The SDP Value of Random 2CSPs

Amulya Musipatla ✉
Carnegie Mellon University, Pittsburgh, PA, USA

Ryan O'Donnell ✉
Carnegie Mellon University, Pittsburgh, PA, USA

Tselil Schramm ✉
Stanford University, CA, USA

Xinyu Wu ✉
Carnegie Mellon University, Pittsburgh, PA, USA

Abstract

We consider a very wide class of models for sparse random Boolean 2CSPs; equivalently, degree-2 optimization problems over $\{\pm 1\}^n$. For each model $\mathcal{M}$, we identify the "high-probability value" $s^*_{\mathcal{M}}$ of the natural SDP relaxation (equivalently, the quantum value). That is, for all $\epsilon > 0$ we show that the SDP optimum of a random n-variable instance is (when normalized by n) in the range $(s^*_{\mathcal{M}} - \epsilon, s^*_{\mathcal{M}} + \epsilon)$ with high probability. Our class of models includes non-regular CSPs, and ones where the SDP relaxation value is strictly smaller than the spectral relaxation value.

2012 ACM Subject Classification Theory of computation → Random network models; Theory of computation → Approximation algorithms analysis

Keywords and phrases Random constraint satisfaction problems

Digital Object Identifier 10.4230/LIPIcs.ICALP.2022.97

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: https://arxiv.org/abs/2108.01038

Funding *Ryan O'Donnell, Xinyu Wu*: Supported by NSF grant FET-1909310.

1 Introduction

A large number of important algorithmic tasks can be construed as constraint satisfaction problems (CSPs): finding an assignment to Boolean variables to optimize the number of satisfied constraints. Almost every form of constraint optimization is NP-complete; thus one is led to questions of efficiently finding near-optimal solutions, or understanding the complexity of average-case rather than worst-case instances. Indeed, understanding the complexity of random sparse CSPs is of major importance not just in traditional algorithms theory, but also in, e.g., cryptography [25], statistical physics [27], and learning theory [14].

Suppose we fix the model $\mathcal{M}$ for a random sparse CSP on n variables (e.g., random k-SAT with a certain clause density). Then it is widely believed that there should be a constant $c^*_{\mathcal{M}}$ such that the optimal value of a random instance is $c^*_{\mathcal{M}} \pm o_{n\to\infty}(1)$ with high probability (whp). (Here we define the optimal value to mean the maximum number of simultaneously satisfiable constraints, divided by the number of variables.) Unfortunately, it is extremely difficult to prove this sort of result; indeed, it was considered a major breakthrough when Bayati, Gamarnik, and Tetali [6] established it for one of the simplest possible cases: Max-Cut on random d-regular graphs (which we will denote by $\mathcal{MC}_d$). Actually "identifying" the value $c^*_{\mathcal{M}}$ (beyond just proving its existence) is even more challenging. It is generally possible to estimate $c^*_{\mathcal{M}}$ using heuristic methods from statistical physics, but making these estimates

49th International Colloquium on Automata, Languages, and Programming (ICALP 2022).
Editors: Mikołaj Bojańczyk, Emanuela Merelli, and David P. Woodruff;
Article No. 97; pp. 97:1–97:19
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

rigorous is beyond the reach of current methods. Taking again the example of Max-Cut on d-regular random graphs, it was only recently [16] that the value $c^*_{\mathcal{MC}_d}$ was determined up to a factor of $1 \pm o_{d\to\infty}(1)$. The value for any particular d, e.g. $c^*_{\mathcal{MC}_3}$, has yet to be established.

Returning to algorithmic questions, we can ask about the *computational feasibility* of optimally solving sparse random CSPs. There are two complementary questions to ask: given a random instance from model $\mathcal{M}$ (with presumed optimal value $c^*_{\mathcal{M}} \pm o_{n\to\infty}(1)$), can one efficiently *find* a solution achieving value $\gtrapprox c^*_{\mathcal{M}}$, and can one efficiently *certify* that every solution achieves value $\lessapprox c^*_{\mathcal{M}}$? The former question is seemingly a bit more tractable; for example, a very recent breakthrough of Montanari [29] gives an efficient algorithm for (whp) finding a cut in a random graph $\mathcal{G}(n,p)$ graph of value at least $(1-\epsilon)c^*_{\mathcal{G}(n,p)}$. On the other hand, we do not know any algorithm for efficiently certifying (whp) that a random instance has value at most $(1+\epsilon)c^*_{\mathcal{G}(n,p)}$. Indeed, it reasonable to conjecture that no such algorithm exists, leading to an example of a so-called "information-computation gap".

To bring evidence for this we can consider *semidefinite programming* (SDP), which provides efficient algorithms for certifying an upper bound on the optimal value of a CSP [20]. Indeed, it is known [35] that, under the Unique Games Conjecture, the basic SDP relaxation provides essentially optimal certificates for CSPs in the *worst case.* In this paper we in particular consider Boolean 2CSPs – more generally, optimizing a homogeneous degree-2 polynomial over the hypercube – as this is the setting where semidefinite programming is most natural. Again, for a fixed model $\mathcal{M}$ of random sparse Boolean 2CSPs, one expects there should exist a constant $s^*_{\mathcal{M}}$ such that the optimal SDP-value of an instance from $\mathcal{M}$ is whp $s^*_{\mathcal{M}} \pm o_{n\to\infty}(1)$. Philosophically, since semidefinite programming is doable in polynomial time, one may be more optimistic about proving this and explicitly identifying $s^*_{\mathcal{M}}$. Indeed, some results in this direction have recently been established.

1.1 Prior work on identifying high-probability SDP values

Let us consider the most basic case: $\mathcal{MC}_d$, Max-Cut on random d-regular graphs. For ease of notation, we will consider the equivalent problem of maximizing $\frac{1}{n}x^\top(-\boldsymbol{A})x$ over $x \in \{\pm 1\}^n$, where $\boldsymbol{A}$ is the adjacency matrix of a random n-vertex d-regular graph.[1] Although $s^*_{\mathcal{MC}_d}$, the high-probability SDP relaxation value, was pursued as early as 1987 [8] (see also [18]), it was not until 2015 that Montanari and Sen [30] established the precise result $s^*_{\mathcal{MC}_d} = 2\sqrt{d-1}$. That is, in a random d-regular graph, whp the basic SDP relaxation value [8, 37, 15, 34] for the size of the maximum cut is $(\frac{d}{4} + \sqrt{d-1} \pm o_{n\to\infty}(1))n$. Here the special number $2\sqrt{d-1}$ is the maximum eigenvalue of the d-regular infinite tree.

The proof of this result has two components: showing $\textsc{Sdp}(-\boldsymbol{A}) \geq 2\sqrt{d-1} - \epsilon$ whp, and showing $\textsc{Sdp-Dual}(-\boldsymbol{A}) \leq 2\sqrt{d-1} + \epsilon$ whp. Here $\textsc{Sdp}(A) = \max\{\langle \rho, A\rangle : \rho \succeq 0,\ \rho_{ii} = \frac{1}{n}\ \forall i\}$ denotes the "primal" SDP value on matrix A (commonly associated with the Goemans–Williamson rounding algorithm [23]), and $\textsc{Sdp-Dual}(A) = \min\{\lambda_{\max}(A + \operatorname{diag}(\zeta)) : \operatorname{avg}_i(\zeta_i) = 0\}$ denotes the (equal) "dual" SDP value on A. To show the latter bound, it is sufficient to observe that $\textsc{Sdp-Dual}(-\boldsymbol{A}) \leq \lambda_{\max}(-\boldsymbol{A})$, the "eigenvalue bound", and $\lambda_{\max}(-\boldsymbol{A}) \leq 2\sqrt{d-1} + o_n(1)$ whp by Friedman's Theorem [21]. As for lower-bounding $\textsc{Sdp}(-\boldsymbol{A})$, Montanari and Sen used the "Gaussian wave" method [19, 13, 24] to construct primal SDP solutions achieving at least $2\sqrt{d-1} - \epsilon$ (whp). The idea here is essentially to build the SDP solutions using an approximate eigenvector (of finite support) of the infinite d-regular tree achieving eigenvalue $2\sqrt{d-1} - \epsilon$; the fact that SDP constraint "$\rho_{ii} = \frac{1}{n}\ \forall i$" can be satisfied relies heavily on the regularity of the graph.

[1] Throughout this work, **boldface** is used to denote random variables.

▶ Remark 1. The Montanari–Sen result in passing establishes that the (high-probability) eigenvalue and SDP bounds coincide for random regular graphs. This is consistent with a known theme, that the two bounds tend to be the same (or nearly so) for graphs where "every vertex looks similar" (in particular, for regular graphs). This theme dates back to Delorme and Poljak [15], who showed that $\textsc{Sdp-Dual}(-A) = \lambda_{\max}(-A)$ whenever A is the adjacency matrix of a vertex-transitive graph.

Subsequently, the high-probability SDP value $s^*_{\mathcal{M}}$ was established for a few other models of random regular 2CSPs. Deshpande, Montanari, O'Donnell, Schramm, and Sen [17] showed that for $\mathcal{M} = \mathcal{NAE}3_c$ – meaning random regular instances of NAE-3SAT (not-all-equals 3Sat) with each variable participating in c clauses – we have $s^*_{\mathcal{M}} = \frac{9}{8} - \frac{3}{8} \cdot \frac{\sqrt{c-1}-\sqrt{2}}{c}$. We remark that NAE-3SAT is effectively a 2CSP, as the predicate $\text{NAE}_3 : \{\pm 1\}^3 \to \{0,1\}$ may be expressed as $\frac{3}{4} - \frac{1}{4}(xy + yz + zx)$, supported on the "triangle" formed by variables x, y, z. The analysis in this paper is somewhat similar to that in [30], but with the infinite graph $\mathfrak{X} = K_3 \star K_3 \star \cdots \star K_3$ (c times) replacing the d-regular infinite tree. This $\mathfrak{X}$ is the $2c$-regular infinite "tree of triangles" depicted (partly, in the case $c = 3$) in Figure 1. More generally, [17] established the high-probability SDP value for large random (edge-signed) graphs that locally resemble $K_r \star K_r \star \cdots \star K_r$, the $(r-1)c$-regular infinite "tree of cliques K_r". (The $r = 2$ case essentially generalizes [30].) As in [30], $s^*_{\mathcal{M}}$ coincides with the (high-probability) eigenvalue bound. The upper bound on $s^*_{\mathcal{M}}$ is shown by using Bordenave's proof [9] of Friedman's Theorem for random (c, r)-biregular graphs. The lower bound on $s^*_{\mathcal{M}}$ is shown using the Gaussian wave technique, relying on the distance-regularity of the graphs $K_r \star K_r \star \cdots \star K_r$ (indeed, it is known that every infinite distance-regular graph is of this form).

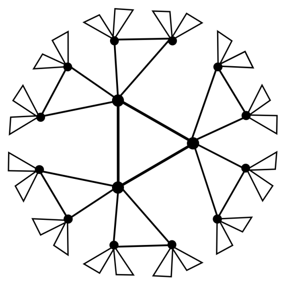

Figure 1 The 6-regular infinite graph $K_3 \star K_3 \star K_3$, modeling random 3-regular NAE3-SAT.

Mohanty, O'Donnell, and Paredes [28] generalized the preceding two results to the case of "two-eigenvalue" 2CSPs. Roughly speaking, these are 2CSPs formed by placing copies of a small weighted "constraint graph" $\mathcal{H}$ – required to have just two distinct eigenvalues – in a random regular fashion onto n vertices/variables. (This is indeed a generalization [17], as cliques have just two distinct eigenvalues.) As two-eigenvalue examples, [28] considered CSPs with the "CHSH constraint" – and its generalizations, the "Forrelation$_k$" constraints – which are important in quantum information theory [11, 1]. Here the SDP value of an instance is particularly relevant as it is precisely the optimal "quantum entangled value" of the 2CSP [12]. Once again, it is shown in [28] that the high-probability SDP and eigenvalue bounds coincide for these types of CSPs. The two-eigenvalue condition is used at a technical level in both the variant of Bordenave's theorem proven for the eigenvalue upper bound, and in the Gaussian wave construction in the SDP lower bound.

Most recently, O'Donnell and Wu [33] used the results of Bordenave and Collins [10] (a substantial generalization of [9]) to establish the high-probability *eigenvalue relaxation bound* for very wide range of random 2CSPs, encompassing all those previously mentioned: namely, any quadratic optimization problem defined by random "matrix polynomial lifts" over literals.

1.2 Our work

In this work, we establish the high-probability SDP value $s^*_{\mathcal{M}}$ for random instances of any 2CSP model $\mathcal{M}$ arising from lifting matrix polynomials (as in [33]). This generalizes all previously described work on SDP values, and covers many more cases, including random lifts of any base 2CSP (as used in certain community detection models) and random graphs modeled on any free/additive/amalgamated product. Such graphs have seen numerous applications within theoretical computer science, for example the zig-zag product in derandomization (e.g. [36]) and lifts of 2CSPs in the study of the stochastic block model (e.g [2]). At the end of this section we derive a corollary of our main theorem with applications to the latter. See Section 2 for more details and definitions, and see [33] for a thorough description of the kinds of random graph/2CSP models that can arise from matrix polynomials.

Very briefly, a matrix polynomial p is a small, explicit "recipe" for producing random n-vertex edge-weighted graphs, each of which "locally resembles" an associated *infinite* graph $\mathfrak{X}_p$. For example, $p_3(Y_1, Y_2, Y_3) \coloneqq Y_1 + Y_2 + Y_3$ is a recipe for random (edge-signed) 3-regular n-vertex graphs, and here $\mathfrak{X}_{p_3}$ is the infinite 3-regular tree. As another example, if $p_{333}(Z_{1,1}, \dots, Z_{3,3})$ denotes the following matrix polynomial –

$$\begin{pmatrix} 0 & Z_{1,1}Z^*_{1,2} + Z_{2,1}Z^*_{2,2} + Z_{3,1}Z^*_{3,2} & Z_{1,1}Z^*_{1,3} + Z_{2,1}Z^*_{2,3} + Z_{3,1}Z^*_{3,3} \\ Z_{1,2}Z^*_{1,1} + Z_{2,2}Z^*_{2,1} + Z_{3,2}Z^*_{3,1} & 0 & Z_{1,2}Z^*_{1,3} + Z_{2,2}Z^*_{2,3} + Z_{3,2}Z^*_{3,3} \\ Z_{1,3}Z^*_{1,1} + Z_{2,3}Z^*_{2,1} + Z_{3,3}Z^*_{3,1} & Z_{1,3}Z^*_{1,2} + Z_{2,3}Z^*_{2,2} + Z_{3,3}Z^*_{3,2} & 0 \end{pmatrix} \tag{1}$$

– then p_{333} is a recipe for random (edge-signed) 6-regular n-vertex graphs where every vertex participates in 3 triangles. In this case, $\mathfrak{X}_{p_{333}}$ is the infinite graph (partly) depicted in Figure 1. The Bordenave–Collins theorem [10] shows that if $\boldsymbol{A}$ is the adjacency matrix of a random *unsigned* n-vertex graph produced from a matrix polynomial p, then whp the "nontrivial" spectrum of $\boldsymbol{A}$ will be within ϵ (in Hausdorff distance) of the spectrum of $\mathfrak{X}_p$. In the course of derandomizing this theorem, O'Donnell and Wu [33] established that for random *edge-signed* graphs, the modifier "nontrivial" should be dropped. As a consequence, in the signed case one gets $\lambda_{\max}(\boldsymbol{A}) \approx \lambda_{\max}(\mathfrak{X}_p)$ up to an additive ϵ, whp; i.e., the high-probability eigenvalue bound for CSPs derived from p is precisely $\lambda_{\max}(\mathfrak{X}_p)$. We remark that for simple enough p there are formulas for $\lambda_{\max}(\mathfrak{X}_p)$; regarding our two example above, it is $2\sqrt{2}$ for $p = p_3$, and it is 5 for $p = p_{333}$. In particular, if p is a linear matrix polynomial, $\lambda_{\max}(\mathfrak{X}_p)$ may be determined numerically with the assistance of a formula of Lehner [26] (see also [22] for the case of standard random lifts of a fixed base graph).

In this paper we investigate the high-probability *SDP* value – denote it s^*_p – of a large random 2CSP (Boolean quadratic optimization problem) produced by a matrix polynomial p. Critically, our level of generality lets us consider *non-regular* random graph models, in contrast to all previous work. Because of this, we see cases in contrast to Remark 1, where (whp) the SDP value is strictly smaller than the eigenvalue relaxation bound. As a simple example, for random edge-signed $(2, 3)$-biregular graphs, the high-probability eigenvalue bound is $\sqrt{2-1} + \sqrt{3-1} = 1 + \sqrt{2} \approx 2.414$, but our work establishes that the high-probability SDP value is $\sqrt{\frac{13}{4} + 2\sqrt{2} - \frac{1}{10}} \approx 2.365$.

An essential part of our work is establishing the appropriate notion of the "SDP value" of an infinite graph $\mathfrak{X}_p$, with adjacency operator A_∞. While the eigenvalue bound $\lambda_{\max}(A_\infty)$ makes sense for the infinite-dimensional operator A_∞, the SDP relaxation does not. The definition $\text{SDP}(A_\infty) = \max\{\langle \rho, A_\infty\rangle : \rho \succeq 0,\ \rho_{ii} = \frac{1}{n}\ \forall i\}$ does not make sense, since "n" is ∞. Alternatively, if one tries the normalization $\rho_{ii} = 1$, any such ρ will have infinite trace, and hence $\langle \rho, A_\infty\rangle$ may be ∞. Indeed, since the only control we have on A_∞'s "size" will be an operator norm ("∞-norm") bound, the expression $\langle \rho, A_\infty\rangle$ is only guaranteed to make sense if ρ is restricted to be trace-class (i.e., have finite "1-norm").

On the other hand, we know that the eigenvalue bound $\lambda_{\max}(A_\infty)$ is too weak, intuitively because it does not properly "rebalance" graphs $\mathfrak{X}_p$ that are not regular/vertex-transitive. The key to obtaining the correct bound is introducing a new notion, intermediate between the eigenvalue and SDP bounds, that is appropriate for graphs $\mathfrak{X}_p$ arising from matrix polynomial recipes. Although these graphs may be irregular, their definition also allows them to be viewed as *vertex-transitive* infinite graphs with $r \times r$ *matrix* edge-weights. In light of their vertex-transitivity, Remark 1 suggests that a "maximum eigenvalue"-type quantity – suitably defined for matrix-edge-weighted graphs – might serve as the sharp replacement for SDP value. We introduce such a quantity, calling it the *partitioned SDP* bound. Let G be an n-vertex graph with $r \times r$ matrices as edge-weights, and let A be its adjacency matrix, thought of as a Hermitian $n \times n$ matrix whose entries are $r \times r$ matrices. We will define

$$\text{PARTSDP}(A) = \sup\{\langle \rho, A\rangle : \rho \succeq 0,\ \text{tr}(\rho)_{ii} = \tfrac{1}{r}\}, \tag{2}$$

where here $\text{tr}(\rho)$ refers to the $r \times r$ matrix obtained by summing the entries on A's main diagonal (themselves $r \times r$ matrices), and $\text{tr}(\rho)_{ii}$ denotes the scalar in the (i,i)-position of $\text{tr}(\rho)$. This partitioned SDP bound may indeed be regarded as intermediate between the maximum eigenvalue and the SDP value. On one hand, given a scalar-edge-weighted n-vertex graph with adjacency matrix A, we may take $r = 1$ and then it is easily seen that $\text{PARTSDP}(A)$ coincides with $\lambda_{\max}(A)$. On the other hand, if we regard A as a 1×1 matrix and take $r = n$ (so that we a have single vertex with a self-loop weighted by all of A), then $\text{PARTSDP}(A) = \text{SDP}(A)$.

In the full version, we define $\text{PARTSDP}(A)$ even for bounded-degree infinite graphs with $r \times r$ edge-weights. Furthermore, it has the following SDP dual:

$$\text{PARTSDP-DUAL}(A) = \inf\{\lambda_{\max}(A + \mathbb{1}_{n\times n} \otimes \text{diag}(\zeta)) : \text{avg}(\zeta_1, \ldots, \zeta_r) = 0\}.$$

We show in the full version that there is no SDP duality gap between $\text{PARTSDP}(A)$ and $\text{PARTSDP-DUAL}(A)$, even in the case of infinite graphs. It is precisely the common value of $\text{PARTSDP}(\mathfrak{X}_p)$ and $\text{PARTSDP-DUAL}(\mathfrak{X}_p)$ that is the high-probability SDP value of large random 2CSPs produced from p; our main theorem is the following:

▶ **Theorem 2.** *Let p be a matrix polynomial with $r \times r$ coefficients. Let A_∞ be the adjacency operator (with $r \times r$ entries) of the associated infinite lift $\mathfrak{X}_p$, and write $s_p^* = \text{PARTSDP}(A_\infty) = \text{PARTSDP-DUAL}(A_\infty)$. Then for any $\epsilon, \beta > 0$ and sufficiently large n, if $\boldsymbol{A}_n$ is the adjacency matrix of a random edge-signed n-lift of p, it holds that $s_p^* - \epsilon \le \text{SDP}(\boldsymbol{A}_n) \le s_p^* + \epsilon$ except with probability at most β.*

Note that $\text{PARTSDP}(A_\infty)$ is a fixed value only dependent on the polynomial p, a finitary object.

The upper bound $\text{SDP}(\boldsymbol{A}_n) \le \text{PARTSDP-DUAL}(A_\infty) + \epsilon$ in this theorem can be derived from the results of [10, 33]. Our main work is to prove the lower bound $\text{SDP}(\boldsymbol{A}_n) \ge \text{PARTSDP}(A_\infty) - \epsilon$. For this, our approach is inspired by the Gaussian Wave construction of [30, 17] for d-regular graphs (in the random lifts model), which can be viewed as constructing a feasible $\textsc{Sdp}(\boldsymbol{A}_n)$ solution of value $\lambda_{\max}(A_\infty) - \epsilon$ using a truncated eigenfunction of A_∞. Since local neighborhoods in A_∞ look like local neighborhoods in $\boldsymbol{A}_n$ with high probability, the eigenfunction can be "pasted" almost everywhere into the graph $\boldsymbol{G}_n$, which gives an SDP solution of value near $\lambda_{\max}(A_\infty)$.

This approach runs into clear obstacles in our setting. Indeed, the raw eigenfunctions are of no use to us, as *the SDP value may be smaller than the spectral relaxation value.* Instead, we first show that there is a ρ_0 with only finitely many nonzero entries that achieves the sup in Equation (2) up to ϵ. This is effectively a finite $r \times r$ matrix-edge-weighted graph. We then show that this ρ_0 can (just as in the regular case) whp be "pasted" almost everywhere into the graph $\boldsymbol{G}_n$ defined by $\boldsymbol{A}_n$, which gives an SDP solution of value close to $\textsc{PartSdp}(A_\infty)$. The fact that $\mathfrak{X}_p$ and $\boldsymbol{G}_n$ are regarded as regular tree-like graphs with matrix edge-weights (rather than as irregular graphs with scalar edges-weights) is crucially used to show that the "pasted solution" satisfies the finite SDP's constraints "$\rho_{ii} = \frac{1}{n}\ \forall i$".

1.2.1 Application to hypothesis testing in block models

Montanari and Sen's aforementioned work [30] was motivated in part by the following question: can the basic SDP value serve as a *hypothesis test* for distinguishing Erdős-Rényi block models from random graphs with a planted partition/max-cut structure? Though polynomial-time hypothesis tests are known to exist whenever the task is information-theoretically possible [31] (even for the robust version of the problem, using a non-basic SDP [4]), the question of whether the canonical basic (Goemans–Williamson) SDP value can be used to hypothesis test is still interesting. Montanari and Sen show that the answer is affirmative for average degree $d = \Theta(1)$ large enough, but for small constant d their question remains open.

Using our results, we resolve the analogue of the question of Montanari and Sen in the large-planted-cut regime of the random-lift version of the block model, the so-called "equitable block model" [2]. The n-vertex equitable 2-community block model with internal degree a and external degree b is the distribution $\textsc{Block}(n, a, b)$ defined as a random $n/2$-lift of the 2-vertex graph with b parallel edges and a self-loops on each vertex.

▶ **Corollary 3.** *For any $\epsilon, \beta > 0$ and sufficiently large n, if $\boldsymbol{G}_n \sim \textsc{Block}(n, a, b)$ and $\boldsymbol{A}_n$ is the adjacency matrix of $\boldsymbol{G}_n$, then*

$$\left| \textsc{Sdp}(-\boldsymbol{A}_n) - \max(b - a,\ 2\sqrt{a + b - 1}) \right| \le \epsilon$$

except with probability at most β.

The bound $b-a$ is achieved by the integral cut which partitions the vertices according to their preimage in the 2-vertex graph. The upper bound $\textsc{Sdp}(-\boldsymbol{A}_n) \le \max(b - a, 2\sqrt{a + b - 1}) + \epsilon$ follows from the eigenvalue bound on $-\boldsymbol{A}_n$. As above, what is new is our lower bound which shows that $\textsc{Sdp}(-\boldsymbol{A}_n) \ge 2\sqrt{a + b - 1} - \epsilon$.

In an $(a + b)$-regular random graph, the result of Montanari and Sen [30] (see also [3]) implies that $\textsc{Sdp}(-\boldsymbol{A}_n) = 2\sqrt{a + b - 1} \pm \epsilon$ with high probability. Hence, Corollary 3 proves that the basic (Goemans–Williamson) SDP value is the same in both models and so does not furnish a hypothesis testing algorithm when $b - a < 2\sqrt{a + b - 1}$. This is consistent with the prediction of [2], who conjecture that when $b - a < 2\sqrt{a + b - 1}$, no polynomial-time algorithm can robustly (under the addition of random noise) hypothesis test between random regular graphs and the equitable block model.[2] In [2], the authors give a lower bound for a different SDP relaxation, but they do not characterize the basic SDP value.

[2] We remark that in the Erdős-Rényi version of the 2-community block model, the threshold for information-

2 Preliminaries

To preface the following definitions and concepts, we remark that our "real" interest is in graphs/2CSPs with real scalar weights. The introduction of matrix edge-weights facilitates two things: helping us define a wide variety of interesting scalar-weighted graphs via matrix polynomial lifts; and, facilitating the definition of PARTSDP(·), which we use to bound the SDP relaxation value of the associated 2CSPs. Our use of potentially complex matrices is also not really essential; we kept them because prior work that we utilize ([10], the tools in Section 3) is stated in terms of complex matrices. However the reader will not lose anything by assuming that all Hermitian matrices are in fact symmetric real matrices.

2.1 Matrix-weighted graphs

In the following definitions, we'll restrict attention to graphs with at-most-countable vertex sets V and bounded degree. We also often use bra-ket notation, with $(|v\rangle)_{v\in V}$ denoting the standard orthonormal basis for the complex vector space $\ell_2(V)$.

▶ **Definition 4** (Matrix-weighted graph). *Fix any $r \in \mathbb{N}^+$. A* matrix-weighted graph *will refer to a directed simple graph $G = (V, E)$ with self-loops allowed, in which each directed edge e has an associated* weight *$a_e \in \mathbb{C}^{r\times r}$. If $(v, w) \in E \implies (w, v) \in E$ and $a_{(w,v)} = a^*_{(v,w)}$, we say that G is an* undirected *matrix-weighted graph. The* adjacency matrix *of G is the operator A, acting on $\ell_2(V) \otimes \mathbb{C}^r$, given by*

$$\sum_{(v,w)\in E} |w\rangle\langle v| \otimes a_{(v,w)}.$$

It can be helpful to think to think of A in matrix form, as a $|V| \times |V|$ matrix whose entries are themselves $r \times r$ edge-weight matrices. Note that if G is undirected if and only if A is self-adjoint, $A = A^*$.

▶ **Definition 5** (Extension of a matrix-weighted adjacency matrix). *Given a $|V| \times |V|$ matrix A with $r \times r$ entries, we may also view it as a $|V|r \times |V|r$ matrix with scalar entries. When we wish to explicitly call attention to the distinction, we will call the latter matrix the* extension *of A, and denote it by $\widetilde{A}$.*

2.2 Matrix polynomials

▶ **Definition 6** (Matrix polynomial). *Let $Y_1, \dots, Y_d$ be formal indeterminates that are their own inverses, and let $Z_1, \dots, Z_e$ be formal indeterminates with formal inverses $Z^*_1, \dots, Z^*_e$. For a fixed r, we define a* matrix polynomial *to be a formal noncommutative polynomial p over the indeterminates $Y_1, \dots, Z^*_e$, with coefficients in $\mathbb{C}^{r\times r}$. In particular, we may write*

$$p = \sum_w a_w w,$$

where the sum is over words w on the alphabet of indeterminates, each a_w is in $\mathbb{C}^{r\times r}$, and only finitely many a_w are nonzero. Here we call a word reduced *if it has no adjacent Y_iY_i or $Z_iZ^*_i$ pairs. We will denote the empty word by $\mathbb{1}$.*

---

theoretic and computational distinguishability coincide. However, in the equitable case, the models are always information-theoretically distinguishable, since in BLOCK(n, a, b) there is always an integral solution of value exactly $b - a$, and this will not occur in the d-regular case with high probability.

As we will shortly describe, we will always be considering substituting unitary operators for the Z_i's, and unitary involution operators for the Y_i's. Thus we can think of Z_i^* as both the inverse *and* the "adjoint" of indeterminate Z_i, and similarly we think of $Y_i^* = Y_i$.

▶ **Definition 7** (Adjoint of a polynomial). *Given a matrix polynomial $p = \sum_w a_w w$ as above, we define its* adjoint *to be*

$$p^* = \sum_w a_w^* w^*,$$

where a^* is the usual adjoint of $a \in \mathbb{C}^{r\times r}$, and w^* is the adjointed reverse of w. That is, if $w = w_1 \cdots w_k$ then $w^* = w_k^* \cdots w_1^*$, where $\mathbb{1}^* = \mathbb{1}$, $Y_i^* = Y_i$, and $Z_i^{**} = Z_i$. We say p is self-adjoint *if $p^* = p$ formally.*

Note that in any self-adjoint polynomial, some terms will be self-adjoint, and others will come in self-adjoint pairs. In this work, we will only be considering self-adjoint polynomials.

2.3 Lifts of matrix polynomials

▶ **Definition 8** (n-lift). *Given a matrix polynomial over the indeterminates $Y_1, \dots, Z_e^*$, we define an n-*lift *to be a sequence $\mathcal{L} = (M_1, \dots, M_d, P_1, \dots, P_e)$ of $n \times n$ matrices, where each P_i is a signed permutation matrix and each M_i is a signed matching matrix.*[3] *A* random n-lift $\boldsymbol{\mathcal{L}} = (\boldsymbol{M}_1, \dots, \boldsymbol{M}_d, \boldsymbol{P}_1, \dots, \boldsymbol{P}_e)$ *is one where the matrices are chosen independently and uniformly at random. A* random unsigned n-lift $\boldsymbol{\mathcal{L}} = (\boldsymbol{M}_1, \dots, \boldsymbol{M}_d, \boldsymbol{P}_1, \dots, \boldsymbol{P}_e)$ *is one where the permutation and matching matrices are chosen independently and uniformly at random, and all the signs are $+1$.*

▶ **Definition 9** (Evaluation/substitution of lifts.). *Given an n-lift $\mathcal{L}$ and a word w, we define $\mathcal{L}^w$ to be the $n \times n$ operator obtained by substituting appropriately into w: namely, $Y_i = M_i$, $Z_i = P_i$, and $Z_i^* = P_i^*$ for each i (and substituting the empty word with the $n \times n$ identity operator). Given also a matrix polynomial $p = \sum_w a_w w$, we define the* evaluation of p at $\mathcal{L}$ *to be the following operator on $\mathbb{C}^n \otimes \mathbb{C}^r$:*[4]

$$p(\mathcal{L}) = \sum_w \mathcal{L}^w \otimes a_w.$$

▶ Remark 10. Note that each P_i is unitary and each M_i a unitary involution (as promised), so $p^*(\mathcal{L}) = p(\mathcal{L})^*$. Thus $p(\mathcal{L})$ is a self-adjoint operator whenever p^* is a self-adjoint polynomial. In this case we also have that $p(\mathcal{L})$ may be viewed as the adjacency matrix of an undirected graph on vertex set $[n]$ with $r \times r$ edge-weights.

Note that the evaluation $p(\mathcal{L})$ of a matrix polynomial may be viewed as the adjacency matrix of a undirected graph on $[n]$ with $r \times r$ edge-weights; or, its extension may be viewed as the adjacency matrix of an undirected graph on $[n] \times [r]$ with scalar edge-weights. In this way, each fixed matrix polynomial p, when applied to a random lift, gives rise to a random (undirected, scalar-weighted) graph model.

---

[3] A signed matching matrix is the adjacency matrix of a perfect matching with ± 1 edge-signs. If $d > 0$ then we must restrict to even n.

[4] Note that coefficients a_w are written on the left in p, as is conventional, but we take the tensor/Kronecker product on the right so that the matrix form of $p(\mathcal{L})$ may be more naturally regarded as an $n \times n$ matrix with $r \times r$ entries.

▶ **Example 11.** *A simple example is the matrix polynomial*

$$p(Y_1, Y_2, Y_3) = Y_1 + Y_2 + Y_3.$$

Here $r = 1$ and each coefficient is just the scalar 1*. This p gives rise to a model of random edge-signed* 3*-regular graphs on* $[n]$*.*

By moving to actual matrix coefficients with $r > 1$, one can get the random (signed) graph model given by randomly n-lifting any base r-vertex graph H.

▶ **Example 12.** *As a simple example,*

$$p(Z_1, Z_2, Z_3) = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} Z_1 + \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix} Z_1^* + \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} Z_2 + \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix} Z_2^* + \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} Z_3 + \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix} Z_3^*$$

is the recipe for random 3*-regular (edge-signed)* $(n+n)$*-vertex bipartite graphs. The reader may like to view this as a* 2×2 *matrix of polynomials,*

$$p(Z_1, Z_2, Z_3) = \begin{pmatrix} 0 & Z_1 + Z_2 + Z_3 \\ Z_1^* + Z_2^* + Z_3^* & 0 \end{pmatrix},$$

but recall that we actually Kronecker-product the coefficient matrices "on the other side". So rather than as a 2×2 *block-matrix with* $n \times n$ *blocks, we think of the resulting adjacency matrix as an* $n \times n$ *block-matrix with* 2×2 *blocks; equivalently, an* n*-vertex graph with* 2×2 *matrix edge-weights.*

▶ **Example 13.** *The matrix polynomial* p_{333} *mentioned in* (1) *gives an example of a* nonlinear *polynomial with matrix coefficients. Again, we wrote it there as a* 3×3 *matrix of polynomials for compactness, but for analysis purposes we will view it as a degree-*2 *polynomial with* 3×3 *coefficients.*

▶ **Definition 14** (∞-lift)**.** *Formally, we extend Definition 8 to the case of* $n = \infty$ *as follows. Let* V_∞ *denote the free product of groups* $\mathbb{Z}_2^{\star d} \star \mathbb{Z}^{\star e}$*, with its components generated by* $g_1, \dots, g_d$*,* $h_1, h_1^{-1}, \dots, h_e, h_e^{-1}$*. Thus the elements of* V_∞ *are in one-to-one correspondence with the reduced words over indeterminates* $Y_1, \dots, Z_e^*$*. The generators* $g_1, \dots, g_d, h_1, \dots, h_e$ *act as permutations on* V_∞ *by left-multiplication, with the first* d *in fact being matchings. We write* $\sigma_1, \dots, \sigma_{d+e}$ *for these permutations, and we also identify them with their associated permutation operators on* $\ell_2(V_\infty)$*. Finally, we write* $\mathfrak{L}_\infty = (\sigma_0, \dots, \sigma_{d+2e})$ *for* "the" ∞-lift *associated to* p*. (Note that this lift is "unsigned".)*

▶ **Definition 15** (Evaluation at the ∞-lift, and $\mathfrak{X}_p$.)**.** *The evaluation of a matrix polynomial* p *at the infinity lift* $\mathcal{L}_\infty$ *is now defined just as in Definition 9; the resulting operator* $p(\mathcal{L}_\infty)$ *operates on* $\ell_2(V_\infty) \otimes \mathbb{C}^r$*. We may think of the result as a matrix-weighted graph on vertex set* V_∞*, and we will sometimes denote this graph by* $\mathfrak{X}_p$*. When* p *is understood, we often write* $A_\infty = p(\mathcal{L}_\infty)$ *for the adjacency operator of* $\mathfrak{X}_p$*, which can be thought of as an infinite matrix with rows/columns indexed by* V_∞ *and entries from* $\mathbb{C}^{r \times r}$*, or as its "extension"* $\widetilde{A}_\infty$*, an infinite matrix with rows/columns indexed by* $V_\infty \times [r]$ *and scalar entries.*

▶ **Example 16.** *For the polynomial* $p = Y_1 + \cdots + Y_d$*, the corresponding graph* $\mathfrak{X}_p$ *is the infinite (unweighted)* d*-regular tree.*

We may now state a theorem which is essentially the main result ("Theorem 2") of [10]. The small difference is that our notion of random n-lifts, which includes ± 1 *signs* on the matchings/permutations, lets one eliminate mention of "trivial" eigenvalues (see [33, Thms. 1.9, 10.10]).

▶ **Theorem 17.** *Let p be a self-adjoint matrix polynomial with coefficients from $\mathbb{C}^{r\times r}$ on indeterminates $Y_1, \dots, Z_e^*$. Then for all $\epsilon, \beta > 0$ and sufficiently large n, the following holds:*

Let $\boldsymbol{A}_n = p(\boldsymbol{\mathcal{L}}_n)$, where $\boldsymbol{\mathcal{L}}_n$ is a random n-lift, and let $A_\infty = p(\mathcal{L}_\infty)$. Then except with probability at most β, the spectra $\operatorname{spec}(\boldsymbol{A}_n)$ and $\operatorname{spec}(A_\infty)$ are at Hausdorff distance at most ϵ

2.4 Random lifts as optimization problems

Given a Hermitian (i.e., self-adjoint) matrix $A \in \mathbb{C}^{n\times n}$, we are interested in the task of maximizing $x^\top Ax$ over all Boolean vectors $x \in \{\pm 1\}^n$. (Since A is Hermitian, the quantity $x^\top Ax$ is always real, so this maximization problem makes sense.) This is the same as maximizing the homogeneous degree-2 (commutative) polynomial $\sum_{i,j} A_{ij}x_ix_j$ over $x \in \{\pm 1\}^n$, and it is also essentially the same task as the Max-Cut problem on (scalar-)weighted undirected graphs. More precisely, if G is a weighted graph on vertex set $[n]$ with adjacency matrix A, then G's maximum cut is indicated by the $x \in \{\pm 1\}^n$ that maximizes $x^\top(-A)x$. For the sake of scaling we will also include a factor of $\frac{1}{n}$ in this optimization problem, leading to the following definition:

▶ **Definition 18** (Optimal value)**.** *Given a Hermitian matrix $A \in \mathbb{C}^{n\times n}$, we define*

$$\textsc{Opt}(A) = \sup_{x\in\{\pm 1\}^n} \left\{\tfrac{1}{n} x^\top Ax\right\} = \sup_{x\in\left\{\pm\frac{1}{\sqrt{n}}\right\}^n} \left\{x^\top Ax\right\}.$$

(For finite-dimensional A, the sup*s and* inf*s mentioned in this section are all achieved.)*

We remark that

$$x^\top Ax = \operatorname{tr}(x^\top Ax) = \operatorname{tr}(xx^\top A) = \langle xx^\top, A\rangle,$$

where we use the notation $\langle B, C\rangle = \operatorname{tr}(BC)$. Thus we also have

$$\textsc{Opt}(A) = \sup_{\rho\in \mathrm{Cut}_n} \left\{\tfrac{1}{n}\langle \rho, A\rangle\right\},$$

where Cut_n is the "cut polytope", the convex hull of all matrices of the form $xx^\top$ for $x \in \{\pm 1\}^n$. (Since $\langle \rho, A\rangle$ is linear in ρ, maximizing over the convex hull is the same as maximizing over the extreme points, which are just those matrices of the form $xx^\top$.)

The above optimization problem has a natural relaxation: maximizing $\frac{1}{n}x^\top Ax$ over all *unit* vectors x. This leads to the following efficiently computable upper bound on $\textsc{Opt}(A)$:

▶ **Definition 19** (Eigenvalue bound)**.** *Given a Hermitian matrix $A \in \mathbb{C}^{n\times n}$, we define the* eigenvalue bound *to be*

$$\textsc{Eig}(A) = \sup\{\langle \rho, A\rangle : \rho \succeq 0,\ \operatorname{tr}(\rho) = 1\},$$

where here $\rho \succeq 0$ denotes that ρ is (Hermitian and) positive semidefinite.

The matrices ρ being optimized over in $\textsc{Eig}(A)$ are known as *density matrices*; i.e., $\textsc{Eig}(A)$ is the maximal inner product between A and any density matrix. Note that if $\varrho \in \mathrm{Cut}_n$, then $\rho = \frac{1}{n}\varrho$ is a density matrix. Thus, $\textsc{Eig}(A)$ is a relaxation of $\textsc{Opt}(A)$, or in other words, $\textsc{Opt}(A) \le \textsc{Eig}(A)$.

The set of density matrices is convex, and it's well known that its extreme points are all the rank-1 density matrices; i.e., those ρ of the form $xx^\top$ for $x \in \mathbb{C}^n$ with $\|x\|_2^2 = 1$. Thus in $\textsc{Eig}(A)$ it is equivalent to just maximize over these extreme points:

$$\textsc{Eig}(A) = \sup_{\substack{x\in\mathbb{C}^n \\ \|x\|_2^2=1}} \left\{\langle xx^\top, A\rangle\right\} = \sup_{\substack{x\in\mathbb{C}^n \\ \|x\|_2^2=1}} \left\{x^\top Ax\right\}.$$

From this formula we see that $\text{EIG}(A)$ is also equal to $\lambda_{\max}(A)$, the maximum eigenvalue of A; hence the terminology "eigenvalue bound". One may also think of $\text{EIG}(A)$ and $\lambda_{\max}(A)$ as SDP duals of one another.

We now mention another well known, tighter, upper bound on $\text{OPT}(A)$.

▶ **Definition 20** (Basic SDP bound). *Given a Hermitian matrix $A \in \mathbb{C}^{n\times n}$, the* basic SDP bound *is defined to be*

$$\text{SDP}(A) = \sup\{\langle \rho, A\rangle : \rho \succeq 0,\ \rho_{ii} = \tfrac{1}{n}, \forall i\}.$$

Recall that an $n \times n$ matrix ϱ is a *correlation matrix* [38] if it is PSD and has all diagonal entries equal to 1. Thus $\text{SDP}(A)$ is equivalently maximizing $\frac{1}{n}\langle \varrho, A\rangle$ over all correlation matrices ϱ. We also note that any cut matrix is a correlation matrix, and any correlation matrix is a density matrix, hence so

$$\text{OPT}(A) \le \text{SDP}(A) \le \text{EIG}(A).$$

▶ **Definition 21** (Dual SDP bound). *The semidefinite dual of $\text{SDP}(A)$ is the following [15]:*

$$\text{SDP-DUAL}(A) = \inf_{\substack{\zeta\in\mathbb{R}^n \\ \text{avg}(\zeta_1,\dots,\zeta_n)=0}} \{\lambda_{\max}(A + \text{diag}(\zeta))\}.$$

Despite the fact that the usual "Slater condition" for strong SDP duality fails in this case (because the set of correlation matrices isn't full-dimensional), one can still show [34] that $\text{SDP}(A) = \text{SDP-DUAL}(A)$ indeed holds for finite-dimensional A.

▶ Remark 22. In this work we frequently consider matrix-weighted graphs with adjacency matrices A, thought of as $n\times n$ matrices with entries from $\mathbb{C}^{r\times r}$. For such matrices, whenever we write $\text{OPT}(A)$, we mean $\text{OPT}(\widetilde{A})$ for the $nr \times nr$ "extension" matrix $\widetilde{A}$ (see Definition 5), and similarly for $\text{EIG}(A)$, $\lambda_{\max}(A)$, $\text{SDP}(A)$, $\text{SDP-DUAL}(A)$.

As mentioned in Section 1, the eigenvalue bound $\lambda_{\max}(A)$ makes sense when A is the adjacency matrix of an infinite graph (with bounded degree). However $\text{SDP}(A)$ does not extend to the infinite case, as the number "n" appearing in its definition is not finite. On the other hand, we now introduce a new, intermediate, "maximum eigenvalue-like" bound that is appropriate for matrix-weighted graphs. This is the "partitioned SDP bound" appearing in the statement of our main Theorem 2. In the following Section 3, we will show that it generalizes well to the case of infinite graphs.

▶ **Definition 23** (Partitioned SDP bound). *Let A be an $n \times n$ Hermitian matrix with entries from $\mathbb{C}^{r\times r}$. We define its* partitioned SDP bound *to be*

$$\text{PARTSDP}(A) = \sup\{\langle \rho, A\rangle : \rho \succeq 0,\ \text{tr}(\rho) \in \tfrac{1}{r}\text{Corr}_r\},$$

where:

- *the matrices ρ are also thought of as $n \times n$ matrices with entries from $\mathbb{C}^{r\times r}$;*
- *$\langle \rho, A\rangle$ is interpreted as $\langle \widetilde{\rho}, \widetilde{A}\rangle$;*
- *$\text{tr}(\rho)$ denotes the sum of the diagonal entries of ρ, which is an $r \times r$ matrix;*
- *Corr_r is the set of $r \times r$ correlation matrices;*
- *in other words, the final condition is that $\text{tr}(\rho)_{ii} = \frac{1}{r}$ for all $i \in [r]$.*

▶ Remark 24. As mentioned, the partitioned SDP bound can be viewed as "intermediate" between the eigenvalue bound and the SDP bound. To explain this, suppose A is an $n \times n$ Hermitian matrix. On one hand, we can regard A as an $n \times n$ matrix with 1×1 matrix entries ($r = 1$); in this viewpoint, $\text{PARTSDP}(A) = \text{EIG}(A)$. On the other hand, we can regard A as a 1×1 matrix with a single $n \times n$ matrix entry ($r = n$); in this viewpoint, $\text{PARTSDP}(A) = \text{SDP}(A)$.

It is easy to see that the partitioned SDP bound indeed has an SDP formulation, and we now state its SDP dual:

▶ **Definition 25.** *The SDP dual of* $\textsc{PartSdp}(A)$ *is the following:*

$$\textsc{PartSdp-Dual}(A) = \inf_{\substack{\zeta\in\mathbb{R}^r\\ \operatorname{avg}(\zeta_1,\dots,\zeta_r)=0}} \{\lambda_{\max}(A + \mathbb{1}_{n\times n} \otimes \operatorname{diag}(\zeta))\}.$$

Weak SDP duality, $\textsc{PartSdp}(A) \le \textsc{PartSdp-Dual}(A)$, holds as always, but again it is not obvious that strong SDP duality holds. In fact, not only does strong duality hold, it even holds in the case of *infinite* matrices A. This fact is crucial for our work, and proving it the subject of the upcoming technical section.

3 The infinite SDPs

This technical section is deferred to the full version and we only state some results here. First, we show that strong duality $\textsc{PartSdp}(A) = \textsc{PartSdp-Dual}(A)$ holds, even for infinite matrices A with $r \times r$ entries. Even in the finite case this is not trivial, as the feasible region for the SDP $\textsc{PartSdp}(A)$ is not full-dimensional, and hence the Slater condition ensuring strong duality does not apply. The infinite case involves some additional technical considerations, needed so that we may eventually apply the strong duality theorem for conic linear programming of Bonnans and Shapiro [7, Thm. 2.187]. Second, we show that in the optimization problem $\textsc{PartSdp}(A)$, values arbitrarily close to the optimum can be achieved by matrices ρ of finite support (i.e., with only finitely many nonzero entries). Indeed (though we don't need this fact), these finite-rank ρ need only have rank at most r. This fact is familiar from the case of $r = 1$, where the optimizer in the eigenvalue bound Definition 19 is achieved by a ρ of rank 1 (namely $|\psi\rangle\langle\psi|$ for any maximum eigenvector $|\psi\rangle$). Finally, we consolidate all these results into a theorem statement suitable for use with graphs produced by infinite lifts of matrix polynomials.

3.1 SDP duality for matrix edge-weighted graphs

Let V_∞ be a countable set of nodes and let G_∞ be a bounded-degree graph on V_∞ with matrix edge-weights from $\mathbb{C}^{r\times r}$. Let A_∞ be the adjacency operator for G_∞, acting on $\ell_2(V_\infty) \otimes \mathbb{C}^r$ and assumed self-adjoint; we may think of it as an infinite matrix with rows and columns indexed by V_∞, and with entries from $\mathbb{C}^{r\times r}$.

For any $\epsilon > 0$, there exists:

- $\hat\zeta \in \mathbb{R}^r$ with $\operatorname{avg}_j(\hat\zeta_j) = 0$;
- a finite subset $F \subset V_\infty$;
- a PSD matrix ρ with rows/columns indexed by V_∞ and entries from $\mathbb{C}^{r\times r}$, supported on the rows/columns F, with

$$\operatorname{tr}(\rho)_{jj} = \tfrac{1}{r}, \quad j = 1 \dots r;$$

such that for

$$\hat A = \widetilde A_\infty + \mathbb{1}_{V_\infty} \otimes \operatorname{diag}(\hat\zeta),$$

we have

$$s^* := \lambda_{\max}(\hat A) = \textsc{PartSdp-Dual}(A_\infty) = \textsc{PartSdp}(A_\infty) \ge \langle\rho, A_\infty\rangle = \langle\rho_F, A_F\rangle \ge s^* - \epsilon, \tag{3}$$

where ρ_F, A_F denote ρ, A_∞ (respectively) restricted to the rows/columns F.

4 The SDP value of random matrix polynomial lifts

In this section we prove our main Theorem 2. To that end, let p be any self-adjoint matrix polynomial over indeterminates $Y_1, \dots, Y_d, Z_1, \dots, Z_e^*$ with $r \times r$ coefficients. Let $A_\infty = p(\mathcal{L}_\infty)$ denote the adjacency operator of the infinite lift $\mathfrak{X}_p$, and write $s^* = \textsc{PartSdp}(A_\infty) = \textsc{PartSdp-Dual}(A_\infty)$ as in Equation (3). Fix any $\epsilon, \beta > 0$, and let $\boldsymbol{A}_n = p(\boldsymbol{\mathcal{L}})$ denote the adjacency matrix of a corresponding n-lift, formed from $\boldsymbol{\mathcal{L}} = (\boldsymbol{M}_1, \dots, \boldsymbol{M}_d, \boldsymbol{P}_1, \dots, \boldsymbol{P}_e)$. Our goal is to show that except with probability at most β (assuming n is sufficiently large),

$$s^* - \epsilon \leq \textsc{Sdp}(\boldsymbol{A}_n) = \textsc{Sdp-Dual}(\boldsymbol{A}_n) \leq s^* + \epsilon.$$

Given our setup, the upper bound follows easily from prior work, namely Theorem 17. Let $\hat{\zeta}$ and $\hat{A}$ be as in Section 3.1, and consider the matrix polynomial p' defined by

$$p' = p + \operatorname{diag}(\hat{\zeta})\mathbb{1}.$$

Then on one hand, the ∞-lift of p' has adjacency operator precisely $\hat{A}$; on the other hand,

$$\boldsymbol{A}'_n := p'(\boldsymbol{\mathcal{L}}) = \boldsymbol{A}_n + \mathbb{1}_{n \times n} \otimes \operatorname{diag}(\hat{\zeta}).$$

Thus Theorem 17 tells us that except with probability $\beta/2$ (provided n is large enough), the spectra $\operatorname{spec}(\boldsymbol{A}'_n)$ and $\operatorname{spec}(\hat{A})$ are at Hausdorff distance at most ϵ, from which it follows that

$$\lambda_{\max}(\boldsymbol{A}'_n) \leq \lambda_{\max}(\hat{A}) + \epsilon = s^* + \epsilon.$$

But this indeed proves $\textsc{Sdp-Dual}(\boldsymbol{A}_n) \leq s^* + \epsilon$, because $\hat{\zeta}$ has $\operatorname{avg}(\hat{\zeta}) = 0$ and hence is feasible for $\textsc{Sdp-Dual}(\boldsymbol{A}_n)$.

It therefore remains to prove $\textsc{Sdp}(\boldsymbol{A}_n) \geq s^* - \epsilon$.

4.1 A lower bound on the basic SDP value

In this section we complete the proof of our main theorem by showing that $\textsc{Sdp}(\boldsymbol{A}_n) \geq s^* - \epsilon$ except with probability at most $o(1) = o_{n\to\infty}(1)$ (which is at most $\beta/2$ as needed, provided n is large enough).

Let F, ρ, ρ_F, A_F be as in Section 3.1, except with that section's "ϵ" replaced by $\epsilon/2$, so that $\langle \rho_F, A_F \rangle \geq s^* - \epsilon/2$. Adding finitely many vertices to F if necessary, we may assume that it consists of all reduced words over $Y_1, \dots, Y_d, Z_1, \dots, Z_e^*$ of length at most some finite f_0. We also make the following definition:

▶ **Definition 26** (Cycle in a lift). *Given an n-lift $\mathcal{L}$, a* cycle of length $\ell > 0$ *is a pair (i, w), where $i \in [n]$ and w is a reduced word of length ℓ such that $\mathcal{L}^w |i\rangle = \pm |i\rangle$, and $\langle i|\mathcal{L}^{w'}|i\rangle = 0$ (i.e., $\mathcal{L}^{w'} |i\rangle \neq \pm |i\rangle$) for all proper prefixes w' of w.*

We will employ the following basic random graph result, [10, Lem. 23], stated in our language:

▶ **Lemma 27.** *For the random n-lift $\boldsymbol{\mathcal{L}}$, the expected number of cycles of length ℓ is $O(\ell(d + 2e - 1)^\ell)$.*

Applying this for all $\ell \leq f := 2f_0 + \deg(p)$ and using Markov's inequality, we conclude:

▶ **Corollary 28.** *Except with probability at most $n^{-.99}$, the random n-lift $\boldsymbol{\mathcal{L}}$ has at most $O(n^{.99})$ cycles of length at most f. In this case, we can exclude a set of "bad" vertices $B \subseteq [n]$ with $|B|/n \leq O(n^{-.01}) = o(1)$ so that:*

$$\forall i \notin B, \quad \forall \text{ reduced words } w \text{ with } 0 < |w| \leq 2f_0 + \deg(p), \qquad \langle i|\boldsymbol{\mathcal{L}}^w|i\rangle = 0.$$

We henceforth fix an outcome $\boldsymbol{\mathcal{L}} = \mathcal{L}$ (and hence $\boldsymbol{A}_n = A_n$) such that the conclusion of Corollary 28 holds, accruing our $o(1)$ probability of failure. Under this assumption, we will show $\textsc{Sdp}(A_n) \geq s^* - \epsilon/2 - o(1)$, which is sufficient to complete the proof.

▶ Remark 29. Note that our choice of $\boldsymbol{\mathcal{L}}$ depends only on the structure (i.e. the non-zero entries of A_n), not the signs. Therefore, we can take the signs in $\boldsymbol{\mathcal{L}}$ to be arbitrary; in particular, $\boldsymbol{\mathcal{L}}$ can be unsigned.

Our plan will be to first construct a "provisional" near-feasible PSD solution σ for $\textsc{Sdp}(A_n)$, with rows/columns indexed by $[n]$ and with $r \times r$ entries, such that:

- $|\langle \sigma, A_n \rangle - \langle \rho_F, A_F \rangle| \leq o(1)$, and hence $\langle \sigma, A_n \rangle \geq s^* - \epsilon/2 - o(1)$;
- for $i \notin B$, the $r \times r$ matrix σ_{ii} has diagonal entries $\frac{1}{nr}$.

Then, we will show how to "fix" σ to a some σ' that is truly feasible for $\textsc{Sdp}(A_n)$, while still having $\langle \sigma', A_n \rangle \geq \langle \sigma, A_n \rangle - o(1) \geq s^* - \epsilon/2 - o(1)$.

4.1.1 Constructing a near-feasible solution

▶ **Definition 30.** *For each $i \in [n]$, we define a linear operator $\Phi_i : \mathbb{C}^F \to \mathbb{C}^n$ by*

$$\Phi_i = \sum_{v \in F} \mathcal{L}^v \, |i\rangle\langle v| \, ,$$

and also $\widetilde{\Phi}_i = \Phi_i \otimes \mathbb{1}_{r \times r}$. We furthermore define σ_i to be the $n \times n$ matrix with $r \times r$ entries whose extension $\widetilde{\sigma}_i$ is

$$\widetilde{\sigma}_i = \widetilde{\Phi}_i \cdot \widetilde{\rho}_F \cdot \widetilde{\Phi}_i^{\mathsf{T}}.$$

▶ Remark 31. $\widetilde{\sigma}_i$ is PSD, being the conjugation by $\widetilde{\Phi}_i$ of the PSD operator $\widetilde{\rho}_F$.

▶ **Proposition 32.** *If $i \notin B$, then $\widetilde{\Phi}_i^{\mathsf{T}} \widetilde{A}_n \widetilde{\Phi}_i = \widetilde{A}_F$.*

Proof. Thinking of $\widetilde{\Phi}_i^{\mathsf{T}} \widetilde{A}_n \widetilde{\Phi}_i$ as an $F \times F$ matrix of $r \times r$ matrices, it follows that its (u, v) entry is given by

$$\sum_{\text{term } a_w w \text{ in } p} \langle i | \mathcal{L}^{v^* w u} | i \rangle \, a_w.$$

On the other hand, the (u, v) entry of A_F is by definition

$$\sum_{\text{term } a_w w \text{ in } p} 1[v^* w u = \mathbb{1}] a_w,$$

where "$v^* w u = \mathbb{1}$" denotes that the reduced form of word $v^* w u$ is the empty word. We therefore have equality for all u, v provided $\langle i | \mathcal{L}^{v^* w u} | i \rangle = 0$ whenever $v^* w u \neq \emptyset$. But Corollary 28 tells us this indeed holds for $i \notin B$, because $|v^* w u| \leq 2 f_0 + \deg(p)$. ◀

▶ **Corollary 33.** *For $i \notin B$ we have $\langle \sigma_i, A_n \rangle = \langle \rho_F, A_F \rangle$.*

Proof. When $i \notin B$,

$$\langle \sigma_i, A_n \rangle = \operatorname{tr}(\widetilde{\sigma}_i \widetilde{A}_n) = \operatorname{tr}(\widetilde{\Phi}_i \widetilde{\rho}_F \widetilde{\Phi}_i^{\mathsf{T}} \widetilde{A}_n) = \operatorname{tr}(\widetilde{\rho}_F \widetilde{\Phi}_i^{\mathsf{T}} \widetilde{A}_n \widetilde{\Phi}_i) = \operatorname{tr}(\widetilde{\rho}_F \widetilde{A}_F) = \langle \rho_F, A_F \rangle,$$

where the last equality used Proposition 32 ◀

We now define our "provisional" SDP solution σ via

$$\sigma = \operatorname*{avg}_{i\in[n]}\{\sigma_i\};$$

this is indeed PSD, being the average of PSD operators. Using Corollary 33, $|B|/n = o(1)$, and the fact that $|\langle\sigma_i, A_n\rangle| \le O(1)$ for every i (since σ_i only has $O(1)$ nonzero entries, each bounded in magnitude by $O(1)$), we conclude:

▶ **Proposition 34.** *$|\langle\sigma, A_n\rangle - \langle\rho_F, A_F\rangle| \le o(1)$, and hence $\langle\sigma, A_n\rangle \ge s^* - \epsilon/2 - o(1)$.*

Now similar to Proposition 35 we have the following:

▶ **Proposition 35.** *If $j \notin B$, then the (j, j) entry of σ is $\frac{1}{n}\operatorname{tr}(\rho_F)$ (and hence is an $r \times r$ matrix with diagonal entries equal to $\frac{1}{nr}$).*

Proof. By definition, the (j, j) entry of σ is

$$\begin{aligned}
&= \operatorname*{avg}_{i\in[n]} \sum_{u,v\in F} \langle j|\mathcal{L}^u|i\rangle \langle u|\rho_F|v\rangle \langle i|\mathcal{L}^{v^*}|j\rangle \\
&= \tfrac{1}{n} \sum_{u,v\in F} \langle u|\rho_F|v\rangle \cdot \sum_{i\in[n]} \langle j|\mathcal{L}^u|i\rangle \langle i|\mathcal{L}^{v^*}|j\rangle \\
&= \tfrac{1}{n} \sum_{u,v\in F} \langle u|\rho_F|v\rangle \cdot \langle j|\mathcal{L}^{uv^*}|j\rangle && \text{(since } \textstyle\sum_i |i\rangle\langle i| = \mathbb{1})
\end{aligned}$$

where we are writing $\langle u|\rho_F|v\rangle$ for the $r \times r$ matrix at the (u, v) entry of ρ_F. Now when $j \notin B$, we have that $\langle j|\mathcal{L}^{vw^*}|j\rangle = 1[vw^* = \emptyset]$ by Corollary 28, since $|vw^*| \le 2f_0$. Thus all summands above drop out, except for the ones with $u = v$; this indeed gives $\frac{1}{n}\operatorname{tr}(\rho_F)$. ◀

4.1.2 Fixing σ

Finally, we slightly fix σ to make it truly feasible for $\textsc{Sdp}(A)$. Let σ' be the $n \times n$ matrix, with entries from $\mathbb{C}^{r\times r}$, defined as follows:

$$\sigma'_{ij} = \begin{cases} \sigma_{ij} & \text{if } i, j \notin B, \\ \frac{1}{nr}\mathbb{1}_{r\times r} & \text{if } i = j \in B, \\ 0 & \text{else.} \end{cases}$$

This σ' is easily seen to be PSD, being a principal submatrix of the PSD matrix σ, direct-summed with the PSD matrix $\frac{1}{nr}\mathbb{1}_{r\times r}$. As well, the $nr \times nr$ extension matrix $\widetilde{\sigma}'$ has all diagonal entries equal to $\frac{1}{nr}$, by Proposition 35. Thus $\widetilde{\sigma}'$ is feasible for $\textsc{Sdp}(A_n) = \textsc{Sdp}(\widetilde{A}_n)$, and it remains for us to show that

$$\langle\sigma, A_n\rangle - \langle\sigma', A_n\rangle \le o(1); \tag{4}$$

this will imply $\langle\sigma', A_n\rangle \ge s^* - \epsilon/2 - o(1)$ by Proposition 34, and hence $\textsc{Sdp}(A_n) \ge s^* - \epsilon$ (for sufficiently large n), as desired.

We have

$$\langle\sigma, A_n\rangle - \langle\sigma', A_n\rangle = \langle\sigma - \sigma', A_n\rangle \le \|\widetilde{\sigma} - \widetilde{\sigma}'\|_1 \|\widetilde{A}_n\|_\infty.$$

Next,

$$\|\widetilde{A}_n\|_\infty = \|p(\mathcal{L})\|_\infty = \|\sum_w \mathcal{L}^w \otimes a_w\|_\infty \le \sum_w \|\mathcal{L}^w\|_\infty \cdot \|a_w\|_\infty = \sum_w \|a_w\|_\infty \le O(1).$$

Here the final equality is because each $\mathcal{L}^w$ is a signed permutation matrix (hence has $\|\mathcal{L}^w\|_\infty = 1$), and the final inequality is because p has only constantly many coefficients, of constant size. Thus to establish Inequality (4), it remains to show $\|\widetilde{\sigma} - \widetilde{\sigma}'\|_1 \leq o(1)$.

Define the orthogonal projection matrices $\Pi_B = \sum_{i \in B} |i\rangle\langle i| \otimes \mathbb{1}_{r\times r}$ and similarly $\Pi_{\overline{B}}$, where $\overline{B} = [n] \setminus B$. Observe that

$$\sigma' = \tfrac{1}{nr}\Pi_B + \Pi_{\overline{B}}\sigma\Pi_{\overline{B}},$$

and thus

$$\|\widetilde{\sigma} - \widetilde{\sigma}'\|_1 = \|\widetilde{\sigma} - \widetilde{\Pi}_{\overline{B}}\widetilde{\sigma}\widetilde{\Pi}_{\overline{B}} - \tfrac{1}{nr}\widetilde{\Pi}_B\|_1 \leq \|\widetilde{\sigma} - \widetilde{\Pi}_{\overline{B}}\widetilde{\sigma}\widetilde{\Pi}_{\overline{B}}\|_1 + \tfrac{1}{nr}\|\widetilde{\Pi}_B\|_1.$$

But $\frac{1}{nr}\|\widetilde{\Pi}_B\|_1 = \frac{|B|}{n} = o(1)$, so it remains to show

$$\|\widetilde{\sigma} - \widetilde{\Pi}_{\overline{B}}\widetilde{\sigma}\widetilde{\Pi}_{\overline{B}}\|_1 \leq o(1).$$

Note that $\widetilde{\sigma} \in \mathbb{C}^{nr\times nr}$ is nearly a density matrix: it is PSD, and all but a $1 - o(1)$ fraction of its diagonal entries are $\frac{1}{nr}$, with the remaining ones being bounded in magnitude by $O(\frac{1}{n})$. Thus $\operatorname{tr}(\widetilde{\sigma}) = 1 \pm o(1)$, and we can therefore scale $\widetilde{\sigma}$ by a $1 \pm o(1)$ factor to produce a true density matrix $\hat{\sigma}$. Clearly it now suffices to show

$$\|\hat{\sigma} - \widetilde{\Pi}_{\overline{B}}\hat{\sigma}\widetilde{\Pi}_{\overline{B}}\|_1 \leq o(1).$$

But this follows from Winter's Gentle Measurement Lemma [39, Lem. 9], which bounds the quantity on the left by $\sqrt{8\lambda}$, where $\lambda = 1 - \operatorname{tr}(\hat{\sigma}\widetilde{\Pi}_{\overline{B}}) = o(1)$. This completes the proof.

5 Application to block models

In this section we will describe the application of our results to the hypothesis testing problem in block models; that is, using the (Goemans–Williamson) SDP value of the negated adjacency matrix $\textsc{Sdp}(-A)$ to distinguish between uniformly regular d-regular graphs and the "equitable stochastic block model" described by Bandeira et al and others [32, 5, 2].

In the n-vertex equitable 2-community block model, the n vertices are divided into two groups, and each vertex has a edges to vertices of the same group, and b edges to vertices of the other group (and $a + b = d$). A random graph of this form is generated by taking a random unsigned $n/2$-lift of the 2-vertex graph which has b parallel edges, and a self-loops on each vertex. The corresponding polynomial which describes the lift is

$$p(Y_1, \dots, Y_{2a}, Z_1, \dots, Z_b) = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}\left(\sum_{i=1}^{a} Y_i\right) + \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}\left(\sum_{i=a+1}^{2a} Y_i\right) + \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}\left(\sum_{i=1}^{b} Z_i\right) + \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}\left(\sum_{i=1}^{a} Z_i^*\right). \tag{5}$$

By evaluating p at a random unsigned $n/2$-lift $\boldsymbol{\mathcal{L}} = (\boldsymbol{M}_1, \dots, \boldsymbol{M}_{2a}, \boldsymbol{P}_1, \dots, \boldsymbol{P}_b)$ we obtain (the adjacency matrix of) a random graph $\boldsymbol{A}_n$ in the n-vertex equitable 2-community block model. We note that the infinite lift $p(\mathcal{L}_\infty)$ is (multiple copies of) the infinite d-regular graph; thus in particular we have $\operatorname{spec}(A_\infty) = [-2\sqrt{d-1}, 2\sqrt{d-1}]$.

To upper bound $\textsc{Sdp}(-\boldsymbol{A}_n)$, we use the eigenvalue bound $\textsc{Sdp}(-\boldsymbol{A}_n) \leq \lambda_{\max}(-\boldsymbol{A}_n)$. In contrast to the previous sections of our paper, the finite random lift here is unsigned, hence we need to take some care with the "trivial eigenvalues" of the resulting lift.

▶ **Definition 36** (Trivial eigenvalues). *Given a matrix polynomial p with coefficients from $\mathbb{C}^{r\times r}$, the associated* trivial eigenvalues *are simply the (multiset of) eigenvalues of $p(1,\dots,1) \in \mathbb{C}^{r\times r}$.*

Note that the trivial eigenvalues of our polynomial in Equation (5) are those of $\begin{pmatrix} a & b \\ b & a \end{pmatrix}$, namely $a+b$ and $a-b$. In the case of unsigned matrix polynomial lifts, Bordenave and Collins [10] show the following:

▶ **Theorem 37.** *Let p be a self-adjoint matrix polynomial with coefficients from $\mathbb{C}^{r\times r}$ on indeterminates $Y_1,\dots,Z_e^*$. Write T for the trivial eigenvalues of p. Then for all $\epsilon,\beta > 0$ and sufficiently large n, the following holds:*

Let $\boldsymbol{A}_n = p(\boldsymbol{\mathcal{L}}_n)$, where $\boldsymbol{\mathcal{L}}_n$ is a random unsigned n-lift, and let $A_\infty = p(\mathcal{L}_\infty)$. Then except with probability at most β, the spectra $\operatorname{spec}(\boldsymbol{A}_{n,\perp}) \setminus T$ and $\operatorname{spec}(A_\infty)$ are at Hausdorff distance at most ϵ.

Applying this to the equitable stochastic block model, it means that the nontrivial eigenvalues of graphs in the model are between $-2\sqrt{d-1}-\epsilon$ and $2\sqrt{d-1}+\epsilon$ with high probability (i.e., the same range as for random d-regular graphs). Recalling that the trivial eigenvalues of $-\boldsymbol{A}_n$ are $-(a+b)$ and $-(a-b)$, we indeed get $\lambda_{\max}(\boldsymbol{A}_n) \leq \max(b-a,\ 2\sqrt{a+b-1})+\epsilon$ with high probability, implying the upper bound in Corollary 3

As for the lower bound in Corollary 3, we refer to Remark 29 to see that the proof in Section 4.1 holds regardless of signs on the lift used to construct the SDP solution. Thus Section 4.1 also establishes a lower bound of $\textsc{Sdp}(-\boldsymbol{A}_n) \geq 2\sqrt{a+b-1}-\epsilon$. Meanwhile, in the $\textsc{Block}(n,a,b)$ we can always take an integral solution that is the ± 1-indicator for the two communities in the partition; this will have value exactly $b-a$ and hence we also always have $\textsc{Sdp}(-\boldsymbol{A}_n) \geq b-a$. This completes the proof of the lower bound in Corollary 3.

References

1 Scott Aaronson and Andris Ambainis. Forrelation: a problem that optimally separates quantum from classical computing. *SIAM J. Comput.*, 47(3):982–1038, 2018. doi:10.1137/15M1050902.

2 Afonso S Bandeira, Jess Banks, Dmitriy Kunisky, Christopher Moore, and Alex Wein. Spectral planting and the hardness of refuting cuts, colorability, and communities in random graphs. In *Proceedings of the 34th Annual Conference on Learning Theory*, pages 410–473. PMLR, 2021.

3 Jess Banks, Robert Kleinberg, and Cristopher Moore. The lovász theta function for random regular graphs and community detection in the hard regime. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, 2017.

4 Jess Banks, Sidhanth Mohanty, and Prasad Raghavendra. Local statistics, semidefinite programming, and community detection. In *soda21*, pages 1298–1316. SIAM, 2021.

5 Paolo Barucca. Spectral partitioning in equitable graphs. *Physical Review E*, 95(6):062310, 2017.

6 Mohsen Bayati, David Gamarnik, and Prasad Tetali. Combinatorial approach to the interpolation method and scaling limits in sparse random graphs. *Annals of Probability*, 41(6):4080–4115, 2013.

7 J. Frédéric Bonnans and Alexander Shapiro. *Perturbation Analysis of Optimization Problems*. Springer Series in Operations Research. Springer-Verlag, New York, 2000. doi:10.1007/978-1-4612-1394-9.

8 Ravi Boppana. Eigenvalues and graph bisection: An average-case analysis. In *Proceedings of the 28th Annual IEEE Symposium on Foundations of Computer Science*, pages 280–285, 1987.

9 Charles Bordenave. A new proof of Friedman's second eigenvalue theorem and its extension to random lifts. *Ann. Sci. Éc. Norm. Supér. (4)*, 53(6):1393–1439, 2020. doi:10.24033/asens.245.

10 Charles Bordenave and Benoît Collins. Eigenvalues of random lifts and polynomials of random permutation matrices. *Annals of Mathematics*, 190(3):811–875, 2019. doi:10.4007/annals.2019.190.3.3.

11 John Clauser, Michael Horne, Abner Shimony, and Richard Holt. Proposed experiment to test local hidden-variable theories. *Physical Review Letters*, 23(15):880–884, 1969.

12 Richard Cleve, Peter Høyer, Benjamin Toner, and John Watrous. Consequences and limits of nonlocal strategies. In *Proceedings of the 19th Annual Computational Complexity Conference*, pages 236–249. IEEE, 2004.

13 Endre Csóka, Balázs Gerencsér, Viktor Harangi, and Bálint Virág. Invariant Gaussian processes and independent sets on regular graphs of large girth. *Random Structures Algorithms*, 47(2):284–303, 2015. doi:10.1002/rsa.20547.

14 Amit Daniely, Nati Linial, and Shai Shalev-Shwartz. From average case complexity to improper learning complexity. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 441–448, 2014.

15 Charles Delorme and Svatopluk Poljak. Laplacian eigenvalues and the maximum cut problem. *Mathematical Programming*, 62:557–574, 1993.

16 Amir Dembo, Andrea Montanari, Subhabrata Sen, et al. Extremal cuts of sparse random graphs. *Annals of Probability*, 45(2):1190–1217, 2017.

17 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 30th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2305–2321. SIAM, 2019.

18 Wilm Donath and Alan Hoffman. Lower bounds for the partitioning of graphs. *IBM J. Res. Develop.*, 17:420–425, 1973. doi:10.1147/rd.175.0420.

19 Yehonatan Elon. Gaussian waves on the regular tree. Technical report, arXiv, 2009. arXiv:0907.5065.

20 Uriel Feige and László Lovász. Two-prover one-round proof systems: Their power and their problems. In *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*, pages 733–744, 1992.

21 Joel Friedman. A proof of Alon's second eigenvalue conjecture and related problems. *Memoirs of the American Mathematical Society*, 195(910):viii+100, 2008.

22 Jorge Garza-Vargas and Archit Kulkarni. Spectra of infinite graphs via freeness with amalgamation. Technical report, arXiv, 2021. arXiv:1912.10137.

23 Michel Goemans and David Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. Assoc. Comput. Mach.*, 42(6):1115–1145, 1995. doi:10.1145/227683.227684.

24 Viktor Harangi and Bálint Virág. Independence ratio and random eigenvectors in transitive graphs. *Ann. Probab.*, 43(5):2810–2840, 2015. doi:10.1214/14-AOP952.

25 Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability Obfuscation from well-founded assumptions. Technical report, arXiv, 2020. arXiv:2008.09317.

26 Franz Lehner. Computing norms of free operators with matrix coefficients. *Amer. J. Math.*, 121(3):453–486, 1999. URL: http://muse.jhu.edu/journals/american_journal_of_mathematics/v121/121.3lehner.pdf.

27 Marc Mézard and Andrea Montanari. *Information, physics, and computation*. Oxford Graduate Texts. Oxford University Press, Oxford, 2009. doi:10.1093/acprof:oso/9780198570837.001.0001.

28 Sidhanth Mohanty, Ryan O'Donnell, and Pedro Paredes. The SDP value for random two-eigenvalue CSPs. In Christophe Paul and Markus Bläser, editors, *Proceedings of the 37th Annual Symposium on Theoretical Aspects of Computer Science*, volume 154 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 50:1–50:45, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.STACS.2020.50.

29 Andrea Montanari. Optimization of the Sherrington–Kirkpatrick hamiltonian. *SIAM Journal on Computing*, FOCS19, 2021.

30 Andrea Montanari and Subhabrata Sen. Semidefinite programs on sparse random graphs and their application to community detection. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing*, pages 814–827. ACM, New York, 2016. `doi:10.1145/2897518.2897548`.

31 Elchanan Mossel, Joe Neeman, and Allan Sly. Reconstruction and estimation in the planted partition model. *Probability Theory and Related Fields*, 162(3):431–461, 2015.

32 MEJ Newman and Travis Martin. Equitable random graphs. *Physical Review E*, 90(5):052824, 2014.

33 Ryan O'Donnell and Xinyu Wu. Explicit near-fully X-Ramanujan graphs. In *Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science*, pages 1045–1056, 2020.

34 Svatopluk Poljak and Franz Rendl. Nonpolyhedral relaxations of graph-bisection problems. *SIAM Journal on Optimization*, 5(3):467–487, 1995.

35 Prasad Raghavendra. *Approximating NP-hard problems: efficient algorithms and their limits*. PhD thesis, University of Washington, 2009.

36 Omer Reingold, Salil Vadhan, and Avi Wigderson. Entropy waves, the zig-zag graph product, and new constant-degree expanders. *Ann. of Math. (2)*, 155(1):157–187, 2002. `doi:10.2307/3062153`.

37 Franz Rendl and Henry Wolkowicz. A projection technique for partitioning the nodes of a graph. *Ann. Oper. Res.*, 58:155–179, 1995. Applied mathematical programming and modeling, II (APMOD 93) (Budapest, 1993). `doi:10.1007/BF02032130`.

38 George Styan. Hadamard products and multivariate statistical analysis. *Linear Algebra Appl.*, 6:217–240, 1973. `doi:10.1016/0024-3795(73)90023-2`.

39 Andreas Winter. Coding theorem and strong converse for quantum channels. *Institute of Electrical and Electronics Engineers. Transactions on Information Theory*, 45(7):2481–2485, 1999. `doi:10.1109/18.796385`.