

Sum-of-Squares Lower Bounds for Sparse Independent Set

Chris Jones* Aaron Potechin[†] Goutham Rajendran[‡] Madhur Tulsiani[§]
Jeff Xu[¶]

December 20, 2021

Abstract

The Sum-of-Squares (SoS) hierarchy of semidefinite programs is a powerful algorithmic paradigm which captures state-of-the-art algorithmic guarantees for a wide array of problems. In the average case setting, SoS lower bounds provide strong evidence of algorithmic hardness or information-computation gaps. Prior to this work, SoS lower bounds have been obtained for problems in the “dense” input regime, where the input is a collection of independent Rademacher or Gaussian random variables, while the sparse regime has remained out of reach. We make the first progress in this direction by obtaining strong SoS lower bounds for the problem of Independent Set on sparse random graphs. We prove that with high probability over an Erdős-Rényi random graph $G \sim G_{n, \frac{d}{n}}$ with average degree $d > \log^2 n$, degree- D_{SoS} SoS fails to refute the existence of an independent set of size $k = \Omega\left(\frac{n}{\sqrt{d}(\log n)(D_{\text{SoS}})^{c_0}}\right)$ in G (where c_0 is an absolute constant), whereas the true size of the largest independent set in G is $O\left(\frac{n \log d}{d}\right)$.

Our proof involves several significant extensions of the techniques used for proving SoS lower bounds in the dense setting. Previous lower bounds are based on the pseudo-calibration heuristic of Barak et al. [FOCS 2016] which produces a candidate SoS solution using a planted distribution indistinguishable from the input distribution via low-degree tests. In the sparse case the natural planted distribution *does* admit low-degree distinguishers, and we show how to adapt the pseudo-calibration heuristic to overcome this.

Another notorious technical challenge for the sparse regime is the quest for matrix norm bounds. In this paper, we obtain new norm bounds for *graph matrices* in the sparse setting. While in the dense setting the norms of graph matrices are characterized by the size of the minimum vertex separator of the corresponding graph, this turns not to be the case for sparse graph matrices. Another contribution of our work is developing a new combinatorial understanding of structures needed to understand the norms of sparse graph matrices.

*University of Chicago. csj@uchicago.edu. Supported in part by NSF grant CCF-2008920.

[†]University of Chicago. potechin@uchicago.edu. Supported in part by NSF grant CCF-2008920.

[‡]University of Chicago. goutham@uchicago.edu. Supported in part by NSF grant CCF-1816372.

[§]Toyota Technological Institute at Chicago. madhurt@ttic.edu. Supported by NSF grant CCF-1816372.

[¶]Carnegie Mellon University. jeffxusichao@cmu.edu. Supported by NSF CAREER Award #2047933.

Contents

1	Introduction	1
1.1	Our main results	2
1.2	Our approach	3
1.3	Related work	4
1.4	Organization of the paper	5
2	Technical Preliminaries	5
2.1	The Sum-of-Squares hierarchy	5
2.2	p -biased Fourier analysis	5
2.3	Ribbons and graph matrices	6
3	An overview of the proof techniques	8
3.1	Modified pseudocalibration	8
3.2	Approximate PSD decompositions, norm bounds and conditioning	10
4	Pseudocalibration with connected truncation	12
4.1	Pseudo-calibration	12
4.2	The failure of "Just try pseudo-calibration"	13
4.3	Salvaging the doomed	14
5	Norm bounds and factoring	17
5.1	Norm bounds	17
5.2	Factoring ribbons and shapes	18
6	PSD-ness	19
6.1	Overview of the approximate PSD decomposition	19
6.2	Informal sketch for bounding τ and τ_p	21
6.3	Decomposition in terms of ribbons	26
6.4	Factor out Π and edges incident to $U_\alpha \cap V_\alpha$	28
6.5	Conditioning I: reduction to sparse shapes	31
6.6	Shifting to shapes	35
6.7	Counting shapes with the tail bound function $c(\alpha)$	40
6.8	Statement of main lemmas	42
6.9	Conditioning II: Frobenius norm trick	44
6.10	Norm bounds	45
6.11	Proof of main lemmas	47
7	Open Problems	57

A Trace Power Calculations	61
B Reserving edges to preserve left/middle shapes	65
C Second proof of Lemma 6.9	66
D Formalizing the PSD Decomposition	68

1 Introduction

The Sum-of-Squares (SoS) hierarchy is a powerful convex programming technique that has led to successful approximation and recovery algorithms for various problems in the past decade. SoS captures the best-known approximation algorithms for several classical combinatorial optimization problems. Some of the additional successes of SoS also include Tensor PCA [HSS15, MSS16] and Constraint Satisfaction Problems with additional structure [BRS11, GS11]. SoS is a family of convex relaxations parameterized by *degree*; by taking larger degree, one gets a better approximation to the true optimum at the expense of a larger SDP instance. Thus we are interested in the tradeoff between degree and approximation quality. For an introduction to Sum-of-Squares algorithms, see [BS16, FKP19].

The success of SoS on the upper bound side has also conferred on it an important role for the investigation of algorithmic hardness. Lower bounds for the SoS hierarchy provide strong unconditional hardness results for several optimization problems and are of particular interest when NP-hardness results are unavailable. An important such setting is the study of average case complexity of optimization problems, where relatively few techniques exist for establishing NP-hardness results [ABB19]. In this setting, a study of the SoS hierarchy not only provides a powerful benchmark for average-case complexity, but also helps in understanding the structural properties of the problems: what makes them algorithmically challenging? Important examples of such results include an improved understanding of sufficient conditions for average-case hardness of CSPs [KMOW17] and lower bounds for the planted clique problem [BHK⁺16].

An important aspect of previous lower bounds for the SoS hierarchy is that they apply for the so-called *dense setting*, which corresponds to cases when the input distribution can be specified by a collection of independent Rademacher or Gaussian variables. In the case of planted clique, this corresponds to the case when the input is a random graph distributed according to $G_{n, \frac{1}{2}}$ i.e. specified by a collection of $\binom{n}{2}$ independent Rademacher variables. In the case of CSPs, one fixes the structure of the lower bound instance and only considers an instance to be specified by the signs of the literals, which can again be taken as uniformly random $\{-1, 1\}$ variables. Similarly, recent results by a subset of the authors [PR20] for tensor PCA apply when the input tensor has independent Rademacher or Gaussian entries. The techniques used to establish these lower bounds have proved difficult to extend to the case when the input distribution naturally corresponds to a *sparse graph* (or more generally, when it is specified by a collection of independent sub-gaussian variables, with Orlicz norm $\omega(1)$ instead of $O(1)$).

In this paper we are interested in extending lower bound technology for SoS to the *sparse setting*, where the input is a graph with average degree $d \leq n/2$. We use as a case study the fundamental combinatorial optimization problem of independent set. For the dense case $d = n/2$, finding an independent set is equivalent to finding a clique and the paper [BHK⁺16] shows an average-case lower bound against the Sum-of-Squares algorithm. We extend the techniques introduced there, namely pseudocalibration, graph matrices, and the approximate decomposition into positive semidefinite matrices, in order to show the first average-case lower bound for the sparse setting. We hope that the techniques developed in this paper offer a gateway for the analysis of SoS on other sparse problems. Section 7 lists several such problems that are likely to benefit from an improved understanding of the sparse setting.

Sample $G \sim G_{n, \frac{d}{n}}$ as an Erdős-Rényi random graph¹ with average degree d , where we think of $d \ll n$. Specializing to the problem of independent set, a maximum independent set in G has size:

¹Unfortunately our techniques do not work for a random d -regular graph. See the open problems (Section 7).

Fact 1.1 ([COE15, DM11, DSS16]). *W.h.p. the max independent set in G has size $(1 + o_d(1)) \cdot \frac{2 \ln d}{d} \cdot n$.*

The value of the degree-2 SoS relaxation for independent set equals the Lovász ϑ function, which is an upper bound on the independence number $\alpha(G)$, by virtue of being a relaxation. For random graphs $G \sim G_{n,d/n}$ this value is larger by a factor of about \sqrt{d} than the true value of $\alpha(G)$ with high probability.

Fact 1.2 ([CO05]). *W.h.p. $\vartheta(G) = \Theta(\frac{n}{\sqrt{d}})$.*

We will prove that the value of higher-degree SoS is also on the order of n/\sqrt{d} , rather than n/d , and thereby demonstrate that the information-computation gap against basic SDP/spectral algorithms persists against higher-degree SoS.

1.1 Our main results

The solution to the convex relaxations obtained via the SoS hierarchy can be specified by the so-called “pseudoexpectation operator”.

Definition 1.3 (Pseudoexpectation). *A degree- D pseudoexpectation operator $\tilde{\mathbb{E}}$ is a linear functional on polynomials of degree at most D (in n variables) such that $\tilde{\mathbb{E}}[1] = 1$ and $\tilde{\mathbb{E}}[f^2] \geq 0$ for every polynomial f with degree at most $D/2$. A pseudoexpectation is said to satisfy a polynomial constraint $g = 0$ if $\tilde{\mathbb{E}}[f \cdot g] = 0$ for all polynomials f when $\deg(f \cdot g) \leq D$.*

In considering relaxations for independent set of a graph $G = (V, E)$, with variables x_v being the 0/1 indicators of the independent set, the SoS relaxation searches for pseudoexpectation operators satisfying the polynomial constraints

$$\forall v \in V. \quad x_v^2 = x_v \quad \text{and} \quad \forall (u, v) \in E. \quad x_u x_v = 0.$$

The objective value of the convex relaxation is given by the quantity $\tilde{\mathbb{E}}[\sum_{v \in V} x_v] = \sum_{v \in V} \tilde{\mathbb{E}}[x_v]$. For the results below, we say that an event occurs with high probability (w.h.p.) when it occurs with probability at least $1 - O(1/n^c)$ for some $c > 0$. The following theorem states our main result.

Theorem 1.4. *There is an absolute constant $c_0 \in \mathbb{N}$ such that for sufficiently large $n \in \mathbb{N}$ and $d \in [(\log n)^2, n^{0.5}]$, and parameters k, D_{SoS} satisfying*

$$k \leq \frac{n}{D_{\text{SoS}}^{c_0} \cdot \log n \cdot d^{1/2}},$$

it holds w.h.p. for $G = (V, E) \sim G_{n, d/n}$ that there exists a degree- D_{SoS} pseudoexpectation satisfying

$$\forall v \in V. \quad x_v^2 = x_v \quad \text{and} \quad \forall (u, v) \in E. \quad x_u x_v = 0,$$

and objective value $\tilde{\mathbb{E}}[\sum_{v \in V} x_v] \geq (1 - o(1))k$.

Remark 1.5. *This is a non-trivial lower bound whenever $D_{\text{SoS}} \leq \left(\frac{d^{1/2}}{\log n}\right)^{1/c_0}$.*

Remark 1.6. *It suffices to set $c_0 = 20$ for our current proof. We did not optimize the tradeoff in D_{SoS} with k , but we did optimize the log factor (with the hope of eventually removing it).*

Remark 1.7. Using the same technique, we can prove an $n^{\Omega(\varepsilon)}$ SoS-degree lower bound for all $d \in [\sqrt{n}, n^{1-\varepsilon}]$.

For $n^\varepsilon \leq d \leq n^{0.5}$, the theorem gives a polynomial n^δ SoS-degree lower bound. For smaller d , the bound is still strong against low-degree SoS, but it becomes trivial as D_{SoS} approaches $(d^{1/2}/\log n)^{1/c_0}$ or d approaches $(\log n)^2$ since k matches the size of the maximum independent set in G , hence there is an actual distribution over independent sets of this size (the expectation operator for which is trivially is also a pseudoexpectation operator).

The above bound says nothing about the “almost dense” regime $d \in [n^{1-\varepsilon}, n/2]$. To handle this regime, we observe that our techniques, along with the ideas from the $\Omega(\log n)$ -degree SoS bound from [BHK⁺16] for the dense case, prove a lower bound for any degree $d \geq n^\varepsilon$.

Theorem 1.8. For any $\varepsilon_1, \varepsilon_2 > 0$ there is $\delta > 0$, such that for $d \in [n^{\varepsilon_1}, n/2]$ and $k \leq \frac{n}{d^{1/2+\varepsilon_2}}$, it holds w.h.p. for $G = (V, E) \sim G_{n, d/n}$ that there exists a degree- $(\delta \log d)$ pseudoexpectation satisfying

$$\forall v \in V. \quad x_v^2 = x_v \quad \text{and} \quad \forall (u, v) \in E. \quad x_u x_v = 0,$$

and objective value $\tilde{\mathbb{E}}[\sum_{v \in V} x_v] \geq (1 - o(1))k$.

In particular, these theorems rule out polynomial-time certification (i.e. constant degree SoS) for any $d \geq \text{polylog}(n)$.

1.2 Our approach

Proving lower bounds for the case of sparse graphs requires extending the previous techniques for SoS lower bounds in multiple ways. The work closest to ours is the planted clique lower bound of [BHK⁺16]. The idea there is to view a random graph $G \sim G_{n, 1/2}$ as a random input in $\{-1, 1\}^{\binom{n}{2}}$, and develop a canonical method called “pseudocalibration” for obtaining the pseudoexpectation $\tilde{\mathbb{E}}$ as a function of G . The pseudocalibration method takes the low-degree Fourier coefficients of μ based on a *different* distribution on inputs G (with large planted cliques), and takes higher degree coefficients to be zero. This is based on the heuristic that distribution $G_{n, 1/2}$ and the planted distribution are indistinguishable by low-degree tests. The pseudoexpectation obtained via this heuristic is then proved to be PSD (i.e., to satisfy $\tilde{\mathbb{E}}[f^2] \geq 0$) by carefully decomposing its representation as a (moment) matrix Λ . One then needs to estimate the norms of various terms in this decomposition, known as “graph matrices”, which are random matrices with entries as low-degree polynomials (in $\{-1, 1\}^{\binom{n}{2}}$), and carefully group terms together to form PSD matrices.

Each of the above components require a significant generalization in the sparse case. To begin with, there is no good planted distribution to work with, as the natural planted distribution (with a large planted independent set) is *distinguishable from $G_{n, d/n}$ via low-degree tests!* While we still use the natural planted distribution to compute *some* pseudocalibrated Fourier coefficients, we also truncate (set to zero) several low-degree Fourier coefficients, in addition to the high-degree coefficients as in [BHK⁺16]. In particular, when the Fourier coefficients correspond to subgraphs where certain vertex sets are disconnected (viewed as subsets of $\binom{n}{2}$), we set them to zero. This is perhaps the most conceptually interesting part of the proof, and we hope that the same “connected truncation” will be useful for other integrality gap constructions.

The technical machinery for understanding norm bounds, and obtaining PSD terms, also requires a significant update in the sparse case. Previously, norm bounds for graph matrices were understood in terms of minimum vertex separators for the corresponding graphs, and arguments

for obtaining PSD terms required working with the combinatorics of vertex separators [AMP20]. However, the number of vertices in a vertex separator turns out to be insufficient to control the relevant norm bounds in the sparse case. This is because of the fact that unlike random ± 1 variables, their p -biased analogs no longer have Orlicz norm $O(1)$ but instead $O(\frac{1}{\sqrt{p}})$, which results in both the vertices as well as edges in the graph playing a role in the norm bounds. We then characterize the norms of the relevant random matrices in terms of vertex separators, where the cost of a separator depends on the number of vertices and also the number of induced edges. Moreover, the estimates on spectral norms obtained via the trace power method can fluctuate significantly due to rare events (presence of some dense subgraphs), and we need to carefully condition on the absence of these events. A more detailed overview of our approach is presented in Section 3.

1.3 Related work

Several previous works prove SoS lower bounds in the dense setting, when the inputs can be viewed as independent Gaussian or Rademacher random variables. Examples include the planted clique lower bound of Barak et al. [BHK⁺16], CSP lower bounds of Kothari et al. [KMOW17], and the tensor PCA lower bounds [HKP⁺17, PR20]. The technical component of decomposing the moment matrix in the dense case, as a sum of PSD matrices, is developed into a general “machinery” in a recent work by a subset of the authors [PR20]. A different approach than the ones based on pseudocalibration, which also applies in the dense regime, was developed by Kunisky [Kun20].

For the case of independent set in random sparse graphs, many works have considered the *search* problem of finding a large independent set in a random sparse graph. Graphs from $G_{n,d/n}$ are known to have independent sets of size $(1 + o_d(1)) \cdot \frac{2 \ln d}{d} \cdot n$ with high probability, and it is possible to find an independent set of size $(1 + o_d(1)) \cdot \frac{\ln d}{d} \cdot n$, either by greedily taking a maximal independent set in the dense case [GM75] or by using a local algorithm in the sparse case [Wor95]. This is conjectured to be a computational phase transition, with concrete lower bounds against search beyond $\frac{\ln d}{d} \cdot n$ for local algorithms [RV17] and low-degree polynomials [Wei20]. The game in the search problem is all about the constant 1 vs 2, whereas our work shows that the integrality gap of SoS is significantly worse, on the order of \sqrt{d} . Lower bounds against search work in the regime of constant d (though in principle they could be extended to at least some $d = \omega(1)$ with additional technical work), while our techniques require $d \geq \log(n)$. For search problems, the *overlap distribution* of two high-value solutions has emerged as a heuristic indicator of computational hardness, whereas for certification problems it is unclear how the overlap distribution plays a role.

Norm bounds for sparse graph matrices were also obtained using a different method of matrix deviation inequalities, by a subset of the authors [RT20].

The work [BBK⁺20] constructs a *computationally quiet* planted distribution that is a candidate for pseudocalibration. However, their distribution is not quite suitable for our purposes.^{2 3}

A recent paper by Pang [Pan21] fixes a technical shortcoming of [BHK⁺16] by constructing a pseudoexpectation operator that satisfies “ $\sum_{v \in V} x_v = k$ ” as a polynomial constraint (whereas the

²[BBK⁺20] provide evidence that their distribution is hard to distinguish from $G_{n,d/n}$ with probability $1 - o(1)$ (it is not “strongly detectable”). However, their distribution *is* distinguishable with probability $\Omega(1)$, via a triangle count (it is “weakly detectable”). In SoS pseudocalibration, this manifests as $\tilde{\mathbb{E}}[1] = \Theta_d(1)$. We would like the low-degree distinguishing probability to be $o(1)$ i.e. $\tilde{\mathbb{E}}[1] = 1 + o_d(1)$ so that normalizing by $\tilde{\mathbb{E}}[1]$ does not affect the objective value.

³Another issue is that their planted distribution introduces noise by adding a small number of edges inside the planted independent set.

shortcoming was $\tilde{\mathbb{E}}[\sum_{v \in V} x_v] \geq (1 - o(1))k$ like we have here).

1.4 Organization of the paper

In [Section 2](#), we introduce the technical preliminaries and notation required for our arguments. A technical overview of our proof strategy is presented in [Section 3](#). [Section 4](#) describes our modification of the pseudo-calibration heuristic using additional truncations. Norm bounds for sparse graph matrices are obtained in [Section 5](#). [Section 6](#) then proves the PSD-ness of our SoS solutions broadly using the machinery developed in [\[BHK⁺16\]](#) and [\[PR20\]](#).

2 Technical Preliminaries

2.1 The Sum-of-Squares hierarchy

The Sum-of-Squares (SoS) hierarchy is a hierarchy of semidefinite programs parameterized by its degree D . We will work with two equivalent definitions of a degree- D SoS solution: a pseudoexpectation operator $\tilde{\mathbb{E}}$ ([Definition 1.3](#)) and a moment matrix. For a degree- D solution to be well defined, we need D to be at least the maximum degree of all constraint polynomials. The degree- D SoS algorithm checks feasibility of a polynomial system by checking whether or not a degree- D pseudoexpectation operator exists. This can be done in time $n^{O(D)}$ via semidefinite programming (ignoring some issues of bit complexity [\[RW17\]](#)). To show an SoS lower bound, one must construct a pseudoexpectation operator that exhibits the desired integrality gap.

2.1.1 Moment matrix

We define the moment matrix associated with a degree- D pseudoexpectation $\tilde{\mathbb{E}}$.

Definition 2.1 (Moment Matrix of $\tilde{\mathbb{E}}$). *The moment matrix $\Lambda = \Lambda(\tilde{\mathbb{E}})$ associated to a pseudoexpectation $\tilde{\mathbb{E}}$ is a $\binom{[n]}{\leq D/2} \times \binom{[n]}{\leq D/2}$ matrix with rows and columns indexed by subsets of $I, J \subseteq [n]$ of size at most $D/2$ and defined as*

$$\Lambda[I, J] := \tilde{\mathbb{E}} \left[x^I \cdot x^J \right].$$

To show that a candidate pseudoexpectation satisfies $\tilde{\mathbb{E}}[f^2] \geq 0$ in [Definition 1.3](#), we will rely on the following standard fact.

Fact 2.2. *In the definition of pseudoexpectation, [Definition 1.3](#), the condition $\tilde{\mathbb{E}}[f^2] \geq 0$ for all $\deg(f) \leq D/2$ is equivalent to $\Lambda \succeq 0$.*

2.2 p -biased Fourier analysis

Since we are interested in sparse Erdős-Rényi graphs in this work, we will resort to p -biased Fourier analysis [[O'D14](#), Section 8.4]. Formally, we view the input graph $G \sim G_{n,p}$ as a vector in $\{0, 1\}^{\binom{[n]}{2}}$ indexed by sets $\{i, j\}$ for $i, j \in [n], i \neq j$, where each entry is independently sampled from the p -biased Bernoulli distribution, $\text{Bernoulli}(p)$. Here, by convention $G_e = 1$ indicates the edge e is present, which happens with probability p . The Fourier basis we use for analysis on G is the set of p -biased Fourier characters (which are naturally indexed by graphs H on $[n]$).

Definition 2.3. χ denotes the p -biased Fourier character,

$$\chi(0) = \sqrt{\frac{p}{1-p}}, \quad \chi(1) = -\sqrt{\frac{1-p}{p}}.$$

For H a subset or multi-subset of $\binom{[n]}{2}$, let $\chi_H(G) := \prod_{e \in H} \chi(G_e)$.

We will also need the function $1 - G_e$ which indicates that an edge is not present.

Definition 2.4. For $H \subseteq \binom{[n]}{2}$, let $1_{\overline{H}}(G) = \prod_{e \in H} (1 - G_e)$.

When H is a clique, this is the *independent set indicator* for the vertices in H .

Proposition 2.5. For $e \in \{0, 1\}$, $1 + \sqrt{\frac{p}{1-p}}\chi(e) = \frac{1}{1-p}(1 - e)$. Therefore, for any $H \subseteq \binom{[n]}{2}$,

$$\sum_{T \subseteq H} \left(\frac{p}{1-p} \right)^{|T|/2} \chi_T(G) = \frac{1}{(1-p)^{|H|}} \cdot 1_{\overline{H}}(G).$$

2.3 Ribbons and graph matrices

A degree- D pseudoexpectation operator is a vector in $\mathbb{R}^{\binom{[n]}{\leq D}}$. The matrices we consider will have rows and columns indexed by all subsets of $[n]$. We express the moment matrix Λ in terms of the Fourier basis on $G_{n,p}$. A particular Fourier character in a particular matrix entry is identified by a combinatorial structure called a *ribbon*.

Definition 2.6 (Ribbon). A ribbon is a tuple $R = (V(R), E(R), A_R, B_R)$, where $(V(R), E(R))$ is an undirected multigraph without self-loops, $V(R) \subseteq [n]$, and $A_R, B_R \subseteq V(R)$. Let $C_R := V(R) \setminus (A_R \cup B_R)$.

Definition 2.7 (Matrix for a ribbon). For a ribbon R , the matrix M_R has rows and columns indexed by all subsets of $[n]$ and has a single nonzero entry,

$$M_R[I, J] = \begin{cases} \chi_{E(R)}(G) & I = A_R, J = B_R \\ 0 & \text{Otherwise} \end{cases}$$

Definition 2.8 (Ribbon isomorphism). Two ribbons R, S are isomorphic, or have the same shape, if there is a bijection between $V(R)$ and $V(S)$ which is a multigraph isomorphism between $E(R), E(S)$ and is a bijection from A_R to A_S and B_R to B_S . Equivalently, letting S_n permute the vertex labels of a ribbon, the two ribbons are in the same S_n -orbit.

If we ignore the labels on the vertices of a ribbon, what remains is the *shape* of the ribbon.

Definition 2.9 (Shape). A shape is an equivalence class of ribbons with the same shape. Each shape has associated with it a representative $\alpha = (V(\alpha), E(\alpha), U_\alpha, V_\alpha)$, where $U_\alpha, V_\alpha \subseteq V(\alpha)$. Let $W_\alpha := V(\alpha) \setminus (U_\alpha \cup V_\alpha)$.

Definition 2.10 (Embedding). Given a shape α and an injective function $\varphi : V(\alpha) \rightarrow [n]$, we let $\varphi(\alpha)$ be the ribbon obtained by labeling α in the natural way.

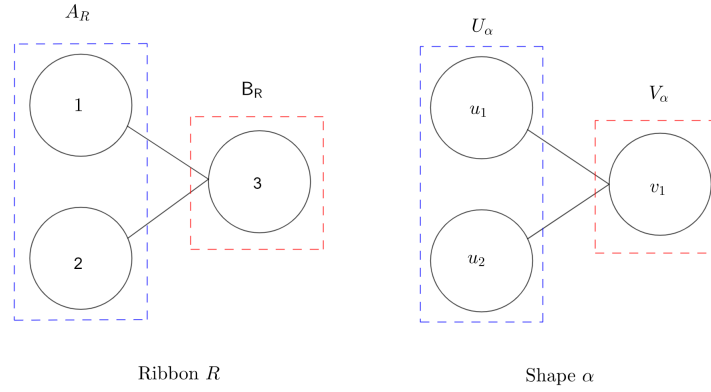


Figure 1: An example of a ribbon and a shape

Definition 2.11 (Graph matrix). For a shape α , the graph matrix M_α is

$$M_\alpha = \sum_{\text{injective } \varphi: V(\alpha) \rightarrow [n]} M_{\varphi(\alpha)}.$$

Injectivity is an important property of graph matrices. On the one hand, we have a finer partition of ribbons than allowing all assignments, and this allows more control. On the other hand, injectivity introduces technically challenging “intersection terms” into graph matrix multiplication. A graph matrix is essentially a sum over all ribbons with shape α (this is not entirely accurate as each ribbon will be repeated $|\text{Aut}(\alpha)|$ times).

Definition 2.12 (Automorphism). For a shape α , $\text{Aut}(\alpha)$ is the group of bijections from $V(\alpha)$ to itself such that U_α and V_α are fixed as sets and the map is a multigraph automorphism on $E(\alpha)$. Equivalently, $\text{Aut}(\alpha)$ is the stabilizer subgroup (of S_n) of any ribbon of shape α .

Fact 2.13.

$$M_\alpha = \sum_{\text{injective } \varphi: V(\alpha) \rightarrow [n]} M_{\varphi(\alpha)} = |\text{Aut}(\alpha)| \sum_{R \text{ ribbon of shape } \alpha} M_R$$

Example 2.14 (Ribbon). As an example, consider the ribbon in Fig. 1. We have $A_R = \{1, 2\}$, $B_R = \{3\}$, $V(R) = \{1, 2, 3\}$, $E(R) = \{\{1, 2\}, \{2, 3\}\}$. The Fourier character is $\chi_{E_R} = \chi_{1,3}\chi_{2,3}$. And finally, M_R is a matrix with rows and columns indexed by subsets of $[n]$, with exactly one nonzero entry $M_R(\{1, 2\}, \{3\}) = \chi_{\{1,3\}}\chi_{\{2,3\}}$. Succinctly,

$$M_R = \begin{matrix} & \text{column } \{3\} \\ & \downarrow \\ \text{row } \{1, 2\} \rightarrow & \begin{pmatrix} 0 & \vdots & 0 \\ \cdots \cdots \chi_{1,3}\chi_{2,3} \cdots \cdots \\ 0 & \vdots & 0 \end{pmatrix} \end{matrix}$$

Example 2.15 (Shape). In Fig. 1, consider the shape α as shown. We have $U_\alpha = \{u_1, u_2\}$, $V_\alpha = \{v_1\}$, $W_\alpha = \emptyset$, $V(\alpha) = \{u_1, u_2, v_1\}$, and $E(\alpha) = \{\{u_1, v_1\}, \{u_2, v_1\}\}$. M_α is a matrix with rows and

columns indexed by subsets of $[n]$. The nonzero entries will have rows and columns indexed by $\{a_1, a_2\}$ and b_1 respectively for all distinct $a_1, a_2, b_1 \in [n]$, with the corresponding entry being $M_\alpha(\{a_1, a_2\}, \{b_1\}) = \chi_{a_1, b_1} \chi_{a_2, b_1}$. Succinctly,

$$M_\alpha = \begin{matrix} & & \text{column } \{b_1\} \\ & & \downarrow \\ & & \vdots \\ \text{row } \{a_1, a_2\} \rightarrow & \left(\begin{matrix} \cdots \cdots \chi_{a_1, b_1} \chi_{a_2, b_1} \cdots \cdots \\ \vdots \end{matrix} \right) \\ & & \vdots \end{matrix}$$

Definition 2.16 (Proper). A ribbon or shape is proper if it has no multi-edges. Otherwise, it is improper. Let $\text{mul}_\alpha(e)$ be the multiplicity of edge e in ribbon or shape α .

An improper ribbon or shape with an edge e of multiplicity 2, e.g., has a squared Fourier character χ_e^2 . Since this is a function on $\{0, 1\}$, by expressing it in the Fourier basis an improper ribbon or shape can be decomposed in a unique way into a linear combination of proper ones, which we call *linearizations*.

Definition 2.17 (Linearization). Given an improper ribbon or shape α , a linearization β is a proper ribbon or shape such that $\text{mul}_\beta(e) \leq \text{mul}_\alpha(e)$ for all $e \in E(\alpha)$.

Definition 2.18 (Isolated vertex). For a shape α , an isolated vertex is a degree-0 vertex in W_α . Let I_α denote the set of isolated vertices in α . Similarly, for a ribbon R , the isolated vertices are denoted I_R .

We stress that an isolated vertex never refers to degree-0 vertices inside $U_\alpha \cup V_\alpha$.

Definition 2.19 (Trivial shape). A shape α is trivial if $V(\alpha) = U_\alpha = V_\alpha$ and $E(\alpha) = \emptyset$.

M_α for a trivial α is the identity matrix restricted to the degree- $|U_\alpha|$ block.

Definition 2.20 (Transpose). Given a ribbon R or shape α , we define its transpose by swapping A_R and B_R (resp. U_α and V_α). Observe that this transposes the matrix for the ribbon/shape.

3 An overview of the proof techniques

Here, we will give a sketch of the proof techniques that we utilize in our SoS lower bound. Recall that we are given a graph $G \sim G_{n,p}$ where $d = pn$ is the average degree and our goal is to show that for any constant $\varepsilon > 0$, $D_{\text{SoS}} \approx n^\delta$ for some $\delta > 0$, degree D_{SoS} SoS thinks there exists an independent set of size $k \approx \frac{n}{d^{1/2+\varepsilon}(D_{\text{SoS}} \log n)^{c_0}}$ whereas the true independent set has size $\approx \frac{n \log d}{d}$ for some absolute constant c_0 .

To prove the lower bound, we review the Planted Clique lower bound [BHK⁺16] and describe the obstacles that need to be overcome in the sparse setting.

3.1 Modified pseudocalibration

Since SoS is a convex program, the goal of an SoS lower bound is to construct a dual object: a set of *pseudomoments* $\tilde{\mathbb{E}}[x^S]$ for each small $S \subseteq V(G)$, which are summarized in the *moment matrix*. The

moment matrix must (i) obey the problem constraints (ii) be SoS-symmetric, and (iii) be positive semidefinite (PSD). Following the recipe of pseudocalibration introduced by [BHK⁺16], we can produce a candidate moment matrix which is guaranteed to satisfy the first two conditions, while like all other SoS lower bounds, the hard work remains in verifying the PSDness of the moment matrix. Pseudocalibration has been successfully exploited in a multitude of SoS lower bound applications, e.g., [BHK⁺16, KMOW17, MRX20, GJJ⁺20, PR20]. So, this is a natural starting point for us.

Failure of pseudocalibration The first obstacle we overcome is the lack of a planted distribution. Pseudocalibration requires a planted and random distribution which are hard to distinguish using the *low-degree, likelihood ratio test* (i.e. $\tilde{\mathbb{E}}[1]$ is bounded whp) [HKP⁺17, Hop18]. In the case of sparse independent set, we have the following natural hypothesis testing problem with which one may hope to pseudocalibrate.

- Null Hypothesis: Sample a graph $G \sim G_{n,p}$.
- Alternate Hypothesis: Sample a graph $G \sim G_{n,p}$. Then, sample a subset $S \subseteq [n]$ where each vertex is chosen with probability $\frac{k}{n}$. Then, plant an independent set in S , i.e. remove all the edges inside S .

In the case of sparse independent set, the naïve planted distribution *is* distinguishable from a random instance via a simple low-degree test – counting 4-cycles. In all uses of pseudocalibration that we are aware of, the two distributions being compared are conjecturally hard to distinguish by all polynomial-time algorithms. We are still searching for a suitable planted distribution for sparse independent set, and we believe this is an interesting question on its own.

Fixing pseudocalibration via connected truncation To get around with this issue, we close our eyes and “pretend” the planted distribution is quiet, ignoring the obvious distinguisher, and make a “connected truncation” of the moment matrix to remove terms which correspond to counting subgraphs in G . What remains is that $\tilde{\mathbb{E}}[x^S]$ is essentially independent of the global statistics of G . It should be pointed out here that this is inherently distinct from the local truncation for weaker hierarchies (e.g. Sherali-Adams) where the moment matrix is an entirely local function [CM18]. In contrast, our $\tilde{\mathbb{E}}[x^S]$ may depend on parts of the graph that are far away from S , in fact, even up to radius n^δ , exceeding the diameter of the random graph!

At this point, the candidate moment matrix can be written as follows.

$$\Lambda := \sum_{\alpha \in \mathcal{S}} \left(\frac{k}{n}\right)^{|V(\alpha)|} \cdot \left(\frac{p}{1-p}\right)^{\frac{|E(\alpha)|}{2}} M_\alpha.$$

Here, \mathcal{S} ranges over all proper shapes α of appropriately bounded size such that *all vertices of α are connected to $U_\alpha \cup V_\alpha$* . The latter property is the important distinction from standard pseudocalibration and will turn out to be quite essential for our analysis.

Using connected objects to take advantage of correlation decay is also a theme in the cluster expansion from statistical physics (see Chapter 5 of [FV18]). Although not formally connected with connected truncation, the two methods share some similar characteristics.

3.2 Approximate PSD decompositions, norm bounds and conditioning

Continuing, to show the moment matrix is PSD, Planted Clique [BHK⁺16] performs an approximate factorization of the moment matrix in terms of graph matrices. A crucial part of this approach is to identify "dominant" and "non-dominant" terms in the approximate PSD decomposition. Then, the dominant terms are shown to be PSD and the non-dominant terms are bounded against the dominant terms. In this approach, a crucial component in the latter step is to control the norms of graph matrices.

Tighter norm bounds for sparse graph matrices Existing norm bounds in the literature [MP16, AMP20] for graph matrices have focused exclusively on the dense setting $G_{n,1/2}$. Unfortunately, while these norm bounds apply for the sparse setting, they're too weak to be useful. Consider the case where we sample $G \sim G_{n,p}$ and try to bound the spectral norm of the centered adjacency matrix. Existing norm bounds give a bound of $\tilde{O}\left(\frac{\sqrt{n(1-p)}}{\sqrt{p}}\right)$ whereas the true norm is $O(\sqrt{n})$ regardless of d . This is even more pronounced when we use shapes with more vertices. So, our first step is to tighten the existing norm bounds in the literature for sparse graph matrices.

For a shape α , a vertex separator S is a subset of vertices such that there are no paths from U_α to V_α in $\alpha \setminus S$. It is known from previous works that in the dense case the spectral norm is controlled by the number of vertices in the minimum vertex separator between U_α and V_α . Assuming α does not have isolated vertices and $U_\alpha \cap V_\alpha = \emptyset$ for simplicity, the norm of M_α is given by the following expression, up to polylog factors and the leading coefficient of at most $|V(\alpha)|^{|V(\alpha)|}$,

$$\|M_\alpha\| \leq \max_{\text{vertex separator } S} \sqrt{n}^{|V(\alpha)|-|V(S)|}$$

However, it turns out this is no longer the controlling quantity if the underlying input matrix is sparse, and tightly determining this quantity arises as a natural task for our problem, and for future attack on SoS lower bounds for other problems in the sparse regime. To motivate the difference, we want to point out this is essentially due to the following simple observation. For $k \geq 2$,

$$\begin{aligned} |\mathbb{E}[X^k]| &= 1 \\ |\mathbb{E}[Y^k]| &\leq \left(\sqrt{\frac{1-p}{p}}\right)^{k-2} \approx \left(\sqrt{\frac{n}{d}}\right)^{k-2} \end{aligned}$$

for X a uniform ± 1 bit and Y a p -biased random variable $\text{Ber}(p)$. This suggests that in the trace power method, there will be a preference among vertex separators of the same size if some contain more edges inside the separator (because vertices inside vertex separators are "fixed" in the dominant term in the trace calculation, and thus edges within the separator will contribute some large power of $\sqrt{\frac{n}{d}}$, creating a noticeable influence on the final trace). Finally, this leads us to the following characterization for sparse matrix norm bounds, up to polylog factors and the leading coefficient of at most $|V(\alpha)|^{|V(\alpha)|}$,

$$\|M_\alpha\| \leq \max_{\text{vertex separator } S} \sqrt{n}^{|V(\alpha)|-|V(S)|} \left(\sqrt{\frac{1-p}{p}}\right)^{|E(S)|}$$

The formal statement is given in [Section 6.10](#). We prove this via an application of the trace method followed by a careful accounting of the large terms.

The key conceptual takeaway is that we need to redefine the weight of a vertex separator to also incorporate the edges within the separator, as opposed to only considering the number of vertices. We clearly distinguish these with the terms *Dense Minimum Vertex Separator* (DMVS) and *Sparse Minimum Vertex Separator* (SMVS). When $p = \frac{1}{2}$, these two bounds are the same up to lower order factors.

Approximate PSD decomposition We then perform an approximate PSD decomposition of the graph matrices that make up Λ . The general factoring strategy is the same as [BHK⁺16], though in the sparse regime we must be very careful about what kind of combinatorial factors we allow. Each shape comes with a natural “vertex decay” coefficient arising from the fractional size of the independent set and an “edge decay” coefficient arising from the sparsity of the graph. The vertex decay coefficients can be analyzed in a method similar to Planted Clique (which only has vertex decay). For the edge decay factors, we use novel charging arguments. At this point, the techniques are strong enough to prove [Theorem 1.8](#), an SoS-degree $\Omega(\log n)$ lower bound for $d \geq n^\epsilon$. The remaining techniques are needed to push the SoS degree up and the graph degree down.

Conditioning In our analysis, it turns out that to obtain strong SoS lower bounds in the sparse regime, a norm bound from the vanilla trace method is not quite sufficient. Sparse random matrices’ spectral norms are fragile with respect to the influence of an unlikely event, exhibiting deviations away from the expectation with polynomially small probability (rather than exponentially small probability, like what is obtained from a good concentration bound). These “bad events” are small dense subgraphs present in a graph sampled from $G_{n,p}$.

To get around this, we condition on the high probability event that G has no small dense subgraphs. For example, for $d = n^{1-\epsilon}$ whp every small subgraph S has $O(|S|)$ edges (even up to size n^δ). For a **shape** which is dense (i.e. v vertices and more than $O(v)$ edges) we can show that its norm falls off extremely rapidly under this conditioning. This allows us to throw away dense shapes, which is critical for controlling combinatorial factors that would otherwise dominate the moment matrix.

This type of conditioning is well-known: a long line of work showing tight norm bounds for the simple adjacency matrix appeals to a similar conditioning argument within the trace method [BLM15, Bor19, FM17, DMO⁺19]. We instantiate the conditioning in two ways. The first is through the following identity.

Observation 3.1. *Given a set of edges $E \subseteq \binom{[n]}{2}$, if we know that not all of the edges of E are in $E(G)$ then*

$$\chi_E(G) = \sum_{E' \subseteq E: E' \neq E} \left(\sqrt{\frac{p}{1-p}} \right)^{|E|-|E'|} \chi_{E'}(G)$$

This simple observation whose proof is deferred to [Section 6.5](#) can be applied recursively to replace a dense shape α by a sum of its sparse subshapes $\{\beta\}$. The second way we eliminate dense shapes is by using a bound on the Frobenius norm which improves on the trace calculation for dense shapes. After conditioning, we can restrict our attention to sparse shapes, which allows us to avoid several combinatorial factors which would otherwise overwhelm the “charging” argument.

Handling the subshapes $\{\beta\}$ requires some care. Destroying edges from a shape can cause its norm to either go up or down: the vertex separator gets smaller (increasing the norm), but if we remove edges from inside the SMVS, the norm goes down. An important observation is we do

not necessarily have to apply [Observation 3.1](#) on the entire set of edges of a shape, but we can also just apply it on some of the edges. We will choose a set of edges $\text{Res}(\alpha) \subseteq E(\alpha)$ that “protects the minimum vertex separator” and only apply conditioning on edges outside $\text{Res}(\alpha)$. In this way the norm of subshapes β will be guaranteed to be less than α . The fact that it’s possible to reserve such edges is shown separately for the different kind of shapes we encounter in our analysis (see [Appendix B](#)).

Finally, we are forced to include certain dense shapes that encode indicator functions of independent sets. These shapes must be factored out and tracked separately throughout the analysis.⁴ After handling all of these items we have shown that Λ is PSD.

4 Pseudocalibration with connected truncation

4.1 Pseudo-calibration

Pseudo-calibration is a heuristic introduced by [\[BHK⁺16\]](#) to construct candidate moment matrices to exhibit SoS lower bounds. Here, we formally describe the heuristic briefly.

To pseudocalibrate, a crucial starting point is to come up with a hard hypothesis testing problem. Let ν denote the null distribution and μ denote the alternative distribution. Let v denote the input and x denote the variables for our SoS relaxation. The main idea of pseudo-calibration is that, for an input v sampled from ν and any polynomial $f(x)$ of degree at most the SoS degree, and for any low-degree test $g(v)$, the correlation of $\tilde{\mathbb{E}}[f(x)]$ with g as functions of v should match in the planted and random distributions. That is,

$$\mathbb{E}_{v \sim \nu}[\tilde{\mathbb{E}}[f(x)]g(v)] = \mathbb{E}_{(x,v) \sim \mu}[f(x)g(v)]$$

Here, the notation $(x, v) \sim \mu$ means that in the planted distribution μ , the input is v and x denotes the planted structure in that instance. For example, in independent set, x would be the indicator vector of the planted independent set.

Let \mathcal{F} denote the Fourier basis of polynomials for the input v . By choosing g from \mathcal{F} such that the degree is at most the truncation parameter $D_V = CD_{\text{SoS}} \log n$ (hence the term “low-degree test”), the pseudo-calibration heuristic specifies all lower order Fourier coefficients for $\tilde{\mathbb{E}}[f(x)]$ as a function of v . The heuristic suggests that the higher order coefficients are set to be 0. In total the candidate pseudoexpectation operator can be written as

$$\tilde{\mathbb{E}}[f(x)] = \sum_{\substack{g \in \mathcal{F}: \\ \text{deg}(g) \leq D_V}} \mathbb{E}_{v \sim \nu}[\tilde{\mathbb{E}}[f(x)]g(v)]g(v) = \sum_{\substack{g \in \mathcal{F}: \\ \text{deg}(g) \leq D_V}} \mathbb{E}_{(x,v) \sim \mu}[f(x)g(v)]g(v)$$

The Fourier coefficients $\mathbb{E}_{(x,v) \sim \mu}[f(x)g(v)]$ can be explicitly computed in many settings, which therefore gives an explicit pseudoexpectation operator $\tilde{\mathbb{E}}$.

An advantage of pseudo-calibration is that this construction automatically satisfies some nice properties that the pseudoexpectation $\tilde{\mathbb{E}}$ should satisfy. It’s linear in v by construction. For all polynomial equalities of the form $f(x) = 0$ that are satisfied in the planted distribution, it’s true that $\tilde{\mathbb{E}}[f(x)] = 0$. For other polynomial equalities of the form $f(x, v) = 0$ that are satisfied in the

⁴One may ask if we could definitionally get rid of dense shapes, like we did for disconnected shapes via the connected truncation, but these dense shapes are absolutely necessary.

planted distribution, the equality $\tilde{\mathbb{E}}[f(x, v)] = 0$ is approximately satisfied. In many cases, $\tilde{\mathbb{E}}$ can be mildly adjusted to satisfy these exactly.

The condition $\tilde{\mathbb{E}}[1] = 1$ is not automatically satisfied. In all previous successful applications of pseudo-calibration, $\tilde{\mathbb{E}}[1] = 1 \pm o(1)$. Once we have this, we simply set our final pseudoexpectation operator to be $\tilde{\mathbb{E}}'$ defined as $\tilde{\mathbb{E}}'[f(x)] = \tilde{\mathbb{E}}[f(x)] / \tilde{\mathbb{E}}[1]$. We remark that the condition $\tilde{\mathbb{E}}[1] = 1 \pm o(1)$ has been quite successful in predicting distinguishability of the planted/null distributions [HKP⁺17, Hop18].

4.2 The failure of "Just try pseudo-calibration"

Despite the power of pseudocalibration in guiding the construction of a pseudomoment matrix in SoS lower bounds, it heavily relies upon a planted distribution that is hard to distinguish from the null distribution. Unfortunately, a "quiet" planted distribution remains on the search in our setting.

Towards this end, we will consider the following "naïve" planted distribution \mathcal{D}_{pl} that likely would have been many peoples' first guess:

- (1) Sample a random graph $G \sim G_{n,p}$;
- (2) Sample a subset $S \subseteq [n]$ by picking each vertex with probability $\frac{k}{n}$;
- (3) Let \tilde{G} be G with edges inside S removed, and output (S, \tilde{G}) .

Proposition 4.1. *For all $d = O(\sqrt{n})$ and $k = \Omega(n/d)$, the naïve planted distribution \mathcal{D}_{pl} is distinguishable in polynomial time from $G_{n,p}$ with $\Omega(1)$ probability.*

Proof. The number of labeled 4-cycles in $G_{n,d/n}$ has expectation $\mathbb{E}[C_4] = \frac{d^4}{n^4} \cdot n(n-1)(n-2)(n-3)$ and variance $\text{Var}[C_4] = O(d^4)$. In \mathcal{D}_{pl} the expected number of labeled 4-cycles is

$$\begin{aligned} \mathbb{E}_{pl}[C_4] &= \frac{d^4}{n^4} \cdot n(n-1)(n-2)(n-3) \cdot \left(\left(1 - \frac{k}{n}\right)^4 + 4\frac{k}{n} \left(1 - \frac{k}{n}\right)^3 + 2\left(\frac{k}{n}\right)^2 \left(1 - \frac{k}{n}\right)^2 \right) \\ &= \frac{d^4}{n^4} \cdot n(n-1)(n-2)(n-3) \cdot (1 - O(k^2/n^2)) \\ &= \mathbb{E}[C_4] - O(d^4 k^2/n^2) = \mathbb{E}[C_4] - O(d^2) \end{aligned}$$

Since this is less than $\mathbb{E}[C_4]$ by a factor on the order of $\sqrt{\text{Var}[C_4]}$, counting 4-cycles succeeds with constant probability. ■

Remark 4.2. *To beat distinguishers of this type, it may be possible to construct the planted distribution from a sparse quasirandom graph (in the sense that all small subgraph counts match $G_{n,p}$ to leading order) which has an independent set size of $\Omega(n/\sqrt{d})$. In the dense setting, the theory of quasirandom graphs states that if the planted distribution and $G_{n,1/2}$ match the subgraph count of C_4 , this is sufficient for all subgraph counts to match; in the sparse setting this is no longer true and the situation is more complicated [CG02].*

Remark 4.3. *Coja-Oghlan and Efthymiou [COE15] show that a slight modification of this planted distribution (correct the expected number of edges to be $p\binom{n}{2}$) is indistinguishable from the random distribution provided k is slightly smaller than the size of the maximum independent set in $G_{n,p}$. This is not useful for pseudocalibration because we are trying to plant an independent set with larger-than-expected size.*

The Fourier coefficients of the planted distribution are:

Lemma 4.4. *Let $x_T(S)$ be the indicator function for T being in the planted solution i.e. $T \subseteq S$. Then, for all $T \subseteq [n]$ and $\alpha \subseteq \binom{[n]}{2}$,*

$$\mathbb{E}_{(S, \tilde{G}) \sim \mathcal{D}_{pl}} [x_T(S) \cdot \chi_\alpha(\tilde{G})] = \left(\frac{k}{n}\right)^{|V(\alpha) \cup T|} \left(\frac{p}{1-p}\right)^{\frac{|E(\alpha)|}{2}}$$

Proof. First observe that if any vertex of $V(\alpha) \cup T$ is outside S , then the expectation is 0. This is because either T is outside S , in which case $x_T = 0$, or a vertex of α is outside S , in which case the expectation of any edge incident on this vertex is 0 so the entire expectation is 0 using independence. Now, each vertex of $V(\alpha) \cup T$ is in S independently with probability $\left(\frac{k}{n}\right)^{|V(\alpha) \cup T|}$.

Conditioned on this event happening, the character is simply $\chi_\alpha(0) = \left(\frac{p}{1-p}\right)^{\frac{|E(\alpha)|}{2}}$. Putting them together gives the result. ■

This planted distribution motivates the following incorrect definition of the pseudo-expectation operator.

Definition 4.5 (Incorrect definition of $\tilde{\mathbb{E}}$). *For $S \subseteq [n], |S| \leq D_{SoS}$,*

$$\tilde{\mathbb{E}}[x^S] := \sum_{\substack{\alpha \subseteq \binom{[n]}{2}: \\ |\alpha| \leq D_V}} \left(\frac{k}{n}\right)^{|V(\alpha) \cup S|} \left(\frac{p}{1-p}\right)^{\frac{|E(\alpha)|}{2}} \chi_\alpha(G)$$

For this operator, $\tilde{\mathbb{E}}[1] = 1 + \Omega(1)$. More generally, tail bounds of graph matrix sums are not small, which ruins our analysis technique.

4.3 Salvaging the doomed

We will now discuss our novel truncation heuristic. The next paragraphs are for discussion and the technical proof resumes at [Definition 4.7](#). Letting \mathcal{S} be the set of shapes α that contribute to the moment $\tilde{\mathbb{E}}[x^I]$, the ‘‘connected truncation’’ restricts \mathcal{S} to shapes α such that all vertices in α have a path to I (or in the shape view, to $U_\alpha \cup V_\alpha$).

Why might this be a good idea? Consider the planted distribution. The only tests we know of to distinguish the random/planted distributions are counting the number of edges or counting occurrences of small subgraphs. These tests cannot be implemented using only connected Fourier characters; shapes with disconnected middle vertices are needed to count small subgraphs. For example, suppose we fix a particular vertex $v \in V(G)$, and we consider the set of functions on G which are allowed to depend on v but are otherwise symmetric in the vertices of G . A basis for these functions can be made by taking shapes α such that U_α has a single vertex, $V_\alpha = \emptyset$, then fixing $U_\alpha = v$ (i.e. take the vector entry in row v). Let sym_v denote this set of functions.

Proposition 4.6. *Let $T(G)$ be the triangle counting function on G . Let $\text{conn}(\text{sym}_v)$ be the subset of sym_v such that all vertices have a path to v . Then $T(G) \notin \text{span}(\text{conn}(\text{sym}_v))$.*

Proof. $T(G)$ has a unique representation in sym_v which requires disconnected Fourier characters. For example, one component of the function is $T_{\bar{v}}(G) = \text{number of triangles not containing } v$. This is three vertices x, y, z outside of U_α with $1_{(x,y) \in E(G)} 1_{(x,z) \in E(G)} 1_{(y,z) \in E(G)}$. The edge indicator function is implemented by [Proposition 2.5](#). But there are no edges between x, y, z and $U_\alpha = v$ in these shapes.

It's not even possible to implement $T_v(G) = \text{number of triangles containing } v$, as a required shape is two vertices x, y outside of U_α connected with an edge. ■

Despite the truncation above, our pseudocalibration operator is not a "local function" of the graph. Our $\tilde{\mathbb{E}}[x^S]$ can depend on vertices that are far away from S , but in an attenuated way. The graph matrix for α is a sum of all ways of overlaying the vertices of α onto G . The edges do not need to be overlaid. If an edge "misses" in G , then we can use this edge to get far away from S , but we take a decay factor of $\chi_e(0) = \sqrt{\frac{p}{1-p}} = \sqrt{\frac{d}{n-d}}$.

The "local function" property of $\tilde{\mathbb{E}}[x^S]$ is also a connected truncation, but it is a connected truncation in a different basis. The basis is the 0/1 basis of $1_H(G)$ for $H \subseteq \binom{[n]}{2}$. A reasonable definition of graph matrices in this basis is

$$M_\alpha = \sum_{\text{injective } \sigma: V(\alpha) \rightarrow [n]} 1_{\sigma(E(\alpha))}(G)$$

which sums all ways to embed α into G . $\tilde{\mathbb{E}}[x^S]$ is a local function if and only if in this basis it is a sum of shapes α satisfying the (same) condition that all vertices are connected to $S = U_\alpha \cup V_\alpha$.

For sparse graphs, the two bases are somewhat heuristically interchangeable since:

$$\frac{1}{p} 1_e(G) = 1 - \sqrt{\frac{1-p}{p}} \chi_e(G) \approx -\sqrt{\frac{1-p}{p}} \chi_e(G).$$

Comparing the 0/1 basis and the Fourier basis, the 0/1 basis expresses combinatorial properties such as subgraph counts more nicely, while spectral analysis is only feasible in the Fourier basis. In the proof (see [Definition 6.21](#)), we will augment ribbons so that they may also contain 0/1 indicators, and this flexibility helps us overcome both the spectral and combinatorial difficulties in the analysis.

Formally we define the candidate moment matrix as:

Definition 4.7 (Moment matrix).

$$\Lambda := \sum_{\alpha \in \mathcal{S}} \binom{k}{n}^{|V(\alpha)|} \cdot \left(\frac{p}{1-p} \right)^{\frac{|E(\alpha)|}{2}} \frac{M_\alpha}{|\text{Aut}(\alpha)|}.$$

where \mathcal{S} is the set of proper shapes such that

1. $|U_\alpha|, |V_\alpha| \leq D_{\mathcal{S}oS}$.
2. $|V(\alpha)| \leq D_V$ where $D_V = CD_{\mathcal{S}oS} \log n$ for a sufficiently large constant C .
3. Every vertex in α has a path to $U_\alpha \cup V_\alpha$.

We refer to $\left(\frac{k}{n}\right)^{|V(\alpha)|}$ as the “vertex decay factor” and $\left(\frac{p}{1-p}\right)^{\frac{|E(\alpha)|}{2}}$ as the “edge decay factor”. The candidate moment matrix is the principal submatrix indexed by $(\leq D_{\text{SoS}}/2)$. The non-PSDness properties of Λ are easy to verify:

Lemma 4.8. Λ is SoS-symmetric.

Proof. This is equivalent to the fact that the coefficient of α does not depend on how $U_\alpha \cup V_\alpha$ is partitioned into U_α and V_α for $|U_\alpha \cup V_\alpha| \leq D_{\text{SoS}}$. ■

Lemma 4.9. $\tilde{\mathbb{E}}[1] = 1$.

Proof. For $U_\alpha = V_\alpha = \emptyset$, the only shape in which all vertices are connected to $U_\alpha \cup V_\alpha$ is the empty shape. Therefore $\tilde{\mathbb{E}}[1] = 1$. ■

Lemma 4.10. $\tilde{\mathbb{E}}$ satisfies the feasibility constraints.

Proof. We must show that $\tilde{\mathbb{E}}[x^S] = 0$ whenever S is not an independent set in G . Observe that if ribbon R contributes to $\tilde{\mathbb{E}}[x^S]$, then if we modify the set of edges inside S , the resulting ribbon still contributes to $\tilde{\mathbb{E}}[x^S]$. In fact, each edge also comes with a factor of $\sqrt{\frac{p}{1-p}}$. By [Proposition 2.5](#), we can group these ribbons into an indicator function $\frac{1}{(1-p)^{\binom{|S|}{2}}} 1_{E(S)}$. That is, $\tilde{\mathbb{E}}[x^S] = 0$ if S has an edge. ■

Lemma 4.11. With probability at least $1 - o_n(1)$, $\tilde{\mathbb{E}}$ has objective value $\tilde{\mathbb{E}}[\sum x_i] \geq (1 - o(1))k$.

Proof.

Claim 4.12.

$$\mathbb{E}[\tilde{\mathbb{E}}[x_i]] = \frac{k}{n}$$

Proof. The only shape that survives under expectation is the shape with one vertex, and it comes with coefficient $\frac{k}{n}$. ■

Claim 4.13.

$$\text{Var} \left[\tilde{\mathbb{E}}[x_i] \right] \leq d^{-\Omega(\epsilon)}$$

Proof. Let $\text{Count}(v, e)$ be the number of shapes with $|U_\alpha| = 1, V_\alpha = \emptyset, v$ vertices and e edges,

$$\begin{aligned} \text{Var} \left[\tilde{\mathbb{E}}[x_i] \right] &= \sum_{\alpha \neq \emptyset: \text{connected to } i} \left(\left(\frac{k}{n} \right)^{|V(\alpha)|} \left(\frac{p}{1-p} \right)^{|E(\alpha)|/2} \right)^2 \\ &= \sum_{v=2}^{D_V} \sum_{v-1 \leq e \leq v^2} \left(\frac{k}{n} \right)^{2v} \left(\frac{p}{1-p} \right)^e \cdot \text{Count}(v, e) \cdot n^v \\ &\leq O(1) \sum_{v=2}^{D_V} \left(\frac{4k\sqrt{d}}{(1-p)n} \right)^v \\ &\leq O \left(\frac{k\sqrt{d}}{n} \right) \end{aligned}$$

where the first inequality follows by observing the dominant term is tree-like and bounding the number of trees (up-to-isomorphism) with v vertices by 2^{2v} . ■

Hence,

$$\Pr\left[\left|\sum \tilde{\mathbb{E}}[x_i] - k\right| \geq t\right] \leq \frac{\text{Var}[\sum \tilde{\mathbb{E}}[x_i]]}{t^2} = \frac{k\sqrt{d}}{t^2} \leq o_n(1)$$

where the last inequality follows by picking $t = n^{\frac{1}{2}+\delta} = o(k)$ for $\delta > 0$. ■

What remains is to show $\Lambda \succeq 0$. To do this it's helpful to renormalize the matrix entries by multiplying the degree- (k, l) block by a certain factor.

Definition 4.14 (Shape coefficient). *For all shapes α , let*

$$\lambda_\alpha := \left(\frac{k}{n}\right)^{|V(\alpha)| - \frac{|U_\alpha| + |V_\alpha|}{2}} \cdot \left(\frac{p}{1-p}\right)^{\frac{|E(\alpha)|}{2}}.$$

It suffices to show $\sum_{\alpha \in \mathcal{S}} \lambda_\alpha \frac{M_\alpha}{|\text{Aut}(\alpha)|} \succeq 0$ because left and right multiplying by a rank-1 matrix and its transpose returns Λ .

Proposition 4.15. *If α, β are composable shapes, then $\lambda_\alpha \lambda_\beta = \lambda_{\alpha \circ \beta}$.*

5 Norm bounds and factoring

As explained in [Section 3](#), existing norm bounds in the literature are not tight enough for sparse graph matrices. We first obtain tighter norm bounds for sparse graph matrices.

5.1 Norm bounds

In this section, we show the spectral norm bounds for M_α in terms of simple combinatorial factors involving α . With only log factor loss the norm bounds hold with very high probability (all but probability $n^{-\Omega(\log n)}$). This is too tight of a probabilistic bound since it allows for polynomially rare events such as small dense subgraphs. We will need to use conditioning to improve the norm bound for shapes with a lot of edges.

Definition 5.1. *For a shape α , define a vertex separator (or a separator) S to be a subset of vertices such that there is no path from U_α to V_α in the graph $\alpha \setminus S$.*

Roughly, the norm bounds for a proper shape α are:

$$\|M_\alpha\| \leq \tilde{O} \left(\max_{\text{vertex separator } S} \left\{ \sqrt{n}^{|V(\alpha)| - |S|} \sqrt{\frac{1-p}{p}}^{|E(S)|} \right\} \right).$$

The maximizer of the above is called the *sparse minimum vertex separator*. The proof of this bound uses the trace method, which also underlies graph matrix norm bounds in the dense case. We defer the proof to [Appendix A](#).

To get norm bounds for improper shapes, we linearize the shape and take the largest norm bound among its linearizations.

Definition 5.2. For a linearization β of a shape α , $E_{\text{phantom}}(\beta)$ is the set (not multiset) of “phantom edges” of α which are not in β .

For an improper shape α :

$$\|M_\alpha\| \leq \tilde{O} \left(\max_{\beta, S} \left\{ \sqrt{n}^{|V(\alpha)| - |S| + |I_\beta|} \left(\sqrt{\frac{1-p}{p}} \right)^{|E(\alpha)| - |E(\beta)| - 2|E_{\text{phantom}}(\beta)| + |E(S)|} \right\} \right)$$

where β is a linearization of α and S is a separator of β .

5.2 Factoring ribbons and shapes

Because on the fact that norm bounds depend on vertex separators, we will need to do some combinatorics on vertex separators of shapes. The essential ideas presented in this section have appeared in prior works such as [BHK⁺16, AMP20, PR20], but we redefine them for convenience and to set up the notation for the rest of the paper.

Definition 5.3 (Composing ribbons). Two ribbons R, S are composable if $B_R = A_S$. The composition $R \circ S$ is the (possibly improper) ribbon $T = (V(R) \cup V(S), E(R) \sqcup E(S), A_R, B_S)$.

Fact 5.4. If R, S are composable ribbons, then $M_{R \circ S} = M_R M_S$.

Definition 5.5 (Composing shapes). Two shapes α, β are composable if $|V_\alpha| = |U_\beta|$. Given a bijection $\varphi : V_\alpha \rightarrow U_\beta$, the composition $\alpha \circ_\varphi \beta$ is the (possibly improper) shape ζ whose multigraph is the result of gluing together the graphs for α, β along V_α and U_β using φ . Set $U_\zeta = U_\alpha$ and $V_\zeta = V_\beta$.

If we write $\alpha \circ \beta$ then we will implicitly assume that α and β are composable and the bijection φ is given. We would like to say that the graph matrix $M_{\alpha \circ \beta}$ also factors as $M_\alpha M_\beta$, but this is not quite true. There are *intersection terms*.

Definition 5.6 (Intersection pattern). For composable shapes $\alpha_1, \alpha_2, \dots, \alpha_k$, let $\alpha = \alpha_1 \circ \alpha_2 \circ \dots \circ \alpha_k$. An intersection pattern P is a partition of $V(\alpha)$ such that for all i and $v, w \in V(\alpha_i)$, v and w are not in the same block of the partition.⁵ We say that a vertex “intersects” if its block has size at least 2 and let $\text{Int}(P)$ denote the set of intersecting vertices.

Let $\mathcal{P}_{\alpha_1, \alpha_2, \dots, \alpha_k}$ be the set of intersection patterns between $\alpha_1, \alpha_2, \dots, \alpha_k$.

Definition 5.7 (Intersection shape). For composable shapes $\alpha_1, \alpha_2, \dots, \alpha_k$ and an intersection pattern $P \in \mathcal{P}_{\alpha_1, \alpha_2, \dots, \alpha_k}$, let $\alpha_P = \alpha_1 \circ \alpha_2 \circ \dots \circ \alpha_k$ then identify all vertices in blocks of P , i.e. contract them into a single super vertex. Keep all edges (and hence α_P may be improper).

Composable ribbons R_1, \dots, R_k with shapes $\alpha_1, \dots, \alpha_k$ induce an intersection pattern $P \in \mathcal{P}_{\alpha_1, \dots, \alpha_k}$ based on which vertices are equal. When multiplying graph matrices, by casing on which vertices are equal we have:

Proposition 5.8. For composable shapes $\alpha_1, \alpha_2, \dots, \alpha_k$,

$$M_{\alpha_1} \cdots M_{\alpha_k} = \sum_{P \in \mathcal{P}_{\alpha_1, \dots, \alpha_k}} M_{\alpha_P}.$$

⁵The intersection pattern also specifies the bijections φ for composing the shapes $\alpha_1, \dots, \alpha_k$.

Note that different P may give the same α_P , and hence the total coefficient on M_{α_P} for a given shape α_P is more complicated.

Definition 5.9 (Minimum vertex separators). *For a shape α , a vertex separator S is a minimum vertex separator (MVS) if it has the smallest possible size. MVS S is the leftmost minimum vertex separator (LMVS) if $\alpha \setminus S$ cuts all paths from U_α to any other MVS. The rightmost minimum vertex separator (RMVS) likewise cuts paths from V_α .*

The LMVS and RMVS can be easily shown to be uniquely defined [BHK⁺16].

Definition 5.10 (Left shape). *A shape σ is a left shape if it is proper, the unique minimum vertex separator is V_σ , there are no edges with both endpoints in V_σ , and every vertex is connected to U_σ .*

Definition 5.11 (Middle shape). *A shape τ is a middle shape if U_τ is the leftmost minimum vertex separator of τ , and V_τ is the rightmost minimum vertex separator of τ . If τ is proper, we say it is a proper middle shape.*

Definition 5.12 (Right shape). *σ' is a right shape if σ'^\top is a left shape.*

We also extend the definition of (L/R)MVS, left, middle, and right to ribbons. Every proper shape admits a canonical decomposition into left, right, and middle parts.

Proposition 5.13. *Every proper shape α has a unique decomposition $\alpha = \sigma \circ \tau \circ \sigma'^\top$, where σ is a left shape, τ is a middle shape, and σ'^\top is a right shape.*

The decomposition takes σ to be the set of vertices reachable from U_α via paths that do not pass through the LMVS, similarly for σ' vis-à-vis the RMVS, and then τ is the remainder. This decomposition respects the connected truncation:

Proposition 5.14. *If all vertices in α have a path to $U_\alpha \cup V_\alpha$, then decomposing $\alpha = \sigma \circ \tau \circ \sigma'^\top$, all vertices in τ have a path to both U_τ and V_τ .*

Proof. It suffices to show that there is a path to $U_\tau \cup V_\tau$. In this case, say there is a path to vertex $u \in U_\tau$, then u must have a path to V_τ . Otherwise, $U_\tau \setminus \{u\}$ would be a smaller vertex separator of τ than U_τ , a contradiction to τ being a middle shape.

Let $v \in V(\tau) \subseteq V(\alpha)$. By assumption there is a path from v to $u \in U_\alpha \cup V_\alpha$; without loss of generality, $u \in U_\alpha$. Since v is not in σ , which was constructed by taking all vertices reachable from U_α without passing through U_τ , the path must pass through U_τ . ■

6 PSD-ness

To prove PSD-ness, we perform an approximate PSD decomposition of the moment matrix Λ . This type of decomposition originates from the planted clique problem [BHK⁺16]. Our usage of it will be similar at a high level, but the accounting of graph matrices is more complicated.

6.1 Overview of the approximate PSD decomposition

Recall that our goal is to show that the matrix $\sum_{\alpha \in \mathcal{S}} \lambda_\alpha M_\alpha$ is PSD. Each shape α decomposes into left, middle, and right parts $\alpha = \sigma \circ \tau \circ \sigma'^\top$. We have a corresponding approximate decomposition of the graph matrix for α , up to intersection terms,

$$M_\alpha = M_{\sigma \circ \tau \circ \sigma'^\top} \approx M_\sigma M_\tau M_{\sigma'}^\top.$$

The dominant term in the PSD decomposition of Λ (call it \mathcal{H}) collects together α such that its middle shape τ is trivial. Since the coefficients λ_α also factor, this term is automatically PSD.

$$\sum_{\alpha: \text{trivial}} \lambda_\alpha M_\alpha = \sum_{\sigma, \sigma'} \lambda_{\sigma \circ \sigma'^\top} M_{\sigma \circ \sigma'^\top} \approx \sum_{\sigma, \sigma'} \lambda_\sigma \lambda_{\sigma'} M_\sigma M_{\sigma'}^\top = \left(\sum_{\sigma} \lambda_\sigma M_\sigma \right) \left(\sum_{\sigma} \lambda_\sigma M_\sigma \right)^\top =: \mathcal{H}.$$

There are two types of things to handle: α with nontrivial τ , and intersection terms. For α with a fixed middle shape τ , these should all be charged directly to \mathcal{H} (include τ^\top so the matrix is symmetric):

$$\left(\sum_{\sigma} \lambda_\sigma M_\sigma \right) \lambda_\tau (M_\tau + M_\tau^\top) \left(\sum_{\sigma} \lambda_\sigma M_\sigma \right)^\top \preceq \frac{1}{c(\tau)} \left(\sum_{\sigma} \lambda_\sigma M_\sigma \right) \left(\sum_{\sigma} \lambda_\sigma M_\sigma \right)^\top$$

where we leave some space $c(\tau)$. Intuitively this is possible because nontrivial τ have smaller coefficients λ_τ due to vertex/edge decay, and the norm of M_τ is controlled due to the factorization of α into left, middle, and right parts. This check amounts to $\sum_{\tau \text{ nontrivial}} \lambda_\tau M_\tau \preceq o(1)\text{Id}$.

The intersection terms need to be handled in a recursive way so that their norms can be kept under control: if $\sigma, \tau, \sigma'^\top$ intersect to create a shape ζ , then we need to factor out the non-intersecting parts of σ and σ' from the intersecting parts γ, γ' . Informally writing

$$\zeta = (\sigma - \gamma) \circ \tau_p \circ (\sigma' - \gamma')^\top,$$

where τ_p is the intersection of $\gamma, \tau, \gamma'^\top$, we now perform a further factorization

$$M_\zeta \approx M_{\sigma - \gamma} M_{\tau_p} M_{\sigma' - \gamma'}^\top$$

which recursively creates more intersection terms. The point is that the sum over $\sigma - \gamma, \sigma' - \gamma'$ is equivalent to the sum over σ, σ' (up to truncation error).

$$\begin{aligned} \sum_{\sigma - \gamma} \sum_{\sigma' - \gamma'} \lambda_{\sigma \circ \tau \circ \sigma'^\top} M_{\sigma - \gamma} M_{\tau_p} M_{\sigma' - \gamma'}^\top &= \left(\sum_{\sigma - \gamma} \lambda_{\sigma - \gamma} M_{\sigma - \gamma} \right) \lambda_{\gamma \circ \tau \circ \gamma'^\top} M_{\tau_p} \left(\sum_{\sigma - \gamma} \lambda_{\sigma - \gamma} M_{\sigma - \gamma}^\top \right) \\ &= \left(\sum_{\sigma} \lambda_\sigma M_\sigma \right) \lambda_{\gamma \circ \tau \circ \gamma'^\top} M_{\tau_p} \left(\sum_{\sigma} \lambda_\sigma M_\sigma^\top \right) + \text{truncation error.} \end{aligned}$$

In summary, we have the following informal decomposition of the moment matrix,

$$\Lambda = \left(\sum_{\sigma} \lambda_\sigma M_\sigma \right) \left(\text{Id} + \sum_{\tau \text{ nontrivial}} \lambda_\tau M_\tau + \sum_{\substack{\text{intersection terms} \\ \tau_p \in \mathcal{P}_{\gamma, \tau, \gamma'}}} \lambda_{\gamma \circ \tau \circ \gamma'^\top} M_{\tau_p} \right) \left(\sum_{\sigma} \lambda_\sigma M_\sigma \right)^\top + \text{truncation error.}$$

We then need to compare the intersection terms M_{τ_p} with Id . The vertex decay factors are compared with τ via the ‘‘intersection tradeoff lemma’’. For the edge decay factors, we give new charging arguments. The number of ways to produce τ_p as an intersection pattern needs to be bounded combinatorially.

Finally, there is the issue of truncation. The total norm of the truncation error should be small. This can be accomplished by taking a large truncation parameter.

The matrix

$$\text{Id} + \sum_{\tau \text{ nontrivial}} \lambda_{\tau} M_{\tau} + \sum_{\substack{\text{intersection terms} \\ \tau_P \in \mathcal{P}_{\gamma, \tau, \gamma'}}} \lambda_{\gamma \circ \tau \circ \gamma'} M_{\tau_P}$$

actually does have a nullspace, which prevents the above strategy from working perfectly. This is because Λ needs to satisfy the independent set constraints ($\tilde{\mathbb{E}}[x^S] = 0$ if S has an edge). For independent set it's easy to factor out these constraints. Instead of Id , the leading term is a diagonal projection matrix $\pi := \sum_{\tau: V(\tau)=U_{\tau}=V_{\tau}} \lambda_{\tau} M_{\tau}$ and the non-dominant middle shapes are τ with $|V(\tau)| > \frac{|U_{\tau}|+|V_{\tau}|}{2}$.

6.2 Informal sketch for bounding τ and τ_P

The most important part of the overview given in the previous section is showing that middle shapes τ and intersection terms τ_P can be charged to the identity matrix, i.e. a type of “graph matrix tail bound” for middle shapes and intersection terms. In this subsection, we describe the properties that τ and τ_P satisfy, their coefficients, and their norm bounds. Using this, we show that for each individual non-trivial τ , $\lambda_{\tau} \|M_{\tau}\| \ll 1$ and for each intersection pattern P , $\lambda_{\gamma \circ \tau \circ \gamma'} \|M_{\tau_P}\| \ll 1$. We then explain how these arguments are used to prove [Theorem 1.8](#). The combinatorial arguments in this section crucially rely on the connected truncation.

6.2.1 Middle shapes

Proposition 6.1. *For each middle shape τ such that $|V(\tau)| > \frac{|U_{\tau}|+|V_{\tau}|}{2}$ and every vertex in τ is connected to $U_{\tau} \cup V_{\tau}$, $\lambda_{\tau} \|M_{\tau}\| \ll 1$.*

Proof. The coefficient λ_{τ} is

$$\lambda_{\tau} = \left(\frac{k}{n}\right)^{|V(\tau)| - \frac{|U_{\tau}|+|V_{\tau}|}{2}} \left(-\sqrt{\frac{p}{1-p}}\right)^{|E(\tau)|}$$

The norm bound on M_{τ} is

$$\|M_{\tau}\| \leq \tilde{O} \left(n^{\frac{|V(\tau) \setminus S|}{2}} \left(\sqrt{\frac{1-p}{p}}\right)^{|E(S)|} \right)$$

where S is the sparse minimum vertex separator of τ . We now have that $|\lambda_{\tau}| \|M_{\tau}\|$ is

$$\tilde{O} \left(\left(\frac{k}{\sqrt{n}} \sqrt{\frac{p}{1-p}}\right)^{|V(\tau)| - \frac{|U_{\tau}|+|V_{\tau}|}{2}} \left(\sqrt{\frac{1-p}{np}}\right)^{|S| - \frac{|U_{\tau}|+|V_{\tau}|}{2}} \left(\sqrt{\frac{p}{1-p}}\right)^{|E(\tau) \setminus E(S)| - |V(\tau) \setminus S|} \right).$$

We claim that each of the three terms is upper bounded by 1. This will prove the claim, since the first term provides a decay for $|V(\tau)| > \frac{|U_{\tau}|+|V_{\tau}|}{2}$. The base of the first term is less than 1 for $k \lesssim n/\sqrt{d}$. The exponent of the second term is nonnegative: since τ is a middle shape, both U_{τ} and V_{τ} are minimum vertex separators of τ , and since S is a vertex separator, it must have larger size. The exponent of the third term is nonnegative by the following lemma.

Lemma 6.2. For a middle shape τ such that every vertex is connected to $U_\tau \cup V_\tau$, and any vertex separator S of τ ,

$$|E(\tau) \setminus E(S)| \geq |V(\tau) \setminus S|.$$

Proof. First, we claim that every vertex in τ is connected to S . Let $v \in V(\tau)$, and by the connected truncation, v is connected to some $u \in U_\tau$. Since τ is a middle shape, there must be a path from u to V_τ (else $U_\tau \setminus \{u\}$ would be a smaller vertex separator than U_τ). The path necessarily passes through S since S is a vertex separator, therefore we now have a path from v to S .

Now, we can assign an edge of $E(\tau) \setminus E(S)$ to each vertex of $V(\tau) \setminus S$ to prove the claim. To do this, run a breadth-first search from S , and assign an edge to the vertex that it explores. ■

6.2.2 Intersection terms

Recall that intersection terms are formed by the intersection of $\sigma, \tau, \sigma'^\top$. Then we factor out the non-intersecting parts, leaving the *middle intersection* τ_P , which is an intersection of some portion γ of σ , the middle shape τ , and some portion γ' of σ' , which we now make formal.

Definition 6.3 (Middle intersection). Let γ, γ' be left shapes and τ be a shape such that $\gamma \circ \tau \circ \gamma'^\top$ are composable. We say that an intersection pattern $P \in \mathcal{P}_{\gamma, \tau, \gamma'^\top}$ is a middle intersection if U_γ is a minimum vertex separator in γ of U_γ and $V_\gamma \cup \text{Int}(P)$. Similarly, $U_{\gamma'}$ is a minimum vertex separator in γ' of $U_{\gamma'}$ and $V_{\gamma'} \cup \text{Int}(P)$. Finally, we also require that P has at least one intersection.

Let $\mathcal{P}_{\gamma, \tau, \gamma'}^{\text{mid}}$ denote the set of middle intersections.

Remark 6.4. For middle intersections we use the notation τ_P to denote the resulting shape, as compared to α_P which is used for an arbitrary intersection pattern.

Remark 6.5. In fact this definition also captures recursive intersection terms which are created from later rounds of factorization. We say that $P \in \mathcal{P}_{\gamma_k, \dots, \gamma_1, \tau, \gamma_1^\top, \dots, \gamma_k^\top}$ is a middle intersection, denoted $P \in \mathcal{P}_{\gamma_k, \dots, \gamma_1, \tau, \gamma_1', \dots, \gamma_k'}^{\text{mid}}$ if for all $j = 0, \dots, k-1$, letting τ_j be the shape of intersections so far between $\gamma_j, \dots, \gamma_1, \tau, \gamma_1^\top, \dots, \gamma_j^\top$, the intersection $\gamma_{j+1}, \tau_j, \gamma_{j+1}^\top$ is a middle intersection.

We need the following structural property of middle intersections.

Proposition 6.6. For a middle intersection $P \in \mathcal{P}_{\gamma, \tau, \gamma'}^{\text{mid}}$ such that every vertex of τ is connected to both U_τ and V_τ , every vertex in τ_P is connected to both U_{τ_P} and V_{τ_P} .

Proof. By assumption, all vertices in τ have a path to both U_τ and V_τ . Since $U_\tau = V_{\gamma'}$, and all vertices in γ have a path to U_γ by definition of a left shape, all vertices in τ have a path to $U_\gamma = U_{\tau_P}$. Similarly, $V_\tau = U_{\gamma'}$ is connected to $V_{\gamma'} = V_{\tau_P}$. Thus we have shown that all vertices in τ are connected to both U_{τ_P} and V_{τ_P} .

Vertices in γ have a path to U_{τ_P} by definition; we must show that they also have a path to V_{τ_P} . A similar argument will hold for vertices in γ' . Since all vertices in γ have paths to U_γ , it suffices to show that all $u \in U_\gamma$ have a path to V_{τ_P} . There must be a path from u to some $v \in V_\gamma \cup \text{Int}(P)$, else this would violate the definition of a middle intersection by taking $U_\gamma \setminus \{u\}$. The vertex v is in either τ or γ' , both of which are connected to V_{τ_P} , hence we are done. ■

Now we show that middle intersections have small norm. We focus on the first level of intersection terms; the general case of middle intersections in [Remark 6.5](#) follows by induction.

Proposition 6.7. *For left shapes γ, γ' , proper middle shape τ such that every vertex in τ is connected to $U_\tau \cup V_\tau$, and a middle intersection $P \in \mathcal{P}_{\gamma, \tau, \gamma'}^{\text{mid}}$, $|\lambda_{\gamma \circ \tau \circ \gamma'}| \|M_{\tau_P}\| \ll 1$.*

Proof. For an intersection pattern P , the coefficient is

$$\lambda_{\gamma \circ \tau \circ \gamma'} = \left(\frac{k}{n}\right)^{|V(\tau_P)| + i_P - \frac{|U_{\tau_P}| + |V_{\tau_P}|}{2}} \left(-\sqrt{\frac{p}{1-p}}\right)^{|E(\tau_P)|}$$

where $i_P := |V(\gamma \circ \tau \circ \gamma')| - |V(\tau_P)|$ is the number of intersections in P .

The norm bound on M_{τ_P} is

$$\tilde{O} \left(\max_{\beta, S} \left\{ n^{\frac{|V(\beta)| + |I_\beta| - |S|}{2}} \left(\sqrt{\frac{1-p}{p}} \right)^{|E(S)| + |E(\tau_P)| - |E(\beta)| - 2|E_{\text{phantom}}|} \right\} \right)$$

where the maximization is taken over all linearizations β of τ_P and all separators S for β .

We now have that $|\lambda_{\gamma \circ \tau \circ \gamma'}| \|M_{\tau_P}\|$ is

$$\tilde{O} \left(\max_{\beta, S} \left\{ \left(\frac{k}{\sqrt{n}} \sqrt{\frac{p}{1-p}} \right)^{|V(\gamma \circ \tau \circ \gamma')| - \frac{|U_\beta| + |V_\beta|}{2}} \left(\sqrt{\frac{1-p}{np}} \right)^{i_P - |I_\beta| + |S| - \frac{|U_\beta| + |V_\beta|}{2}} \right. \right. \\ \left. \left. \left(\sqrt{\frac{p}{1-p}} \right)^{|E(\beta) \setminus E(S)| - |V(\beta) \setminus S| - |I_\beta| + 2|E_{\text{phantom}}|} \right\} \right).$$

We claim that each of the three terms is upper bounded by 1, which proves the proposition since the first term provides a decay (since an intersection is nontrivial). The base of the first term (vertex decay) is less than 1 for $k \lesssim n/\sqrt{d}$. The second term has a nonnegative exponent by the *intersection tradeoff lemma* from [\[BHK⁺16\]](#). The version we cite is Lemma 9.32 in [\[PR20\]](#), with the simplification that τ is a proper middle shape, so $I_\tau = \emptyset$ and $|S_{\tau, \min}| = |U_\tau| = |V_\tau|$ (the full form is used for intersection terms with $\gamma_1, \dots, \gamma_k, k > 1$).

Lemma 6.8 (Intersection tradeoff lemma). *For all left shapes γ, γ' and proper middle shapes τ , let $P \in \mathcal{P}_{\gamma, \tau, \gamma'}$ be a middle intersection, then*

$$|V(\tau)| - \frac{|U_\tau| + |V_\tau|}{2} + |V(\gamma)| - |U_\gamma| + |V(\gamma')| - |U_{\gamma'}| \geq |V(\tau_P)| + |I_{\tau_P, \min}| - |S_{\tau_P, \min}|$$

where $S_{\alpha, \min}$ is defined to be a minimum vertex separator of α with all multi-edges deleted, and $I_{\alpha, \min}$ is the set of isolated vertices in α with all multi-edges deleted.

The above inequality is equivalent to

$$i_P + \frac{|U_\tau| + |V_\tau|}{2} - |U_\gamma| - |U_{\gamma'}| \geq |I_{\tau_P, \min}| - |S_{\tau_P, \min}|.$$

Since $|U_\tau| \leq |U_\gamma| = |U_\beta|, |V_\tau| \leq |U_{\gamma'}| = |V_\beta|, |I_{\tau_P, \min}| \geq |I_\beta|, |S| \geq |S_{\tau_P, \min}|$, this implies

$$i_P - |I_\beta| + |S| - \frac{|U_\beta| + |V_\beta|}{2} \geq 0.$$

The third term has a nonnegative exponent, as the following lemma shows.

Lemma 6.9. For any linearization β of τ_P and all separators S for β ,

$$|E(\beta)| + 2|E_{phantom}| - |E(S)| \geq |V(\tau_P)| + |I_\beta| - |S|.$$

Proof. We give a proof using induction. An alternate proof that uses an explicit charging scheme is given in [Appendix C](#).

We start with $E(\beta)$ and add in pairs of phantom edges one by one. The final graph will have $|E(\beta)| + 2|E_{phantom}|$ edges. The inductive claim is that at any time in the process, for all connected components C in the current graph,

1. If C is not connected to $U_{\tau_P} = U_\beta$ or $V_{\tau_P} = V_\beta$, $|E(C) \setminus E(S)| + 2|E_{phantom}(C)| \geq |C \setminus S| + |I_\beta \cap C| - 2$. For example, C may consist of a single isolated vertex or a set of isolated vertices connected by phantom edges.
2. If C is connected to U_{τ_P} or V_{τ_P} but not both, $|E(C) \setminus E(S)| + 2|E_{phantom}(C)| \geq |C \setminus S| + |I_\beta \cap C| - 1$.
3. If C is connected to both U_{τ_P} and V_{τ_P} , $|E(C) \setminus E(S)| + 2|E_{phantom}(C)| \geq |C \setminus S| + |I_\beta \cap C|$

To see that the base case holds for $E(\beta)$ and $E_{phantom}(C) = \emptyset$, let C be a component C of $E(\beta)$. If C is a single isolated vertex, we are in case (1), and the inequality is a tight equality. Otherwise, C has no isolated vertices, and therefore $|E(C) \setminus E(S)| \geq |C \setminus S| - 1$ holds by connectivity, which is sufficient for C in either case (1) or (2). If C is in case (3), C is connected to both U_{τ_P} and V_{τ_P} . Since S is a separator, in this case C must include at least one vertex in S and the stronger inequality $|E(C) \setminus E(S)| \geq |C \setminus S|$ holds as required for the base case.

Now consider adding another phantom edge e . If e is entirely contained within a component C , the right-hand side of the inequality is unchanged, therefore the inequality for C continues to hold. If e goes between two components C_1 and C_2 , there are six cases depending on which case C_1 and C_2 are in. Let C denote the new joined component. We show the argument when C_1, C_2 are both case (1) or both case (2). The remaining cases are similar.

(Both are case (1)) C is connected to neither U_{τ_P} nor V_{τ_P} , so we must show the inequality for case (1) still holds. Adding the inequalities for C_1 and C_2 ,

$$\begin{aligned} & |E(C_1) \setminus E(S)| + 2|E_{phantom}(C_1)| + |E(C_2) \setminus E(S)| + 2|E_{phantom}(C_2)| \\ & \geq |C_1 \setminus S| + |I_\beta \cap C_1| + |C_2 \setminus S| + |I_\beta \cap C_2| - 4. \end{aligned}$$

Equivalently, using $2|E_{phantom}(C)| = 2|E_{phantom}(C_1)| + 2|E_{phantom}(C_2)| + 2$,

$$\begin{aligned} & |E(C) \setminus E(S)| + 2|E_{phantom}(C)| - 2 \geq |C \setminus S| + |I_\beta \cap C| - 4 \\ \Leftrightarrow & |E(C) \setminus E(S)| + 2|E_{phantom}(C)| \geq |C \setminus S| + |I_\beta \cap C| - 2. \end{aligned}$$

(Both are case (2)) C is connected to either one or both of U_{τ_P} and V_{τ_P} ; we show the tighter inequality required by both, case (3). Adding the inequalities for C_1 and C_2 ,

$$\begin{aligned} & |E(C_1) \setminus E(S)| + 2|E_{phantom}(C_1)| + |E(C_2) \setminus E(S)| + 2|E_{phantom}(C_2)| \\ & \geq |C_1 \setminus S| + |I_\beta \cap C_1| + |C_2 \setminus S| + |I_\beta \cap C_2| - 2. \end{aligned}$$

Again using $2|E_{\text{phantom}}(C)| = 2|E_{\text{phantom}}(C_1)| + 2|E_{\text{phantom}}(C_2)| + 2$,

$$\begin{aligned} & |E(C) \setminus E(S)| + 2|E_{\text{phantom}}(C)| - 2 \geq |C \setminus S| + |I_\beta \cap C| - 2 \\ \Leftrightarrow & |E(C) \setminus E(S)| + 2|E_{\text{phantom}}(C)| \geq |C \setminus S| + |I_\beta \cap C|. \end{aligned}$$

At the end of the process, the connectivity of the graph is the same as the graph τ_P (though nonzero edge multiplicities may be modified). By [Proposition 5.14](#), every vertex in $V(\tau_P)$ is connected to both U_{τ_P} and V_{τ_P} . Therefore only case (3) holds. Summing the inequalities proves the lemma. \blacksquare

6.2.3 Proof of [Theorem 1.8](#)

The preceding lemmas are informal, but they actually show something more that justifies [Theorem 1.8](#). For technical reasons, our formal proof that implements [Section 6.1](#) will break for $d \geq n^{0.5}$, and thus it doesn't cover [Theorem 1.8](#), though we show that the arguments above plus the argument of [\[BHK⁺16\]](#) are already sufficient to prove [Theorem 1.8](#).

In the proof of [Proposition 6.1](#) the term $\lambda_\tau \|M_\tau\|$ was broken into three parts, each less than 1. Using just the first two parts,

$$|\lambda_\tau| \|M_\tau\| \leq \tilde{O} \left(\left(\frac{k}{\sqrt{n}} \sqrt{\frac{p}{1-p}} \right)^{|V(\tau) \setminus S|} \left(\frac{k}{n} \right)^{|S| - \frac{|U_\tau| + |V_\tau|}{2}} \right).$$

Letting $k = \frac{n}{d^{1/2+\varepsilon}}$,

$$\begin{aligned} |\lambda_\tau| \|M_\tau\| &\leq \tilde{O} \left(\left(\frac{1}{d^\varepsilon} \right)^{|V(\tau) \setminus S|} \left(\frac{1}{d^{1/2+\varepsilon}} \right)^{|S| - \frac{|U_\tau| + |V_\tau|}{2}} \right) \\ &\leq \tilde{O} \left(\left(\frac{1}{d^\varepsilon} \right)^{|V(\tau) \setminus S_\tau|} \left(\frac{1}{d^{1/2+\varepsilon}} \right)^{|S_\tau| - \frac{|U_\tau| + |V_\tau|}{2}} \right) \end{aligned}$$

where S_τ is the dense MVS of τ . Observe that *this equals $\lambda_\tau \|M_\tau\|$ in the dense case ($p = 1/2$) for a random graph of size d* . That is, suppose we performed the pseudocalibration and norm bounds from [\[BHK⁺16\]](#) for a random graph $G \sim G_{d,1/2}$. Then we would get $\lambda_\tau = \left(\frac{1}{d^{1/2+\varepsilon}} \right)^{|V(\tau) \setminus S_\tau|}$

and $\|M_\tau\| \leq (\log d)^{O(1)} \sqrt{d}^{|V(\tau) \setminus S_\tau|}$, and $\lambda_\tau \|M_\tau\|$ is the same as the above. The main point is: by the analysis from [\[BHK⁺16\]](#), we know that for shapes up to size $\Omega(\log d)$, the sum of these norms is $o(1)$. Therefore⁶ our matrices sum to $o(1)$ for SoS degree $\Omega(\log d)$.

The same phenomenon occurs for the intersection terms. Essentially we are neglecting the effect of the edges and considering only the vertex factors. By taking the first two out of three

⁶The log factor is $(\log n)^{O(1)}$ for our norm bound, vs $(\log d)^{O(1)}$ for $G \sim G_{d,1/2}$, but these are the same up to a constant for $d \geq n^\varepsilon$.

terms used in the proof of [Proposition 6.7](#) we have the bound ($k = \frac{n}{d^{1/2+\varepsilon}}$)

$$\begin{aligned} |\lambda_{\gamma \circ \tau \circ \gamma' \tau}| \|M_{\tau_P}\| &\leq \tilde{O} \left(\left(\frac{1}{d^\varepsilon} \right)^{|V(\tau_P)|+|I_\beta|-|S|} \left(\frac{1}{d^{1/2+\varepsilon}} \right)^{|V(\gamma \circ \tau \circ \gamma' \tau)|-|V(\tau_P)|-|I_\beta|+|S|-\frac{|U_\beta|+|V_\beta|}{2}} \right) \\ &\leq \tilde{O} \left(\left(\frac{1}{d^\varepsilon} \right)^{|V(\tau_P)|+|I_{\tau_P, \min}|-|S_{\tau_P, \min}|} \left(\frac{1}{d^{1/2+\varepsilon}} \right)^{|V(\gamma \circ \tau \circ \gamma' \tau)|-|V(\tau_P)|-|I_{\tau_P, \min}|+|S_{\tau_P, \min}|-\frac{|U_{\tau_P}|+|V_{\tau_P}|}{2}} \right) \end{aligned}$$

where $S_{\tau_P, \min}$ and $I_{\tau_P, \min}$ are the dense MVS and isolated vertices respectively in the graph τ_P after deleting all multiedges. For $G \sim G_{d, 1/2}$, the pseudocalibration coefficient and norm bounds are

$$\begin{aligned} \lambda_{\gamma \circ \tau \circ \gamma' \tau} &= \left(\frac{1}{d^{1/2+\varepsilon}} \right)^{|V(\gamma \circ \tau \circ \gamma' \tau)|-\frac{|U_\gamma|+|U_{\gamma'}|}{2}} \\ \|M_{\tau_P}\| &\leq (\log d)^{O(1)} \sqrt{d}^{|V(\tau_P)|-|S_{\tau_P, \min}|+|I_{\tau_P, \min}|}. \end{aligned}$$

Multiplying these together, one can see that $\lambda_{\gamma \circ \tau \circ \gamma' \tau} \|M_{\tau_P}\|$ is the same. Therefore, the analysis of [\[BHK⁺16\]](#) also shows that the sum of all intersection terms is negligible.

By passing to the ‘‘dense proxy graph’’ $G \sim G_{d, 1/2}$, we can essentially use [\[BHK⁺16\]](#) to deduce a degree- $\Omega(\log d)$ lower bound in the sparse case just using connected truncation with the sparse norm bounds. The analysis of [\[BHK⁺16\]](#) is limited to SoS degree $\Omega(\log d)$ and $d \geq n^\varepsilon$. In order to push the SoS degree up and the graph degree down in the remaining sections, we need to utilize conditioning and handle the combinatorial terms more carefully.

6.3 Decomposition in terms of ribbons

We now work towards the proof of [Theorem 1.4](#). Recall that our goal is to show PSD-ness of the matrix $\sum_{\alpha \in \mathcal{S}} \lambda_\alpha \frac{M_\alpha}{|\text{Aut}(\alpha)|}$.

Definition 6.10 (Properly composable). *Composable ribbons R_1, \dots, R_k are properly composable if there are no intersections beyond the necessary ones $B_{R_i} = A_{R_{i+1}}$.*

Definition 6.11 (Proper composition). *Given composable ribbons R_1, \dots, R_k which are not necessarily properly composable, with shapes $\alpha_1, \dots, \alpha_k$, let $R_1 \odot R_2 \odot \dots \odot R_k$ be the shape of $\alpha_1 \circ \dots \circ \alpha_k$.*

That is, $R_1 \odot \dots \odot R_k$ is the shape obtained by concatenating the ribbons but not collapsing vertices which repeat between ribbons, as if they were properly composable. Compare this with $R_1 \circ \dots \circ R_k$, in which repetitions collapse.

Definition 6.12 (\mathcal{L} and \mathcal{M}). *Let \mathcal{L} and \mathcal{M} be the set of left and middle shapes in \mathcal{S} . Shapes in these sets will be denoted $\sigma, \gamma \in \mathcal{L}$ and $\tau \in \mathcal{M}$. We abuse notation and let $L, G \in \mathcal{L}, T \in \mathcal{M}$ denote ribbons of shapes in \mathcal{L}, \mathcal{M} (following the convention of using Greek letters for shapes and Latin letters for ribbons).*

Formalizing the process described in [Section 6.1](#), we have the following lemma. It is easier to formally manipulate unsymmetrized objects, ribbons, for the PSD decomposition and switch to symmetrized objects, shapes, only when we need to invoke norm bounds. The details are carried out in [Appendix D](#).

Lemma 6.13. (Decomposition in terms of ribbons).

$$\sum_{\alpha \in \mathcal{S}} \lambda_\alpha \frac{M_\alpha}{|\text{Aut}(\alpha)|} = \sum_{R \in \mathcal{S}} \lambda_R M_R =$$

$$\left(\sum_{L \in \mathcal{L}} \lambda_L M_L \right) \left(\sum_{j=0}^{2D_{\text{SoS}}} (-1)^j \sum_{G_j, \dots, G_1, T, G'_1, \dots, G'_j} \lambda_{G_j \circ \dots \circ G_1 \circ T \circ G'_1 \circ \dots \circ G'_j} M_{G_j \circ \dots \circ G_1 \circ T \circ G'_1 \circ \dots \circ G'_j} \right) \left(\sum_{L' \in \mathcal{L}} \lambda_{L'} M_{L'} \right)^\top$$

+ truncation error_{too many vertices} + truncation error_{too many edges in one part}

where

1.

$$\text{truncation error}_{\text{too many vertices}} = - \sum_{\substack{L, T, L': \\ |V(L \circ T \circ L')| > D_V, \\ L, T, L' \text{ are properly composable}}} \lambda_{L \circ T \circ L'} M_{L \circ T \circ L'}$$

$$+ \sum_{j=1}^{2D_{\text{SoS}}} (-1)^{j+1} \sum_{\substack{L, G_j, \dots, G_1, T, G'_1, \dots, G'_j, L': \\ |V(L \circ G_j \circ \dots \circ G_1)| > D_V \text{ or} \\ |V(G'_1 \circ \dots \circ G'_j \circ L')| > D_V}} \lambda_{L \circ G_j \circ \dots \circ G_1 \circ T \circ G'_1 \circ \dots \circ G'_j} M_L M_{G_j \circ \dots \circ G_1 \circ T \circ G'_1 \circ \dots \circ G'_j} M_{L'}^\top$$

$$\text{truncation error}_{\text{too many edges in one part}} = \sum_{\substack{L, T, L': \\ |V(L \circ T \circ L')| \leq D_V, \\ |E_{\text{mid}}(T)| - |V(T)| > CD_{\text{SoS}}, \\ L, T, L' \text{ are properly composable}}} \lambda_{L \circ T \circ L'} M_{L \circ T \circ L'}$$

$$+ \sum_{j=1}^{2D_{\text{SoS}}} (-1)^j \sum_{\substack{L, G_j, \dots, G_1, T, G'_1, \dots, G'_j, L': \\ |V(L \circ G_j \circ \dots \circ G_1)| \leq D_V \text{ and } |V(G'_1 \circ \dots \circ G'_j \circ L')| \leq D_V \\ |E_{\text{mid}}(G_j)| - |V(G_j)| > CD_{\text{SoS}} \text{ or } |E_{\text{mid}}(G'_j)| - |V(G'_j)| > CD_{\text{SoS}} \\ L, (G_j \circ \dots \circ G_1 \circ T \circ G'_1 \circ \dots \circ G'_j), L' \text{ are properly composable}}} \lambda_{L \circ G_j \circ \dots \circ G_1 \circ T \circ G'_1 \circ \dots \circ G'_j} M_{L \circ G_j \circ \dots \circ G_1 \circ T \circ G'_1 \circ \dots \circ G'_j}$$

where for a ribbon R , $E_{\text{mid}}(R)$ is the set of edges of R which are not contained in A_R , not contained in B_R , and are not incident to any vertices in $A_R \cap B_R$.

2. in all of these sums, the ribbons $L, G_j, \dots, G_1, T, G'_1, \dots, G'_j, L'$ satisfy the following conditions:

- (a) $T \in \mathcal{M}$, $L, L' \in \mathcal{L}$, and each $G_i, G'_i \in \mathcal{L}$.
- (b) $L, G_j, \dots, G_1, T, G'_1, \dots, G'_j, L'$ are composable.
- (c) The intersection pattern induced by $G_j, \dots, G_1, T, G'_1, \dots, G'_j$ is a middle intersection pattern.
- (d) $|V(T)| \leq D_V$, $|V(G_j \circ \dots \circ G_1)| \leq D_V$, and $|V(G'_1 \circ \dots \circ G'_j)| \leq D_V$
- (e) Except when noted otherwise (which only happens for truncation error_{too many edges in one part}), all of the ribbons $G_j, \dots, G_1, T, G'_1, \dots, G'_j$ (but not necessarily L, L') satisfy the constraint that $|E_{\text{mid}}(R)| - |V(R)| \leq CD_{\text{SoS}}$.

6.4 Factor out Π and edges incident to $U_\alpha \cap V_\alpha$

We need to augment our graph matrices to include missing edge indicators. There are two reasons. First, the dominant term in the decomposition is a projection matrix with these indicators, rather than the identity matrix. Second, we need to factor out dependence on $U_\tau \cap V_\tau$. These vertices do not essentially participate in the graph matrix, since the matrix is block diagonal with a block for each assignment to $U_\alpha \cap V_\alpha$. (When proving things with graph matrices, it is smart to start by assuming $U_\alpha \cap V_\alpha = \emptyset$, then handle the case $U_\alpha \cap V_\alpha \neq \emptyset$).

For the first issue, define the matrix Π .

Definition 6.14. Let $\pi \in \mathbb{R}^{(\leq D_{\text{SOS}}^n) \times (\leq D_{\text{SOS}}^n)}$ be the projector to the independent set constraints. π is a diagonal matrix with entries

$$\pi[S, S] = \mathbb{1} [S \text{ is an independent set in } G]$$

Definition 6.15. Let $\Pi \in \mathbb{R}^{(\leq D_{\text{SOS}}^n) \times (\leq D_{\text{SOS}}^n)}$ be a rescaling of π by

$$\Pi[S, S] = \left(\frac{1}{1-p} \right)^{\binom{|S|}{2}} \mathbb{1} [S \text{ is an independent set in } G]$$

The rescaling satisfies $\mathbb{E}_{G \sim G_{n,p}}[\Pi[S, S]] = 1$. Recall that in [Proposition 2.5](#), we had the function $1 + \sqrt{\frac{p}{1-p}} \chi_{\{e\}} = \frac{1}{1-p} \mathbb{1}_{e \notin E(G)}$ which is $\frac{1}{1-p}$ times the indicator function for e being absent from the input graph G . Since the coefficients λ_R come with $\sqrt{\frac{p}{1-p}}$ for each edge, we can group the edges inside A_R, B_R into missing edge indicators for all edges inside A_R, B_R – that is, independent set indicators on A_R, B_R which form the matrix Π . For all the sums of ribbons we consider, after grouping:

$$\sum_R \lambda_R M_R = \Pi^{1/2} \left(\sum_{\substack{R \\ E(A_R) = E(B_R) = \emptyset}} \left(\frac{1}{1-p} \right)^{\binom{|A_R|}{2} + \binom{|B_R|}{2} - \binom{|A_R \cap B_R|}{2}} \lambda_R M_R \right) \Pi^{1/2}.$$

For the second issue, consider the function $1 + \sqrt{\frac{p}{1-p}} \chi_{\{e\}} = \frac{1}{1-p} \mathbb{1}_{e \notin E(G)}$ again. The magnitude of this function is clearly bounded by $\frac{1}{1-p}$. However, if we try to bound the magnitude of this function term by term, we instead get a bound of 2 because 1 always has magnitude 1 and if $e \in E(G)$ then $\sqrt{\frac{p}{1-p}} \chi_{\{e\}} = 1$, so the best bound we can give on the magnitude of $\sqrt{\frac{p}{1-p}} \chi_{\{e\}}$ is also 1.

This can become a problem if $U_\tau \cap V_\tau$ is large, in which case there are many possible subsets of edges incident to $U_\tau \cap V_\tau$ even if the rest of the shape is small. If we try to bound things term by term, we will get a factor of $2^{|U_\tau \cap V_\tau|}$ which may be too large. To handle this, our strategy will be to group terms into missing edge indicators, which gives a factor of $\left(\frac{1}{1-p} \right)^{|U_\tau \cap V_\tau|}$ per vertex instead. This almost works, however due to the connected truncation there are some edge cases of [Proposition 2.5](#) where we do not have the term where T is empty. We handle this using the following lemma.

Lemma 6.16. For any set of potential edges E ,

$$\sum_{E' \subseteq E: E' \neq \emptyset} \left(\sqrt{\frac{p}{1-p}} \right)^{|E'|} \chi_{E'} = \sum_{e \in E} \left(\sqrt{\frac{p}{1-p}} \right) \chi_{\{e\}} \sum_{E' \subseteq E \setminus \{e\}} \frac{1}{|E| \binom{|E|-1}{|E'|}} \left(\frac{1}{1-p} \right)^{|E'|} \mathbf{1}_{\forall e' \in E', e' \notin E(G)}$$

Proof. Observe that if we give an ordering to the edges of E then we can write

$$\sum_{E' \subseteq E: E' \neq \emptyset} \left(\sqrt{\frac{p}{1-p}} \right)^{|E'|} \chi_{E'} = \sum_{e \in E} \left(\sqrt{\frac{p}{1-p}} \right) \chi_{\{e\}} \left(\prod_{e' \in E: e' > e} \left(\frac{1}{1-p} \right) \mathbf{1}_{e' \notin E(G)} \right)$$

Taking the average over all orderings of the edges of E gives the result. To see this, observe that if we take a random ordering, the probability of having a term $\left(\sqrt{\frac{p}{1-p}} \right) \chi_{\{e\}} \left(\frac{1}{1-p} \right)^{|E'|} \mathbf{1}_{\forall e' \in E', e' \notin E(G)}$ is $\frac{1}{|E| \binom{|E|-1}{|E'|}}$ as there must be exactly $|E'|$ elements after e (so e must be in the correct position) and these elements must be E' . \blacksquare

Remark 6.17. We take the average over all orderings of the edges E so that our expression will be symmetric with respect to permuting the indices in $U_\tau \cap V_\tau$.

Based on this lemma, we make the following definition.

Definition 6.18 (Quasi-indicator function). Given a set of edges $E \subseteq \binom{[n]}{2}$, we define the quasi-missing edge indicator function q_E to be

$$q_E = (1-p)^{|E|+1} \sum_{E' \subseteq E} \frac{1}{(|E|+1) \binom{|E|}{|E'|}} \left(\frac{1}{1-p} \right)^{|E'|} \mathbf{1}_{\forall e \in E', e \notin E(G)}$$

Proposition 6.19. $\sum_{E' \subseteq E: E' \neq \emptyset} \left(\sqrt{\frac{p}{1-p}} \right)^{|E'|} \chi_{E'} = \left(\frac{1}{1-p} \right)^{|E|} \sum_{e \in E} \left(\sqrt{\frac{p}{1-p}} \right) \chi_{\{e\}} q_{E \setminus \{e\}}$

Remark 6.20. q_E is a linear combination of terms where for each edge $e \in E$, either there is a missing edge indicator for e or e is not mentioned at all. Moreover, we add the factor of $(1-p)^{|E|+1}$ to q_E so that the coefficients in this linear combination are non-negative and have sum at most 1. This is the only fact that we will use about q_E when we analyze the norms of the resulting graph matrices.

Definition 6.21 (Augmented ribbon). An augmented ribbon is a vertex set $V(R) \subseteq [n]$ with two subsets $A_R, B_R \subseteq V(R)$, as well as a multiset of edge-functions. In our augmentation, each edge-function is either a (single edge) Fourier character, a (single edge) missing edge indicator, or a (subset of edges) quasi-missing edge indicator.

Definition 6.22 (Matrix for an augmented ribbon). The matrix M_R has rows and columns indexed by all subsets of $[n]$, with entries:

$$M_R[I, J] = \begin{cases} \prod_{\text{edge-functions } f \text{ on } S \subseteq V(R)} f(G|_S) & I = A_R, J = B_R \\ 0 & \text{Otherwise} \end{cases}$$

Two augmented ribbons are isomorphic if they are in the same S_n -orbit (i.e. they are equal after renumbering the vertices), and we define an augmented shape for each orbit. Equivalently,

augmented shapes are equivalence classes of augmented ribbons under *type-preserving* multi-hypergraph isomorphism, where the type of a hyperedge is the associated function. As before, we can specify the shape by a representative graph with edge-functions. The graph matrix M_α for an augmented shape α is still defined as the sum of M_R over injective embeddings of α into $[n]$.

Definition 6.23 (Permissible). *We say that an augmented ribbon R is permissible if the following conditions are satisfied:*

1. All vertices $v \in V(R)$ are reachable from $A_R \cup B_R$ using the Fourier character edges.
2. There is a missing edge indicator for all edges within A_R and a missing edge indicator for all edges within B_R .
3. For any vertex $v \in V(R) \setminus (A_R \cap B_R)$ which is reachable from $(A_R \cup B_R) \setminus (A_R \cap B_R)$ without passing through $(A_R \cap B_R)$, there is a missing edge indicator for all edges between $(A_R \cap B_R)$ and v .
4. For each connected component C of $R \setminus (A_R \cap B_R)$ which is disconnected from $A_R \cup B_R$, there is precisely one pair of vertices $u \in A_R \cap B_R$ and $w \in C$ such that $(u, w) \in E(R)$ and R contains the quasi-missing edge indicator $q_{((A_R \cap B_R) \times C) \setminus \{u, w\}}$ for all other edges between $A_R \cap B_R$ and C .
5. There are no other (quasi-)missing edge indicators, and the Fourier character edges are disjoint from the (quasi-)missing edge indicators.

A shape is permissible if any (equivalently, all) of its ribbons are permissible.

To explain the conditions, the first condition is the connected truncation. The second condition follows because we factored out independent set indicators. The third condition observes that if we have a vertex $v \in W_\alpha$ that is connected to $(U_\alpha \cup V_\alpha) \setminus (U_\alpha \cap V_\alpha)$, then we can factor out a missing edge indicator for all edges that go from v to $U_\alpha \cap V_\alpha$ because such edges do not affect the connectivity properties of α . The fourth condition picks out components of W_α that do affect the connectivity and handles them via the definition of quasi-missing edge indicators as explained above.

Remark 6.24. *For the remainder of Section 6, ribbon means “augmented ribbon” and shape means “augmented shape”. $E(\alpha)$ refers to the multiset of Fourier characters edges in α . We will write “permissible $R \in \mathcal{S}$ ” to mean a permissible ribbon such that the corresponding non-augmented ribbon that includes all edges involved in any edge function is in \mathcal{S} (and similarly for other sets of ribbons/shapes).*

Lemma 6.25. *If α is permissible, then $\pi M_\alpha = M_\alpha \pi = M_\alpha$.*

Proof. α has missing edge indicators for all edges inside U_α and V_α , hence M_α is zero on any rows or columns that are not independent sets. ■

We count the extra factors of $1/(1-p)$ with the following definition.

Definition 6.26. *For a permissible shape or ribbon R , let:*

$$m(R) = \left(\frac{1}{1-p} \right)^{\binom{|A_R|}{2}/2 + \binom{|B_R|}{2}/2 - \binom{|A_R \cap B_R|}{2} + |V(R) \setminus (A_R \cup B_R)| \cdot |A_R \cap B_R|}.$$

Lemma 6.27.

$$\begin{aligned}
& \sum_{R \in \mathcal{S}} \lambda_R M_R = \\
& \Pi^{1/2} \left(\sum_{\substack{\text{permissible} \\ L \in \mathcal{L}}} m(L) \lambda_L M_L \right) \cdot \\
& \left(\sum_{j=0}^{2D_{\text{SoS}}} (-1)^j \sum_{\substack{\text{permissible} \\ G_j, \dots, G_1, T, G'_1, \dots, G'_j}} m(T) \left(\prod_{i=1}^j m(G_i) m(G'_i) \right) \lambda_{G_j \circ \dots \circ G_1 \circ T \circ G'_1 \circ \dots \circ G'_j} M_{G_j \circ \dots \circ G_1 \circ T \circ G'_1 \circ \dots \circ G'_j} \right) \cdot \\
& \left(\sum_{\substack{\text{permissible} \\ L \in \mathcal{L}}} m(L) \lambda_L M_L \right)^\top \Pi^{1/2} \\
& + \text{truncation error}
\end{aligned}$$

Proof. The factoring process is completely independent of edges inside $A_T, B_T, A_{G_i}, B_{G_i}$ or edges incident to $A_R \cap B_R$ for any of these shapes, as noted in [Lemma D.4](#). All subsets of edges inside A_R or B_R are present, but not all subsets of edges incident to $A_R \cap B_R$ are present; because of the connected truncation, we do not have the empty set of edges between $A_R \cap B_R$ and components of $R \setminus (A_R \cap B_R)$ which are disconnected from $A_R \cup B_R$. Now group these Fourier characters into missing edge indicators and quasi-missing edge indicators as explained above. ■

We can show the factor $m(R)$ is a negligible constant:

Lemma 6.28. *For any ribbon R with degree at most D_{SoS} ,*

$$m(R) \leq \left(\frac{1}{(1-p)^{2D_{\text{SoS}}}} \right)^{|V(R) \setminus (A_R \cap B_R)|}.$$

Proof. The claim follows from the following two inequalities:

$$\begin{aligned}
& \binom{|A_R|}{2} / 2 + \binom{|B_R|}{2} / 2 - \binom{|A_R \cap B_R|}{2} \leq D_{\text{SoS}} (|A_R \setminus (A_R \cap B_R)| + |B_R \setminus (A_R \cap B_R)|) \\
& |A_R \cap B_R| |V(R) \setminus (A_R \cup B_R)| \leq D_{\text{SoS}} |V(R) \setminus (A_R \cap B_R)|.
\end{aligned}$$

This will be handled later by the vertex decay. ■

6.5 Conditioning I: reduction to sparse shapes

To improve norm bounds for dense shapes, we can replace them by sparse ones. The replacement of dense ribbons is accomplished by the following lemmas.

Lemma 6.29. $1_{e \in E(G)} \chi_e = -\sqrt{p(1-p)} + (1-p) \chi_e$

Proof. Follows by explicit computation, as in [Proposition 2.5](#). ■

Lemma 6.30. *Given a set of edges E , if we know that not all of the edges of E are in $E(G)$ then*

$$\chi_E = - \sum_{E' \subseteq E: E' \neq E} \left(-\sqrt{\frac{p}{1-p}} \right)^{|E|-|E'|} \chi_{E'}$$

Proof. Since not all of the edges of E are in $E(G)$, $(1 - 1_{E \subseteq E(G)})\chi_E = \chi_E$. Now observe that

$$1_{E \subseteq E(G)}\chi_E = \prod_{e \in E} (1_{e \in E(G)}\chi_e) = \prod_{e \in E} (-\sqrt{p(1-p)} + (1-p)\chi_e)$$

Thus,

$$\chi_E = (1 - 1_{E \subseteq E(G)})\chi_E = (1 - (1-p)^{|E|})\chi_E - \sum_{E' \subseteq E: E' \neq E} (-\sqrt{p(1-p)})^{|E|-|E'|} (1-p)^{|E'|} \chi_{E'}$$

Solving for χ_E gives

$$\chi_E = - \sum_{E' \subseteq E: E' \neq E} \left(-\sqrt{\frac{p}{1-p}} \right)^{|E|-|E'|} \chi_{E'}$$

Corollary 6.31. *Given a set of edges E , if we know that at most $k < |E|$ of the edges of E are in $E(G)$, then*

$$\chi_E = (-1)^{|E|-k} \sum_{E' \subseteq E: |E'| \leq k} \binom{|E|-1-|E'|}{|E|-1-k} \left(-\sqrt{\frac{p}{1-p}} \right)^{|E|-|E'|} \chi_{E'}$$

Proof. Apply [Lemma 6.30](#) repeatedly. ■

Corollary 6.32. *Given a set of edges E , if we know that at most $k < |E|$ of the edges of E are in $E(G)$, then*

$$\chi_E = \sum_{E' \subseteq E: |E'| \leq k} c_{E'} \chi_{E'}$$

where $|c_{E'}| \leq (2|E|\sqrt{p})^{|E|-|E'|}$.

We can do the replacement if the graph of the shape does not occur in our random sample of $G \sim G_{n,p}$. We formalize a bound on the density of subgraphs of $G_{n,p}$. In expectation the number of occurrences of a small subgraph H is essentially $n^{|V(H)|} p^{|E(H)|}$ and hence H is unlikely to appear in a random graph sample if it has too many edges (specifically, morally all subgraphs are sparse, $|E(H)| \leq |V(H)| \log_{1/p}(n)$). To translate from expectation to concentration, we use a simple first moment calculation to bound the densest- k -subgraph in $G_{n,p}$.

Proposition 6.33 (Sparsity of small subgraphs of $G_{n,p}$). *For $G \sim G_{n,p}$, $p \leq \frac{1}{2}$, constant $\eta > 0$, with probability at least $1 - O(1/n^\eta)$, every subgraph S of G such that $|V(S)| \leq \sqrt[3]{\frac{n}{d}}$ satisfies:*

$$|E(S)| \leq 3|V(S)| \log_{1/p}(n) + 3\eta \log_{1/p}(n).$$

Proof. Let $e^*(v) := 3v \log_{1/p}(n) + 3\eta \log_{1/p}(n)$.

$$\begin{aligned}
\Pr[\exists \text{too dense subgraph}] &\leq \sum_{v=2}^n \sum_{e=e^*(v)}^{\binom{v}{2}} \mathbb{E}[\text{number of subgraphs with } v \text{ vertices, } e \text{ edges}] \\
&\leq \sum_{v=2}^n \sum_{e=e^*(v)}^{\binom{v}{2}} \binom{n}{v} v^{2e} p^e \\
&\leq \sum_{v=2}^n \sum_{e=e^*(v)}^{\binom{v}{2}} \frac{n^v}{v!} (p^{1/3})^e \\
&= \sum_{v=2}^n \sum_{e=e^*(v)}^{\binom{v}{2}} \frac{n^v p^{e^*(v)/3}}{v!} (p^{1/3})^{e-e^*(v)} \\
&= \sum_{v=2}^n \sum_{e=e^*(v)}^{\binom{v}{2}} \frac{1}{n^\eta \cdot v!} (p^{1/3})^{e-e^*(v)} \\
&\leq O(1/n^\eta).
\end{aligned}$$

■

Definition 6.34 (Sparse). A graph G is C' -sparse if $|E(G)| \leq C' |V(G)|$. A shape α is C' -sparse if the underlying graph (not multigraph) is. We will generally say that a graph/shape is sparse if it is C' -sparse for some C' .

Corollary 6.35. For $d \leq n^{0.5}$, w.h.p. every subgraph of $G \sim G_{n, \frac{d}{n}}$ of size at most $n^{0.16}$ is 7-sparse.

Definition 6.36 (Forbidden subgraph). We call a graph H forbidden if it does not occur as a subgraph of G .

The assumptions in [Theorem 1.4](#) ensure that every graph of size at most D_V that is not sparse is forbidden per [Corollary 6.35](#). This is the only class of forbidden graphs that we will need in this work.

Remark 6.37. For $d = n^{1-\varepsilon}$, subgraphs up to size $n^{\Omega(\varepsilon)}$ will be $O(1/\varepsilon)$ -sparse. Using this bound and our techniques, [Theorem 1.4](#) can be extended to show a $n^{\Omega(\varepsilon)}$ SoS-degree lower bound for $d = n^{1-\varepsilon}$.

Remark 6.38. For $d \ll n^\varepsilon$, subgraphs of $G_{n,p}$ will actually be significantly sparser. For example, it is well-known that $o(\log n)$ -radius neighborhoods in $G_{n,d/n}$ for constant d are trees with at most one extra cycle.

Definition 6.39 (Subshape/subribbon, supershape/superribbon). We call shape β a sub-shape of shape α if $V(\alpha) = V(\beta)$, $U_\alpha = U_\beta$, $V_\alpha = V_\beta$, and $E(\beta) \subseteq E(\alpha)$. Furthermore, the (quasi-)missing edge indicators of β and α must be equal. Supershapes, subribbons, and superribbons are defined similarly. We write $\alpha \subseteq \beta$ if α is a subshape of β .

Definition 6.40. For a sparse subribbon R of a forbidden ribbon U , let $c(R, U)$ be the coefficient on R after applying conditioning [Corollary 6.31](#) to the non-reserved edges of U .

Definition 6.41. Let λ'_R incorporate the constant factors into λ_R ,⁷

$$\lambda'_R := m(R) \left(\sum_{U \supseteq R} c(R, U) \lambda_U \right) \lambda_R,$$

$$\lambda'_{R_1 \odot \dots \odot R_k} := \prod_{i=1}^k \lambda'_{R_i}.$$

Lemma 6.42.

$$\begin{aligned} & \sum_{R \in \mathcal{S}} \lambda_R M_R = \\ & \Pi^{1/2} \left(\sum_{\substack{\text{sparse} \\ \text{permissible} \\ L \in \mathcal{L}}} \lambda'_L M_L \right) \cdot \\ & \left(\sum_{j=0}^{2D_{\text{SoS}}} (-1)^j \sum_{\substack{\text{sparse} \\ \text{permissible} \\ G_j, \dots, G_1, T, G'_1, \dots, G'_j}} \lambda'_{G_j \odot \dots \odot G_1 \odot T \odot G'_1 \top \dots \odot G'_j \top} M_{G_j \odot \dots \odot G_1 \odot T \odot G'_1 \top \dots \odot G'_j \top} \right) \\ & \left(\sum_{\substack{\text{sparse} \\ \text{permissible} \\ L \in \mathcal{L}}} \lambda'_L M_L \right)^\top \Pi^{1/2} \\ & + \text{truncation error} \end{aligned}$$

The proof follows mostly by definition of λ'_R but there is an important detail. When we condition a left L /middle T /intersection G ribbon, the result may no longer be a left/middle/intersection ribbon. For example, the MVS might change drastically. It turns out that we can avoid this by “reserving” $O(|V(R)|)$ edges in the ribbon and only allowing the removal of edges outside of the reserved set. The reserved edges guarantee that the subribbon will continue to be a left/middle/intersection ribbon, at the cost of increasing the sparsity. We prove that $O(|V(R)|)$ edges suffice in [Appendix B](#).

We can show that the conditioning negligibly affects the coefficients of the ribbons.

Lemma 6.43. Under the conditions of [Theorem 1.4](#), for any ribbon $R \in \mathcal{S}$,

$$\sum_{U \supseteq R} \lambda_U c(R, U) = o(1) \lambda_R.$$

Proof. Let U be a superribbon with k more edges than R . We have

$$\lambda_U = \left(\sqrt{\frac{p}{1-p}} \right)^k \lambda_U \leq (2\sqrt{p})^k \lambda_R. \quad (1)$$

⁷When R is a left L /middle T /intersecting G ribbon, the sum over U should be restricted to U with the same type, and which actually appear in the decomposition. The stated sum is an upper bound.

Using the bound from [Corollary 6.32](#),

$$|c(R, U)| \leq (2|E(U)|\sqrt{p})^k \leq (2|V(R)|^2\sqrt{p})^k. \quad (2)$$

The number of ribbons U that are superribbons of R with k extra edges is at most

$$|V(R)|^{2k} \quad (3)$$

because (see [Appendix B](#)) due to the connectivity of the reserved set, the vertex set of U and R is the same. Combining [Eq. \(1\)](#), [Eq. \(2\)](#), [Eq. \(3\)](#), the change in coefficient is at most

$$\sum_{k=1}^{\infty} \lambda_R (4|V(R)|^4 p)^k \leq \sum_{k=1}^{\infty} \lambda_R \left(\frac{4D_V^4 d}{n} \right)^k \leq o(1)\lambda_R.$$

The last inequality uses $d \leq n^{0.5}$ and $D_V \leq \tilde{O}(D_{\text{SoS}}) \leq \tilde{O}(n^{0.1})$. ■

6.6 Shifting to shapes

For a proper middle shape τ and left shapes $\gamma_j, \dots, \gamma_1, \gamma'_1, \dots, \gamma'_j$, recall the notation $\mathcal{P}_{\gamma_j, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_j}^{\text{mid}}$ for the set of middle intersection patterns ([Definition 6.3](#) and [Remark 6.5](#)). We also let $\mathcal{P}_{\tau}^{\text{mid}}$ contain a single term, the non-intersecting singleton partition.

We now analyze

$$\sum_{\substack{\text{sparse} \\ \text{permissible} \\ G_j, \dots, G_1, T, G'_1, \dots, G'_j}} \lambda'_{G_j \circ \dots \circ G_1 \circ T \circ G'_1 \circ \dots \circ G'_j} M_{G_j \circ \dots \circ G_1 \circ T \circ G'_1 \circ \dots \circ G'_j}$$

We first partition this sum based on the shapes $\gamma_j, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_j$ of $G_j, \dots, G_1, T, G'_1, \dots, G'_j$. We then partition this sum further based on the intersection pattern $P \in \mathcal{P}_{\gamma_j, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_j}^{\text{mid}}$.

Definition 6.44. Define $N_P(\tau_P)$ to be the number of ways to choose ribbons $G_j, \dots, G_1, T, G'_1, \dots, G'_j$ of shapes $\gamma_j, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_j$ so that they have intersection pattern P and $G_j \circ \dots \circ G_1 \circ T \circ G'_1 \circ \dots \circ G'_j \circ \dots \circ G'_j \circ \dots \circ G_1 \circ T \circ G'_1 \circ \dots \circ G'_j = T_P$ for a given ribbon T_P of shape τ_P .

By symmetry, this is independent of the choice of T_P .

Recall that an intersection pattern specifies both intersections and how the shapes should be glued together via bijections between the V of each shape and the U of the following shape ([Definition 5.6](#)).

Definition 6.45. We say that two intersection patterns are equivalent if there is an element of

$$\text{Aut}(\gamma_j) \times \dots \times \text{Aut}(\gamma_j) \times \text{Aut}(\tau) \times \text{Aut}(\gamma'_1 \circ \dots \circ \gamma'_1) \times \dots \times \text{Aut}(\gamma'_j \circ \dots \circ \gamma'_j)$$

which maps one intersection pattern to the other (where the gluing maps are permuted accordingly).

This gives us the following equation:

$$\begin{aligned} & \sum_{\substack{\text{sparse} \\ \text{permissible} \\ G_j, \dots, G_1, T, G'_1, \dots, G'_j}} \lambda'_{G_j \circ \dots \circ G_1 \circ T \circ G'_1 \circ \dots \circ G'_j} M_{G_j \circ \dots \circ G_1 \circ T \circ G'_1 \circ \dots \circ G'_j} \\ &= \sum_{\substack{\text{sparse} \\ \text{permissible} \\ \gamma_j, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_j}} \sum_{\substack{\text{nonequivalent} \\ P \in \mathcal{P}^{\text{mid}} \\ \gamma_j, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_j}} N_P(\tau_P) \lambda'_{\gamma_j \circ \dots \circ \gamma_1 \circ \tau \circ \gamma'_1 \circ \dots \circ \gamma'_j} \frac{M_{\tau_P}}{|\text{Aut}(\tau_P)|} \end{aligned}$$

Grouping into shapes, here is the full decomposition:

Lemma 6.46 (Decomposition in terms of shapes).

$$\begin{aligned} & \sum_{R \in \mathcal{S}} \lambda_R M_R = \\ & \Pi^{1/2} \left(\sum_{\substack{\text{sparse} \\ \text{permissible} \\ \sigma \in \mathcal{L}}} \lambda'_\sigma \frac{M_\sigma}{|\text{Aut}(\sigma)|} \right) \cdot \\ & \left(\sum_{j=0}^{2D_{\text{SoS}}} (-1)^j \sum_{\substack{\text{sparse} \\ \text{permissible} \\ \gamma_j, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_j}} \sum_{\substack{\text{nonequivalent} \\ P \in \mathcal{P}^{\text{mid}} \\ \gamma_j, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_j}} N_P(\tau_P) \lambda'_{\gamma_j \circ \dots \circ \gamma_1 \circ \tau \circ \gamma'_1 \circ \dots \circ \gamma'_j} \frac{M_{\tau_P}}{|\text{Aut}(\tau_P)|} \right) \cdot \\ & \left(\sum_{\substack{\text{sparse} \\ \text{permissible} \\ \sigma \in \mathcal{L}}} \lambda'_\sigma \frac{M_\sigma}{|\text{Aut}(\sigma)|} \right)^\top \Pi^{1/2} \\ & + \Pi^{1/2} \text{truncation error}_{\text{too many vertices}} \Pi^{1/2} + \Pi^{1/2} \text{truncation error}_{\text{too many edges in one part}} \Pi^{1/2} \end{aligned}$$

where

1.

$$\begin{aligned} & \text{truncation error}_{\text{too many vertices}} = - \sum_{\substack{\text{sparse, permissible} \\ \sigma, \tau, \sigma': \\ |V(\sigma \circ \tau \circ \sigma')| > D_V}} \lambda'_{\sigma \circ \tau \circ \sigma'} \frac{M_{\sigma \circ \tau \circ \sigma'}}{|\text{Aut}(\sigma \circ \tau \circ \sigma')|} \\ & + \sum_{j=1}^{2D_{\text{SoS}}} (-1)^{j+1} \sum_{\substack{\text{sparse, permissible} \\ \sigma, \gamma_j, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_j, \sigma': \\ |V(\sigma \circ \gamma_j \circ \dots \circ \gamma_1)| > D_V \text{ or} \\ |V(\gamma_1 \circ \dots \circ \gamma'_j \circ \sigma')| > D_V}} \sum_{\substack{\text{nonequiv.} \\ P \in \mathcal{P}^{\text{mid}} \\ \gamma_j, \dots, \gamma'_j}} N_P(\tau_P) \lambda'_{\sigma \circ \gamma_j \circ \dots \circ \gamma'_j \circ \sigma'} \frac{M_\sigma}{|\text{Aut}(\sigma)|} \frac{M_{\tau_P}}{|\text{Aut}(\tau_P)|} \frac{M_{\sigma'}^\top}{|\text{Aut}(\sigma')|} \end{aligned}$$

$$\begin{aligned}
\text{truncation error}_{\text{too many edges in one part}} = & \sum_{\substack{\text{sparse, permissible} \\ \sigma, \tau, \sigma': \\ |V(\sigma \circ \tau \circ \sigma')| \leq D_V, \\ |E_{\text{mid}}(\tau) - |V(\tau)|| > CD_{\text{SoS}}}} \lambda'_{\sigma \circ \tau \circ \sigma'} \frac{M_{\sigma \circ \tau \circ \sigma'}}{|\text{Aut}(\sigma \circ \tau \circ \sigma')|} \\
+ \sum_{j=1}^{2D_{\text{SoS}}} (-1)^j & \sum_{\substack{\text{sparse, permissible} \\ \sigma, \gamma_j, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_j, \sigma': \\ |V(\sigma \circ \gamma_j \circ \dots \circ \gamma_1)| \leq D_V \text{ and } |V(\gamma'_1 \circ \dots \circ \gamma'_j \circ \sigma')| \leq D_V \\ |E_{\text{mid}}(\gamma_j) - |V(\gamma_j)|| > CD_{\text{SoS}} \text{ or } |E_{\text{mid}}(\gamma'_j) - |V(\gamma'_j)|| > CD_{\text{SoS}}}} \sum_{\substack{\text{nonequiv.} \\ P \in \mathcal{P}_{\gamma_j, \dots, \gamma'_j}^{\text{mid}}}} \lambda'_{\sigma \circ \gamma_j \circ \dots \circ \gamma'_j \circ \sigma'} \frac{M_{\sigma \circ \tau_P \circ \sigma'}}{|\text{Aut}(\sigma \circ \tau_P \circ \sigma')|}
\end{aligned}$$

2. in all of these sums, the shapes $\sigma, \gamma_j, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_j, \sigma'$ satisfy the following conditions:

- (a) $\tau \in \mathcal{M}, \sigma, \sigma' \in \mathcal{L}$, and each $\gamma_i, \gamma'_i \in \mathcal{L}$.
- (b) $\sigma, \gamma_j, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_j, \sigma'$ are composable.
- (c) $|V(\tau)| \leq D_V, |V(\gamma_j \circ \dots \circ \gamma_1)| \leq D_V$, and $|V(\gamma'_1 \circ \dots \circ \gamma'_j)| \leq D_V$
- (d) Except when noted otherwise (which only happens for truncation error_{too many edges in one part}), all of the shapes $\gamma_j, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_j$ (but not necessarily σ, σ') satisfy $|E_{\text{mid}}(\alpha) - |V(\alpha)|| \leq CD_{\text{SoS}}$.

To analyze these expressions, we need upper bounds on the number of possible intersection patterns, and on $N_P(\tau_P)$. To gain slightly more control we furthermore partition the intersection patterns based on how many intersections occur.

Lemma 6.47. *There are at most*

$$(4D_{\text{SoS}})^{|V(\tau_P)| - \frac{|U_{\tau_P}| + |V_{\tau_P}|}{2}} + k(3D_V)^k$$

non-equivalent intersection patterns $P \in \mathcal{P}_{\gamma_j, \dots, \gamma'_j}^{\text{mid}}$ which have exactly k intersections.

Proof. We first show a naïve bound. To specify an intersection pattern, it is sufficient to specify how the shapes $\gamma_j, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_j$ of $G_j, \dots, G_1, T, G'_1, \dots, G'_j$ glue together and which vertices intersect with each other. There are $\left(\prod_{i=1}^j |V_{\gamma_i}|\right) \left(\prod_{i=1}^j |V_{\gamma'_i}|\right)$ ways to glue the shapes together. To specify the intersections, we can give the following data.

1. For each i and each vertex $v \in V(\gamma_i) \setminus V_{\gamma_i}$, specify whether or not v intersects with a vertex in $\gamma_{i-1}, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_i$. Similarly, for each vertex in $v \in V(\gamma'_i) \setminus U_{\gamma'_i}$, specify whether it intersects with a vertex in $\gamma_{i-1}, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_{i-1}$.

The number of possibilities for this is at most

$$2^{\sum_{i=1}^j |V(\gamma_i) \setminus V_{\gamma_i}| + |V(\gamma'_i) \setminus V_{\gamma'_i}|}$$

2. For each vertex which intersects with another vertex, specify a vertex which it intersects with. There are at most $3D_V$ choices for this and this happens k times where k is the number of intersections.

This gives a total bound of

$$2^{\sum_{i=1}^j |V(\gamma_i) \setminus V_{\gamma_i}| + |V(\gamma'_i) \setminus V_{\gamma'_i}|} \left(\prod_{i=1}^j |V_{\gamma_i}|! \right) \left(\prod_{i=1}^j |V_{\gamma'_i}|! \right) (3D_V)^k$$

To improve this bound, we observe that there may be some redundancy in the choice of gluing maps. In particular, given a permissible shape α , let $Dormant(\alpha)$ be the set of vertices in $U_\alpha \cap V_\alpha$ which are only involved in missing edge indicators and quasi-missing edge indicators (and are thus indistinguishable). Instead of specifying how the vertices in $Dormant(\alpha)$ match up with the vertices in the previous shape and how these vertices match up with the vertices in the next shape, we can make an arbitrary choice for one of these gluings and make the appropriate choice for the other gluing. This saves us a factor of $|Dormant(\alpha)|!$

Thus, we can divide our naïve bound by

$$\left(\prod_{i=1}^j |Dormant(\gamma_i)|! \right) |Dormant(\tau)|! \left(\prod_{i=1}^j |Dormant(\gamma'_i)|! \right)$$

which gives a bound of

$$\frac{2^{\sum_{i=1}^j |V(\gamma_i) \setminus V_{\gamma_i}| + |V(\gamma'_i) \setminus V_{\gamma'_i}|} \left(\prod_{i=1}^j |V_{\gamma_i}|! \right) \left(\prod_{i=1}^j |V_{\gamma'_i}|! \right) (3D_V)^k}{\left(\prod_{i=1}^j |Dormant(\gamma_i)|! \right) |Dormant(\tau)|! \left(\prod_{i=1}^j |Dormant(\gamma'_i)|! \right)}$$

Since each connected component of $\alpha \setminus (U_\alpha \cap V(\alpha))$ can only have an edge to one vertex in $(U_\alpha \cap V(\alpha))$ and the components connected to $(U_\alpha \cup V(\alpha)) \setminus (U_\alpha \cap V(\alpha))$ have no such edges, we have that $|(U_\alpha \cap V_\alpha) \setminus Dormant(\alpha)| \leq |V(\alpha) \setminus (U_\alpha \cup V_\alpha)|$. Thus, our bound is at most

$$\frac{2^{\sum_{i=1}^j |V(\gamma_i) \setminus V_{\gamma_i}| + |V(\gamma'_i) \setminus V_{\gamma'_i}|} \left(\prod_{i=1}^j |V_{\gamma_i}|! \right) \left(\prod_{i=1}^j |V_{\gamma'_i}|! \right) (3D_V)^k}{\left(\prod_{i=1}^j \frac{|U_{\gamma_i} \cap V_{\gamma_i}|!}{D_{SoS}^{|V(\gamma_i) \setminus (U_{\gamma_i} \cup V_{\gamma_i})|}} \right) \left(\frac{|U_\tau \cap V_\tau|!}{D_{SoS}^{|V(\tau) \setminus (U_\tau \cup V_\tau)|}} \right) \left(\prod_{i=1}^j \frac{|U_{\gamma'_i} \cap V_{\gamma'_i}|!}{D_{SoS}^{|V(\gamma'_i) \setminus (U_{\gamma'_i} \cup V_{\gamma'_i})|}} \right)}$$

To cancel out the $\left(\prod_{i=1}^j |V_{\gamma_i}|! \right) \left(\prod_{i=1}^j |V_{\gamma'_i}|! \right)$ factors, we use the following proposition.

Proposition 6.48. *For any composable shapes $\alpha_1, \dots, \alpha_j$,*

$$\prod_{i=1}^{j-1} |V_{\alpha_i}|! \leq D_{SoS}^{\sum_{i=1}^j \frac{|U_{\alpha_i} \setminus V_{\alpha_i}| + |V_{\alpha_i} \setminus U_{\alpha_i}|}{2}} \prod_{i=1}^j |U_{\alpha_i} \cap V_{\alpha_i}|!$$

Proof. Observe that for all $i \in [j-1]$, $|V_{\alpha_i}|! \leq D_{SoS}^{|V_{\alpha_i} \setminus U_{\alpha_i}|} |U_{\alpha_i} \cap V_{\alpha_i}|!$ so

$$\prod_{i=1}^{j-1} |V_{\alpha_i}|! \leq D_{SoS}^{\sum_{i=1}^{j-1} |V_{\alpha_i} \setminus U_{\alpha_i}|} \prod_{i=1}^j |U_{\alpha_i} \cap V_{\alpha_i}|!$$

Similarly, for all $i \in [j-1]$, $|V_{\alpha_{i+1}}|! = |U_{\alpha_{i+1}}|! \leq D_{SoS}^{|U_{\alpha_{i+1}} \setminus V_{\alpha_{i+1}}|} |U_{\alpha_{i+1}} \cap V_{\alpha_{i+1}}|!$ so

$$\prod_{i=1}^{j-1} |V_{\alpha_i}|! \leq D_{SoS}^{\sum_{i=1}^{j-1} |U_{\alpha_{i+1}} \setminus V_{\alpha_{i+1}}|} \prod_{i=1}^j |U_{\alpha_i} \cap V_{\alpha_i}|!$$

Multiplying these two inequalities together and taking the square root gives the result. ■

We now observe that for any shape α ,

$$|V(\alpha) \setminus (U_\alpha \cup V_\alpha)| + \frac{|U_\alpha \setminus V_\alpha| + |V_\alpha \setminus U_\alpha|}{2} = |V(\alpha)| - \frac{|U_\alpha| + |V_\alpha|}{2}$$

We also observe that for any shape α , $V(\alpha) \setminus |V_\alpha| \leq 2 \left(|V(\alpha)| - \frac{|U_\alpha| + |V_\alpha|}{2} \right)$. Putting everything together, our upper bound is

$$(4D_{SoS}) \left(\sum_{i=1}^j |V(\gamma_i)| - \frac{|U_{\gamma_i}| + |V_{\gamma_i}|}{2} \right) + (|V(\tau)| - \frac{|U_\tau| + |V_\tau|}{2}) + \left(\sum_{i=1}^j |V(\gamma'_i)| - \frac{|U_{\gamma'_i}| + |V_{\gamma'_i}|}{2} \right) (3D_V)^k$$

Since

$$\begin{aligned} & \left(\sum_{i=1}^j |V(\gamma_i)| - \frac{|U_{\gamma_i}| + |V_{\gamma_i}|}{2} \right) + \left(|V(\tau)| - \frac{|U_\tau| + |V_\tau|}{2} \right) + \left(\sum_{i=1}^j |V(\gamma'_i)| - \frac{|U_{\gamma'_i}| + |V_{\gamma'_i}|}{2} \right) \\ &= |V(\tau_P)| - \frac{|U_{\tau_P}| + |V_{\tau_P}|}{2} + k \end{aligned}$$

the result follows. ■

Lemma 6.49. *For any intersection pattern P which has exactly k intersections,*

$$N_P(\tau_P) \leq \frac{D_{SoS}^{2k + |V(\tau_P)| - \frac{|U_{\tau_P}| + |V_{\tau_P}|}{2}}}{|U_{\tau_P} \cap V_{\tau_P}|!} |Aut(\tau_P)|$$

Proof. Observe that if we are given a ribbon of shape τ_P , an element of $Aut(\tau_P)$, the intersection pattern P allows us to recover the original ribbons. This implies that $N_P(\tau_P) \leq |Aut(\tau_P)|$. However, this bound is not quite good enough as we need a factor of $\frac{1}{|U_{\tau_P} \cap V_{\tau_P}|!}$

To obtain this factor, we observe that there may be some redundancy in the information provided by an element of $Aut(\tau_P)$. If we let $Dormant(P)$ be the set of vertices u in $U_{\tau_P} \cap V_{\tau_P}$ such that

1. u is not involved in any intersections
2. u is only incident to missing edge indicators and quasi-missing edge indicators.

then the vertices in $Dormant(P)$ are indistinguishable from each other, so permuting these vertices will not change the ribbons of shapes $\gamma_j, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_j$ we end up with. Thus, there are at least $|Dormant(P)|!$ elements in $Aut(\tau_P)$ which will result in the same ribbons which implies that $N_P(\tau_P) \leq \frac{|Aut(\tau_P)|}{|Dormant(P)|!}$

We now observe that for each vertex $v \in U_{\tau_P} \cap V_{\tau_P}$, at least one of the following 3 cases must hold:

1. v is involved in an intersection.
2. v is in the intersection of U and V for all of the shapes
3. v is dormant and is in the intersection of U and V for all of the shapes

This implies that

$$|Dormant(P)| \geq |U_{\tau_P} \cap V_{\tau_P}| - k - \left(\sum_{i=1}^j |V(\gamma_i) \setminus (U_{\gamma_i} \cup V_{\gamma_i})| \right) \\ - |V(\tau) \setminus (U_{\tau} \cup V_{\tau})| - \left(\sum_{i=1}^j |V(\gamma'_i) \setminus (U_{\gamma'_i} \cup V_{\gamma'_i})| \right)$$

Since

$$\left(\sum_{i=1}^j |V(\gamma_i) \setminus (U_{\gamma_i} \cup V_{\gamma_i})| \right) + |V(\tau) \setminus (U_{\tau} \cup V_{\tau})| + \left(\sum_{i=1}^j |V(\gamma'_i) \setminus (U_{\gamma'_i} \cup V_{\gamma'_i})| \right) \leq k + |V(\tau_P)| - \frac{|U_{\tau_P}| + |V_{\tau_P}|}{2}$$

we have that $\frac{1}{|Dormant(P)|!} \leq \frac{D_{SoS}^{2k+|V(\tau_P)| - \frac{|U_{\tau_P}| + |V_{\tau_P}|}{2}}}{|U_{\tau_P} \cap V_{\tau_P}|!}$ and the result follows. \blacksquare

6.7 Counting shapes with the tail bound function $c(\alpha)$

We have collected the unsymmetrized ribbons into symmetrized shapes, and we must upper bound the number of possible shapes σ , γ_i and τ that are summed over. We will do this by introducing a “tail bound function” $c(\alpha)$, such that $\sum_{\text{nontrivial permissible shapes } \alpha} 1/c(\alpha) \ll 1$. The function will be of the form $c(\alpha) = f(\alpha)^{|V(\alpha) \setminus (U_{\alpha} \cap V_{\alpha})|} g(|E(\alpha)|)$. That is, we need a vertex decay of $f(\alpha)$ per vertex not in $U_{\alpha} \cap V_{\alpha}$, as well as a decay from conditioning when the shape α has too many edges.

Lemma 6.50. *For all $v \in \mathbb{N}$ and all $e, k \in \mathbb{N} \cup \{0\}$, there are at most*

$$2D_{SoS}^{k+1} 2^v v^{2e}$$

permissible shapes α with degree at most D_{SoS} such that $|V(\alpha) \setminus (U_{\alpha} \cap V_{\alpha})| = v$, there are e edges in $E(\alpha)$ which are not incident to a vertex in $U_{\alpha} \cap V_{\alpha}$, and $\alpha \setminus (U_{\alpha} \cap V_{\alpha})$ has k connected components which are disconnected from $U_{\alpha} \cup V_{\alpha}$.

Proof. We can choose such a shape α as follows:

1. First choose the number of vertices in $U_{\alpha} \cap V_{\alpha}$. There are at most $D_{SoS} + 1$ possibilities for this as $0 \leq |U_{\alpha} \cap V_{\alpha}| \leq D_{SoS}$.
2. Indicate whether $U_{\alpha} \setminus V_{\alpha}$ and $V_{\alpha} \setminus U_{\alpha}$ are non-empty. There are 4 choices for this.
3. Order the vertices in $V(\alpha) \setminus (U_{\alpha} \cap V_{\alpha})$ so that the vertices in $U_{\alpha} \setminus V_{\alpha}$ come first (if there are any), followed by the vertices in $V_{\alpha} \setminus U_{\alpha}$ (if there are any), followed by the vertices in each of the k connected components of $\alpha \setminus (U_{\alpha} \cap V_{\alpha})$ which are disconnected from $U_{\alpha} \cup V_{\alpha}$ (where we put the vertex which has an edge to $U_{\alpha} \cap V_{\alpha}$ first), followed by any remaining vertices.
4. For each vertex except for the first vertex, indicate whether it is in the same group as the previous vertex or is the first vertex of the next group. There are at most 2^{v-1} choices for this.
5. For the first vertex w_i in each connected component W_i of $\alpha \setminus (U_{\alpha} \cap V_{\alpha})$ which is disconnected from $U_{\alpha} \cup V_{\alpha}$, choose which vertex in $U_{\alpha} \cap V_{\alpha}$ w_i is adjacent to. There are at most D_{SoS}^k choices for this.

6. For each edge which is not incident to a vertex in $U_\alpha \cap V_\alpha$, choose its two endpoints. There are at most v^{2e} choices for this.
7. The missing edge indicators and quasi-missing edge indicators are fixed because α is a permissible shape. ■

This bound works well when $v \leq C \cdot D_{\text{SoS}}$. For larger shapes, we will need a different bound.

Lemma 6.51. *For all $v \in \mathbb{N}$ and all $e, k, C \in \mathbb{N} \cup \{0\}$, there are at most*

$$2(D_{\text{SoS}} + 1)^{k+1} 8^v 4^e ((C + 4)D_{\text{SoS}})^{(C+2)D_{\text{SoS}}} v^{\max\{0, e+k-v-CD_{\text{SoS}}\}}$$

permissible shapes α with degree at most D_{SoS} such that $|V(\alpha) \setminus (U_\alpha \cap V_\alpha)| = v$, there are e edges in $E(\alpha)$ which are not incident to a vertex in $U_\alpha \cap V_\alpha$, and $\alpha \setminus (U_\alpha \cap V_\alpha)$ has k connected components which are disconnected from $U_\alpha \cup V_\alpha$.

Proof. This can be proved in the same way as Lemma 6.50 except that we handle the edges which are not incident to a vertex in $U_\alpha \cap V_\alpha$ in a different way. We order these based on a breadth first search where we start with the vertices in $U_\alpha \cup V_\alpha$ in the queue.

1. For each new vertex which we reach, indicate whether it will be reached again by one of the first CD_{SoS} edges which do not go to a new vertex. Also indicate whether it has edges incident to it which lead to a new vertex (in which case it should be added to the queue) or not.
2. For each edge, first indicate whether it starts from the same vertex as the previous edge or starts from the next vertex in the queue. Then indicate whether this edge reaches a new vertex or goes to a vertex which has already been reached. If it is one of the first $(C + 2)D_{\text{SoS}}$ edges which goes to a vertex which has already been reached, indicate which vertex this is. Note that there are at most $(C + 4)D_{\text{SoS}}$ choices for this as this vertex is either in $U_\alpha \cup V_\alpha$ or is one of the indicated destinations for these edges (of which there are at most $(C + 2)D_{\text{SoS}}$). If this is a later edge which goes to a vertex which has already been reached, specify this vertex (there are at most v choices for this).

We now observe that at least $v - k - 2D_{\text{SoS}}$ edges which are not incident to a vertex in $U_\alpha \cap V_\alpha$ are needed to connect $U_\alpha \cup V_\alpha$ to the rest of α , so there are at most $e + k - v + 2D_{\text{SoS}}$ edges which do not go to a new vertex. After we take the first $(C + 2)D_{\text{SoS}}$ such edges, there are at most $e + k - v - CD_{\text{SoS}}$ such edges left. ■

Definition 6.52. *Let C' be an upper bound on the sparsity of shapes in Lemma 6.46. Given a permissible shape α such that $|V(\alpha) \setminus (U_\alpha \cap V_\alpha)| = v \geq 1$, there are e edges in $E(\alpha)$ which are not incident to a vertex in $U_\alpha \cap V_\alpha$, and $\alpha \setminus (U_\alpha \cap V_\alpha)$ has k connected components which are disconnected from $U_\alpha \cup V_\alpha$, we define $c(\alpha)$ by:*

1. If $v \leq CD_{\text{SoS}}$ and $e \leq C'v$ then

$$c(\alpha) = 40(2CD_{\text{SoS}}^{4C'+2})^{|V(\alpha)| - \frac{|U_\alpha| + |V_\alpha|}{2}}$$

2. If $v \leq CD_{S_0S}$ and $e > C'v$ then

$$c(\alpha) = 40(2CD_{S_0S}^{4C'+2})^{|V(\alpha)| - \frac{|U_\alpha| + |V_\alpha|}{2}} (2C^2D_{S_0S}^2)^{(e-C'v)}$$

3. If $v > CD_{S_0S}$ and $e \leq v - k + CD_{S_0S}$ then

$$c(\alpha) = 40(16CD_{S_0S}^4)^{|V(\alpha)| - \frac{|U_\alpha| + |V_\alpha|}{2}}$$

4. If $v > CD_{S_0S}$ and $e > v - k + CD_{S_0S}$ then

$$c(\alpha) = 40(16CD_{S_0S}^4)^{|V(\alpha)| - \frac{|U_\alpha| + |V_\alpha|}{2}} (8D_V)^{e+k-v-CD_{S_0S}}$$

If α is a trivial shape then we define $c(\alpha) = 1$.

Corollary 6.53.

$$\sum_{\substack{\alpha \in \mathcal{S}: \\ \text{non-trivial,} \\ \text{permissible}}} \frac{1}{c(\alpha)} \leq \frac{1}{10}.$$

6.8 Statement of main lemmas

Now that we have wrangled the moment matrix into the correct form, we will show that all the error terms are small. Recall the expression for the moment matrix from [Lemma 6.46](#):

$$\begin{aligned} & \Pi^{1/2} \left(\sum_{\substack{\text{sparse} \\ \text{permissible} \\ \sigma \in \mathcal{L}}} \lambda'_\sigma \frac{M_\sigma}{|\text{Aut}(\sigma)|} \right) \cdot \\ & \left(\sum_{j=0}^{2D_{S_0S}} (-1)^j \sum_{\substack{\text{sparse} \\ \text{permissible} \\ \gamma_j, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_j}} \sum_{\substack{\text{nonequivalent} \\ P \in \mathcal{P}^{mid} \\ \gamma_j, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_j}} N_P(\tau_P) \lambda'_{\gamma_j \circ \dots \circ \gamma_1 \circ \tau \circ \gamma'_1 \circ \dots \circ \gamma'_j} \frac{M_{\tau_P}}{|\text{Aut}(\tau_P)|} \right) \cdot \\ & \left(\sum_{\substack{\text{sparse} \\ \text{permissible} \\ \sigma \in \mathcal{L}}} \lambda'_\sigma \frac{M_\sigma}{|\text{Aut}(\sigma)|} \right)^\top \Pi^{1/2} \\ & + \Pi^{1/2} \cdot \text{truncation error} \cdot \Pi^{1/2} \end{aligned}$$

The following lemmas, for which the intuition was given in [Section 6.2](#), are sufficient to prove [Theorem 1.4](#).

Lemma 6.54. (Non-trivial Middle Shapes) For all sparse permissible $\tau \in \mathcal{M}$ such that $|V(\tau)| > \frac{|U_\tau| + |V_\tau|}{2}$ and $|E_{mid}(\tau)| - |V(\tau)| \leq CD_{S_0S}$,

$$\lambda'_\tau \frac{\|M_\tau\|}{|\text{Aut}(\tau)|} \leq \frac{1}{c(\tau)}.$$

Lemma 6.55. (*Intersection Terms*) For all $j \geq 1$ and sparse permissible $\gamma_j, \dots, \tau, \dots, \gamma'_j$ such that for each shape $|E_{mid}(\alpha)| - |V(\alpha)| \leq CD_{SoS}$,

$$\sum_{\substack{\text{nonequivalent} \\ P \in \mathcal{P}^{mid} \\ \gamma_j, \dots, \gamma'_j}} N_P(\tau_P) \lambda'_{\gamma_j \circ \dots \circ \gamma'_j} \frac{\|M_{\tau_P}\|}{|\text{Aut}(\tau_P)|} \leq \frac{1}{c(\tau) \prod_{i=1}^j c(\gamma_i) c(\gamma'_i)}.$$

Lemma 6.56. (*Truncation Error*)

$$\text{truncation error} \preceq n^{-\Omega(CD_{SoS})} \pi.$$

Lemma 6.57. (*Sum of left shapes is well-conditioned*)

$$\left(\sum_{\substack{\text{sparse,} \\ \text{permissible} \\ \sigma \in \mathcal{L}}} \lambda'_\sigma \frac{M_\sigma}{|\text{Aut}(\sigma)|} \right) \left(\sum_{\substack{\text{sparse,} \\ \text{permissible} \\ \sigma \in \mathcal{L}}} \lambda'_\sigma \frac{M_\sigma}{|\text{Aut}(\sigma)|} \right)^\top \succeq n^{-O(D_{SoS})} \pi$$

Proof of Theorem 1.4 assuming Lemma 6.54, Lemma 6.55, Lemma 6.56, Lemma 6.57. For ease of exposition we omit the automorphism groups. In the approximate PSD decomposition, the term $j = 0$ can be broken up into the leading term π and remaining nontrivial middle shapes,

$$\begin{aligned} \sum_{\substack{\text{sparse,} \\ \text{permissible} \\ \tau \in \mathcal{M}}} \lambda'_\tau M_\tau &= \sum_{\substack{\text{permissible} \\ \tau \in \mathcal{M}: \\ U_\tau = V_\tau = V(\tau)}} \lambda_\tau M_\tau &+ \sum_{\substack{\text{sparse,} \\ \text{permissible} \\ \tau \in \mathcal{M}: \\ |V(\tau)| > \frac{|U_\tau| + |V_\tau|}{2}}} \lambda'_\tau M_\tau \\ &= \pi &+ \sum_{\substack{\text{sparse,} \\ \text{permissible} \\ \tau \in \mathcal{M}: \\ |V(\tau)| > \frac{|U_\tau| + |V_\tau|}{2}}} \lambda'_\tau M_\tau. \end{aligned}$$

Using Lemma 6.25, the middle shapes ($j = 0$) and intersection terms ($j \geq 1$) are

$$\begin{aligned} &\pi \left(\text{Id} + \sum_{\substack{\text{sparse,} \\ \text{permissible} \\ \tau \in \mathcal{M}: \\ |V(\tau)| > \frac{|U_\tau| + |V_\tau|}{2}}} \lambda'_\tau M_\tau \right. \\ &\quad \left. + \sum_{j=1}^{2D_{SoS}} (-1)^j \sum_{\substack{\text{sparse,} \\ \text{permissible} \\ \gamma_j, \dots, \gamma'_j}} \sum_{\substack{\text{nonequivalent:} \\ P \in \mathcal{P}^{mid} \\ \gamma_j, \dots, \gamma'_j}} N_P(\tau_P) \lambda'_{\gamma_j \circ \dots \circ \gamma'_j} M_{\tau_P} \right). \end{aligned}$$

By Lemma 6.54 and Lemma 6.55, (summed with Corollary 6.53) this is at least $\Omega(1)\pi$.

Now plugging this into the PSD decomposition, we have:

$$\sum_{\alpha \in \mathcal{S}} \lambda_\alpha \frac{M_\alpha}{|\text{Aut}(\alpha)|} \succeq \Pi^{1/2} \left(\sum_{\substack{\text{sparse} \\ \sigma \in \mathcal{L}}} \lambda'_\sigma M_\sigma \right) \Omega(1)\pi \left(\sum_{\substack{\text{sparse} \\ \sigma \in \mathcal{L}}} \lambda'_\sigma M_\sigma \right)^\top \Pi^{1/2} + \text{truncation error}$$

By Lemma 6.25 again,

$$= \Omega(1)\Pi^{1/2} \left(\sum_{\substack{\text{sparse} \\ \text{permissible} \\ \sigma \in \mathcal{L}}} \lambda'_\sigma M_\sigma \right) \left(\sum_{\substack{\text{sparse} \\ \text{permissible} \\ \sigma \in \mathcal{L}}} \lambda'_\sigma M_\sigma \right)^\top \Pi^{1/2} + \text{truncation error}$$

By Lemma 6.57, Lemma 6.56, and taking C sufficiently large,

$$\begin{aligned} &\succeq n^{-O(D_{\text{SoS}})}\Pi - n^{-\Omega(CD_{\text{SoS}})}\Pi \\ &\succeq 0 \end{aligned}$$

■

6.9 Conditioning II: Frobenius norm trick

Shapes with a large number of excess edges should be handled using their Frobenius norm. The Frobenius norm trick improves on the first moment method, Proposition 6.33, so it could be used to also handle $O(1)$ -sparse subgraphs in this work. On the other hand, the forbidden subgraph method can handle forbidden graphs beyond the first moment method, which may lead to improvements for smaller d .

There are two parts to the Frobenius norm trick. First, we compute the Frobenius norm of a graph matrix.

Lemma 6.58. *For all proper shapes α ,*

$$\mathbb{E}[\text{tr}(M_\alpha M_\alpha^\top)] \leq |\text{Aut}(\alpha)| n^{|\mathcal{V}(\alpha)| + |I_\alpha|}.$$

Proof. Any term contributing to $\mathbb{E}[\text{tr}(M_\alpha M_\alpha^\top)]$ is a labeling of α and α^\top s.t. every edge appears exactly twice. After choosing the labels for α in $n^{|\mathcal{V}(\alpha)|}$ ways, the labeling of α^\top gives an isomorphism between α and α^\top (except for isolated vertices, which give an extra factor of n each). ■

For proper shapes α , this gives a norm bound independent of the number of edges! MVS with a large number of induced edges can't hurt us. We have $\lambda_\alpha \|M_\alpha\| \leq \left(\frac{k}{n}\right)^{|\mathcal{V}(\alpha)| - \frac{|U_\alpha| + |V_\alpha|}{2}} \sqrt{p}^{|\mathcal{E}(\alpha)|} \sqrt{n}^{|\mathcal{V}(\alpha)|}$. The number of edges could be smaller than the number of vertices, but only by D_{SoS} (since everything must be connected to $U_\alpha \cup V_\alpha$). If there are at least CD_{SoS} excess edges, then the norm is small enough to sum.

The Frobenius norm is small for too-dense subgraphs of G . The second part of the Frobenius norm trick allows us to use these bounds with high probability instead of in expectation, which we can do using Markov's inequality once instead of using Markov's inequality on each one.

Lemma 6.59. *For any random matrices M_1, \dots, M_k ,*

$$\mathbb{E} \left[\text{tr} \left(\left(\sum_{i=1}^k M_i \right) \left(\sum_{i=1}^k M_i \right)^\top \right) \right] \leq \left(\sum_{i=1}^k \sqrt{\mathbb{E}[\text{tr}(M_i M_i^\top)]} \right)^2$$

Proof. Observe that by Cauchy-Schwarz,

$$\begin{aligned} \mathbb{E} [\text{tr}(M_i M_j^\top)] &= \mathbb{E} \left[\sum_{a,b} (M_i)_{ab} (M_j)_{ab} \right] \\ &\leq \sqrt{\mathbb{E} \left[\sum_{a,b} (M_i)_{ab}^2 \right]} \sqrt{\mathbb{E} \left[\sum_{a,b} (M_j)_{ab}^2 \right]} \\ &= \sqrt{\mathbb{E} [\text{tr}(M_i M_i^\top)]} \sqrt{\mathbb{E} [\text{tr}(M_j M_j^\top)]} \end{aligned}$$

Applying this inequality for all $i, j \in [k]$ gives the result. \blacksquare

6.10 Norm bounds

For the trace method, it is easier to use graph matrices with left/right sides indexed by *ordered* tuples.

Theorem 6.60. *If $D_V \geq \lceil D_{SOS} \ln(n) \rceil$ then for all proper shapes α such that $|V(\alpha)| \leq 3D_V$ and all $\varepsilon' > 0$, taking M_α to be the graph matrix where the rows and columns are indexed by ordered tuples rather than sets,*

$$\Pr \left(\|M_\alpha\| > 10 \left(\frac{1}{\varepsilon'} \right)^{\frac{1}{2D_V}} \max_{\text{separator } S} \left\{ n^{\frac{|V(\alpha)|-|S|}{2}} (12D_V)^{|V(\alpha)|-\frac{|U_\alpha|+|V_\alpha|}{2}+|S|-\frac{|L_S|+|R_S|}{2}} \left(3\sqrt{\frac{1-p}{p}} \right)^{|E(S)|} \right\} \right) < \varepsilon'$$

Proof. By [Theorem A.7](#),

$$\begin{aligned} \mathbb{E} [\text{tr}((M_\alpha M_\alpha^\top)^q)] &\leq 4^{2q} \binom{|V(\alpha)|-\frac{|U_\alpha|+|V_\alpha|}{2}}{2q} \left(\sqrt{|V(\alpha)|} \right)^{2q|V(\alpha)\setminus(U_\alpha \cup V_\alpha)|} n^{q|V(\alpha)|} \\ &\quad \left(\max_{\text{separator } S} n^{-\frac{|S|}{2}} \left(3\sqrt{\frac{1-p}{p}} \right)^{|E(S)|} \left(\sqrt{|V(\alpha)|} \right)^{|S|-|U_\alpha \cap V_\alpha|} (2q)^{|S|-\frac{|L_S|+|R_S|}{2}+\frac{c(\alpha,S)}{2}} \right)^{2q-2} \end{aligned}$$

Since $|V(\alpha)|$ and q will both be $O(D_V) = O(D_{SOS} \log(n))$, for convenience we combine the factors of $|V(\alpha)|$ and q using the following proposition:

Proposition 6.61.

$$\frac{1}{2}|V(\alpha) \setminus (U_\alpha \cup V_\alpha)| + \frac{1}{2}(|S| - |U_\alpha \cap V_\alpha|) + \frac{c(\alpha, S)}{2} \leq |V(\alpha)| - \frac{|U_\alpha| + |V_\alpha|}{2}$$

Proof. Observe that since $c(\alpha, S)$ counts the number of connected components of $\alpha \setminus S$ which are disconnected from $U_\alpha \cup V_\alpha$, we have that $|S| - |U_\alpha \cap V_\alpha| + c(\alpha, S) \leq |V(\alpha)| - |U_\alpha \cap V_\alpha|$. We now observe that $\frac{|V(\alpha) \setminus (U_\alpha \cup V_\alpha)| + |V(\alpha)| - |U_\alpha \cap V_\alpha|}{2} = |V(\alpha)| - \frac{|U_\alpha| + |V_\alpha|}{2}$ \blacksquare

Since $|V(\alpha)| \leq 3D_V$, for all $q \leq \frac{3}{2}D_V$ we have that

$$\begin{aligned} \mathbb{E} [\text{tr}((M_\alpha M_\alpha^\top)^q)] &\leq \\ &\max_{\text{separator } S} \left\{ n^{|S|} \left(n^{\frac{|V(\alpha)|-|S|}{2}} (12D_V)^{|V(\alpha)|-\frac{|U_\alpha|+|V_\alpha|}{2}+|S|-\frac{|L_S|+|R_S|}{2}} \left(3\sqrt{\frac{1-p}{p}} \right)^{|E(S)|} \right)^{2q} \right\} \end{aligned}$$

Using the fact that for all $q \in \mathbb{N}$ and all $\varepsilon' > 0$,

$$\Pr \left(\|M_\alpha\| > \sqrt[2q]{\frac{\mathbb{E} [\text{tr} ((M_\alpha M_\alpha^\top)^q)]}{\varepsilon'}} \right) < \varepsilon'$$

and taking $q = D_V$, we have that

$$\Pr \left(\|M_\alpha\| > 10 \left(\frac{1}{\varepsilon'} \right)^{\frac{1}{2D_V}} \max_{\text{separator } S} \left\{ n^{\frac{|V(\alpha)|-|S|}{2}} (12D_V)^{|V(\alpha)|-\frac{|U_\alpha|+|V_\alpha|}{2}+|S|-\frac{|L_S|+|R_S|}{2}} \left(3\sqrt{\frac{1-p}{p}} \right)^{|E(S)|} \right\} \right) < \varepsilon'$$

Remark 6.62. We will take $\varepsilon' = \frac{\varepsilon}{c(\alpha)}$ where ε is the overall probability that our SoS bound fails. However, we did not want to hard code this choice into the theorem statement.

For improper shapes, we linearize them. The coefficients in the linearization are:

Proposition 6.63. For a shape α ,

$$M_\alpha = \sum_{\text{linearizations } \beta \text{ of } \alpha} \left(\sqrt{\frac{1-p}{p}} \right)^{|E(\alpha)|-|E(\beta)|-2|E_{\text{phantom}}(\beta)|} M_\beta.$$

Proof. Each Fourier character can be linearized as $\chi_e^k = \mathbb{E}[\chi_e^k] + \mathbb{E}[\chi_e^{k+1}]\chi_e$. Then use [Proposition A.1](#). \blacksquare

Corollary 6.64. If $D_V \geq \lceil D_{\text{SoS}} \ln(n) \rceil$ then for any shape α (including improper shapes, shapes with missing edge indicators, and shapes with quasi-missing edge indicators), taking M_α to be the graph matrix where the rows and columns are indexed by sets (i.e. the definition we use throughout the paper),

$$\Pr \left(\|M_\alpha\| > 20 \left(\frac{2^{|E(\alpha)|-|E_{\text{no repetitions}}(\alpha)|}}{\varepsilon'} \right)^{\frac{1}{2D_V}} 2^{|E(\alpha)|-|E_{\text{no repetitions}}(\alpha)|} \sqrt{|U_\alpha|!|V_\alpha|!} \max_{\beta, S} \left\{ n^{\frac{|V(\alpha)|-|S|+|I_\beta|}{2}} (12D_V)^{|V(\alpha)|-\frac{|U_\alpha|+|V_\alpha|}{2}+|S|-\frac{|L_S|+|R_S|}{2}} \left(\sqrt{\frac{1-p}{p}} \right)^{|E(\alpha)|-|E(\beta)|-2|E_{\text{phantom}}(\beta)|} \left(3\sqrt{\frac{1-p}{p}} \right)^{|E(S)|} \right\} \right) < \varepsilon'$$

where $E_{\text{no repetitions}}(\alpha)$ is the set (rather than the multi-set) of edges of α , β is a linearization of α , S is a separator of β , I_β is the set of vertices of β which are isolated, and $E_{\text{phantom}}(\beta)$ is the set of edges of α which are not in β . As a special case, when α is a proper shape with no isolated vertices (which may still contain missing edge indicators and quasi-missing edge indicators),

$$\Pr \left(\|M_\alpha\| > 20 \left(\frac{1}{\varepsilon'} \right)^{\frac{1}{2D_V}} \sqrt{|U_\alpha|!|V_\alpha|!} \max_S \left\{ n^{\frac{|V(\alpha)|-|S|}{2}} (12D_V)^{|V(\alpha)|-\frac{|U_\alpha|+|V_\alpha|}{2}+|S|-\frac{|L_S|+|R_S|}{2}} \left(3\sqrt{\frac{1-p}{p}} \right)^{|E(S)|} \right\} \right) < \varepsilon'$$

Proof. We make the following observations:

1. There are at most $2^{|E(\alpha)|-|E_{\text{no repetitions}}(\alpha)|}$ linearizations β of α , so to ensure that the bound on M_α fails with probability at most ε' , it is sufficient to ensure that the bound for each M_β fails with probability at most $\frac{\varepsilon'}{2^{|E(\alpha)|-|E_{\text{no repetitions}}(\alpha)|}}$.

2. The coefficient for each linearization β of α is at most $\left(\sqrt{\frac{1-p}{p}}\right)^{|E(\alpha)|-|E(\beta)|-2|E_{\text{phantom}}(\beta)|}$ per Proposition 6.63.
3. To handle quasi-missing edge indicators, we observe that each quasi-missing edge is a convex combination of missing edge indicators, so we can express M_β as a convex combination of $M_{\beta'}$ where β' only has missing edge indicators rather than quasi-missing edge indicators. We now use Lemma 5.14 of [AMP20], which says that when we use the trace power method on a matrix M which is a convex combination of matrices M' , we can obtain a bound on $\|M\|$ which is at most twice the maximum over M' of the bound we obtain on $\|M'\|$.
4. Notice in the trace power method, the length- q walk on the matrix indexed by sets follows the same block structure that each α, α^T is specified by an embedding of vertices in $V(\alpha)$ to $[n]$. The only change is vertices in the boundary only agree as a set instead of ordered tuples, to address this, it suffices to additionally identify a permutation between the boundaries, and this requires $|U_\alpha|!$ when we go from $V_{\alpha_t^T}$ to $U_{\alpha_{t+1}}$, and it requires $|V_\alpha|!$ when go from V_{α_t} to $U_{\alpha_t^T}$. Taking the square root gives us the desired bound. ■

One more proposition is needed to handle the automorphism terms that arise when switching between ribbons/shapes/shapes indexed by ordered tuples,

Proposition 6.65. *For a permissible shape α of degree at most D_{SoS} ,*

$$\frac{\sqrt{|U_\alpha|!|V_\alpha|!}}{|\text{Aut}(\alpha)|} \leq D_{\text{SoS}}^{|V(\alpha) \setminus (U_\alpha \cap V_\alpha)|/2 + |V(\alpha) \setminus (U_\alpha \cup V_\alpha)|}$$

Proof. Write

$$\frac{\sqrt{|U_\alpha|!|V_\alpha|!}}{|\text{Aut}(\alpha)|} = \frac{\sqrt{|U_\alpha|!|V_\alpha|!}}{|U_\alpha \cap V_\alpha|!} \cdot \frac{|U_\alpha \cap V_\alpha|!}{|\text{Aut}(\alpha)|}$$

The first ratio is upper bounded by $D_{\text{SoS}}^{|V(\alpha) \setminus (U_\alpha \cap V_\alpha)|/2}$.

For the second ratio, in a permissible shape, $|\text{Aut}(\alpha)| \geq (|U_\alpha \cap V_\alpha| - |V(\alpha) \setminus (U_\alpha \cup V_\alpha)|)!$. This is because in a permissible shape, the number of edges incident to $U_\alpha \cap V_\alpha$ is at most the number of connected components of $V(\alpha) \setminus (U_\alpha \cup V_\alpha)$, which is upper bounded by $|V(\alpha) \setminus (U_\alpha \cup V_\alpha)|$. The remaining vertices of $U \cap V$ that are not incident to an edge are completely interchangeable. Therefore this ratio is upper bounded by $D_{\text{SoS}}^{|V(\alpha) \setminus (U_\alpha \cup V_\alpha)|}$. ■

6.11 Proof of main lemmas

6.11.1 Lemma 6.54: nontrivial middle shapes

Lemma 6.54. *(Non-trivial Middle Shapes) For all sparse permissible $\tau \in \mathcal{M}$ such that $|V(\tau)| > \frac{|U_\tau| + |V_\tau|}{2}$ and $|E_{\text{mid}}(\tau)| - |V(\tau)| \leq CD_{\text{SoS}}$,*

$$\lambda'_\tau \frac{\|M_\tau\|}{|\text{Aut}(\tau)|} \leq \frac{1}{c(\tau)}.$$

Proof. The remaining “core” shapes are $\tau \in \mathcal{M}$ nontrivial, sparse, permissible with $|E(\alpha)| \leq |V(\alpha)| + CD_{\text{SoS}}$. For these shapes, the norm bound in [Corollary 6.64](#) with $\varepsilon' = \frac{1}{nc(\tau)}$ holds whp, via [Corollary 6.53](#).

$$\begin{aligned} & \lambda'_\tau \frac{\|M_\tau\|}{|\text{Aut}(\tau)|} \\ & \leq m(\tau) \left(\lambda_\tau + \sum_{U \supset \tau} \lambda_U c(\tau, U) \right) \cdot \\ & \quad C \left(2^{|E(\tau)|} nc(\tau) \right)^{\frac{1}{2D_V}} 2^{|E(\tau)|} \max_{\text{separator } S} \left\{ n^{\frac{|V(\alpha)|-|S|}{2}} (12D_V)^{|V(\alpha)| - \frac{|U_\alpha|+|V_\alpha|}{2} + |S| - \frac{|L_S|+|R_S|}{2}} \left(3\sqrt{\frac{1-p}{p}} \right)^{|E(S)|} \right\} \\ & \quad \cdot \frac{\sqrt{|U_\tau|!|V_\tau|!}}{|\text{Aut}(\tau)|} \end{aligned}$$

We have:

$$\begin{aligned} m(\tau) & \leq \left(\frac{1}{(1-p)^{2D_{\text{SoS}}}} \right)^{|V(\tau) \setminus (U_\tau \cap V_\tau)|} \leq 2^{|V(\tau) \setminus (U_\tau \cap V_\tau)|} \quad (\text{Lemma 6.28}) \\ \sum_{U \supset \tau} \lambda_U c(\tau, U) & = o(1)\lambda_\tau \quad (\text{Lemma 6.43}) \\ 2^{|E(\alpha)|/2D_V} & \leq C \quad (\tau \text{ is sparse}) \\ n^{1/2D_V} & \leq 2 \\ c(\tau) & \leq CD_{\text{SoS}}^{c|V(\tau) \setminus (U_\tau \cap V_\tau)|} \quad (\text{Definition 6.52, } \tau \text{ is sparse}) \\ c(\tau)^{\frac{1}{2D_V}} & \leq D_{\text{SoS}}^c \quad (\text{From previous line}) \\ 3^{|E(S)|} & \leq C^{|V(\tau) \setminus (U_\tau \cap V_\tau)|} \quad (\tau \text{ is sparse}) \\ \frac{\sqrt{|U_\tau|!|V_\tau|!}}{|\text{Aut}(\tau)|} & \leq CD_{\text{SoS}}^{c|V(\tau) \setminus (U_\tau \cap V_\tau)|} \quad (\text{Proposition 6.65}) \end{aligned}$$

Going through the charging argument for middle shapes in [Proposition 6.1](#),

$$\lambda_\tau n^{\frac{|V(\alpha)|-|S|}{2}} \left(\sqrt{\frac{1-p}{p}} \right)^{|E(S)|} \leq \left(\frac{k\sqrt{d}}{n} \right)^{|V(\tau)| - \frac{|U_\tau|+|V_\tau|}{2}} \cdot \left(\frac{1}{\sqrt{d}} \right)^{|S| - \frac{|U_\tau|+|V_\tau|}{2}}.$$

Using $D_{\text{SoS}} \leq \frac{d^{1/2}}{\log n}$, we have $\sqrt{d} \geq D_V$ and the last term cancels $D_V^{|S| - \frac{|L_S|+|R_S|}{2}}$ from the norm bound after the following claim:

Claim 6.66. For a middle shape τ and any vertex separator S , $|L_S| \geq |U_\tau|$ and $|R_S| \geq |V_\tau|$.

Proof. L_S and R_S are both vertex separators of τ , and since τ is a middle shape, U_τ, V_τ are MVs of τ . ■

Putting it together, the vertex decay of $\frac{k\sqrt{d}}{n}$ on $|V(\tau)| - \frac{|U_\tau|+|V_\tau|}{2}$ needs to be $CD_V D_{\text{SoS}}^c$ to overcome the combinatorial factors, as stated in [Theorem 1.4](#).

$$\lambda'_\tau \frac{\|M_\tau\|}{|\text{Aut}(\tau)|} \leq \frac{1}{c(\tau)}.$$

■

6.11.2 Lemma 6.55: bounding intersection terms

We will need the following proposition.

Proposition 6.67. *For any C' -sparse permissible composable shapes $\alpha_1, \dots, \alpha_k$ and any intersection pattern P on $\alpha_1, \dots, \alpha_k$, letting τ_P be the resulting shape,*

$$|E(\tau_P)| \leq (2C' + 1) \sum_{i=1}^k \left(|V(\alpha_i)| - \frac{|U_{\alpha_i}| + |V_{\alpha_i}|}{2} \right)$$

Proof. Observe that since α_i is permissible, the only edges incident to $U_{\alpha_i} \cap V_{\alpha_i}$ are one for each connected component of $V(\alpha_i) \setminus (U_{\alpha_i} \cup V_{\alpha_i})$. Therefore, there are at most $|V(\alpha_i) \setminus (U_{\alpha_i} \cup V_{\alpha_i})|$ of these edges. Since α_i is C' -sparse,

$$|E(\alpha_i)| \leq |V(\alpha_i) \setminus (U_{\alpha_i} \cup V_{\alpha_i})| + C'|V(\alpha_i) \setminus (U_{\alpha_i} \cap V_{\alpha_i})| \leq (2C' + 1) \left(|V(\alpha_i)| - \frac{|U_{\alpha_i}| + |V_{\alpha_i}|}{2} \right)$$

Since $|E(\tau_P)| = \sum_{i=1}^k |E(\alpha_i)|$, summing this equation over all $i \in [k]$ gives the result. \blacksquare

Lemma 6.55. (Intersection Terms) *For all $j \geq 1$ and sparse permissible $\gamma_j, \dots, \tau, \dots, \gamma'_j$ such that for each shape $|E_{mid}(\alpha)| - |V(\alpha)| \leq CD_{SoS}$,*

$$\sum_{\substack{\text{nonequivalent} \\ P \in \mathcal{P}_{\gamma_j \dots \gamma'_j}^{mid}}} N_P(\tau_P) \lambda'_{\gamma_j \circ \dots \circ \gamma'_j} \frac{\|M_{\tau_P}\|}{|\text{Aut}(\tau_P)|} \leq \frac{1}{c(\tau) \prod_{i=1}^j c(\gamma_i) c(\gamma'_i)}.$$

Proof. We index the intersecting shapes as $\alpha_i = \gamma_k, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_k$. For each α_i , by Lemma 6.28 and Lemma 6.43,

$$\lambda'_{\alpha_i} \leq 2^{|V(\alpha_i) \setminus (U_{\alpha_i} \cap V_{\alpha_i})|} \lambda_{\alpha_i}.$$

Let $p_\ell(\alpha_i)$ be the number of nonequivalent intersection patterns $P \in \mathcal{P}_{\gamma_j \dots \gamma'_j}^{mid}$ which have exactly ℓ intersections.

Apply the norm bound in Corollary 6.64 with $\varepsilon' = \frac{1}{n^C \prod_{i=1}^{2k+1} c(\alpha_i) p_\ell(\alpha_i)}$, which holds whp via Corollary 6.53.

$$\begin{aligned} & \sum_{\text{non equiv. } P} N_P(\tau_P) \cdot \lambda_{\gamma_k \circ \dots \circ \gamma_1 \circ \tau \circ \gamma'_1 \circ \dots \circ \gamma'_k} \frac{\|M_{\tau_P}\|}{|\text{Aut}(\tau_P)|} \\ & \leq \lambda_{\gamma_k \circ \dots \circ \gamma_1 \circ \tau \circ \gamma'_1 \circ \dots \circ \gamma'_k} \cdot 20 \left(2^{|E(\tau_P)|} n^C \prod_{i=1}^{2k+1} c(\alpha_i) p_\ell(\alpha_i) \right)^{\frac{1}{2D_V}} 2^{|E(\tau_P)|} \cdot \frac{\sqrt{|U_{\tau_P}|! |V_{\tau_P}|!}}{|\text{Aut}(\tau_P)|} \\ & \max_{\beta, S} \left\{ n^{\frac{|V(\tau_P)| - |S| + I_\beta|}{2}} (12D_V)^{|V(\tau_P)| - \frac{|U_{\tau_P}| + |V_{\tau_P}|}{2} + |S| - \frac{|L_S| + |R_S|}{2}} \left(\sqrt{\frac{1-p}{p}} \right)^{|E(\tau_P)| - |E(\beta)| - 2|E_{phantom}(\beta)|} \left(3\sqrt{\frac{1-p}{p}} \right)^{|E(S)|} \right\}. \end{aligned}$$

Bounds:

$$12^{|E(\tau_P)|} \leq C^{\sum_{i=1}^{2k+1} |V(\alpha_i)| - \frac{|U_{\alpha_i}| + |V_{\alpha_i}|}{2}} \quad (\text{Proposition 6.67})$$

$$n^{C/2D_V} \leq 2$$

$$c(\alpha_i) \leq C(D_{\text{SoS}}^c)^{|V(\alpha_i) \setminus (U_{\alpha_i} \cap V_{\alpha_i})|} \quad (\text{Definition 6.52, } \alpha_i \text{ is sparse})$$

$$c(\alpha_i)^{1/2D_V} \leq D_{\text{SoS}}^c \quad (\text{From previous line})$$

$$p_\ell(\alpha_i) \leq (D_{\text{SoS}}^c)^{|V(\tau_P)| - \frac{|U_{\tau_P}| + |V_{\tau_P}|}{2}} D_V^\ell \quad (\text{Lemma 6.47})$$

$$\sum_{\text{non equiv. P}} N_P(\tau_P) \leq \frac{(D_{\text{SoS}}^c)^{\sum_{i=1}^{2k+1} |V(\alpha_i)| - \frac{|U_{\alpha_i}| + |V_{\alpha_i}|}{2}}}{|U_{\tau_P} \cap V_{\tau_P}|!} |\text{Aut}(\tau_P)| \quad (\text{Lemma 6.49})$$

$$\frac{\sqrt{|U_{\tau_P}|! |V_{\tau_P}|!}}{|\text{Aut}(\tau_P)|} \leq C(D_{\text{SoS}}^c)^{|V(\tau_P) \setminus (U_{\tau_P} \cap V_{\tau_P})|} \quad (\text{Proposition 6.65})$$

Following the charging argument for intersection terms in [Proposition 6.7](#),

$$\begin{aligned} & \lambda_{\gamma_k \circ \dots \circ \gamma_1 \circ \tau \circ \gamma_1^T \circ \dots \circ \gamma_k^T} n^{\frac{|V(\tau_P)| - |S| + |I_\beta|}{2}} \left(\sqrt{\frac{1-p}{p}} \right)^{|E(S)| + |E(\alpha)| - |E(\beta)| - 2|E_{\text{phantom}}(\beta)|} \\ & \leq \left(\frac{k\sqrt{d}}{n} \right)^{\sum_{i=1}^{2k+1} |V(\alpha_i)| - \frac{|U_{\alpha_i}| + |V_{\alpha_i}|}{2}} \cdot \left(\frac{1}{\sqrt{d}} \right)^{i_P - |I_\beta| + |S| - \frac{|U_\beta| + |V_\beta|}{2}}. \end{aligned}$$

In [Proposition 6.7](#) we used the intersection tradeoff lemma, [Lemma 6.8](#), to prove that the exponent of the second term is nonnegative. In that proof, use the following claim to get a lower bound of $|S| - \frac{|L_S| + |R_S|}{2}$ on the exponent.

Claim 6.68. $\frac{|L_S| + |R_S|}{2} \geq |S_{\tau_P, \min}|$.

Proof. Both L_S and R_S are vertex separators of β , hence they are larger than $S_{\tau_P, \min}$, which is the smallest separator among all linearizations of τ_P . ■

Since $\sqrt{d} \geq D_V$, the second term cancels $D_V^{|S| - \frac{|L_S| + |R_S|}{2}}$ in the norm bound.

The remaining factors are mostly easily seen to be under control by the vertex decay. One exception is the term D_V^ℓ in $p_\ell(\alpha_i)$, the count of nonequivalent intersection patterns. By combining this with $D_V^{|V(\tau_P)| - \frac{|U_{\tau_P}| + |V_{\tau_P}|}{2}}$ from the norm bound we get at most $D_V^{\sum_{i=1}^{2k+1} |V(\alpha_i)| - \frac{|U_{\alpha_i}| + |V_{\alpha_i}|}{2}}$, which is controlled by one factor of D_V in the vertex decay. In total, a vertex decay of $CD_V D_{\text{SoS}}^c$ per vertex is sufficient.

Since each additional level of intersections is non-trivial, the total number of vertex decay factors is at least k , which is sufficient to handle CD_{SoS}^c per shape in addition to per vertex. ■

6.11.3 [Lemma 6.56](#): truncation error for shapes with too many vertices

The goal of the next two subsections is to upper bound the truncation error.

Lemma 6.56. (Truncation Error)

$$\text{truncation error} \preceq n^{-\Omega(CD_{\text{SOS}})} \pi.$$

In this section we upper bound the first part of the truncation error, where shapes have many extra vertices. The counting is similar to the previous section while we use simpler (and looser) bounds for matrix norms. These will be handled by the fact that large shapes provide lots of vertex decay factors.

We recall

$$\begin{aligned} \text{truncation error}_{\text{too many vertices}} = & - \sum_{\substack{\text{sparse, permissible} \\ \sigma, \tau, \sigma': \\ |V(\sigma \circ \tau \circ \sigma')| > D_V}} \lambda'_{\sigma \circ \tau \circ \sigma'} \frac{M_{\sigma \circ \tau \circ \sigma'}}{|\text{Aut}(\sigma \circ \tau \circ \sigma')|} \\ + \sum_{j=1}^{2D_{\text{SOS}}} (-1)^{j+1} & \sum_{\substack{\text{sparse, permissible} \\ \sigma, \gamma_1, \dots, \gamma_j, \tau, \gamma'_1, \dots, \gamma'_j, \sigma': \\ |V(\sigma \circ \gamma_1 \circ \dots \circ \gamma_j)| > D_V \text{ or} \\ |V(\gamma'_1 \circ \dots \circ \gamma'_j \circ \sigma')| > D_V}} \sum_{\substack{\text{nonequiv.} \\ P \in \mathcal{P}^{\text{mid}} \\ \gamma_j, \dots, \gamma'_j}} N_P(\tau_P) \lambda'_{\sigma \circ \gamma_1 \circ \dots \circ \gamma_j \circ \tau \circ \sigma'} \frac{M_\sigma}{|\text{Aut}(\sigma)|} \frac{M_{\tau_P}}{|\text{Aut}(\tau_P)|} \frac{M_{\sigma'}^\top}{|\text{Aut}(\sigma')|} \end{aligned}$$

To bound this term, we use a tail bound on $c(\alpha)$ plus the main norm bound of this section.

Proposition 6.69.

$$\sum_{\substack{\alpha \in \mathcal{S}: \\ \text{non-trivial,} \\ \text{permissible,} \\ |V(\alpha)| - \frac{|U_\alpha| + |V_\alpha|}{2} \geq k}} \frac{1}{c(\alpha)} \leq \frac{1}{2^k}.$$

Proof. In [Definition 6.52](#) we may introduce a slack of 2 for each power of $|V(\alpha)| - \frac{|U_\alpha| + |V_\alpha|}{2}$. \blacksquare

Lemma 6.70. For all sparse permissible $\sigma, \sigma' \in \mathcal{L}$, $\tau \in \mathcal{M}$, such that $|E_{\text{mid}}(\alpha)| - |V(\alpha)| \leq CD_{\text{SOS}}$ for each of σ, τ, σ' ,

$$\lambda'_{\sigma \circ \tau \circ \sigma'} \frac{\|M_{\sigma \circ \tau \circ \sigma'}\|}{|\text{Aut}(\sigma \circ \tau \circ \sigma')|} \leq \frac{n^{2D_{\text{SOS}}}}{c(\sigma)c(\tau)c(\sigma')}$$

Proof. Let $\alpha = \sigma \circ \tau \circ \sigma'$. As in previous sections, the vertex decay is enough to replace λ'_α by λ_α . Since all shapes obey the edge bound, the factor of $c(\sigma)c(\tau)c(\sigma')$ (which is the same as $c(\alpha)$ up to D_{SOS}^C per vertex) is handled by vertex decay as in previous sections. Unpacking the LHS, the powers of n, d and D_V are:

$$\left(\frac{k}{n}\right)^{|V(\alpha)| - \frac{|U_\alpha| + |V_\alpha|}{2}} \left(\sqrt{\frac{p}{1-p}}\right)^{|E(\alpha)|} \max_S D_V^{|V(S)| - \frac{|L_S| + |R_S|}{2}} \sqrt{n}^{|V(\alpha)| - |V(S)|} \left(\sqrt{\frac{1-p}{p}}\right)^{E(S)}$$

We show this is bounded by $n^{|U_\alpha \cup V_\alpha| + \frac{|U_\alpha| + |V_\alpha|}{2}} \cdot \left(\frac{k\sqrt{d}}{n}\right)^{|V(\alpha)| - \frac{|U_\alpha| + |V_\alpha|}{2}}$ for any separator S . As this is bounded by $n^{2D_{\text{SOS}}}$ times vertex decay, this will complete the proof.

$$= \left(\frac{k\sqrt{d}}{n} \right)^{V(\alpha) - \frac{|U_\alpha| + |V_\alpha|}{2}} \left(\sqrt{\frac{p}{1-p}} \right)^{|E(\alpha)| - |E(S)|} \frac{D_V^{|V(S)| - \frac{|L_S| + |R_S|}{2}}}{\sqrt{d}^{|V(S)| - \frac{|U_\alpha| + |V_\alpha|}{2}}} \left(\frac{1}{\sqrt{p}} \right)^{|V(\alpha)| - |V(S)|}$$

The \sqrt{d} cancels out the factor of D_V , except for a factor of at most $\sqrt{d}^{\frac{|U_\alpha| + |V_\alpha|}{2}}$. This is at most $n^{\frac{|U_\alpha| + |V_\alpha|}{2}}$.

The factors of \sqrt{p} follow a similar strategy as in charging middle shapes. Consider a BFS process from S . For each vertex connected to S , we can charge it to the edge leading to it. For vertices not connected to S , they must be connected to $U_\alpha \cup V_\alpha$ by the connected truncation and we can consider a BFS from $U_\alpha \cup V_\alpha$. This cancels all but $|U_\alpha \cup V_\alpha|$ factors of $1/\sqrt{p}$. This is at most $n^{|U_\alpha \cup V_\alpha|}$. ■

We show how to use [Lemma 6.70](#) to bound truncation error_{too many vertices}. In the first term, for example, at least one of the three parts σ, τ, σ' must have size at least $D_V/3$, and therefore we can upper bound that sum by (let $\alpha = \sigma \circ \tau \circ \sigma'^\tau$ for shorthand)

$$\begin{aligned} & \sum_{\substack{\text{sparse, permissible} \\ \sigma, \tau, \sigma': \\ |V(\sigma)| \geq D_V/3}} \lambda'_\alpha \frac{\|M_\alpha\|}{|\text{Aut}(\alpha)|} + \sum_{\substack{\text{sp., per.} \\ \sigma, \tau, \sigma': \\ |V(\tau)| \geq D_V/3}} \lambda'_\alpha \frac{\|M_\alpha\|}{|\text{Aut}(\alpha)|} + \sum_{\substack{\text{sp., per.} \\ \sigma, \tau, \sigma': \\ |V(\sigma')| \geq D_V/3}} \lambda'_\alpha \frac{\|M_\alpha\|}{|\text{Aut}(\alpha)|} \\ & \leq \sum_{\substack{\text{sp., per.} \\ \sigma, \tau, \sigma': \\ |V(\sigma)| \geq D_V/3}} \frac{n^{2D_{\text{SoS}}}}{c(\sigma)c(\tau)c(\sigma')} + \sum_{\substack{\text{sp., per.} \\ \sigma, \tau, \sigma': \\ |V(\tau)| \geq D_V/3}} \frac{n^{2D_{\text{SoS}}}}{c(\sigma)c(\tau)c(\sigma')} + \sum_{\substack{\text{sp., per.} \\ \sigma, \tau, \sigma': \\ |V(\sigma')| \geq D_V/3}} \frac{n^{2D_{\text{SoS}}}}{c(\sigma)c(\tau)c(\sigma')} \quad (\text{Lemma 6.70}) \\ & \leq \frac{3n^{2D_{\text{SoS}}}}{2^{D_V/3}} \quad (\text{Proposition 6.69}) \\ & = \frac{3n^{2D_{\text{SoS}}}}{2^{\frac{1}{3}CD_{\text{SoS}} \log n}} \leq n^{-\Omega(CD_{\text{SoS}})}. \end{aligned}$$

There is one detail that we have skipped. Some terms have σ or σ' with $|E_{\text{mid}}(\sigma)| - |V(\sigma)| > CD_{\text{SoS}}$. For these terms, the Frobenius norm trick already shows that they are at most $\frac{1}{c(\sigma)c(\tau)c(\sigma')}$.

Lemma 6.71. *For any permissible proper shape α such that $|E(\alpha)| - |V(\alpha)| \geq CD_{\text{SoS}}$,*

$$\lambda'_\alpha \frac{\sqrt{\mathbb{E}[\text{tr}(M_\alpha M_\alpha^\top)]}}{|\text{Aut}(\alpha)|} \leq \frac{1}{c(\alpha)}.$$

Proof.

$$\begin{aligned} \lambda_\alpha \frac{\sqrt{\mathbb{E}[\text{tr}(M_\alpha M_\alpha^\top)]}}{|\text{Aut}(\alpha)|} & \leq \left(\frac{k}{n} \right)^{|V(\alpha)| - \frac{|U_\alpha| + |V_\alpha|}{2}} \sqrt{p}^{|E(\alpha)|} \sqrt{n}^{|V(\alpha)|} \\ & = \left(\frac{k\sqrt{d}}{n} \right)^{|V(\alpha)| - \frac{|U_\alpha| + |V_\alpha|}{2}} \sqrt{d}^{\frac{|U_\alpha| + |V_\alpha|}{2}} \sqrt{p}^{|E(\alpha)| - |V(\alpha)|} \\ & \leq \frac{1}{c(\alpha)} \end{aligned}$$
■

The terms with $j \geq 1$ are bounded by $n^{-\Omega(CD_{SoS})}$ by using [Lemma 6.70](#) once on σ alone and once on σ' alone, using [Lemma 6.55](#) on τ_P , and using the Frobenius norm trick on σ or σ' with too many edges. This completes the proof.

6.11.4 [Lemma 6.56](#): truncation error for shapes with too many edges in one part

Lemma 6.72. $\|truncation\ error_{too\ many\ edges\ in\ one\ part}\| \leq n^{-\Omega(CD_{SoS})}$.

Recall the quantity of our interest for this section,

$$\begin{aligned} truncation\ error_{too\ many\ edges\ in\ one\ part} = & \sum_{\substack{\text{sparse, permissible} \\ \sigma, \tau, \sigma': \\ |V(\sigma \circ \tau \circ \sigma'^T)| \leq D_V, \\ |E_{mid}(\tau)| - |V(\tau)| > CD_{SoS}}} \lambda'_{\sigma \circ \tau \circ \sigma'^T} \frac{M_{\sigma \circ \tau \circ \sigma'^T}}{|\mathbf{Aut}(\sigma \circ \tau \circ \sigma'^T)|} \\ + \sum_{j=1}^{2D_{SoS}} (-1)^j & \sum_{\substack{\text{sparse, permissible} \\ \sigma, \gamma_j, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_j, \sigma': \\ |V(\sigma \circ \gamma_j \circ \dots \circ \gamma_1)| \leq D_V \text{ and } |V(\gamma'_1 \circ \dots \circ \gamma'_j \circ \sigma'^T)| \leq D_V \\ |E_{mid}(\gamma_j)| - |V(\gamma_j)| > CD_{SoS} \text{ or } |E_{mid}(\gamma'_j)| - |V(\gamma'_j)| > CD_{SoS}}} \sum_{\substack{\text{nonequiv.} \\ P \in \mathcal{P}^{mid}_{\gamma_j, \dots, \gamma'_j}}} \lambda'_{\sigma \circ \gamma_j \circ \dots \circ \gamma'_j \circ \sigma'^T} \frac{M_{\sigma \circ \tau_P \circ \sigma'^T}}{|\mathbf{Aut}(\sigma \circ \tau_P \circ \sigma'^T)|} \end{aligned}$$

We will utilize the Frobenius norm trick to bound these norms.

Proposition 6.73. *For all shapes α (including improper shapes and shapes with (quasi-)missing edge indicators),*

$$\mathbb{E}[\text{tr}(M_\alpha M_\alpha^T)] \leq 4^{|E(\alpha)|} |\mathbf{Aut}(\alpha)| \max_{\beta: \text{linearization of } \alpha} \left(\frac{1-p}{p} \right)^{|E(\alpha)| - |E(\beta)| - 2|E_{phantom}|} n^{|V(\alpha)| + |I_\beta|}.$$

Proof. There are at most $2^{|E(\alpha)|}$ linearizations of α . Using [Proposition 6.63](#), the coefficient on each one is at most $\left(\sqrt{\frac{1-p}{p}} \right)^{|E(\alpha)| - |E(\beta)| - 2|E_{phantom}|}$. Applying [Lemma 6.58](#) on each linearization gives the bound. \blacksquare

Lemma 6.74. *If some part γ_i , τ or γ_i^T has CD_{SoS} excess edges then τ_P has at least $(C-2)D_{SoS}$ excess unique edges (i.e. $|E_{unique}(\tau_P)| - |V(\tau_P)| \geq (C-2)D_{SoS}$).*

Proof. Observe that for the part with CD_{SoS} excess edges, we can assign a cycle to each excess edge e so that the cycle contains e and does not contain any other excess edge e' .

Even after intersections, we will still have these cycles. This implies that we can delete the CD_{SoS} excess edges and we will still have that every vertex of τ_P is connected to either U_{τ_P} or V_{τ_P} . In turn, this implies that $|E(\tau_P)| - |V_{\tau_P}| \geq CD_{SoS} - |U_{\tau_P}| - |V_{\tau_P}| \geq (C-2)D_{SoS}$, as needed. \blacksquare

Lemma 6.75. *For all $j \geq 0$ and sparse permissible $\sigma, \gamma_j, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_j, \sigma'$, $|V(\gamma'_1 \circ \dots \circ \gamma'_j \circ \sigma'^T)| \leq D_V$, and $|V(\gamma_1 \circ \dots \circ \gamma_j \circ \sigma^T)| \leq D_V$, while $|E_{mid}(\gamma_j)| - |V(\gamma_j)| > CD_{SoS}$ or $|E_{mid}(\gamma'_j)| - |V(\gamma'_j)| > CD_{SoS}$,*

$$\mathbb{E} \left[\sum_{\text{non-equiv } P} N_P(\tau_P) \lambda'_{\sigma \circ \gamma_j \circ \dots \circ \gamma'_j \circ \sigma'^T} \frac{\|M_{\sigma \circ \tau_P \circ \sigma'^T}\|_F}{|\mathbf{Aut}(\sigma \circ \tau_P \circ \sigma'^T)|} \right] \leq \frac{n^{-\Omega(CD_{SoS})}}{c(\sigma)c(\tau)c(\sigma') \prod_{i=1}^j c(\gamma_i)c(\gamma'_i)}$$

Proof. Let τ_P be any shape that may arise from intersection of $\gamma_j, \dots, \gamma_1, \tau, \gamma'_1, \dots, \gamma'_j$. By Lemma 6.74, we have $|E_{\text{unique}}(\sigma \circ \tau \circ \sigma')| - |V(\sigma \circ \tau \circ \sigma')| \geq CD_{\text{SoS}}$ for some $C > 0$. To apply the Frobenius norm trick on possibly improper shapes, we will again first linearize the shape and then apply Frobenius norm bound,

$$\begin{aligned} & \mathbb{E} \left[\sum_{\text{non-equiv } P} N_P(\tau_P) \lambda'_{\sigma \circ \gamma_j \circ \dots \circ \gamma'_j \circ \sigma' \circ \tau} \frac{\|M_{\sigma \circ \tau_P \circ \sigma' \circ \tau}\|_F}{|\text{Aut}(\sigma \circ \tau_P \circ \sigma' \circ \tau)|} \right] \\ & \leq \sum_{\text{non-equiv } P} N_P(\tau_P) \lambda'_{\sigma \circ \gamma_j \circ \dots \circ \gamma'_j \circ \sigma' \circ \tau} \frac{\mathbb{E}[\|M_{\sigma \circ \tau_P \circ \sigma' \circ \tau}\|_F]}{|\text{Aut}(\sigma \circ \tau_P \circ \sigma' \circ \tau)|} \end{aligned}$$

As in previous sections, the difference between λ' and λ , the count of non-equivalent P , and $N_P(\tau_P)$ are under control from vertex decay, so we drop them. Continuing,

$$\begin{aligned} & \leq \left(\frac{k\sqrt{d}}{n} \right)^{|V(\sigma \circ \gamma_j \circ \dots \circ \gamma'_j \circ \sigma' \circ \tau)| - \frac{|U_\sigma| + |U_{\sigma'}|}{2}} \sqrt{p}^{|E_{\text{unique}}(\sigma \circ \tau_P \circ \sigma' \circ \tau)| - |V(\sigma \circ \tau_P \circ \sigma' \circ \tau)|} \\ & \quad \cdot \sqrt{d}^{\frac{|U_\sigma| + |U_{\sigma'}|}{2}} \sqrt{n}^{|I_\beta| - 2|E_{\text{phantom}}|} \tag{Proposition 6.73} \\ & \leq \left(\frac{k\sqrt{d}}{n} \right)^{|V(\sigma \circ \gamma_j \circ \dots \circ \gamma'_j \circ \sigma' \circ \tau)| - \frac{|U_\sigma| + |U_{\sigma'}|}{2}} \sqrt{p}^{|E_{\text{unique}}(\sigma \circ \tau_P \circ \sigma' \circ \tau)| - |V(\sigma \circ \tau_P \circ \sigma' \circ \tau)|} \sqrt{d}^{D_{\text{SoS}}} \\ & \leq \frac{n^{-\Omega(CD_{\text{SoS}})}}{c(\sigma)c(\tau)c(\sigma') \prod_{i=1}^j c(\gamma_i)c(\gamma'_i)} \end{aligned}$$

where $n^{-\Omega(CD_{\text{SoS}})}$ comes from $\sqrt{p}^{|E_{\text{unique}}(\sigma \circ \tau_P \circ \sigma' \circ \tau)| - |V(\sigma \circ \tau_P \circ \sigma' \circ \tau)|}$. In the second-to-last inequality we observe that any isolated vertex is incident to at least one phantom edge, hence the isolated vertices are under control. We can then restrict our attention to the norm bound factor from the non-isolated vertices in the linearization of τ_P , and the last inequality follows by the decay from excess edges. ■

Now we use Lemma 6.59 to bound all their Frobenius norms together.

$$\mathbb{E}[\|\text{truncation error}_{\text{too many edges in one part}}\|] \leq \sum_{\sigma, \gamma_j, \dots, \gamma'_j, \sigma'} \frac{n^{-\Omega(CD_{\text{SoS}})}}{c(\sigma)c(\tau)c(\sigma') \prod_{i=1}^j c(\gamma_i)c(\gamma'_i)}.$$

By a single application of Markov's inequality, the bound holds.

6.11.5 Lemma 6.57: sum of left shapes is well-conditioned

Lemma 6.57. (*Sum of left shapes is well-conditioned*)

$$\left(\sum_{\substack{\text{sparse,} \\ \text{permissible} \\ \sigma \in \mathcal{L}}} \lambda'_\sigma \frac{M_\sigma}{|\text{Aut}(\sigma)|} \right) \left(\sum_{\substack{\text{sparse,} \\ \text{permissible} \\ \sigma \in \mathcal{L}}} \lambda'_\sigma \frac{M_\sigma}{|\text{Aut}(\sigma)|} \right)^\top \succeq n^{-O(D_{\text{SoS}})} \pi$$

Proof. Observe that

$$\begin{aligned} & \left(\sum_{\substack{\text{sparse,} \\ \text{permissible} \\ \sigma \in \mathcal{L}}} \lambda'_\sigma \frac{M_\sigma}{|Aut(\sigma)|} \right) \left(\sum_{\substack{\text{sparse,} \\ \text{permissible} \\ \sigma' \in \mathcal{L}}} \lambda'_{\sigma'} \frac{M_{\sigma'}}{|Aut(\sigma')|} \right)^\top \\ &= \sum_{j=0}^{D_{\text{SoS}}} \left(\sum_{\substack{\text{sparse,} \\ \text{permissible} \\ \sigma \in \mathcal{L}: |V_\sigma|=j}} \lambda'_\sigma \frac{M_\sigma}{|Aut(\sigma)|} \right) \left(\sum_{\substack{\text{sparse,} \\ \text{permissible} \\ \sigma' \in \mathcal{L}: |V_{\sigma'}|=j}} \lambda'_{\sigma'} \frac{M_{\sigma'}}{|Aut(\sigma')|} \right)^\top \end{aligned}$$

Thus, it is sufficient to show that

$$\sum_{j=0}^{D_{\text{SoS}}} \left(\sum_{\substack{\text{sparse,} \\ \text{permissible} \\ \sigma \in \mathcal{L}: |V_\sigma|=j}} w_j \lambda'_\sigma \frac{M_\sigma}{|Aut(\sigma)|} \right) \left(\sum_{\substack{\text{sparse,} \\ \text{permissible} \\ \sigma' \in \mathcal{L}: |V_{\sigma'}|=j}} w_j \lambda'_{\sigma'} \frac{M_{\sigma'}}{|Aut(\sigma')|} \right)^\top \succeq n^{-O(D_{\text{SoS}})} \pi$$

for some weights $w_0, \dots, w_{D_{\text{SoS}}} \in [0, 1]$. To show this, we will show that if we choose the weights $w_0, \dots, w_{D_{\text{SoS}}} \in [0, 1]$ carefully then all of the terms where σ is non-trivial or σ' is non-trivial can be charged to π .

Definition 6.76. Define π_j to be the matrix such that $M_{AB} = 1$ if $A = B$ and A is an independent set of size j in G and $M_{AB} = 0$ otherwise.

Proposition 6.77. $\pi = \sum_{j=0}^{D_{\text{SoS}}} \pi_j$

Definition 6.78. Define $\lambda'_{\pi_j} \approx \left(\frac{k}{n}\right)^{\frac{j}{2}}$ to be the coefficient λ'_σ of the shape σ where $V(\sigma) = U_\sigma = V_\sigma$, $|U_\sigma| = j$, and there is a missing edge indicator for every possible edge in $V(\sigma) = U_\sigma = V_\sigma$.

We will choose our weights so that the following condition is satisfied. For all C' -sparse σ and σ' such that σ is non-trivial or σ' is non-trivial, letting $i = |U_\sigma|$ and letting $i' = |U_{\sigma'}|$,

$$-w_j \lambda'_\sigma \lambda'_{\sigma'} \frac{M_\sigma M_{\sigma'}^\top + M_{\sigma'} M_\sigma^\top}{|Aut(\sigma)| \cdot |Aut(\sigma')|} \preceq \frac{1}{c(\sigma)c(\sigma')D_{\text{SoS}}} \left(w_i \lambda'^2_{\pi_i} \pi_i + w_{i'} \lambda'^2_{\pi_{i'}} \pi_{i'} \right)$$

To see why this condition is sufficient, observe that since $\sum_\alpha \frac{1}{c(\alpha)} \leq \frac{1}{4}$,

1. For all non-trivial σ such that $|V_\sigma| = i$ and $|V_{\sigma'}| = j$,

$$\sum_{\sigma': |V_{\sigma'}|=j} \frac{1}{c(\sigma)c(\sigma')D_{\text{SoS}}} w_i \lambda'^2_{\pi_i} \pi_i \preceq \frac{2}{c(\sigma)D_{\text{SoS}}} w_i \lambda'^2_{\pi_i} \pi_i$$

Summing this over all non-trivial σ such that $|V_\sigma| = i$ and $|V_{\sigma'}| = j$, we have that

$$\sum_{\sigma: |U_\sigma|=i, |V_\sigma|=j, \sigma \text{ is non-trivial}} \sum_{\sigma': |V_{\sigma'}|=j} \frac{1}{c(\sigma)c(\sigma')D_{\text{SoS}}} w_i \lambda'^2_{\pi_i} \pi_i \preceq \frac{1}{2D_{\text{SoS}}} w_i \lambda'^2_{\pi_i} \pi_i$$

2. When σ is trivial, $\sum_{\sigma':|V_{\sigma'}|=j,\sigma' \text{ is non-trivial}}$ $\frac{1}{c(\sigma')D_{\text{SoS}}}w_j\lambda_{\pi_j}'^2\pi_j \preceq \frac{1}{4D_{\text{SoS}}}w_j\lambda_{\pi_j}'^2\pi_j$

Thus, we have that

$$\sum_{j=0}^{D_{\text{SoS}}} \sum_{\sigma,\sigma':|V_{\sigma}|=|V_{\sigma'}|=j,\sigma \text{ or } \sigma' \text{ is non-trivial}} \frac{1}{c(\sigma)c(\sigma')D_{\text{SoS}}}w_{|U_{\sigma}|}\lambda_{\pi_{|U_{\sigma}|}}'^2\pi_{|U_{\sigma}|} \preceq \sum_{i=0}^{D_{\text{SoS}}} \frac{D_{\text{SoS}}+1}{2D_{\text{SoS}}}w_i\lambda_{\pi_i}'^2\pi_i$$

The $w_{i'}\lambda_{\pi_{i'}}'^2\pi_{i'}$ terms can be bounded in the same way. This implies that

$$\sum_{j=0}^{D_{\text{SoS}}} \left(\sum_{\substack{\text{sparse,} \\ \text{permissible} \\ \sigma \in \mathcal{L}:|V_{\sigma}|=j}} \lambda_{\sigma}' \frac{M_{\sigma}}{|Aut(\sigma)|} \right) \left(\sum_{\substack{\text{sparse,} \\ \text{permissible} \\ \sigma' \in \mathcal{L}:|V_{\sigma'}|=j}} \lambda_{\sigma'}' \frac{M_{\sigma'}}{|Aut(\sigma')|} \right)^{\top} \succeq \sum_{i=0}^{D_{\text{SoS}}} \frac{1}{4}w_i\lambda_{\pi_i}'^2\pi_i \succeq \frac{1}{4} \left(\min_{i \in [0, D_{\text{SoS}}]} \{w_i\lambda_{\pi_i}'^2\} \right) \pi$$

We now choose the weights so that for all σ and σ' such that σ is non-trivial or σ' is non-trivial, letting $i = |U_{\sigma}|$ and letting $i' = |U_{\sigma'}|$,

$$-w_j\lambda_{\sigma}'\lambda_{\sigma'}' \frac{M_{\sigma}M_{\sigma'}^{\top} + M_{\sigma'}M_{\sigma}^{\top}}{|Aut(\sigma)| \cdot |Aut(\sigma')|} \preceq \frac{1}{c(\sigma)c(\sigma')D_{\text{SoS}}} \left(w_i\lambda_{\pi_i}'^2\pi_i + w_{i'}\lambda_{\pi_{i'}}'^2\pi_{i'} \right)$$

Observe that for all $a, b > 0$

$$\left(a \frac{M_{\sigma}}{|Aut(\sigma)|} - b \frac{M_{\sigma'}}{|Aut(\sigma')|} \right) \left(a \frac{M_{\sigma}^{\top}}{|Aut(\sigma)|} - b \frac{M_{\sigma'}^{\top}}{|Aut(\sigma')|} \right) \succeq 0$$

and

$$a^2 \frac{M_{\sigma}M_{\sigma}^{\top}}{|Aut(\sigma)|^2} + b^2 \frac{M_{\sigma'}M_{\sigma'}^{\top}}{|Aut(\sigma')|^2} \preceq a^2 \left\| \frac{M_{\sigma}}{|Aut(\sigma)|} \right\|^2 \pi_i + b^2 \left\| \frac{M_{\sigma'}}{|Aut(\sigma')|} \right\|^2 \pi_{i'}$$

Choosing $a = \frac{\sqrt{\frac{w_j}{c(\sigma)c(\sigma')D_{\text{SoS}}}}\lambda_{\pi_i}'|Aut(\sigma)|}{\|M_{\sigma}\|}$ and $b = \frac{\sqrt{\frac{w_{i'}}{c(\sigma)c(\sigma')D_{\text{SoS}}}}\lambda_{\pi_{i'}}'|Aut(\sigma')|}{\|M_{\sigma'}\|}$, we require that

$$ab = \frac{\sqrt{w_i w_{i'}} \lambda_{\pi_i}' \lambda_{\pi_{i'}}' |Aut(\sigma)| \cdot |Aut(\sigma')|}{c(\sigma)c(\sigma')D_{\text{SoS}} \|M_{\sigma}\| \|M_{\sigma'}\|} \geq w_j \lambda_{\sigma}' \lambda_{\sigma'}'$$

Rearranging, we need to choose the weights so that

$$w_j \leq \frac{\sqrt{w_i w_{i'}} \lambda_{\pi_i}' \lambda_{\pi_{i'}}' |Aut(\sigma)| \cdot |Aut(\sigma')|}{c(\sigma)c(\sigma')D_{\text{SoS}} \lambda_{\sigma}' \lambda_{\sigma'}' \|M_{\sigma}\| \|M_{\sigma'}\|}$$

We have that $\frac{\lambda_{\sigma}' \|M_{\sigma}\|}{\lambda_{\pi_i}' |Aut(\sigma)|}$ is $\tilde{O} \left(\frac{k^3}{n} \right)^{\frac{|U_{\sigma}|-|V_{\sigma}|}{2}}$. To see this, we make the following observations:

1. $\lambda_{\sigma}' \approx \left(\frac{k}{n} \right)^{|V(\sigma)| - \frac{|V_{\sigma}|}{2}} \left(\sqrt{\frac{p}{1-p}} \right)^{|E(\sigma)|}$
2. $\lambda_{\pi_i}' \approx \left(\frac{k}{n} \right)^{\frac{|U_{\sigma}|}{2}}$
3. $\frac{\|M_{\sigma}\|}{|Aut(\sigma)|}$ is $\tilde{O} \left(n^{\frac{|V(\sigma)|S|}{2}} \left(\sqrt{\frac{1-p}{p}} \right)^{|E(S)|} \right)$ where S is the sparse minimum vertex separator of σ .

4. $|E(\sigma) \setminus E(S)| \geq |V(\sigma)| - |S| + |V_\sigma| - |U_\sigma|$ as there are at most $|U_\sigma| - |V_\sigma|$ connected components of σ which do not contain a vertex in S .
5. Since σ is a left shape and S is a vertex separator, $|S| \geq |V_\sigma|$.
6. Since $\sqrt{\frac{np}{1-p}} > 1$ and $\frac{k}{n} \sqrt{\frac{np}{1-p}} \leq 1$,

$$\begin{aligned}
& \left(\frac{k}{n}\right)^{|V(\sigma)| - \frac{|U_\sigma| + |V_\sigma|}{2}} n^{\frac{|V(\sigma) \setminus S|}{2}} \left(\sqrt{\frac{p}{1-p}}\right)^{|V(\sigma)| - |S| + |V_\sigma| - |U_\sigma|} \\
& \leq \left(\sqrt{\frac{np}{1-p}}\right)^{|V_\sigma| - |S|} \left(\frac{k}{n} \sqrt{\frac{np}{1-p}}\right)^{|V(\sigma)| - |U_\sigma|} \left(\frac{k}{n}\right)^{\frac{|U_\sigma| - |V_\sigma|}{2}} n^{\frac{|U_\sigma| - |V_\sigma|}{2}} \\
& \leq k^{\frac{|U_\sigma| - |V_\sigma|}{2}}
\end{aligned}$$

Thus, if we take $w_j = \tilde{O}(k^{D_{\text{SoS}} - j})$ then this equation will be satisfied for all C' -sparse σ and σ' . Note that as long as $E_{\text{mid}}(\sigma) \leq V(\sigma) + CD_{\text{SoS}}$, the factors of $\frac{k}{n} \sqrt{\frac{np}{1-p}}$ are enough to handle $c(\sigma)$. By analyzing their Frobenius norm, it can be shown that with high probability, the terms where $E_{\text{mid}}(\sigma) > V(\sigma) + CD_{\text{SoS}}$ or $E_{\text{mid}}(\sigma') > V(\sigma') + CD_{\text{SoS}}$ have norm $n^{-\Omega(CD_{\text{SoS}})}$ and are thus negligible. ■

7 Open Problems

Several other problems on sparse graphs are conjectured to be hard for SoS and it is our hope that the techniques here can help prove that these problems are hard for SoS. These problems include MaxCut, k -Coloring, and Densest- k -Subgraph. For MaxCut in particular, since there are no constraints other than booleanity of the variables it may be possible to truncate away dense shapes, which we could not do here due to the presence of independent set indicator functions.

Another direction for further research is to handle random graphs which are not Erdős-Rényi. Since the techniques here depend on graph matrix norms, one would hope that they generalize to distributions such as d -regular graphs for which low-degree polynomials are still concentrated. However, in the non-iid setting, it is not clear what the analogue of graph matrices should be used due to the lack of a Fourier basis that is friendly to work with.

The polynomial constraint " $\sum_{v \in V} x_v = k$ " is not satisfied exactly by our pseudoexpectation operator. It's possible that techniques from [Pan21] can be used to fix this.

The parameters in this paper can likely be improved. One direction is to remove the final factor of $\log n$ from our bound. This would allow us to prove an SoS lower bound for the "ultra-sparse regime" $d = O(1)$ rather than $d \geq \log^2 n$. This setting is interesting as there is a nontrivial algorithm that finds an independent set of half optimal size [GS17, RV17]. Furthermore, this algorithm is local in a sense that we don't define here. It would be extremely interesting if this algorithm could be converted into a rounding algorithm for constant-degree SoS.

Another direction is to improve the dependence on D_{SoS} . While our bound has a $\frac{1}{\text{poly}(D_{\text{SoS}})}$ dependence on D_{SoS} , we conjecture that the dependence should actually be $(1-p)^{O(D_{\text{SoS}})}$. If so, this would provide strong evidence for the prevailing wisdom in parameterized complexity and

proof complexity that a maximum independent set of size k requires $n^{\Omega(k)}$ time to find/certify (corresponding to SoS degree $\Omega(k)$).

References

- [ABB19] Enric Boix Adserà, Matthew Brennan, and Guy Bresler. The average-case complexity of counting cliques in Erdős-Rényi hypergraphs. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science*, pages 1256–1280. IEEE Comput. Soc. Press, Los Alamitos, CA, [2019] ©2019. 1
- [AMP20] Kwangjun Ahn, Dhruv Medarametla, and Aaron Potechin. Graph matrices: Norm bounds and applications. abs/1604.03423, 2020. URL: <https://arxiv.org/abs/1604.03423>, arXiv:1604.03423. 4, 10, 18, 47, 62
- [BBK⁺20] Afonso S. Bandeira, Jess Banks, Dmitriy Kunisky, Cristopher Moore, and Alexander S. Wein. Spectral planting and hardness of refuting cuts, colorability, and communities in random graphs. *arXiv preprint arXiv:2008.12237*, 2020. 4
- [BHK⁺16] B. Barak, S. B. Hopkins, J. Kelner, P. Kothari, A. Moitra, and A. Potechin. A nearly tight sum-of-squares lower bound for the planted clique problem. In *Proceedings of the 57th IEEE Symposium on Foundations of Computer Science*, pages 428–437, 2016. 1, 3, 4, 5, 8, 9, 10, 11, 12, 18, 19, 23, 25, 26
- [BLM15] Charles Bordenave, Marc Lelarge, and Laurent Massoulié. Non-backtracking spectrum of random graphs: Community detection and non-regular ramanujan graphs. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 1347–1357, 2015. doi:10.1109/FOCS.2015.86. 11
- [Bor19] Charles Bordenave. A new proof of Friedman’s second eigenvalue theorem and its extension to random lifts. Technical Report 1502.04482v4, arXiv, 2019. To appear in *Annales scientifiques de l’École normale supérieure*. 11
- [BRS11] Boaz Barak, Prasad Raghavendra, and David Steurer. Rounding semidefinite programming hierarchies via global correlation. In *Proceedings of the 52nd IEEE Symposium on Foundations of Computer Science*, pages 472–481, 2011. 1
- [BS16] Boaz Barak and David Steurer. Proofs, beliefs, and algorithms through the lens of sum-of-squares. *Course notes: http://www.sumofsquares.org/public/index.html*, 2016. 1
- [CG02] Fan Chung and Ronald Graham. Sparse quasi-random graphs. volume 22, pages 217–244. 2002. Special issue: Paul Erdős and his mathematics. doi:10.1007/s004930200010. 13
- [CM18] Eden Chlamtác and Pasin Manurangsi. Sherali-adams integrality gaps matching the log-density threshold. *CoRR*, abs/1804.07842, 2018. URL: <http://arxiv.org/abs/1804.07842>, arXiv:1804.07842. 9
- [CO05] Amin Coja-Oghlan. The Lovász number of random graphs. *Combinatorics, Probability and Computing*, 14(4):439–465, 2005. 2

- [COE15] Amin Coja-Oghlan and Charilaos Efthymiou. On independent sets in random graphs. *Random Structures & Algorithms*, 47(3):436–486, 2015. 2, 13
- [DM11] Varsha Dani and Cristopher Moore. Independent sets in random graphs from the weighted second moment method. In *Approximation, randomization, and combinatorial optimization*, volume 6845 of *Lecture Notes in Comput. Sci.*, pages 472–482. Springer, Heidelberg, 2011. URL: https://doi.org/10.1007/978-3-642-22935-0_40, doi: 10.1007/978-3-642-22935-0_40. 2
- [DMO⁺19] Yash Deshpande, Andrea Montanari, Ryan O’Donnell, Tselil Schramm, and Subhabrata Sen. The threshold for sdp-refutation of random regular nae-3sat. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2305–2321. SIAM, 2019. 11
- [DSS16] Jian Ding, Allan Sly, and Nike Sun. Maximum independent sets on random regular graphs. *Acta Math.*, 217(2):263–340, 2016. doi:10.1007/s11511-017-0145-9. 2
- [FKP19] Noah Fleming, Pravesh Kothari, and Toniann Pitassi. Semialgebraic proofs and efficient algorithm design. *Foundations and Trends® in Theoretical Computer Science*, 14(1-2):1–221, 2019. URL: <http://dx.doi.org/10.1561/04000000086>, doi:10.1561/04000000086. 1
- [FM17] Zhou Fan and Andrea Montanari. How well do local algorithms solve semidefinite programs? In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 604–614. ACM, 2017. 11
- [FV18] S. Friedli and Y. Velenik. *Statistical mechanics of lattice systems*. Cambridge University Press, Cambridge, 2018. 9
- [GJJ⁺20] Mrinalkanti Ghosh, Fernando Granha Jeronimo, Chris Jones, Aaron Potechin, and Goutham Rajendran. Sum-of-squares lower bounds for Sherrington-Kirkpatrick via planted affine planes. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science—FOCS 2020*, pages 954–965. IEEE Computer Soc., Los Alamitos, CA, [2020] ©2020. 9
- [GM75] G. R. Grimmett and C. J. H. McDiarmid. On colouring random graphs. *Math. Proc. Cambridge Philos. Soc.*, 77:313–324, 1975. doi:10.1017/S0305004100051124. 4
- [GS11] Venkatesan Guruswami and Ali Kemal Sinop. Lasserre hierarchy, higher eigenvalues, and approximation schemes for graph partitioning and quadratic integer programming with psd objectives. In *FOCS*, pages 482–491, 2011. 1
- [GS17] David Gamarnik and Madhu Sudan. Limits of local algorithms over sparse random graphs. *Ann. Probab.*, 45(4):2353–2376, 2017. doi:10.1214/16-AOP1114. 57
- [HKP⁺17] Samuel B Hopkins, Pravesh K Kothari, Aaron Potechin, Prasad Raghavendra, Tselil Schramm, and David Steurer. The power of sum-of-squares for detecting hidden structures. In *Proceedings of the 58th IEEE Symposium on Foundations of Computer Science*, pages 720–731. IEEE, 2017. 4, 9, 13
- [Hop18] Samuel Brink Klevit Hopkins. Statistical inference and the sum of squares method. 2018. 9, 13

- [HSS15] Samuel B Hopkins, Jonathan Shi, and David Steurer. Tensor principal component analysis via sum-of-squares proofs. In *Conference on Learning Theory*, pages 956–1006, 2015. 1
- [KMOW17] Pravesh Kothari, Ryuhei Mori, Ryan O’Donnell, and David Witmer. Sum of squares lower bounds for refuting any CSP. In *Proceedings of the 49th ACM Symposium on Theory of Computing*, 2017. 1, 4, 9
- [Kun20] Dmitriy Kunisky. Positivity-preserving extensions of sum-of-squares pseudomoments over the hypercube. *arXiv preprint arXiv:2009.07269*, 2020. 4
- [MP16] Dhruv Medarametla and Aaron Potechin. Bounds on the norms of uniform low degree graph matrices. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2016)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2016. 10
- [MRX20] Sidhanth Mohanty, Prasad Raghavendra, and Jeff Xu. Lifting sum-of-squares lower bounds: degree-2 to degree-4. In *Proceedings of the 52nd ACM Symposium on Theory of Computing*, pages 840–853, 2020. 9
- [MSS16] Tengyu Ma, Jonathan Shi, and David Steurer. Polynomial-time tensor decompositions with sum-of-squares. In *Proceedings of the 57th IEEE Symposium on Foundations of Computer Science*, pages 438–446. IEEE, 2016. 1
- [O’D14] Ryan O’Donnell. *Analysis of Boolean functions*. Cambridge University Press, New York, 2014. doi:10.1017/CB09781139814782. 5
- [Pan21] Shuo Pang. SOS lower bound for exact planted clique. In *36th Computational Complexity Conference*, volume 200 of *LIPICs. Leibniz Int. Proc. Inform.*, pages Art. 26, 63. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2021. 4, 57
- [PR20] Aaron Potechin and Goutham Rajendran. Machinery for proving sum-of-squares lower bounds on certification problems. abs/2011.04253, 2020. URL: <https://arxiv.org/abs/2011.04253>, arXiv:2011.04253. 1, 4, 5, 9, 18, 23
- [RT20] Goutham Rajendran and Madhur Tulsiani. Nonlinear concentration via matrix efronstein. Manuscript, 2020. 4
- [RV17] Mustazee Rahman and Bálint Virág. Local algorithms for independent sets are half-optimal. *Ann. Probab.*, 45(3):1543–1577, 2017. doi:10.1214/16-AOP1094. 4, 57
- [RW17] Prasad Raghavendra and Benjamin Weitz. On the bit complexity of sum-of-squares proofs. In *Proceedings of the 44th International Colloquium on Automata, Languages and Programming*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017. 5
- [Wei20] Alexander S Wein. Optimal low-degree hardness of maximum independent set. *arXiv preprint arXiv:2010.06563*, 2020. 4
- [Wor95] Nicholas C. Wormald. Differential equations for random processes and random graphs. *Ann. Appl. Probab.*, 5(4):1217–1235, 1995. URL: [http://links.jstor.org/sici?sici=1050-5164\(199511\)5:4<1217:DEFRPA>2.0.CO;2-A&origin=MSN](http://links.jstor.org/sici?sici=1050-5164(199511)5:4<1217:DEFRPA>2.0.CO;2-A&origin=MSN). 4

A Trace Power Calculations

Let Ω be the distribution which is $\sqrt{\frac{1-p}{p}}$ with probability p and $-\sqrt{\frac{p}{1-p}}$ with probability $1-p$.

Proposition A.1. For all $k \in \mathbb{N}$ and all $p \leq \frac{1}{2}$, $|\mathbb{E}_\Omega[x^k]| \leq \left(\sqrt{\frac{1-p}{p}}\right)^{k-2}$ and $\left|\mathbb{E}_\Omega \left[x^k \mathbf{1}_{x=-\sqrt{\frac{p}{1-p}}} \right]\right| \leq \left(\sqrt{\frac{1-p}{p}}\right)^{k-2}$.

Proof. For the first part, observe that for $k = 1$, $E_\Omega[x^k] = 0$ and for $k \geq 2$,

$$E_\Omega[x^k] = p \left(\sqrt{\frac{1-p}{p}} \right)^k + (1-p) \left(-\sqrt{\frac{p}{1-p}} \right)^k = (1-p) \left(\sqrt{\frac{1-p}{p}} \right)^{k-2} + p \left(-\sqrt{\frac{p}{1-p}} \right)^{k-2}$$

and since $\sqrt{\frac{p}{1-p}} \leq \sqrt{\frac{1-p}{p}}$ and $1-p \geq p$, $0 \leq (1-p) \left(\sqrt{\frac{1-p}{p}} \right)^{k-2} + p \left(-\sqrt{\frac{p}{1-p}} \right)^{k-2} \leq \left(\sqrt{\frac{1-p}{p}} \right)^{k-2}$

For the second part, observe that for all $k \geq 1$

$$\left| \mathbb{E}_\Omega \left[x^k \mathbf{1}_{x=-\sqrt{\frac{p}{1-p}}} \right] \right| = (1-p) \left(\sqrt{\frac{p}{1-p}} \right)^k \leq \left(\sqrt{\frac{1-p}{p}} \right)^{k-2}$$

■

Before we embark upon the trace method calculation, it is necessary to note that throughout our PSD analysis, our graphical matrix for shape has rows/columns indexed by sets (as it is a sum of ribbons which have rows/columns indexed by set); however, for the trace method calculation, it is easier to work with graph matrices indexed by ordered tuples and then identify the change in trace method encoding that allows us to shift to set-indexed matrices.

Definition A.2 (Ordered-ribbon). Given a ribbon R , let \mathcal{O} be an ordering for vertices in $U_R \cup V_R$, we can define graphical matrix for an ordered-ribbon to have rows and columns indexed by subset of $[n]$ with a single non-zero entry

$$M_{R,\mathcal{O}}^{(\text{ordered})}[(I, \mathcal{O}_I), (J, \mathcal{O}_J)] = \chi_{E(R)}(G)$$

for $(I, \mathcal{O}_I) = (A_R, \mathcal{O}_{A_R}), (J, \mathcal{O}_J) = (B_R, \mathcal{O}_{B_R})$.

It is immediate from the above definition that for $S, T \subseteq [n]$, ribbon can be obtained from ordered-ribbons by summing over ordering of the labeled vertices in $U_R \cup V_R$,

$$M_R[S, T] = \sum_{\mathcal{O}: \text{ordering for vertices in } U_R \cup V_R} M_{R,\mathcal{O}}^{(\text{ordered})}[(S, \mathcal{O}_S), (T, \mathcal{O}_T)]$$

Definition A.3 (Graphical matrix with (unsymmetrized) ordered tuples). We analogously define graphical matrix indexed by ordered tuples by summing over ribbons with ordered tuples, for a shape α ,

$$M_\alpha^{(\text{ordered})} = \sum_{\substack{R: \text{ribbon of shape } \alpha, \\ \mathcal{O}: \text{ordering}}} M_{(R,\mathcal{O})}^{(\text{ordered})}$$

Definition A.4 (Graphical matrix with sym-ordered tuples).

$$M_\alpha^{(\text{sym-ordered})}[(I, \mathcal{O}_I), (J, \mathcal{O}_J)] = \sum_{\mathcal{O}: \text{ordering for } U_\alpha \cup V_\alpha} M_{R, \mathcal{O}}^{(\text{ordered})}$$

Remark A.5. The graphical matrix for shape indexed by sets is a principle submatrix of the the graphical matrix with sym-ordered tuples.

We now proceed onto the analysis for graph matrices indexed by (unsymmetrized) ordered tuples.

Definition A.6. Given a shape α and a separator S between U_α and V_α , we make the following definitions:

1. We define $E(S)$ to be the set of edges with both endpoints in S .
2. We define L_S to be the set of vertices in S which are reachable from U_α without passing through any other vertices in S .
3. We define R_S to be the set of vertices in S which are reachable from V_α without passing through any other vertices in S .
4. Define $c(S)$ to be the number of connected components of $\alpha \setminus S$ which are not reachable from U_α or V_α .

Theorem A.7. For any proper shape α (including shapes with missing edge indicators) and any $q \in \mathbb{N}$, letting M_α be the graph matrix where the rows and columns are indexed by ordered tuples rather than sets,

$$\begin{aligned} \mathbb{E} [\text{tr} ((M_\alpha M_\alpha^T)^q)] &\leq 4^{2q \left(|V(\alpha)| - \frac{|U_\alpha| + |V(\alpha)|}{2} \right)} \left(\sqrt{|V(\alpha)|} \right)^{2q |V(\alpha) \setminus (U_\alpha \cup V_\alpha)|} n^{q |V(\alpha)|} \\ &\left(\max_{\text{separator } S} n^{-\frac{|S|}{2}} \left(3 \sqrt{\frac{1-p}{p}} \right)^{|E(S)|} \left(\sqrt{|V(\alpha)|} \right)^{|S| - |U_\alpha \cap V_\alpha|} (2q)^{|S| - \frac{|L_S| + |R_S|}{2} + \frac{c(\alpha, S)}{2}} \right)^{2q-2} \end{aligned}$$

Proof. We use the idea from Appendix B of [AMP20]. Given a constraint graph on $V(\alpha, 2q)$, for each vertex in $V(\alpha, 2q)$ which is not a copy of a vertex in $U_\alpha \cap V_\alpha$, we note whether it has a constraint edge to the left and whether it has a constraint edge to the right. Note that there are at most $4^{2q \left(|V(\alpha)| - \frac{|U_\alpha| + |V(\alpha)|}{2} \right)}$ possibilities for this (note that vertices in $(U_\alpha \cup V_\alpha) \setminus (U_\alpha \cap V_\alpha)$ only have q copies in $V(\alpha, 2q)$). For each copy of α , we then construct a separator S as follows:

1. We include any vertex which has both a constraint edge to the left and a constraint edge to the right in S .
2. We include any vertex in U_α which has a constraint edge to the right in S .
3. We include any vertex in V_α which has a constraint edge to the left in S .
4. We include all vertices in $U_\alpha \cap V_\alpha$.

We choose separators for the copies of α^T in the same way. The theorem follows from the following lemma.

Lemma A.8. *Given any separators $S_3, S_5, \dots, S_{2q-1}$ for the copies of α (except the first one) and any separators $S_2, S_4, \dots, S_{2q-2}$ for the copies of α^\top , there are at most*

$$\left(\sqrt{|V(\alpha)|}\right)^{2q|V(\alpha)\setminus(U_\alpha \cup V_\alpha)|} n^{q|V(\alpha)|} \left(\prod_{j=2}^{2q-1} n^{-\frac{|S_j|}{2}} 3^{|E(S_j)|}\right) \left(\sqrt{|V(\alpha)|}\right)^{|S_j|-|U_\alpha \cap V_\alpha|} (2q)^{|S_j|-\frac{|L_{S_j}|+|R_{S_j}|}{2}+\frac{c(S_j)}{2}}$$

terms in $\mathbb{E}[\text{tr}((M_\alpha M_\alpha^\top)^q)]$ which have these separators and have nonzero expected value and for each of these terms the expected value is at most $\prod_{j=2}^{2q-1} \left(\sqrt{\frac{1-p}{p}}\right)^{|E(S_j)|}$.

Proof. To see that the expected value is at most $\prod_{j=2}^{2q-1} \left(\sqrt{\frac{1-p}{p}}\right)^{|E(S_j)|}$, observe that by Proposition A.1, the expected value of a term is at most

$$\prod_e \left(\sqrt{\frac{1-p}{p}}\right)^{(\text{multiplicity of } e)-2}$$

and this holds even if α contains missing edge indicators. We can view this as giving each edge which is making a middle appearance (i.e. it has appeared before and will appear again) a factor of $\left(\sqrt{\frac{1-p}{p}}\right)$ and giving edges which are appearing for the first or last time a factor of 1 (if an edge appears exactly once, it would instead get a factor of $\sqrt{\frac{p}{1-p}}$, but this only helps us). Since edges must be appearing for the first or last time unless both endpoints are in the separator, the result follows.

To count the number of terms, for each copy of α , we can go through the copies of α and α^\top and identify their vertices one by one. We give each vertex in $U_\alpha \cap V_\alpha$ a label in $[n]$ and we only need to do this once. For the other vertices:

1. If the vertex v has no constraint edge to the left and is thus appearing for the first time, we give v a label in $[n]$.
2. If the vertex v is adjacent to a vertex u which has been identified and has no constraint edge to the right, since u is appearing for the last time, there are $\deg(u) \leq |V(\alpha)|$ edges incident to u which have not appeared for the last time before this copy of α or α^\top . To identify v , it is sufficient to identify which of these edges is (u, v) .
3. Otherwise, we can pay a factor of $2q|V(\alpha)|$ to identify the vertex as it has appeared before.

For each edge in $E(S)$, we note whether the edge is appearing for the first time, appearing for the last time, or is making a middle appearance. For edges which are not in $E(S)$, one of its endpoints will not have a constraint edge to the left or a constraint edge to the right and this determines whether this is the first or last time the edge appears.

We now consider the factors required by each vertex in the $\frac{j+1}{2}$ -th copy of α (which is not a copy of a vertex in $U_\alpha \cap V_\alpha$) for $j = 3, 5, \dots, 2q-1$. Similar reasoning applies to copies of α^\top .

1. Vertices in U_α are already identified.
2. For vertices which are not in U_α but are reachable from U_α without passing through S_j (including vertices in L_{S_j}), when we want to identify these vertices, they will have a neighbor which is already identified and is appearing for the last time. Thus, each of these vertices requires a factor of at most $|V(\alpha)|$.

3. For vertices which are reachable from V_α without passing through S_j (and which are not in S_j themselves), these vertices have no constraint edge to the left, so this is the first time these vertices appear. Thus, each of these vertices requires a factor of n .
4. For vertices in $S_j \setminus L_{S_j}$, these vertices have appeared before but are not adjacent to a vertex which has been identified and has no constraint edge to the right, so we pay a factor of $2q|V(\alpha)|$ for these vertices
5. For a component of $V(\alpha) \setminus S_j$ which is not reachable from U_α or V_α , all of the vertices in this component will either have no constraint edges to the left or no constraint edges to the right. If the vertices in this component have no constraint edges to the left, they are appearing for the first time and require a factor of n . If the vertices in this component have no constraint edges to the right, we can pay a factor of $2q|V(\alpha)|$ for the first vertex in this component and a factor of $|V(\alpha)|$ for all of the other vertices in this component.

We can also try identifying vertices by going from the right to the left. Following similar logic, this gives us the following factors.

1. Vertices in V_α are already identified.
2. For vertices which are not in V_α but are reachable from V_α without passing through S_j (including vertices in R_{S_j}), we need a factor of at most $|V(\alpha)|$.
3. For vertices which are reachable from U_α without passing through S_j (and which are not in S_j themselves), we need a factor of n .
4. For vertices in $S_j \setminus R_{S_j}$, we need a factor of $2q|V(\alpha)|$.
5. For a component of $V(\alpha) \setminus S_j$ which is not reachable from U_α or V_α , all of the vertices in this component will either have no constraint edges to the left or no constraint edges to the right. If the vertices in this component have no constraint edges to the right, they are appearing for the first time (when going from the right to the left) and require a factor of n . If the vertices in this component have no constraint edges to the left, we can pay a factor of $2q|V(\alpha)|$ for the first vertex in this component and a factor of $|V(\alpha)|$ for all of the other vertices in this component.

We can choose whichever direction gives a smaller bounds. Since we don't know which direction this will be, we take the square root of the product of the two bounds. We now consider how many factors of n , $|V(\alpha)|$, and $2q$ this gives us. For S_2, \dots, S_{2q-1} ,

1. Factors of n : Each vertex in S_j does not give us a factor of n . Each vertex outside of S_j gives a factor of n for exactly one of the two directions.
2. Factors of $|V(\alpha)|$: Vertices which are copies of vertices in $U_\alpha \cap V_\alpha$ do not need any factors of $|V(\alpha)|$. For the other vertices, we have the following cases. For each vertex $v \in S_j$,
 - (a) If $v \notin U_\alpha \cup V_\alpha$ then v requires a factor of $|V(\alpha)|$ in both directions.
 - (b) If $v \in U_\alpha \setminus V_\alpha$ or $v \in V_\alpha \setminus U_\alpha$ then v requires a factor of $|V(\alpha)|$ in one of the two directions.

For each vertex $v \notin S_j$,

- (a) If $v \notin U_\alpha \cup V_\alpha$ then v requires a factor of $|V(\alpha)|$ in one of the two directions.
 - (b) If $v \in U_\alpha \setminus V_\alpha$ or $v \in V_\alpha \setminus U_\alpha$ then v does not require a factor of $|V(\alpha)|$ as in one direction it will already be specified and in the other direction it will be a new vertex and will thus require a factor of n rather than $|V(\alpha)|$.
3. Factors of $2q$: Vertices in $S_j \setminus (L_{S_j} \cup R_{S_j})$ require a factor of q in both directions. Vertices in $L_{S_j} \setminus R_{S_j}$ and $L_{S_j} \setminus R_{S_j}$ require a factor of q for exactly one of the two directions. Vertices in $L_{S_j} \cap R_{S_j}$ do not require any factors of q . Each component of $V(\alpha) \setminus S_j$ which is not reachable from U_α or V_α requires a factor of q for one of its vertices for exactly one of the two directions.

For the first copy of α , we need a factor of n for each vertex as it is the first time all of the vertices appear (this also accounts for the vertices in $U_\alpha \cap V_\alpha$). For the final copy of α^T , the vertices which are copies of vertices in $U_\alpha \cup V_\alpha$ have already been identified and we need a factor of at most $|V(\alpha)|$ to specify each vertex which is not a copy of a vertex in $U_\alpha \cup V_\alpha$ as each such vertex will be adjacent to a vertex which has been identified and has no constraint edge to the right.

Putting everything together gives the result. ■

B Reserving edges to preserve left/middle shapes

Given a shape α , and a subset of edges $\text{Res}(\alpha)$ that we are to reserve, we will only apply conditioning on edges $E(\alpha) \setminus \text{Res}(\alpha)$.

Proposition B.1. *Given any middle shape τ such that all vertices have a path to $U_\tau \cup V_\tau$, there is a set of edges $\text{Res}(\tau)$ of size at most $|V(\tau) \setminus (U_\tau \cap V_\tau)|$ such that every vertex has a path to $U_\tau \cup V_\tau$ in $\text{Res}(\tau)$, and removing any subset of the edges $E(\tau) \setminus \text{Res}(\tau)$ does not change the size of the minimum vertex separator of the shape.*

Proof. Select $|U_\tau|$ vertex-disjoint paths between U_τ and V_τ . This is enough to ensure that the minimum vertex separator will not decrease. To ensure that every vertex has a path to $U_\tau \cup V_\tau$, also add in edges connecting to these paths until all vertices are connected. ■

Proposition B.2. *For a middle intersection γ, τ, γ' we can reserve at most $|V(G) \setminus (U_\gamma \cap V_\gamma)|$ edges in γ so that the reserved ribbons are still a middle intersection.*

Proof. Select $|U_\gamma|$ vertex-disjoint paths between U_γ and $\text{Int}(\gamma) \cup V_\gamma$. ■

Proposition B.3. *For a left shape γ , there is a set $\text{Res}(\gamma) \subseteq E(\gamma)$ of edges such that $|\text{Res}(\gamma)| \leq 2|V(\gamma)|$ and reserving $\text{Res}(\gamma)$ is sufficient to guarantee γ remains a left shape under removal of any subset of edges in $E(\gamma) \setminus \text{Res}(\gamma)$.*

Proof. Observe that if γ is a left part, for each vertex $v \in V_\gamma$, if we add another copy of v (duplicating all of the edges incident to v) to V_γ (and U_γ if $v \in U_\gamma \cap V_\gamma$), then there are $|V_\gamma| + 1$ vertex disjoint paths from U_γ to $V_\gamma \cup \{v_{\text{duplicated}}\}$ where $v_{\text{duplicated}}$ is the copy of v . Moreover, this condition is sufficient to guarantee that γ is a left part.

To see this, assume that there is a separator S of size $|V_\gamma|$ between U_γ and $V_\gamma \cup \{v_{\text{duplicated}}\}$. Note that $S \cap V(\gamma)$ is a separator between U_γ and V_γ and V_γ is the only such separator of size at

most $|V_\gamma|$, so we must have that $S = V_\gamma$. However, this is impossible as there is a path from U_γ to $v_{\text{duplicated}}$ which does not pass through S .

To see the moreover part, assume there is a separator S between U_γ and V_γ of size at most $|V_\gamma|$ which is not equal to V_γ . Since $S = V_\gamma$, there is a vertex $v \in V(\gamma) \setminus S$. This condition says that there are $|V_\gamma| + 1$ paths from U_γ to V_γ which are vertex disjoint except that two of the paths end at v . At least one of these paths does not pass through S , so S is not a separator between U_γ and V_γ . Thus, there is no such separator S and γ is a left part.

Now consider the following graph G which is designed so that there is an integer-valued flow of value k from s to t in G corresponding to k vertex disjoint paths from U_γ to V_γ .

1. For each vertex $u \in V(\gamma)$, create two copies u_{in} and u_{out} of u . Add a directed edge with capacity 1 from u_{in} to u_{out} .
2. Add directed edges with capacity 1 from s to u_{in} for each $u \in U_\gamma$. Similarly, add directed edges with capacity 1 from v_{out} to t for each $v \in V_\gamma$.
3. For each edge $(u, v) \in E(\gamma)$, create a directed edge with capacity 1 from u_{out} to v_{in} and a directed edge with capacity 1 from v_{out} to u_{in} .

Choose $|V_\gamma|$ vertex disjoint paths from U_γ to V_γ and take the corresponding paths in G . Now observe that for each $v \in V_\gamma$, if we increased the capacity of the edges from v_{in} to v_{out} and from v_{out} to t to 2, this would increase the max flow value to $|V_\gamma| + 1$. This implies that if we run an iteration of Ford-Fulkerson on the residual graph, we must be able to reach v_{in} for each $v \in V_\gamma$. If we use DFS or BFS to explore the residual graph, we need at most $|V(\alpha)|$ of the non u_{in} to u_{out} edges to reach all of these destinations.

Adding these edges to the $|V_\gamma|$ vertex disjoint paths from U_γ to V_γ gives a set of edges of size at most $2|V(\gamma)|$ which can be reserved to guarantee that γ remains a left part. ■

C Second proof of Lemma 6.9

Lemma 6.9. *For any linearization β of τ_P and all separators S for β ,*

$$|E(\beta)| + 2|E_{\text{phantom}}| - |E(S)| \geq |V(\tau_P)| + |I_\beta| - |S|.$$

We prove this by exhibiting an explicit charging argument.

Proof. Let τ_P^{phant} be the multigraph formed from β plus two edges for each phantom edge in E_{phantom} . We need to assign each vertex of $V(\tau_P^{\text{phant}}) \setminus V(S)$ an edge of $E(\tau_P^{\text{phant}}) \setminus E(S)$, and we need to assign isolated vertices in I_β two edges. Note that the connectivity of τ_P^{phant} is exactly the same as τ_P , just the nonzero edge multiplicities are modified.

For vertices that are in the same connected component of τ_P^{phant} as a vertex in S , consider running a breadth-first search from S , and assign each edge to the vertex it explores. Vertices in I_β must be explored via a double edge, in order for them to become isolated during linearization, so assign both.

For vertices that are not in the same component of τ_P^{phant} as S , the edge assignment is more complicated. Let C be a component. Using Proposition 6.6, component C must intersect both U_β

and V_β . For the isolated vertices in $I_\beta \cap C$, order them by their distance from $U_\beta \cup V_\beta$. Charge them to the two edges along the shortest path to $U_\beta \cup V_\beta$. For the remaining non-isolated vertices, we claim that there is at least one additional double edge in C that has not yet been charged. As noted just above using [Proposition 6.6](#), there is a path P between U_β and V_β in this component. However, since C does not intersect S , which is by assumption a separator in β , the path cannot be entirely in β i.e. it must contain at least one double edge. If P does not pass through an isolated vertex, this double edge evidently has not yet been charged. If P passes through an isolated vertex, because all edges incident to isolated vertices are double edges, there must be more double edges than isolated vertices in P . Since each isolated vertex only charges one incident double edge, they can't all be charged. In either case, P contains an uncharged double edge.

Now contract the edges in τ_p^{phant} that were charged for isolated vertices, and order the non-isolated vertices by their distance from this double edge. The double edge can be used to charge its two endpoints (which are not isolated), and the other vertices can be charged using the next edge in the shortest path to the double edge. This completes the charging, and the proof of the lemma. ■

D Formalizing the PSD Decomposition

Here we prove [Lemma 6.13](#) (which we restate here for convenience) by formally going through the approximate PSD decomposition described in [Section 6.1](#).

Lemma 6.13. (*Decomposition in terms of ribbons*).

$$\begin{aligned} \sum_{\alpha \in \mathcal{S}} \lambda_\alpha \frac{M_\alpha}{|\text{Aut}(\alpha)|} &= \sum_{R \in \mathcal{S}} \lambda_R M_R = \\ &\left(\sum_{L \in \mathcal{L}} \lambda_L M_L \right) \left(\sum_{j=0}^{2D_{\text{SoS}}} (-1)^j \sum_{G_j, \dots, G_1, T, G'_1, \dots, G'_j} \lambda_{G_j \odot \dots \odot G_1 \odot T \odot G'_1{}^\top \odot \dots \odot G'_j{}^\top} M_{G_j \odot \dots \odot G_1 \odot T \odot G'_1{}^\top \odot \dots \odot G'_j{}^\top} \right) \left(\sum_{L' \in \mathcal{L}} \lambda_{L'} M_{L'} \right)^\top \\ &+ \text{truncation error}_{\text{too many vertices}} + \text{truncation error}_{\text{too many edges in one part}} \end{aligned}$$

where

1.

$$\begin{aligned} \text{truncation error}_{\text{too many vertices}} &= - \sum_{\substack{L, T, L': \\ |V(L \circ T \circ L')| > D_V, \\ L, T, L' \text{ are properly composable}}} \lambda_{L \circ T \circ L'} M_{L \circ T \circ L'} \\ &+ \sum_{j=1}^{2D_{\text{SoS}}} (-1)^{j+1} \sum_{\substack{L, G_j, \dots, G_1, T, G'_1, \dots, G'_j, L': \\ |V(L \odot G_j \odot \dots \odot G_1)| > D_V \text{ or} \\ |V(G'_1{}^\top \odot \dots \odot G'_j{}^\top \odot L')| > D_V}} \lambda_{L \odot G_j \odot \dots \odot G_1 \odot T \odot G'_1{}^\top \odot \dots \odot G'_j{}^\top \odot L'} M_L M_{G_j \odot \dots \odot G_1 \odot T \odot G'_1{}^\top \odot \dots \odot G'_j{}^\top} M_{L'}^\top \end{aligned}$$

$$\begin{aligned} \text{truncation error}_{\text{too many edges in one part}} &= \sum_{\substack{L, T, L': \\ |V(L \circ T \circ L')| \leq D_V, \\ |E_{\text{mid}}(T)| - |V(T)| > CD_{\text{SoS}}, \\ L, T, L' \text{ are properly composable}}} \lambda_{L \circ T \circ L'} M_{L \circ T \circ L'} \\ &+ \sum_{j=1}^{2D_{\text{SoS}}} (-1)^j \sum_{\substack{L, G_j, \dots, G_1, T, G'_1, \dots, G'_j, L': \\ |V(L \odot G_j \odot \dots \odot G_1)| \leq D_V \text{ and } |V(G'_1{}^\top \odot \dots \odot G'_j{}^\top \odot L')| \leq D_V \\ |E_{\text{mid}}(G_j)| - |V(G_j)| > CD_{\text{SoS}} \text{ or } |E_{\text{mid}}(G'_j)| - |V(G'_j)| > CD_{\text{SoS}} \\ L, (G_j \odot \dots \odot G_1 \odot T \odot G'_1{}^\top \odot \dots \odot G'_j{}^\top), L' \text{ are properly composable}}} \lambda_{L \odot G_j \odot \dots \odot G_1 \odot T \odot G'_1{}^\top \odot \dots \odot G'_j{}^\top \odot L'} M_{L \odot G_j \odot \dots \odot G_1 \odot T \odot G'_1{}^\top \odot \dots \odot G'_j{}^\top} M_{L'}^\top \end{aligned}$$

where for a ribbon R , $E_{\text{mid}}(R)$ is the set of edges of R which are not contained in A_R , not contained in B_R , and are not incident to any vertices in $A_R \cap B_R$.

2. in all of these sums, the ribbons $L, G_j, \dots, G_1, T, G'_1, \dots, G'_j, L'$ satisfy the following conditions:

- (a) $T \in \mathcal{M}$, $L, L' \in \mathcal{L}$, and each $G_i, G'_i \in \mathcal{L}$.
- (b) $L, G_j, \dots, G_1, T, G'_1{}^\top, \dots, G'_j{}^\top, L'{}^\top$ are composable.
- (c) The intersection pattern induced by $G_j, \dots, G_1, T, G'_1, \dots, G'_j$ is a middle intersection pattern.
- (d) $|V(T)| \leq D_V$, $|V(G_j \odot \dots \odot G_1)| \leq D_V$, and $|V(G'_1{}^\top \odot \dots \odot G'_j{}^\top)| \leq D_V$

- (e) Except when noted otherwise (which only happens for truncation error *too many edges in one part*), all of the ribbons $G_j, \dots, G_1, T, G'_1, \dots, G'_j$ (but not necessarily L, L') satisfy the constraint that $|E_{mid}(R)| - |V(R)| \leq CD_{SoS}$.

Proof. We abuse notation and write the sum over ribbons in \mathcal{S} as

$$\sum_{\alpha \in \mathcal{S}} \lambda_\alpha \frac{M_\alpha}{|\text{Aut}(\alpha)|} = \sum_{R \in \mathcal{S}} \lambda_R M_R.$$

(The automorphism group disappears as it was only there to ensure each ribbon is represented once anyway.)

From [Proposition 5.13](#), every ribbon $R \in \mathcal{S}$ decomposes into $R = L \circ T \circ L'^\top$ where $L, L' \in \mathcal{L}$ and $T \in \mathcal{M}$. We would like that

$$\sum_{R \in \mathcal{S}} \lambda_R M_R = \sum_{\substack{L, L' \in \mathcal{L}, T \in \mathcal{M}: \\ |V(L \circ T \circ L'^\top)| \leq D_V}} \lambda_{L \circ T \circ L'^\top} M_{L \circ T \circ L'^\top}.$$

However, this is not quite correct, as the ribbons on the right hand side may intersect.

Proposition D.1.

$$\sum_{R \in \mathcal{S}} \lambda_R M_R = \sum_{\substack{L, L' \in \mathcal{L}, T \in \mathcal{M}: \\ |V(L \circ T \circ L'^\top)| \leq D_V, \\ L, T, L'^\top \text{ are properly composable}}} \lambda_{L \circ T \circ L'^\top} M_{L \circ T \circ L'^\top}.$$

To start, we uncorrelate the sizes of L, T, L' . Add to the moment matrix

$$\sum_{\substack{L, L' \in \mathcal{L}, T \in \mathcal{M}: \\ |V(L \circ T \circ L'^\top)| > D_V, \\ L, T, L'^\top \text{ are properly composable}}} \lambda_{L \circ T \circ L'^\top} M_{L \circ T \circ L'^\top}.$$

This term is added to truncation error *too many vertices*.

We define matrices for intersection terms I_k , factored terms F_k , and truncation error T_k at level k recursively via the following process. We will maintain that the moment matrix satisfies for all k ,

$$\sum_{R \in \mathcal{S}} \lambda_R M_R = \left(\sum_{L \in \mathcal{L}} \lambda_L M_L \right) \left(\sum_{i=1}^k (-1)^{i+1} F_i \right) \left(\sum_{L \in \mathcal{L}} \lambda_L M_L \right)^\top + (-1)^k I_k + \sum_{i=1}^k (-1)^i T_i.$$

In order to make this equation hold, given I_k we will choose F_{k+1} , T_{k+1} , and I_{k+1} so that

$$I_k = \left(\sum_{L \in \mathcal{L}} \lambda_L M_L \right) F_{k+1} \left(\sum_{L \in \mathcal{L}} \lambda_L M_L \right)^\top - T_{k+1} - I_{k+1}$$

At the start of the $(k+1)$ th iteration, we will have that I_k is equal to a sum over middle ribbons T , “middle intersecting ribbons” $(G_k, \dots, G_1, G'_1, \dots, G'_k) \in \text{Mid}_T^{(k)}$, and non-intersecting ribbons L, L' . Furthermore, the ribbons $G_k, \dots, T, \dots, G'_k$ will satisfy $|E_{mid}(R)| - |V(R)| \leq CD_{SoS}$, in which case we say that “edge bounds hold”. Any time that a violation of this bound appears, we will

throw the term into the truncation error for too many edges in one part. For example, initially we throw away T with too many edges. Initially,

$$\begin{aligned}
I_0 &:= \sum_{\substack{L, L' \in \mathcal{L}, T \in \mathcal{M}: \\ |V(L)| \leq D_V, |V(T)| \leq D_V, |V(L'^\top)| \leq D_V, \\ L, T, L'^\top \text{ are properly composable,} \\ \text{edge bounds hold}}} \lambda_{L \odot T \odot L'^\top} M_{L \odot T \odot L'^\top} \\
T_0 &:= \sum_{\substack{L, L' \in \mathcal{L}, T \in \mathcal{M}: \\ |V(L)| \leq D_V, |V(T)| \leq D_V, |V(L'^\top)| \leq D_V, \\ |E_{\text{mid}}(T)| - |V(T)| > CD_{\text{SoS}}, \\ L, T, L'^\top \text{ are properly composable}}} \lambda_{L \odot T \odot L'^\top} M_{L \odot T \odot L'^\top} \\
F_0 &:= 0
\end{aligned}$$

Definition D.2. For a middle ribbon T , let $\text{Mid}_T^{(k)}$ be the collection of tuples of ribbons $G_k, \dots, G_1, G'_1, \dots, G'_k$ such that the ribbons are composable and the induced intersection pattern of $G_k \circ \dots \circ G_1 \circ T \circ G'_1 \circ \dots \circ G'_k$ is a middle intersection pattern.

The inductive hypothesis for I_k is,

$$\begin{aligned}
I_k &= \sum_{\substack{T \in \mathcal{M}, (G_k, \dots, G'_k) \in \text{Mid}_T^{(k)}, L, L' \in \mathcal{L}: \\ |V(L \odot G_k \odot \dots \odot G_1)| \leq D_V, |V(T)| \leq D_V, |V(G'_1{}^\top \odot \dots \odot G'_k{}^\top \odot L'^\top)| \leq D_V, \\ L, G_k \circ \dots \circ G_1 \circ T \circ G'_1{}^\top \circ \dots \circ G'_k{}^\top, L'^\top \text{ are properly composable,} \\ \text{edge bounds satisfied}}} \lambda_{L \odot G_k \odot \dots \odot G_1 \odot T \odot G'_1{}^\top \odot \dots \odot G'_k{}^\top \odot L'^\top} M_{L \odot G_k \odot \dots \odot G_1 \odot T \odot G'_1{}^\top \odot \dots \odot G'_k{}^\top \odot L'^\top}.
\end{aligned}$$

We can approximate I_k by

$$\left(\sum_{L \in \mathcal{L}} \lambda_L M_L \right) F_{k+1} \left(\sum_{L \in \mathcal{L}} \lambda_L M_L \right)^\top$$

where

$$\begin{aligned}
F_{k+1} &= \sum_{\substack{T \in \mathcal{M}, (G_k, \dots, G'_k) \in \text{Mid}_T^{(k)}: \\ |V(G_k \odot \dots \odot G_1)| \leq D_V, |V(T)| \leq D_V, |V(G'_1{}^\top \odot \dots \odot G'_k{}^\top)| \leq D_V, \\ \text{edge bounds satisfied}}} \lambda_{G_k \odot \dots \odot G_1 \odot T \odot G'_1{}^\top \odot \dots \odot G'_k{}^\top} M_{G_k \odot \dots \odot G_1 \odot T \odot G'_1{}^\top \odot \dots \odot G'_k{}^\top}
\end{aligned}$$

In order to make it so that $I_k = (\sum_{L \in \mathcal{L}} \lambda_L M_L) F_{k+1} (\sum_{L \in \mathcal{L}} \lambda_L M_L)^\top - T_{k+1} - I_{k+1}$, we need to handle the following issues.

1. I_k only contains terms where $|V(L \odot G_k \odot \dots \odot G_1)| \leq D_V$ and $|V(G'_1{}^\top \odot \dots \odot G'_k{}^\top \odot L'^\top)| \leq D_V$, whereas some larger terms appear in the approximation. To handle this, we add

$$\begin{aligned}
&\sum_{\substack{T \in \mathcal{M}, (G_k, \dots, G'_k) \in \text{Mid}_T^{(k)}, L, L' \in \mathcal{L}: \\ |V(G_k \odot \dots \odot G_1)| \leq D_V, |V(T)| \leq D_V, |V(G'_1{}^\top \odot \dots \odot G'_k{}^\top)| \leq D_V, \\ |V(L \odot G_k \odot \dots \odot G_1)| > D_V \text{ or } |V(G'_1{}^\top \odot \dots \odot G'_k{}^\top \odot L'^\top)| > D_V, \\ \text{edge bounds satisfied}}} \lambda_{L \odot G_k \odot \dots \odot G_1 \odot T \odot G'_1{}^\top \odot \dots \odot G'_k{}^\top \odot L'^\top} M_L M_{G_k \odot \dots \odot G_1 \odot T \odot G'_1{}^\top \odot \dots \odot G'_k{}^\top} M_{L'^\top}
\end{aligned}$$

to T_{k+1} . This is the source of the truncation error terms with too many vertices.

2. I_k only contains terms where $L, G_k \circ \dots \circ G_1 \circ T \circ G_1^\top \circ \dots \circ G_k^\top, L^\top$ are properly composable. The remaining terms are $L, G_k \circ \dots \circ G_1 \circ T \circ G_1^\top \circ \dots \circ G_k^\top, L^\top$ such that they are not properly composable, and these will mostly be put into I_{k+1} .

$$\sum_{\substack{T \in \mathcal{M}, (G_k, \dots, G_k') \in \text{Mid}_T^{(k)}, L, L' \in \mathcal{L}: \\ |V(L \circ G_k \circ \dots \circ G_1)| \leq D_V, |V(T)| \leq D_V, |V(G_1^\top \circ \dots \circ G_k^\top \circ L^\top)| \leq D_V, \\ L, G_k \circ \dots \circ G_1 \circ T \circ G_1^\top \circ \dots \circ G_k^\top, L^\top \text{ are not properly composable,} \\ \text{edge bounds satisfied}}} \lambda_{L \circ G_k \circ \dots \circ G_k' \circ L^\top} M_L M_{G_k \circ \dots \circ G_k'} M_{L^\top}$$

These terms do not yet match our inductive hypothesis for I_{k+1} ; for each $L, G_k \circ \dots \circ G_1 \circ T \circ G_1^\top \circ \dots \circ G_k^\top, L^\top$ which are not properly composable, we must separate out the intersecting portions G_{k+1}, G_{k+1}' . To do this, decompose L as $L_2 \circ G_{k+1}$ where $B_{L_2} = A_{G_{k+1}}$ is the leftmost minimum vertex separator between A_L and $B_L \cup \{\text{intersected vertices}\}$. We decompose L' as $L_2' \circ G_{k+1}'$ in a similar way. This is exactly the definition that $G_{k+1}, \dots, G_1, T, G_1', \dots, G_{k+1}'$ are a middle intersection, [Definition 6.3](#). That is, $(G_{k+1}, \dots, G_{k+1}') \in \text{Mid}^{(k+1)}$.

Claim D.3. L_2 and L_2' are left ribbons.

Proof. Definitionally, B_{L_2} is the unique minimum vertex separator of L_2 . Furthermore, reachability of all vertices in L_2 is inherited from L . \blacksquare

We record a lemma that will be needed later:

Lemma D.4. *Factoring $L = L_2 \circ G_{k+1}$ is oblivious to edges inside A_L, B_L , edges incident to $A_L \cap B_L$, or edges not in L .*

Proof. Edges inside A_L or B_L do not affect the connectivity between A_L, B_L ; the same is true for edges incident to $A_L \cap B_L$ since all vertices of $A_L \cap B_L$ must be taken in any separator of A_L and B_L . The last claim is clear because the factoring depends only on which vertices of L intersected, as well as the structure of L . \blacksquare

Observe that now instead of summing over $L, \text{Mid}^{(k)}, L'$, we may sum over $L_2, \text{Mid}^{(k+1)}, L_2'$. That is, if we fix T and the ribbons G_k, \dots, G_k' , then L determines the pair L_2, G_{k+1} , and vice versa.

We take I_{k+1} to be the subset of G_{k+1} that satisfy the edge bound $|E_{\text{mid}}(G_{k+1})| - |V(G_{k+1})| \leq CD_{\text{SoS}}$. (L_2 and L_2' are written as L, L' here so as to more clearly match the inductive hypothesis.)

$$\sum_{\substack{T \in \mathcal{M}, (G_{k+1}, \dots, G_{k+1}') \in \text{Mid}_T^{(k+1)}, L, L' \in \mathcal{L}: \\ |V(L \circ G_{k+1} \circ \dots \circ G_1)| \leq D_V, |V(T)| \leq D_V, |V(G_1^\top \circ \dots \circ G_{k+1}'^\top \circ L^\top)| \leq D_V, \\ L, G_{k+1} \circ \dots \circ G_1 \circ T \circ G_1^\top \circ \dots \circ G_{k+1}'^\top, L^\top \text{ are properly composable,} \\ \text{edge bounds satisfied}}} \lambda_{L \circ G_{k+1} \circ \dots \circ G_{k+1}' \circ L^\top} M_{L \circ G_{k+1} \circ \dots \circ G_{k+1}' \circ L^\top}$$

3. For technical reasons, we want to stop this process if $|E_{\text{mid}}(G_{k+1})| - |V(G_{k+1})| > CD_{\text{SoS}}$ or $|E_{\text{mid}}(G_{k+1}')| - |V(G_{k+1}')| > CD_{\text{SoS}}$. Thus we take such terms G_{k+1} and G_{k+1}' and add them to T_{k+1} instead of to I_{k+1} .

$$\sum_{\substack{T \in \mathcal{M}, (G_{k+1}, \dots, G_{k+1}') \in \text{Mid}_T^{(k+1)}, L, L' \in \mathcal{L}: \\ |V(L \circ G_k \circ \dots \circ G_1)| \leq D_V, |V(T)| \leq D_V, |V(G_1^\top \circ \dots \circ G_k^\top \circ L^\top)| \leq D_V, \\ |E_{\text{mid}}(G_{k+1})| - |V(G_{k+1})| > CD_{\text{SoS}} \text{ or } |E_{\text{mid}}(G_{k+1}')| - |V(G_{k+1}')| > CD_{\text{SoS}}, \\ L, G_{k+1} \circ \dots \circ G_1 \circ T \circ G_1^\top \circ \dots \circ G_{k+1}'^\top, L^\top \text{ are properly composable,} \\ \text{edge bounds satisfied up to } k}}} \lambda_{L \circ G_{k+1} \circ \dots \circ G_{k+1}' \circ L^\top} M_{L \circ G_{k+1} \circ \dots \circ G_{k+1}' \circ L^\top}$$

This is the source of the truncation error terms with too many edges.

Iteratively applying this procedure gives the result.

Proposition D.5. *The recursion terminates within $2D_{\text{SoS}}$ steps and $I_{2D_{\text{SoS}}} = 0$.*

Proof. Each additional intersection G_{k+1}, G'_{k+1} is nontrivial and increases $|A_{G_{k+1}}| + |A_{G'_{k+1}}|$. Since $|A_{G_k}|$ is upper bounded by D_{SoS} , the recursion ends within $2D_{\text{SoS}}$ steps. ■
■