

ARTICLE

Spectral gap in random bipartite biregular graphs and applications

Gerandy Brito¹, Ioana Dumitriu² and Kameron Decker Harris^{3,*}

¹College of Computing, Georgia Institute of Technology, Atlanta, Georgia 30332, USA, ²Department of Mathematics, University of California San Diego, La Jolla, California 92093, USA and ³Department of Computer Science, Western Washington University, Bellingham, Washington 98225, USA

*Corresponding author. Email: kameron.harris@wwu.edu

(Received 16 May 2020; revised 10 December 2020; accepted 9 April 2021; first published online 23 July 2021)

Abstract

We prove an analogue of Alon’s spectral gap conjecture for random bipartite, biregular graphs. We use the Ihara–Bass formula to connect the non-backtracking spectrum to that of the adjacency matrix, employing the moment method to show there exists a spectral gap for the non-backtracking matrix. A by-product of our main theorem is that random rectangular zero-one matrices with fixed row and column sums are full rank with high probability. Finally, we illustrate applications to community detection, coding theory, and deterministic matrix completion.

Keywords: Spectral gap; expanders; random graphs; bipartite

2020 MSC Codes: Primary 15A18; Secondary 60B20, 05C50

1. Introduction

Random regular graphs, where each vertex has the same degree d , are among the most well-known examples of *expanders*: graphs with high connectivity and which exhibit rapid mixing. Expanders are of particular interest in computer science, from sampling and complexity theory to design of error-correcting codes. For an extensive review of their applications, see [35]. What makes random regular graphs particularly interesting expanders is the fact that they exhibit all three existing types of expansion properties: edge, vertex, and spectral.

The study of regular random graphs took off with the work of [7, 8, 10], and slightly later [52] and [64]. Most often, their expanding properties are described in terms of the existence of the *spectral gap*, which we define below.

Let A be the adjacency matrix of a simple graph, where $A_{ij} = 1$ if i and j are connected and zero otherwise. Denote $\sigma(A) = \{\lambda_1 \geq \lambda_2 \geq \dots\}$ as its spectrum. For a random d -regular graph, $\lambda_1 = \max_i |\lambda_i| = d$, but the second largest eigenvalue $\eta = \max(|\lambda_2|, |\lambda_n|)$ is asymptotically almost surely of much smaller order, leading to a spectral gap. Note that we will always use η to be the second largest eigenvalue of the adjacency matrix A . For a list of important symbols see Appendix A.

Spectral expansion properties of a graph are, strictly speaking, defined with respect to the smallest nonzero eigenvalue of the normalised Laplacian, $\mathcal{L} = I - D^{-1/2}AD^{-1/2}$, where I is the identity and D is the diagonal matrix of vertex degrees. In the case of a d -regular graph, $\sigma(\mathcal{L})$ is a scaled and shifted version of $\sigma(A)$. Thus, a spectral gap for A translates directly into one for \mathcal{L} .

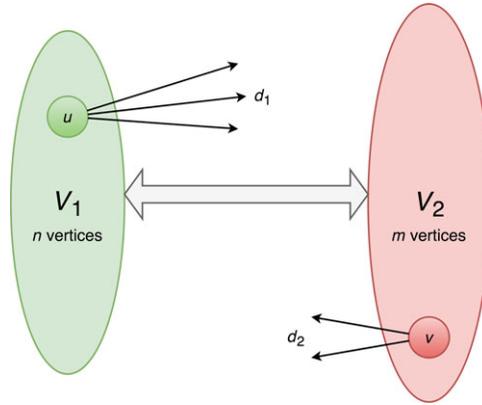


Figure 1. The structure of a bipartite, biregular graph. There are $n = |V_1|$ left vertices, $m = |V_2|$ right vertices, each of degree d_1 and d_2 , with the constraint that $nd_1 = md_2$. The distribution $\mathcal{G}(n, m, d_1, d_2)$ is taken uniformly over all such graphs.

The study of the second largest eigenvalue in regular graphs had a first breakthrough in the Alon-Boppana bound [1], which states that the second largest eigenvalue satisfies

$$\eta \geq 2\sqrt{d-1} - \frac{c_d}{\log n}.$$

Graphs for which the Alon-Boppana bound is attained are called *Ramanujan*. Friedman [27] proved the conjecture of [1] that almost all d -regular graphs have $\eta \leq 2\sqrt{d-1} + \epsilon$ for any $\epsilon > 0$ with high probability as the number of vertices goes to infinity. This result was simultaneously simplified and deepened in [29]. More recently, [12] gave a different proof that $\eta \leq 2\sqrt{d-1} + \epsilon_n$ for a sequence $\epsilon_n \rightarrow 0$ as n , the number of vertices tends to infinity; the new proof is based on the non-backtracking operator and the Ihara-Bass identity.

1.1. Bipartite biregular model

In this paper, we prove the analog of Friedman and Bordenave’s result for bipartite, biregular random graphs. These are graphs for which the vertex set partitions into two independent sets V_1 and V_2 , such that all edges occur between the sets. In addition, all vertices in set V_i have the same degree d_i . See Figure 1 for a schematic of such a graph. Along the way, we also bound the smallest positive eigenvalue and the rank of the adjacency matrix.

Let $\mathcal{G}(n, m, d_1, d_2)$ be the uniform distribution of simple, bipartite, biregular random graphs. Any $G \sim \mathcal{G}(n, m, d_1, d_2)$ is sampled uniformly from the set of simple bipartite graphs with vertex set $V = V_1 \cup V_2$, with $|V_1| = n$, $|V_2| = m$ and where every vertex in V_i has degree d_i . Note that we must have $nd_1 = md_2 = |E|$. Without any loss of generality, we will assume $n \leq m$ and thus $d_1 \geq d_2$ when necessary. Sometimes we will write that G is a (d_1, d_2) -regular graph, when we want to explicitly state the degrees. Let X be the $n \times m$ matrix with entries $X_{ij} = 1$ if and only if there is an edge between vertices $i \in V_1$ and $j \in V_2$. Using the block form of the adjacency matrix

$$A = \begin{pmatrix} 0 & X \\ X^* & 0 \end{pmatrix}, \tag{1}$$

It is well known that $\mathcal{G}(n, m, d_1, d_2)$ is connected with high probability, as long as $d_i \geq 3$. From (1), it can be verified that all eigenvalues of A occur in pairs λ and $-\lambda$, where $|\lambda|$ is a singular value of X , along with at least $|n - m|$ zero eigenvalues. For these reasons, the second largest eigenvalue is

$\eta = \lambda_2(A) = -\lambda_{n+m-1}(A)$. Furthermore, the leading or Perron eigenvalue of A is always $\sqrt{d_1 d_2}$, matched to the left by $-\sqrt{d_1 d_2}$, which reduces to the result for d -regular when $d_1 = d_2$.

We will focus on the spectrum of the adjacency matrix. Similar to the case of the d -regular graph, in the bipartite, biregular graph, the spectrum of the normalized Laplacian is a scaled and shifted version of the adjacency matrix: Because of the structure of the graph, $D^{-1/2}AD^{-1/2} = \frac{1}{\sqrt{d_1 d_2}}A$. Therefore, a spectral gap for A again implies that one exists for \mathcal{L} .

Previous work on bipartite, biregular graphs includes the work of [25] and [43], who proved the analog of the Alon–Boppana bound. For every $\epsilon > 0$,

$$\eta \geq \sqrt{d_1 - 1} + \sqrt{d_2 - 1} - \epsilon \tag{2}$$

as the number of vertices goes to infinity. This bound also follows immediately from the fact that the second largest eigenvalue cannot be asymptotically smaller than the right limit of the asymptotic support for the eigenvalue distribution, which is $\sqrt{d_1 - 1} + \sqrt{d_2 - 1}$ and was first computed by [32]. They found the spectral measure $\mu(\lambda)$ has a point mass at $\lambda = 0$ of size $\frac{1}{2}|d_1 - d_2|/(d_1 + d_2)$ and a continuous part given by the density

$$d\mu(\lambda) = \frac{d_1 d_2 \sqrt{(-\lambda^2 + d_1 d_2 - (z - 1)^2)(\lambda^2 - d_1 d_2 + (z + 1)^2)}}{\pi(d_1 + d_2)(d_1 d_2 - \lambda^2)|\lambda|}, \tag{3}$$

supported on

$$|\sqrt{d_1 - 1} - \sqrt{d_2 - 1}| \leq |\lambda| \leq \sqrt{d_1 - 1} + \sqrt{d_2 - 1}, \text{ where } z = \sqrt{(d_1 - 1)(d_2 - 1)}.$$

Graphs where η attains the Alon–Boppana bound, equation (2), are also called Ramanujan. Complete graphs are always Ramanujan but not sparse, whereas d -regular or bipartite (d_1, d_2) -regular graphs are sparse. Our results show that almost every (d_1, d_2) -regular graph is ‘almost’ Ramanujan.

Beyond the first two eigenvalues, we should mention that [13] studied the limiting spectral distribution of large sparse graphs. They obtained a set of two coupled equations that can be solved for the eigenvalue distribution of any (d_1, d_2) -regular random graph. The solution of the coupled equations for fixed d_1 and d_2 shows convergence of the spectral distribution of a random regular bipartite graph to the Marčenko–Pastur law. This was first observed by [32]. For $d_1, d_2 \rightarrow \infty$ with d_1/d_2 converging to a constant, [24] showed that the limiting spectral distribution converges to a transformed version of the Marčenko–Pastur law. When $d_1 = d_2 = d$, this is equal to the Kesten–McKay distribution [51], which becomes the semicircular law as $d \rightarrow \infty$ [32, 24]. Notably, [53] obtained the same results when they calculated the asymptotic distribution of eigenvalues for bipartite, biregular graphs of high girth. However, their results are not applicable to random bipartite biregular graphs as these asymptotically almost surely have low girth [24].

Our techniques borrow heavily from the results of [14] and [12], who simplified the trace method of [27] by counting non-backtracking walks built up of segments with at most one cycle, and by relating the eigenvalues of the adjacency matrix to the eigenvalues of the non-backtracking one via the Ihara–Bass identity. The combinatorial methods we use to bound the number of such walks are similar to how [15] counted self-avoiding walks in the context of community recovery in a regular stochastic block model.

Finally, we should mention that similar techniques have been employed by [20] to study the spectral gap of the Markov matrix of a random directed multigraph. The non-backtracking operator of a bipartite biregular graph could be seen as the adjacency matrix of a directed multigraph, whose eigenvalues are a simple scaling away from the eigenvalues of the Markov matrix of the same. However, the block structure of our non-backtracking matrix means that the corresponding multigraph is bipartite, and this makes it different from the model used in [20].

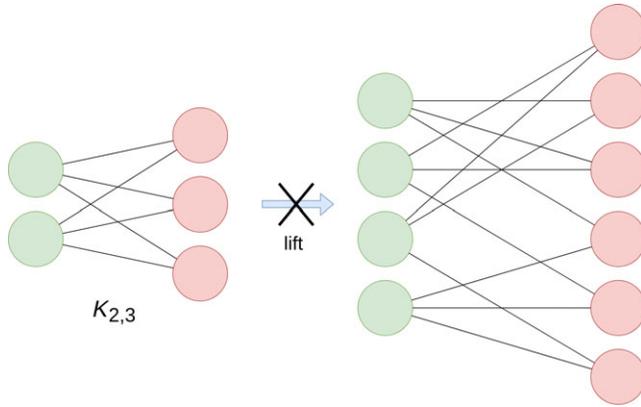


Figure 2. The (3,2)-bipartite, biregular graph on the right is not a 2-lift of $K_{2,3}$. Every pair of left vertices shares a neighbour on the right.

1.2. Configuration versus random lift model

Random lifts are a model that allows the construction of large, random graphs by repeatedly lifting the vertices of a base graph and permuting the endpoints of copied edges. See [12] for a recent overview. A number of spectral gap results have been obtained for random lift models, e.g. [27, 2, 29, 12].

Random lift models are contiguous with the configuration model in very particular cases. See Section 4.1 for a definition of the configuration model; this is a useful substitute for the uniform model and is practically equivalent. For even d , random n -lifts of a single vertex with $d/2$ self-loops are equivalent to the d -regular configuration model. For odd d , no equivalent lift construction is known or even believed to exist.

For (d_1, d_2) -biregular, bipartite graphs, the situation is more complicated. A celebrated result due to [48] showed the existence of infinite families of (d_1, d_2) -regular bipartite graphs that are Ramanujan. That is, with $\eta = \sqrt{d_1 - 1} + \sqrt{d_2 - 1}$ by taking repeated lifts of the complete bipartite graph on d_1 left and d_2 right vertices K_{d_1, d_2} . If $d_1 = d_2 = d$, then the configuration model is contiguous to the random lift of the multigraph with two vertices and d edges connecting them. Certainly, for a biregular bipartite graph with $n/d_2 = m/d_1 = k$ not an integer, we cannot construct it by lifting K_{d_1, d_2} as considered by [48]. But even for k integer, it seems likely the two models are not contiguous, for the reasons we now explain.

Suppose there were a base graph G that could be lifted to produce any (3,2)-biregular, bipartite graph. Consider another graph H which is a union of 2 complete bipartite graphs $K_{2,3}$. Then H is a (3,2)-biregular, bipartite graph and occurs in the configuration model with nonzero probability. The only G that H could be a lift of is $K_{2,3}$, because it is a disconnected union and $K_{2,3}$ itself is not a lift of any graph or multigraph (note that $2 + 3 = 5$ is prime). Therefore, G would have to be $K_{2,3}$. Figure 2 shows an example of another graph H' with the same number of vertices as H which is (3,2)-biregular, bipartite but is *not* a lift of $K_{2,3}$. Now, H and H' both occur in the configuration model with equal, nonzero probability. Therefore, we cannot construct *every* example of a (3,2)-biregular, bipartite graph by repeatedly lifting a *single* base graph G .

Since the eventual goal of any argument based on lifts that also applies to the configuration model would have to show that almost all bipartite, biregular graphs can be obtained by lifting and are sampled asymptotically uniformly from the lift model, the above considerations suggest this argument would be highly non-trivial. We in fact doubt such an argument can be made. Intuitively, in the configuration model edges occur ‘nearly independently,’ whereas for random lifts there are strong dependencies due to the fact that many edges are not allowed; see [12].

1.3. Structure of the paper

Briefly, we now lay out the method of proof that the bipartite, biregular random graph is Ramanujan. The proof outline is given in detail in Section 5.1, after some important preliminary terms and definitions given in Section 4. The bulk of our work builds to Theorem 3.1, which is actually a bound on the second eigenvalue of the non-backtracking matrix B , as explained in Section 2. The Ramanujan bound on the second eigenvalue of A then follows as Theorem 3.2. As a side result, we find that row- and column-regular, rectangular matrices (the off-diagonal block X of the adjacency matrix in equation (1)) with aspect ratio smaller than one ($d_1 \neq d_2$) have full rank with high probability.

To find the second eigenvalue of B , we subtract from it a matrix S that is formed from the leading eigenvectors and examine the spectral norm of the ‘almost-centered’ matrix $\bar{B} = B - S$. We then proceed to use the trace method to bound the spectral norm of the matrix \bar{B}^ℓ by its trace. However, since \bar{B} is not positive definite, this leads us to consider

$$\mathbb{E} \left(\|\bar{B}^\ell\|^{2k} \right) \leq \mathbb{E} \left(\text{Tr} \left((\bar{B}^\ell)(\bar{B}^\ell)^* \right)^k \right).$$

On the right-hand side, the terms in \bar{B}^ℓ refer to circuits built up of $2k$ segments, each of length $\ell + 1$, since an entry B_{ef} is a walk on two edges. Because the degrees are bounded, it turns out that, for $\ell = O(\log(n))$, the depth ℓ neighbourhoods of every vertex contain at most one cycle—they are ‘tangle-free.’ Thus, we can bound the trace by computing the expectation of the circuits that contribute, along with an upper bound on their multiplicity, taking each segment to be ℓ -tangle-free.

Finally, to demonstrate the usefulness of the spectral gap, we highlight three applications of our bound. In Section 6, we show a community detection application. Finding communities in networks is important for the areas of social network, bioinformatics, neuroscience, among others. Random graphs offer tractable models to study when detection and recovery are possible.

We show here how our results lead to community detection in regular stochastic block models with arbitrary numbers of groups, using a very general theorem by [62]. Previously, [54] studied the spectral density of such models, and the community detection problem of the special case of two groups was previously studied by [15] and [15].

In Section 7, we examine the application to linear error correcting codes built from sparse expander graphs. This concept was first introduced by [30] who explicitly used random bipartite biregular graphs. These ‘low-density parity-check’ codes enjoyed a renaissance in the 1990s, when people realised they were well suited to modern computers. For an overview, see [57, 58]. Our result yields an explicit lower bound on the minimum distance of such codes, i.e. the number of errors that can be corrected.

The final application, in Section 8, leads to generalised error bounds for matrix completion. Matrix completion is the problem of reconstructing a matrix from observations of a subset of entries. Heiman et al. [33] gave an algorithm for reconstruction of a square matrix with low complexity as measured by a norm γ_2 , which is similar to the trace norm (sum of the singular values, also called the nuclear norm or Ky Fan n -norm). The entries that are observed are at the non-zero entries of the adjacency matrix of a bipartite, biregular graph. The error of the reconstruction is bounded above by a factor which is proportional to the ratio of the leading two eigenvalues, so that a graph with larger spectral gap has a smaller generalization error. We extend their results to rectangular graphs, along the way strengthening them by a constant factor of two. The main result of the paper gives an explicit bound in terms of d_1 and d_2 .

As this paper was being prepared for submission, we became aware of the work of [22]. In their interesting paper, they use the *smallest* positive eigenvalue of a random bipartite lift to study convex relaxation techniques for random not-all-equal-3SAT problems. It seems that our main result addresses the configuration model version of this constraint satisfaction problem, the first open question listed at the end of [22].

2. Non-backtracking matrix B

Given $G \sim \mathcal{G}(n, m, d_1, d_2)$, we define the non-backtracking operator B . This operator is a linear endomorphism of $\mathbb{R}^{|\vec{E}|}$, where \vec{E} is the set of oriented edges of G and $|\vec{E}| = 2|E|$. Throughout this paper, we will use $V(H)$, $E(H)$, and $\vec{E}(H)$ to denote the vertices, edges, and oriented or directed edges of a graph, subgraph, or path H . For oriented edges $e = (u, v)$, where u and v are the starting and ending vertices of e , and $f = (s, t)$, define:

$$B_{ef} = \begin{cases} 1, & \text{if } v = s \text{ and } u \neq t; \\ 0, & \text{otherwise.} \end{cases}$$

We order the elements of \vec{E} as $\{e_1, e_2, \dots, e_{2|E|}\}$, so that the first $|E|$ have end point in the set V_2 . In this way, we can write

$$B = \begin{pmatrix} 0 & B^{(12)} \\ B^{(21)} & 0 \end{pmatrix}.$$

for $|E| \times |E|$ matrices $B^{(12)}$ and $B^{(21)}$ with entries equal to 0 or 1.

We are interested in the spectrum of B . Denote by $\mathbf{1}_\alpha$ the vector with first $|E|$ coordinates equal to 1 and the last $|E|$ equal to $\alpha = \sqrt{d_1 - 1} / \sqrt{d_2 - 1}$. We can check that

$$B\mathbf{1}_\alpha = B^* \mathbf{1}_\alpha = \lambda \mathbf{1}_\alpha$$

for $\lambda = \sqrt{(d_1 - 1)(d_2 - 1)}$. By the Perron–Frobenius Theorem, we conclude that $\lambda_1 = \lambda$ and the associated eigenspace has dimension one. Also, one can check that if λ is an eigenvalue of B with eigenvector $v = (v_1, v_2)$, $v_i \in \mathbb{R}^{|\vec{E}|}$, then $-\lambda$ is also an eigenvalue with eigenvector $v' = (-v_1, v_2)$. Thus, $\sigma(B) = -\sigma(B)$ and $\lambda_{2|E|} = -\lambda_1$.

2.1 Connecting the spectra of A and B

Understanding the spectrum of B turns out to be a challenging question. A useful result in this direction is the following theorem proved by [6], and subsequently in [63] and [40]; see also Theorem 3.3 in [3].

Theorem 2.1 (Ihara–Bass formula). *Let $G = (V, E)$ be any finite graph and B be its non-backtracking matrix. Then*

$$\det(B - \lambda I) = (\lambda^2 - 1)^{|E| - |V|} \det(D - \lambda A + \lambda^2 I),$$

where D is the diagonal matrix with $D_{vv} = d_v - 1$ and A is the adjacency matrix of G .

We use the Ihara–Bass formula to analyse the relationship of the spectrum of B to the spectrum of A in the case of a bipartite biregular graph. It will turn out that this relationship can be completely unpacked. From Theorem 2.1, we get that

$$\sigma(B) = \{\pm 1\} \cup \{\lambda : D - \lambda A + \lambda^2 I \text{ is not invertible}\}.$$

Note that there are precisely $2(m + n)$ eigenvalues of B that are determined by A , and that $\lambda = 0$ is not in the spectrum of B , since the graph has no isolated vertices ($\det(D) \neq 0$).

We use the special structure of G to get a more precise description of $\sigma(B)$. The matrices A and D are equal to:

$$A = \begin{pmatrix} 0 & X \\ X^* & 0 \end{pmatrix}, \quad D = \begin{pmatrix} (d_1 - 1)I_n & 0 \\ 0 & (d_2 - 1)I_m \end{pmatrix},$$

where I_k is the $k \times k$ identity matrix. Let $\lambda \in \sigma(B) \setminus \{-1, 1\}$. Then there exists a nonzero vector v such that

$$(D - \lambda A + \lambda^2 I)v = 0.$$

Writing $v = (v_1, v_2)$ with $v_1 \in \mathbb{C}^n, v_2 \in \mathbb{C}^m$, we obtain:

$$Xv_2 = \frac{d_1 - 1 + \lambda^2}{\lambda} v_1, \tag{4}$$

$$X^*v_1 = \frac{d_2 - 1 + \lambda^2}{\lambda} v_2. \tag{5}$$

The above imply that, provided that the right-hand side is non-zero,

$$\xi^2 = \frac{(d_1 - 1 + \lambda^2)(d_2 - 1 + \lambda^2)}{\lambda^2} \tag{6}$$

is a nonzero eigenvalue of both XX^* with eigenvector v_1 and X^*X , with eigenvector v_2 . We can rewrite equation (6) as

$$\lambda^4 - (\xi^2 - d_1 - d_2 + 2)\lambda^2 + (d_1 - 1)(d_2 - 1) = 0. \tag{7}$$

We will now detail how the eigenvalues of A (denoted ξ here) map to eigenvalues of B and vice-versa. Let us examine the special case $\xi = 0$. Assume $n \leq m$ for simplicity. Assume that the rank of X is r . Then X has $m - r$ independent vectors in its nullspace. Let u be one such vector. Now, if we pick $v_2 = u, v_1 = 0$, and $\lambda = \pm i\sqrt{d_2 - 1}$, equations (4) and (5) are satisfied. Hence, $\pm i\sqrt{d_2 - 1}$ are eigenvalues of B , both with multiplicity $m - r$.

Since the rank of X is r , it follows that the nullity of X^* is $n - r$, so there are $n - r$ independent vectors w for which $X^*w = 0$. Now, note that picking $v_1 = w, v_2 = 0$, and $\lambda = \pm i\sqrt{d_1 - 1}$, we satisfy equations (4) and (5). Thus, $\pm i\sqrt{d_1 - 1}$ are eigenvalues of B , both with multiplicity $n - r$.

The remaining $4r$ eigenvalues of B determined by A come from nonzero eigenvalues of A . For each ξ^2 with ξ a nonzero eigenvalue of A , we will have precisely 4 complex solutions to equation (6). Since there are $2r$ such eigenvalues, coming in pairs $\pm\xi$, they determine a total of $4r$ eigenvalues of B , and the count is complete. To summarise the discussion above, we have the following Lemma:

Lemma 2.2. *Any eigenvalue of B belongs to one of the following categories:*

1. ± 1 are both eigenvalues with multiplicities $|E| - |V| = nd_1 - m - n$,
2. $\pm i\sqrt{d_1 - 1}$ are eigenvalues with multiplicities $n - r$, where r is the rank of the matrix X ,
3. $\pm i\sqrt{d_2 - 1}$ are eigenvalues with multiplicities $m - r$, and
4. every pair of non-zero eigenvalues $(-\xi, \xi)$ of A generates exactly 4 eigenvalues of B .

3. Main result

We spend the bulk of this paper in the proof of the following:

Theorem 3.1. *If B is the non-backtracking matrix of a bipartite, biregular random graph $G \sim \mathcal{G}(n, m, d_1, d_2)$, then its second largest eigenvalue*

$$|\lambda_2(B)| \leq ((d_1 - 1)(d_2 - 1))^{1/4} + \epsilon_n$$

asymptotically almost surely, with $\epsilon_n \rightarrow 0$ as $n \rightarrow \infty$. Equivalently, there exists a sequence $\epsilon_n \rightarrow 0$ as $n \rightarrow \infty$ so that

$$\mathbb{P} [|\lambda_2(B)| - ((d_1 - 1)(d_2 - 1))^{1/4} > \epsilon_n] \rightarrow 0 \text{ as } n \rightarrow \infty.$$

Remark. For the random lift model, Theorem 3.1 was proved by [12], which applies to random bipartite graphs only when $d_1 = d_2 = d$ as discussed in Section 1.2.

We combine Theorems 2.1 and 3.1 to prove our main result concerning the spectrum of A .

Theorem 3.2 (Spectral gap). *Let $A = \begin{pmatrix} 0 & X \\ X^* & 0 \end{pmatrix}$ be the adjacency matrix of a bipartite, biregular random graph $G \sim \mathcal{G}(n, m, d_1, d_2)$. Without loss of generality, assume $d_1 \geq d_2$ or, equivalently, $n \leq m$. Then:*

i. *Its second largest eigenvalue $\eta = \lambda_2(A)$ satisfies*

$$\eta \leq \sqrt{d_1 - 1} + \sqrt{d_2 - 1} + \epsilon'_n$$

asymptotically almost surely, with $\epsilon'_n \rightarrow 0$ as $n \rightarrow \infty$.

ii. *Its smallest positive eigenvalue $\eta_{min}^+ = \min(\{\lambda \in \sigma(A) : \lambda > 0\})$ satisfies*

$$\eta_{min}^+ \geq \sqrt{d_1 - 1} - \sqrt{d_2 - 1} - \epsilon''_n$$

asymptotically almost surely, with $\epsilon''_n \rightarrow 0$ as $n \rightarrow \infty$. (Note that this will be almost surely positive if $d_1 > d_2$; no further information is gained if $d_1 = d_2$.)

iii. *If $d_1 \neq d_2$, the rank of X is n with high probability.*

Remark. Since the first draft of this work came out, considerable advances have been made regarding the question of singularity of random regular graphs. It was conjectured in [21] that, for $3 \leq d \leq n - 3$, the adjacency matrix of uniform d -regular graphs is not singular with high probability as n grows to infinity. For directed d -regular graphs and growing d , this is now known to be true, following the results of [19] and [45, 46]. For constant degree d , [36, 37] proved the asymptotic non-singularity of the adjacency matrix for both undirected and directed d -regular graphs. The last case can be interpreted as singularity of the adjacency matrix of random d -regular bipartite graph. To the best of our knowledge, Theorem 3.2(iii) is the first result concerning the rank of rectangular random matrices with d_1 nonzero entries in each row and d_2 in each column.

Remark. The analysis of the Ihara-Bass formula for Markov matrices of bipartite biregular graph appeared before in [39]. We have independently proven Lemma 2.2 and extracted from it more information than is given in [39], including Theorem 3.2(iii).

Proof. Equation (7), describing those eigenvalues of B which are neither ± 1 and do not correspond to 0 eigenvalues of A is equivalent to

$$0 = x^2 + \alpha\beta - (y - \alpha - \beta)x, \tag{8}$$

where $x = \lambda^2$, $y = \xi^2$, $\alpha = d_1 - 1$, and $\beta = d_2 - 1$. A simple discriminant calculation and analysis of equation (8), keeping in mind that $y \neq 0$, leads to a number of cases in terms of y :

Case 1: $y \in ((\sqrt{\alpha} - \sqrt{\beta})^2, (\sqrt{\alpha} + \sqrt{\beta})^2)$, i.e. roughly speaking, η is in the bulk, means that x is on the circle of radius $\sqrt{\alpha\beta}$ and the corresponding pair of eigenvalues λ are on a circle of radius $(\alpha\beta)^{1/4}$.

Case 2: $y \in (0, (\sqrt{\alpha} - \sqrt{\beta})^2]$ means that x is real and negative, so λ is purely imaginary.

In this case, one may also show that the smaller of the two possible values for x is increasing as a function of y and $x_- \in (-\alpha, -\sqrt{\alpha\beta}]$. The larger of the two values of

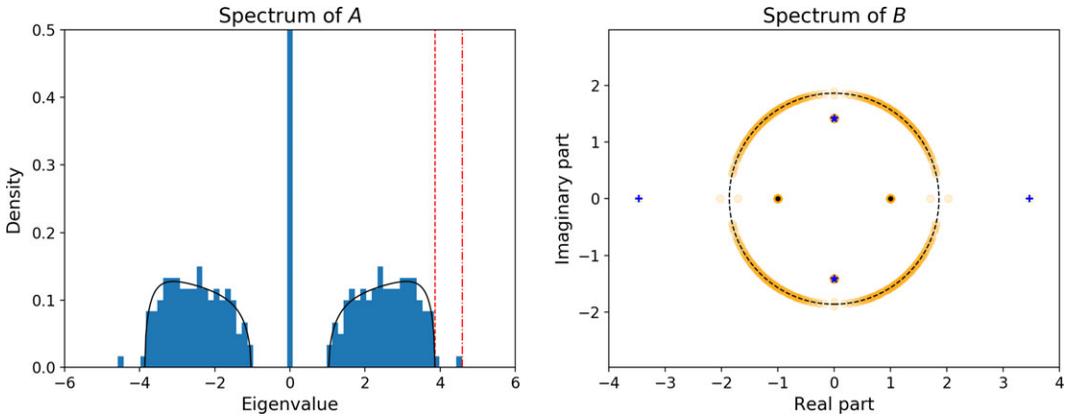


Figure 3. Example spectra for a sample graph $G \sim \mathcal{G}(120, 280, 7, 3)$. Left, we depict the spectrum of the adjacency matrix A . The dash-dotted line marks the leading eigenvalue, while the dashed line marks our bound for the second eigenvalue, Theorem (3.2). The Marčenko–Pastur limiting spectral density, equation (3), is shown in black. Right, we depict the spectrum of the non-backtracking matrix B for the same graph. Each eigenvalue is shown as a transparent orange circle, the leading eigenvalues are marked with blue crosses, and the eigenvalues arising from zero eigenvalues of A are marked with blue stars. Our main result, Theorem 3.1, proves that with high probability the non-leading eigenvalues are inside, on, or very close to the black dashed circle. In this case, there are 8 outliers of the circle, which arise from 2 pairs of eigenvalues below and above the Marčenko–Pastur bulk.

x is decreasing and $x_+ \in [-\sqrt{\alpha\beta}, -\beta)$. Correspondingly, the largest in absolute value that λ could be in this case is $\pm i\alpha^{1/4} = \pm i(d_1 - 1)^{1/4}$.

Case 3: $y \geq (\sqrt{\alpha} + \sqrt{\beta})^2$ means that both solutions x_{\pm} are real, and the larger of the two is larger than $\sqrt{\alpha\beta}$.

Note that equation (8) shows there is a continuous dependence between x and y , and consequently between ξ and λ . Putting these cases together with Lemma 2.2, a few things become apparent:

1. $\xi > \sqrt{d_1 - 1} + \sqrt{d_2 - 1}$ means that $\lambda > ((d_1 - 1)(d_2 - 1))^{1/4}$.
2. $|\lambda_2| \leq ((d_1 - 1)(d_2 - 1))^{1/4} + \epsilon$ implies that all eigenvalues except for the largest two will be either 0, or in a small neighbourhood $[\sqrt{\alpha} - \sqrt{\beta} - \delta, \sqrt{\alpha} + \sqrt{\beta} + \delta]$ of the bulk, with δ small if ϵ is small since the dependence of δ on ϵ can be deduced from equation (8).
3. $|\lambda_2| \leq ((d_1 - 1)(d_2 - 1))^{1/4} + \epsilon$ with high probability implies that if $d_1 \neq d_2$, $r = n$ with high probability. Otherwise, we would have eigenvalues of B with absolute value $\sqrt{d_1 - 1}$ and this is larger than $((d_1 - 1)(d_2 - 1))^{1/4}$.

This completes the proof, with results (i) and (ii) following from the point 2 and (iii) following from point 3. □

In Figure 3, we depict the spectra of A and B for a sample graph $G \sim \mathcal{G}(120, 280, 7, 3)$. Looking at the non-backtracking spectrum, we observe the two leading eigenvalues $\pm\sqrt{(d_1 - 1)(d_2 - 1)}$ (blue crosses) outside the circle of radius $((d_1 - 1)(d_2 - 1))^{1/4}$ along with a number of zero eigenvalues (black dots). There are also multiple purely imaginary eigenvalues which can arise from $|\xi| \in (0, \sqrt{d_1 - 1} - \sqrt{d_2 - 1}]$ as well as $\xi = 0$. However, due to Theorem 3.2, only the smaller of $i\sqrt{d_1 - 1}$ and $i\sqrt{d_2 - 1}$ is observed with non-negligible probability, implying that X has rank $r = n$ with high probability (shown as blue stars). Furthermore, we observe two pairs of real eigenvalues of B which are connected to a pair of eigenvalues of A from ‘above’ the bulk, as well as two pairs of imaginary eigenvalues of B which are connected to a pair of eigenvalues of A from ‘below’ the bulk.

4. Preliminaries

We describe the standard configuration model for constructing such graphs. We then define the ‘tangle-free’ property of random graphs. Since small enough neighbourhoods are tangle-free with high probability, we only need to count tangle-free paths when we eventually employ the trace method.

4.1. The configuration model

The configuration or permutation model is a practical procedure to sample random graphs with a given degree distribution. Let us recall its definition for bipartite biregular graphs. Let $V_1 = \{v_1, v_2, \dots, v_n\}$ and $V_2 = \{w_1, w_2, \dots, w_m\}$ be the vertices of the graph. We define the set of half edges out of V_1 to be the collection of ordered pairs

$$\vec{E}_1 = \{(v_i, j) \text{ for } 1 \leq i \leq n \text{ and } 1 \leq j \leq d_1\}$$

and analogously the set of half edges out of V_2 :

$$\vec{E}_2 = \{(w_i, j) \text{ for } 1 \leq i \leq m \text{ and } 1 \leq j \leq d_2\},$$

see Figure 1. Note that $|\vec{E}_1| = |\vec{E}_2| = nd_1 = md_2 = |E|$. To sample a graph, we choose a random permutation π of $[nd_1]$. We put an edge between v_i and w_j in the G whenever

$$\pi((i - 1)d_1 + s) = (j - 1)d_2 + t$$

for any pair of values $1 \leq s \leq d_1, 1 \leq t \leq d_2$. For specific half edges $e = (v_i, j)$ and $f = (w_s, t)$, we use the notation $\pi(e) = f$ as shorthand for $\pi((i - 1)d_1 + j) = (s - 1)d_2 + t$ and say that “ e matches to f .”

The graph obtained may not be simple, since multiple half edges may be matched between any pair of vertices. However, conditioning on a simple graph outcome, the distribution is uniform in the set of all simple bipartite biregular graphs. Furthermore, for fixed d_1, d_2 and $n, m \rightarrow \infty$, the probability of getting a simple graph is bounded away from zero [11].

Consider the random $B \in \mathbb{R}^{2|E| \times 2|E|}$ whose first $|E|$ rows are indexed by the elements of \vec{E}_1 and the last $|E|$ rows are indexed by those of \vec{E}_2 , in lexicographic order. Columns are indexed in the same way. Entry B_{ef} with $e = (v_i, j) \in \vec{E}_1$ and $f = (w_s, t) \in \vec{E}_2$ is defined as

$$B_{ef} = \begin{cases} 1, & \text{if } \pi(e) = f' = (w_s, t') \text{ and } t' \neq t; \\ 0, & \text{otherwise.} \end{cases}$$

This defines the upper half of B . We define the lower half similarly, by putting

$$B_{fe} = \begin{cases} 1, & \text{if } \pi^{-1}(f) = e' = (v_i, j') \text{ and } j' \neq j; \\ 0, & \text{otherwise.} \end{cases}$$

This is the same definition used in [12]. In words, it says that the directed edge given by e followed by the directed edge given by f are connected by some half edge $f' = (w_s, t')$, and the path they form does not backtrack. This is therefore the same matrix introduced in Section 2, ordered according to the half edges. Notice that the randomness comes from the matching only.

We consider two symmetric matrices $M = M(\pi)$ and N , indexed the same as B , and defined by:

$$M_{ef} = 1_{\{\pi(e)=f\}} \text{ and } M_{fe} = 1_{\{\pi^{-1}(f)=e\}} \text{ for } e \in \vec{E}_1 \text{ and } f \in \vec{E}_2$$

and

$$N_{gh} = 1_{\{u=v \text{ and } i \neq j\}} \text{ for } g = (u, i), h = (v, j) \in \vec{E}_1 \cup \vec{E}_2.$$

We see that a term like $M_{eg}N_{gf}$ corresponds to M matching the directed edge e to g by π , and N taking us out of the vertex of g along the directed edge f , which is different from g . Thus, the rule of matrix multiplication means that

$$B = MN. \tag{9}$$

This equality will be useful in Section 5.2 when working with products of the matrix B .

4.2. Tangle-free paths

Sparse random graphs, including bipartite graphs, have the important property of being ‘tree-like’ in the neighbourhood of a typical vertex. Formally, consider a vertex $v \in V_1 \cup V_2$. For a natural number ℓ , we define the ball of radius ℓ centered at v to be:

$$B_\ell(v) = \{w \in V_1 \cup V_2 : d_G(v, w) \leq \ell\}$$

where $d_G(\cdot, \cdot)$ is the graph distance.

Definition 4.1. A graph G is ℓ -tangle-free if $B_\ell(v)$ contains at most one cycle for any vertex v .

The next lemma says that most bipartite biregular graphs are ℓ -tangle-free up to logarithmic sized neighbourhoods.

Lemma 4.2. Let $G \sim \mathcal{G}(n, m, d_1, d_2)$ be a bipartite, biregular random graph. Let $\ell < \frac{1}{8} \log_d(n)$, for $d = \max\{d_1, d_2\}$. Then G is ℓ -tangle-free with probability at least $1 - n^{-1/2}$.

Proof. This is essentially the proof given in [47], Lemma 2.1. Fix a vertex v . We will use the so called *exploration process* to discover the ball $B_\ell(v)$. More precisely, we order the set \vec{E}_1 lexicographically: $(v_i, j) < (v_{i'}, j')$ if $i \leq i'$ and $j \leq j'$. The exploration process reveals π one edge at the time, by doing the following:

- A uniform element is chosen from \vec{E}_2 , and it is declared equal to $\pi(1)$.
- A second element is chosen uniformly, now from the set $\vec{E}_2 \setminus \{\pi(1)\}$ and set equal to $\pi(2)$.
- Once we have determined $\pi(i)$ for $i \leq k$, we set $\pi(k + 1)$ equal to a uniform element sampled from the set $\vec{E}_2 \setminus \{\pi(1), \pi(2), \dots, \pi(k)\}$.

We use the final π to output a graph as we did in the configuration model. The law of these graphs is the same. With the exploration process, we expose first the neighbours of v , then the neighbours of these vertices, and so on. This breadth-first search reveals all vertices in $B_k(v)$ before any vertices in $B_{j>k}(v)$. Note that, although our bound is for the family $\mathcal{G}(n, m, d_1, d_2)$, the neighbourhood sizes are bounded above by those of the d -regular graph with $d = \max(d_1, d_2)$.

Consider the matching of half edges attached to vertices in the ball $B_i(v)$ at depth i (thus revealing vertices at depth $i + 1$). In this process, we match a maximum $m_i \leq d^{i+1}$ pairs of half edges total. Let $\mathcal{F}_{i,k}$ be the filtration generated by matching up to the k th half edge in $B_i(v)$, for $1 \leq k \leq m_i$. Denote by $A_{i,k}$ the event that the k th matching creates a cycle at the current depth. For this to happen, the matched vertex must have appeared among the $k - 1$ vertices already revealed at depth $i + 1$. The number of unmatched half edges is at least $nd - 2d^{i+1}$. We then have that:

$$\mathbb{P}(A_{i,k}) \leq \frac{(k - 1)(d - 1)}{nd - 2d^{i+1}} \leq \frac{(d - 1)m_i}{(1 - 2d^{i+1}n^{-1})nd} \leq \frac{m_i}{n}.$$

So, we can stochastically dominate the sum

$$\sum_{i=1}^{\ell-1} \sum_{k=1}^{m_i} A_{i,k}$$

by $Z \sim \text{Bin}(d^{\ell+1}, n^{-1}d^\ell)$. So the probability that $B_\ell(v)$ is ℓ -tangle-free has the bound:

$$\mathbb{P}(B_\ell(v) \text{ is not } \ell\text{-tangle-free}) = \mathbb{P}\left(\sum_{i=1}^{\ell-1} \sum_{k=1}^{m_i} A_{i,k} > 1\right) \leq \mathbb{P}(Z > 1) = O\left(\frac{d^{4\ell+1}}{n^2}\right) = O(n^{-3/2}),$$

which follows using that $\ell = c \log_d n$ with $c < 1/8$. The Lemma follows by taking a union bound over all vertices. □

5. Proof of Theorem 3.1

5.1. Outline

We are now prepared to explain the main result. To study the second largest eigenvalue of the non-backtracking matrix, we examine the spectral radius of the matrix obtained by subtracting off the dominant eigenspace. We use for this:

Lemma 5.1 ([14], **Lemma 3**). *Let T and R be matrices such that $\text{Im}(T) \subset \text{Ker}(R)$, $\text{Im}(T^*) \subset \text{Ker}(R)$. Then all eigenvalues λ of $T + R$ that are not eigenvalues of T satisfy:*

$$|\lambda| \leq \max_{x \in \text{Ker}(T)} \frac{\|(T + R)x\|}{\|x\|}.$$

Throughout the text, $\|\cdot\|$ is the spectral norm for matrices and ℓ^2 -norm for vectors. Recall that the leading eigenvalues of B , in magnitude, are $\lambda_1 = \sqrt{(d_1 - 1)(d_2 - 1)}$ and $\lambda_{2|E|} = -\lambda_1$ with corresponding eigenvectors $\mathbf{1}_\alpha$ and $\mathbf{1}_{-\alpha}$. Applying Lemma 5.1 with $T = \frac{\lambda_1^\ell}{\mathbf{1}_\alpha^* \mathbf{1}_\alpha} (\mathbf{1}_\alpha \mathbf{1}_\alpha^* + (-1)^\ell \mathbf{1}_{-\alpha} \mathbf{1}_{-\alpha}^*)$ and $R = B^\ell - T$, we get that

$$\lambda_2(B) \leq \max_{\substack{x \in \text{Ker}(T) \\ \|x\| = 1}} \left(\|B^\ell x\|\right)^{1/\ell}. \tag{10}$$

It will be important later to have a more precise description of the set $\text{Ker}(T)$. It is not hard to check that

$$\begin{aligned} \text{Ker}(T) &= \{x: \langle x, \mathbf{1}_\alpha \rangle = \langle x, \mathbf{1}_{-\alpha} \rangle = 0\} \\ &= \{(v, w) \in \mathbb{R}^{2|E|}: \langle v, \mathbf{1} \rangle = \langle w, \mathbf{1} \rangle = 0\}. \end{aligned}$$

In the last line, the vectors v, w and $\mathbf{1}$ are $|E|$ -dimensional, and $\mathbf{1}$ is the vector of all ones.

In order to use equation (10), we must bound $\|B^\ell x\|$ for large powers ℓ and $x \in \text{Ker}(T)$. This amounts to counting the contributions of certain non-backtracking walks. We will use the tangle-free property in order to only count ℓ -tangle-free walks. We break up B^ℓ into two parts in Section 5.2, an ‘almost’ centered matrix \bar{B}^ℓ and the remainder $\sum_j R^{\ell,j}$ and we bound each term independently.

To compute these bounds, we need to count the contributions of many different non-backtracking walks. We will use the trace technique, so only circuits which return to the starting vertex will contribute. In Section 5.3, we compute the expected contribution of products of B along such circuits, employing a result from [12].

Section 5.4 covers the combinatorial component of the proof. The total contributions $\|B^\ell x\|$ come from many non-backtracking circuits of different flavours, depending on their number of vertices, edges, cycles, etc. Each circuit is broken up into $2k$ segments of tangle-free walks of length ℓ . We need to compute not only the expectation along the circuit, but also upper-bound the number of circuits of each flavor. We introduce an injective encoding of such circuits that depends on

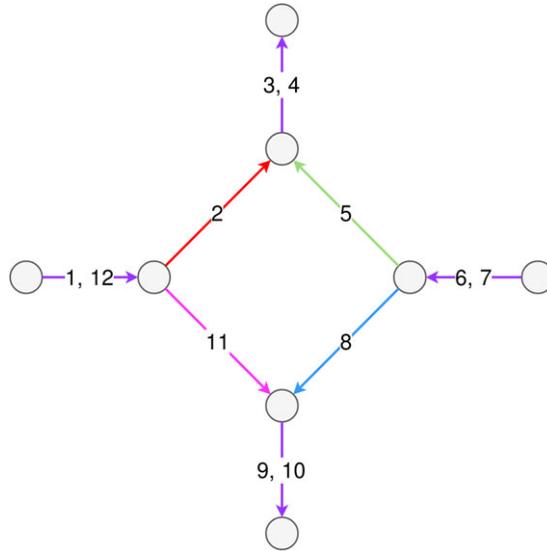


Figure 4. An example circuit that contributes to the trace in equation (11), for $k = 2$ and $\ell = 2$. Edges are numbered as they occur in the circuit. Each segment $\{\gamma_i\}_{i=1}^4$ is of length $\ell + 1 = 3$ and made up of edges $3(i - 1) + 1$ through $3i$. The last edge of each γ_i is the first edge of γ_{i+1} , and these are shown in purple. Every path γ_i with i even follows the edges backwards due to the matrix transpose. However, this detail turns out not to make any difference since the underlying graph is undirected. Our example has no cycles in each segment for clarity, but, in general, each segment can have up to one cycle, and the overall circuit may be tangled.

the number of vertices, length of the circuit, and, crucially, the tree excess of the circuit. An important part of these calculations is to keep track of the imbalance between left and right vertices visited in the circuit, since this controls the powers of d_1 and d_2 in the result.

Finally, in Section 5.8, we put all of these ingredients together and use Markov’s inequality to bound each matrix norm with high probability. We find that $\|\bar{B}^\ell\|$ contributes a factor that goes as $((d_1 - 1)(d_2 - 1))^{\ell/4}$, whereas $\|R^{\ell,j}\|$ contributes only a factor of $(d - 1)^\ell/n$, up to polylogarithmic factors in n . Thus, the main contribution to the circuit counts comes from the mean and, in fact, comes from circuits which are exactly trees traversed forwards and backwards. Interestingly, this is analogous to what happens when using the trace method on random matrices of independent entries.

In the proof, we are forced to consider tangled paths but which are built up of tangle-free components. This delicate issue was first made clear by [28] who introduced the idea of tangles and a ‘selective trace.’ Bordenave et al. [14], who we follow closely in this part of our analysis, also has a good discussion of these issues and their history. We use the fact that

$$\mathbb{E} \left(\|\bar{B}^\ell\|^{2k} \right) \leq \mathbb{E} \left(\text{Tr} \left[\left((\bar{B}^\ell)(\bar{B}^\ell)^* \right)^k \right] \right), \tag{11}$$

and so deal with circuits built up of $2k$ segments which are ℓ -tangle-free. Notice that the first segment comes from \bar{B}^ℓ , the second from $(\bar{B}^\ell)^*$, etc. Because of this, the directionality of the edges along each segment alternates. See Figure 4 for an illustration of a path which contributes for $k = 2$ and $\ell = 2$. Also, while each segment is ℓ -tangle-free, the overall circuit may be tangled.

5.2. Matrix decomposition

We start this section by defining the set of paths that will be relevant to bound the norm of $\|\bar{B}^\ell\|$. We closely follow [12].

Definition 5.2. Define Γ_{ef}^ℓ to be the set of all non-backtracking paths of $2\ell + 1$ half edges, starting at e and ending at f . A path in this set will be denoted by $\gamma = (e_1, e_2, \dots, e_{2\ell}, e_{2\ell+1})$, where $e_1 = e$ and $e_{2\ell+1} = f$. The non-backtracking property means that, for all $1 \leq i \leq \ell$, e_{2i} and e_{2i+1} share the same vertex but $e_{2i} \neq e_{2i+1}$. Similarly, let $\Gamma^\ell = \bigcup_{e,f} \Gamma_{ef}^\ell$.

Each path in Γ_{ef}^ℓ uses $2\ell + 1$ half edges, corresponding to $\ell + 1$ edges in the graph. To be clear, the above definition counts all possible non-backtracking sequences of half edges. These are different than the usual non-backtracking paths *and do not necessarily exist in the graph*. Some of these paths might backtrack along a duplicate edge which utilizes a different half edge.

We now have

$$(B^\ell)_{ef} = \sum_{\gamma \in \Gamma_{ef}^\ell} \prod_{t=1}^{\ell} B_{e_{2t-1}e_{2t+1}} = \sum_{\gamma \in \Gamma_{ef}^\ell} \prod_{t=1}^{\ell} M_{e_{2t-1}e_{2t}} N_{e_{2t}e_{2t+1}} = \sum_{\gamma \in \Gamma_{ef}^\ell} \prod_{t=1}^{\ell} M_{e_{2t-1}e_{2t}} \tag{12}$$

where we used equation (9) and the fact that Γ_{ef}^ℓ is non-backtracking, so $N_{e_{2t}e_{2t+1}} = 1$.

Recall that we will use equation (10) and Lemma 5.1 to bound λ_2 . Denote by \bar{B} the matrix with entries equal to $\bar{B} = B - S$, where

$$S = \frac{1}{|E|} \begin{pmatrix} 0 & (d_2 - 1)\mathbf{1}\mathbf{1}^* \\ (d_1 - 1)\mathbf{1}\mathbf{1}^* & 0 \end{pmatrix}.$$

Note that \bar{B} is an *almost* centred version of B , and $\text{Ker}(S) = \text{Ker}(T) = \text{span}(\mathbf{1}_\alpha, \mathbf{1}_{-\alpha})$, where T is the matrix from Lemma 5.1. To apply the lemma, we wish to get an expression like equation (12) for \bar{B}^ℓ . To do so, we write:

$$\bar{B} = B - S = MN - S = (M - S')N$$

the above matrix equation in the unknown S' can be solved by simple manipulations. We get

$$S' = \frac{1}{|E|} \begin{pmatrix} 0 & \mathbf{1}\mathbf{1}^* \\ \mathbf{1}\mathbf{1}^* & 0 \end{pmatrix}.$$

Using again that N is identically one over the elements of the set Γ_{ef}^ℓ , we find a similar formula to equation (12):

$$(\bar{B}^\ell)_{ef} = \sum_{\gamma \in \Gamma_{ef}^\ell} \prod_{t=1}^{\ell} (B - S)_{e_{2t-1}e_{2t+1}} = \sum_{\gamma \in \Gamma_{ef}^\ell} \prod_{t=1}^{\ell} \bar{M}_{e_{2t-1}e_{2t}}, \tag{13}$$

where $\bar{M} = M - S'$.

The following telescoping sum formula is a simple algebraic manipulation and appears in [50] and [14]:

$$\prod_{s=1}^{\ell} x_s = \prod_{s=1}^{\ell} y_s + \sum_{j=1}^{\ell} \prod_{s=1}^{j-1} y_s (x_j - y_j) \prod_{t=j+1}^{\ell} x_t.$$

Using this, with $x_s = B_{e_{2s-1}e_{2s+1}}$ and $y_s = \bar{B}_{e_{2s-1}e_{2s+1}}$, we obtain the following relation:

$$(B^\ell)_{ef} = (\bar{B}^\ell)_{ef} + \sum_{\gamma \in \Gamma_{ef}^\ell} \sum_{j=1}^{\ell} \prod_{s=1}^{j-1} \bar{B}_{e_{2s-1}e_{2s+1}} S_{e_{2j-1}e_{2j+1}} \prod_{t=j+1}^{\ell} B_{e_{2t-1}e_{2t+1}}. \tag{14}$$

This decomposition breaks the elements in Γ_{ef}^ℓ into two subpaths, also non-backtracking, of length j and $\ell - j$, respectively.

Definition 5.3. Let $F_{ef}^\ell \subset \Gamma_{ef}^\ell$ denote the subset of paths which are tangle-free, with $F^\ell = \bigcup_{e,f} F_{ef}^\ell$.

We will take the parameter ℓ to be small enough so that the path γ is tangle-free with high probability. Thus, the sums in equations (12) or (13) need only be over the paths $\gamma \in F_{ef}^\ell$. However, to recover the matrices B and \bar{B} by rearranging equation (14), we need to also count those tangle-free subpaths that arise from splitting tangled paths. While breaking a tangle-free path will necessarily give us two new tangle-free subpaths, the converse is not always true. This extra term generates a remainder that we define now.

Definition 5.4. Let $T_{ef}^{\ell,j}$ be the set of non-backtracking paths containing $2\ell + 1$ half edges, starting at e and ending at f , such that overall the path is tangled but the first $2j - 1$, middle three, and last $2(\ell - j) + 1$ half edges form tangle-free subpaths: $\gamma = (e_1, \dots, e_{2\ell+1}) \in T_{ef}^{\ell,j}$ if and only if $\gamma' = (e_1, \dots, e_{2j-1}) \in F^{j-1}$, $\gamma'' = (e_{2j-1}, e_{2j}, e_{2j+1}) \in F^1$, and $\gamma''' = (e_{2j+1}, \dots, e_{2\ell+1}) \in F^{\ell-j}$. Set $T_{ef}^{\ell,j} = \bigcup_{e,f} T_{ef}^{\ell,j}$.

Set the remainder

$$\begin{aligned}
 R_{ef}^{\ell,j} &= \sum_{\gamma \in T_{ef}^{\ell,j}} \prod_{j=1}^{\ell} \prod_{s=1}^{j-1} \bar{B}_{e_{2s-1}e_{2s+1}} S_{e_{2j-1}e_{2j+1}} \prod_{t=j+1}^{\ell} B_{e_{2t-1}e_{2t+1}} \\
 &= \sum_{\gamma \in T_{ef}^{\ell,j}} \prod_{j=1}^{\ell} \prod_{s=1}^{j-1} \bar{M}_{e_{2s-1}e_{2s}} S_{e_{2j-1}e_{2j+1}} \prod_{t=j+1}^{\ell} M_{e_{2t-1}e_{2t}}.
 \end{aligned}
 \tag{15}$$

Since the paths are non-backtracking, the N terms are all unity.

Adding and subtracting $\sum_{j=1}^{\ell} R_{ef}^{\ell,j}$ to equation (14) and rearranging the sums, we obtain

$$B^{(\ell)} = \bar{B}^{(\ell)} + \sum_{j=1}^{\ell} \bar{B}^{(j)} S B^{(\ell-j)} - \sum_{j=1}^{\ell} R_{ef}^{\ell,j},
 \tag{16}$$

where the matrices $B^{(\ell)}$ and $\bar{B}^{(\ell)}$ are tangle-free versions of B^ℓ and \bar{B}^ℓ , i.e. element ef in both matrices only counts paths $\gamma \in F_{ef}^\ell$. Multiplying equation (16) on the right by $x \in \text{Ker}(T)$ and using that $B^{(\ell-j)}x$ is also within $\text{Ker}(S)$, since it is just the space spanned by the leading eigenvectors, we find that the middle term is identically zero. Thus, for $x \in \text{Ker}(T)$,

$$\|B^{(\ell)}x\| \leq \|\bar{B}^{(\ell)}x\| + \left\| \sum_{j=1}^{\ell} R_{ef}^{\ell,j}x \right\| \leq \|\bar{B}^{(\ell)}\| + \sum_{j=1}^{\ell} \|R_{ef}^{\ell,j}\|.
 \tag{17}$$

5.3 Expectation bounds

Our goal is to find a bound on the expectation of certain random variables which are products of \bar{B}_{ef} along a circuit. To do this, we will need to bound the probabilities of different subgraphs when exploring G . This requires us to introduce the concept of consistent edges and their multiplicity.

Definition 5.5. Let $\gamma = (e_1, \dots, e_{2k})$ be a sequence of half edges of even length, with $\vec{E}(\gamma)$ its set of half edges, and $E(\gamma) = \{\{e_{2i-1}, e_{2i}\} \text{ for } i \in [k]\}$ its set of edges (unordered pairs, thus undirected).

- The multiplicity of a half edge $e \in \vec{E}(\gamma)$ is $m_\gamma(e) = \sum_{t=1}^{2k} 1_{\{e_t=e\}}$.
- The multiplicity of an edge $\{h_1, h_2\} \in E(\gamma)$, is $m_\gamma(\{h_1, h_2\}) = \sum_{t=1}^k 1_{\{\{e_{2t-1}, e_{2t}\}=\{h_1, h_2\}\}}$.
- An edge $\{h_1, h_2\}$ is consistent if $m_\gamma(h_1) = m_\gamma(h_2) = m_\gamma(\{h_1, h_2\})$.

Lemma 5.6. *Let $\gamma = (e_1, \dots, e_{2k})$ be a sequence of half edges of even length, with M and \bar{M} the matching matrix and its centred version generated by a uniform matching in the configuration model. Then for $1 \leq k \leq \sqrt{|E|}$ and $0 \leq t_0 \leq k$ we have that*

$$\left| \mathbb{E} \prod_{t=1}^{t_0} \bar{M}_{e_{2t-1}e_{2t}} \prod_{t=t_0+1}^k M_{e_{2t-1}e_{2t}} \right| \leq C \cdot 2^b \cdot \left(\frac{1}{|E|}\right)^\mathcal{E} \left(\frac{3k}{\sqrt{|E|}}\right)^{\mathcal{E}_1}$$

where $b =$ number of inconsistent edges of multiplicity one occurring before t_0 , $\mathcal{E}_1 =$ number of consistent edges with multiplicity one occurring before t_0 , $\mathcal{E} = |E(\gamma)|$, and C is a universal constant.

Proof. Recall the form of the matrices

$$M = \begin{pmatrix} 0 & M_1 \\ M_1^* & 0 \end{pmatrix} \quad \text{and} \quad \bar{M} = M - \frac{1}{|E|} \begin{pmatrix} 0 & \mathbf{1}\mathbf{1}^* \\ \mathbf{1}\mathbf{1}^* & 0 \end{pmatrix}.$$

Matrix $M_1 \in \mathbb{R}^{|E| \times |E|}$ is a random permutation matrix between $nd_1 = |E|$ and $md_2 = |E|$ half edges. Therefore, M_1 is distributed exactly the same as a matching matrix of a random $|E|$ -lift of a single edge, and the same holds for its centered version $M_1 - \frac{1}{|E|}\mathbf{1}\mathbf{1}^*$. The only paths γ that contribute in this bipartite setting must alternate between the bipartite sets and avoid the 0 blocks, otherwise the bound holds trivially. For one of these paths γ assume, without loss of generality, that the path starts in set V_1 . Then define the transformed path $\gamma' = (e'_1, \dots, e'_{2k}) = (e_1, e_2, e_4, e_3, e_5, \dots)$, i.e. with every other pair in γ in reverse order. Note that

$$\prod_{t=1}^{t_0} \bar{M}_{e_{2t-1}e_{2t}} \prod_{t=t_0+1}^k M_{e_{2t-1}e_{2t}} = \prod_{t=1}^{t_0} (\bar{M}_1)_{e'_{2t-1}e'_{2t}} \prod_{t=t_0+1}^k (M_1)_{e'_{2t-1}e'_{2t}}. \tag{18}$$

Then the Lemma holds by [12], Proposition 28. □

5.4. Path counting

This section is devoted to counting the number of ways non-backtracking walks can be concatenated to obtain a circuit as in Section 5.2. We will follow closely the combinatorial analysis used in [16]. In that paper, the authors needed a similar count for self-avoiding walks. We make the necessary adjustments to our current scenario.

Our goal is to find a reasonable bound for the number of circuits that contribute to the trace bound, equation (11) and shown graphically in Figure 4. Define $\mathcal{C}_{\mathcal{V}, \mathcal{E}}^{\mathcal{R}}$ as those circuits which visit exactly $\mathcal{V} = |V(\gamma)|$ different vertices, $\mathcal{R} = |V(\gamma) \cap V_2|$ of them in the right set, and $\mathcal{E} = |E(\gamma)|$ different edges. Note, these are undirected edges in $E(G)$. This is a set of circuits of length $2k\ell$ obtained as the concatenation of $2k$ non-backtracking, tangle-free walks of length ℓ . We denote such a circuit as $\gamma = (\gamma_1, \gamma_2, \dots, \gamma_{2k})$, where each γ_j is a length ℓ walk.

To bound $\mathcal{C}_{\mathcal{V}, \mathcal{E}}^{\mathcal{R}} = |\mathcal{C}_{\mathcal{V}, \mathcal{E}}^{\mathcal{R}}|$, we will first choose the set of vertices and order them. The circuits that contribute are indeed directed non-backtracking walks. However, by considering undirected walks along a fixed ordering of vertices, that ordering sets the orientation of the first and thus the rest of the directed edges in γ . Thus, we are counting the directed walks which contribute to equation (11). We relabel the vertices as $1, 2, \dots, \mathcal{V}$ as they appear in γ . Denote by \mathcal{T}_γ the

spanning tree of those edges leading to new vertices as induced by the path γ . The enumeration of the vertices tells us how we traverse the circuit and thus defines \mathcal{T}_γ uniquely.

We encode each walk γ_j by dividing it into sequences of subpaths of three types, which in our convention *must always occur* as type 1 \rightarrow type 2 \rightarrow type 3, although some may be empty subpaths. Each type of subpath is encoded with a number, and we use the encoding to upper bound the number of such paths that can occur. Given our current position on the circuit, i.e. the label of the current vertex, and the subtree of \mathcal{T}_γ already discovered (over the whole circuit γ not just the current walk γ_j), we define the types and their encodings:

- Type 1: These are paths with the property that all of their edges are edges of \mathcal{T}_γ and have been traversed already in the circuit. These paths can be encoded by their end vertex. Because this is a path contained in a tree, there is a unique path connecting its initial and final vertex. We use 0 if the path is empty.
- Type 2: These are paths with all of their edges in \mathcal{T}_γ but which are traversed for the first time in the circuit. We can encode these paths by their length, since they are traversing new edges, and we know in what order the vertices are discovered. We use 0 if the path is empty.
- Type 3: These paths are simply a single edge, not belonging to \mathcal{T}_γ , that connects the end of a path of type 1 or 2 to a vertex that has been already discovered. Given our position on the circuit, we can encode an edge by its final vertex. Again, we use 0 if the path is empty.

Now, we decompose γ_j into an ordered sequence of triples to encode its subpaths:

$$(p_1, q_1, r_1)(p_2, q_2, r_2) \cdots (p_t, q_t, r_t),$$

where each p_i characterises subpaths of type 1, q_i characterises subpaths of type 2 and r_i characterizes subpaths of type 3. These subpaths occur in the order given by the triples. We perform this decomposition using the minimal possible number of triples.

Now, p_i and r_i are both numbers in $\{0, 1, \dots, \mathcal{V}\}$, since our cycle has \mathcal{V} vertices. On the other hand, $q_i \in \{0, 1, \dots, \ell\}$ since it represents the length of a subpath of a non-backtracking walk of length ℓ . Hence, there are $(\mathcal{V} + 1)^2(\ell + 1)$ possible triples. Next, we want to bound how many of these triples occur in γ_j . We will use the following lemma.

Lemma 5.7. *Let $(p_1, q_1, r_1)(p_2, q_2, r_2) \cdots (p_t, q_t, r_t)$ be a minimal encoding of a non-backtracking walk γ_j , as described above. Then $r_i = 0$ can only occur in the last triple $i = t$.*

Proof. We can check this case by case. Assume that for some $i < t$ we have $(p_i, q_i, 0)$ and consider the concatenation with $(p_{i+1}, q_{i+1}, r_{i+1})$. First, notice that both p_{i+1} and q_{i+1} cannot be zero since then we will have $(p_i, q_i, 0)(0, 0, v^*)$ which can be written as (p_i, q_i, v^*) . If $q_i \neq 0$, then we must have $p_{i+1} \neq 0$. Otherwise, we split a path of new edges (type 2), and the decomposition is not minimal. This implies that we visit new edges and move to edges already visited, hence we need to go through a type 3 edge, implying that $r_i \neq 0$. Finally, if $p_i \neq 0$ and $q_i = 0$, then we must have $p_{i+1} = 0$; otherwise, we split a path of old edges (type 1). We also require $q_{i+1} \neq 0$, but $(p_i, 0, 0)(0, q_{i+1}, r_{i+1})$ is the same as (p_i, q_{i+1}, r_{i+1}) , which contradicts the minimality condition. This covers all possibilities and finishes the proof. \square

Using the lemma, any encoding of a non-backtracking walk γ_j has at most one triple with $r_i = 0$. All other triples indicate the traversing of a type 3 edge. We now give a very rough upper bound for how many of such encodings there can be. To do so, we will use the tangle-free property and slightly modify the encoding of the paths with cycles. Consider the two cases:

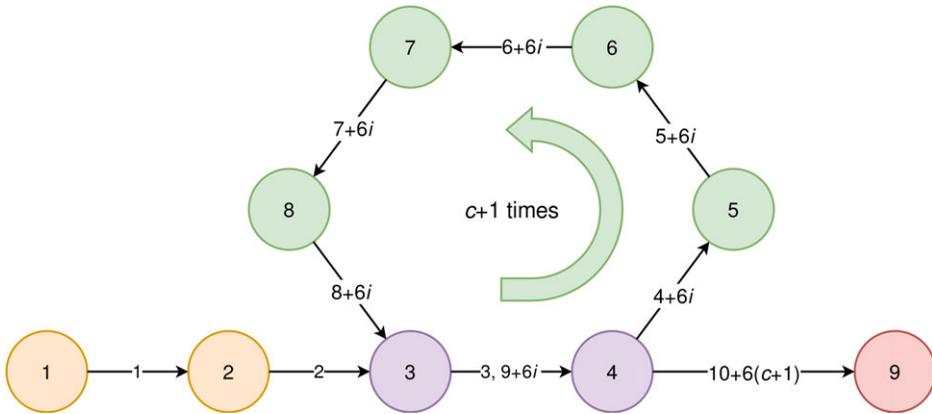


Figure 5. Encoding an ℓ -tangle-free walk, in this case the first walk in the circuit γ_1 , when it contains a cycle. The vertices and edges are labeled in the order of their traversal. The segments γ^a, γ^b , and γ^c occur on edges numbered $(1, 2, 3); (4 + 6i, 5 + 6i, 6 + 6i, 7 + 6i, 8 + 6i, 9 + 6i)$ for $i = 0, 1, \dots, c$; and $(10 + 6c)$, respectively. The encoding is $(0, 3, 0)|(0, 4, 3)(4, 0, 0)|(0, 1, 0)$. Suppose $c = 1$. Then $\ell = 22$ and the encoding is of length $3 + (4 + 1 + 1)(c + 1) + 1$, we can back out c to find that the cycle is repeated twice. The encodings become more complicated later in the circuit as vertices see repeat visits.

Case 1: Path γ_j contains no cycle. This implies that we traverse each edge within γ_j once. Thus, we can have at most $\chi = \mathcal{E} - \mathcal{V} + 1$ many triples with $r_i \neq 0$. This gives a total of at most

$$((\mathcal{V} + 1)^2(\ell + 1))^{x+1}$$

many ways to encode one of these paths.

Case 2: Path γ_j contains a cycle. Since we are dealing with non-backtracking, tangle-free walks, we enter the cycle once, loop around some number of times, and never come back. We change the encoding of such paths slightly. Let γ_j^a, γ_j^b , and γ_j^c be the segments of the path before, during, and after the cycle. We mark the start of the cycle with $|$ and its end with $\|$. The new encoding of the path is

$$(p_1^a, q_1^a, r_1^a) \cdots (p_{t^a}^a, q_{t^a}^a, r_{t^a}^a) | (p_1^b, q_1^b, r_1^b) \cdots (p_{t^b}^b, q_{t^b}^b, r_{t^b}^b) \| (p_1^c, q_1^c, r_1^c) \cdots (p_{t^c}^c, q_{t^c}^c, r_{t^c}^c),$$

where we encode the segments separately. Observe that each a subpath is connected and self-avoiding. The above encoding tells us all we need to traverse γ_j , including how many times to loop around the cycle: since the total length is ℓ , we can back out the number of circuits around the cycle from the lengths of γ_j^a, γ_j^b , and γ_j^c , see Figure 5. Following the analysis made for Case 1, the subpaths $\gamma_j^a, \gamma_j^b, \gamma_j^c$ are encoded by at most $\chi + 1$ triples, but we also have at most ℓ choices each for our marks $|$ and $\|$. We are left with at most

$$\ell^2 ((\mathcal{V} + 1)^2(\ell + 1))^{x+1}$$

ways to encode any path of this kind.

Together, these two cases mean there are less than $2\ell^2 ((\mathcal{V} + 1)^2(\ell + 1))^{x+1}$ such paths.

Now we conclude by encoding the entire circuit $\gamma = (\gamma_1, \dots, \gamma_{2k})$. We first choose \mathcal{V} vertices, \mathcal{R} in the set V_2 , and order them, which can occur in $(m)_{\mathcal{R}}(n)_{\mathcal{V}-\mathcal{R}} \leq m^{\mathcal{R}}n^{\mathcal{V}-\mathcal{R}}$ different ways. Finally, in the whole path γ we are counting concatenations of $2k$ paths which are ℓ -tangle-free. Therefore, we conclude with the following Lemma:

Lemma 5.8. Let $C_{\mathcal{V},\mathcal{E}}^{\mathcal{R}}$ be the set of circuits $\gamma = (\gamma_1, \dots, \gamma_{2k})$ of length $2k\ell$ obtained as the concatenation of $2k$ non-backtracking, tangle-free walks of length ℓ , i.e. $\gamma_s \in F^\ell$ for all $s \in [2k]$, which visit exactly $\mathcal{V} = |V(\gamma)|$ different vertices, $\mathcal{R} = |V(\gamma) \cap V_2|$ of them in the right set, and $\mathcal{E} = |E(\gamma)|$ different edges. If $C_{\mathcal{V},\mathcal{E}}^{\mathcal{R}} = |C_{\mathcal{V},\mathcal{E}}^{\mathcal{R}}|$, then

$$C_{\mathcal{V},\mathcal{E}}^{\mathcal{R}} \leq m^{\mathcal{R}} n^{\mathcal{V}-\mathcal{R}} (2\ell)^{4k} ((\mathcal{V} + 1)^2(\ell + 1))^{2k(\chi+1)}, \tag{19}$$

where $\chi = \mathcal{E} - \mathcal{V} + 1$.

The circuits that contribute to the remainder term $R^{\ell,j}$ are slightly different. In this case, each length ℓ segment is an element of $T^{\ell,j}$ rather than F^ℓ . We have to slightly modify the previous argument for this case.

Lemma 5.9. Let $D_{\mathcal{V},\mathcal{E}}^{\mathcal{R}}$ be the set of circuits $\gamma = (\gamma_1, \dots, \gamma_{2k})$ of length $2k\ell$ obtained as the concatenation of $2k$ elements $\gamma_s \in T^{\ell,j}$ for $s = 1, \dots, 2k$, that visit exactly \mathcal{V} vertices, \mathcal{R} of which are in V_2 , and \mathcal{E} different edges. Then for $D_{\mathcal{V},\mathcal{E}}^{\mathcal{R}} = |D_{\mathcal{V},\mathcal{E}}^{\mathcal{R}}|$, we have

$$D_{\mathcal{V},\mathcal{E}}^{\mathcal{R}} \leq m^{\mathcal{R}} n^{\mathcal{V}-\mathcal{R}} (2\ell)^{6k} ((\mathcal{V} + 1)^2(\ell + 1))^{6k(\chi+1)}. \tag{20}$$

Proof. Since each $\gamma^s = (e_1, \dots, e_{2\ell+1}) \in T^{\ell,j}$, we have that $\gamma' = (e_1, \dots, e_{2j-1}) \in F^{j-1}$, $\gamma'' = (e_{2j-1}, e_{2j}, e_{2j+1}) \in F^1$, and $\gamma''' = (e_{2j+1}, \dots, e_{2\ell+1}) \in F^{\ell-j}$. Encoding γ' , γ'' , and γ''' as before, we have the generous upper bound of at most

$$[(2\ell)((\mathcal{V} + 1)^2(\ell + 1))^{\chi+1}]^3$$

many encodings for each γ_s . Choosing and ordering the vertices, then concatenating $2k$ of these paths gives the final result. □

5.5. Half edge isomorphism counting

We have constructed the circuits in $C_{\mathcal{V},\mathcal{E}}^{\mathcal{R}}$ and $D_{\mathcal{V},\mathcal{E}}^{\mathcal{R}}$ by choosing the vertices and edges that participate in them. However, the expectation bound applies to matchings of half-edges in the configuration model. Since there are multiple ways to configure the half edges into such a circuit, this must be taken into account in the combinatorics.

Lemma 5.10. Let $I_{\mathcal{V},\mathcal{E}}^{\mathcal{R}}$ be the number of half edge choices for the graph induced by $\gamma \in C_{\mathcal{V},\mathcal{E}}^{\mathcal{R}} \cup D_{\mathcal{V},\mathcal{E}}^{\mathcal{R}}$. Then,

$$I_{\mathcal{V},\mathcal{E}}^{\mathcal{R}} \leq d_1^{\mathcal{V}-\mathcal{R}} (d_1 - 1)^{\mathcal{E}-\mathcal{V}+\mathcal{R}} d_2^{\mathcal{R}} (d_2 - 1)^{\mathcal{E}-\mathcal{R}}. \tag{21}$$

Proof. For every left vertex v , with degree g_v on the graph induced by γ , the number of choices of half edges is $d_1(d_1 - 1) \dots (d_1 - g_v + 1) \leq d_1(d_1 - 1)^{g_v-1}$. Note that the choice of half edges are independent for the left vertices. We then get that there are $d_1^{\mathcal{V}-\mathcal{R}} (d_1 - 1)^{\mathcal{E}-\mathcal{V}+\mathcal{R}}$ many choices, where we used that the sum of all the degrees on one component of a bipartite graph equals the number of edges: $\mathcal{E} = \sum_v g_v$. Similarly, for right vertices we get $d_2^{\mathcal{R}} (d_2 - 1)^{(\mathcal{E}-\mathcal{R})}$. □

Corollary 5.11. We have that

$$I_{\mathcal{V},\mathcal{E}}^{\mathcal{R}} \leq (d_1(d_1 - 1))^{\mathcal{V}-\mathcal{R}} (d_2(d_2 - 1))^{\mathcal{R}} (d - 1)^{2(\chi-1)}. \tag{22}$$

5.6. Bounding the imbalance ψ

We focus now on the quantity defined as $\psi = \mathcal{R} - \mathcal{E}/2$. Informally, ψ captures the *imbalance* between the number of vertices on each partition of the bipartite graph visited by the circuit γ . We show that this imbalance is not too large.

Lemma 5.12. *Let $\ell < \frac{1}{32} \log_d(n)$, then $\psi \leq 16k^2$ with high probability.*

Proof. For any subgraph H define $\psi(H) = \mathcal{R}(H) - \mathcal{E}(H)/2$. We set $\psi = \psi(\gamma)$. To bound this quantity, we analyse the subgraph $\gamma_{\leq i}$, obtained by the concatenation of the first i walks in γ , i.e. the union of the graphs induced by $\gamma_1, \dots, \gamma_i$. Our choice of ℓ implies that every neighbourhood of radius 4ℓ is tangle-free with high probability. Hence, every non-backtracking walk γ_j is either a path or a path with exactly one loop. It is not hard to conclude that $\psi(\gamma_j) \leq 2$ for all j and $\psi(\gamma_{\leq 1}) = \psi(\gamma_1) \leq 2$. We now proceed inductively to add walks to our graph, one by one, as they appear on the circuit. We will upper bound the increment $\psi(\gamma_{\leq i+1}) - \psi(\gamma_{\leq i})$ by looking at how the addition of γ_{i+1} changes the imbalance.

To analyse this, consider the intersection of γ_{i+1} and each $\gamma_j, 1 \leq j \leq i$. Notice that ψ may increase only if there are vertices at which the two walks split apart. We claim that there are at most two such vertices. Assume that at v_1, v_2 and v_3 the two walks split. Then there are two disjoint cycles in the union of γ_{i+1} and γ_j , obtained by following each first from v_1 to v_2 and then from v_2 to v_3 . But this is a contradiction, since the diameter of this union is less than $2\ell < \frac{1}{8} \log(n)$, which implies that their union is tangle-free. We conclude that $\psi(\gamma_{i+1} \cup \gamma_j) \leq 8$ since there are at most two splits and each split contributes at most four to the imbalance. Then $\psi(\gamma_{\leq i+1}) \leq \psi(\gamma_{\leq i}) + 8i + 2$, which implies that $\psi(\gamma) \leq 16k^2$, as desired. \square

5.7. Bounding the inconsistent edges

We will need a bound on the number of inconsistent edges of multiplicity one, which we get in the following lemma. Recall Definition 5.5, which introduced inconsistent edges.

Lemma 5.13. *Let b_C denote the number of inconsistent edges of multiplicity one on a circuit $\gamma = (\gamma_1, \dots, \gamma_{2k})$ consisting of $2k$ non-backtracking walks of length ℓ each. It holds that $b_C \leq 4(k + \chi)$, where $\chi = \mathcal{E} - \mathcal{V} + 1$.*

Proof. Let $\{e, f\}$ be an inconsistent edge of multiplicity one, where e and f are its half edges. For inconsistency and without loss of generality, there must exist another edge $\{e, f'\}$ in γ , so that $m_\gamma(e) \neq 1$. We may assume that $\{e, f\}$ is traversed before $\{e, f'\}$. Let v be the vertex of e and consider the two possible scenarios:

- Case 1: There is no cycle containing v in γ . Then the edge $\{e, f\}$ may only be inconsistent if v is visited at the end of one of the $2k$ non-backtracking walks γ_i and $\{e, f'\}$ is at the beginning of γ_{i+1} . Hence, in this case, we have at most two inconsistent edges of multiplicity one. This yields at most $4k$ such edges.
- Case 2: There is a cycle passing through v . For each such cycle, there is an edge that does not belong to the tree T_γ (defined in Section 5.4). Furthermore, each cycle creates at most four inconsistent edges. Combining these two facts we get at most 4χ and the proof follows.

Lemma 5.14. *Consider a circuit $\gamma = (\gamma_1, \dots, \gamma_{2k})$ with $\gamma_s \in T^{\ell, j}$ for $j \in [2k]$, with γ_s decomposed as $\gamma'_s, \gamma''_s, \gamma'''_s$ as in the Definition 5.4. Let b_D denote the number of inconsistent edges of multiplicity one in the union of segments $\bigcup_{s=1}^{2k} (\gamma'_s \cup \gamma''_s)$. Then $b_D \leq 16k + 4\chi$, where $\chi = \mathcal{E} - \mathcal{V} + 1$.*

Proof. The argument is similar to the above; however, now there are $4k$ segments in $\bar{\gamma} = \bigcup_{s=1}^{2k} (\gamma'_s \cup \gamma''_s)$, counting γ'_s and γ''_s separately. As above, each of these $4k$ segments may yield

at most 2 inconsistent edges. Furthermore, the graph induced by $\bar{\gamma}$ may not be connected; let C be the number of connected components. Each edge that creates a cycle may yield at most 4 inconsistent edges, and there are at most $\mathcal{E} - \mathcal{V} + C$ non-tree edges. Then, we have that

$$b_{\mathcal{D}} \leq 8k + 4(\mathcal{E} - \mathcal{V} + C) \leq 8k + 4(\mathcal{E} - \mathcal{V} + 2k) \leq 16k + 4\chi,$$

as claimed. □

5.8. Bounds on the norm of \bar{B}^ℓ and $R^{\ell j}$.

All of the ingredients are gathered to bound the matrix norms.

Theorem 5.15. *Let $\ell = \lfloor c \log(n) \rfloor$ where $c < \frac{1}{32}$ is a universal constant. It holds that*

$$\|\bar{B}^{(\ell)}\| \leq ((d_1 - 1)(d_2 - 1))^{\ell/4} \exp(\log^{3/4} n)$$

asymptotically almost surely.

Proof. The following holds for any natural number k , but for our proof, we will take

$$k = \lfloor \log(n)^{1/3} \rfloor \quad \text{and} \quad \ell = \lfloor c \log n \rfloor \text{ for some } c < \frac{1}{32}. \tag{23}$$

We have

$$\mathbb{E} \left(\|\bar{B}^{(\ell)}\|^{2k} \right) \leq \mathbb{E} \left(\text{Tr} \left[\left(\bar{B}^{(\ell)} \bar{B}^{(\ell)*} \right)^k \right] \right) = \mathbb{E} \left(\sum_{\gamma \in \mathcal{C}} \prod_{t=1}^{2k\ell} \bar{B}_{e_{2t-1}e_{2t+1}} \right). \tag{24}$$

The sum is taken over the set \mathcal{C} of all circuits γ of length $2k\ell$, where $\gamma = (\gamma_1, \gamma_2, \dots, \gamma_{2k})$ is formed by concatenation of $2k$ tangle-free segments $\gamma_s \in F^\ell$, with the convention $e_1^{s+1} = e_{\ell+1}^s$. Again, refer to Figure 4 for clarification.

As in Section 5.4, we will break these into circuits which visit exactly $\mathcal{V} = |V(\gamma)|$ different vertices, $\mathcal{R} = |V(\gamma) \cap V_2|$ of them in the right set, and $\mathcal{E} = |E(\gamma)|$ different edges. We define three disjoint sets of circuits:

- $\mathcal{C}_1 = \{\gamma \in \mathcal{C} : \text{all edges in } \gamma \text{ are traversed at least twice}\},$
- $\mathcal{C}_2 = \{\gamma \in \mathcal{C} : \text{at least one edge in } \gamma \text{ is traversed exactly once and } \mathcal{V} \leq k\ell + 1\}, \text{ and}$
- $\mathcal{C}_3 = \{\gamma \in \mathcal{C} : \text{at least one edge in } \gamma \text{ is traversed exactly once and } \mathcal{V} > k\ell + 1\}.$

Define the quantities

$$I_j = \left| \mathbb{E} \sum_{\gamma \in \mathcal{C}_j} \prod_{t=1}^{2k\ell} \bar{B}_{e_{2t-1}e_{2t+1}} \right| \leq \sum_{\gamma \in \mathcal{C}_j} \left| \mathbb{E} \prod_{t=1}^{2k\ell} \bar{M}_{e_{2t-1}e_{2t}} \right|$$

for $j = 1, 2$ and 3 , so that (24) can be bounded as

$$\mathbb{E} \left(\|\bar{B}^{(\ell)}\|^{2k} \right) \leq I_1 + I_2 + I_3. \tag{25}$$

We will bound each term on the right-hand side above. The reason for this division is that, by Theorem 5.6, when we have any two-path traversed exactly once, the expectation of the corresponding circuit is smaller, because the matrix \bar{B} is nearly centered. We will see that the

leading order terms in equation (24) will come from circuits in \mathcal{C}_1 . From Lemmas 5.6 and 5.8 and Corollary 5.11, we get that

$$\begin{aligned}
 I_j &\leq \sum_{\mathcal{V}, \mathcal{E}, \mathcal{R}} \sum_{\gamma \in \mathcal{C}_j \cap \mathcal{C}_{\mathcal{V}, \mathcal{E}}^{\mathcal{R}}} \left| \mathbb{E} \prod_{t=1}^{2k\ell} \bar{M}_{e_{2t-1}e_{2t}} \right| \\
 &\leq \sum_{\mathcal{V}, \mathcal{E}, \mathcal{R}} C_{\mathcal{V}, \mathcal{E}}^{\mathcal{R}} I_{\mathcal{V}, \mathcal{E}}^{\mathcal{R}} \left| \mathbb{E} \prod_{t=1}^{2k\ell} \bar{M}_{e_{2t-1}e_{2t}} \right| \\
 &\leq \sum_{\mathcal{V}, \mathcal{E}, \mathcal{R}} n^{\mathcal{V}-\mathcal{R}} m^{\mathcal{R}} (2\ell)^{4k} ((\mathcal{V} + 1)^2 (\ell + 1))^{2k(\chi+1)} \\
 &\quad \cdot (d_1(d_1 - 1))^{\mathcal{V}-\mathcal{R}} (d_2(d_2 - 1))^{\mathcal{R}} (d - 1)^{2(\chi-1)} \\
 &\quad \cdot C 2^{bc} \left(\frac{1}{|E|} \right)^{\mathcal{E}} \left(\frac{6k\ell}{\sqrt{|E|}} \right)^{\mathcal{E}_1} \\
 &\leq C \sum_{\mathcal{V}, \mathcal{E}, \mathcal{R}} (2\ell)^{4k} ((\mathcal{V} + 1)^2 (\ell + 1))^{2k(\chi+1)} (d_1 - 1)^{\mathcal{V}-\mathcal{R}} (d_2 - 1)^{\mathcal{R}} (d - 1)^{2(\chi-1)} \\
 &\quad \cdot 2^{bc} \left(\frac{1}{|E|} \right)^{\mathcal{E}-\mathcal{V}} \left(\frac{6k\ell}{\sqrt{|E|}} \right)^{\mathcal{E}_1} \\
 &\leq C \sum_{\mathcal{V}, \mathcal{E}, \mathcal{R}} (2\ell)^{4k} ((\mathcal{V} + 1)^2 (\ell + 1))^{2k(\chi+1)} ((d_1 - 1)(d_2 - 1))^{\mathcal{E}/2} \\
 &\quad \cdot \left(\frac{d_2 - 1}{d_1 - 1} \right)^{\psi} (d_1 - 1)^{1-\chi} (d - 1)^{2(\chi-1)} 2^{bc} \left(\frac{1}{|E|} \right)^{\mathcal{E}-\mathcal{V}} \left(\frac{6k\ell}{\sqrt{|E|}} \right)^{\mathcal{E}_1} \\
 &\leq c_0 n \ell^{4k} c_1^{k^2} c_2^k \sum_{\mathcal{V}, \mathcal{E}} \alpha^{2\mathcal{E}} \mathcal{V} ((\mathcal{V} + 1)^2 (\ell + 1))^{2k(\chi+1)} \left(\frac{c_3}{n} \right)^{\chi} \left(\frac{c_4 k \ell}{\sqrt{n}} \right)^{\mathcal{E}_1} \tag{26}
 \end{aligned}$$

We use $C, c_0, c_1, c_2, c_3, c_4$ to denote constant terms and set $\alpha = ((d_1 - 1)(d_2 - 1))^{1/4}$. In the last line, we used Lemmas 5.12 and 5.13 to bound ψ and b_C in terms of k and χ and remove the sum over \mathcal{R} , which contains at most \mathcal{V} terms. We will use equation (26) to bound each I_j .

5.8.1 Bounding I_1

Here, $\mathcal{E}_1 = 0$ since every edge is traversed twice. We then have that $\mathcal{V} - 1 \leq \mathcal{E} \leq k\ell$. Since γ is connected, we have $1 \leq \mathcal{V} \leq k\ell + 1$. Thus, on the right-hand side of equation (26) we get

$$\begin{aligned}
 I_1 &\leq c_0 n \ell^{4k} c_1^{k^2} c_2^k \sum_{\mathcal{V}=1}^{k\ell+1} \sum_{\mathcal{E}=\mathcal{V}-1}^{k\ell} \alpha^{2\mathcal{E}} \mathcal{V} ((\mathcal{V} + 1)^2 (\ell + 1))^{2k(\chi+1)} \left(\frac{c_3}{n} \right)^{\chi} \\
 &= c_0 n \ell^{4k} c_1^{k^2} c_2^k \sum_{\mathcal{V}=1}^{k\ell+1} \alpha^{2(\mathcal{V}-1)} \mathcal{V} ((\mathcal{V} + 1)^2 (\ell + 1))^{2k} \sum_{\chi=0}^{k\ell-\mathcal{V}+1} \left(\alpha^2 ((\mathcal{V} + 1)^2 (\ell + 1))^{2k} \frac{c_3}{n} \right)^{\chi}
 \end{aligned}$$

The second sum is upper bounded by $\sum_{\chi=0}^{\infty} \left(C \frac{((\mathcal{V}+1)^2(\ell+1))^{2k}}{n} \right)^{\chi}$, where C is some constant, which we show next is bounded by a common constant for our choices of k and ℓ and all $\mathcal{V} \in [k\ell + 1]$. To see this, it will suffice to show that $((\mathcal{V} + 1)^2 (\ell + 1))^{2k} = o(n)$. But $((\mathcal{V} + 1)^2 (\ell + 1)) = O(k^2 \ell^3)$ which, for our choices of k and ℓ , yields

$$2k \log ((\mathcal{V} + 1)^2 (\ell + 1)) = O(\log (n)^{1/3} \log (\log (n))) = o(\log (n))$$

as desired.

Finally, the first summand is maximised for $\mathcal{V} = k\ell + 1$ and there are at $k\ell + 1$ many terms in that sum. Therefore, modifying the constant c_0 yields

$$I_1 \leq c'_0 n \ell^{4k} c_1^{k^2} c_2^k (k\ell + 1)^2 ((k\ell + 2)^2(\ell + 1))^{2k} \alpha^{2k\ell}. \tag{27}$$

5.8.2 Bounding I_2

Here there is at least one edge traversed exactly once, so we have $\mathcal{E} \geq \mathcal{V}$ for $\gamma \in \mathcal{C}_2$. Taking $\mathcal{E}_1 = 0$ only increases the right-hand side on equation (26); it becomes

$$\begin{aligned} I_2 &\leq c_0 n \ell^{4k} c_1^{k^2} c_2^k \sum_{\mathcal{V}=1}^{k\ell+1} \sum_{\mathcal{E}=\mathcal{V}}^{2k\ell} \alpha^{2\mathcal{E}} \mathcal{V}((\mathcal{V} + 1)^2(\ell + 1))^{2k(\chi+1)} \left(\frac{c_3}{n}\right)^\chi \\ &= c_0 n \ell^{4k} c_1^{k^2} c_2^k \sum_{\mathcal{V}=1}^{k\ell+1} \alpha^{2(\mathcal{V}-1)} \mathcal{V}((\mathcal{V} + 1)^2(\ell + 1))^{2k} \sum_{\chi=1}^{2k\ell-\mathcal{V}+1} \left(\alpha^2((\mathcal{V} + 1)^2(\ell + 1))^{2k} \frac{c_3}{n}\right)^\chi \end{aligned}$$

Notice that this last term is almost identical to the one in the bound of I_1 , except that now we start the second sum at $\chi = 1$, which leads to an extra factor of $O(((\mathcal{V} + 1)^2(\ell + 1))^{2k}/n)$. This allow us to factor out another geometric series and proceed as we did for I_1 . This yields

$$I_2 \leq c'_0 \ell^{4k} c_1^{k^2} c_2^k (k\ell + 1)^2 ((k\ell + 2)^2(\ell + 1))^{4k} \alpha^{2k\ell}, \tag{28}$$

since there are $k\ell + 1$ terms in the first sum.

5.8.3 Bounding I_3

This set will require more delicate treatment, since circuits in \mathcal{C}_3 visit potentially many vertices and edges, yet we need to keep the power of α at most $2k\ell$.

We first show that, in this case, \mathcal{E}_1 is also large. We have $\mathcal{E} \geq \mathcal{V}$, and let $\mathcal{V} = k\ell + t$. Define \mathcal{E}'_1 as the number of edges traversed once in γ , so that $\mathcal{E}'_1 = b + \mathcal{E}_1$. Since γ has length $2k\ell$, we deduce that $2(\mathcal{E} - \mathcal{E}'_1) + \mathcal{E}'_1 \leq 2k\ell$, which implies that $\mathcal{E}'_1 \geq 2t$. Finally, Lemma 5.13 yields $\mathcal{E}_1 \geq (2t - 4(\chi + k))_+$. equation (26) then gives,

$$\begin{aligned} I_3 &\leq c_0 n \ell^{4k} c_1^{k^2} c_2^k \sum_{\mathcal{V}=k\ell+1}^{2k\ell} \sum_{\mathcal{E}=\mathcal{V}}^{2k\ell} \alpha^{2\mathcal{E}} \mathcal{V}((\mathcal{V} + 1)^2(\ell + 1))^{2k(\chi+1)} \left(\frac{c_3}{n}\right)^\chi \left(\frac{(c_4 k \ell)^2}{n}\right)^{(\mathcal{V}-k\ell-2(\chi+k))_+} \\ &= c_0 n \ell^{4k} c_1^{k^2} c_2^k \sum_{t=1}^{k\ell} \sum_{\chi=1}^{k\ell-t+1} \alpha^{2(k\ell+\chi+t-1)} (k\ell + t)((k\ell + t + 1)^2(\ell + 1))^{2k(\chi+1)} \\ &\quad \cdot \left(\frac{c_3}{n}\right)^\chi \left(\frac{(c_4 k \ell)^2}{n}\right)^{(t-2(\chi+k))_+} \\ &= c_0 n \ell^{4k} c_1^{k^2} c_2^k \alpha^{2(k\ell-1)} \sum_{t=1}^{k\ell} (k\ell + t)((k\ell + t + 1)^2(\ell + 1))^{2k} \alpha^{2t} \\ &\quad \cdot \sum_{\chi=1}^{k\ell-t+1} \left(\alpha^2((k\ell + t + 1)^2(\ell + 1))^{2k} \frac{c_3}{n}\right)^\chi \left(\frac{(c_4 k \ell)^2}{n}\right)^{(t-2(\chi+k))_+}. \end{aligned}$$

To simplify our notation, we will write

$$F(k, \ell) = c_0 n \ell^{4k} c_1^{k^2} c_2^k \alpha^{2(k\ell-1)} (2k\ell)((2k\ell + 1)^2(\ell + 1))^{2k}.$$

Observe now that we can write

$$I_3 \leq F(k, \ell) \sum_{t=1}^{k\ell} \alpha^{2t} \sum_{\chi=1}^{k\ell-t+1} \left(\frac{c(k, \ell)}{n}\right)^\chi \left(\frac{(c_4 k \ell)^2}{n}\right)^{(t-2k-2\chi)_+}$$

where $c(k, \ell) = c_3 \alpha^2 ((2k\ell + 1)^2 (\ell + 1))^{2k}$. To bound the double sum on the right-hand side above, we start by removing a factor of $\frac{c(k, \ell)}{n}$, which leaves

$$I_3 \leq F(k, \ell) \frac{c(k, \ell)}{n} \sum_{t=1}^{k\ell} \alpha^{2t} \sum_{\chi=0}^{k\ell-t} \left(\frac{c(k, \ell)}{n}\right)^\chi \left(\frac{(c_4 k \ell)^2}{n}\right)^{(t-2k-2-2\chi)_+} \tag{29}$$

The n in the denominator is crucial to cancel the linear term in $F(k, \ell)$, keeping the upper bound for I_3 small. We focus on bounding the double sum. We split the sum in t in two parts.

Case 1: $t < 2k + 2$. For these values of t , we have $(t - 2k - 2 - 2\chi)_+ = 0$, hence

$$\sum_{t=1}^{2k+1} \alpha^{2t} \sum_{\chi=0}^{k\ell-t} \left(\frac{c(k, \ell)}{n}\right)^\chi \left(\frac{(c_4 k \ell)^2}{n}\right)^{(t-2k-2-2\chi)_+} = \sum_{t=1}^{2k+1} \alpha^{2t} \sum_{\chi=0}^{k\ell-t} \left(\frac{c(k, \ell)}{n}\right)^\chi = O(\alpha^{4k}) \tag{30}$$

where the last equality uses again the same geometric upper bound for the sum over χ .

Case 2: $t \geq 2k + 2$. We split the second sum, from $\chi = 0$ to $N = \lfloor t/2 - k - 1 \rfloor$ and the terms with $\chi > N$ and analyse the two separately. The first can be upper bounded by

$$\sum_{t=2k+2}^{k\ell} \alpha^{2t} \sum_{\chi=0}^N \left(\frac{c(k, \ell)}{n}\right)^\chi \left(\frac{c_4 k \ell}{\sqrt{n}}\right)^{(t/2-k-1-\chi)_+} \leq \alpha^{4k+4} \sum_{t=0}^{k\ell-2k-2} \alpha^{2t(N+1)} \left(\frac{c_4 k \ell}{\sqrt{n}}\right)^{\frac{t}{2}}.$$

The last inequality can be checked in two steps: We first factor out the power of α^{4k+4} , and then use that $\frac{c(k, \ell)}{n} \leq \frac{c_4 k \ell}{\sqrt{n}}$, which holds for large enough n , to simplify the second sum to the addition of $N + 1$ equal terms. To bound the right hand side, we one more time upper bound by a geometric series of ratio less than one to get

$$\sum_{t=2k+2}^{k\ell} \alpha^{2t} \sum_{\chi=0}^N \left(\frac{c(k, \ell)}{n}\right)^\chi \left(\frac{c_4 k \ell}{\sqrt{n}}\right)^{(t/2-k-1-\chi)_+} \leq O(k\ell \alpha^{4k}). \tag{31}$$

We are left with the terms $\chi > N$. In this case, we get $(t - 2k - 2 - 2\chi)_+ = 0$, so

$$\sum_{t=2k+2}^{k\ell} \alpha^{2t} \sum_{\chi=N+1}^{k\ell-t} \left(\frac{c(k, \ell)}{n}\right)^\chi \left(\frac{c_4 k \ell}{\sqrt{n}}\right)^{(t/2-k-1-\chi)_+} = \sum_{t=2k+2}^{k\ell} \alpha^{2t} \sum_{\chi=N+1}^{k\ell-t} \left(\frac{c(k, \ell)}{n}\right)^\chi.$$

The sum over χ is of the order of $\left(\frac{c(k, \ell)}{n}\right)^{N+1} \leq \left(\frac{c(k, \ell)}{n}\right)^{t/2-k-1}$. Substituting this into the above, we are left with

$$\sum_{t=2k+2}^{k\ell} \alpha^{2t} \sum_{\chi=N+1}^{k\ell-t} \left(\frac{c(k, \ell)}{n}\right)^\chi = C \sum_{t=2k+2}^{k\ell} \alpha^{2t} \left(\frac{c(k, \ell)}{n}\right)^{t/2-k-1}$$

for some universal constant C . After factoring α^{4k+4} and changing variables in the summation, we conclude that

$$\sum_{t=2k+2}^{k\ell} \alpha^{2t} \sum_{\chi=N+1}^{k\ell-t} \left(\frac{c(k, \ell)}{n}\right)^\chi = C \alpha^{4k+4} \sum_{t=0}^{k\ell-2k-2} \alpha^{2t} \left(\frac{c(k, \ell)}{n}\right)^{t/2} = O(\alpha^{4k}). \tag{32}$$

Using (29) and the results for case 1 (30) and case 2 (31) and (32), we conclude that

$$I_3 \leq F(k, \ell) \frac{c(k, \ell)}{n} O(k\ell\alpha^{4k}) = c'_0 \ell^{4k} c_1^{k^2} (c'_2)^k (k\ell)^2 ((2k\ell + 1)^2 (\ell + 1))^{4k} \alpha^{2k\ell} \tag{33}$$

5.8.4 Finishing the proof of Theorem 5.15

We have bounded the three pieces we need to prove the theorem. From (27), (28), and (33), with n sufficiently large, we get

$$\begin{aligned} \mathbb{E} \left(\|\bar{B}^{(\ell)}\|^{2k} \right) &\leq I_1 + I_2 + I_3 \\ &\leq \alpha^{2k\ell} n \cdot C c_1^{k^2} c_5^k \ell^{4k} (k\ell)^4 (k^2 \ell^3)^{4k} \\ &\leq \alpha^{2k\ell} n \cdot C' c_1^{(\log n)^{2/3}} c_6^{(\log n)^{1/3}} (\log n)^{4(\log n)^{1/3}} (\log n)^{16/3} ((\log n)^{11/3})^4 (\log n)^{1/3} \\ &\leq \alpha^{2k\ell} n \cdot C'' c_1^{(\log n)^{2/3}} c_6^{(\log n)^{1/3}} (\log n)^{20(\log n)^{1/3}+6} \\ &:= \alpha^{2k\ell} n \cdot f(n), \end{aligned}$$

where c_5, c_6, C, C', C'' are universal constants.

Take any $\epsilon > 0$. It can be checked that $(\log n)^{20(\log n)^{1/3}+6} = o(n^\epsilon)$, and $f(n) = o(n^\epsilon)$ as well. Let $g(n) = \exp((\log n)^{3/4})$; then $g(n) = o(n^\epsilon)$ but $g(n)^{2k} \gg n^\epsilon$. We apply Markov's inequality, so that

$$\mathbb{P} \left[\|B^{(\ell)}\| \geq \alpha^\ell g(n) \right] \leq \frac{\mathbb{E} \left(\|\bar{B}^{(\ell)}\|^{2k} \right)}{\alpha^{2k\ell} g(n)^{2k}} \leq \frac{nf(n)}{g(n)^{2k}} = o(1), \tag{34}$$

which is the statement of the theorem. □

Theorem 5.16. *Let $1 \leq j \leq \ell = \lfloor c \log(n) \rfloor$ where $c < \frac{1}{32}$ is a universal constant. Then*

$$\|R^{\ell,j}\| \leq \frac{(d-1)^\ell}{n} \exp((\log n)^{3/4})$$

asymptotically almost surely.

Proof. The proof is analogous to the proof of Theorem 5.15. Recall the definition of $R^{\ell,j}$ in equation (15). For any integer k , we have that

$$\begin{aligned} \mathbb{E} \left(\|R^{\ell,j}\|^{2k} \right) &\leq \mathbb{E} \left(\text{Tr} \left[\left((R^{\ell,j})(R^{\ell,j})^* \right)^k \right] \right) \\ &\leq \sum_{\gamma \in \mathcal{D}} \prod_{s=1}^{2k} \left| \mathbb{E} \prod_{i=1}^{j-1} \bar{B}_{e_{2i-1}^s e_{2i+1}^s} S_{e_{2j-1}^s e_{2j+1}^s} \prod_{i=j+1}^{\ell} B_{e_{2i-1}^s e_{2i+1}^s} \right| \\ &= \sum_{\gamma \in \mathcal{D}} \prod_{s=1}^{2k} \left| \mathbb{E} \prod_{i=1}^{j-1} \bar{M}_{e_{2i-1}^s e_{2i}^s} S_{e_{2j-1}^s e_{2j+1}^s} \prod_{i=j+1}^{\ell} M_{e_{2i-1}^s e_{2i}^s} \right| \end{aligned} \tag{35}$$

Now, the sum is over the set \mathcal{D} of circuits $\gamma = (\gamma_1, \gamma_2, \dots, \gamma_{2k})$ of length $2k\ell$ formed from $2k$ elements of $T^{\ell,j}$, $\gamma_s = (e_1^s, e_2^s, \dots, e_{2\ell+1}^s)$ for $s \in [2k]$, with the convention $e_1^{s+1} = e_{2\ell+1}^s$ (see Definition 5.4).

Proceeding as before, using Lemma 5.9 and Corollary 5.11, we get that

$$\begin{aligned} \mathbb{E} \left(\|R^{\ell,j}\|^{2k} \right) &\leq \sum_{\mathcal{V}, \mathcal{E}, \mathcal{R}} \sum_{\gamma \in \mathcal{D}_{\mathcal{V}, \mathcal{E}}^{\mathcal{R}}} \prod_{s=1}^{2k} \left| \mathbb{E} \prod_{i=1}^{j-1} \bar{M}_{e_{2i-1}^s e_{2i}^s} S_{e_{2j-1}^s e_{2j+1}^s} \prod_{i=j+1}^{\ell} M_{e_{2i-1}^s e_{2i}^s} \right| \\ &\leq \sum_{\mathcal{V}, \mathcal{E}, \mathcal{R}} D_{\mathcal{V}, \mathcal{E}}^{\mathcal{R}} I_{\mathcal{V}, \mathcal{E}}^{\mathcal{R}} \prod_{s=1}^{2k} \left| \mathbb{E} \prod_{i=1}^{j-1} \bar{M}_{e_{2i-1}^s e_{2i}^s} S_{e_{2j-1}^s e_{2j+1}^s} \prod_{i=j+1}^{\ell} M_{e_{2i-1}^s e_{2i}^s} \right| \\ &\leq \sum_{\mathcal{V}, \mathcal{E}, \mathcal{R}} m^{\mathcal{R}} n^{\mathcal{V}-\mathcal{R}} (2\ell)^{6k} ((\mathcal{V} + 1)^2 (\ell + 1))^{6k(\chi+1)} \\ &\quad \cdot (d_1(d_1 - 1))^{\mathcal{V}-\mathcal{R}} (d_2(d_2 - 1))^{\mathcal{R}} (d - 1)^{2(\chi-1)} \\ &\quad \cdot \prod_{s=1}^{2k} \left(\frac{d-1}{|E|} \right) C 2^{b_s} \left(\frac{1}{|E|} \right)^{\mathcal{E}_s} \left(\frac{6\ell}{\sqrt{|E|}} \right)^{\mathcal{E}_{1,s}}. \end{aligned}$$

The last inequality uses Lemma 5.6 on each path γ_s , with \mathcal{E}_s , b_s , and $\mathcal{E}_{1,s}$ defined analogous to the quantities \mathcal{E} , b , and \mathcal{E}_1 in the same lemma. Note that $S_{ef} \leq \frac{d-1}{|E|}$ for any ef and $\sum_{s=1}^{2k} b_s \leq b_{\mathcal{D}} \leq 16(k + \chi)$ by Lemma 5.14. Setting $\mathcal{E} = \sum_{s=1}^{2k} \mathcal{E}_s$, taking $\mathcal{E}_1 = \sum_{s=1}^{2k} \mathcal{E}_{1,s} \geq 0$, using $d_1, d_2 \leq d$, and combining terms, we find that

$$\begin{aligned} \mathbb{E} \left(\|R^{\ell,j}\|^{2k} \right) &\leq \sum_{\mathcal{V}, \mathcal{E}, \mathcal{R}} (2\ell)^{6k} ((\mathcal{V} + 1)^2 (\ell + 1))^{6k(\chi+1)} \\ &\quad \cdot (d_1 - 1)^{\mathcal{V}-\mathcal{R}} (d_2 - 1)^{\mathcal{R}} (d - 1)^{2(\chi-1)} \\ &\quad \cdot \left(\frac{C'}{|E|} \right)^{2k} 2^{b_{\mathcal{D}}} \left(\frac{1}{|E|} \right)^{\mathcal{E}-\mathcal{V}} \\ &\leq c_0 \ell^{6k} c_1^k n^{1-2k} \sum_{\mathcal{V}, \mathcal{E}} \mathcal{V} ((\mathcal{V} + 1)^2 (\ell + 1))^{6k(\chi+1)} (d - 1)^{\mathcal{V}} \left(\frac{c_2}{n} \right)^{\chi} \\ &\leq c_0 \ell^{6k} c_1^k n^{1-2k} \sum_{\mathcal{V}=1}^{2k\ell} \sum_{\mathcal{E}=\mathcal{V}-1}^{2k\ell} \mathcal{V} ((\mathcal{V} + 1)^2 (\ell + 1))^{6k(\chi+1)} (d - 1)^{\mathcal{V}} \left(\frac{c_2}{n} \right)^{\chi} \\ &\leq c_0 \ell^{6k} c_1^k n^{1-2k} \sum_{\mathcal{V}=1}^{2k\ell} \mathcal{V} ((\mathcal{V} + 1)^2 (\ell + 1))^{6k} (d - 1)^{\mathcal{V}} \sum_{\chi=0}^{2k\ell} \left(((\mathcal{V} + 1)^2 (\ell + 1))^{6k} \frac{c_2}{n} \right)^{\chi}, \end{aligned}$$

for some constants c_0, c_1, c_2 . Note that the important n^{-2k} comes from the $|E|^{-2k}$ deterministic term. Now, like before we have that $((\mathcal{V} + 1)^2 (\ell + 1))^{6k} = o(n)$, so for sufficiently large n we can bound the sum over χ by a constant. The sum over \mathcal{V} is easily bounded as before, leading to

$$\begin{aligned} \mathbb{E} \left(\|R^{\ell,j}\|^{2k} \right) &\leq (d - 1)^{2k\ell} n^{1-2k} \cdot c_0' \ell^{6k} c_1^k (2k\ell)^2 ((2k\ell + 1)^2 (\ell + 1))^{6k} \\ &\leq (d - 1)^{2k\ell} n^{1-2k} \cdot c_0'' c_1^{(\log n)^{1/3}} (\log n)^{6(\log n)^{1/3}} (\log n)^{8/3} ((\log n)^{11/3})^{6(\log n)^{1/3}} \\ &:= (d - 1)^{2k\ell} n^{1-2k} \cdot f(n). \end{aligned}$$

To finish the proof, we note that $f(n) = o(n^\epsilon)$ for any $\epsilon > 0$. Take $g(n) = \exp((\log n)^{3/4})$. Then, applying Markov's inequality,

$$\mathbb{P} \left[\|R^{\ell,j}\| \geq \frac{(d - 1)^\ell}{n} g(n) \right] \leq \frac{\mathbb{E} \left(\|R^{\ell,j}\|^{2k} \right) n^{2k}}{(d - 1)^{2k\ell} g(n)^{2k}} \leq \frac{nf(n)}{g(n)^{2k}} = o(1). \tag{36}$$

5.9. Proof of the main result, Theorem 3.1

We will again take $\ell = \lfloor c \log n \rfloor$, with c chosen so that the graph is ℓ -tangle-free with high probability. By equations (10) and (17),

$$|\lambda_2|^\ell \leq \|\bar{B}^{(\ell)}\| + \sum_{k=1}^{\ell} \|R^{\ell,j}\|.$$

Notice that $(d - 1)^\ell \leq (d - 1)^{c \log n} \leq n^{c \log d}$, so take $c < \min\left(\frac{1}{32}, \frac{1}{\log d}\right)$. Then $(d - 1)^\ell = O(n^\epsilon)$ for some $0 < \epsilon < 1$, and recall that $\exp((\log n)^{3/4}) = o(n^{\epsilon'})$ for any $\epsilon' > 0$. We apply Theorems 5.15 and 5.16 to get

$$\begin{aligned} |\lambda_2| &\leq \left(\exp((\log n)^{3/4}) ((d_1 - 1)(d_2 - 1))^{\ell/4} + \ell \exp((\log n)^{3/4}) \frac{(d - 1)^\ell}{n} \right)^{1/\ell} \\ &= ((d_1 - 1)(d_2 - 1))^{1/4} + \epsilon_n, \end{aligned}$$

where $\epsilon_n \rightarrow 0$ as $n \rightarrow \infty$. □

6. Application: community detection

In many cases, such as online networks, we would like to be able to recover specific communities in those graphs. In the typical setup, a community is a set of vertices that are more densely connected together than to the rest of the graph.

The model we present here is inspired by the planted partition or stochastic blockmodel (SBM, [34]). In the SBM, each vertex belongs to a class or community, and the probability that two vertices are connected is a function of the classes of the vertices. It is a generalisation of the Erdős-Rényi random graph. The classes or blocks in the SBM make it a good model for graphs with community structure, where nodes preferentially connect to other nodes depending on their communities [55].

There are many methods for detecting a community given a graph. For an overview of the topic, see [26]. Spectral clustering is a common method which can be applied to any set of data $\{\zeta_i\}_{i=1}^n$. Given a symmetric and non-negative similarity function S , the similarity is computed for every pair of data points, forming a matrix $A_{ij} = S(\zeta_i, \zeta_j) = S(\zeta_j, \zeta_i) \geq 0$. The spectral clustering technique is to compute the leading eigenvectors of A , or matrices related to it, and use the eigenvectors to cluster the data. In our case, the matrix in question is just the Markov matrix of a graph, defined soon. We will show that we can guarantee the success of the technique if the degrees are large enough.

Our graph model is a regular version of the SBM. We build it on a ‘frame,’ which is a small, weighted graph that defines the community structure present in the larger, random graph. Each class is represented by a vertex in the frame. The edge weights in the frame define the number of edges between classes. What makes our model differ from the SBM is that the connections between classes are described by a regular random graph rather than an Erdős-Rényi random graph. However, the graph itself is not necessarily regular.

A number of authors have studied similar models. Our model is a generalisation of a random lift of the frame, which is said to *cover* the random graph (e.g. [49, 3, 14]). This type of random graph was also studied by [54], who called it an equitable random graph, since the community structure is equivalent to an equitable partition. This partition induces a number of symmetries across vertices in each community which are useful when studying the eigenvalues of the graph. Barrett et al. [4] studied the effect of these symmetries from a group theory standpoint. The work of [5] is closest to ours: they consider spectral properties of such graphs and their implications for spectral community clustering. In particular, they show that the spectrum of what we call the

‘frame’ (in their words, the discrete spectrum, which is deterministic) is contained in that of the random graph. They use the resolvent method (called the cavity method in the physics community) to analyse the continuous part of the spectrum in the limit of large graph size and argue that community detection is possible when the deterministic frame eigenvalues all lie outside the bulk. However, this analysis assumes that there are no stochastic eigenvalues outside the bulk, which will only hold with high probability if the graph is Ramanujan. Our analysis shows that, if a set of pairwise spectral gaps hold between all communities, then this will be the case.

6.1. The frame model

We define the *random regular frame graph* distribution $\mathcal{G}(n, H)$ as a distribution of simple graphs on n vertices parametrised by the ‘frame’ $H = (V, E, p, D)$ is a weighted, directed graph. Here, V is the vertex set, $E \subseteq \{(i, j): i, j \in V\}$ is the directed edge set, the vertex weights are p , and the edge weights are D . Note that we drop the arrows on the edge set in this section, since it will always be directed. The vertex weight vector $p \in \mathbb{R}^{|V|}$, where $\sum_{i \in V} p_i = 1$, sets the relative sizes of the classes. The edge weights are a matrix of degrees $D \in \mathbb{N}^{|V| \times |V|}$. These assign the number of edges between each class in the random graph: D_{ij} is the number of edges from each vertex in class i to vertices in class j . The degrees must satisfy the balance condition

$$p_i D_{ij} = p_j D_{ji} \tag{37}$$

for all $i, j \in V$ where (i, j) or (j, i) are in E . This requires that, for every edge $e \in E$, its reverse orientation also exists in H . We also require that $n_i = np_i \in \mathbb{N}$ for every $i \in V$, so that the number of vertices in each type is integer.

Given the frame H , a random regular frame graph $G \sim \mathcal{G}(n, H)$ is a simple graph on n vertices with n_i vertices in class i . It is chosen uniformly among graphs with the constraint that each vertex in class i makes D_{ij} connections among the vertices in class j . In other words, if $i = j$, we sample that block of the adjacency matrix as the adjacency matrix of a D_{ii} -regular random graph on n_i vertices. For off-diagonal blocks $i \neq j$, these are sampled as bipartite, biregular random graphs $\mathcal{G}(n_i, n_j, D_{ij}, D_{ji})$.

Sampling from $\mathcal{G}(n, H)$ can be performed similar to the configuration model, where each node is assigned as many half-edges as its degree, and these are wired together with a random matching [55]. The detailed balance condition equation (37) ensures that this matching is possible. Practically, we often have to generate many candidate matchings before the resulting graph is simple, but the probability of a simple graph is bounded away from zero for fixed D .

An example of a random regular frame graph is the bipartite, biregular random graph. The family $\mathcal{G}(n, m, d_1, d_2)$ is a random regular frame graph $\mathcal{G}(n + m, H)$, where the frame H is the directed path on two vertices: $V = \{1, 2\}$ and $E = \{(1, 2), (2, 1)\}$. The weights are taken as $p_1 = n/(n + m)$, $p_2 = m/(n + m)$, $D_{12} = d_1$, and $D_{21} = d_2$.

Another example random regular frame graph is shown in Figure 6. In this case, the frame H has $V = \{A, B, C\}$ and $E = \{(A, B), (A, C), (B, A), (B, C), (C, A), (C, B)\}$ with weights p and D as shown in Figure 6A. We see that this generates a random tripartite graph with regular degrees between vertices in different independent sets, shown in Figure 6B.

6.2. Markov and related matrices of frame graphs

Now, we define a number of matrices associated with the frame and the sample of the random regular frame graph.

Let G be a simple graph. Define $D_G = \text{diag}(d_G)$, the diagonal matrix of degrees in G . The Markov matrix $P = P(G)$ is defined as

$$P = D_G^{-1}A,$$

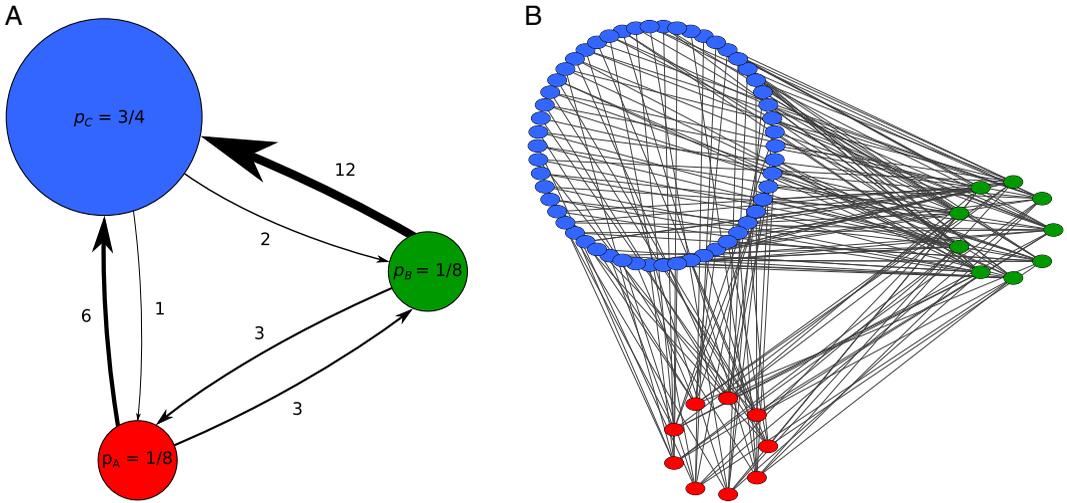


Figure 6. Schematic and realisation of a random regular frame graph. **A**, The frame graph. The vertices of the frame (red = A, green = B, blue = C) are weighted according to their proportions p in the random regular frame graph. The edge weights D_{ij} set the between-class vertex degrees in the random regular frame graph. This frame will yield a random tripartite graph. **B**, Realisation of the graph on 72 vertices. In this instance, there are $1/8 \times 72 = 9$ green and red vertices and $3/4 \times 72 = 54$ blue vertices. Each blue vertex connects to $k_{CA} = 1$ red vertex and $k_{CB} = 2$ green vertices. This is actually a multigraph; with so few vertices, the probability that the configuration model algorithm yields parallel edges is high.

where $A = A(G)$ is the adjacency matrix. The Markov matrix is the row-normalized adjacency matrix, and it contains the transition probabilities of a random walker on the graph G . Let $L = I - \mathcal{L} = D_G^{-1/2} A D_G^{-1/2}$ be a matrix simply related to the normalized Laplacian. We call this the *symmetrized Markov matrix*. Then P and L have the same eigenvalues, but L is symmetric, since $L_{ij} = \frac{A_{ij}}{\sqrt{d_i d_j}}$.

Suppose $G \sim \mathcal{G}(n, H)$, where the frame $H = (V, E, p, D)$. Another matrix that will be useful is what we call the *Markov matrix of the frame* R , where $R_{ij} = \frac{D_{ij}}{\sum_j D_{ij}}$. Thus, R is a row-normalized D , in the same way that the Markov matrix P is the row-normalized adjacency matrix A . Furthermore, R is invariant under any uniform scaling of the degrees. Because of this equitable partition property of random regular frame graphs, eigenvectors of the frame matrices $D = D(H)$ or $R = R(H)$ lift to eigenvectors of $A = A(G)$ or $P = P(G)$, respectively. Suppose $Dx = \lambda x$, then it is a straightforward exercise to check that $A\tilde{x} = \lambda\tilde{x}$ for the piecewise constant vector

$$\tilde{x} = \begin{bmatrix} \mathbf{1}_{n_1} x_1 \\ \mathbf{1}_{n_2} x_2 \\ \vdots \end{bmatrix}.$$

Using the same procedure, we can lift any eigenpair of R to an eigenpair of P with the same eigenvalue.

6.2.1 Bounds on the eigenvalues of frame graphs in terms of blocks

The following result is due to [62]:

Proposition 6.1. *Let G be a random regular frame graph $G(n, H)$, P its Markov matrix, and L the Laplacian with vertices ordered by class in both cases. Let R be the Markov matrix of the frame*

$H = (V, E, p, D)$, with $|V(H)| = K$ classes. Define the matrices $L^{(kl)}$ as the (k, l) block of L with respect to the clustering of vertices by class. For $l \neq k$, let

$$M^{(kl)} = \begin{pmatrix} 0 & L^{(kl)} \\ L^{(kl)} & 0 \end{pmatrix} = \begin{pmatrix} 0 & L^{(kl)} \\ L^{(lk)*} & 0 \end{pmatrix}.$$

For $l = k$, let $M^{(kk)} = L^{(kk)}$. Assume that all eigenvalues of D are nonzero and pick a constant C such that

$$\frac{|\lambda_2^{(kl)}|}{\lambda_1^{(kl)}} \leq C < 1$$

for every $k, l = 1, \dots, K$, where $\lambda_1^{(kl)}$ and $\lambda_2^{(kl)}$ are the leading and second eigenvalues of $M^{(kl)}$. Under these conditions, the eigenvalues of P which are not eigenvalues of R are bounded by

$$C \max_{k=1, \dots, K} \left(R_{kk} + \sum_{l \neq k} \sqrt{R_{kl}R_{lk}} \right) \leq \frac{C}{2} \left(1 + \max_{k=1, \dots, K} \sum_{l=1}^K R_{lk} \right).$$

The spectrum of the Markov matrix $\sigma(P)$ enjoys a simple connection to $\sigma(A)$ when A is the adjacency matrix of a graph drawn from $G(n, m, d_1, d_2)$. In this case, $L = \frac{A}{\sqrt{d_1 d_2}}$, so the eigenvalues of P are just the scaled eigenvalues of A . This and the spectral gap for bipartite, biregular random graphs, Theorem 3.2, lead to the following remark:

Remark. For a random regular frame graph, $M^{(kl)}$ corresponds to the symmetrized Markov matrix L of a bipartite biregular graph $G(n_k, n_l, D_{kl}, D_{lk})$. Thus,

$$\frac{|\lambda_2^{(kl)}|}{\lambda_1^{(kl)}} \leq \frac{\sqrt{D_{kl} - 1} + \sqrt{D_{lk} - 1}}{\sqrt{D_{kl}D_{lk}}} + \epsilon.$$

Suppose we are given a frame that fits the conditions of Proposition 6.1; namely, D cannot have any zero eigenvalues. Then we can uniformly grow the degrees, which leaves R invariant, but allows us to reach an arbitrarily small C . This ensures that the leading K eigenvalues of P are equal to the eigenvalues of R . Note that this actually means that the entire random regular frame graph satisfies a weak Ramanujan property. We now show that this guarantees spectral clustering.

6.3. Spectral clustering

Spectral clustering is a popular method of community detection. Because some eigenvectors of P , the Markov matrix of a random regular frame graph, are piecewise constant on classes, we can use them to recover the communities so long as those eigenvectors can be identified. Suppose there are K total classes in our random regular frame graph. Then, given the eigenvectors x^1, x^2, \dots, x^K , which are piecewise constant across classes, we can cluster vertices by class. For each vertex $v \in V(G)$, associate the vector $y^v \in \mathbb{R}^K$ where $y_j^v = x_j^v$. Then if $y^v = y^u$ for $u, v \in V(G)$, vertices u and v belong to the same class.^a It is simple to recover these piecewise constant vectors x^1, x^2, \dots, x^K when they are the leading eigenvectors. These facts lead to the following theorem:

^aIn the SBM case, the eigenvectors are not piecewise constant, but they are aligned with the eigenvectors of R and thus highly correlated across vertices in the same class. A more flexible clustering method such as K -means must be applied to the vectors y in that case.

Theorem 6.2 (Spectral clustering guarantee in frame graphs). *Let G be a random regular frame graph $G(n, H)$ and P its Markov matrix. Let R be the Markov matrix of the frame $H = (V, E, p, D)$, with $|V(H)| = K$ classes and $\lambda_1 \geq \dots \geq \lambda_K$ the eigenvalues of R and $|\lambda_K| > 0$. Then we can scale the degrees by some $\kappa \in \mathbb{N}$, $D \rightarrow \kappa D$, so that the vertex classes are recoverable by spectral clustering of the leading K eigenvectors of P .*

Remark. The conditions of Theorem 6.2, while very general, are also weaker than may be expected using more sophisticated methods tailored to the specific frame model. We illustrate this with the following example.

6.3.1 Example: The regular stochastic block model

Brito et al. [16] and Barucca [5] studied a regular stochastic block model, which can be seen as a special case of our frame model. Let the frame H be the complete directed graph on two vertices, including self loops, where

$$D = \begin{pmatrix} d_1 & d_2 \\ d_2 & d_1 \end{pmatrix}$$

and $p = (1/2, 1/2)$. Define the regular stochastic block model as $\mathcal{G}(2n, H)$. This is a graph with two classes of equal size, representing two communities of vertices, with within-class degree d_1 and between-class degree d_2 . We assume $d_1 > d_2$, since communities are more strongly connected within. Brito et al. [16] proved the following theorem:

Theorem 6.3. *If $(d_1 - d_2)^2 > 4(d_1 + d_2 - 1)$, then there is an efficient algorithm for strong recovery, i.e. recovery of the exact communities with high probability as $n \rightarrow \infty$.*

Theorem 6.3 gives a sharp bound on the degrees for recovery, which we can compare to our spectral clustering results. The eigenvalues of D are $d_1 + d_2$ and $d_1 - d_2$, and the Markov matrix of the frame R has eigenvalues 1 and $(d_1 - d_2)/(d_1 + d_2)$. The diagonal blocks $L^{(11)}$ and $L^{(22)}$ each correspond to the Laplacian matrix of a d_1 -regular random graph on n vertices, whereas the off-diagonal block term $M^{(12)}$ corresponds to the Laplacian of a d_2 -regular bipartite graph on $2n$ vertices. Using our results and the previously known results for regular random graphs [27, 28, 14], we can pick some $C > 2\sqrt{d_2 - 1}/d_2$ since $d_1 > d_2$ and we will eventually take the degrees to be large. Using Proposition 6.1, we find that the spurious eigenvalues of P come after the leading 2 eigenvalues if

$$\frac{2\sqrt{d_2 - 1}}{d_2} < \frac{d_1 - d_2}{d_1 + d_2},$$

to leading order in the degrees. Rearranging, we obtain the condition

$$(d_1 - d_2)^2 > 4(d_2 - 1) \left(\frac{d_1 + d_2}{d_2} \right)^2.$$

Assuming $d_2/d_1 = \beta < 1$ fixed, and taking the limit $d_1, d_2 \rightarrow \infty$, we find that the result of [16] becomes

$$d_1 > 4 \frac{1 + \beta}{(1 - \beta)^2} + o(1),$$

whereas our result becomes

$$d_1 > \frac{4}{\beta} \left(\frac{1 + \beta}{1 - \beta} \right)^2 + o(1),$$

illustrating that the spectral threshold is a factor of $(1 + \beta)/\beta$ weaker.

7. Application: low density parity check or expander codes

Another useful application of random graphs is as expanders, loosely defined as graphs where the neighbourhood of a small set of nodes is large. Expander codes, also called low-density parity-check (LDPC) codes, were first introduced by Gallager in his PhD thesis [30]. These are a family of linear error correcting codes whose parity-check matrix is encoded in an expander graph. A linear code is a set $\mathcal{C} \subset \Sigma^L$, where a length L codeword $x \in \mathcal{C}$ if and only if $Hx = 0$. The alphabet Σ is typically a finite field, and $H \in \Sigma^{P \times L}$ is the parity check matrix. In the simplest case, $\Sigma = \mathbb{F}_2$ and each row of H can be interpreted as a parity constraint on codewords. The performance of such codes depends on how good an expander that graph is, which in turn can be shown to depend on the separation of eigenvalues. For a good introduction and overview of the subject, see [58].

Following [61], we construct a code \mathcal{C} from a (d_1, d_2) -regular bipartite graph G on $n + m$ vertices and two smaller linear codes \mathcal{C}_1 and \mathcal{C}_2 of length d_1 and d_2 , respectively. We write $\mathcal{C}_1 = [d_1, k_1, \delta_1]$ and $\mathcal{C}_2 = [d_2, k_2, \delta_2]$ with the usual convention of length, dimension, and minimum distance. We assume that the codes are all binary, using the finite field \mathbb{F}_2 the codeword is $x \in \mathcal{C} \subset \mathbb{F}_2^{|E|}$ where $|E| = nd_1 = md_2$. That is, we associate a bit to each edge in the graph bipartite graph G . Let $(e_i(v))_{i=1}^{d_1}$ represent the set of edges incident to a vertex v in some arbitrary, fixed order. Then the vector $x \in \mathcal{C}$ if and only if the vectors $(x_{e_1(u)}, x_{e_2(u)}, \dots, x_{e_{d_1}(u)})^T \in \mathcal{C}_1$ for all $u \in V_1$ and $(x_{e_1(v)}, x_{e_2(v)}, \dots, x_{e_{d_2}(v)})^T \in \mathcal{C}_2$ for all $v \in V_2$. The final code \mathcal{C} is also linear. With this construction, the code \mathcal{C} has rate at least $k_1/d_1 + k_2/d_2 - 1$ [61].

Furthermore, [38] proved the following bound on the minimum distance of the resulting code:

Theorem 7.1. *Suppose $\delta_1 \geq \delta_2 > \eta/2$, where η is the second largest eigenvalue of the adjacency matrix of G . Then the code \mathcal{C} has minimum distance*

$$\delta \geq \frac{n}{d_2} \left(\delta_1 \delta_2 - \frac{\eta}{2} (\delta_1 + \delta_2) \right).$$

Corollary 7.2. *Suppose the code \mathcal{C} is constructed from a biregular, bipartite random graph $G \sim \mathcal{G}(n, m, d_1, d_2)$ and the conditions of Theorem 7.1 hold. Then the minimum distance of \mathcal{C} satisfies*

$$\delta \geq \frac{n}{d_2} \left(\delta_1 \delta_2 - \frac{\sqrt{d_1 - 1} + \sqrt{d_2 - 1}}{2} (\delta_1 + \delta_2) - \epsilon_n \right).$$

We see that these Tanner codes will have maximal distance for smallest η and used our main result, Theorem 3.2, to obtain the explicit bound in Corollary 7.2. By growing the graph, the above shows a way to construct arbitrarily large codes whose minimum distance remains proportional to the code size nd_1 . That is, the relative distance $\delta/(nd_1)$ is bounded away from zero as $n \rightarrow \infty$. However, the above bound will only be useful if it yields a positive result, which depends on the codes \mathcal{C}_1 and \mathcal{C}_2 as well as the degrees.

Remark. In general, the performance guarantees on LDPC codes that are obtainable from graph eigenvalues are weaker than those that come from other methods. Although our method does guarantee high distance for some high degree codes, analysis of specific decoding algorithms or a probabilistic expander analyses yield better bounds that work for lower degrees [58].

7.1 Example: An unbalanced code based on a (14,9)-regular bipartite graph

We illustrate the applicability of our distance bound with an example. Let $\mathcal{C}_1 = [14, 8, 7]$ and $\mathcal{C}_2 = [9, 4, 6]$. These can be achieved by using a Reed–Salomon code on the common field \mathbb{F}_q for any $q > 14$ [58]. We take $q = 2^4 = 16$ for inputs that are actually binary, and this means each edge in the graph actually contains 4 bits of information. Employing Corollary 7.2, the Tanner code \mathcal{C}

will have relative minimum distance $\delta/(nd_1) \geq 0.0014$ and rate at least 0.016. Taking $n = 216$ and $m = 336$ gives the code a minimum distance of at least 4.

8. Application: matrix completion

Assume that we have some matrix $Y \in \mathbb{R}^{n \times m}$ which has low ‘complexity.’ Perhaps it is low-rank or simple by some other measure. If we observe Y_{ij} for a limited set of entries $(i, j) \in E \subset [n] \times [m]$, then *matrix completion* is any method which constructs a matrix \hat{Y} so that $\|\hat{Y} - Y\|$ is small, or even zero. Matrix completion has attracted significant attention in recent years as a tractable algorithm for making recommendations to users of online systems based on the tastes of other users (a.k.a. the Netflix problem). We can think of it as the matrix version of compressed sensing [18, 17].

Recently, a number of authors have studied the performance of matrix completion algorithms where the index set E is the edge set of a regular random graph [33, 9, 31]. Heiman et al. [33] describe a *deterministic* method of matrix completion, where they can give performance guarantees for a fixed observation set E over many input matrices Y . The error of their reconstruction depends on the spectral gap of the graph. We expand upon the result of [33], extending it to rectangular matrix and improving their bounds in the process.

8.1. Matrix norms as measures of complexity and their relationships

We will employ a number of different matrix and vector norms in this section. These are all related by the properties of the underlying Banach spaces. The complexity of Y is measured using a factorisation norm (also called the max-norm):

$$\gamma_2(Y) = \min_{UV^*=Y} \|U\|_{\ell_2 \rightarrow \ell_\infty^n} \|V\|_{\ell_2 \rightarrow \ell_\infty^m}.$$

The minimum is taken over all possible factorisations of $Y = UV^*$, and the norm $\|X\|_{\ell_2 \rightarrow \ell_\infty^n} = \max_i \sqrt{\sum_j X_{ij}^2}$ returns the largest ℓ_2 norm of a row. So, equivalently,

$$\gamma_2(Y) = \min_{UV^*=Y} \max_{i,j} \|u_i\|_2 \|v_j\|_2,$$

where u_i and v_j are the rows of U and V . See [44] for a number of results about the norm γ_2 . In particular, note that

$$\gamma_2(A \circ B) \leq \gamma_2(A)\gamma_2(B) \tag{38}$$

$$\frac{1}{\sqrt{nm}} \|Y\|_{\text{Tr}} \leq \gamma_2(Y) \tag{39}$$

$$\gamma_2(Y) \leq \sqrt{\text{rank}(Y)} \|Y\|_\infty. \tag{40}$$

Property (38) says that γ_2 is sub-multiplicative under the Hadamard product [42, 23] and will be used in our proof. Properties (39) and (40) relate γ_2 to two common complexity measures of matrices, the trace norm (sum of singular values, i.e. the $\ell_2^m \rightarrow \ell_2^n$ nuclear norm) and rank. Note also the well-known fact that

$$\|Y\|_{\text{Tr}} = \min_{UV^*=Y} \|U\|_F \|V\|_F,$$

where $\|X\|_F = \sqrt{\sum_{ij} X_{ij}^2}$ is the Frobenius norm. We see that the trace norm constrains factors U and V to be small on average via $\|\cdot\|_F$, whereas the norm γ_2 is similar but constrains factors uniformly via $\|\cdot\|_{\ell_2 \rightarrow \ell_\infty}$. However, we should note that computing $\gamma_2(Y)$ is more costly than the

trace norm, which can be performed with just the singular value decomposition, although still possible in polynomial time with convex programming [33].

8.2. Matrix completion generalisation bounds

The method of matrix completion that we study is to return the matrix X which is the solution to:

$$\begin{aligned} &\underset{X}{\text{minimise}} && \gamma_2(X) \\ &\text{subject to} && X_{ij} = Y_{ij}, (i, j) \in E. \end{aligned} \tag{41}$$

The γ_2 norm was first proposed by [60, 59] as a robust complexity measure, and it was shown to be an effective and practical regularisation on real data sets [41, 56].

Heiman et al. [33] analyse the performance of the convex program (41) for a square matrix Y using an expander argument, assuming that E is the edge set of a d -regular graph with second eigenvalue η . They obtain the following theorem:

Theorem 8.1 ([33]). *Let E be the set of edges of a d -regular graph with second eigenvalue bound η . For every $Y \in \mathbb{R}^{n \times n}$, if \hat{Y} is the output of the optimization problem (38), then*

$$\frac{1}{n^2} \|\hat{Y} - Y\|_F^2 \leq c\gamma_2(Y)^2 \frac{\eta}{d},$$

where $c = 8K_G \leq 14.3$ is a universal constant and $\|\cdot\|_F$ is the Frobenius norm.

Considering sampling following the biadjacency matrix of a bipartite graph, we find a similar result which also applies to rectangular matrices. If $n = m$ and $d_1 = d_2 = d$, our bound is equivalent to that of Theorem 8.1, but with constants improved by a factor of two due to stronger mixing in bipartite graphs. Intuitively, using a biadjacency matrix is a ‘more random’ way of sampling than using an adjacency matrix, since it is not symmetric.

Theorem 8.2. *Let E be the set of edges of a (d_1, d_2) -regular graph with second eigenvalue bound η . For every $Y \in \mathbb{R}^{n \times m}$, if \hat{Y} is the output of the optimization problem (41), then*

$$\frac{1}{nm} \|\hat{Y} - Y\|_F^2 \leq c\gamma_2(Y)^2 \frac{\eta}{\sqrt{d_1 d_2}},$$

where $c = 4K_G \leq 7.13$.

Proof. We start by considering a rank-1 sign matrix $S = uv^*$, where $u \in \{-1, 1\}^n, v \in \{-1, 1\}^m$. Let $S' = \frac{1}{2}(S + J)$, where J is the all-ones matrix, so that S' has the entries of -1 in S replaced by zeros. Then $S' = 1_A 1_B^* + 1_{A^c} 1_{B^c}^*$ for subsets $A \subset V_1 = [n]$ and $B \subset V_2 = [m]$, where $A = \{i: u_i = 1\}$ and $B = \{j: v_j = 1\}$. Consider the expression

$$\begin{aligned} &\left| \frac{1}{nm} \sum_{ij} s_{ij} - \frac{1}{|E|} \sum_{(i,j) \in E} s_{ij} \right| = \left| \frac{1}{nm} \sum_{ij} (2s'_{ij} - 1) - \frac{1}{|E|} \sum_{(i,j) \in E} (2s'_{ij} - 1) \right| \\ &= 2 \left| \frac{1}{nm} \sum_{ij} s'_{ij} - \frac{1}{|E|} \sum_{(i,j) \in E} s'_{ij} \right| \\ &= 2 \left| \frac{|A||B| + |A^c||B^c|}{nm} - \frac{E(A, B) + E(A^c, B^c)}{|E|} \right| \\ &\leq 2 \left| \frac{|A||B|}{nm} - \frac{E(A, B)}{|E|} \right| + 2 \left| \frac{|A^c||B^c|}{nm} - \frac{E(A^c, B^c)}{|E|} \right|. \end{aligned}$$

The following is a bipartite version of the expander mixing lemma [23]:

$$\left| \frac{E(A, B)}{|E|} - \frac{|A||B|}{nm} \right| \leq \frac{\eta}{\sqrt{d_1 d_2}} \sqrt{\frac{|A||B|}{nm} \left(1 - \frac{|A|}{n}\right) \left(1 - \frac{|B|}{m}\right)} = \frac{\eta}{\sqrt{d_1 d_2}} \sqrt{\frac{|A||B||A^c||B^c|}{(nm)^2}}.$$

We find that

$$\begin{aligned} \left| \frac{1}{nm} \sum_{ij} s_{ij} - \frac{1}{|E|} \sum_{(i,j) \in E} s_{ij} \right| &\leq \frac{4\eta}{\sqrt{d_1 d_2}} \sqrt{\frac{|A||B||A^c||B^c|}{(nm)^2}} \\ &= \frac{4\eta}{\sqrt{d_1 d_2}} \sqrt{xy(1-x)(1-y)} \\ &\leq \frac{\eta}{\sqrt{d_1 d_2}}, \end{aligned}$$

since $xy(1-x)(1-y)$ attains a maximal value of 2^{-4} for $0 \leq x, y \leq 1$.

The rest of the proof develops identical to the results of [33], which we include for completeness. We apply the result for rank-1 sign matrices to any matrix R . Let $R = \sum_i \alpha_i S^i$, where S^i is a rank-1 sign matrix and $\alpha_i \in \mathbb{R}$. For a general matrix R , this might require many rank-1 sign matrices. Define the sign nuclear norm $\nu(R) = \sum_i |\alpha_i|$. Then,

$$\left| \frac{1}{nm} \sum_{ij} r_{ij} - \frac{1}{|E|} \sum_{(i,j) \in E} r_{ij} \right| \leq \nu(R) \frac{\eta}{\sqrt{d_1 d_2}}.$$

It is a consequence of Grothendieck’s inequality, a well-known theorem in functional analysis, that there exists a universal constant $1.5 \leq K_G \leq 1.8$ so that $\gamma_2(X) \leq \nu(X) \leq K_G \gamma_2(X)$ for any real matrix X ; see [33].

Now, let the matrix of residuals $R = (\hat{Y} - Y) \circ (\hat{Y} - Y)$, where \circ is the Hadamard entry-wise product of two matrices, so that $R_{ij} = (\hat{Y}_{ij} - Y_{ij})^2$. Since

$$\frac{1}{|E|} \sum_{(i,j) \in E} r_{ij} = 0,$$

we conclude that

$$\frac{1}{nm} \sum_{ij} r_{ij} \leq \nu(R) \frac{\eta}{\sqrt{d_1 d_2}} \leq K_G \gamma_2(R) \frac{\eta}{\sqrt{d_1 d_2}}.$$

Furthermore, $\gamma_2(R) \leq \gamma_2(\hat{Y} - Y)^2 \leq (\gamma_2(\hat{Y}) + \gamma_2(Y))^2$ by (38) and the triangle inequality. Since \hat{Y} is the output of the algorithm and Y is a feasible solution, $\gamma_2(\hat{Y}) \leq \gamma_2(Y)$. Thus, $\gamma_2(R) \leq 4\gamma_2(Y)^2$ and the proof is finished. \square

8.3. Noisy matrix completion bounds

Furthermore, our analysis easily extends to the case where the matrix we observe is corrupted with noise. As mentioned in the above remark, similar results will hold for the trace norm. In the noisy case, we solve the problem

$$\begin{aligned} &\underset{X}{\text{minimize}} && \gamma_2(X) \\ &\text{subject to} && \frac{1}{|E|} \sum_{(i,j) \in E} (X_{ij} - Z_{ij})^2 \leq \delta^2 \end{aligned} \tag{42}$$

and obtain the following theorem:

Theorem 8.3. *Suppose we observe $Z_{ij} = Y_{ij} + \epsilon_{ij}$ with bounded error*

$$\frac{1}{|E|} \sum_{(i,j) \in E} \epsilon_{ij}^2 \leq \delta^2.$$

Then solving the optimization problem (42) will yield a bound of

$$\frac{1}{nm} \|\hat{Y} - Y\|_F^2 \leq c\gamma_2(Y)^2 \frac{\eta}{\sqrt{d_1 d_2}} + 4\delta^2,$$

where $c = 4K_G \leq 7.13$.

Proof. Denote \hat{Y} the solution to P42. It will be useful to introduce the sampling operator $\mathcal{P}_E: \mathbb{R}^{n \times m} \rightarrow \mathbb{R}^{n \times m}$, where $(\mathcal{P}_E(X))_{ij} = X_{ij}$ if $(i, j) \in E$ and 0 otherwise. Again let $R = (\hat{Y} - Y) \circ (\hat{Y} - Y)$ be the matrix of squared errors, then

$$\begin{aligned} \left| \frac{1}{nm} \|\hat{Y} - Y\|_F^2 - \frac{1}{|E|} \|\mathcal{P}_E(\hat{Y} - Y)\|_F^2 \right| &= \left| \frac{1}{nm} \sum_{ij} (\hat{Y}_{ij} - Y_{ij})^2 - \frac{1}{|E|} \sum_{(i,j) \in E} (\hat{Y}_{ij} - Y_{ij})^2 \right| \\ &\leq K_G \gamma_2(R) \frac{\eta}{\sqrt{d_1 d_2}}. \end{aligned}$$

However, since Y is a feasible solution to P42, we have

$$\gamma_2(\hat{Y}) \leq \gamma_2(Y).$$

Applying (38) and the triangle inequality,

$$\gamma_2(R) \leq \left(\gamma_2(\hat{Y} - Y) \right)^2 \leq \left(\gamma_2(\hat{Y}) + \gamma_2(Y) \right)^2 \leq 4\gamma_2(Y)^2.$$

Using the triangle inequality again gives

$$\|\mathcal{P}_E(\hat{Y} - Y)\|_F \leq \|\mathcal{P}_E(\hat{Y} - Z)\|_F + \|\mathcal{P}_E(Z - Y)\|_F \leq 2\delta\sqrt{|E|},$$

taking into account the bound on the observation errors. Because

$$\frac{1}{nm} \|\hat{Y} - Y\|_F^2 \leq \left| \frac{1}{nm} \|\hat{Y} - Y\|_F^2 - \frac{1}{|E|} \|\mathcal{P}_E(\hat{Y} - Y)\|_F^2 \right| + \frac{1}{|E|} \|\mathcal{P}_E(\hat{Y} - Y)\|_F^2$$

we get the final bound

$$\frac{1}{nm} \|\hat{Y} - Y\|_F^2 \leq 4K_G \gamma_2(Y)^2 \frac{\eta}{\sqrt{d_1 d_2}} + 4\delta^2.$$

□

8.4. Application of the spectral gap

Theorem 8.2 provides a bound on the mean squared error of the approximation X . Directly applying Theorem 3.2, we obtain the following bound on the generalisation error of the algorithm using a random biregular, bipartite graph:

Corollary 8.4. *Let E be sampled from a $\mathcal{G}(n, m, d_1, d_2)$ random graph. For every $Y \in \mathbb{R}^{n \times m}$, if \hat{Y} is the output of the optimization problem (38), then*

$$\frac{1}{nm} \|\hat{Y} - Y\|_F^2 \leq c\gamma_2(Y)^2 \frac{\sqrt{d_1 - 1} + \sqrt{d_2 - 1} + \epsilon_n}{\sqrt{d_1 d_2}},$$

where $c = 4K_G \leq 7.13$ is a universal constant.

Acknowledgements

We would like to thank Marina Meilă for sharing Proposition 6.1 and for suggestions and comments. Thank you also to Pierre Youssef, Simon Coste, and Subhabrata Sen for helpful comments and connections. We are grateful to our anonymous reviewers for useful comments and suggestions, and at least in one case for pointing out errors in an earlier version of this manuscript (which have since been fixed). K. D. H. was supported by the Big Data for Genomics and Neuroscience NIH training grant, Washington Research Foundation postdoctoral fellowship, as well as NSF grants DMS-1122105 and DMS-1514743. G. B. was partially supported by NSF CAREER award DMS-1552267. I. D. was supported by NSF DMS-1712630 and NSF CAREER award DMS-0847661.

References

- [1] Alon, N. (1986) Eigenvalues and expanders. *Combinatorica* **6** (2) 83–96. ISSN 0209-9683, 1439-6912. doi: [10.1007/BF02579166](https://doi.org/10.1007/BF02579166).
- [2] Angel, O., Friedman, J. and Hoory, S. (2007) The Non-Backtracking Spectrum of the Universal Cover of a Graph. arXiv:0712.0192 [math].
- [3] Angel, O., Friedman, J. and Hoory, S. (2015) The non-backtracking spectrum of the universal cover of a graph. *Trans. Am. Math. Soc.* **367** (6) 4287–4318. ISSN 0002-9947, 1088-6850. doi: [10.1090/S0002-9947-2014-06255-7](https://doi.org/10.1090/S0002-9947-2014-06255-7).
- [4] Barrett, W., Francis, A. and Webb, B. (2017) Equitable decompositions of graphs with symmetries. *Linear Algebra Appl.* **513** (Supplement C) 409–434. ISSN 0024-3795. doi: [10.1016/j.laa.2016.10.017](https://doi.org/10.1016/j.laa.2016.10.017).
- [5] Barucca, P. (2017) Spectral partitioning in equitable graphs. *Phys. Rev. E* **95** (6) 062310. doi: [10.1103/PhysRevE.95.062310](https://doi.org/10.1103/PhysRevE.95.062310).
- [6] Bass, H. (1992) The Ihara-Selberg zeta function of a tree lattice. *Int. J. Math.*, **03** (06) 717–797. ISSN 0129-167X. doi: [10.1142/S0129167X92000357](https://doi.org/10.1142/S0129167X92000357).
- [7] Bender, E. A. (1974) The Asymptotic Number of Non-Negative Integer Matrices with Given Row and Column Sums. *Discrete Math.* **10** (2) 217–223. ISSN 0012-365X. doi: [10.1016/0012-365X\(74\)90118-6](https://doi.org/10.1016/0012-365X(74)90118-6).
- [8] Bender, E. A. and Rodney Canfield, E. (1978) The asymptotic number of labeled graphs with given degree sequences. *J. Comb. Theory Ser. A* **24** (3) 296–307. ISSN 0097-3165. doi: [10.1016/0097-3165\(78\)90059-6](https://doi.org/10.1016/0097-3165(78)90059-6).
- [9] Bhojanapalli, S. and Jain, P. (2014) Universal matrix completion. In *PMLR*, pp. 1881–1889.
- [10] Bollobás, B. (1980) A probabilistic proof of an asymptotic formula for the number of labelled regular graphs. *Eur. J. Comb.* **1** (4) 311–316. ISSN 0195-6698. doi: [10.1016/S0195-6698\(80\)80030-8](https://doi.org/10.1016/S0195-6698(80)80030-8).
- [11] Bollobás, B. (2001) *Random Graphs*. Cambridge University Press, Cambridge. ISBN 978-0-511-81406-8.
- [12] Bordenave, C. (2015) A New Proof of Friedman’s Second Eigenvalue Theorem and its Extension to Random Lifts. arXiv:1502.04482 [math].
- [13] Bordenave, C. and Lelarge, M. (2010) Resolvent of large random graphs. *Random Struct. Algorithms* **37** (3) 332–352. ISSN 10429832. doi: [10.1002/rsa.20313](https://doi.org/10.1002/rsa.20313).
- [14] Bordenave, C., Lelarge, M. and Massoulié, L. (2015) Non-Backtracking Spectrum of Random Graphs: Community Detection and Non-Regular Ramanujan Graphs. arXiv:1501.06087 [cs, math].
- [15] Brito, G., Dumitriu, I., Ganguly, S., Hoffman, C. and Tran, L. V. (2015) Recovery and Rigidity in a Regular Stochastic Block Model. arXiv:1507.00930 [math].
- [16] Brito, G., Dumitriu, I., Ganguly, S., Hoffman, C. and Tran, L. V. (2016) Recovery and rigidity in a regular stochastic block model. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, Philadelphia, PA, USA. Society for Industrial and Applied Mathematics, pp. 1589–1601. ISBN 978-1-61197-433-1.
- [17] Candès, E. J. and Plan, Y. (2010) Matrix completion with noise. *Proc. IEEE* **98** (6) 925–936. ISSN 0018-9219. doi: [10.1109/JPROC.2009.2035722](https://doi.org/10.1109/JPROC.2009.2035722).
- [18] Candès, E. J. and Tao, T. (2010) The power of convex relaxation: near-optimal matrix completion. *IEEE Trans. Inf. Theory* **56** (5) 2053–2080. ISSN 0018-9448. doi: [10.1109/TIT.2010.2044061](https://doi.org/10.1109/TIT.2010.2044061).
- [19] Cook, N. A. (2017) On the singularity of adjacency matrices for random regular digraphs. *Prob. Theory Related Fields* **167** (1–2) 143–200. ISSN 0178-8051, 1432-2064. doi: [10.1007/s00440-015-0679-8](https://doi.org/10.1007/s00440-015-0679-8).
- [20] Coste, S. (2017) The Spectral Gap of Sparse Random Digraphs. arXiv:1708.00530 [math].
- [21] Costello, K. P. and Vu, V. H. (2008) The rank of random graphs. *Random Struct. Algorithms* **33** (3) 269–285. ISSN 1098-2418. doi: [10.1002/rsa.20219](https://doi.org/10.1002/rsa.20219).
- [22] Deshpande, Y., Montanari, A., O’Donnell, R., Schramm, T. and Sen, S. (2018) The Threshold for SDP-Refutation of Random Regular NAE-3SAT. arXiv:1804.05230 [cs, math].
- [23] De Winter, S., Schillewaert, J. and Verstraete, J. (2012) Large incidence-free sets in geometries. *Electr. J. Comb.* **19** (4) P24. ISSN 1077-8926.
- [24] Dumitriu, I. and Johnson, T. (2016) The Marčenko-Pastur law for sparse random bipartite biregular graphs. *Random Struct. Algorithms* **48** (2) 313–340. ISSN 1098-2418. doi: [10.1002/rsa.20581](https://doi.org/10.1002/rsa.20581).

- [25] Feng, K. and Li, W.-C. W. (1996) Spectra of hypergraphs and applications. *J. Number Theory* **60** (1) 1–22. ISSN 0022-314X. doi: [10.1006/jnth.1996.0109](https://doi.org/10.1006/jnth.1996.0109).
- [26] Fortunato, S. (2010) Community detection in graphs. *Phys. Rep.* **486** (3–5) 75–174. ISSN 0370-1573. doi: [10.1016/j.physrep.2009.11.002](https://doi.org/10.1016/j.physrep.2009.11.002).
- [27] Friedman, J. (2003) Relative expanders or weakly relatively Ramanujan graphs. *Duke Math. J.* **118** (1) 19–35. ISSN 0012-7094, 1547-7398. doi: [10.1215/S0012-7094-03-11812-8](https://doi.org/10.1215/S0012-7094-03-11812-8).
- [28] Friedman, J. (2004) A Proof of Alon’s Second Eigenvalue Conjecture and Related Problems. arXiv:cs/0405020.
- [29] Friedman, J. and Kohler, D.-E. (2014) The Relativized Second Eigenvalue Conjecture of Alon. arXiv:1403.3462 [cs, math].
- [30] Gallager, R. (1962) Low-density parity-check codes. *IRE Trans. Inf. Theory* **8** (1) 21–28. ISSN 0096-1000. doi: [10.1109/TIT.1962.1057683](https://doi.org/10.1109/TIT.1962.1057683).
- [31] Gamarnik, D., Li, Q. and Zhang, H. (2017) Matrix Completion from $O(n)$ Samples in Linear Time. arXiv:1702.02267 [cs, math, stat].
- [32] Godsil, C. D. and Mohar, B. (1988) Walk generating functions and spectral measures of infinite graphs. *Linear Algebra Appl.* **107** (Supplement C) 191–206. ISSN 0024-3795. doi: [10.1016/0024-3795\(88\)90245-5](https://doi.org/10.1016/0024-3795(88)90245-5).
- [33] Heiman, E., Schechtman, G. and Shraibman, A. (2014) Deterministic algorithms for matrix completion. *Random Struct. Algorithms* **45** (2) 306–317. ISSN 1098-2418. doi: [10.1002/rsa.20483](https://doi.org/10.1002/rsa.20483).
- [34] Holland, P. W., Laskey, K. B. and Leinhardt, S. (1983) Stochastic blockmodels: first steps. *Soc. Networks* **5** (2) 109–137. ISSN 0378-8733. doi: [10.1016/0378-8733\(83\)90021-7](https://doi.org/10.1016/0378-8733(83)90021-7).
- [35] Hoory, S., Linial, N. and Wigderson, A. (2006) Expander graphs and their applications. *Bull. Am. Math. Soc.* **43** (4) 439–561. ISSN 0273-0979, 1088-9485. doi: [10.1090/S0273-0979-06-01126-8](https://doi.org/10.1090/S0273-0979-06-01126-8).
- [36] Huang, J. (2018 a) Invertibility of Adjacency Matrices for Random d -Regular Graphs. arXiv preprint arXiv:1807.06465.
- [37] Huang, J. (2018 b) Invertibility of Adjacency Matrices for Random d -Regular Directed Graphs. arXiv preprint arXiv:1806.01382.
- [38] Janwa, H. and Lal, A. K. (2003) On tanner codes: minimum distance and decoding. *Appl. Algebra Eng. Commun. Comput.* **13** (5) 335–347. ISSN 0938-1279, 1432-0622. doi: [10.1007/s00200-003-0098-4](https://doi.org/10.1007/s00200-003-0098-4).
- [39] Kempton, M. (2016) Non-backtracking random walks and a weighted Ihara’s theorem. *Open J. Discrete Math.* **06** (04) 207–226. ISSN 2161-7635, 2161-7643. doi: [10.4236/ojdm.2016.64018](https://doi.org/10.4236/ojdm.2016.64018).
- [40] Kotani, M. and Sunada, T. (2000) Zeta functions of finite graphs. *J. Math. Sci. Univ. Tokyo* **7** (1) 7–25. ISSN 13405705.
- [41] Lee, J. D., Recht, B., Srebro, N., Tropp, J. and Salakhutdinov, R. R. (2010) Practical large-scale optimization for max-norm regularization. In *Advances in Neural Information Processing Systems*, Vol. 23 (J. Lafferty, C. Williams, J. Shawe-Taylor, R. Zemel and A. Culotta, eds), Curran Associates, Inc., pp. 1297–1305.
- [42] Lee, T., Shraibman, A. and Špalek, R. (2008) A direct product theorem for discrepancy. In *2008 23rd Annual IEEE Conference on Computational Complexity*, pp. 71–80. doi: [10.1109/CCC.2008.25](https://doi.org/10.1109/CCC.2008.25).
- [43] Li, W.-C. W. and Solé, P. (1996) Spectra of regular graphs and hypergraphs and orthogonal polynomials. *Eur. J. Comb.* **17** (5) 461–477. ISSN 0195-6698. doi: [10.1006/eujc.1996.0040](https://doi.org/10.1006/eujc.1996.0040).
- [44] Linial, N., Mendelson, S., Schechtman, G. and Shraibman, A. (2007) Complexity measures of sign matrices. *Combinatorica* **27** (4) 439–463. ISSN 0209-9683, 1439-6912. doi: [10.1007/s00493-007-2160-5](https://doi.org/10.1007/s00493-007-2160-5).
- [45] Litvak, A. E., Lytova, A., Tikhomirov, K., Tomczak-Jaegermann, N. and Youssef, P. (2016) Anti-concentration property for random digraphs and invertibility of their adjacency matrices. *Comptes Rendus Mathématique* **354** (2) 121–124. ISSN 1631-073X. doi: [10.1016/j.crma.2015.12.002](https://doi.org/10.1016/j.crma.2015.12.002).
- [46] Litvak, A. E., Lytova, A., Tikhomirov, K., Tomczak-Jaegermann, N. and Youssef, P. (2017) Adjacency matrices of random digraphs: singularity and anti-concentration. *J. Math. Anal. Appl.* **445** (2) 1447–1491. ISSN 0022-247X. doi: [10.1016/j.jmaa.2016.08.020](https://doi.org/10.1016/j.jmaa.2016.08.020).
- [47] Lubetzky, E. and Sly, A. (2010) Cutoff phenomena for random walks on random regular graphs. *Duke Math. J.* **153** (3) 475–510. ISSN 0012-7094, 1547-7398. doi: [10.1215/00127094-2010-029](https://doi.org/10.1215/00127094-2010-029).
- [48] Marcus, A., Spielman, D. A. and Srivastava, N. (2013 a) Interlacing families I: bipartite Ramanujan graphs of all degrees. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pp. 529–537. doi: [10.1109/FOCS.2013.63](https://doi.org/10.1109/FOCS.2013.63).
- [49] Marcus, A., Spielman, D. A. and Srivastava, N. (2013 b) Interlacing Families II: Mixed Characteristic Polynomials and the Kadison-Singer Problem. arXiv:1306.3969 [math].
- [50] Massoulié, L. (2013) Community Detection Thresholds and the Weak Ramanujan Property. arXiv:1311.3085 [cs].
- [51] McKay, B. D. (1981) Subgraphs of random graphs with specified degrees. *Congressus Numerantium* **33** 213–223.
- [52] McKay, B. D. (1984) Asymptotics for 0-1 matrices with prescribed line sums. In *Enumeration and Design*, pp. 225–238.
- [53] Mizuno, H. and Sato, I. (2003) The semicircle law for semiregular bipartite graphs. *J. Comb. Theory Ser. A* **101** (2) 174–190. ISSN 0097-3165. doi: [10.1016/S0097-3165\(02\)00010-9](https://doi.org/10.1016/S0097-3165(02)00010-9).
- [54] Newman, M. E. J. and Martin, T. (2014) Equitable random graphs. *Phys. Rev. E* **90** (5) 052824. doi: [10.1103/PhysRevE.90.052824](https://doi.org/10.1103/PhysRevE.90.052824).
- [55] Newman, M. (2010) *Networks: An Introduction*. Oxford University Press. ISBN 978-0-19-150070-1.
- [56] Recht, B. and Ré, C. (2013) Parallel stochastic gradient algorithms for large-scale matrix completion. *Math. Program. Comput.* **5** (2) 201–226. ISSN 1867-2957. doi: [10.1007/s12532-013-0053-8](https://doi.org/10.1007/s12532-013-0053-8).
- [57] Richardson, T. and Urbanke, R. (2003) The renaissance of Gallager’s low-density parity-check codes. *IEEE Commun. Mag.* **41** (8) 126–131. ISSN 0163-6804. doi: [10.1109/MCOM.2003.1222728](https://doi.org/10.1109/MCOM.2003.1222728).

- [58] Richardson, T. and Urbanke, R. (2008) *Modern Coding Theory*. Cambridge University Press, New York, NY, USA. ISBN 978-0-521-85229-6.
- [59] Srebro, N., Rennie, J. and Jaakkola, T. S. (2005) Maximum-margin matrix factorization. In *Advances in Neural Information Processing Systems 17* (L. K. Saul, Y. Weiss and L. Bottou, eds), MIT Press, pp. 1329–1336.
- [60] Srebro, N. and Shraibman, A. (2005) Rank, trace-norm and max-norm. In *Learning Theory, Lecture Notes in Computer Science*, Springer, Berlin, Heidelberg, pp. 545–560. ISBN 978-3-540-26556-6 978-3-540-31892-7. doi: [10.1007/11503415_37](https://doi.org/10.1007/11503415_37).
- [61] Tanner, R. (1981) A recursive approach to low complexity codes. *IEEE Trans. Inf. Theory* 27 (5) 533–547. ISSN 0018-9448. doi: [10.1109/TIT.1981.1056404](https://doi.org/10.1109/TIT.1981.1056404).
- [62] Wan, Y. and Meilă, M. (2015) A class of network models recoverable by spectral clustering. In *Advances in Neural Information Processing Systems 28* (C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama and R. Garnett, eds), Curran Associates, Inc., pp. 3285–3293.
- [63] Watanabe, Y. and Fukumizu, K. (2009) Graph zeta function in the bethe free energy and loopy belief propagation. In *Advances in Neural Information Processing Systems*.
- [64] Wormald, N. C. (1981) The asymptotic connectivity of labelled regular graphs. *J. Comb. Theory Ser. B* 31 (2) 156–167. ISSN 0095-8956. doi: [10.1016/S0095-8956\(81\)80021-4](https://doi.org/10.1016/S0095-8956(81)80021-4).

Appendix A. List of symbols

A	adjacency matrix
B	non-backtracking matrix
X^*	transpose of the matrix X
$\sigma(X)$	the eigenvalue spectrum of a matrix X
η	second-largest eigenvalue of the adjacency matrix A
$\lambda_i(X)$	the i th largest eigenvalue, in absolute value, of a matrix X
$\mathcal{G}(n, m, d_1, d_2)$	family of bipartite d_1, d_2 -regular random graphs on n, m vertices
d	the maximum degree: $d = \max\{d_1, d_2\} = d_1$ without loss of generality
$V(G)$	vertex set of graph or subgraph G
$E(G)$	edge set of a graph or subgraph G
$\vec{E}(G)$	oriented edge set of a graph or subgraph G
χ	tree excess of a graph or subgraph G : $\chi(G) = V(G) - E(G) + 1$
γ	a path
Γ_{ef}^ℓ	non-backtracking paths of length $\ell + 1$ from oriented edge e to f
F_{ef}^ℓ	non-backtracking, tangle-free paths of length $\ell + 1$ from oriented edge e to f
$T_{ef}^{\ell, j}$	non-backtracking paths of length $\ell + 1$, from e to f , such that the overall path is tangled but the first j and last $\ell - j$ form tangle-free subpaths
$\ \cdot\ $	ℓ^2 -norm of a vector or spectral norm of a matrix
$\ \cdot\ _F$	Frobenius norm of a matrix