including spiked random matrix models [AKJ18, PWBM16], compressed sensing [ZK16, Ran11, KGR11] and generalized linear models [BKM⁺19] (to name only a few) there are conjectured computational thresholds where the underlying problem goes from being efficiently solvable to computationally intractable, and the algorithms which are proven or conjectured attain this threshold are often not robust. We hope that the local statistics hierarchy can be harnessed to design robust algorithms up to these computational thresholds, as well as to provide evidence for computational intractability in the conjectured hard regime. The relation (if any) between the local statistics SDP hierarchy and iterative methods such as belief propagation or AMP is also worth investigating.

The remainder of the paper will be laid out as follows. In Section 4 we will collect some preliminary results, including several standard and useful observations on non-backtracking walks and reversible Markov chains. Section 5 contains the proof that our SDP can distinguish the null and planted models above the KS threshold, and Section 6 adapts this proof to show that spectral distinguishing is possible in this regime as well. In Section 7 we prove the other half of Theorem 2.2, namely that no constant level of our hierarchy succeeds below this threshold. Section 8 concerns the robustness guarantees of our algorithm. Finally, in Appendix B, we will perform several calculations on the DRBM, including the first moment bound of Proposition 3.1, and the explicit computation of the local statistics appearing in the LoSt hierarchy.

4 Preliminaries

4.1 Nonbacktracking Walks and Orthogonal Polynomials

The central tool in our proofs will be *non-backtracking walks* on G —these are walks which on every step are forbidden from visiting the vertex they were at two steps previously. We will collect here some known results on these walks specific to the case of d -regular graphs. Write $A_G^{(s)}$ for the $n \times n$ matrix whose (v, w) entry counts the number of length- s non-backtracking walks between vertices v and w in G . One can check that the $A_G^{(s)}$ satisfy a two-term linear recurrence,

$$\begin{aligned} A_G^{(0)} &= \mathbb{1} \\ A_G^{(1)} &= A_G \\ A_G^{(2)} &= A_G^2 - d\mathbb{1} \\ A_G^{(s)} &= AA_G^{(s-1)} - (d-1)A_G^{(s-2)} \quad s > 2, \end{aligned}$$

since to enumerate non-backtracking walks of length s , we can first extend each such walk of length $s-1$ in every possible way, and then remove those extensions that backtrack.

On d -regular graphs, the above recurrence immediately shows that $A_G^{(s)} = q_s(A_G)$ for a family of monic, scalar ‘non-backtracking polynomials’ $\{q_s\}_{s \geq 0}$, where $\deg q_s = s$. To avoid a collision of symbols, we will use z as the variable in all univariate polynomials appearing in the paper. It is well known that these polynomials are an orthogonal polynomial sequence with respect to the Kesten-McKay measure

$$d\mu_{\text{KM}}(z) = \frac{1}{2\pi} \frac{d}{\sqrt{d-1}} \frac{\sqrt{4(d-1)-z^2}}{d^2-z^2} dz \mathbf{1}_{\left[|z| < 2\sqrt{d-1}\right]},$$

with its associated inner product

$$\langle f, g \rangle_{\text{KM}} \triangleq \int f(z)g(z) d\mu_{\text{KM}}(z)$$

on the vector space of square integrable functions on $(-2\sqrt{d-1}, 2\sqrt{d-1})$. One can again check that

$$\|q_s\|_{\text{KM}}^2 \triangleq \int q_s(z)^2 d\mu_{\text{KM}} = q_s(d) = \begin{cases} 1 & s = 0 \\ d(d-1)^{s-1} & s \geq 1 \end{cases} = \frac{1}{n} (\# \text{ length-}s \text{ n.b. walks on } G)$$

in the normalization we have chosen [ABLS07]. Thus any function f in this vector space can be expanded as

$$f = \sum_{s \geq 0} \frac{\langle f, q_s \rangle_{\text{KM}}}{\|q_s\|_{\text{KM}}^2} q_s.$$

We will also need the following lemma of Alon et al. [ABLS07, Lemma 2.3] bounding the size of the polynomials q_s :

Lemma 4.1. *For any $\varepsilon > 0$, there exists an $\eta > 0$ such that for $z \in [-2\sqrt{d-1} - \eta, 2\sqrt{d-1} + \eta]$,*

$$|q_s(z)| \leq 2(s+1)\|q_s\|_{\text{KM}} + \varepsilon.$$

The behavior of the non-backtracking polynomials with respect to the inner product $\langle \cdot, \cdot \rangle_{\text{KM}}$ idealizes that of the $A_G^{(s)} = q_s(A_G)$ under the trace inner product. In particular, if $s + t < \text{girth}(G)$

$$\langle A_G^{(s)}, A_G^{(t)} \rangle = n \langle q_s, q_t \rangle_{\text{KM}} = \begin{cases} n (\# \text{ length-}s \text{ n.b. walks on } G) & s = t \\ 0 & s \neq t \end{cases}.$$

This is because the diagonal entries of $A_G^{(s)} A_G^{(t)}$ count pairs of non-backtracking walks with length s and t respectively: if $s \neq t$ any such pair induces a cycle of length at most $s + t$, or perhaps is a pair of identical walks in the case $s = t$. Above the girth, if we can control the number of cycles, we can quantify how far the $A_G^{(s)}$ are from orthogonal in the trace inner product.

Luckily for us, sparse random graphs have very few cycles. To make this precise, call a vertex *bad* if it is at most L steps from a cycle of length at most C . These are exactly the vertices for which the diagonal entries of $A_G^{(s)} A_G^{(t)}$ are nonzero, when $s + t < C + L$.

Lemma 4.2. *For any constant C and L , with high probability any graph $G \sim \mathcal{P}$ has at most $O(\log n)$ bad vertices.*

We will defer the proof of this lemma to the appendix, but two nice facts follow from it immediately. First, from the above discussion,

$$\langle A_G^{(s)}, A_G^{(t)} \rangle = O(\log n)$$

for any $s, t = O(1)$. The second useful corollary is more or less that in random graphs we can use non-backtracking walks as a proxy for self-avoiding ones.

Lemma 4.3. *Write $A_G^{(s)}$ for the $n \times n$ matrix whose i, j entry is a one exactly when i and j are connected by a self-avoiding walk of length s . Then with high probability, for any graph $G \sim \mathcal{P}$,*

$$\left\| A_G^{(s)} - A_G^{(s)} \right\|_{\text{F}}^2 = O(\log n) \tag{7}$$

Proof. Every row of both $A_s^{(G)}$ and $A_s^{(G)}$ have L^2 norm $O(1)$, and they differ only in the rows corresponding to, say, the $2s$ -bad vertices, of which there are here are only $O(\log n)$. \square

4.2 Reversible Markov Chains

We will need standard fact about reversible Markov chains. Let us maintain the notation for M , its eigenvalues $1 = \lambda_1 \geq |\lambda_2| \geq \dots \geq |\lambda_k|$, and its stationary distribution π . Recall from above that $Me = e$, $\pi^T M = \pi^T$, and the reversibility condition on M means $\text{Diag}(\pi)M$ is symmetric.

Lemma 4.4. *Let F be the matrix of right eigenvectors, normalized so that the columns have unit norm (note that the first column of F is, up to scaling, the all-ones vector). Then $F^{-1} \text{Diag}(\pi)F = \mathbb{1}$.*

Proof. First, reversibility tells us $\text{Diag}(\pi)^{1/2}M \text{Diag}(\pi)^{-1/2}$ is symmetric, and thus by the spectral theorem that it satisfies

$$\text{Diag}(\pi)^{1/2}M \text{Diag}(\pi)^{1/2}O = O\Lambda$$

for some orthogonal O . It is readily seen that $M \text{Diag}(\pi)^{-1/2}O = \text{Diag}(\pi)^{-1/2}O\Lambda$, so $\text{Diag}(\pi)^{-1/2}O$ contains, up to scaling, the right eigenvectors of M . \square

4.3 Local Statistics in the Planted Model

The Local Statistics SDP that we are studying includes constraints that our pseudoexpectation match certain low-degree moments in the planted distribution. As we discussed in the technical overview, these correspond to the counts of partially labelled subgraphs in G . To set some notation, a *partially labelled graph* (H, S, τ) is a graph $H = (V(H), E(H))$, together with a distinguished subset of vertices $S \subset V(H)$, and a labelling $\tau : S \rightarrow [k]$ of these distinguished vertices. We'll say a graph is *unlabelled* or *fully labelled* if $S = \emptyset$ or $S = V(H)$, and in these cases abuse notation and simply refer to H or (H, τ) respectively. At times it will also be useful to refer to graphs with distinguished vertices, but no labelling; we will write these as (H, S) . An *occurrence* of a partially labelled graph (H, S, τ) in a fully labelled one (G, σ) is an injective homomorphism $H \rightarrow G$, that agrees on labels, i.e. vertices in S are mapped to ones in $V(G)$ with the same label.

The low-degree moment constraints in $\text{LoSt}(2, m)$ are exactly the counts of occurrences of partially labelled subgraphs (H, S, τ) in a graph $G \sim \mathcal{P}$, for which H has at most m edges and 2 distinguished vertices. The following theorem characterizes these counts in any planted model; we will discuss it briefly below and remit the proof to the appendix.

Definition 4.5. Let (H, S) be a connected graph on $O(1)$ edges, with distinguished vertices S . Define $C_{H,S,d}$ to be the number of occurrences of (H, S) in an infinite d -regular tree in which some vertex in S is mapped to the root. If $S = \emptyset$, choose some distinguished vertex arbitrarily—the count will be the same no matter which one is chosen; we will at times use $C_{H,d}$ as shorthand in this case. Finally, if $(H, S) = (H_1, S_1) \sqcup \dots \sqcup (H_\ell, S_\ell)$ has ℓ connected components, take $C_{H,S,d} = C_{H_1,S_1} \dots C_{H_\ell,S_\ell}$. We note for later use that if H contains a cycle, $C_{H,S} = 0$, and if it is a path of length s with endpoints distinguished, $C_{H,S} = \|q_s\|_{\text{KM}}^s$, the number of vertices at depth s in the tree.

Theorem 4.6 (Local Statistics). *If (H, S, τ) is a partially labelled graph with $O(1)$ edges, then in any planted model $\mathcal{P}_{d,k,M,\pi}$*

1. If H is unlabelled, i.e. $S = \emptyset$, then $n^{-\ell} \mathbb{E} p_{H,S,\tau}(\mathbf{x}, \mathbf{G}) \rightarrow C_{H,S,d}$
2. If H is labelled, with $S = \{\alpha, \beta\}$, $\tau(\alpha) = i$, and $\tau(\beta) = j$, then

$$n^{-\ell} \mathbb{E} p_{H,S,\tau}(\mathbf{x}, \mathbf{G}) \rightarrow \pi(i) M_{i,j}^{\text{dist}(\alpha,\beta)} C_{H,S,d},$$

and $p_{H,S,\tau}(\mathbf{x}, \mathbf{G})$ enjoys concentration up to an additive $\pm o(n^\ell)$. We say that $\text{dist}(\alpha, \beta) = \infty$ if these two vertices lie in disjoint components of H , and we interpret $M_{i,j}^\infty = \pi(j)$.

Remark 4.7. In our Local Statistics SDP 3.2, we promised to formalize the symbol \simeq appearing in the affine moment-matching constraints on the pseudoexpectation; let's do so now. Throughout the paper, fix a very small error tolerance $0 < \delta$, and write \simeq_ℓ to mean "equal up to $\pm \delta n^\ell$ ". Then the constraint for each partially labelled subgraph with ℓ connected components should read $\mathbb{E} p_{H,S,\tau}(\mathbf{x}, \mathbf{G}) \simeq_\ell \mathbb{E} p_{H,S,\tau}(\mathbf{x}, \mathbf{G})$. We will write \simeq instead of \simeq_1 whenever there is no chance for confusion. Finally, because we have defined our model quite rigidly, whenever (H, S, τ) consists of a single vertex with label $i \in [k]$, $p_{H,S,\tau}(\mathbf{x}, \mathbf{G}) = \pi(i)n$. Similarly when (H, S) consists of two distinguished vertices with labels $i, j \in [k]$ respectively,

$$p_{H,S,\tau}(\mathbf{x}, \mathbf{G}) = \begin{cases} \pi(i)\pi(j)n^2 & i \neq j \\ \pi(i)^2 n^2 - \pi(i)n & i = j \end{cases}$$

and the moment-matching constraints in our SDP will accordingly include $=$ instead of \simeq .

Let's take a moment and get a feel for Theorem 4.6. As a warm-up, consider the case when (H, S, τ) is a path of length $s \leq m$ with the endpoints labelled as $i, j \in [k]$, and we simply need to count the number of pairs of vertices in G with labels i and j respectively that are connected by a path of length s . As d -regular random graphs from models like \mathcal{P} have very few short cycles, assume for simplicity that the girth is in fact much larger than m , so that the depth- s neighborhood about every vertex is a tree. If we start from a vertex i and follow a uniformly random edge, the parameter matrix M from our model says that, on average at least, the probability of arriving at a vertex in group j is roughly $M_{i,j}$, and similarly if we take s (non-backtracking) steps, this probability is roughly $M_{i,j}^s$. There are $\pi(i)n$ starting vertices in group i , and $d(d-1)^{s-1}$ vertices at distance s from any such vertex.

If (H, S, τ) is a tree in which the two distinguished vertices are at distance s , then we can enumerate occurrences of (H, S, τ) in G by first choosing the image of the path connecting these two, and then counting the ways to place the remaining vertices. If we again assume that the girth is sufficiently large, it isn't too hard to see that the number of ways to do this second step is a constant independent of the number of ways to place the path, so we've reduced to the case above. The idea for the cases $|S| = 0, 1$ is similar. We'll prove Theorem 4.6 in Appendix B.1.

5 Distinguishing with Local Statistics

Throughout this section, fix the parameters (d, k, M, π) of a planted model \mathcal{P} . We'll prove half of our main theorem, namely that for any $\epsilon > 0$, if

$$d > d_{KS} + \epsilon = 1 + \frac{1}{\lambda_2^2} + \epsilon$$

then there exists some m so that the $\text{LoSt}(2, m)$ SDP can distinguish the planted and null models. When $(\mathbf{x}, \mathbf{G}) \sim \mathcal{P}$, the SDP is surely feasible as we can simply set

$$\tilde{\mathbb{E}} p(\mathbf{x}, \mathbf{G}) = p(\mathbf{x}, \mathbf{G})$$

for any polynomial we choose. We will thus be done if we can show *infeasibility* when \mathcal{P} is above the KS threshold, m is sufficiently large, and $\mathbf{G} \sim \mathcal{N}$. Our strategy will be to first reduce to the problem of designing a univariate polynomial with particular properties, and then to solve this design problem using some elementary results from Section 4.

Let $\mathbf{G} \sim \mathcal{P}$, and assume we had a viable pseudoexpectation $\tilde{\mathbb{E}}$ for the $\text{LoSt}(2, m)$ SDP. Write $X \succeq 0$ for the $nk \times nk$ matrix whose $(u, i), (v, j)$ entry is $\tilde{\mathbb{E}} x_{u,i} x_{v,j}$ (it is routine that $\tilde{\mathbb{E}} \succeq 0$ implies positive semidefiniteness of X). It will at times be useful to think of X as a $k \times k$ matrix of $n \times n$ blocks $X_{i,j}$, and at others as an $n \times n$ matrix of $k \times k$ blocks $X_{u,v}$. Recall also the matrices $A_{\mathbf{G}}^{(s)}$ from Section 4 that count self-avoiding walks of length s . Our strategy will be to first write the moment-matching constraints on $\tilde{\mathbb{E}}$ as affine constraints of the form $\langle X_{i,j}, Y \rangle \simeq C$, and then combine these affine constraints to contradict feasibility of X .

Lemma 5.1. *For any i, j , and any $s = 0, \dots, m$, recalling that $A_{\mathbf{G}}^{(s)}$ is the matrix counting non-backtracking walks of length s , and \mathbb{J} is the all-ones matrix,*

$$\begin{aligned} \langle X_{i,j}, A_{\mathbf{G}}^{(s)} \rangle &\simeq \pi(i) M_{i,j}^s \|q_s\|_{KM}^2 n \\ \langle X_{i,j}, \mathbb{J} \rangle &= \pi(i)\pi(j)n^2. \end{aligned}$$

Proof. For the first assertion, let (H, S, τ) be the path of length s whose endpoints are labelled $i, j \in [k]$. Each *self-avoiding* walk of length s in \mathbf{G} is an occurrence of H , so from Theorem 4.6

$$\langle X_{i,j}, A_{\mathbf{G}}^{(s)} \rangle = \tilde{\mathbb{E}} p_{H,S,\tau}(\mathbf{x}, \mathbf{G}) \simeq \pi(i) M_{i,j}^s \|q_s\|_{KM}^2.$$

We can now use Lemma 4.3 to replace the self-avoiding walk matrices $A_{\mathbf{G}}^{(s)}$ with their non-backtracking counterparts. The matrix X has diagonal elements $X_{(u,i),(u,i)} = \tilde{\mathbb{E}} x_{u,i}^2 = \tilde{\mathbb{E}} x_{i,u}$ by the Boolean constraint, and $\tilde{\mathbb{E}} (x_{u,1} + \dots + x_{u,k}) = 1$ by the Single Color constraint. By PSD-ness of X , every $\tilde{\mathbb{E}} x_{u,i}^2 = \tilde{\mathbb{E}} x_{u,i}$ is nonnegative, so each is between zero and one. It is a standard fact that the off-diagonal entries of such a PSD matrix have magnitude at most one, so from Lemma 4.1

$$\langle X_{i,j}, A_{\mathbf{G}}^{(s)} \rangle = \langle X_{i,j}, A_{\mathbf{G}}^{(s)} \rangle + \langle X_{i,j}, A_{\mathbf{G}}^{(s)} - A_{\mathbf{G}}^{(s)} \rangle = \langle X_{i,j}, A_{\mathbf{G}}^{(s)} \rangle \pm O(\log n) \simeq \pi(i) M_{i,j}^s \|q_s\|_{KM}^2$$

for $s = 0, \dots, m$. For the second assertion, when $i \neq j$ take (H, S, τ) to be the partially labelled graph on two disconnected vertices, with labels i and j respectively. From Remark 4.7 we have

$$\langle X_{i,j}, \mathbb{J} \rangle = \tilde{\mathbb{E}} p_{H,S,\tau}(\mathbf{x}, \mathbf{G}) = \pi(i)\pi(j)n^2.$$

When $i = j$, take (H, S, τ) as above and (H', S', τ') to be a single vertex labelled i . □

We will now apply a fortuitous change of basis furnished to us by the parameter matrix M . Recall that F is the matrix whose columns are the right eigenvectors of M , satisfying $MF = F\Lambda$ and $F^T \text{Diag}(\pi)F = \mathbb{1}$. Now define a matrix $\check{X} \triangleq (F^T \otimes \mathbb{1})X(F \otimes \mathbb{1})$, by which we mean that

$$\check{X} = \begin{pmatrix} F_{1,1}\mathbb{1} & \cdots & F_{1,k}\mathbb{1} \\ \vdots & \ddots & \vdots \\ F_{k,1}\mathbb{1} & \cdots & F_{k,k}\mathbb{1} \end{pmatrix} \begin{pmatrix} X_{1,1} & \cdots & X_{1,k} \\ \vdots & \ddots & \vdots \\ X_{k,1} & \cdots & X_{k,k} \end{pmatrix} \begin{pmatrix} F_{1,1}\mathbb{1} & \cdots & F_{1,k}\mathbb{1} \\ \vdots & \ddots & \vdots \\ F_{k,1}\mathbb{1} & \cdots & F_{k,k}\mathbb{1} \end{pmatrix}.$$

We will think of \check{X} , analogous to X , as a $k \times k$ matrix of $n \times n$ blocks $\check{X}_{i,j}$. Note that we can also think of this as a change of basis $x \mapsto F^T x$ directly on the variables appearing in polynomials accepted by our pseudoexpectation.

Lemma 5.2. *For any $s = 0, \dots, m$, if $i \neq j$ $\langle \check{X}_{i,j}, \mathcal{A}_{\mathbf{G}}^{(s)} \rangle \simeq 0$, and*

$$\langle \check{X}_{i,i}, \mathcal{A}_{\mathbf{G}}^{(s)} \rangle \simeq \lambda_i^s \|q_s\|_k^2 n.$$

Furthermore,

$$\langle \check{X}_{i,j}, \mathbb{J} \rangle = \begin{cases} n^2 & i = j = 1 \\ 0 & \text{else} \end{cases}.$$

Proof. Our block-wise change of basis commutes with taking inner products between the blocks $X_{i,j}$ and the non-backtracking walk matrices. In other words,

$$\begin{aligned} \begin{pmatrix} \langle \check{X}_{1,1}, \mathcal{A}_{\mathbf{G}}^{(s)} \rangle & \cdots & \langle \check{X}_{1,k}, \mathcal{A}_{\mathbf{G}}^{(s)} \rangle \\ \vdots & \ddots & \vdots \\ \langle \check{X}_{k,1}, \mathcal{A}_{\mathbf{G}}^{(s)} \rangle & \cdots & \langle \check{X}_{k,k}, \mathcal{A}_{\mathbf{G}}^{(s)} \rangle \end{pmatrix} &= F^T \begin{pmatrix} \langle X_{1,1}, \mathcal{A}_{\mathbf{G}}^{(s)} \rangle & \cdots & \langle X_{1,k}, \mathcal{A}_{\mathbf{G}}^{(s)} \rangle \\ \vdots & \ddots & \vdots \\ \langle X_{k,1}, \mathcal{A}_{\mathbf{G}}^{(s)} \rangle & \cdots & \langle X_{k,k}, \mathcal{A}_{\mathbf{G}}^{(s)} \rangle \end{pmatrix} F \\ &\simeq F^T \text{Diag}(\pi) M^s F \cdot \|q_s\|_{KM}^s n \\ &= F^T \text{Diag}(\pi) F \Lambda^s \cdot \|q_s\|_{KM}^s n \\ &= \Lambda^s \cdot \|q_s\|_{KM}^s n \end{aligned}$$

A parallel calculation gives us

$$\begin{aligned} \begin{pmatrix} \langle \check{X}_{1,1}, \mathbb{J} \rangle & \cdots & \langle \check{X}_{1,k}, \mathbb{J} \rangle \\ \vdots & \ddots & \vdots \\ \langle \check{X}_{k,1}, \mathbb{J} \rangle & \cdots & \langle \check{X}_{k,k}, \mathbb{J} \rangle \end{pmatrix} &= F^T \begin{pmatrix} \langle X_{1,1}, \mathbb{J} \rangle & \cdots & \langle X_{1,k}, \mathbb{J} \rangle \\ \vdots & \ddots & \vdots \\ \langle X_{k,1}, \mathbb{J} \rangle & \cdots & \langle X_{k,k}, \mathbb{J} \rangle \end{pmatrix} F \\ &= F^T \pi \pi^T F \cdot n^2 \\ &= e_1 e_1^T n^2, \end{aligned}$$

where e_1 is the first standard basis vector. The final line comes since π , being the left eigenvector associated to $\lambda_1 = 1$, is (up to scaling) the first row of F^{-1} . \square

The remainder of the proof will amount to combining the constraints on the diagonal blocks of \check{X} . As X is PSD, \check{X} is as well, so any PSD linear combination $0 \preceq c_0 \mathbb{1} + \cdots + c_m \mathcal{A}_{\mathbf{G}}^{(s)}$ must satisfy

$$0 \leq \frac{1}{n} \left\langle \sum_{s=0}^m c_s \mathcal{A}_{\mathbf{G}}^{(s)}, \check{X}_{i,i} \right\rangle \simeq \sum_{s=0}^m c_s \lambda_i^s \|q_s\|_{KM}^2.$$

We can show that no \check{X} satisfying the given constraints, and thus that the SDP is infeasible, by producing such constants c_s as to make the right hand side of the above equation negative for at least one of $\lambda_1, \dots, \lambda_k$. Notice also

$$\sum_{s=0}^m c_s \mathcal{A}_{\mathbf{G}}^{(s)} = \sum_{s=0}^m c_s q_s(\mathcal{A}_{\mathbf{G}}) \triangleq f(\mathcal{A}_{\mathbf{G}})$$

for some polynomial $f \in \mathbb{R}[x]$ of degree m . Because $f(A_G)$ is a scalar polynomial in A_G , its eigenvalues are f applied to those of A_G , and we get $f(A_G) \succeq 0$ when f is nonnegative on $\text{Spec } A_G$. By Friedman's Theorem [Fri08], this spectrum consists of the 'trivial' eigenvalue d , together with $n - 1$ remaining eigenvalues whose magnitudes with high probability are at most $2\sqrt{d-1} + \eta$ for any $\eta > 0$. In fact, it is not necessary even that $f(d) > 0$. To see this, note that from our discussion above,

$$\langle f(A_G - \frac{d}{n}J), \check{X}_{i,i} \rangle = \langle f(A_G) - \frac{f(d)}{n}J, \check{X}_{i,i} \rangle = \langle f(A_G), \check{X}_{i,i} \rangle$$

for $i = 2, \dots, k$. Since $A_G - \frac{d}{n}J$ is the projection of A away from the eigenspace corresponding to $d \in \text{Spec } A_G$, $f(A_G - \frac{d}{n}J) \succeq 0$ whenever f is positive on the remainder of the spectrum.

From our discussion Section 4, for any $\lambda \in [-1, 1]$,

$$\begin{aligned} \sum_{s=0}^m c_s \lambda^2 \|q_s\|_{KM}^2 &= \sum_{s=0}^m \lambda^s \frac{\langle f, q_s \rangle_{KM}}{\|q_s\|_{KM}^2} \|q_s\|_{KM}^2 \\ &= \langle f, \sum_{s=0}^m \lambda^s q_s \rangle_{KM} \end{aligned}$$

Thus we've reduced the proof to the construction of a degree m polynomial f , nonnegative on $[-2\sqrt{d-1} - \eta, 2\sqrt{d-1} + \eta]$ for some $\eta > 0$ and satisfying $\langle f, \sum_{s=0}^m \lambda^s q_s \rangle_{KM} < 0$. The following extremely simple choice of f will finish things up.

Call m' the largest even number less than or equal to m , let $\varepsilon > 0$ be a very small number (unrelated to the distance ε of our model to the KS-threshold) and take

$$f(z) = -q_{m'}(z) + 2m' \|q_{m'}\|_{KM} + \varepsilon$$

which by Lemma 4.1 has the desired positivity property. This choice of f satisfies

$$\langle f(z), \sum_{s=0}^m \lambda^s q_s \rangle = -\|q_{m'}\|_{KM}^2 |\lambda|^{m'} + 2m' \|q_{m'}\|_{KM} + \varepsilon,$$

which is negative when

$$|\lambda| > \left(\frac{2m' + \varepsilon / \|q_{m'}\|_{KM}}{\|q_{m'}\|_{KM}} \right)^{1/m'} \rightarrow_m \frac{1}{\sqrt{d-1}}.$$

Remark 5.3. We can choose constants a and b such that the LocalStatistic SDP (3.2) is infeasible on G drawn from \mathcal{N} if we set the distance from the Kesten-Stigum bound ε and global error tolerance δ as $(\varepsilon, \delta) = (a, b)$, and also if we choose these as $(\varepsilon, \delta) = (a, 2b)$. In particular, this means that when $\delta = b$, for any PSD matrix X with an all-ones diagonal, there is a polynomial f such that the constraint

$$\langle f(A_G), X \rangle = \|q_s\|_{KM}^2 \lambda^s n \pm \delta n$$

is violated by a margin of $\Omega(n)$.

6 Interlude: Spectral Distinguishing

Our argument in the previous section can be recast to prove Corollary 2.4, namely that above the Kesten-Stigum threshold the spectrum of the adjacency matrix can also be used to distinguish the null and planted distributions.

Let $(\mathbf{x}, \mathbf{G}) \sim \mathcal{P}_{d,k,M,\pi}$, and write $\mathbf{X} \triangleq \mathbf{x}\mathbf{x}^\top$, and

$$\check{\mathbf{X}} = (\mathbf{F}^\top \otimes \mathbb{1})\mathbf{X}(\mathbf{F} \otimes \mathbb{1}) = (\mathbf{F}^\top \mathbf{x})(\mathbf{F}^\top \mathbf{x})^\top \triangleq \check{\mathbf{x}}\check{\mathbf{x}}^\top.$$

Think of $\check{\mathbf{X}}$ as a block matrix $(\check{\mathbf{X}}_{i,j})_{i,j \in [k]}$, as we did \mathbf{X} in the previous section, and $\check{\mathbf{x}}$ as a block vector $(\check{\mathbf{x}}_i)_{i \in [k]}$. Applying Theorem 4.6 and repeating the calculations in Lemmas 5.1 and 5.2 *mutatis mutandis* with $\check{\mathbf{X}}$ instead of \mathbf{X} , we can show that w.h.p.

$$\langle \check{\mathbf{X}}_{i,j}, \mathbf{A}_{\mathbf{G}}^{(s)} \rangle \simeq \lambda_i \|\mathbf{q}_s\|_{\mathbf{K}_M}^2 n \text{ if } i = j$$

and zero otherwise, for every $s = O(1)$ and

$$\langle \check{\mathbf{X}}_{1,1}, \mathbb{J} \rangle = \begin{cases} n^2 & i = j = 1 \\ 0 & \text{else} \end{cases}.$$

Because $\mathbf{A}_{\mathbf{G}}^{(s)} = \mathbb{1}$, we know

$$\check{\mathbf{x}}_i^\top \check{\mathbf{x}}_j = \langle \check{\mathbf{X}}_{i,j}, \mathbb{1} \rangle = 0$$

when $i \neq j$. In other words, the k vectors $\check{\mathbf{x}}_1, \dots, \check{\mathbf{x}}_k$ are orthogonal.

Now let $d > d_{\text{KS}} + \epsilon$. We can show that $\mathbf{A}_{\mathbf{G}}$ has an eigenvalue with a separation $\eta > 0$ from the bulk spectrum by proving

$$\check{\mathbf{x}}_i^\top f(\mathbf{A}_{\mathbf{G}}) \check{\mathbf{x}}_i = \langle \check{\mathbf{X}}_{i,i}, f(\mathbf{A}_{\mathbf{G}}) \rangle < 0$$

for some polynomial $f(x)$ positive on $(-2\sqrt{d-1}-\eta, 2\sqrt{d-1}+\eta)$. The same polynomial from Section 5 works here. As the $\check{\mathbf{x}}_i$ are orthogonal, we get one distinct eigenvalue outside the bulk for each eigenvalue of M for which $d > 1/\lambda_i^2 + 1 + \epsilon$.

Remark 6.1. To distinguish the null model from the planted one using the spectrum of $\mathbf{A}_{\mathbf{G}}$, simply return `PLANTED` if $\mathbf{A}_{\mathbf{G}}$ has a single eigenvalue other than d whose magnitude is bigger than $2\sqrt{d-1} + \delta$ for any error tolerance δ you choose, and `NULL` otherwise. Unfortunately, this distinguishing algorithm is not robust to adversarial edge insertions and deletions. For instance, given a graph $\mathbf{G} \sim \mathcal{N}$, the adversary can create a disjoint copy of \mathbf{K}_{d+1} , the complete graph on $d+1$ vertices, whose eigenvalues are all $\pm d$. The spectrum of the perturbed graph is the disjoint union of $\pm d$ and the eigenvalues of the other component(s), so the algorithm will be fooled. We will show in Section 8 that the Local Statistics SDP is robust to this kind of perturbation.

7 Lower Bounds Against Local Statistics SDPs

In this section, we prove the complementary bound to Theorem 2.2, namely that if every one of $\lambda_2, \dots, \lambda_k$ has modulus at most $1/\sqrt{d-1}$ there exists some feasible solution to the Local Path Statistics SDP for every $m \geq 1$. We can specify a pseudoexpectation completely by way of an $(nk+1) \times (nk+1)$ positive semidefinite matrix

$$\begin{pmatrix} 1 & \tilde{\mathbb{E}} \mathbf{x}^\top \\ \tilde{\mathbb{E}} \mathbf{x} & \tilde{\mathbb{E}} \mathbf{x}^\top \mathbf{x} \end{pmatrix} \triangleq \begin{pmatrix} 1 & \mathbf{l}^\top \\ \mathbf{l} & \mathbf{X} \end{pmatrix}.$$

After first writing down the general properties required of *any* quadratic pseudoexpectation satisfying \mathcal{B}_k , we'll show that in order for $\tilde{\mathbb{E}}$ to match every moment asked of it by the $\text{LoSt}(2, m)$ SDP, it suffices for it to satisfy

$$\tilde{\mathbb{E}} p_{H,S,\tau}(\mathbf{x}, \mathbf{G}) \simeq \mathbb{E} p_{H,S,\tau}(\mathbf{x}, \mathbf{G})$$

when (H, S, τ) is a path of length $0, \dots, m$ with labelled endpoints. Finally, we'll construct a pseudoexpectation matching these path moments by way of some elementary properties of the non-backtracking polynomials from Section 4.

Lemma 7.1. *The set of \mathcal{B}_k -satisfying pseudoexpectations is parameterized by pairs $(X, \mathbf{l}) \in \mathbb{R}^{nk \times nk} \times \mathbb{R}^{nk}$ for which*

$$\begin{pmatrix} 1 & \mathbf{l}^\top \\ \mathbf{l} & X \end{pmatrix} \succeq 0 \quad (8)$$

$$\text{diag}(X) = \mathbf{l} \quad (9)$$

$$\text{tr } X_{u,u} = \mathbf{e}^\top \mathbf{l} = 1 \quad \forall u \in [n] \quad (10)$$

$$X_{u,v} \mathbf{e} = \mathbf{l}_u \quad \forall u, v \in [n] \quad (11)$$

Proof. Recall that the set \mathcal{B}_k is defined by the polynomial equations

$$\begin{array}{lll} \text{Boolean} & x_{u,i}^2 = x_{u,i} & \forall u \in [n] \text{ and } i \in [k] \\ \text{Single Color} & \sum_i x_{u,i} = 1 & \forall u \in [n] \end{array}$$

That a degree-two pseudoexpectation *satisfies* these constraints means

$$\begin{aligned} \tilde{\mathbb{E}} p(x) x_{u,i}^2 &= \tilde{\mathbb{E}} p(x) x_{u,i} & \forall p \text{ s.t. } \deg p = 0 \\ \tilde{\mathbb{E}} p(x) \sum_i x_{u,i} &= \tilde{\mathbb{E}} p(x) & \forall p \text{ s.t. } \deg p \leq 1. \end{aligned}$$

Writing $X = \tilde{\mathbb{E}} x^\top x$ and $\mathbf{l} = \tilde{\mathbb{E}} x$ as above, the first constraint is equivalent to $\mathbf{l} = \text{diag}(X)$, since the degree-zero polynomials are just constants, and we can guarantee that the second holds for every polynomial of degree at most one by requiring it on $p = 1$ and $p = x_{v,j}$ for all v and j . The Lemma is simply a concise packaging of these facts, using the block notation $X = (X_{u,v})_{u,v \in [n]}$ and $\mathbf{l} = (\mathbf{l}_u)_{u \in [n]}$. \square

Proposition 7.2. *It suffices to check*

$$\tilde{\mathbb{E}} p_{H,S,\tau}(x, \mathbf{G}) \simeq \mathbb{E} p_{H,S,\tau}(x, \mathbf{G})$$

in the cases (i) (H, S, τ) is a path of length $s = 0, \dots, m$ with labelled endpoints, and (ii) when (H, S, τ) is a graph with no edges on one or two labelled vertices.

We will defer the proof of Proposition 7.2 to Appendix C. Its conclusion in hand, we can now set about constructing a pseudoexpectation. We'll construct $\mathbf{l} \in \mathbb{R}^{nk}$ and $X \in \mathbb{R}^{nk \times nk}$ so that (i) the \mathcal{B}_k constraints in Lemma 7.1 hold, and (ii)

$$\begin{aligned} \langle \mathbf{e}, \mathbf{l}_i \rangle &= \pi(i)n \\ \langle X_{i,j}, \mathbf{A}_{\mathbf{G}}^{(s)} \rangle &\simeq \pi(i) M_{i,j}^s n \\ \langle X_{i,j}, \mathbb{J} \rangle &= \pi(i)\pi(j)n^2. \end{aligned}$$

It will simplify things immensely to use the same change of basis as we did in Section 5. Namely, letting F be the matrix of right eigenvectors, we will produce a pair $\check{\mathbf{l}} \in \mathbb{R}^{nk}$ and $\check{X} \in \mathbb{R}^{nk \times nk}$ so that $\mathbf{l} = (F^{-\top} \otimes \mathbb{1}) \check{\mathbf{l}}$

and $X = (F^{-T} \otimes \mathbb{1})\check{X}(F^{-1} \otimes \mathbb{1})$ satisfy the above conditions. Recycling the relevant calculations from Section 5, the above moment conditions translate to

$$\begin{aligned}\langle e, \check{l}_i \rangle &= \begin{cases} n & i = 1 \\ 0 & \text{else} \end{cases} \\ \langle \check{X}_{i,j}, A_G^{(s)} \rangle &\simeq \lambda_i^s \|q_s\|_{KM}^2 n \\ \langle \check{X}_{i,j}, \mathbb{J} \rangle &= \begin{cases} n^2 & i = j = 1 \\ 0 & \text{else} \end{cases}\end{aligned}$$

The key steps in designing \check{X} are as follows.

Proposition 7.3. *For every ϵ , every m , and every λ such that $|\lambda|^2(d-1) < 1 - \epsilon$, there exists a polynomial y nonnegative on $(-2\sqrt{d-1}, 2\sqrt{d-1})$ and satisfying*

$$\langle q_s, y \rangle_{KM} = \lambda^s \|q_s\|_{KM}^2.$$

Proposition 7.4. *Let $G \sim \mathcal{N}$. If there exists a polynomial y meeting the conditions of Proposition 7.3 for some $\lambda \in (-1, 1)$, then there exists a PSD matrix $Y(\lambda)$ for which*

$$\begin{aligned}Y(\lambda)_{u,u} &= 1 & \forall u \in [n] \\ \langle Y(\lambda), A_G^{(s)} \rangle &\simeq \lambda^s \|q_s\|_{KM}^2 n & \forall s \in [m] \\ \langle Y(\lambda), \mathbb{J} \rangle &= 0\end{aligned}$$

With these propositions in hand, define \check{X} to be the $k \times k$ block diagonal matrix

$$\check{X} = \begin{pmatrix} \mathbb{J} & & & \\ & Y(\lambda_2) & & \\ & & \ddots & \\ & & & Y(\lambda_k), \end{pmatrix}$$

i.e. $\check{X}_{i,j} = 0$ when $i \neq j$, and the diagonal blocks are as above, and similarly let $\check{l} = (e, 0, \dots, 0)^T$. This way, certainly

$$\begin{pmatrix} 1 & \check{l}^T \\ \check{l} & \check{X} \end{pmatrix} \succeq 0 \tag{12}$$

(by taking a Schur complement), and the three inner product conditions above are satisfied on every block. We now need to check carefully that

$$\begin{pmatrix} 1 & \check{l}^T \\ \check{l} & X \end{pmatrix} \triangleq \begin{pmatrix} 1 & & \\ & F^{-T} \otimes \mathbb{1} & \\ & & \end{pmatrix} \begin{pmatrix} 1 & \check{l}^T \\ \check{l} & \check{X} \end{pmatrix} \begin{pmatrix} 1 & & \\ & F^{-1} \otimes \mathbb{1} & \\ & & \end{pmatrix}$$

is a pseudoexpectation satisfying \mathcal{B}_k . The above construction guarantees PSD-ness, since we have multiplied a matrix and its transpose on the right and left respectively of a PSD matrix. Since π is the first row of F^{-1} ,

we know $\mathbf{l}_i = \pi(i)\mathbf{e}$. On the other hand, X is obtained by changing basis block-wise, the diagonal of X depends only on the diagonals of \mathbb{J} and the $Y(\lambda_i)$, all of which are all ones, so

$$\begin{aligned}\text{diag } X &= \text{diag} \left((F^{-T} \otimes \mathbb{1}) \text{Diag } \text{diag } \check{X}(F^{-1} \otimes \mathbb{1}) \right) \\ &= \text{diag} \left((F^{-T} \otimes \mathbb{1})(F^{-1} \otimes \mathbb{1}) \right) \\ &= \text{diag} \left(F^{-T}F^{-1} \otimes \mathbb{1} \right) \\ &= \text{diag} \left(\text{Diag } \pi \otimes \mathbb{1} \right) \\ &= (\pi(1)\mathbf{e}, \dots, \pi(k)\mathbf{e})\end{aligned}$$

as desired. Similarly, because \check{X} is diagonal, $\check{X}_{u,u} = \mathbb{1}$, and

$$\text{tr } X_{u,u} = \text{tr } F^{-T} \check{X}_{u,u} F^{-1} = \text{tr } F^{-T} F^{-1} = \text{tr } \text{Diag } \pi = 1.$$

Finally, the top row of each $\check{X}_{u,v}$ is the vector \mathbf{e}_1^T , so

$$X_{u,v}\mathbf{e} = F^{-T} \check{X}_{u,v} F^{-1} \mathbf{e} = F^{-T} \check{X}_{u,v} \mathbf{e}_1 = F^{-T} \mathbf{e}_1 = \pi = \mathbf{l}_u.$$

This completes the construction of our pseudoexpectation.

Remark 7.5. By correctly setting the ϵ from Proposition 7.3 and the global error tolerance δ , we can once again choose $(\epsilon, \delta) = (a, b)$ with the property that whenever the $\text{LoSt}(2, m)$ SDP is feasible $\delta = b$, it is feasible as well with $\delta = 2b$. Thus every one of the \simeq constraints—i.e. those that depend on the observed graph \mathbf{G} —are satisfied with slack $\Omega(n)$.

Proof of Proposition 7.3. Such a polynomial y is exactly of the form

$$y = \sum_{s=0}^m \lambda^s q_s + \text{terms with larger } q_s\text{'s}.$$

We will use the extremely simple construction of letting the coefficients on the terms q_{m+1}, q_{m+1}, \dots also be powers of λ . The idea here is that, whenever $|\lambda|^2(d-1) < 1$, this series converges to a positive function on $(-2\sqrt{d-1}, 2\sqrt{d-1})$, so by taking a long enough initial segment, we can get a positive approximant.

In particular, let $p \gg m$ be even, and set

$$y = \sum_{s=0}^p \lambda^s q_s.$$

It is a standard calculation, employing the recurrence relation on the polynomials q_s , that

$$y(z) = \frac{1 - \lambda^2 + \lambda^{p+2}(d-1)q_p(z) - \lambda^{p+1}q_{p+1}(z)}{(d-1)\lambda^2 - \lambda z + 1}.$$

One can quickly verify that

$$\frac{1 - \lambda^2}{(d-1)\lambda^2 - \lambda z + 1} > 0 \quad \forall |z| \leq 2\sqrt{d-1},$$

so all we need to verify is that $\lambda^2(d-1) < 1$ ensures $\lambda^{p+2}(d-1)q_p - \lambda^{p+1}q_{p+1} \rightarrow_p 0$. This follows immediately from Lemma 4.1, as $|q_p| \leq 2p\sqrt{d-1}^p$. \square

Proof of Proposition 7.4. Let y be the polynomial guaranteed in the theorem statement; our strategy will be to modify the matrix $y(A_G)$. First note that by expanding y in the q_s basis, we have

$$y(A_G) = \sum_{s=0}^m q_s(A_G) \lambda^s + \dots = \sum_{s=0}^m A_G^{(s)} \lambda^s + \dots,$$

so it is clear that $y(A_G)$ satisfies the affine constraints against the $A_G^{(s)}$ matrices. Moreover, as y is strictly positive on $[-2\sqrt{d-1}, 2\sqrt{d-1}]$, it is (by continuity) nonnegative on a constant size fattening of this interval, and by Friedman's theorem the spectrum of A_G other than the eigenvalue at d is contained w.h.p. in such a set. Thus $y(A_G)$ is positive, except perhaps the eigenvalue $y(d)$, which we will fix in a moment.

However, $y(A_G)$ does not have the right inner product with the all ones matrix, and—unless $2 \deg y + 1 < \text{girth}(G)$ —its diagonal entries need not be ones. Our corrective to these issues will exploit two fortunate facts. First, those diagonal entries different from one are exactly those within $\deg y$ steps of a constant length cycle; from Lemma 4.2 we know that there are at most $O(\log n)$ of these. We'll keep the terminology from that Lemma, calling such vertices *bad* and the remaining ones *good*. Second, as y is a scalar polynomial and \mathbb{J} commutes with A_G ,

$$\langle y(A_G), \mathbb{J}/n \rangle = y(d) = O_n(1).$$

Thus we can correct the inner product with \mathbb{J} , and at the same time resolve the possible negativity of the eigenvalue $y(d)$, by passing to

$$\begin{aligned} \tilde{Y}(\lambda) &= \frac{1}{1 - y(d)/n} (\mathbb{1} - \mathbb{J}/n) y(A_G) (\mathbb{1} - \mathbb{J}/n) \\ &= \frac{1}{1 - y(d)/n} (y(A_G) - y(d)\mathbb{J}/n); \end{aligned}$$

since $\langle A_s^{(G)}, \mathbb{J}/n \rangle = q_s(d) = O(1)$ the result will still satisfy the inner product constraints with the matrices $A_s^{(G)}$ up to an additive $\pm \delta n$. This new matrix is certainly PSD, for instance by writing out $y(A_G)$ in its eigenvalue basis, and observing that left and right multiplication by $(\mathbb{1} - \mathbb{J}/n)$ simply projects away the eigenspace of the $y(d)$ eigenvalue. Thus we can write the $\tilde{Y}(\lambda)_{u,v} = \alpha_u^T \alpha_v$ for some vectors $\alpha_1, \dots, \alpha_n \in \mathbb{R}^n$. The scale factor we applied above makes sure that for every good vertex u , $\|\alpha_u\| = 1$, and being orthogonal to the all-ones matrix is equivalent to $\sum_u \alpha_u = 0$.

The remaining diagonal elements are at worst some constant C dependent on d and y , since the diagonal entries of each $A_G^{(s)}$ are all $O(1)$. Thus, writing Γ for the set of good vertices, we know

$$\left\| \sum_{u \in \Gamma} \alpha_u \right\| = \left\| \sum_{u \notin \Gamma} \alpha_u \right\| \leq C \log n$$

It is clear that by removing at most $C \log n$ vertices from Γ to create a new set Γ' we can choose a collection of unit vectors β_u for each $u \in \Gamma'$ so that

$$\sum_{u \notin \Gamma'} \beta_u = \sum_{u \in \Gamma'} \alpha_u.$$

Our final matrix $Y(\lambda)$ will be the Gram matrix of these new β and remaining α vectors. We must finally check that the affine constraints against the $A_G^{(s)}$ matrices are still approximately satisfied. However, even

starting from a bad vertex, there are at most a constant number of vertices within s steps of it, and at most a constant number of non-backtracking walks to any such vertex. Thus

$$\begin{aligned} \left| \langle Y(\lambda), \mathbf{A}_{\mathbf{G}}^{(s)} \rangle - \langle \tilde{Y}(\lambda), \mathbf{A}_{\mathbf{G}}^{(s)} \rangle \right| &= \left| 2 \sum_{u \notin \Gamma', v \in \Gamma'} (\mathbf{A}_{\mathbf{G}}^{(s)})_{u,v} \alpha_u^T (\alpha_v - \beta_v) + \sum_{u,v \notin \Gamma'} (\mathbf{A}_{\mathbf{G}}^{(s)})_{u,u} (\|\alpha_u\| - \|\beta_u\|) \right| \\ &= O(\log n) \end{aligned}$$

where we have used that $\max_u \|\alpha_u\| = O(1)$ and broken up both summations by first enumerating the $O(\log n)$ vertices in U' and then the at most $O(1)$ vertices in its depth s neighborhood. \square

8 Robustness Guarantees

In this section, again fix a planted model \mathcal{P} with parameters (d, k, M, π) , and let $\kappa(n)$ be some slowly growing function of n . Assume that we observe a graph $\tilde{\mathbf{H}}$ on n vertices, which we are promised was drawn from one of \mathcal{N} or \mathcal{P} and then corrupted by $\kappa(n)$ adversarial edge insertions or deletions. Our goal is to decide, upon seeing $\tilde{\mathbf{H}}$, from which model the unperturbed graph \mathbf{G} was sampled. We present an algorithm below that works for an appropriate regime of κ .

Algorithm 8.1. Given a graph $\tilde{\mathbf{H}}$ and $m \in \mathbb{N}$ as input, do the following. Delete all edges incident to vertices that have degree greater than d in $\tilde{\mathbf{H}}$, and then greedily add edges connecting any vertices with degree less than d to obtain a d -regular graph \mathbf{H} . Run the distinguishing SDP (3.2) at level m on \mathbf{H} and output `NULL` if the SDP is infeasible, and `PLANTED` otherwise.

Theorem 8.2. *Let δ be any positive constant. Supposing $d > d_{\text{KS}} + \epsilon$, there exist $\rho > 0$ satisfying $\kappa(n) \leq \rho n$, and $m \in \mathbb{N}$, such that Algorithm 8.1, on input $\tilde{\mathbf{H}}$ and m , correctly distinguishes whether \mathbf{G} was drawn from \mathcal{N} , or from \mathcal{P} with probability $1 - o(1)$.*

Proof. Note that \mathbf{H} is can be obtained by taking \mathbf{G} and making up to $\xi \rho n$ edge insertions and deletions for an absolute constant ξ . It suffices to show that (i) the SDP is feasible on \mathbf{H} as input if \mathbf{G} is drawn from the planted distribution, and (ii) the SDP is infeasible on \mathbf{H} as input if \mathbf{G} is drawn from the null distribution.

Call a vertex $v \in [n]$ *corrupted* if its $(m+1)$ -neighborhood in \mathbf{H} differs from its $(m+1)$ -neighborhood in \mathbf{G} . We begin by analyzing the difference $\mathbf{A}_{\mathbf{G}}^{(s)} - \mathbf{A}_{\mathbf{H}}^{(s)}$ for $s \in [m]$. Suppose v is not a corrupted vertex, then $\mathbf{A}_{\mathbf{G}}^{(s)}$ and $\mathbf{A}_{\mathbf{H}}^{(s)}$ agree on the v -th row and column, which means $(\mathbf{A}_{\mathbf{G}}^{(s)} - \mathbf{A}_{\mathbf{H}}^{(s)})_{v,-} = 0$. On the other hand, if v is a corrupted vertex,

$$\begin{aligned} \left\| (\mathbf{A}_{\mathbf{G}}^{(s)} - \mathbf{A}_{\mathbf{H}}^{(s)})_{v,-} \right\|_1 &\leq \left\| \mathbf{A}_s^{(\mathbf{G})} \right\|_1 + \left\| \mathbf{A}_s^{(\mathbf{H})} \right\|_1 \\ &\leq 2d(d-1)^{s-1} \end{aligned}$$

In particular, this means the entrywise 1-norm of $\mathbf{A}_s^{(\mathbf{G})} - \mathbf{A}_s^{(\mathbf{H})}$, is bounded by $2\xi \rho n \cdot 2d(d-1)^{\ell-1}$ since there are at most $2\xi \rho n$ corrupted vertices (i.e. if all corrupted edges had disjoint endpoints).

From Remark 7.5, if \mathbf{G} is drawn from the planted distribution, the matrices $Y(\lambda_i)$ are PSD and satisfy the affine constraints regarding inner products with the $\mathbf{A}_{\mathbf{G}}^{(s)}$ matrices with slack $\Omega(n)$. Every diagonal entry

of $Y(\lambda_i)$ is one, so by PSDness their off-diagonal entries have modulus at most one. Thus

$$\left| \langle \mathbf{A}_s^{(\mathbf{G})} Y(\lambda_i) \rangle - \langle \mathbf{A}_s^{(\mathbf{H})} \rangle \right| = \left| \langle \mathbf{A}_s^{(\mathbf{G})} - \mathbf{A}_s^{(\mathbf{H})} \rangle \right| \leq \left\| \mathbf{A}_s^{(\mathbf{G})} - \mathbf{A}_s^{(\mathbf{H})} \right\|_1 \leq 2\xi\rho d(d-1)^{s-1}.$$

Because of the $\Omega(n)$ slack, if we construct $Y(\lambda_i)$ from \mathbf{H} instead of \mathbf{G} , the constraints will *still* be satisfied for small enough ρ . We can then use these $Y(\lambda_i)$ to build the full feasible solution as before.

On the other hand, when \mathbf{G} is drawn from the null model, we noted in Remark 5.3 that any pseudoexpectation satisfying the Boolean and Single Color constraints violates some linear combination of the above affine constraints by a margin of $\Omega(n)$, and this constraint will still be violated for ϵ sufficiently small. \square

Remark 8.3. The parameter ρ controlling the number of adversarial edge insertions and deletions made to random input \mathbf{G} that the level- m LocalStatistic SDP can tolerate can be seen to decrease with m , which is indicative of a tradeoff between how close to the threshold an algorithm in this hierarchy works and how robust it is to perturbations.

Corollary 8.4. *Algorithm 8.1 correctly distinguishes between whether \mathbf{G} was drawn from the null distribution, or from the planted distribution, from input \mathbf{H} , with probability $1 - o(1)$ when $\kappa(n) = o(n)$.*

References

- [Abb17] Emmanuel Abbe, *Community detection and stochastic block models: recent developments*, The Journal of Machine Learning Research **18** (2017), no. 1, 6446–6531.
- [ABH16] Emmanuel Abbe, Afonso S Bandeira, and Georgina Hall, *Exact recovery in the stochastic block model*, IEEE Transactions on Information Theory **62** (2016), no. 1, 471–487.
- [ABLS07] Noga Alon, Itai Benjamini, Eyal Lubetzky, and Sasha Sodin, *Non-backtracking random walks mix faster*, Communications in Contemporary Mathematics **9** (2007), no. 04, 585–603.
- [AKJ18] Ahmed El Alaoui, Florent Krzakala, and Michael I Jordan, *Fundamental limits of detection in the spiked wigner model*, arXiv preprint arXiv:1806.09588 (2018).
- [AS12] Pranjali Awasthi and Or Sheffet, *Improved spectral-norm bounds for clustering*, Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, Springer, 2012, pp. 37–49.
- [AS15] Emmanuel Abbe and Colin Sandon, *Community detection in general stochastic block models: Fundamental limits and efficient algorithms for recovery*, 2015 IEEE 56th Annual Symposium on Foundations of Computer Science, IEEE, 2015, pp. 670–688.
- [BDG⁺16] Gerandy Brito, Ioana Dumitriu, Shirshendu Ganguly, Christopher Hoffman, and Linh V Tran, *Recovery and rigidity in a regular stochastic block model*, Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms, Society for Industrial and Applied Mathematics, 2016, pp. 1589–1601.
- [BKM17] Jess Banks, Robert Kleinberg, and Cristopher Moore, *The Lovász theta function for random regular graphs and community detection in the hard regime*, arXiv preprint arXiv:1705.01194 (2017).

- [BKM⁺19] Jean Barbier, Florent Krzakala, Nicolas Macris, Léo Miolane, and Lenka Zdeborová, *Optimal errors and phase transitions in high-dimensional generalized linear models*, Proceedings of the National Academy of Sciences **116** (2019), no. 12, 5451–5460.
- [BKW19] Afonso S Bandeira, Dmitriy Kunisky, and Alexander S Wein, *Computational hardness of certifying bounds on constrained pca problems*, arXiv preprint arXiv:1902.07324 (2019).
- [BLM15] Charles Bordenave, Marc Lelarge, and Laurent Massoulié, *Non-backtracking spectrum of random graphs: community detection and non-regular ramanujan graphs*, 2015 IEEE 56th Annual Symposium on Foundations of Computer Science, IEEE, 2015, pp. 1347–1357.
- [BS95] Avrim Blum and Joel Spencer, *Coloring random and semi-random k -colorable graphs*, Journal of Algorithms **19** (1995), no. 2, 204–234.
- [BSW01] Eli Ben-Sasson and Avi Wigderson, *Short proofs are narrowfiresolution made simple*, Journal of the ACM (JACM) **48** (2001), no. 2, 149–169.
- [BT19] Jess Banks and Luca Trevisan, *Vector colorings of random, ramanujan, and large-girth irregular graphs*, arXiv preprint arXiv:1907.02539 (2019).
- [CJSX14] Yudong Chen, Ali Jalali, Sujay Sanghavi, and Huan Xu, *Clustering partially observed graphs via convex optimization*, The Journal of Machine Learning Research **15** (2014), no. 1, 2213–2238.
- [CL⁺15] T Tony Cai, Xiaodong Li, et al., *Robust and computationally feasible community detection in the presence of arbitrary outlier nodes*, The Annals of Statistics **43** (2015), no. 3, 1027–1059.
- [CO04] Amin Coja-Oghlan, *Coloring semirandom graphs optimally*, International Colloquium on Automata, Languages, and Programming, Springer, 2004, pp. 383–395.
- [CO07] ———, *Solving np -hard semirandom graph problems in polynomial expected time*, Journal of Algorithms **62** (2007), no. 1, 19–46.
- [CSV17] Moses Charikar, Jacob Steinhardt, and Gregory Valiant, *Learning from untrusted data*, Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, ACM, 2017, pp. 47–60.
- [DKMZ11] 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** (2011), no. 6, 065701.
- [DMS⁺17] Amir Dembo, Andrea Montanari, Subhabrata Sen, et al., *Extremal cuts of sparse random graphs*, The Annals of Probability **45** (2017), no. 2, 1190–1217.
- [Fei02] Uriel Feige, *Relations between average case complexity and approximation complexity*, Proceedings of the thirty-fourth annual ACM symposium on Theory of computing, ACM, 2002, pp. 534–543.
- [FK00] Uriel Feige and Robert Krauthgamer, *Finding and certifying a large hidden clique in a semirandom graph*, Random Structures & Algorithms **16** (2000), no. 2, 195–208.
- [FK01] Uriel Feige and Joe Kilian, *Heuristics for semirandom graph problems*, Journal of Computer and System Sciences **63** (2001), no. 4, 639–671.

- [FP07] Delphine Féral and Sandrine Péché, *The largest eigenvalue of rank one deformation of large wigner matrices*, *Communications in mathematical physics* **272** (2007), no. 1, 185–228.
- [Fri08] Joel Friedman, *A proof of Alon’s second eigenvalue conjecture and related problems*.
- [Gri01] Dima Grigoriev, *Linear lower bound on degrees of positivstellensatz calculus proofs for the parity*, *Theoretical Computer Science* **259** (2001), no. 1, 613–622.
- [GV16] Olivier Guédon and Roman Vershynin, *Community detection in sparse networks via grothendieckfs inequality*, *Probability Theory and Related Fields* **165** (2016), no. 3-4, 1025–1049.
- [HS17] Samuel B Hopkins and David Steurer, *Efficient bayesian estimation from few samples: community detection and related problems*, *Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on*, IEEE, 2017, pp. 379–390.
- [HWX16] Bruce Hajek, Yihong Wu, and Jiaming Xu, *Achieving exact cluster recovery threshold via semidefinite programming*, *IEEE Transactions on Information Theory* **62** (2016), no. 5, 2788–2797.
- [KB19] Dmitriy Kunisky and Afonso S Bandeira, *A tight degree 4 sum-of-squares lower bound for the sherrington-kirkpatrick hamiltonian*, arXiv preprint arXiv:1907.11686 (2019).
- [KGR11] Ulugbek Kamilov, Vivek K Goyal, and Sundeep Rangan, *Optimal quantization for compressive sensing under message passing reconstruction*, 2011 IEEE International Symposium on Information Theory Proceedings, IEEE, 2011, pp. 459–463.
- [KK10] Amit Kumar and Ravindran Kannan, *Clustering with spectral norm and the k-means algorithm*, 2010 IEEE 51st Annual Symposium on Foundations of Computer Science, IEEE, 2010, pp. 299–308.
- [KV06] Michael Krivelevich and Dan Vilenchik, *Semirandom models as benchmarks for coloring algorithms*, 2006 Proceedings of the Third Workshop on Analytic Algorithmics and Combinatorics (ANALCO), SIAM, 2006, pp. 211–221.
- [Mas14] Laurent Massoulié, *Community detection thresholds and the weak ramanujan property*, Proceedings of the forty-sixth annual ACM symposium on Theory of computing, ACM, 2014, pp. 694–703.
- [MMV12] Konstantin Makarychev, Yury Makarychev, and Aravindan Vijayaraghavan, *Approximation algorithms for semi-random partitioning problems*, Proceedings of the forty-fourth annual ACM symposium on Theory of computing, ACM, 2012, pp. 367–384.
- [MMV16] ———, *Learning communities in the presence of errors*, Conference on Learning Theory, 2016, pp. 1258–1291.
- [MNS18] Elchanan Mossel, Joe Neeman, and Allan Sly, *A proof of the block model threshold conjecture*, *Combinatorica* **38** (2018), no. 3, 665–708.
- [Moi12] Ankur Moitra, *A singly-exponential time algorithm for computing nonnegative rank*, arXiv preprint arXiv:1205.0044 (2012).
- [MS15] Andrea Montanari and Subhabrata Sen, *Semidefinite programs on sparse random graphs and their application to community detection*, arXiv preprint arXiv:1504.05910 (2015).

- [NM14] MEJ Newman and Travis Martin, *Equitable random graphs*, Physical Review E **90** (2014), no. 5, 052824.
- [PWBM16] Amelia Perry, Alexander S Wein, Afonso S Bandeira, and Ankur Moitra, *Optimality and sub-optimality of pca for spiked random matrices and synchronization*, arXiv preprint arXiv:1609.05573 (2016).
- [Ran11] Sundeep Rangan, *Generalized approximate message passing for estimation with random linear mixing*, 2011 IEEE International Symposium on Information Theory Proceedings, IEEE, 2011, pp. 2168–2172.
- [SVC16] Jacob Steinhardt, Gregory Valiant, and Moses Charikar, *Avoiding imposters and delinquents: Adversarial crowdsourcing and peer prediction*, Advances in Neural Information Processing Systems, 2016, pp. 4439–4447.
- [ZK16] Lenka Zdeborová and Florent Krzakala, *Statistical physics of inference: Thresholds and algorithms*, Advances in Physics **65** (2016), no. 5, 453–552.

A Recovery

As discussed in the introduction, this paper will not settle fully the question of recovering the planted communities. However, we can at least reduce some key aspects of this problem to Conjecture 2.5 regarding the spectrum of A_G when $\mathbf{G} \sim \mathcal{P}_{(d,k,M,\pi)}$.

There are numerous ways to pose the recovery task, and as many metrics of success, but let us set ourselves the modest goal of, given \mathbf{G} drawn from a planted model with $\lambda_1^2, \dots, \lambda_\ell^2 > (d-1)^{-1}$ and knowledge of the parameters (d, k, M, π) , recovering a vector in \mathbb{R}^n with constant correlation to each of the vectors $\check{\mathbf{x}}_1, \dots, \check{\mathbf{x}}_\ell$ from the prior section. If $\ell = k$, we can use this and our knowledge of M to apply the change-of-basis F^{-1} and recover vectors correlated to the indicators $\mathbf{x}_1, \dots, \mathbf{x}_k$ for each of the k communities.

Our first claim is that, assuming Conjecture 2.5, the eigenvectors of A_G can be used to approximate the $\check{\mathbf{x}}_i$'s. In the prior section we showed that there exists a polynomial f strictly positive on $(-2\sqrt{d-1}, 2\sqrt{d-1}) \cup \{d\}$ with the property that

$$\check{\mathbf{x}}_i^\top f(A) \check{\mathbf{x}}_i < -\delta n$$

for some constant δ . Writing μ_1, \dots, μ_n for the eigenvalues of A_G and Π_1, \dots, Π_n for the orthogonal projectors onto their associated eigenspaces, we can expand this as

$$\begin{aligned} -\delta n &> \sum_{u \in [n]} f(\mu_u) \check{\mathbf{x}}_i^\top \Pi_u \check{\mathbf{x}}_i \\ &= \sum_{|\mu_u| < 2\sqrt{d-1}} f(\mu_u) \check{\mathbf{x}}_i^\top \Pi_u \check{\mathbf{x}}_i + \sum_{|\mu_u| \geq 2\sqrt{d-1}} f(\mu_u) \check{\mathbf{x}}_i^\top \Pi_u \check{\mathbf{x}}_i \\ &\geq \sum_{|\mu_u| \geq 2\sqrt{d-1}} f(\mu_u) \check{\mathbf{x}}_i^\top \Pi_u \check{\mathbf{x}}_i && f(x) \text{ positive on } (-2\sqrt{d-1}, 2\sqrt{d-1}) \\ &\geq \inf_{|x| \leq d} f(x) \cdot \check{\mathbf{x}}_i^\top \left(\sum_{|\mu_u| \geq 2\sqrt{d-1}} \Pi_u \right) \check{\mathbf{x}}_i. \end{aligned}$$

Thus, even if there are only constantly many eigenvectors outside the bulk, a (for instance) random vector in their span will have $O(n)$ correlation with each of the \check{x}_i 's.

In order to recover *robustly* we will lean on the results of Section 8. If we begin with \mathbf{G} from the planted model, perform ϵn adversarial edge insertion or deletions, and then run the SDP again, we showed that the old SDP solution will *still* be feasible. Thus, if we take \check{X} from the SDP run on the corrupted graph, we will still have

$$-\delta n > \langle f(\mathbf{A}_{\mathbf{G}}), \check{X}_{i,i} \rangle \geq \inf_{|x| \leq d} f(x) \cdot \left\langle \sum_{|\mu_u| \geq 2\sqrt{d-1}} \Pi_u, \check{X}_{i,i} \right\rangle,$$

so a, say, Gaussian vector with covariance $\check{X}_{i,i}$ will have constant correlation with the subspace spanned by the outside-the-bulk eigenvectors of $\mathbf{A}_{\mathbf{G}}$, the adjacency matrix of the *unperturbed* graph, which we showed above have the same correlation guarantee with the \check{x}_i 's.

B The Degree Regular Block Model

This appendix is devoted to several results on the DRBM including the first moment calculation from the first section and the expectation and concentration of non-backtracking walk counts between vertices of different types. We begin with some standard results on the asymptotics of various combinatorial quantities we'll encounter.

From Stirling's approximation

$$\sqrt{2\pi n} \exp(n \log n - n) \leq n! \leq \sqrt{4\pi n} \exp(n \log n - n),$$

and the identity $(2n-1)!! = \frac{(2n)!}{2^n n!}$, we immediately get

$$\exp(n \log 2n - n) \leq (2n-1)!! \leq 2 \exp(n \log 2n - n).$$

For some nonnegative vector $\alpha = (\alpha_1, \dots, \alpha_k)$ with $\sum \alpha_i = 1$, write $\binom{n}{\alpha n}$ for the multinomial coefficient enumerating the ways to divide n into sets of size $\alpha_1 n, \dots, \alpha_k n$. Then

$$\binom{n}{\alpha n} = \frac{n!}{\prod_i (\alpha_i n)!} \leq \frac{\sqrt{4\pi n}}{\prod_i \sqrt{2\pi \alpha_i n}} \exp nH(\alpha)$$

where $H(\alpha) \triangleq -\sum_i \alpha_i \log \alpha_i$ is the entropy of the distribution described by α .

Fix the parameters (d, k, M, π) . For ease of analysis, we will work in a version of the configuration model, where a graph \mathbf{G} is sampled as follows.

1. Randomly and uniformly select one of the $\binom{n}{\pi n}$ π -balanced partitions of the vertices, and adorn each vertex with d 'stubs' or 'half-edges.'
2. For each $i \in [k]$, randomly and uniformly select which of the $\pi(i)dn$ stubs will connect with every group j ; there are $\binom{\pi(i)dn}{\pi(i)M_{i,-}dn}$ such partitions of each group's stubs.
3. For each $i < j$, randomly and uniformly chose one of the $(\pi(i)M_{i,j}dn)!$ matchings of the stubs between groups i and j .

4. For each $i \in [k]$, randomly and uniformly choose one of the $(\pi(i)M_{i,i}dn - 1)!!$ perfect matchings on the group i stubs.

The result will be a simple graph with probability $1 - o_n(1)$, so any results that hold with high probability in this model will hold with the same guarantee if we choose \mathbf{G} uniformly from all graphs with an M -good partition. To sample in the null model, we simply adorn each vertex with d stubs, and choose one of the $(dn - 1)!!$ perfect matchings uniformly at random.

First Moment Bound. Let Ξ be the random variable counting the number of M -good partitions in a graph \mathbf{G} from the null model. There are $\binom{n}{\pi n}$ possible ways to partition the vertices in accordance with π , and we can read off the probability that a uniformly random matching on the dn half-edges makes each one M -good from the sampling procedure above. In particular, using the fact that M is stochastic with $\pi^T M = \pi$, we have

$$\begin{aligned} \mathbb{P}[\Xi > 0] &\leq \mathbb{E} \Xi \\ &= \frac{1}{(dn - 1)!!} \binom{n}{\pi n} \prod_i \binom{\pi(i)dn}{\pi(i)M_{i,-}dn} \prod_{i < j} (\pi(i)M_{i,j}dn)! \prod_i (\pi(i)M_{i,i}dn - 1)!! \\ &\leq \text{poly}(n) \exp \left(-\frac{dn}{2} \log dn + \frac{dn}{2} + nH(\pi) + \sum_i \pi(i)dnH(M_{i,-}) \right. \\ &\quad \left. + \frac{1}{2} \sum_{i,j} (\pi(i)M_{i,j}dn \log \pi(i)M_{i,j}dn - \pi(i)M_{i,j}dn) \right) \\ &= \text{poly}(n) \exp n \left(\left(1 - \frac{d}{2}\right) H(\pi) + \frac{d}{2} H(\pi, M) \right) \end{aligned}$$

where $H(M, \pi) = \sum_i \pi(i)H(M_{i,-})$ is the average row entropy of M , under the stationary distribution. Thus, the probability that $\Xi > 0$, i.e. that \mathbf{G} has any M -good partitions, is exponentially small whenever

$$d > \frac{2H(\pi)}{H(\pi) - H(\pi, M)}$$

□

Remark B.1. As a sanity check, when M has zero on the diagonal and $\frac{1}{k-1}$ elsewhere, π is the uniform distribution with entropy $\log k$, and the average row entropy of M is $\log(k-1)$, so we have a first moment bound of

$$d > \frac{2 \log k}{\log k - \log(k-1)} \approx 2k \log(k-1) - \log k + 2.$$

An M -good partition is, in this case, a coloring—although not all colorings are M -good, since they might have atypically many edges between each group—and this bound matches roughly the first moment bound for coloring in Erdős-Rényi and d -regular random graphs.

B.1 Local Statistics in the Planted Model

Given graph on constant number of vertices H , a subset of distinguished vertices S and a labeling $\tau : S \rightarrow [k]$ of the distinguished vertices, we are interested in computing $\mathbb{E} p_{H,S,\tau}(\mathbf{x}, \mathbf{G})$. Recall Φ_H was the set of

all injective homomorphisms $\phi : H \rightarrow \mathbf{G}$ and define Φ'_H as the set of all injective homomorphisms $\phi : H \rightarrow K_n$. Finally recall how $p_{H,S,\tau}$ was defined:

$$p_{H,S,\tau}(\mathbf{x}, \mathbf{G}) = \sum_{\phi \in \Phi_H} \prod_{u \in S} \mathbf{x}_{\phi(u), \tau(u)}.$$

This can be rewritten as

$$p_{H,S,\tau}(\mathbf{x}, \mathbf{G}) = \sum_{\phi \in \Phi'_H} \mathbf{1}[\phi(H) \subseteq \mathbf{G}] \prod_{u \in S} \mathbf{x}_{\phi(u), \tau(u)}.$$

For the rest of this section, assume that $V(\mathbf{G})$ is partitioned into an arbitrary π -balanced partition $C_1 \cup \dots \cup C_k$. Each $u \in V$ is adorned with d half-edges. We refer to the collection of half-edges as \tilde{V} . For half-edge v attached to vertex u we use $\sigma(v)$ to denote the cluster in $[k]$ that u is assigned to. Finally, refer to the collection of half-edges that adorn vertices in C_i as S_i .

Proposition B.2. *Let \mathbf{P} be the random perfect matching on the half-edges \tilde{V} that is used to sample \mathbf{G} , and let \mathbf{R} be some matching of size T on the complete graph on \tilde{V} where T is a constant. Then,*

$$\Pr[\mathbf{R} \subseteq \mathbf{P}] = (1 \pm o_n(1)) \cdot \prod_{\{i,j\} \in \mathbf{R}} \frac{M_{\sigma(i)\sigma(j)}}{\pi(\sigma(j)) d_n}.$$

Note that the above is well-defined since $\frac{M_{\sigma(i)\sigma(j)}}{\pi(\sigma(j)) d_n} = \frac{M_{\sigma(j)\sigma(i)}}{\pi(\sigma(i)) d_n}$.

Proof. Let T_a be the number of matched vertices in S_a in \mathbf{R} . And let $T_{a,b}$ be the number of edges in \mathbf{R} between clusters S_a and S_b . We give an expression for the probability of $\mathbf{R} \subseteq \mathbf{P}$. By using $\pi(a)M_{ab} = \pi(b)M_{ba}$, we can write

$$\Pr[\mathbf{R} \subseteq \mathbf{P}] = \frac{\prod_{a < b} \frac{(\pi(a)M_{ab} d_n)!}{(\pi(a)M_{ab} d_n - T_{a,b})!} \cdot \prod_a \frac{(\pi(a)M_{aa} d_n)!}{(\pi(a)M_{aa} d_n - 2T_{a,a})!} \cdot \frac{(\pi(a)M_{aa} d_n - 2T_{a,a} - 1)!!}{(\pi(a)M_{aa} d_n - 1)!!}}{\prod_a \frac{(\pi(a) d_n)!}{(\pi(a) d_n - T_a)!}}.$$

This is in

$$(1 \pm o_n(1)) \frac{\prod_{a < b} (\pi(a)M_{ab} d_n)^{T_{a,b}} \cdot \prod_a (\pi(a)M_{aa} d_n)^{T_{a,a}}}{\prod_a (\pi(a) d_n)^{T_a}} \quad (13)$$

The product in the numerator of (13) has one term per edge; specifically for edge $\{i, j\}$, there is a term equal to $\pi(\sigma(i))M_{\sigma(i)\sigma(j)} d_n$. For every edge $\{i, j\}$, there are two terms in the denominator — $\pi(\sigma(i)) d_n$ and $\pi(\sigma(j)) d_n$. Thus, we can rewrite (13) as

$$(1 \pm o_n(1)) \prod_{\{i,j\} \in \mathbf{R}} \frac{\pi(\sigma(i))M_{\sigma(i)\sigma(j)} d_n}{(\pi(\sigma(i)) d_n)(\pi(\sigma(j)) d_n)} = (1 \pm o_n(1)) \prod_{\{i,j\} \in \mathbf{R}} \frac{M_{\sigma(i)\sigma(j)}}{\pi(\sigma(j)) d_n}.$$

□

For the rest of this section we fix an arbitrary π -balanced labeling $\sigma : [n] \rightarrow [k]$ — i.e., we fix $\mathbf{x}_{i,\sigma(i)}$ to be 1 and the remaining $\mathbf{x}_{i,c}$ to be 0. We first prove Theorem 4.6 in the special case when (H, S) is a self avoiding walk on s vertices v_1, \dots, v_s with $S = \{v_1, v_s\}$ and use this to derive the theorem for all (H, S) .

Towards this end, we begin by computing $\mathbb{E} p_{H,S,\tau}(\mathbf{x}, \mathbf{G})$. (Recall that the polynomial $p_{H,S,\tau}$ was defined in Equation (5).) Let $L_{H,S,\tau}$ be the collection of all labelings γ of v_1, \dots, v_s in $[k]$ such that γ agrees with τ on S . Define Φ'_H as the set of all injective homomorphisms from H to the complete graph on n vertices. We can rewrite $p_{H,S,\tau}$ as

$$\begin{aligned}
p_{H,S,\tau}(\mathbf{x}, \mathbf{G}) &= \sum_{\phi \in \Phi'_H} \#\{\phi(H) \subseteq \mathbf{G}\} \prod_{u \in S} x_{\phi(u), \tau(u)} \\
&= \sum_{\phi \in \Phi'_H} \#\{\phi(H) \subseteq \mathbf{G}\} \prod_{u \in S} \mathbf{1}[\sigma(\phi(u)) = \tau(u)] \\
&= \sum_{\gamma \in L_{H,S,\tau}} \sum_{\phi \in \Phi'_H} \#\{\phi(H) \subseteq \mathbf{G}\} \prod_{u \in V(H)} \mathbf{1}[\sigma(\phi(u)) = \gamma(u)] \\
&= \sum_{\gamma \in L_{H,S,\tau}} \sum_{\phi \in \Phi'_{H,\gamma}} \#\{\phi(H) \subseteq \mathbf{G}\} \tag{14}
\end{aligned}$$

where $\Phi'_{H,\gamma}$ is the collection of all $\phi \in \Phi'_H$ such that $\sigma(\phi(v_i)) = \gamma(v_i)$ for $i = 1, \dots, s$. Thus in expectation:

$$\mathbb{E} p_{H,S,\tau}(\mathbf{x}, \mathbf{G}) = \sum_{\gamma \in L_{H,S,\tau}} \sum_{\phi \in \Phi'_{H,\gamma}} \mathbb{E}[\#\{\phi(H) \subseteq \mathbf{G}\}]. \tag{15}$$

For fixed γ and fixed ϕ in $\Phi'_{H,\gamma}$ we compute $\mathbb{E}[\#\{\phi(H) \subseteq \mathbf{G}\}]$. There are exactly $d^s(d-1)^{s-2}$ partial matchings of half-edges that “collapse” to $\phi(H)$. Fix such a matching R and use \mathbf{P} to denote the random perfect matching picked on half-edges to sample \mathbf{G} . Then

$$\mathbb{E}[\#\{\phi(H) \subseteq \mathbf{G}\}] = \Pr[R \subseteq \mathbf{P}] \cdot d^s(d-1)^{s-2}.$$

By plugging in Proposition B.2 into the above expression we get that it is equal to

$$(1 \pm o_n(1)) \cdot \left(\prod_{i=1}^{s-1} \frac{M_{\gamma(v_i)\gamma(v_{i+1})}}{\pi(\gamma(v_{i+1}))} dn \right) \cdot d^s(d-1)^{s-2}.$$

Since the inner summand only depends on γ , (15) can be written as

$$\begin{aligned}
\mathbb{E} p_{H,S,\tau}(\mathbf{x}, \mathbf{G}) &= (1 \pm o_n(1)) \cdot \sum_{\gamma \in L_{H,S,\tau}} |\Phi'_{H,\gamma}| \cdot \left(\prod_{i=1}^{s-1} \frac{M_{\gamma(v_i)\gamma(v_{i+1})}}{\pi(\gamma(v_{i+1}))} dn \right) \cdot d^s(d-1)^{s-2} \\
&= (1 \pm o_n(1)) \cdot \sum_{\gamma \in L_{H,S,\tau}} \left(\prod_{i=1}^s \pi(\gamma(v_i)) n \right) \cdot \left(\prod_{i=1}^{s-1} \frac{M_{\gamma(v_i)\gamma(v_{i+1})}}{\pi(\gamma(v_{i+1}))} dn \right) \cdot d^s(d-1)^{s-2} \\
&= (1 \pm o_n(1)) \cdot \pi(\gamma(v_1)) n \cdot \sum_{\gamma \in L_{H,S,\tau}} \prod_{i=1}^{s-1} M_{\gamma(v_i)\gamma(v_{i+1})} \cdot d(d-1)^{s-2} \\
&= (1 \pm o_n(1)) \cdot \pi(\tau(v_1)) n \cdot M_{\tau(v_1)\tau(v_s)}^{s-1} \cdot d(d-1)^{s-2} \\
&= (1 \pm o_n(1)) \cdot \pi(\tau(v_1)) n \cdot M_{\tau(v_1)\tau(v_s)}^{\text{dist}(v_1, v_s)} \cdot C_{H,S,d}.
\end{aligned}$$

Next we bound the variance of $p_{H,S,\tau}(\mathbf{x}, \mathbf{G})$. For any subgraph B of K_n , denote the set of partial matchings on dn half-edges that induce B with Z_B . In the expression for $p_{H,S,\tau}(\mathbf{x}, \mathbf{G})$ from (14), we can write $\#\{\phi(H) \subseteq \mathbf{G}\}$ as

$$\sum_{R \in Z_{\phi(H)}} \mathbf{1}[R \subseteq \mathbf{P}].$$

We are interested in computing $\mathbb{E} p_{H,S,\tau}(\mathbf{x}, \mathbf{G})^2 - (\mathbb{E} p_{H,S,\tau}(\mathbf{x}, \mathbf{G}))^2$. Writing $p_{H,S,\tau}$ as

$$\sum_{\gamma \in \mathcal{L}_{H,S,\tau}} \sum_{\phi \in \Phi'_{H,\gamma}} \sum_{M \in Z_{\phi(H)}} \mathbf{1}[R \subseteq \mathbf{P}]$$

tells us that it is a sum of terms of the form $\mathbf{1}[R \subseteq \mathbf{P}]$ over $(s-1)$ -sized matchings in some collection \mathcal{M} . Thus,

$$\mathbb{E} p_{H,S,\tau}(\mathbf{x}, \mathbf{G})^2 - (\mathbb{E} p_{H,S,\tau}(\mathbf{x}, \mathbf{G}))^2 = \sum_{R_1, R_2 \in \mathcal{M}} \mathbb{E}[\mathbf{1}[R_1 \cup R_2 \subseteq \mathbf{P}]] - \mathbb{E}[\mathbf{1}[R_1 \subseteq \mathbf{P}]] \cdot \mathbb{E}[\mathbf{1}[R_2 \subseteq \mathbf{P}]]$$

We can split the sum in the RHS of the above into two sums: a sum over pairs R_1 and R_2 that don't intersect, and a sum over pairs that intersect. If $R_1 \cup R_2$ is not a valid matching, then it has a negative contribution to the above sum. From Proposition B.2, we get:

$$\begin{aligned} & \sum_{\substack{R_1, R_2 \in \mathcal{M} \\ R_1 \cap R_2 = \emptyset \\ R_1 \cup R_2 \text{ valid matching}}} \Pr[R_1 \cup R_2 \subseteq \mathbf{P}] - \Pr[R_1 \subseteq \mathbf{P}] \cdot \Pr[R_2 \subseteq \mathbf{P}] \\ = & \sum_{\substack{R_1, R_2 \in \mathcal{M} \\ R_1 \cap R_2 = \emptyset \\ R_1 \cup R_2 \text{ valid matching}}} (1 \pm o_n(1)) \cdot \prod_{\{i,j\} \in R_1 \cup R_2} \frac{M_{\sigma(i)\sigma(j)}}{\pi(\sigma(j)) d_n} - (1 \pm o_n(1)) \prod_{\{i,j\} \in R_1} \frac{M_{\sigma(i)\sigma(j)}}{\pi(\sigma(j)) d_n} \prod_{\{i,j\} \in R_2} \frac{M_{\sigma(i)\sigma(j)}}{\pi(\sigma(j)) d_n} \\ = & \sum_{\substack{R_1, R_2 \in \mathcal{M} \\ R_1 \cap R_2 = \emptyset \\ R_1 \cup R_2 \text{ valid matching}}} o_n(1) \cdot \prod_{\{i,j\} \in R_1 \cup R_2} \frac{M_{\sigma(i)\sigma(j)}}{\pi(\sigma(j)) d_n} \end{aligned}$$

Since $|R_1 \cup R_2| = 2(s-1)$ and $|\mathcal{M}| \leq (d^2 n)^s$ the above can be upper bounded by $o_n(1) \cdot (d^2 n)^{2s} \cdot \frac{1}{(d n)^{2s-2}} = o_n(n^2)$.

Next we consider the sum over pairs of matchings which have nonempty intersection:

$$\sum_{\substack{R_1, R_2 \in \mathcal{M} \\ R_1 \cap R_2 \neq \emptyset \\ R_1 \cup R_2 \text{ valid matching}}} \Pr[R_1 \cup R_2 \subseteq \mathbf{P}] - \Pr[R_1 \subseteq \mathbf{P}] \cdot \Pr[R_2 \subseteq \mathbf{P}] \leq \sum_{\substack{R_1, R_2 \in \mathcal{M} \\ R_1 \cap R_2 \neq \emptyset}} 2(1 \pm o_n(1)) \prod_{\{i,j\} \in R_1 \cup R_2} \frac{M_{\sigma(i)\sigma(j)}}{\pi(\sigma(j)) d_n}$$

$|R_1 \cup R_2| = 2(s-1) - |R_1 \cap R_2|$ and the number of summands in the RHS of the above expression can be bounded by $(d^2 n)^{2(s-1) - |R_1 \cap R_2| + 1}$ and as a result the RHS of the above can be bounded by $O(n)$ (where the $O(\cdot)$ hides constants depending on d and s). Putting the two bounds together we get:

$$\mathbb{V}[p_{H,S,\tau}] = o_n(n^2).$$

By Chebyshev's inequality,

Proposition B.3. $p_{H,S,\tau}$ is equal to $(1 \pm o_n(1)) \cdot \pi(\tau(v_1)) n \cdot M_{\tau(v_1)\tau(v_s)}^{\text{dist}(v_1, v_s)} \cdot C_{H,S,d}$ with probability $1 - o_n(1)$ thereby confirming Theorem 4.6 for paths with labeled endpoints.

We now use this derive Theorem 4.6 for all graphs.

Proof of Theorem 4.6. Now suppose (H, S, τ) is any partially labelled graph with T edges where T is a constant, let H_1, \dots, H_ℓ denote the connected components of H and assign a distinguished vertex u_i to connected component i (where $u_i \in S$ if H_i contains a vertex in S). We return to expressing $p_{H,S,\tau}$ as it was defined in (5). Recall that Φ_H was the set of all injective homomorphisms $\phi : H \rightarrow \mathbf{G}$.

$$p_{H,S,\tau}(\mathbf{x}, \mathbf{G}) = \sum_{\phi \in \Phi_H} \prod_{u \in S} x_{\phi(u), \tau(u)}$$

We show that except with probability $o_n(1)$ the random variable $p_{H,S,\tau}$ is $n^\ell \pi(\tau(\alpha)) M_{\tau(\alpha)\tau(\beta)}^{\text{dist}(\alpha,\beta)} C_{H,S,d} \pm o(n^\ell)$ when $S = \{\alpha, \beta\}$ and is $n^\ell C_{H,S,d} \pm o(n^\ell)$ when $S = \emptyset$. The statement about $\mathbb{E} p_{H,S,\tau}$ follows from the fact that $p_{H,S,\tau}$ is always nonnegative and is bounded above by n^ℓ .

We break our proof into cases.

Case 1: Some H_i contains a cycle. From an application of Proposition B.2, the number of cycles of length $\leq T$ in \mathbf{G} is $o(n)$ with probability $o_n(1)$, and thus the number of injective homomorphisms from H_i to \mathbf{G} is $o(n)$. Since there are at most $O(n)$ injective homomorphisms from any other connected component H_j to \mathbf{G} , with probability $1 - o_n(1)$, the number of injective homomorphisms is $o(n^\ell)$ and the statement of Theorem 4.6 follows for this case.

In the rest of this proof we will assume that H is a forest.

Case 2: S is empty. There are at most n^ℓ ways to map the collection of u_i to $[n]$. Suppose u_i has been mapped to a vertex v , then there are at most $C_{H_i,S,d}$ ways to map the remaining vertices of H_i to $V(\mathbf{G})$ since \mathbf{G} is d -regular graph. This gives an upper bound of $\prod_{1 \leq i \leq \ell} C_{H_i,S,d}$ on $n^{-\ell} p_{H,S,\tau}(\mathbf{x}, \mathbf{G})$.

To prove a nearly matching lower bound that holds with probability $1 - o_n(1)$, observe that if each u_i is mapped to v_i such that (a) the radius- T neighborhood of v_i is a tree, (b) the distance between any v_i and v_j for $i \neq j$ is at least $2T + 1$, then there are $\prod_{1 \leq i \leq \ell} C_{H_i,S,d} = C_{H,S,d}$ ways to map the rest of the vertices from H to \mathbf{G} .

We inductively show that for every constant $\ell \geq 0$, the number of ways to map u_1, \dots, u_ℓ to v_1, \dots, v_ℓ by satisfying conditions (a) and (b) is at least $(1 - o_n(1))n^\ell$. Since there are $o(n)$ cycles of length $\leq T$ in the graph except with probability $o_n(1)$, there are $n - o(n)$ ways to map u_1 to v_1 while satisfying (a) and (b). Suppose the number of maps from $u_1, \dots, u_{\ell-1}$ to $V(\mathbf{G})$ satisfying (a) and (b) is $(1 - o_n(1))n^{\ell-1}$. If $u_1, \dots, u_{\ell-1}$ have been mapped to $v_1, \dots, v_{\ell-1}$ in a way respecting (a) and (b) the number of ways to map u_ℓ to a vertex in $V(\mathbf{G})$ while satisfying (a) and (b) is $n - o(n) \cdot (\ell - 1) \cdot d(d - 1)^T$, which is $n - o(n)$ since ℓ, T and d are constants. Thus, from our induction hypothesis the number of ways to map u_1, \dots, u_ℓ to v_1, \dots, v_ℓ is $(1 - o_n(1))n^\ell$, and our treatment of this case is complete.

Case 3: $\alpha, \beta \in S$ are in the same connected component. Since α, β are in the same connected component (which, without loss of generality, we say is H_1) and H is a forest there is a unique path $P_{\alpha,\beta}$ between them. From Proposition B.3, $p_{P_{\alpha,\beta},S,\tau} = (1 \pm o_n(1)) \cdot \pi(\tau(\alpha))n \cdot M_{\tau(\alpha)\tau(\beta)}^{\text{dist}(\alpha,\beta)} \cdot C_{P_{\alpha,\beta},S,d}$. There are at most $\frac{C_{H_1,S,d}}{C_{P_{\alpha,\beta},S,d}}$ ways to extend an injective homomorphism $\phi : P_{\alpha,\beta} \rightarrow \mathbf{G}$ to $\phi : H_1 \rightarrow \mathbf{G}$, and thus

$$p_{H_1,S,\tau} = (1 \pm o_n(1)) \cdot \pi(\tau(\alpha))n \cdot M_{\tau(\alpha)\tau(\beta)}^{\text{dist}(\alpha,\beta)} \cdot C_{H_1,S,d}. \quad (16)$$

There are at most $n^{\ell-1} \cdot \prod_{2 \leq i \leq \ell} C_{H_i, \emptyset, d}$ ways to map the rest of H into \mathbf{G} via an injective homomorphism, which establishes the upper bound in this case.

The proof of the matching lower bound of $(1 - o_n(1)) \cdot \pi(\tau(\alpha)) n^\ell \cdot M_{\tau(\alpha)\tau(\beta)}^{\text{dist}(\alpha, \beta)} \cdot C_{H, S, d}$ is identical to the proof in **Case 2**.

Case 4: $\alpha, \beta \in S$ are in different connected components. Without loss of generality, we say $\alpha \in H_1$ and $\beta \in H_2$. There are at most $\pi(\tau(\alpha))\pi(\tau(\beta))n^2$ ways to injectively map $\{\alpha, \beta\}$ to \mathbf{G} and at most $C_{H_1, \{\alpha\}, d} C_{H_2, \{\beta\}, d} \cdot \prod_{3 \leq i \leq \ell} C_{H_i, \emptyset, d}$ ways to extend this to an injective homomorphism from H to \mathbf{G} , which establishes an upper bound of

$$C_{H, S, d} \cdot \pi(\tau(\alpha)) M_{\tau(\alpha)\tau(\beta)}^{\text{dist}(\alpha, \beta)} n^\ell.$$

Once again, the proof of the matching lower bound of $(1 - o_n(1)) \cdot \pi(\tau(\alpha)) n^\ell \cdot M_{\tau(\alpha)\tau(\beta)}^{\text{dist}(\alpha, \beta)} \cdot C_{H, S, d}$ follows the proof from **Case 2**. \square

C Proof of Proposition 7.2

Assume that $\tilde{\mathbb{E}}$ matches all the promised moments, and let (H, S, τ) be an arbitrary partially labelled graph with ℓ connected components, and two distinguished vertices connected by a path of length $s \in \{0, \dots, m\} \cup \{\infty\}$, where $s = 0$ means that the distinguished vertices (and their labels) are identical, and $s = \infty$, means that lie in disjoint connected components connecting them. Let's use the shorthand P_s for a path of length s with distinguished endpoints, and (P_s, i, j) to mean that these endpoints have labels $i, j \in [k]$. Let's also adopt the shorthand $C_{P_s, d}$ for the constant corresponding to P_s in 4.5. All these definitions extend naturally to the corner cases $s = 0$, when there is only one distinguished vertex and $i = j$, and $s = \infty$ when we will interpret it as two disconnected and labelled vertices.

Our central claim will be that, as polynomials, with high probability

$$\left\| n^{\ell-1} \frac{C_{d, H, S}}{C_{d, P_s}} p_{P_s, i, j} - p_{H, S, \tau} \right\|_1 = o(n^\ell) \simeq_\ell 0 \quad (17)$$

where here $\|\cdot\|$ means the coefficient-wise L^1 norm. To get a feel for this, in the case when H is connected and its two distinguished vertices are connected by a path of length s , $\frac{C_{H, S, d}}{C_{P_s, d}}$ counts the number of ways to place the remainder of H in a d -regular tree (or locally-treelike graph), once we commit to the locations of the two distinguished vertices and the path between them.

Once we've shown (17), it is a standard SoS calculation that the Boolean constraint implies $|\tilde{\mathbb{E}} \chi_{u, i} \chi_{v, j}| \leq 1$ for every u, v, i, j . Thus, using the local statistic constraints from Theorem 4.6, we will have

$$\tilde{\mathbb{E}} p_{H, S, \tau} \simeq_\ell n^{\ell-1} \frac{C_{d, H, S}}{C_{d, P_s, T}} \tilde{\mathbb{E}} p_{P_s, i, j} \simeq_\ell n^\ell \pi(i) M_{i, j}^s C_{d, H, S}$$

as desired. To prove (17), we need to open up the subgraph counting polynomials $p_{H, S, \tau}$ a bit more, and we'll require some notation to do this cleanly. Maintain the notation of Φ_H from Section 3, namely that this is the set of all injective homomorphisms from H to \mathbf{G} ; moreover for each (H, S, τ) , write $\Phi_{H, S, u, v}$ for the

set of all injective homomorphisms $H \rightarrow \mathbf{G}$ such that the distinguished vertices are mapped to the vertices u and v of \mathbf{G} . Then

$$\begin{aligned} \left\| n^{\ell-1} \frac{C_{d,H,S}}{C_{d,P_s,T}} p_{P_s,i,j}(x) - p_{H,S,\tau}(x) \right\| &= \left\| \sum_{u,v} \left(n^{\ell-1} \frac{C_{d,H,S}}{C_{d,P_s,T}} |\Phi_{P_s,u,v}| - |\Phi_{H,S,\tau}| \right) x_{u,i} x_{v,j} \right\| \\ &= \sum_{u,v} \left| n^{\ell-1} \frac{C_{d,H,S}}{C_{d,P_s,T}} |\Phi_{P_s,u,v}| - |\Phi_{H,S,u,v}| \right|. \end{aligned}$$

Let's first consider the case when $s < \infty$, so that the two distinguished vertices in H are connected by some path of length at most m . Under this assumption $|\Phi_{H,S,u,v}| = O(n^{\ell-1})$. From Lemma 4.2, there are at least $n - O(\log n)$ vertices whose depth- m neighborhoods contain no cycles; we can restrict the above sum to run over only $u, v \in U$ and pay a cost of no more than $O(n^{\ell-1} \log n) = o(n^\ell)$. Thus

$$\left\| n^{\ell-1} \frac{C_{d,H,S}}{C_{d,P_s,T}} p_{P_s,i,j}(x) - p_{H,S,\tau}(x) \right\|_1 = \sum_{u,v \in U} \left| n^{\ell-1} \frac{C_{d,H,S}}{C_{d,P_s,T}} |\Phi_{P_s,u,v}| - |\Phi_{H,S,u,v}| \right| + o(n^\ell).$$

There are only $O(n)$ nonzero terms in the above sum, because every vertex can reach only constantly many vertices via a path of length $s \leq m$, and both $\Phi_{H,S,u,v}$ and $\Phi_{P_s,u,v}$ are empty if u and v are not connected in this way. This means we've reduced the problem to showing that

$$\left| n^{\ell-1} \frac{C_{d,H,S}}{C_{d,P_s,T}} |\Phi_{P_s,u,v}| - |\Phi_{H,S,u,v}| \right| = o(n^{\ell-1}).$$

Let's write $H = H_1 \sqcup H_2 \cdots \sqcup H_\ell$, and assume that both distinguished vertices are in H_1 . Recall for later use that $C_{H,S,d} = C_{H_1,S,d} C_{H_2,\emptyset} \cdots C_{H_\ell,\emptyset}$. Since each of the components is of constant size, we claim that

$$|\Phi_{H,S,u,v}| = |\Phi_{H_1,S,u,v}| |\Phi_{H_2}| \cdots |\Phi_{H_\ell}| + o(n^{\ell-1})$$

where by the latter terms we mean the number of occurrences of H_2, \dots, H_ℓ respectively in \mathbf{G} , without any constraints of distinguished vertices or labels. The idea is that we can choose one of the occurrences in $\Phi_{H,S,u,v}$ by first choosing where to place (H_1, S) so that the vertices in S map to u and v , then choosing where to place H_2 , etc. Since each component is of constant size, we overcount only by an additive $O(n^{\ell-1})$ by ignoring the requirement that the connected components not collide.

The quantity $|\Phi_{H_i}|$ is the count of a particular subgraph, with no label constraints, so Theorem 4.6 tells us that w.h.p.

$$|\Phi_{H_i}| = n C_{d,H_i} + o(n).$$

Similarly, from our discussion above,

$$|\Phi_{H_1,S_1,u,v}| = \frac{C_{H_1,S_1,d}}{C_{P_s,d}},$$

since once we declare that the two vertices in S are mapped to two vertices whose depth- m neighborhoods are tree-like, our freedom in place the rest of H_1 is exactly the number of ways to do so in a d -regular tree. But now we are done:

$$\begin{aligned} |\Phi_{H,S,u,v}| &= |\Phi_{H_1,S_1,u,v}| |\Phi_{H_2}| \cdots |\Phi_{H_\ell}| + o(n^{\ell-1}) \\ &= n^{\ell-1} \frac{C_{H_1,S_1,d}}{C_{P_s,d}} C_{H_2} \cdots C_{H_\ell} + o(n^{\ell-1}) \\ &= n^{\ell-1} \frac{C_{H,S,d}}{C_{P_s,d}} + o(n^{\ell-1}). \end{aligned}$$

We finally need to treat the case when $s = \infty$ and the distinguished vertices of H are in different connected components; under this assumption $|\Phi_{H,S,u,v}| = O(n^{\ell-2})$, since there are only $\ell - 2$ components to place once we commit to mapping the distinguished vertices to a given u and v . Again we need to consider

$$\sum_{u,v} \left| n^{\ell-1} \frac{C_{H,S,d}}{C_{P_s,d}} |\Phi_{P_s,u,v}| - |\Phi_{H,S,u,v}| \right|,$$

and again we can restrict the sum to vertices without nearby cycles for the price of an additive $o(n^\ell)$, since we are omitting $O(n \log n)$ terms with magnitude $O(n^{\ell-2})$. From this point the calculation continues as above.