

UC Irvine

UC Irvine Previously Published Works

Title

Sparse random graphs: regularization and concentration of the Laplacian

Permalink

<https://escholarship.org/uc/item/3rh827d9>

Authors

Le, Can M
Levina, Elizaveta
Vershynin, Roman

Publication Date

2015-02-10

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

SPARSE RANDOM GRAPHS: REGULARIZATION AND CONCENTRATION OF THE LAPLACIAN

CAN M. LE, ELIZAVETA LEVINA, AND ROMAN VERSHYNIN

ABSTRACT. We study random graphs with possibly different edge probabilities in the challenging sparse regime of bounded expected degrees. Unlike in the dense case, neither the graph adjacency matrix nor its Laplacian concentrate around their expectations due to the highly irregular distribution of node degrees. It has been empirically observed that simply adding a constant of order $1/n$ to each entry of the adjacency matrix substantially improves the behavior of Laplacian. Here we prove that this regularization indeed forces Laplacian to concentrate even in sparse graphs. As an immediate consequence in network analysis, we establish the validity of one of the simplest and fastest approaches to community detection – regularized spectral clustering, under the stochastic block model. Our proof of concentration of regularized Laplacian is based on Grothendieck’s inequality and factorization, combined with paving arguments.

CONTENTS

1. Introduction	1
2. Outline of the argument	7
3. Community detection in sparse networks	9
4. Grothendieck’s theorem and the first core block	12
5. Expansion of the core, and concentration of the adjacency matrix	15
6. Decomposition of the residual	21
7. Concentration of the Laplacian on the core	24
8. Control of the Laplacian on the residual, and proof of Theorem 1.3	27
9. Proof of Corollary 3.1 (community detection)	29
References	31

1. INTRODUCTION

Concentration properties of random graphs have received a substantial attention in the probability literature. In statistics, applications of these results to network analysis have been a particular focus of recent attention,

Date: April 24, 2015.

R. V. is partially supported by NSF grant 1265782 and U.S. Air Force grant FA9550-14-1-0009. E. Levina is partially supported by NSF DMS grants 1106772 and 1159005.

discussed in more detail in Section 3. For dense graphs (with expected node degrees growing with the number of nodes n), a number of results are available [37, 40, 14, 39, 30], mostly when the average expected node degree grows faster than $\log n$. Real networks, however, are frequently very sparse, and no concentration results are available in the regime where the degrees are bounded by a constant. This paper makes one of the first contributions to the study of concentration in this challenging sparse regime.

1.1. Do random graphs concentrate? Geometry of graphs is reflected in matrices canonically associated to them, most importantly the adjacency and Laplacian matrices. Concentration of random graphs can be understood as concentration of these canonical random matrices around their means.

To recall the notion of graph Laplacian, let A be the $n \times n$ adjacency matrix of an undirected finite graph G on the vertex set V , $|V| = n$, with $A_{ij} = 1$ if there is an edge between vertices i and j , and 0 otherwise. The (symmetric, normalized) *Laplacian* is defined as ¹

$$\mathcal{L}(A) = D^{-1/2}(D - A)D^{-1/2} = I - D^{-1/2}AD^{-1/2}. \quad (1.1)$$

Here I is the identity matrix, and $D = \text{diag}(d_i)$ is the diagonal matrix with degrees $d_i = \sum_{j \in V} A_{ij}$ on the diagonal. Graph Laplacians can be thought of as discrete versions of the Laplace-Beltrami operators on Riemannian manifolds; see [16].

The eigenvalues and eigenvectors of the Laplacian matrix $\mathcal{L}(A)$ reflect some fundamental geometric properties of the graph G . The spectrum of $\mathcal{L}(A)$, which is often called the *graph spectrum*, is a subset of the interval $[0, 2]$. The smallest eigenvalue is always zero. The *spectral gap* of G , which is usually defined as the minimum of the second smallest eigenvalue and the gap between 2 and the largest eigenvalue, provides a quantitative measure of connectivity of G ; see [16].

In this paper we will study Laplacians of *random* graphs. A classical and well studied model of random graphs is the Erdős-Rényi model $G(n, p)$, where an undirected graph G on n vertices is constructed by connecting each pair of vertices independently with a fixed probability p . Although the main result of this paper is neither known nor trivial for $G(n, p)$, we shall work with a more general, *inhomogeneous Erdős-Rényi model* $G(n, (p_{ij}))$ in which edges are still generated independently, but with different probabilities p_{ij} ; see e.g. [10]. This includes many popular network models as special cases, including the stochastic block model [26]. We ask the following basic question.

Question 1.1. *When does the Laplacian of a random graph concentrate near a deterministic matrix?*

¹ We first define the Laplacian of the subgraph induced by non-isolated nodes using (1.1) and then extend it for the whole graph by setting the new row and column entries of isolated nodes to zero. However, we will only work with restrictions of the Laplacian on nodes of positive degrees.

More precisely, for a random graph drawn from the inhomogeneous Erdős-Rényi model $G(n, (p_{ij}))$, we are asking whether, with high probability,

$$\|\mathcal{L}(A) - \mathcal{L}(\bar{A})\| \ll \|\mathcal{L}(\bar{A})\| \quad \text{for} \quad \bar{A} = \mathbb{E} A = (p_{ij}).$$

Here $\|\cdot\|$ is the operator norm and $\mathcal{L}(\bar{A})$ is the Laplacian of the weighted graph with adjacency matrix \bar{A} (obtained by simply replacing A with \bar{A} in the definition of the Laplacian). Since the proper scaling for $\|\mathcal{L}(\bar{A})\|$ is $\Omega(1)$ (except for a trivial graph with no edges), we can equivalently restate the question as whether

$$\|\mathcal{L}(A) - \mathcal{L}(\bar{A})\| \ll 1.$$

The answer to this question is crucial in network analysis; see Section 3.

1.2. Dense graphs concentrate, sparse graphs do not. Concentration of relatively *dense* random graphs – those whose with expected degrees grow at least as fast as $\log n$ – is well understood. Both the adjacency matrix and the Laplacian concentrate in this regime. Indeed, Oliveira [37] showed that the inhomogeneous Erdős-Rényi model satisfies

$$\|\mathcal{L}(A) - \mathcal{L}(\bar{A})\| = O\left(\sqrt{\frac{\log n}{d_0}}\right) \quad (1.2)$$

with high probability, where $d_0 = \min_{i \in V} \sum_{j \in V} \bar{A}_{ij}$ denotes the smallest expected degree of the graph. The concentration inequality (1.2) is non-trivial when its right-hand side is $o(1)$, i.e., $d_0 \gg \log n$.

Results like (1.2) for the Laplacian can be deduced from concentration inequalities for the adjacency matrix A , combined with (simple) concentration inequalities for the degrees of vertices. Concentration for adjacency matrices can in turn be deduced either from matrix-valued deviation inequalities (as in [37]) or from bounds for norms of random matrices (as in [25]).

For *sparse* random graphs, with bounded expected degrees, neither the adjacency matrix nor the Laplacian concentrate, due to the high variance of the degree distribution ([3, 21, 18]). High degree vertices make the adjacency matrix unstable, and low degree vertices make the Laplacian unstable. Indeed, a random graph in Erdős-Rényi model $G(n, p)$ has isolated vertices with high probability if the expected degree $d = np$ is $o(\log n)$. In this case, the Laplacian $\mathcal{L}(A)$ has multiple zero eigenvalues, while $\mathcal{L}(\bar{A})$ has a single eigenvalue at zero and all other eigenvalues at 1. This implies that $\|\mathcal{L}(A) - \mathcal{L}(\bar{A})\| \geq 1$, so the Laplacian fails to concentrate. Moreover, there are vertices with degrees $\gg d$ with high probability, which force $\|A\| \gg d$ while $\|\bar{A}\| = d$, so the adjacency matrix does not concentrate either.

1.3. Regularization of sparse graphs. If the concentration of sparse random graphs fails because of the degree distribution is too irregular, we may naturally ask if *regularizing* the graph in some way solves the problem. If such a regularization is to work, it has to enforce spectrum stability and concentration of the Laplacian, but also preserve the graph's geometry.

One simple way to deal with isolated and very low degree nodes, proposed by [6] and analyzed by [27], is to add the same small number $\tau > 0$ to all entries of the adjacency matrix A . That is, we replace A with

$$A_\tau := A + \tau \mathbf{1}\mathbf{1}^\top \quad (1.3)$$

where $\mathbf{1}$ denotes the vector in \mathbb{R}^n whose components are all equal to 1, and then use the Laplacian of A_τ in all subsequent analysis. This regularization creates weak edges (with weight τ) between all previously disconnected vertices, thus increasing all node degrees by $n\tau$. Another way to deal with low degree nodes, proposed by [14] and studied theoretically by [39], is to add a constant $n\tau$ directly to the diagonal of D in the definition (1.1).

Our paper answers the question of whether the regularization (1.3) leads to concentration in the sparse case, by which we mean the case when all node degrees are bounded. Note that for both regularizations described above, the concentration holds trivially if we allow τ to be arbitrarily large. The concentration was obtained if $n\tau$ grows at least as fast as $\log n$ in [39] and [27]. However, when all expected node degrees are bounded, this requirement will lead to $\tau \mathbf{1}\mathbf{1}^\top$ dominating \bar{A} . To apply the concentration results obtained in [39, 27] to community detection, one needs the average of expected node degrees to grow at least as $\log n$, although the minimum expected degree can stay bounded. This is an unavoidable consequence of using Oliveira's result [37], which gives a $\log n$ factor in the bound, and makes the extension of these bounds to our case of all bounded degrees difficult.

To the best of our knowledge, up to this point it has been unknown whether any regularization creates informative concentration for the adjacency matrix or the graph Laplacian in the sparse case. However, a different Laplacian based on non-backtracking random walks was proposed in [29] and analyzed theoretically in [11]; this can be thought of as an alternative and more complicated form of regularization, since introducing non-backtracking random walks also avoids isolated nodes and very low degree vertices such as those attached to the core of the graph by just one edge (which includes dangling trees). Other methods, which are related to the non-backtracking random walks, are the belief propagation algorithm [20, 19] and the spectral algorithm based on the Bethe Hessian matrix [41]. Although these methods have been empirically shown to perform well in sparse case, there is no theoretical analysis available in that regime so far.

1.4. Sparse graphs concentrate after regularization. We will prove that regularization (1.3) does enforce concentration of the Laplacian even for graphs with bounded expected degrees. To formally state our result for the inhomogeneous Erdős-Rényi model, we shall work with random matrices of the following form.

Assumption 1.2. A is an $n \times n$ symmetric random matrix whose binary entries are jointly independent on and above the diagonal, with $\mathbb{E} A = (p_{ij})$.

Let numbers $d \geq e$, $d_0 > 0$ and α be such that

$$\max_{i,j} np_{ij} \leq d, \quad \min_j \sum_{i=1}^n p_{ij} \geq d_0, \quad \frac{d}{d_0} \leq \alpha.$$

Theorem 1.3 (Concentration of the regularized Laplacian). *Let A be a random matrix satisfying Assumption 1.2 and denote $\bar{A}_\tau = \mathbb{E} A_\tau$. Then for any $r \geq 1$, with probability at least $1 - n^{-r}$ we have*

$$\|\mathcal{L}(A_\tau) - \mathcal{L}(\bar{A}_\tau)\| \leq Cr\alpha^2 \log^3(d) \left(\frac{1}{\sqrt{d}} + \frac{1}{\sqrt{n\tau}} \right) \quad \text{for any } \tau > 0.$$

Here C denotes an absolute constant.

We will give a slightly stronger result in Theorem 8.4. The exponents of r , α and of $\log d$ are certainly not optimal, and to keep the argument more transparent, we did not try to optimize them. We do not know if the $\log d$ term can be completely removed; however, in sparse graphs d and thus $\log d$ are of constant order anyway.

Remark 1.4 (Concentration around the original Laplacian). It is important to ask whether regularization does not destroy the original model – in other words, whether $\mathcal{L}(\bar{A}_\tau)$ is close to $\mathcal{L}(\bar{A})$. If we choose the regularization parameter τ so that $d \gg n\tau \gg 1$, it is easy to check that $\|\mathcal{L}(\bar{A}_\tau) - \mathcal{L}(\bar{A})\| \ll 1$, thus regularization almost does not affect the *expected* geometry of the graph. Together with Theorem 1.3 this implies that

$$\|\mathcal{L}(A_\tau) - \mathcal{L}(\bar{A})\| \ll 1.$$

In other words, regularization forces the Laplacian to stay near $\mathcal{L}(\bar{A})$, and this would not happen without regularization.

Remark 1.5 (Weighted graphs). Since our arguments will be based on probabilistic rather than graph-theoretic considerations, the assumption that A has binary entries is not at all crucial. With small modifications, it can be relaxed for matrices with entries that take values in the interval $[0, 1]$, and possibly for more general distributions of entries. We do not pursue such generalizations to make the arguments more transparent.

Remark 1.6 (Directed graphs). Theorem 1.3 also holds for directed graphs (whose adjacency matrices are not symmetric and have all independent entries) for a suitably modified definition of the Laplacian (1.1), with the two appearances of D replaced by matrices of row and column degrees, respectively. In fact, our proof starts from directed graphs and then generalizes to undirected graphs.

1.5. Concentration on the core. As we noted in Section 1.2, sparse random graphs fail to concentrate without regularization. We are going to show that this failure is caused by just a few vertices, n/d of them. On the rest of the vertices, which form what we call *the core*, both the adjacency matrix and the Laplacian concentrate even without regularization. The idea of

constructing a graph core with large spectral gap has been exploited before. Alon and Kahale [3] constructed a core for random 3-colorable graphs by removing vertices with large degrees; Feige and Ofek [21] constructed a core for $G(n, p)$ in a similar way; Coja-Oghlan and Lanka [18] provided a different construction for a somewhat more general model (random graphs with given expected degrees), which in general cannot be used to model networks with communities. Alon and co-authors [2] used Grothendieck’s inequality and SDP duality to construct a core; they showed that the discrepancy of a graph, which measures how much it resembles a random graph with given expected degrees, is determined by the spectral gap of the restriction of the Laplacian on the core (and vice versa).

The following result gives our construction of the core for the general inhomogeneous Erdős-Rényi model $G(n, (p_{ij}))$. As we will discuss further, our method of core construction is very different from the previous works.

Theorem 1.7 (Concentration on the core). *In the setting of Theorem 1.3, there exists a subset J of $[n]$ which contains all but at most n/d vertices, and such that*

(1) *the adjacency matrix concentrates on $J \times J$:*

$$\|(A - \bar{A})_{J \times J}\| \leq Cr\sqrt{d} \log^3 d;$$

(2) *the Laplacian concentrates on $J \times J$:*

$$\|(\mathcal{L}(A) - \mathcal{L}(\bar{A}))_{J \times J}\| \leq \frac{Cr\alpha^2 \log^3 d}{\sqrt{d}}.$$

We will prove this result in Theorems 5.7 and 7.2 below. Note that concentration of the Laplacian (part 2) follows easily from concentration of the adjacency matrix (part 1). This is because most vertices of the graph have degrees $\sim d$, so keeping only such vertices in the core we can relate the Laplacian to the adjacency matrix as $\mathcal{L}(A) \approx I - \frac{1}{d}A$. This makes the deviation of the Laplacian in Theorem 1.7 about d times smaller than the deviation of the adjacency matrix.

The rest of this paper is organized as follows. Section 2 outlines the steps we will take to prove the main Theorem 1.3. Section 3 discusses the application of this result to community detection in networks. The proof is broken up into the following sections: Section 4 states the Grothendieck’s results we will use and applies them to the first core block (which may not yet be as large as we need). Section 5 presents an expansion of the core to the required size and proves the adjacency matrix concentrates there. Section 6 describes a decomposition of the residual of the graph (after extracting the expanded core) that will allow us to control its behavior. Sections 7 and 8 prove the result for the Laplacian, showing, respectively, that it concentrates on the core and can be controlled on the residual, which completes the proof of the main theorem. The proof of the corollary for community detection is given in Section 9.

2. OUTLINE OF THE ARGUMENT

Our approach to proving Theorem 1.3 consists of the following two steps.

1. Remove the few (n/d) problematic vertices from the graph. On the rest of the graph – the *core* – the Laplacian concentrates even without regularization, by Theorem 1.7.
2. Reattach the problematic vertices – the *residual* – back to the core, and show that regularization provides enough stability so that the concentration is not destroyed.

We will address these two tasks separately.

2.1. Construction of the core. We start with the first step and outline the proof of the adjacency part of Theorem 1.7 (the Laplacian part follows easily, as already noted). Our construction of the core is based on the following theorem which combines two results due to Grothendieck, his famous inequality and a factorization theorem. This result states that the operator norm of a matrix can be bounded by the $\ell_\infty \rightarrow \ell_1$ norm on a large sub-matrix. This norm is defined for an $m \times k$ matrix B as

$$\|B\|_{\infty \rightarrow 1} = \max_{x \in \{-1,1\}^m, y \in \{-1,1\}^k} x^\top B y. \quad (2.1)$$

This norm is equivalent to the *cut norm*, which is more frequently used in theoretical computer science community (see [22, 4, 28]).

Theorem 2.1 (Grothendieck). *For every $m \times k$ matrix B and for any $\delta > 0$, there exists a sub-matrix $B_{I \times J}$ with $|I| \geq (1 - \delta)m$ and $|J| \geq (1 - \delta)k$ and such that*

$$\|B_{I \times J}\| \leq \frac{2\|B\|_{\infty \rightarrow 1}}{\delta \sqrt{mk}}.$$

We will deduce and discuss this theorem in Section 4.1. The $\ell_\infty \rightarrow \ell_1$ norm is simpler to deal with than the operator norm, since the maximum of the quadratic form in (2.1) is taken with respect to vectors x, y whose coordinates are all ± 1 . This can be helpful when B is a random matrix. Indeed, for $B = A - \bar{A}$, one can first use standard concentration inequalities (Bernstein's) to control $x^\top B y$ for fixed x and y , and afterwards apply the union bound over the 2^{m+k} possible choices of x, y . This simple argument shows that, while concentration fails in the operator norm, *adjacency matrices of sparse graphs concentrate in the $\ell_\infty \rightarrow \ell_1$ norm*:

$$\|A - \bar{A}\|_{\infty \rightarrow 1} = O(n\sqrt{d}) \quad \text{with high probability.} \quad (2.2)$$

To see this is a concentration result, note that for large d the right hand side is much smaller than $\|\bar{A}\|_{\infty \rightarrow 1}$, which is of order nd . This fact was observed in [24], and we include the proof in this paper as Lemma 4.6.

Next, applying Grothendieck's Theorem 2.1 with $m = k = n$ and $\delta = 1/20$, we obtain a subset J_1 which contains all but $0.1n$ vertices, on which

the adjacency matrix concentrates:

$$\|(A - \bar{A})_{J_1 \times J_1}\| = O(\sqrt{d}). \quad (2.3)$$

Again, to understand this as concentration, note that for large d the right hand side is much smaller than $\|\bar{A}_{J_1 \times J_1}\|$, which is of order d .

We obtained the concentration inequality claimed in the adjacency part of Theorem 1.7, but with a core that may not be as large as we claimed. Our next goal is to reduce the number of residual vertices from $0.1n$ to n/d . To expand the core, we continue to apply the argument above to the remainder of the matrix, thus obtaining new core blocks. We repeat this process until the core becomes as large as required. At the end, all the core blocks constructed this way are combined using the triangle inequality, at the small cost of a factor polylogarithmic in d .

2.2. Controlling the regularized Laplacian on the residual. The second step is to show that regularized Laplacian $\mathcal{L}(A_\tau)$ is stable with respect to adding a few vertices. We will quickly deduce such stability from the following sparse decomposition of the adjacency matrix.

Theorem 2.2 (Sparse decomposition). *In the setting of Theorem 1.3, we can decompose any sub-matrix $A_{I \times J}$ with at most n/d rows or columns into two matrices with disjoint support,*

$$A = A_C + A_{\mathcal{R}},$$

in such a way that each row of $A_{\mathcal{R}}$ and each column of A_C will have at most $10r \log r$ entries that equal 1.

We will obtain a slightly more informative version of this result as Theorem 6.4. The proof is not difficult. Indeed, using a standard concentration argument it is possible to show that there exists at least one sparse row or column of $A_{I \times J}$. Then we can iterate the process – remove this row or column and find another one from the smaller sub-matrix, etc. The removed rows and columns form the $A_{\mathcal{R}}$ and A_C , respectively.

To use Theorem 2.2 for our purpose, it would be easier to drop the identity from the definition of the Laplacian. Thus we consider the averaging operator

$$L(A) := I - \mathcal{L}(A) = D^{-1/2} A D^{-1/2}, \quad (2.4)$$

which is occasionally also called the Laplacian. We show that the regularized averaging operator is small (in the operator norm) on all sub-matrices with small dimensions.

Theorem 2.3 (Residual). *In the setting of Theorem 1.3, any sub-matrix $L(A_\tau)_{I \times J}$ with at most n/d rows or columns satisfies*

$$\|L(A_\tau)_{I \times J}\| \leq \frac{2}{\sqrt{d}} + \frac{\sqrt{10r \log d}}{\sqrt{n\tau}} \quad \text{for any } \tau > 0.$$

We will prove this result as Theorem 8.2 below. The proof is based on the sparse decomposition constructed in Theorem 2.2 and proceeds as follows. It is enough to bound the norm of $L(A_\tau)_\mathcal{R}$. By definition (2.4), $L(A)$ normalizes each entry of A by the sums of entries in its row and column. It is not difficult to see that Laplacians must scale accordingly, namely

$$\|L(A_\tau)_\mathcal{R}\| \leq \sqrt{\varepsilon} \|L(A_\tau)_{I \times J}\| \quad (2.5)$$

if the sum of entries in each column of $L(A_\tau)_\mathcal{R}$ is at most ε times smaller than the corresponding sum for $L(A_\tau)$. Let us assume that $L(A_\tau)_{I \times J}$ has n/d rows. The sum of entries of each column of $L(A_\tau)_\mathcal{R}$ is at most $n\tau/d + 10r \log d$. (The first term here comes from adding the regularization parameter τ to each of n/d entries of the column, and the second term comes from the sparsity of \mathcal{R} .) The sum of entries of each column of $L(A_\tau)_{I \times J}$ is at least $n\tau$ due to regularization. Substituting this into (2.5), we obtain

$$\|L(A_\tau)_\mathcal{R}\| \leq \sqrt{\frac{n\tau/d + 10r \log d}{n\tau}} \|L(A_\tau)_{I \times J}\|.$$

Since the norm of L is always bounded by 1, this leads to the conclusion of Theorem 2.3.

Finally, Theorem 1.3 follows by combining the core part (Theorem 1.7) with the residual part (Theorem 2.3). To do this, we decompose the part of the Laplacian outside the core $J \times J$ into two residual matrices, one on $J^c \times [n]$ and another on $J \times J^c$. We use that the regularized Laplacian concentrates on the core and is small on each of the residual matrices. Combining these bounds by triangle inequality, we obtain Theorem 1.3.

3. COMMUNITY DETECTION IN SPARSE NETWORKS

3.1. Stochastic models of complex networks. Concentration results for random graphs have remarkable implications for network analysis, specifically for understanding the behavior of spectral clustering applied in the community detection problem. Real world networks are often modelled as random graphs, and finding *communities* – groups of nodes that behave similarly to each other. Most of the models proposed for modeling communities to date are special cases of the inhomogeneous Erdős-Rényi model, which we discussed in Section 1.4. In particular, the *stochastic block model* [26] assigns one of K possible community (block) labels to each node i , which we will call $c_i \in \{1, \dots, K\}$, and then assumes that the probability of an edge $p_{ij} = B_{c_i c_j}$, where B is a symmetric $K \times K$ matrix containing the probabilities of edges within and between communities.

For simplicity of presentation, we focus on the simplest version of the stochastic block model, also known as the balanced planted partition model, which assumes $K = 2$, $B_{11} = B_{22} = p$, $B_{12} = q$, and the two communities contain the same number of nodes (we assume that n is an even number and split the set of vertices into two equal parts \mathcal{C}_1 and \mathcal{C}_2). We further assume that $p > q$, and thus on average there are more edges within communities

than between them. (This is called an assortative network model; the disassortative case $p < q$ can in principle be treated similarly but we will not consider it here). We call this model of random graphs $G(n, p, q)$.

3.2. The community detection problem. The community detection problem is to recover the node labels c_i , $i = 1, \dots, n$ from a single realization of the random graph model, in our case $G(n, p, q)$, or in more common notation, $G(n, \frac{a}{n}, \frac{b}{n})$. A large literature exists on both the detection algorithms and the theoretical results establishing when detection is possible, with the latter mostly confined to the simplest $G(n, \frac{a}{n}, \frac{b}{n})$ model. A conjecture was made in the physics literature [19] and rigorous results established in a series of papers by Mossel, Neeman and Sly, as well as independently by two other groups – see [35, 36, 34, 1, 32]. It is now known that no method can do better than random guessing unless

$$(a - b)^2 > 2(a + b).$$

Further, weak consistency (fraction of mislabelled nodes going to 0 with high probability) is achievable if and only if $(a - b)^2/(a + b) \rightarrow \infty$, and strong consistency, or exact recovery (labelling *all* nodes correctly with high probability) requires a stronger necessary and sufficient condition given by [34] in terms of certain binomial probabilities, which is satisfied when the average expected degree $\frac{1}{2}(a + b)$ is of order $\log n$ or larger, and a and b are sufficiently separated. Most existing results on community detection are obtained in the latter regime, showing exact recovery is possible when the degree grows faster than $\log n$ – see e.g., [33, 9].

There are very few existing results about community detection on sparse graphs with bounded average degrees. Consistency is no longer possible, but one can still hope to do better than random guessing above the detection threshold. A (quite complicated) adaptive spectral algorithm by Coja-Oghlan [17] achieves community detection if

$$(a - b)^2 \geq C(a + b) \log(a + b)$$

for a sufficiently large constant C . Recently, two other spectral algorithms based on non-backtracking random walks were proposed by Mossel, Neeman and Sly [35] and Massoulié [32], which perform detection better than random guessing (fraction of misclassified vertices is bounded away from 0.5 as $n \rightarrow \infty$ with high probability) as long as

$$(a - b)^2 > C(a + b) \text{ for } C \geq 2. \quad (3.1)$$

Finally, semi-definite programming approaches to community detection have been discussed and analyzed in the dense regime [15, 13, 7], and very recently Guédon and Vershynin [24] proved that they achieve community detection in the sparse regime under the same condition (3.1), also using Grothendieck's results.

3.3. Regularized spectral clustering in the sparse regime. As an application of the new concentration results, we show that *regularized spectral clustering* [5] can be used for community detection under the $G(n, \frac{a}{n}, \frac{b}{n})$ model in the sparse regime. Strictly speaking, regularized spectral clustering is performed by first computing the leading K eigenvectors and then applying the K -means clustering algorithm to estimate node labels, but since we are focusing on the case $K = 2$, we will simply show that the signs of the elements of the eigenvector corresponding to the second smallest eigenvalue (under the $G(n, \frac{a}{n}, \frac{b}{n})$ model the eigenvector corresponding to the smallest eigenvalue 0 does not contain information about the community structure) match the partition into communities with high probability. Passing from a concentration result on the Laplacian to a result about K -means clustering on its eigenvectors can be done by standard tools such as those used in [39] and is omitted here.

Corollary 3.1 (Community detection in sparse graphs). *Let $\varepsilon \in (0, 1)$ and $r \geq 1$. Let A be the adjacency matrix drawn from the stochastic block model $G(n, \frac{a}{n}, \frac{b}{n})$. Assume that $a > b$, $a \geq e$, and*

$$(a - b)^2 \geq Cr^2\varepsilon^{-2}(a + b)\log^6 a \quad (3.2)$$

for some large constant $C > 0$. Choose $\tau = (d_1 + \dots + d_n)/n^2$, where d_1, \dots, d_n are degrees of the vertices. Denote by v and \bar{v} the unit-norm eigenvectors associated to the second smallest eigenvalues of $\mathcal{L}(A_\tau)$ and $\mathcal{L}(\bar{A}_\tau)$, respectively. Then with probability at least $1 - n^{-r}$, we have

$$\min_{\beta=\pm 1} \|v + \beta\bar{v}\|_2 \leq \varepsilon.$$

In particular, the signs of the elements of v correctly estimate the partition into the two communities, up to at most εn misclassified vertices.

Let us briefly explain how Corollary 3.1 follows from the new concentration results. According to Theorem 1.3 and the standard perturbation results (Davis-Kahan theorem), the eigenvectors of $\mathcal{L}(A_\tau)$ approximate the corresponding eigenvectors of $\mathcal{L}(\bar{A}_\tau)$ and therefore of $\mathcal{L}(\bar{A})$. The latter matrix has rank two. The trivial eigenvector of $\mathcal{L}(\bar{A}_\tau)$ is $\mathbf{1}$, with all entries equal to 1. The first non-trivial eigenvector has entries 1 and -1 , and it is constant on each of the two communities. Since we have a good approximation of that eigenvector, we can recover the communities.

Remark 3.2 (Alternative regularization). A different natural regularization [14, 39] we briefly mentioned in Section 1.3 is to add a constant, say $n\tau$, to the diagonal of the degree matrix D in the definition of the Laplacian rather than to the adjacency matrix A . Thus we have the alternative regularized Laplacian $I - D_\tau^{-1/2}AD_\tau^{-1/2}$, where $D_\tau = D + n\tau I$. One can think of this regularization as *adding a few stars* to the graph. Suppose for simplicity that $n\tau$ is an integer. It is easy to check that this version of regularized Laplacian can also be obtained as follows: add $n\tau$ new vertices, connect each of them

to all existing vertices, compute the (ordinary) Laplacian of the resulting graph, and restrict it to the original n vertices. It is straightforward to show a version of Theorem 1.3 holds for this regularization as well; we omit it here out of space considerations.

4. GROTHENDIECK'S THEOREM AND THE FIRST CORE BLOCK

Our arguments will be easier to develop for non-symmetric adjacency matrices, which have all independent entries. One can think of them as adjacency matrices of *directed random graphs*. So most of our analysis will be concerning directed graphs, but in the end of some sections we will discuss undirected graphs.

We are about to start proving the adjacency part of Theorem 1.7, first for directed graphs. Our final result will be a little stronger, see Theorems 5.6 and 5.7 below, and it will hold under the following weaker assumptions on A .

Assumption 4.1 (Directed graph, bounded expected average degree). A is an $n \times n$ random matrix with independent binary entries, and $\mathbb{E} A = (p_{ij})$. Let number $d \geq e$ be such that

$$\frac{1}{n} \sum_{i,j=1}^n p_{ij} \leq d.$$

In other words, we shall consider a directed random graph whose *expected average degree is bounded by d* .

In this section, we construct the first core block – one that misses $0.1n$ vertices and on which the adjacency matrix is concentrated as we explained in Section 2.1. The construction will be based on two Grothendieck's theorems.

4.1. Grothendieck's theorems. Grothendieck's inequality is a fundamental result, which was originally proved in [23] and formulated in [31] in the form we are going to use in this paper. Grothendieck's inequality has found applications in many areas [2, 38, 28], and most recently in the analysis of networks [24].

Theorem 4.2 (Grothendieck's inequality). *Consider an $m \times k$ matrix of real numbers $B = (B_{ij})$. Assume that for all numbers $s_i, t_i \in \{-1, 1\}$, one has*

$$\left| \sum_{i,j} B_{ij} s_i t_j \right| \leq 1$$

Then, for any Hilbert space H and all vectors x_i, y_i in H with norms at most 1, one has

$$\left| \sum_{i,j} B_{ij} \langle x_i, y_j \rangle \right| \leq K_G.$$

Here K_G is an absolute constant usually called *Grothendieck's constant*. The best value of K_G is still unknown, and the best known bound is $K_G < \pi/(2 \ln(1 + \sqrt{2})) \leq 1.783$.

It will be useful to formulate Grothendieck's inequality in terms of the $\ell_\infty \rightarrow \ell_1$ norm, which is defined as

$$\begin{aligned} \|B\|_{\infty \rightarrow 1} &= \max_{\|t\|_\infty \leq 1} \|Bt\|_1 = \max_{s \in \{-1,1\}^n, t \in \{-1,1\}^k} s^\top Bt \\ &= \max_{s \in \{-1,1\}^m, t \in \{-1,1\}^k} \sum_{i,j} B_{ij} s_i t_j. \end{aligned} \quad (4.1)$$

Grothendieck's inequality then states that for any $m \times k$ matrix B , any Hilbert space H and all vectors x_i, y_i in H with norms at most 1, one has

$$\left| \sum_{i,j} B_{ij} \langle x_i, y_j \rangle \right| \leq K_G \|B\|_{\infty \rightarrow 1}.$$

Remark 4.3 (Cut norm). The $\ell_\infty \rightarrow \ell_1$ norm is equivalent to the *cut norm*, which is often used in theoretical computer science literature (see [4, 28]), and which is defined as the maximal sum of entries over all sub-matrices of B . The cut norm is obtained if we allow s_i and t_j in (4.1) to take values in $\{0,1\}$ as opposed to $\{-1,1\}$. When A is the adjacency matrix of a random graph and $\bar{A} = \mathbb{E} A$, the cut-norm of $A - \bar{A}$ measures the degree of “randomness” of the graph, as it controls the *fluctuation of the number of edges* that run between any two subset of vertices.

We combine Grothendieck's inequality with another result of A. Grothendieck (see [23, 38]), which characterizes the matrices B for which $\sum_{i,j} B_{ij} \langle x_i, y_j \rangle$ is small for all vectors x_i, y_i with norms at most 1.

Theorem 4.4 (Grothendieck's factorization). *Consider an $m \times k$ matrix of real numbers $B = (B_{ij})$. Assume that for any Hilbert space H and all vectors x_i, y_i in H with norms at most 1, one has*

$$\left| \sum_{i,j} B_{ij} \langle x_i, y_j \rangle \right| \leq 1.$$

Then there exist positive weights μ_i and μ'_j that satisfy $\sum_{i=1}^m \mu_i = 1$ and $\sum_{j=1}^k \mu'_j = 1$ and such that

$$\|D_\mu^{-1/2} B D_{\mu'}^{-1/2}\| \leq 1,$$

where $D_\mu = \text{diag}(\mu_i)$ and $D_{\mu'} = \text{diag}(\mu'_j)$ denote the diagonal matrices with the weights on the diagonal.

Combining Grothendieck's inequality and factorization, we deduce a result that allows one to control the usual (operator) norm by the $\ell_\infty \rightarrow \ell_1$ norm on almost all of the matrix. We already mentioned this result as Theorem 2.1. Let us recall it again and give a proof.

Theorem 4.5 (Grothendieck). *For every $m \times k$ matrix B and for any $\delta > 0$, there exists a sub-matrix $B_{I \times J}$ with $|I| \geq (1 - \delta)m$ and $|J| \geq (1 - \delta)k$ and such that*

$$\|B_{I \times J}\| \leq \frac{2\|B\|_{\infty \rightarrow 1}}{\delta\sqrt{mk}}.$$

Proof. Combining Theorems 4.2 and 4.4, we obtain positive weights μ_i and μ'_j which sum to 1 and satisfy

$$\|D_\mu^{-1/2} B D_{\mu'}^{-1/2}\| \leq K_G \|B\|_{\infty \rightarrow 1}. \quad (4.2)$$

Let us choose the set I to contain the indices of the weights μ_i that are bounded below by $(\delta m)^{-1}$. Since all weights sum to one, I contains at least $(1 - \delta)m$ indices as required. Similarly, we define J to contain the indices of the weights μ'_j that are bounded below by $(\delta k)^{-1}$; this set also has the required cardinality.

By construction, all (diagonal) entries of $D_\mu^{-1/2}$ and $D_{\mu'}^{-1/2}$ are positive and bounded above by $\sqrt{\delta m}$ and $\sqrt{\delta k}$ respectively. This implies that

$$\|(D_\mu^{-1/2} B D_{\mu'}^{-1/2})_{\sigma_1 \times \sigma_2}\| \geq \sqrt{\delta m} \sqrt{\delta k} \|B_{\sigma_1 \times \sigma_2}\|.$$

On the other hand, by (4.2) the left hand side of this inequality is bounded above by $K_G \|B\|_{\infty \rightarrow 1}$. This completes the proof, since Grothendieck's constant K_G is bounded by 2. \square

4.2. Concentration of adjacency matrices in $\ell_\infty \rightarrow \ell_1$ norm. As we explained in Section 1.2, the adjacency matrices of sparse random graphs do not concentrate *in the operator norm*. Remarkably, concentration can be enforced by switching to the $\ell_\infty \rightarrow \ell_1$ norm. We stated an informal version of this result in (2.2); now we are ready for a formal statement. It has been proved in [24]; let restate and prove it here for the reader's convenience.

Lemma 4.6 (Concentration of adjacency matrices in $\ell_\infty \rightarrow \ell_1$ norm). *Let A be a random matrix satisfying Assumption 4.1. Then for any $r \geq 1$ the following holds with probability at least $1 - e^{-2rn}$:*

$$\|A - \bar{A}\|_{\infty \rightarrow 1} \leq 5rn\sqrt{d}.$$

Proof. By definition,

$$\|A - \bar{A}\|_{\infty \rightarrow 1} = \max_{x, y \in \{-1, 1\}^n} \sum_{i, j=1}^n (A_{i, j} - \bar{A}_{i, j}) x_i y_j. \quad (4.3)$$

For a fixed pair x, y , the terms $X_{ij} := (A_{i, j} - \bar{A}_{i, j}) x_i y_j$ are independent random variables. So we can use Bernstein's inequality (see Theorem 2.10 in [12]) to control the sum $\sum_{i, j=1}^n X_{ij}$. There are n^2 terms here, all of them are bounded in absolute value by one, and their average variance is at most

d/n . Therefore by Bernstein's inequality, for any $t > 0$ we have

$$\mathbb{P} \left\{ \sum_{i,j=1}^n X_{ij} > tn^2 \right\} \leq \exp \left(-\frac{n^2 t^2 / 2}{d/n + t/3} \right). \quad (4.4)$$

It is easy to check that this is further bounded by e^{-4rn} if we choose $t = 5r\sqrt{d}/n$. Thus, taking a union bound over 4^n choices of pairs x, y and using (4.3) and (4.4), we obtain that

$$\|A - \bar{A}\|_{\infty \rightarrow 1} \leq tn^2 = 5rn\sqrt{d} \quad (4.5)$$

with probability at least $1 - 4^n \cdot e^{-4rn} \geq 1 - e^{-2rn}$. The lemma is proved. \square

Remark 4.7 (Concentration). To better understand Lemma 4.6 as a concentration result, note that $\|\bar{A}\|_{\infty \rightarrow 1} = nd$ if d is the average expected degree of the graph (that is, $d = \frac{1}{n} \sum_{i,j=1}^n p_{ij}$). Then the conclusion of Lemma 4.6 can then be stated as

$$\|A - \bar{A}\|_{\infty \rightarrow 1} \leq \frac{7r}{\sqrt{d}} \|\bar{A}\|_{\infty \rightarrow 1}.$$

For large d , this means that A concentrates near its mean in $\ell_\infty \rightarrow \ell_1$ norm.

4.3. Construction of the first core block. We can now quickly deduce the existence of the first core block – the one on which the adjacency matrix concentrates in the operator norm, as we outlined in (2.3).

To do this, we first apply Lemma 4.6, then use Grothendieck's Theorem 4.5 for $m = k = n$ and $\delta = 1/20$, and finally we intersect the subsets I and J . We conclude the following.

Proposition 4.8 (First core block). *Let A be a matrix satisfying the conclusion of Concentration Lemma 4.6. There exist a subset $J_1 \subseteq [n]$ which contains all but at most $0.1n$ indices, and such that*

$$\|(A - \bar{A})_{J_1 \times J_1}\| \leq Cr\sqrt{d}.$$

Remark 4.9 (Concentration). To better understand Lemma 4.6, one can check that $\|\bar{A}\| \geq d$ if d is the average expected degree of the graph (that is, $d = \frac{1}{n} \sum_{i,j=1}^n p_{ij}$). Then the conclusion of Lemma 4.6 can then be stated as

$$\|(A - \bar{A})_{J_1 \times J_1}\| \leq \frac{Cr}{\sqrt{d}} \|\bar{A}\|.$$

5. EXPANSION OF THE CORE, AND CONCENTRATION OF THE ADJACENCY MATRIX

Our next goal is to expand the core so it contains all but at most n/d (rather than $0.1n$) vertices. As we explained in Section 2.1, this will be done by repeatedly constructing core blocks (using Grothendieck's theorems) in the parts of the matrix not yet in the core. This time we will require a slightly stronger upper bound on the average degrees than in Assumption 4.1.

Assumption 5.1 (Directed graphs, stronger bound on expected density). A is an $n \times n$ random matrix with independent binary entries, and $\mathbb{E} A = (p_{ij})$. Let number $d \geq e$ be such that

$$\max_{i,j} np_{ij} \leq d.$$

5.1. Concentration in $\ell_\infty \rightarrow \ell_1$ norm on blocks. First, we will need to sharpen the concentration inequality of Lemma 4.6 and make it sensitive to the size of the blocks.

Lemma 5.2 (Concentration of adjacency matrices in $\ell_\infty \rightarrow \ell_1$ norm). *Let A be a random matrix satisfying Assumption 5.1. Then for any $r \geq 1$ the following holds with probability at least $1 - n^{-2r}$. Consider a block² $I \times J$ whose dimensions $m \times k$ satisfy $\min(m, k) \geq n/4d$. Then*

$$\|(A - \bar{A})_{I \times J}\|_{\infty \rightarrow 1} \leq 30r\sqrt{mkd}.$$

Proof. The proof is similar to that of Lemma 4.6, except we take a further union bound over the blocks $I \times J$ in the end. Let us fix I and J . Without loss of generality, we may assume that $m \leq k$. By definition,

$$\|(A - \bar{A})_{I \times J}\|_{\infty \rightarrow 1} = \max_{x \in \{-1,1\}^m, y \in \{-1,1\}^k} \sum_{i \in I, j \in J} (A_{ij} - \bar{A}_{ij})x_i y_j. \quad (5.1)$$

For fixed pair x, y , we use Bernstein's inequality like in Lemma 4.6. Denoting $X_{ij} = (A_{ij} - \bar{A}_{ij})x_i y_j$, we obtain

$$\mathbb{P} \left\{ \sum_{i \in I, j \in J} X_{ij} > tmk \right\} \leq \exp \left(-\frac{mkt^2/2}{d/n + t/3} \right). \quad (5.2)$$

Deviating at this point from the proof of Lemma 4.6, we would like this probability to be bounded by $(en/k)^{-8rk}$ in order to make room for the later union bound over I, J . One can easily check that this happens if we choose $t = 15r\sqrt{d/mn} \log(en/k)$; this is the place where we use the assumption $m \geq n/4d$. Thus, taking a union bound over $2^m \cdot 2^k$ choices of pairs x, y and using (5.1) and (5.2), we obtain that

$$\|(A - \bar{A})_{I \times J}\|_{\infty \rightarrow 1} \leq tkm = 15r\sqrt{mkd} \cdot \sqrt{\frac{k}{n}} \log \left(\frac{en}{k} \right) \quad (5.3)$$

with probability at least $1 - 2^{m+k} \cdot (en/k)^{-8rk}$. We continue by taking a union bound over all choices of I and J . Recalling our assumption that $m \leq k$, we obtain that (5.3) holds uniformly for all I, J , as in the statement of the lemma, with probability at least

$$1 - \sum_{m=n/d}^n \sum_{k=m}^n \binom{n}{m} \binom{n}{k} 2^{m+k} \left(\frac{en}{k} \right)^{-8rk} \geq 1 - n^{-2r}.$$

²By block we mean a product set $I \times J$ with arbitrary index subsets $I, J \subseteq [n]$. These subsets are not required to be intervals of successive integers.

Thus we proved a slightly stronger version of the lemma, since the extra term $\sqrt{\frac{k}{n}} \log\left(\frac{en}{k}\right)$ in (5.3) is always bounded by 2. \square

As in Section 4.3, we can combine Lemma 5.2 with Grothendieck's Theorem 4.5. We conclude the following expansion result.

Lemma 5.3 (Weak expansion of core into a block). *Let A be a matrix satisfying the conclusion of Concentration Lemma 5.2. Consider a block $I \times J$ whose dimensions $m \times k$ satisfy $\min(m, k) \geq n/4d$. Then for every $\delta \in (0, 1)$ there exists a sub-block $I' \times J'$ of dimensions at least $(1 - \delta)m \times (1 - \delta)k$ and such that*

$$\|(A - \bar{A})_{I' \times J'}\| \leq Cr\delta^{-1}\sqrt{d}.$$

5.2. Strong expansion of the core into a block. The core sub-block $I' \times J'$ constructed in Lemma 5.3 is still too small for our purposes. For $m \leq k$, we would like J' to miss the number of columns that is a small fraction in m (the smaller dimension!) rather than k . To achieve this, we can apply Lemma 5.3 repeatedly for the parts of the block not yet in the core, until we gain the required number of columns. Let us formally state and prove this result.

Proposition 5.4 (Strong expansion into a block). *Let A be a matrix satisfying the conclusion of Concentration Lemma 5.2. Then any block $I \times [n]$ with $|I| =: m \geq n/4d$ rows contains a sub-block $I' \times J'$ of dimensions at least $(m - m/8) \times (n - m/8)$ and such that*

$$\|(A - \bar{A})_{I' \times J'}\| \leq Cr\sqrt{d}\log^2 d.$$

Proof. Let $\delta \in (0, 1)$ be a small parameter whose value we will chose later. The first application of Lemma 5.3 gives us a sub-block $I_1 \times J_1$ which misses at most δm rows and δn columns of $I \times [n]$, and on which A concentrates nicely:

$$\|(A - \bar{A})_{I_1 \times J_1}\| \leq Cr\delta^{-1}\sqrt{d}.$$

If the number of missing columns is to big, i.e. $\delta n > m/8$, we apply Lemma 5.3 again for the block consisting of the missing columns, that is for $I \times J_1^c$. It has dimensions at least $m \times \delta n$. We obtain a sub-block $I_2 \times J_2$ which misses at most δm rows and $\delta^2 n$ columns, and on which A nicely concentrates:

$$\|(A - \bar{A})_{I_2 \times J_2}\| \leq Cr\delta^{-1}\sqrt{d}.$$

If the number of missing columns is still too big, i.e. $\delta^2 n > m/8$, we continue this process for $I \times (J_1 \cup J_2)^c$, otherwise we stop. Figure 1 illustrates this process.

The process we just described terminates after a finite number of applications of Lemma 5.3, which we denote by T . The termination criterion yields that

$$T \leq \frac{\log(8n/m)}{\log(1/\delta)} \leq \frac{\log(8d)}{\log(1/\delta)}. \quad (5.4)$$




FIGURE 1. Core expansion into a block. First we construct the leftmost core block $I_1 \times J_1$, then the next core block to the right $I_2 \times J_2$, etc. The number of remaining columns reduces exponentially.

(The second inequality follows from the assumption that $m \geq n/d$.) As an outcome of this process, we obtain disjoint blocks $I_t \times J_t \subseteq I \times [n]$ which satisfy

$$|I \setminus I_t| \leq \delta m \quad \text{and} \quad |J \setminus (J_1 \cup \dots \cup J_T)| \leq m/8 \quad (5.5)$$

for all t . The matrix A concentrates nicely on each of these blocks:

$$\|(A - \bar{A})_{I_t \times J_t}\| \leq Cr\delta^{-1}\sqrt{d}.$$

We are ready to choose the index sets I' and J' that would satisfy the required conclusion. We include in I' all rows of I except those left out at each of the block extractions, and we include in J' all columns of each block. Formally, we define

$$I' := I_1 \cap \dots \cap I_T \quad \text{and} \quad J' := J_1 \cup \dots \cup J_T. \quad (5.6)$$

By (5.5), these subsets are adequately large, namely

$$|I \setminus I'| \leq T\delta m \quad \text{and} \quad |J \setminus J'| \leq m/8. \quad (5.7)$$

To check that A concentrates on $I' \times J'$, we can decompose this block into (parts of) the sub-blocks we extracted before, and use the bounds on their norms. Indeed, using (5.6) we obtain

$$\begin{aligned} \|(A - \bar{A})_{I' \times J'}\| &\leq \sum_{t=1}^T \|(A - \bar{A})_{I' \times J_t}\| \leq \sum_{t=1}^T \|(A - \bar{A})_{I_t \times J_t}\| \\ &\leq CT\delta^{-1}\sqrt{d}. \end{aligned} \quad (5.8)$$

It remains to choose the value of δ . We let $\delta = c/\log(8d)$ where $c > 0$ is an absolute constant. Choosing c small enough to ensure that we have $T\delta \leq 1/8$ according to (5.4). This implies that, due to (5.7), the size of the block $I' \times J'$ we constructed is indeed at least $(m - m/8) \times (n - m/8)$ as we claimed. Finally, using our choice of δ and the bound (5.4) on T we

conclude from (5.8) that

$$\|(A - \bar{A})_{I' \times J'}\| \leq Cr\sqrt{d} \frac{\log^2(8d)}{c \log[c^{-1} \log(8d)]}.$$

This is slightly better than we claimed. \square

5.3. Concentration of the adjacency matrix on the core: final result. Recall that our goal is to improve upon Proposition 4.8 by expanding the core set J_1 until it contains all but n/d vertices. With the expansion tool given by Proposition 5.4, we are one step away from this goal. We are going to show that if the core is not yet as large as we want, we can still expand it a bit more.

Lemma 5.5 (Expansion of the core that is not too large). *Let A be a matrix satisfying the conclusion of Concentration Lemma 5.2. Consider a subset J of $[n]$ which contains all but $m \geq n/4d$ indices. Then there exists a subset J' of $[n]$ which contains all but at most $m/2$ indices, and such that*

$$\|(A - \bar{A})_{J' \times J'}\| \leq \|(A - \bar{A})_{J \times J}\| + Cr\sqrt{d} \log^2 d. \quad (5.9)$$

Proof. We can decompose the entire $[n] \times [n]$ into three disjoint blocks – the core block $J \times J$ and the two blocks $J^c \times [n]$ and $J \times J^c$ in which we would like to expand the core; see Figure 2 for illustration.




FIGURE 2. Expansion of the core.

Applying Proposition 5.4 to the $m \times n$ block $J^c \times [n]$, we obtain a sub-block $I_1 \times J_1$ which contains all but at most $m/8$ of its rows and columns, and on which A nicely concentrates:

$$\|(A - \bar{A})_{I_1 \times J_1}\| \leq Cr\sqrt{d} \log^2 d. \quad (5.10)$$

Doing the same for the $(n - m) \times m$ block $J \times J^c$ (after transposing and extending to an $m \times n$ block), we obtain a sub-block $I_2 \times J_2$ which again contains all but at most $m/8$ of its rows and columns, and on which A nicely concentrates:

$$\|(A - \bar{A})_{I_2 \times J_2}\| \leq Cr\sqrt{d} \log^2 d. \quad (5.11)$$

Let I_0 denote the set of all rows in $[n]$ except those $m/8 + m/8$ rows missed in the construction of either of the two sub-blocks $I_1 \times J_1$ or $I_2 \times J_2$. Similarly,

we let J_0 be the set of the columns. The decomposition of $[n] \times [n]$ considered in the beginning of the proof induces a decomposition of $I_0 \times J_0$ into three blocks, which are sub-blocks of $J \times J$, $I_1 \times J_1$ and $I_2 \times J_2$. (This follows since we remove all missing rows and columns.) Therefore, by triangle inequality we have

$$\|(A - \bar{A})_{I_0 \times J_0}\| \leq \|(A - \bar{A})_{J \times J}\| + \|(A - \bar{A})_{I_1 \times J_1}\| + \|(A - \bar{A})_{I_2 \times J_2}\|.$$

Substituting (5.10) and (5.11), we conclude that A nicely concentrates on the block $I_0 \times J_0$ – just as we desired in (5.9). Since I_0 and J_0 may be different sets, we finalize the argument by choosing $J' = I_0 \cap J_0$. Then $J' \times J'$ is a sub-block of $I_0 \times J_0$, so the concentration inequality (5.9) now holds as promised. Moreover, since each of the sets I_0 and J_0 misses at most $m/4$ indices, J' misses at most $m/2$ indices as claimed. \square

Lemma 5.5 allows us to keep expanding the core until it misses all but $(4d)^{-1}n$ vertices.

Theorem 5.6 (Concentration of adjacency matrix on core). *Let A be a random matrix satisfying Assumption 5.1. Then for any $r \geq 1$ the following holds with probability at least $1 - 2n^{-2r}$. There exists a subset J of $[n]$ which contains all but at most $n/4d$ indices, and such that*

$$\|(A - \bar{A})_{J \times J}\| \leq Cr\sqrt{d} \log^3 d.$$

Proof. Fix a realization of the random matrix A which satisfies the conclusions of Proposition 4.8 and Concentration Lemma 5.2. Then Proposition 4.8 gives us the first subset J_1 that misses at most $0.1n$ indices, and such that

$$\|(A - \bar{A})_{J_1 \times J_1}\| \leq Cr\sqrt{d}.$$

If the number of missing indices is smaller than $n/4d$, we stop. Otherwise we apply the Expansion Lemma 5.5. We obtain a subset J_2 which misses twice fewer indices than J_1 , and for which

$$\|(A - \bar{A})_{J_2 \times J_2}\| \leq Cr\sqrt{d} + Cr\sqrt{d} \log^2 d.$$

If the new number of missing indices is smaller than $n/4d$, we stop. Otherwise we keep applying the Expansion Lemma 5.5.

Each application of this lemma results in an additive term $Cr\sqrt{d} \log^2 d$, and it also halves the number of missing indices. By the stopping criterion, the total number of applications is at most $\log d$. Thus, after the process stops, the final set J satisfies

$$\|(A - \bar{A})_{J \times J}\| \leq Cr\sqrt{d} + Cr\sqrt{d} \log^2 d \cdot \log d.$$

This completes the proof. \square

5.4. Extending the result for undirected graphs. Theorem 5.6 can be readily extended for undirected graphs, where the adjacency matrix A is symmetric, with only entries on and above the diagonal that are independent. We claimed such result in the adjacency part of Theorem 1.7; let us restate and prove it.

Theorem 5.7 (Concentration of adjacency matrix on core: undirected graphs). *Let A be a random matrix satisfying the same requirements as in Assumption 5.1, except A is symmetric. Then for any $r \geq 1$ the following holds with probability at least $1 - 2n^{-2r}$. There exists a subset J of $[n]$ which contains all but at most $n/2d$ indices, and such that*

$$\|(A - \bar{A})_{J \times J}\| \leq Cr\sqrt{d} \log^3 d.$$

Proof. We decompose the matrix $A = A^+ + A^-$ so that each of A^+ and A^- has all independent entries. (Consider the parts of A above and below the diagonal.) It remains to apply Theorem 5.6 for A^+ and A^- and intersect the two subsets we obtain. The conclusion follows by triangle inequality. \square

6. DECOMPOSITION OF THE RESIDUAL

In this section we show how to decompose the residual (in fact, any small matrix) into two parts, one with sparse rows and the other with sparse columns. This will lead to Theorem 2.2, which we will obtain in a slightly more informative form as Theorem 6.4 below.

Again, we will work with directed graphs for most of the time, and in the end discuss undirected graphs.

6.1. Square sub-matrices: selecting a sparse row. First we show how to select just one sparse row from square sub-matrices. Then we extend this to rectangular matrices, and finally we iterate the process to construct the required decomposition.

Lemma 6.1 (Selecting a sparse row). *Let A be a random matrix satisfying Assumption 5.1. Then for any $r \geq 1$ the following holds with probability at least $1 - n^{-2r}$. Every square sub-matrix of A with at most n/d rows has a row with at most $10r \log d$ entries that equal 1.*

Proof. The argument consists of a standard application of Chernoff's inequality and a union bound over the square sub-matrices $A_{I \times J}$.

Let us fix the dimensions $m \times m$ and the support $I \times J$ of a sub-matrix $A_{I \times J}$ for a moment, and consider one of its rows. The number of entries that equal 1 in i -th row $S_i = \sum_{j \in J} A_{ij}$ is a sum of m independent Bernoulli random variables A_{ij} . Each A_{ij} has expectation at most d/n by the assumptions on A . Thus the expected number of ones in i -th row is at most 1, since

$$\mathbb{E} S_i \leq \frac{dm}{n} =: \mu$$

which is bounded by 1 by assumption on m .

To upgrade this to a high-probability statement, we can use Chernoff's inequality. It implies that the probability that i -th row is too dense (denser than we are seeking in the lemma) is

$$\mathbb{P}\{S_i > 10r \log d\} \leq \left(\frac{10r \log d}{e\mu}\right)^{-10r \log d} =: p. \quad (6.1)$$

By independence, the probability that all m rows of $A_{I \times J}$ are too dense is at most p^m . Before we take the union bound over $I \times J$, let us simplify the probability p in (6.1). Since $r \geq 1$, we have $10r \log d / (e\mu) \geq 3/\mu = 3n/(dm)$. Therefore

$$\log(1/p) \geq 10r \log(d) \log\left(\frac{3n}{dm}\right). \quad (6.2)$$

By assumption $m \leq n/d$ on the number of rows, both logarithms in the right hand side of (6.2) are bounded below by 1. Then, using the elementary inequality $2ab \geq a + b$ that is valid for all $a, b \geq 1$, we obtain

$$\log(1/p) \geq 5r \left[\log d + \log \frac{3n}{dm} \right] = 5r \log\left(\frac{3n}{m}\right).$$

Summarizing, we have shown that for a fixed support $I \times J$, the probability that all m rows of the sub-matrix $A_{I \times J}$ are too dense is bounded by

$$p^m \leq \left(\frac{3n}{m}\right)^{-5mr}.$$

It remains to take a union bound over all supports $I \times J$. This bounds the failure probability of the conclusion of lemma by

$$\sum_{m=1}^{n/d} \binom{n}{m}^2 p^m \leq \sum_{m=1}^{n/d} \left(\frac{en}{m}\right)^{2m} \left(\frac{3n}{m}\right)^{-5mr} \leq n^{-2r}.$$

This completes the proof. \square

6.2. Rectangular sub-matrices, and iteration. Although we stated Lemma 6.1 for square matrices, it can be easily adapted for rectangular matrices as well. Indeed, consider a $m \times k$ sub-matrix of A . If the matrix is tall, that is $m \geq k$, then we can extend it to a square $m \times m$ sub-matrix by adding arbitrary columns from A . Applying Lemma 6.1, we obtain a sparse row of the bigger matrix – one with at most $10r \log d$ ones in it. Then trivially the same row of the original $m \times k$ sub-matrix will be sparse as well.

The same reasoning can be repeated for fat sub-matrices, that is for $m \leq k$, this time by applying Lemma 6.1 to the transpose of A . This way we obtain a sparse column of a fat sub-matrix. Combining the two cases, we conclude the following result that is valid for all small sub-matrices.

Lemma 6.2 (Selecting a sparse row or column). *Let A be a random matrix satisfying Assumption 5.1. Then for any $r \geq 1$ the following holds with probability at least $1 - 2n^{-2r}$. Every sub-matrix of A whose dimensions $m \times k$ satisfy $\min(m, k) \leq n/d$ has a row (if $m \geq k$) or a column (if $m \leq k$) with at most $10r \log d$ entries that equal 1.*

Iterating this result – selecting rows and columns one by one – we are going to obtain a desired decomposition of the residual. Here we adopt the following convention. Given a subset \mathcal{R} of $[n] \times [n]$, we denote by $A_{\mathcal{R}}$ the $n \times n$ matrix³ that has the same entries as A on \mathcal{C} and zero outside \mathcal{R} .

Theorem 6.3 (Decomposition of the residual). *Let A be a random matrix satisfying Assumption 5.1. Then for any $r \geq 1$ the following holds with probability $1 - 2n^{-2r}$. Every index subset $I \times J$ of $[n] \times [n]$ whose dimensions $m \times k$ satisfy $\min(m, k) \leq n/d$ can be decomposed into two disjoint subsets \mathcal{R} and \mathcal{C} with the following properties:*

- (i) *each row of \mathcal{R} and each column of \mathcal{C} have at most $\min(m, k)$ entries;*⁴
- (ii) *each row of the matrix $A_{\mathcal{R}}$ and each column of the matrix $A_{\mathcal{C}}$ have at most $10r \log d$ entries that equal 1.*

Proof. Let us fix a realization of A for which the conclusion of Lemma 6.2 holds. Suppose we would like to decompose an $m \times k$ sub-matrix $A_{I \times J}$. According to Lemma 6.2, it has a sparse row or column. Remove this row or column, and apply Lemma 6.2 for the remaining sub-matrix. We obtain a sparse row or column of the smaller matrix. Remove it as well, and apply Lemma 6.2 for the remaining sub-matrix. Continue this process until we removed everything from $A_{I \times J}$. Then define \mathcal{R} to be the union of all rows we removed throughout this process, and \mathcal{C} the union of the removed columns. By construction, \mathcal{R} and \mathcal{C} satisfy part (ii) of the conclusion.

Part (i) follows by analyzing the construction of \mathcal{R} and \mathcal{C} . Without loss of generality, let $m \leq k$. The construction starts by removing columns of $A_{I \times J}$ (which obviously have m entries as required) until the aspect ratio reverses, i.e. there remain fewer columns than m . After that point, both dimensions of the remaining sub-matrix are again bounded by m , so part (i) follows. \square

6.3. Extending the result for undirected graphs. Theorem 6.3 can be readily extended for undirected graphs. We stated such result as Theorem 2.3; let us restate it in a somewhat more informative form.

Theorem 6.4 (Decomposition of the residual, undirected graphs). *Let A be a random matrix satisfying the same requirements as in Assumption 5.1, except A is symmetric. Then for any $r \geq 1$ the following holds with probability $1 - 2n^{-2r}$. Every index subset $I \times J$ of $[n] \times [n]$ whose dimensions $m \times k$ satisfy $\min(m, k) \leq n/d$ can be decomposed into two disjoint subsets \mathcal{R} and \mathcal{C} with the following properties:*

- (i) *each row of \mathcal{R} and each column of \mathcal{C} have at most $2 \min(m, k)$ entries;*
- (ii) *each row of the matrix $A_{\mathcal{R}}$ and each column of the matrix $A_{\mathcal{C}}$ have at most $20r \log d$ entries that equal 1.*

³This does not exactly agree with our usage of $A_{I \times J}$ which denotes an $|I| \times |J|$ matrix, but this slight disagreement will not cause confusion.

⁴Formally, for \mathcal{R} this means that $|\{j : (i, j) \in \mathcal{R}\}| \leq \min(m, k)$ for each $i \in [n]$, and similarly for \mathcal{C} .

Proof. We decompose the matrix $A = A^+ + A^-$ so that each of A^+ and A^- has all independent entries. (Consider the parts of A above and below the diagonal.) It remains to apply Theorem 5.6 for A^+ and A^- and choose \mathcal{R} to be the union of the disjoint sets \mathcal{R}^+ and \mathcal{R}^- we obtain this way; similarly for \mathcal{C} . The conclusion follows trivially. \square

7. CONCENTRATION OF THE LAPLACIAN ON THE CORE

In this section we translate the concentration result on the core, Theorems 5.6, from adjacency matrices to Laplacian matrices. This will lead to the second part of Theorem 1.7.

From now on, we will focus on undirected graphs, where A is a symmetric matrix. Throughout this section, it will be more convenient to work with the alternative Laplacian defined in (2.4) as

$$L(A) = I - \mathcal{L}(A) = D^{-1/2} A D^{-1/2}.$$

Clearly, the concentration results are the same for both definitions of Laplacian, since $L(A) - L(\bar{A}) = \mathcal{L}(A) - \mathcal{L}(\bar{A})$ (and similarly for A_τ).

7.1. Concentration of degrees. We will easily deduce concentration of $L(A)$ on the core from concentration of adjacency matrix A (which we already proved in Theorems 5.6) and the degree matrix $D = \text{diag}(d_j)$. The following lemma establishes concentration of D on the core.

Lemma 7.1 (Concentration of degrees on core). *Let A be a random matrix satisfying the same requirements as in Assumption 5.1, except A is symmetric. Then for any $r \geq 1$, the following holds with probability at least $1 - n^{-2r}$. There exists a subset J of $[n]$ which contains all but at most $n/2d$ indices, and such that the degrees $d_j = \sum_{i=1}^n A_{ij}$ satisfy*

$$|d_j - \mathbb{E} d_j| \leq 30r\sqrt{d \log d} \quad \text{for all } j \in J.$$

Proof. Let us fix $j \in [n]$ for a moment. We decompose A into an upper triangular and a lower triangular matrix, each of which has independent entries. This induces the decomposition of the degrees

$$d_j = \sum_{i=1}^n A_{ij} = \sum_{i=j}^n A_{ij} + \sum_{i=1}^{j-1} A_{ij} =: d_j^- + d_j^+.$$

By triangle inequality, it is enough to show that d_j^- and d_j^+ concentrate near their own expected values. Without loss of generality, let us do this for d_j^+ .

By construction, d_j^+ is a sum of n independent Bernoulli random variables (including $n - j + 1$ zeros) whose variances are all bounded by d/n by assumption on A . Thus Bernstein's inequality (see Theorem 2.10 in [12]) yields

$$\mathbb{P} \left\{ |d_j^+ - \mathbb{E} d_j^+| > nt \right\} \leq \exp \left(- \frac{nt^2/2}{d/n + t/3} \right), \quad t > 0.$$

Choosing $t = 15(r/n)\sqrt{d \log d}$ and simplifying the probability bound, we obtain

$$\mathbb{P} \left\{ |d_j^+ - \mathbb{E} d_j^+| > 15r\sqrt{d \log d} \right\} \leq d^{-13r}.$$

We choose J^+ to consist of the indices for which $|d_j^+ - \mathbb{E} d_j^+| \leq 15r\sqrt{d \log d}$. To control the size of the complement $(J^+)^c$, we may view it as a sum of n independent Bernoulli random variables, each with expectation at most d^{-13r} . Thus $\mathbb{E} |(J^+)^c| \leq nd^{-13r}$, and Chernoff's inequality implies that

$$\mathbb{P} \left\{ |(J^+)^c| > n/4d \right\} \leq \left(\frac{n/4d}{e \cdot nd^{-13r}} \right)^{-n/4d}.$$

Simplifying, we see that this probability is bounded by $2n^{-3r}$.

Repeating the argument for d_j^- , we obtain a similar set J^- . Choosing J to be the intersection of J^+ and J^- and combining the two concentration bounds by triangle inequality, we complete the proof. \square

7.2. Concentration of Laplacian on core. We are ready to prove the second part of Theorem 1.7, which we restate as follows.

Theorem 7.2 (Concentration of Laplacian on core). *Let A be a matrix satisfying Assumption 1.2. Then for any $r \geq 1$, the following holds with probability at least $1 - 3n^{-2r}$. There exists a subset J of $[n]$ which contains all but at most n/d indices, and such that*

$$\|(L(A) - L(\bar{A}))_{J \times J}\| \leq \frac{Cr\alpha^2 \log^3 d}{\sqrt{d}}. \quad (7.1)$$

Proof. We need to compare the Laplacians

$$L(A) = D^{-1/2}AD^{1/2} \quad \text{and} \quad L(\bar{A}) = \bar{D}^{-1/2}\bar{A}\bar{D}^{-1/2}$$

on a big core block $J \times J$, where $D = \text{diag}(d_i)$ contains the actual degrees d_i , and $\bar{D} = \text{diag}(\bar{d}_i)$ the expected degrees $\bar{d}_i = \mathbb{E} d_i$.

We get the core set J by intersecting the two corresponding sets on which A concentrates (from Theorem 5.7) and the degree matrix D concentrates (from Lemma 7.1). To keep the notation simple, let us write the Laplacians on the core as

$$L(A)_{J \times J} = BRB \quad \text{and} \quad L(\bar{A})_{J \times J} = \bar{B}\bar{R}\bar{B},$$

where obviously $R = A_{J \times J}$, $\bar{R} = \bar{A}_{J \times J}$, $B = D_{J \times J}^{-1/2}$ and $\bar{B} = \bar{D}_{J \times J}^{-1/2}$. Then we can express the difference of the Laplacians as a telescoping sum

$$(L(A) - L(\bar{A}))_{J \times J} = B(R - \bar{R})B + B\bar{R}(B - \bar{B}) + (B - \bar{B})\bar{R}\bar{B}. \quad (7.2)$$

We will estimate each of the three terms separately.

By the conclusion of Theorem 5.6, we have

$$\|R - \bar{R}\| \leq Cr\sqrt{d} \log^3 d. \quad (7.3)$$

Moreover, since all entries of \bar{A} are bounded by d/n by assumption, we have $\|\bar{A}\| \leq d$, and in particular the sub-matrix \bar{R} must also satisfy

$$\|\bar{R}\| \leq d. \quad (7.4)$$

Next we compare B and \bar{B} , which are diagonal matrices with entries $d_j^{-1/2}$ and $\bar{d}_j^{-1/2}$ on the diagonal, respectively. Since $\bar{d}_j \geq d_0$ by assumption, we have

$$\|\bar{B}\| \leq \frac{1}{\sqrt{d_0}}. \quad (7.5)$$

Moreover, by the conclusion of Lemma 7.1, the degrees satisfy

$$|d_j - \bar{d}_j| \leq 30r\sqrt{d \log d} \quad \text{for all } j \in J. \quad (7.6)$$

We can assume that the right hand side here is bounded by $d_0/2$; otherwise the right hand side in the desired bound (7.1) is greater than two, which makes the bound trivially true. Therefore, in particular, (7.6) implies

$$d_j \geq \bar{d}_j - d_0/2 \geq d_0/2. \quad (7.7)$$

The difference between the corresponding entries of B and \bar{B} is

$$|d_j^{-1/2} - \bar{d}_j^{-1/2}| = \frac{|d_j - \bar{d}_j|}{(d_j^{1/2} + \bar{d}_j^{1/2})(d_j \bar{d}_j)^{1/2}}.$$

Since $d_j \geq d_0$ by definition, $\bar{d}_j \geq d_0/2$ by (7.7), and $|d_j - \bar{d}_j|$ is small by (7.6), this expression is bounded by

$$\frac{30r\sqrt{d \log d}}{d_0^{3/2}} = \frac{30r\sqrt{\alpha \log d}}{d_0}.$$

This and (7.7) implies that

$$\|B - \bar{B}\| \leq \frac{30r\sqrt{\alpha \log d}}{d_0} \quad \text{and} \quad \|B\| \leq \frac{2}{\sqrt{d_0}}. \quad (7.8)$$

It remains to substitute into (7.2) the bounds (7.3) for $R - \bar{R}$, (7.4) for \bar{R} , (7.8) for $B - \bar{B}$ and B , and (7.5) for \bar{B} . Using triangle inequality and recalling that $d = \alpha d_0$, we obtain (7.1) and complete the proof. \square

Remark 7.3 (Regularized Laplacian). We just showed that the Laplacian concentrates on the core even without regularization. It is also true with regularization. Indeed, Theorem 7.2 holds for the regularized Laplacian $L(A_\tau) = I - \mathcal{L}(A_\tau)$, and they state that

$$\|(L(A_\tau) - L(\bar{A}_\tau))_{J \times J}\| \leq \frac{C r \alpha^2 \log^3 d}{\sqrt{d}} \quad \text{for any } \tau \geq 0. \quad (7.9)$$

This is true because Theorem 7.2 is based on concentration of the adjacency matrix A and the degree matrix D on the core. Both of these results trivially hold with regularization as well as without it, as the regularization τ parameter cancels out, e.g. $A_\tau - \bar{A}_\tau = A - \bar{A}$. We leave details to the interested reader.

8. CONTROL OF THE LAPLACIAN ON THE RESIDUAL, AND PROOF OF THEOREM 1.3

8.1. Laplacian is small on the residual. Now we demonstrate how regularization makes Laplacian more stable. We express this as the fact that small sub-matrices of the regularized Laplacian $L(A_\tau)$ have small norms. This fact can be easily deduced from the sparse decomposition of such matrices that we constructed in Theorem 6.3 and the following elementary observation.

Lemma 8.1 (Restriction of Laplacian). *Let B be an $n \times n$ symmetric matrix with non-negative entries, and let \mathcal{C} be a subset of $[n] \times [n]$. Consider the $n \times n$ matrix $B_{\mathcal{C}}$ that has the same entries as B on \mathcal{C} and zero outside \mathcal{C} . Let $\varepsilon \in (0, 1)$. Suppose the sum of entries of each row of $B_{\mathcal{C}}$ is at most ε times the sum of entries of the corresponding row of B . Then*

$$\|(L(B))_{\mathcal{C}}\| \leq \sqrt{\varepsilon}.$$

Proof. Let us denote by $\tilde{L}(B_{\mathcal{C}})$ an analog of the Laplacian for possibly non-symmetric matrix $B_{\mathcal{C}}$, that is

$$\tilde{L}(B_{\mathcal{C}}) = D_r^{-1/2} B_{\mathcal{C}} D_c^{-1/2}.$$

Here $D_r = \text{diag}(B_{\mathcal{C}} \mathbf{1})$ is a diagonal matrix and each diagonal entry $(D_r)_{i,i}$ of D_r is the sum of entries of i -th row of $B_{\mathcal{C}}$; $D_c = \text{diag}(B_{\mathcal{C}}^T \mathbf{1})$ is a diagonal matrix and $(D_c)_{i,i}$ is the sum of entries of i -th column of $B_{\mathcal{C}}$. Note that we can write $L(B)_{\mathcal{C}}$ as

$$L(B)_{\mathcal{C}} = D^{-1/2} B_{\mathcal{C}} D^{-1/2},$$

where $D = \text{diag}(B \mathbf{1}) = \text{diag}(B^T \mathbf{1})$. We have $(D_r)_{i,i} \leq \varepsilon D_{i,i}$ by the assumption and $(D_c)_{i,i} \leq D_{i,i}$ because \mathcal{C} is a subset of $[n] \times [n]$. Since entries of both $\tilde{L}(B_{\mathcal{C}})$ and $(L(B))_{\mathcal{C}}$ are non-negative, we obtain

$$\|(L(B))_{\mathcal{C}}\| \leq \sqrt{\varepsilon} \|\tilde{L}(B_{\mathcal{C}})\|.$$

It remains to prove $\|\tilde{L}(B_{\mathcal{C}})\| \leq 1$. To see this, consider an $2n \times 2n$ symmetric matrix

$$S = \begin{pmatrix} 0_n & B_{\mathcal{C}} \\ B_{\mathcal{C}}^T & 0_n \end{pmatrix},$$

where 0_n is an $n \times n$ matrix whose entries are zero. The Laplacian of S has the form

$$L(S) = \begin{pmatrix} 0_n & \tilde{L}(B_{\mathcal{C}}) \\ \tilde{L}(B_{\mathcal{C}})^T & 0_n \end{pmatrix}.$$

Since $L(S)$ has norm one, it follows that $\|\tilde{L}(B_{\mathcal{C}})\| \leq 1$. This completes the proof. \square

Theorem 8.2 (Regularized Laplacian on residual). *Let A be a random matrix satisfying the same requirements as in Assumption 5.1, except A is symmetric. Then for any $r \geq 1$ the following holds with probability $1 - 2n^{-2r}$.*

Any sub-matrix $L(A_\tau)_{I \times J}$ of the regularized Laplacian $L(A_\tau)$ with at most n/d rows or columns satisfies

$$\|L(A_\tau)_{I \times J}\| \leq \frac{2}{\sqrt{d}} + \frac{\sqrt{40r \log d}}{\sqrt{n\tau}} \quad \text{for any } \tau > 0.$$

Proof. The decomposition $I \times J = \mathcal{R} \cup \mathcal{C}$ we constructed in Theorem 6.4 reduces the problem to bounding $L(A_\tau)_{\mathcal{R}}$ and $L(A_\tau)_{\mathcal{C}}$. Let us focus on $L(A_\tau)_{\mathcal{R}}$. Recall that every row of the index set \mathcal{R} has at most n/d entries, and every row of the matrix $A_{\mathcal{R}}$ has at most $10r \log d$ entries that equal one (while all other entries are zero). This implies that the sum of entries of each row of $(A_\tau)_{\mathcal{R}}$ is bounded by

$$\frac{n\tau}{d} + 10r \log d.$$

We compare this to the sum of the entries of each row of A_τ , which is trivially at least $n\tau$. It is worthwhile to note that this is the only place in the entire argument where regularization is crucially used. Applying the Restriction Lemma 8.1, we obtain

$$\|L(A_\tau)_{\mathcal{R}}\| \leq \sqrt{\frac{1}{d} + \frac{10r \log d}{n\tau}}.$$

Repeating the same reasoning for columns, we obtain the same bound for $L(A_\tau)_{\mathcal{C}}$. Using triangle inequality and simplifying the expression, we conclude the desired bound for $L(A_\tau)_{I \times J}$. \square

Let us notice a similar, and much simpler, bound for the Laplacian of the regularized expected matrix $\bar{A}_\tau = \mathbb{E} A_\tau$.

Lemma 8.3 (Regularized Laplacian of the expected matrix on the residual). *Let A be a random matrix satisfying the same requirements as in Assumption 5.1, except A is symmetric. Then any sub-matrix $L(\bar{A}_\tau)_{I \times J}$ with at most n/d rows or columns satisfies*

$$\|L(\bar{A}_\tau)_{I \times J}\| \leq \frac{2}{\sqrt{d}} + \frac{2}{\sqrt{n\tau}} \quad \text{for any } \tau > 0.$$

Proof. Assume that $L(\bar{A}_\tau)_{I \times J}$ has at most n/d rows. Recall that the matrix \bar{A}_τ has entries $np_{ij} + \tau$. Then the sum of entries of each column of the sub-matrix $(\bar{A}_\tau)_{I \times J}$ is at most

$$\frac{n}{d} \cdot \max_j (p_{ij} + \tau) \leq \frac{n\tau}{d} + 1.$$

We compare this to the sum of entries of each column of A_τ , which is at least $n\tau$. Applying the Restriction Lemma 8.1, we obtain

$$\|L(\bar{A}_\tau)_{I \times J}\| \leq \sqrt{\frac{1}{d} + \frac{1}{n\tau}}.$$

This leads to the desired conclusion. \square

8.2. Concentration of the regularized Laplacian. We are ready to deduce the main Theorem 1.3 in a slightly stronger form.

Theorem 8.4 (Concentration of the regularized Laplacian). *Let A be a random matrix satisfying Assumption 1.2. Then for any $r \geq 1$, with probability at least $1 - n^{-r}$ we have*

$$\|L(A_\tau) - L(\bar{A}_\tau)\| \leq \frac{Cr\alpha^2 \log^3 d}{\sqrt{d}} + \frac{20\sqrt{r \log d}}{\sqrt{n\tau}} \quad \text{for any } \tau > 0.$$

Proof. The proof is a combination of the Concentration Theorem 7.2 and the Restriction Theorem 8.2. Fix a realization of A for which the conclusions of both of these results hold. Theorem 7.2 yields the existence of a core set J that contains all but at most n/d indices from $[n]$, and on which the regularized Laplacian concentrates:

$$\|(L(A_\tau) - L(\bar{A}_\tau))_{J \times J}\| \leq \frac{Cr\alpha^2 \log^3 d}{\sqrt{d}}. \quad (8.1)$$

(Here we used the version (7.9) that is valid for the regularized Laplacian.)

Next, let us decompose the residual $[n] \times [n] \setminus J \times J$ into two blocks $J^c \times [n]$ and $J \times J^c$. The first block has at most n/d rows, so the conclusion of Restriction Theorem 8.2 applies to it. It follows that

$$\|L(A_\tau)_{J^c \times [n]}\| \leq \frac{2}{\sqrt{d}} + \frac{\sqrt{40r \log d}}{\sqrt{n\tau}},$$

An even simpler bound holds for the expected version $L(\bar{A}_\tau)_{J^c \times [n]}$ according to Lemma 8.3. Summing these two bounds by triangle inequality, we conclude that that

$$\|(L(A_\tau) - L(\bar{A}_\tau))_{J^c \times [n]}\| \leq \frac{4}{\sqrt{d}} + \frac{10\sqrt{r \log d}}{\sqrt{n\tau}}.$$

In a similar way we obtain the same bound for the restriction onto the second residual block, $J \times J^c$. Combining these two bounds with (8.1), we complete the proof by triangle inequality. \square

9. PROOF OF COROLLARY 3.1 (COMMUNITY DETECTION)

Proof of Corollary 3.1. Note that $n\tau$ is the average node degree with expectation $(a+b)/2$. Using Bernstein's inequality (see Theorem 2.10 in [12]), it is easy to check that with probability at least $1 - e^{-rn}$, we have

$$\left| n\tau - \frac{a+b}{2} \right| \leq \frac{16r}{\sqrt{a+b}} \frac{a+b}{2}. \quad (9.1)$$

It follows from assumption (3.2) (and increasing the constant C if necessary) that $16r/\sqrt{a+b} \leq 1/2$. Therefore (9.1) implies

$$\left| n\tau - \frac{a+b}{2} \right| \leq \frac{a+b}{4}. \quad (9.2)$$

Let us fix a realization of the random matrix A which satisfies (9.2) and the conclusion of Theorem 1.3. For the model $G(n, \frac{a}{n}, \frac{b}{n})$ we have $d = a$ and $\alpha \leq 2$. From Theorem 1.3 and (9.2) we obtain

$$\begin{aligned} \|L(A_\tau) - L(\bar{A}_\tau)\| &< C' r \log^3(a) \left(\frac{1}{\sqrt{a}} + \frac{1}{\sqrt{n\tau}} \right) \\ &\leq \frac{3C' r \log^3 a}{\sqrt{a}} =: \delta, \end{aligned} \quad (9.3)$$

for some absolute constant $C' > 0$.

We will use Davis-Kahan Theorem (see Theorem VII.3.2 in [8]) and (9.3) to bound the difference between v and \bar{v} . Matrix $L(\bar{A}_\tau)$ has two non-zero eigenvalues: $\lambda_1 = 1$ and $\lambda_2 = (a - b)/(a + b + n\tau)$. By (9.2) we have

$$\frac{4(a - b)}{7(a + b)} \leq \lambda_2 \leq \frac{4(a - b)}{5(a + b)} \leq \frac{4}{5}. \quad (9.4)$$

To upper-bound the gaps in the spectra of $L(A_\tau)$ and $L(\bar{A}_\tau)$, let us denote

$$S = (\lambda_2 - \delta, 4/5 + \delta), \quad S' = (-\delta, \delta) \cup (1 - \delta, 1 + \delta).$$

Then $\lambda_2 \in S$ because $\lambda_2 \leq 4/5$ by (9.4); the remaining eigenvalues of $L(\bar{A}_\tau)$, which are either zero or one, are in S' . Inequality (9.3) implies that eigenvalues of $L(A_\tau)$ are at most δ away from the corresponding eigenvalues of $L(\bar{A}_\tau)$. Therefore the second largest eigenvalue of $L(A_\tau)$ is in S and the remaining eigenvalues of $L(A_\tau)$ are in S' .

Note that S and S' are disjoint because δ is small compared to λ_2 . In fact, from the definition of δ , assumption (3.2) (increasing the constant C if necessary), and (9.4) we have

$$\delta \leq \frac{3C'}{\sqrt{a}} \cdot \frac{\varepsilon(a - b)}{\sqrt{C(a + b)}} \leq \frac{3C'\varepsilon}{\sqrt{C}} \cdot \frac{a - b}{a + b} \leq \frac{\varepsilon\lambda_2}{20}. \quad (9.5)$$

Using (9.4) and (9.5), we bound the distance $\text{dist}(S, S')$ between S and S' as follows:

$$\text{dist}(S, S') = \min \left\{ \lambda_2 - 2\delta, \frac{1}{5} - 2\delta \right\} \geq \frac{\lambda_2}{4} - 2\delta > \frac{\lambda_2}{8}. \quad (9.6)$$

Applying Theorem VII.3.2 in [8] and using (9.3), (9.6), (9.5) we obtain

$$\|vv^T - \bar{v}\bar{v}^T\| \leq \frac{(\pi/2) \|L(A_\tau) - L(\bar{A}_\tau)\|}{\text{dist}(S, S')} \leq \frac{(\pi/2)\delta}{\lambda_2/8} \leq \frac{\pi\varepsilon}{5}. \quad (9.7)$$

It is easy to check that

$$\min_{\beta=\pm 1} \|v + \beta\bar{v}\|_2 \leq \sqrt{2} \|vv^T - \bar{v}\bar{v}^T\|. \quad (9.8)$$

Therefore from (9.7) and (9.8) we have $\min_{\beta=\pm 1} \|v + \beta\bar{v}\|_2 \leq \varepsilon$. The proof is complete. \square

REFERENCES

- [1] E. Abbe, A. S. Bandeira, and G. Hall. Exact recovery in the stochastic block model. *arXiv:1405.3267*, 2014.
- [2] N. Alon, A. Coja-Oghlan, H. Hàn, M. Kang, V. Rödl, and M. Schacht. Quasi-randomness and algorithmic regularity for graphs with general degree distributions. *SIAM J. Comput.*, 39:2336–2362, 2010.
- [3] N. Alon and N. Kahale. A spectral technique for coloring random 3-colorable graphs. *SIAM J. Comput.*, (26):1733–1748, 1997.
- [4] N. Alon and A. Naor. Approximating the cut-norm via grothendiecks inequality. *SIAM Journal on Computing*, 35(4):787803, 2006.
- [5] A. A. Amini, A. Chen, P. J. Bickel, and E. Levina. Fitting community models to large sparse networks. *Annals of Statistics*, 41(4):2097–2122, 2013.
- [6] A. A. Amini, A. Chen, P. J. Bickel, and E. Levina. Pseudo-likelihood methods for community detection in large sparse networks. *The Annals of Statistics*, 41(4):2097–2122, 2013.
- [7] A. A. Amini and E. Levina. On semidefinite relaxations for the block model. *arXiv:1406.5647*, 2014.
- [8] R. Bhatia. *Matrix Analysis*. Springer-Verlag New York, 1996.
- [9] P. J. Bickel and A. Chen. A nonparametric view of network models and Newman-Girvan and other modularities. *Proc. Natl. Acad. Sci. USA*, 106:21068–21073, 2009.
- [10] B. Bollobas, S. Janson, and O. Riordan. The phase transition in inhomogeneous random graphs. *Random Structures and Algorithms*, 31:3–122, 2007.
- [11] C. Bordenave, M. Lelarge, and L. Massoulié. Non-backtracking spectrum of random graphs: community detection and non-regular Ramanujan graphs. *arxiv:1501.06087*, 2015.
- [12] S. Boucheron, G. Lugosi, and P. Massart. *Concentration inequalities: a nonasymptotic theory of independence*. Oxford University Press, 2013.
- [13] T. Cai and X. Li. Robust and Computationally Feasible Community Detection in the Presence of Arbitrary Outlier Nodes. *preprint, arXiv:1404.6000*, 2014.
- [14] K. Chaudhuri, F. Chung, and A. Tsiatas. Spectral clustering of graphs with general degrees in the extended planted partition model. *Journal of Machine Learning Research Workshop and Conference Proceedings*, 23:35.1 – 35.23, 2012.
- [15] Y. Chen, S. Sanghavi, and H. Xu. Clustering Sparse Graphs. *NIPS*, 2012.
- [16] F. R. K. Chung. *Spectral Graph Theory*. CBMS Regional Conference Series in Mathematics, 1997.
- [17] A. Coja-Oghlan. Graph partitioning via adaptive spectral techniques. *J. Combinatorics, Probability and Computing*, 19:2, 2010.
- [18] A. Coja-Oghlan and A. Lanka. The spectral gap of random graphs with given expected degrees. *The electronic journal of combinatorics*, 16(1), 2009.
- [19] A. Decelle, F. Krzakala, C. Moore, and L. Zdeborová. Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Physical Review E*, 84:066106, 2011.
- [20] A. Decelle, F. Krzakala, C. Moore, and L. Zdeborová. Inference and phase transitions in the detection of modules in sparse networks. *Physical Review Letter*, 107:065701, 2011.
- [21] U. Feige and E. Ofek. Spectral techniques applied to sparse random graphs. *Random Structures & Algorithms*, 27:251–275, 2005.
- [22] A. Frieze and R. Kannan. Quick approximation to matrices and applications. *Combinatorica*, 19(2):175–220, 1999.
- [23] A. Grothendieck. Résumé de la théorie métrique des produits tensoriels topologiques. *Bol. Soc. Mat. São Paulo*, 8:1–79, 1953.

- [24] O. Guédon and R. Vershynin. Community detection in sparse networks via grothendieck’s inequality. *arXiv:1411.4686*, 2014.
- [25] B. Hajek, Y. Wu, and J. Xu. Achieving exact cluster recovery threshold via semidefinite programming. *arXiv:1412.6156*, 2014.
- [26] P. W. Holland, K. B. Laskey, and S. Leinhardt. Stochastic blockmodels: first steps. *Social Networks*, 5(2):109–137, 1983.
- [27] A. Joseph and B. Yu. Impact of regularization on spectral clustering. *arXiv:1312.1733*, 2013.
- [28] S. Khot and A. Naor. Grothendieck-type inequalities in combinatorial optimization. *Communications on Pure and Applied Mathematics*, 65(7):992–1035, 2012.
- [29] F. Krzakala, C. Moore, E. Mossel, J. Neeman, A. Sly, L. Zdeborov, and P. Zhang. Spectral redemption in clustering sparse networks. *Proceedings of the National Academy of Sciences*, 110(52):20935–20940, 2013.
- [30] J. Lei and A. Rinaldo. Consistency of spectral clustering in stochastic block models. *arXiv:1312.2050*, 2013.
- [31] J. Lindenstrauss and A. Pełczyński. Absolutely summing operators in L_p -spaces and their applications. *Studia Math.*, 29:275–326, 1968.
- [32] L. Massoulié. Community detection thresholds and the weak Ramanujan property. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, STOC ’14, pages 694–703, 2014.
- [33] McSherry. Spectral partitioning of random graphs. *Proc. 42nd FOCS*, pages 529–537, 2001.
- [34] E. Mossel, J. Neeman, and A. Sly. Consistency thresholds for binary symmetric block models. *arXiv:1407.1591*, 2014.
- [35] E. Mossel, J. Neeman, and A. Sly. A proof of the block model threshold conjecture. *arXiv:1311.4115*, 2014.
- [36] E. Mossel, J. Neeman, and A. Sly. Reconstruction and estimation in the planted partition model. *Probability Theory and Related Fields*, 2014.
- [37] R. Oliveira. Concentration of the adjacency matrix and of the laplacian in random graphs with independent edges. *arXiv:0911.0600*, 2010.
- [38] G. Pisier. Grothendieck’s theorem, past and present. *Bulletin (New Series) of the American Mathematical Society*, 49(2):237–323, 2012.
- [39] T. Qin and K. Rohe. Regularized spectral clustering under the degree-corrected stochastic blockmodel. In *Advances in Neural Information Processing Systems*, pages 3120–3128, 2013.
- [40] K. Rohe, S. Chatterjee, and B. Yu. Spectral clustering and the high-dimensional stochastic block model. *Annals of Statistics*, 39(4):1878–1915, 2011.
- [41] A. Saade, F. Krzakala, and L. Zdeborová. Spectral clustering of graphs with the Bethe Hessian. *Advances in Neural Information Processing Systems 27*, pages 406–414, 2014.

E-mail address: canle@umich.edu, elevina@umich.edu, romanv@umich.edu